

PrivBox: Verifiable Decentralized Reputation System for Online Marketplaces

Muhammad Ajmal Azad, Samiran Bag, Feng Hao

School of Computing Science, Newcastle University, Newcastle Upon Tyne, United Kingdom

Abstract

In online marketplaces (e-commerce, cloud marketplaces), potential buyers/consumers do not have direct access to inspect the quality of products and services offered by retailers and service providers of marketplaces. Therefore, consumers have to trust the reputation system of the marketplace for making a meaningful decision whether they should have interaction with the particular service provider or not. Consumer's feedback plays an important role while evaluating the trustworthiness of the service provider, but it brings challenges to security and the consumer's privacy. Existing centralized reputation systems collect and process consumer's feedback at the centralized trusted system but these systems could leak sensitive information of consumers (such as buying history, likes and dislikes). To ensure the privacy of consumers, in this paper, we present PrivBox, a privacy-preserving decentralized reputation system that computes reputation of retailers or service providers by leveraging feedback from users in a secure and private way. The PrivBox system uses primitives of a homomorphic cryptographic system and non-interactive zero-knowledge proof to achieve objectives of privacy-preservation and well-formedness. PrivBox performs its operations in a decentralized setting, and ensures the following characteristics. 1) It guarantees privacy of consumers without relying on any trusted setup or trusted third party system, 2) it ensures that the consumer's feedback ratings remain within the prescribed range, and 3) it enables consumers and service providers to verify the computed statistics without relying on a trusted third party. To evaluate the performance, we have implemented operations of the PrivBox system. The results demonstrate that the proposed system has a small communication and computation overheads with the essential properties of privacy-preservation and decentralization.

Keywords: Online Marketplaces, Privacy Preservation, Decentralized Reputation Aggregation, E-commerce, Cloud Marketplaces

1. Introduction

In 2017, an estimated 1.61 billion people have purchased products and services over the Internet (online) marketplaces (for example Amazon, eBay, Taobao, Rakuten, Alibaba) [1]. These transactions have resulted in an aggregate revenue of around \$1.9 trillion [2], and is expected to reach an overall revenue of \$4 trillion by the year 2020, approximately a double of 2017 [3]. Other online systems such as Airbnb and Uber have also attracted a large number of users for services other than purchasing products (for example room sharing, room, and car booking). Recently, cloud marketplaces operated by cloud service providers (like AWS Marketplace, Oracle Marketplace, Salesforce, Azure Marketplace) have also attracted a large number of customers by providing software applications and hardware platforms as a service for the computationally expensive tasks. Gartner has predicted that cloud marketplaces will be expected to reach a revenue of around \$71.55 billion by the year 2021 [4].

Although, online marketplaces play an important role in providing business opportunities to small retailers. The growing market-places have also attracted a large number of fraud retailers whom misuse the platform for financial benefits by committing frauds with the consumers. Some of the most common frauds happening today over online marketplaces are: buyers are not receiving goods that they have ordered, receiving products which have inferior value or are significantly different from the original description [5, 6]. The Experian statistics reveals that e-commerce frauds (online auctions, buying products) have increased by 33% since 2015 [7]. Frauds over online marketplaces have resulted in an annual loss of billions of dollars to consumers all over the world [8, 9, 10].

In an electronic marketplace, a consumer¹ does not have an opportunity to physically inspect and evaluate the quality of products and services before purchasing them. Therefore, the consumer has to trust the system that provides information about the trustworthiness of retailers of the marketplace. Similarly, in a cloud-based edge computing system, the consumer may wish to know the trust-

Email addresses: muhammad.azad@ncl.ac.uk (Muhammad Ajmal Azad), samiran.bag@ncl.ac.uk (Samiran Bag), feng.hao@ncl.ac.uk (Feng Hao)

¹The term consumer, user, and buyer are used interchangeably in this paper, and refer to the same entity.

worthiness of nearby mobile edges before outsourcing his sensitive computation to the edge nodes. The marketplace (e-commerce, mobile edge or cloud marketplace) deploys a reputation system that could assist consumers in evaluating the trustworthiness of the retailers and service providers. Consumers then consider this information, while making a meaningful decision whether they should transact with the retailer or not. The trustworthiness score of a retailer is computed by aggregating feedback scores submitted by consumers of the marketplace who already have had transactions with the retailer.

Reputation systems can be classified into two types. 1) a content-driven system that computes reputation of a retailer using text comments left by consumers for the retailer² (for example, whether the product is received on time, tracking is provided) [11]. 2) a user-driven system that utilizes feedback scores (like, dislike, rating score (0-5)) left by the consumer for his past transactions [12, 13, 14]. The system can also be implemented by combining both content-driven and user-driven systems. The reputation system of online marketplaces (e.g. e-Bay, Amazon, Airbnb, Stack-overflow, online dating applications) mainly uses user ratings in a trusted centralized setup [12, 13]. The trusted system has to protect private information of consumers and ensure the privacy, security, and integrity of the consumer's data. However, consumers are reluctant to trust the centralized system, especially in providing a negative rating to the particular entity because of fear of retaliation if their negative ratings are exposed to others [15, 16]. The data anonymization approaches [17, 18, 19] could provide a privacy-preservation layer by hiding real identities of feedback providers, but anonymization is prone to de-anonymization and de-identification attacks [20, 21, 15]. For example, Minkus et al. [15] were able to identify private information of eBay consumers by correlating feedback scores left by consumers for their purchases on the eBay network and information from their Facebook profiles. The reputation system could protect privacy of feedback providers by using cryptographic systems [22, 23, 24, 25]. However, these systems not only require high system resources but also rely on a trusted group of users for the privacy protection.

In this paper, we present PrivBox, a novel decentralized and verifiable reputation system that securely computes reputation of retailers in a marketplace. The system enables feedback providers to provide their rating scores for retailers in an encrypted form. The design of a PrivBox system consists of three major components: the consumer who uses services, the service provider (a retailer on the marketplace or an online seller) who provides services to consumers, and a public bulletin board (a response collector or the tally system) that holds cryptograms of ratings submitted by consumers. Consumers can rate retailers on the binary scale—likes (1) or dislikes (0)—and submit

cryptograms of ratings to the public Bulletin Board (BB). Anyone (collector, analyst, marketplace, seller or buyer) can then compute the reputation of the particular retailer by simply multiplying the cryptograms from the public bulletin board in a secure way. The PrivBox system guarantees privacy in the presence of active and passive adversaries. Furthermore, the PrivBox system also incorporates non-interactive zero-knowledge proofs to prevent malicious consumers from providing out-of-range values for their ratings. We prototyped cryptographic operations of the PrivBox system, and evaluate its performance for the computation and communication overheads. Experimental results demonstrate that the computational and communication overheads of the PrivBox system are reasonable with the inherent property of privacy-preservation.

In summary, this paper makes the following contributions:

- It presents a decentralized system that does not require any trusted setup or trusted set of users for protecting private inputs of participants. The feedback ratings are exchanged to the public bulletin board in an encrypted form. We apply non-interactive zero-knowledge proofs to ensure the well-formedness of feedback ratings while ensuring the privacy of consumers. The proposed system is also publicly verifiable without the use of any trusted third party.
- We prove the security and privacy properties of the system under malicious and honest-but-curious adversarial models.
- We implement the cryptographic operations of the system and analyze the computational and communication overheads.

The PrivBox system can also be applied in other domains such as collaborative intrusion detection systems [26, 27], where multiple Internet service providers could collaborate for the effective and early identification of attackers; collaborative filtering of spammers in online social networks, where multiple service providers collaborate with each other [28, 29, 30, 31] for the quick detection of stealthy and smart spammers; and the private statistics aggregation in the private and secure surveys [32, 33]. Currently, the trustworthiness of a cloud marketplace is estimated based on the service level agreement by the cloud service provider; however, it is necessary that the trustworthiness estimation should be take into account feedback from consumers of the cloud provider. PrivBox can be adopted in this context by providing reputation aggregation as a service to consumers.

This paper is organized as follows. Section 2 describes systems proposed for computing the reputation of entities in commercial and non-commercial systems. Section 3 provides an overview of the PrivBox system, and Section 4 details the cryptographic operations of the proposed system. Section 5 presents the security and privacy properties of

²Terms retailer, seller, and service provider are used interchangeably in this paper.

PrivBox. Section 6 presents a prototype implementation and evaluation of PrivBox. Section 7 presents discussion on defenses against attacks on PrivBox. Section 8 provides discussions on other potential application domains where the PrivBox system could be applied. Finally, Section 9 concludes the paper.

2. Potential Reputation Systems

Broadly, reputation systems can operate in two modes: 1) systems that require a trusted third party for the collection and aggregation of ratings from consumers of the system [12, 13, 34], and 2) systems that compute reputation in a decentralized way without relying on any trusted third party system [35, 36, 37, 38]. Although, a trusted third party system promises to ensure privacy, security, and integrity of the information submitted by its users, but the system poses a serious threat to the privacy of users when the third party itself becomes a malicious entity [39, 40] or maliciously collaborates with other entities for the financial benefits (for example selling information to marketing or insurance companies) [41, 42]. Furthermore, users are also reluctant to exchange any sensitive information to a centralized system because of privacy concerns and fear of the retaliatory effect if their negative ratings are exposed to the rated entity.

The identity of the feedback provider can be anonymized by using a one-time or life-time anonymous identity [43, 18] but this approach has two limitations: 1) the anonymized data can be de-anonymized by correlating the information from different sources [20, 21, 15], and 2) anonymization would not provide any meaningful information which could be used to characterize behaviour of entities and participants. Furthermore, it is more important to protect feedback values rather than hiding or anonymizing source identities. Li et al. [44] proposed a cryptographic system that enables enterprises to divide the outsourced data and separately stores them in a distributed cloud setup. This approach prevents the cloud service provider from inferring information from the outsourced data, but it has a search overhead at the cloud service provider. Gai et al. [45] proposed a fully homomorphic encryption solution to perform blended arithmetic operations over the real numbers using a tensor-based solution. The proposed approach is noise-free as compared to the Gentry’s lattice-based system which is noisy [46].

Several systems have also been proposed for the reputation aggregation in the P2P network [22, 23, 24, 25, 47, 17] that operates in decentralized settings. These systems normally rely on a preselected set of trusted users or peers for ensuring privacy of participants [48, 49, 47, 17]. Further, existing decentralized protocols require that feedback providers should remain online during the aggregation process. Soska et al. [50] proposed a decentralized anonymous marketplace that uses the public ledger based consensus for aggregating the reputation of retailers while preserving the privacy of feedback providers. Clark et al. [51]

proposed a dynamic decentralized reputation system that allows clients to delegate their trust values to other trusted peers before leaving the network. A Bitcoin-based decentralized reputation system is proposed in [52] where the trust of participants is expressed as Bitcoins. Post et al. [53] proposed Bazaar, which leverages the weighted links between buyers and sellers on the marketplace who have a successful transaction. The weights on the link represent the aggregate monetary value transactions between buyers and sellers.

Major commercial reputation systems have a centralized trusted system architecture [12, 13] responsible for the management and processing of the user’s data. Table 1 presents different features of commercial reputation systems. Commercial reputation systems compute reputation of retailers, consumers, and sellers by adding or averaging the rating scores provided by consumers. For example, eBay³ a popular auction site allows buyers and sellers to rate each other as a positive, a negative or a neutral (represented as 1, -1, 0) score after the transaction. The aggregation process is centralized, where the eBay reputation engine computes the aggregated score of sellers and buyers by summing ratings together. The aggregated ratings are then displayed on the page of the seller and the buyer. Epinions.com⁴ is a general consumer product review site (owned by eBay) that allows users to have a review about the quality of different products before buying them. The Epinions registered users provide ratings (on the scale of 1 to 5 stars) to products and other users. Amazon⁵ is a popular website, starting business as an online bookstore in 1994, but now it has become the largest electronic marketplace in the world. Registered users of the site are allowed to rate retailers at the scale of 1-5 stars after the transaction. The system displays the average of all ratings on the web-page about the retailer. Early web search engines simply use the content of the search query to present the top pages to users. However, spammers can evade these systems by simply including the popular search queries in their content. Web search engines now use link-based reputation systems (for example, the PageRank used in a Google search engine) for suggesting reputed pages at the top of the searched query. In PageRank [54], the reputation of the web-page is computed as the number of reputed pages pointing links to the respective page.

In this paper, we present a novel decentralized reputation aggregation system that protects privacy of users without relying on a trusted system or anonymous identities. The protocol enables feedback providers to submit their ratings in the encrypted form to the bulletin board. Participants or any other entity in the system would not learn the actual feedback value of the individual, but would learn the aggregate statistics of the retailer as whole.

³<https://www.ebay.com/>

⁴<http://www.epinions.com/>

⁵<https://www.amazon.com/>

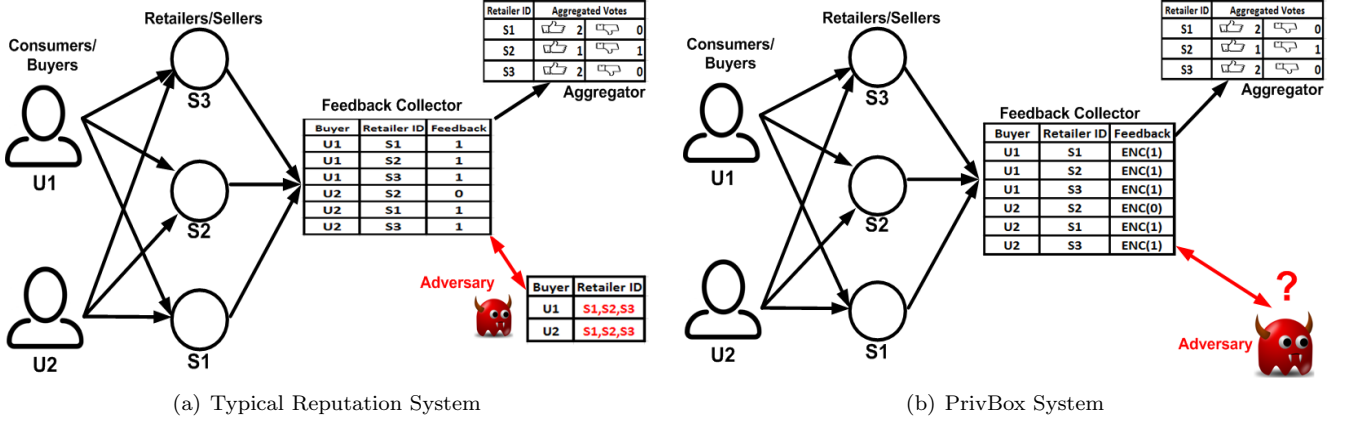


Figure 1: A) A typical reputation aggregation system that could leak how a buyer rated a particular retailer. B) The same aggregation system with the PrivBox system, where the adversary could not link ratings submitted by the buyer.

System	Anonymous Identities	Encrypted Rating	Rating Scale	Architecture	Verifiable
Amazon	✓	✗	(0-5)	Centralized	✗
Ebay	✓	✗	(0,1,-1)	Centralized	✗
Uber	✓	✗	(0-5)	Centralized	✗
Airbnb	✓	✗	(0-5)	Centralized	✗
Epinions	✗	✗	(0-5)	Centralized	✗
OpenBazzar	✓	✗	(0,1 or 0-5)	Decentralized	✗
Yelp	✗	✗	(0-5)	Centralized	✗
PrivBox	✓	✓	(0,1 or 0-5)	decentralized	✓

Table 1: Commercial reputation systems and their attributes.

3. PrivBox Overview

In this section, we present design objectives of the privacy-preserving decentralized reputation system, define the problem, and provide an architecture of the PrivBox system. We also discuss the adversarial model for the proposed system.

3.1. Design Objectives

A typical centralized reputation system is shown in Figure 1.A, where the reputation system of the marketplace collects and aggregates feedback values reported by its consumers. The consumer of the marketplace can have an anonymized or real identity that he uses to submit the feedback value for his interacted retailers. Suppose a consumer $U1$ has purchased the product from the retailer $S1$, and left the feedback value for $S1$ on the marketplace reputation system. Other consumers who transacted with $S1$ have also left positive and negative feedbacks for $S1$. Based on these feedback values, the reputation system updates the reputation of the retailer, which would help new or old consumers to decide whether they should have transaction with the retailer or not. However, the centralized system would know ratings of the particular consumer for his interacted retailers. From Figure 1. A, it can be seen

that the centralized reputation system can learn that consumer $U1$ has rated the retailer $S1$ positively, and also interacted with $S1, S2, S3$. Further, the user of this system does not have the ability to verify the reputation score stated by the reputation system.

The objective of this paper is to present a verifiable decentralized reputation system for the online marketplace. The system provides an opportunity to compute reputation of the retailer while hiding the buyer's feedback scores using a homomorphic cryptographic method. With the placement of the PrivBox solution, a consumer of the marketplace is not required to anonymize his identity; instead, he hides his ratings by presenting cryptograms of ratings. The value of the rating score (0 or 1, like or dislike, rating between 1 to 5 stars) is encrypted using cryptographic primitives as shown in Figure 1.B. To this extent, the adversary on the reputation system or the reputation system itself would not be able to learn how a particular consumer has rated a particular retailer or another interacted consumer. The PrivBox system could provide maximum privacy unless a maximum number of consumers ($n-1$) in the system collude to find the rating score of the target consumer. Furthermore, the design choice of PrivBox ensures two other properties: 1) it limits consumers to provide rating scores within the prescribed range, and 2) it provides public verification of the reputation score stated

by the marketplace.

3.2. Problem Definition

Assume there are n consumers/users who have rated m retailers of the marketplace. Let $F_{ji} \in \{0, 1\}$ be the rating score assigned by the consumer j to the retailer i . The objective is to design a reputation system which takes the secret feedback score from all consumers and computes the reputation of the retailer over the marketplace. The reputation of a given retailer S is computed by summing the secret feedback values. Let $\tau = (\tau_1, \tau_2, \dots, \tau_\eta)$ be the identities of consumers who rated the retailer, such that $\tau_i \in [1, n]$ and $\tau_i \neq \tau_j$, for any $i, j \in [1, \eta], i \neq j$. The PrivBox system computes the reputation RE_i of each retailer $i \in [1, n]$, as follows: $\frac{2 * \sum_{j=1}^{\eta} F_{\tau_j i} - \eta}{\eta + 2}, \forall i \in [1, n]$, where $\sum_{j=1}^{\eta} F_{\tau_j i}$ is the sum of feedback scores, and η is the number of consumers. The reputation system needs to perform its functions in a decentralized and privacy-preserving way such that no trusted third party system and trusted peers are needed for the management of cryptographic parameters. Further, the system should provide correct result even in the presence of malicious participants.

3.3. System Architecture

The PrivBox system consists of three major components as shown in Figure 2.

- 1) **Consumers:** who use services of some service provider or users of the marketplace.
- 2) **Service Providers:** that provide services to users and they can be registered retailers in the marketplace or independent sellers.
- 3) **Bulletin Board (BB):** that holds the cryptograms of rating scores and public cryptographic parameters provided by consumers.

The consumer transacts directly with the marketplace or the retailer (for example buying products from the retailer over the Amazon and eBay networks, or buying products from the independent online store). Once the product is received by the consumer, he is then asked for the feedback about his recent transaction. The consumer provides his feedback to the bulletin board in an encrypted form. The bulletin board acts as a platform for a public authenticated channel, where authenticated consumers can post data, say with a digital signature to prove the data authenticity. In particular, the bulletin board stores the identity of the consumer providing rating, tokens issued by the marketplace (to ensure it is a legitimate transaction), encrypted feedback scores, and the associated zero-knowledge proof to prove the well-formedness of feedback ciphertext. Once the information is published on the BB, the marketplace could use this information to compute the aggregated reputation of the retailer or seller. The marketplace can then put this aggregated reputation score on the web page of the retailer in order to provide information about the trustworthiness of the retailer. Further, a

new or old user can also verify the stated reputation score by accessing the information from the bulletin board in a secure and private way.

3.4. Adversarial Model and Assumptions

With respect to privacy, we assume a malicious adversarial model for feedback providers. In the malicious adversarial model, the participants/consumers try to disrupt the functionality of the system by providing an out-of-range rating. The malicious participant also tries to infer private information about the target user by colluding with others. We assume that users provide encrypted feedback scores to the bulletin board, and use information from the bulletin board for the reputation aggregation and verification of stated reputation. We assume an Honest-but-Curious (HBC) adversary model for the bulletin board. In an HBC model, the party (bulletin board, participant) correctly follow the protocol function but they may use shared information to learn the private information of other participants.

Let G denote a finite cyclic group of a prime order q in which the Decisional Diffie-Hellman (DDH) problem is intractable. Let g be the generator. There are n participants participating in the aggregation process, and they all agreed on (G, g) . Each participant U_i generates a random secret value (private key) and a public key. He then broadcasts the public key to the bulletin board and keeps the private key secret to himself. The public key is publicly available to everyone in the system. We also assume that marketplace issued a token to the participant for providing feedback about the interacted retailer, and this token can be used only once. We assume that the marketplace is honest in providing tokens to those consumers who purchased products from the rated retailers.

3.5. Privacy Preservation Goals

The goal of the PrivBox system is to ensure the integrity of feedback scores provided by participants, such that participants only know their own input and the revealed aggregated score. Further, anyone could publicly verify the aggregated score without relying on any trusted system. Let us assume that there are n participants in the system, and each holds a secret rating score about the retailer. The participant would like to provide this secret score for the aggregation process in such a way that this score should not be revealed to other participants. The PrivBox system ensures the privacy of participants and has the following privacy properties. 1) Only the global reputation score is revealed while the individual feedback remains secret, unlinkable and anonymous throughout the process; 2) the system ensures that adversary participants would not be able to manipulate the scores provided by feedback providers; 3) the system ensures that an encrypted feedback score must be within the prescribed range without actually revealing the feedback

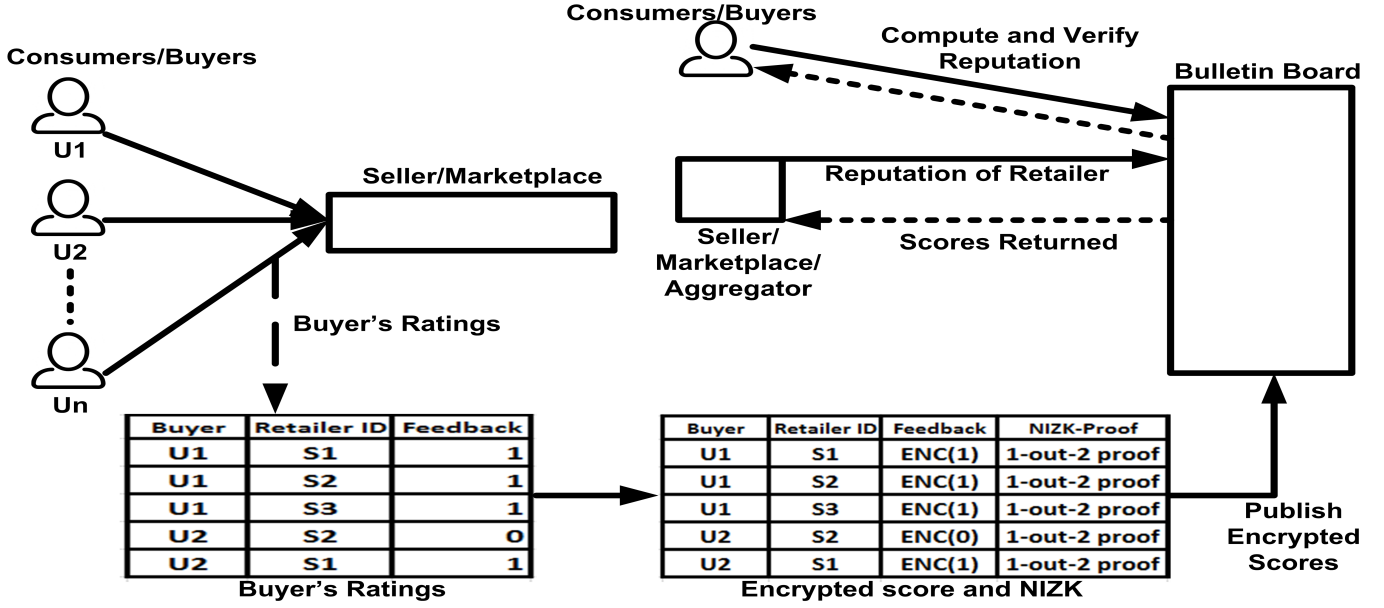


Figure 2: The Architecture of PrivBox System.

value. The system achieves these privacy goals by combining homomorphic encryption and decentralized tallying [55] for the exchange of rating cryptograms, as well as using efficient zero-knowledge proofs for ensuring the cryptogram well-formedness.

4. PrivBox System Operations

In this section, we present cryptographic operations of the PrivBox system. First, we present the system operations for binary responses (Section 4.2) and, second, we extend it to multiple responses (Section 4.7).

4.1. Notations

Table 2 summarizes notations used throughout the paper. We denote the set of all consumers as $U = U_1, U_2, \dots, U_n$. For cryptographic operations, we use x_i to represent the private key of U_i , $X_i = g^{x_i}$ to represent the public key of the U_i and R_i to represent the restructured key of U_i .

4.2. Providing Feedback

Let $U = \{1, 2, \dots, n\}$ be consumers in the network holding the secret feedback scores (0,1) for their interacted retailers and service providers. Assume a multiplicative cyclic group where p and q are large primes that satisfy $q \mid p - 1$. Let there be a subgroup \mathbb{G}_q of order q of the group \mathbb{Z}_p^* , and g is a generator of \mathbb{G}_q . Once the product is purchased from the marketplace and is delivered to the consumer, the marketplace generates a token for the transaction and sends it to the consumer. In order to provide the encrypted feedback for the retailer, the consumer has to generate the cryptographic parameters (public, private and restructured keys). First, the i th consumer generates

Variables	Description
U_i	i 'th user, consumer or participant
$[n]$	the set $\{1, 2, \dots, n\}$
v_{ij}	trust value assigned to seller P_j by U_i
τ	the vector of indices of actual participants
η	no. of elements in τ
G	group of order p in which DDH problem is hard
Y_i	public key of participant $U_i, i \in [\eta]$
x_i	private of participant $U_i, i \in [\eta]$
R_i	Restructured key of participant $U_i, i \in [\eta]$
$C_i = (c_{\tau_i 1}, c_{\tau_i 2}, \dots, c_{\tau_i n})$	vector of encrypted scores of n participants
θ_j	product of cryptograms
θ	$(\theta_1, \theta_2, \dots, \theta_n)$
T'	updated global trust/Reputation vector

Table 2: Notation used in PrivBox.

a random value i.e. a secret key ($x_i \in \mathbb{Z}_q^*$) and computes the public key (X_i) as follows.

$$X_i = g^{x_i} \quad (1)$$

The consumer then publishes the public key X_i on the bulletin board. Table 3 represents the structure of the bulletin board in this phase. Second, the consumer computes the restructured key R_i used for creating the cryptogram as follows.

$$R_i = \prod_{j \in n, j < i} X_j / \prod_{j \in n, j > i} X_j \quad (2)$$

The computation of R_i as above ensures that

$$\prod_{i \in n} R_i^{X_i} = 1. \quad (3)$$

SETUP:G, p, g			
INITIALIZATION:			
$T = (t_1, t_2, \dots, t_n)$			
Key $_{\tau_1}$	Key $_{\tau_2}$	Key $_{\tau_\eta}$
$g^{x_{\tau_1 1}}$	$g^{x_{\tau_2 1}}$	$g^{x_{\tau_\eta 1}}$
$g^{x_{\tau_1 2}}$	$g^{x_{\tau_2 2}}$	$g^{x_{\tau_\eta 2}}$
\vdots	\vdots	\vdots	\vdots
$g^{x_{\tau_1 n}}$	$g^{x_{\tau_2 n}}$	$g^{x_{\tau_\eta n}}$

Table 3: Phase 1: Computing Restructured Key from the public keys published on the public bulletin board.

Finally, the consumer generates the cryptograms of feedback scores using the restructured key and the private key, and presents them to the bulletin board along with the token. The token alone cannot reveal any information about the purchased product, the retailer from whom the products is purchased or the consumer private feedback score. The range of the feedback score can be different on different applications. For example, on Amazon, the feedback score takes the value between 0-5 stars, while on the eBay network, it can be -1,0 or 1. For simplicity, we provide details for the binary responses. This is similar to the way Amazon and eBay ask their consumers for the feedback once the product has been delivered to them. For example, Amazon asks buyers the following questions about their transactions with the retailer: *item arrived by the due date (yes or no)*, *item is the same as described by the retailer (yes or no)*, and *prompt and courteous service (yes or no)*. PrivBox can also be extended to queries having answers on a scale of 0-5 or any other scale with a slight modification in the non-interactive zero-knowledge (NIZK) proof, which proves that only one value has been chosen from the possible responses.

In Privbox, we use +1 for the positive feedback, and 0 to represent the negative feedback, respectively. The feedback submission to the bulletin board is a two-step process. Firstly, the marketplace \mathcal{S} generates a token for the transaction, and sends it to the consumer. Consumers have identities from the array of indices as: $\tau = (\tau_1, \tau_2, \dots, \tau_\eta)$. Obviously, η is the total number of consumers who have been issued a token by the marketplace \mathcal{S} . The consumer then generates the private and the public keys $(x_{\tau_{ij}}, X_{\tau_{ij}} = g^{x_{\tau_{ij}}}) : i \in [1, \eta], j \in [1, n]$, where g is a generator shared among consumers. The consumer then posts the public key $(X_{\tau_{ij}})$ over the public bulletin board. Secondly, the consumer generates the restructured key $(R_{\tau_{ij}} = \prod_{k=1}^{i-1} X_{\tau_{kj}} / \prod_{k=i+1}^{\eta} X_{\tau_{kj}})$ that he can use to generate his encrypted feedback. The consumer creates the cryptogram as follows: $c_{\tau_{ij}} = g^{x_{\tau_{ij}} y_{\tau_{ij}}} g^{v_{\tau_{ij}}} = R_{\tau_{kj}}^{x_{\tau_{kj}}} g^{v_{\tau_{ij}}}$

where $y_{\tau_{ij}} = \sum_{k=1}^{i-1} x_{\tau_{kj}} - \sum_{k=i+1}^{\eta} x_{\tau_{kj}}, \forall i \in [\eta], j \in [n]$, and $v_{\tau_{ij}}$ is either 0 or 1.

In addition to the cryptograms, the consumer \mathcal{U}_{τ_i} also provides the non-interactive zero-knowledge proof (NIZK) to prove the feedback score $v_{\tau_{ij}}$ is either 0 or 1. At the end of this process, the following information is published over the public bulletin board: the identity of the consumer providing the feedback, the encrypted value of the feedback, the identity of the retailer or the service provider for which this response has been provided, the token, and the 1-out-2 NIZK proof. Table 4 presents structure of the bulletin board after consumers have submitted cryptograms to the bulletin board.

4.3. Computing Reputation

Once consumers have submitted their cryptograms and NIZK proofs to the bulletin board, any entity (participant, marketplace, or analyst) can compute the aggregated reputation of the retailer. This can be done by simply multiplying the cryptograms from the bulletin board. The positive score of the retailer or seller $\theta(\theta_1, \theta_2, \dots, \theta_n)$ is computed as below:

$$\theta_r = \prod_{k=1}^{\eta} c_{\tau_k r} \quad (4)$$

$$= \prod_{k=1}^{\eta} g^{x_{\tau_k r} y_{\tau_k r}} g^{v_{\tau_k r}} \quad (5)$$

$$= g^{\sum_{k=1}^{\eta} x_{\tau_k r} y_{\tau_k r} + \sum_{k=1}^{\eta} v_{\tau_k r}} \quad (6)$$

As $\sum_{k=1}^{\eta} x_{\tau_k r} y_{\tau_k r} = 0$, then $\theta_r = \theta_r = g^{\sum_{k=1}^{\eta} v_{\tau_k r}}$. Since, values of $v_{\tau_k r} \in \{0, 1\}, 0 \leq \sum_{k=1}^{\eta} v_{\tau_k r} \leq \eta$, so a brute force search would yield $\sum_{k=1}^{\eta} v_{\tau_k r}$ for all values of $r \in [n]$.

At this point, we already have the aggregate sum of positive ratings i.e. the sum of consumers who have shown trust (1) on the retailer. The number of negative rating can be computed by subtracting the positive ratings from the total number of users who have provided ratings. The simplest approach to compute the reputation of the retailer is to use the negative and positive ratings together, i.e. subtracting negative ratings from the positive ratings [56]. We use the beta reputation system [57] to compute the final aggregated reputation of the business entity or the retailer E on the marketplace. Let n be the number of consumers providing ratings, P_E represents the number of consumers who provided positive ratings about entity E , and N_E represents the number of consumers who rated the entity E as non-trustworthy, then the final reputation RE_E of an entity can be computed as follows:

$$RE_E = \frac{P_E - N_E}{n + 2} \quad (7)$$

The system can be easily extended to other reputation systems, e.g. the average of ratings can be computed by simply averaging the sum of individual ratings over the number of users.

$Token_1$	$Token_2$	$Token_\eta$
C_{τ_1}	C_{τ_2}	C_{τ_η}
$c_{\tau_1 1}$ $P_{WF}[c_{\tau_1 1} : X_{\tau_1 1}, R_{\tau_1 1}]$	$c_{\tau_2 1}$ $P_{WF}[c_{\tau_2 1} : X_{\tau_2 1}, R_{\tau_2 1}]$	$c_{\tau_\eta 1}$ $P_{WF}[c_{\tau_\eta 1} : X_{\tau_\eta 1}, R_{\tau_\eta 1}]$
$c_{\tau_1 2}$ $P_{WF}[c_{\tau_1 2} : X_{\tau_1 2}, R_{\tau_1 2}]$	$c_{\tau_2 2}$ $P_{WF}[c_{\tau_2 2} : X_{\tau_2 2}, R_{\tau_2 2}]$	$c_{\tau_\eta 2}$ $P_{WF}[c_{\tau_\eta 2} : X_{\tau_\eta 2}, R_{\tau_\eta 2}]$
\vdots	\vdots	\vdots	\vdots
\vdots	\vdots	\vdots	\vdots
$c_{\tau_1 n}$ $P_{WF}[c_{\tau_1 n} : X_{\tau_1 n}, R_{\tau_1 n}]$	$c_{\tau_2 n}$ $P_{WF}[c_{\tau_2 n} : X_{\tau_2 n}, R_{\tau_2 n}]$	$c_{\tau_\eta n}$ $P_{WF}[c_{\tau_\eta n} : X_{\tau_\eta n}, R_{\tau_\eta n}]$

Table 4: Phase 2: Providing Feedback and NIZK proof to the public bulletin board.

4.4. Zero Knowledge Proof of Knowledge

A non-interactive zero-knowledge (NIZK) proof is a single zero-knowledge statement that can be sent from the prover (consumer) to the public bulletin board, and could be verified by anyone without interacting with the prover. In PrivBox, we prove knowledge of a secret value of the feedback score assigned by the consumer in a non-interactive manner by applying the Fiat-Shamir heuristic [58].

Here, we show how the consumer would construct a non-interactive zero-knowledge proof to prove the well-formedness for his responses as discussed in Section 4.2. The consumer provides a NIZK proof for a statement of this form: $Z = g^{xy}g^v$, given g^x, g^y and $v \in \{0, 1\}$. In other words, the statement can be written as $(Z = g^{xy}) \vee (Z = g^{xy}g)$, where g^x and g^y are given. Since the prover has a logical ‘OR’ statement to construct a proof, she needs to provide a simulated proof for the false statement and a real proof for the true statement.

Case I: Let us suppose $Z = g^{xy}$. Hence, the prover needs to provide a real proof for the statement $Z = g^{xy}$, and a simulated proof for the second statement $Z = g^{xy}g$. We show how this could be done.

The prover selects $r_1 \in \mathbb{Z}_p$ uniformly at random and generates two commitments $com_1 = (g^y)^{r_1}, com_2 = g^{r_1}$. The prover also generates two other commitments as follows:

$com'_1 = (g^y)^{res'}(Z/g)^{ch_2}, com'_2 = g^{res'}(g^x)^{ch_2}$ where $res', ch_2 \in_R \mathbb{Z}_p$. Now, let ch be the challenge generated by feeding the commitments into a secure hash function. The prover calculates $ch_1 = ch - ch_2$. She generates a response $res = r_1 - ch_1 * x$. Now, the verification equations are listed below.

1. $(g^y)^{res} \stackrel{?}{=} \frac{com_1}{z^{ch_1}}$
2. $g^{res} \stackrel{?}{=} \frac{com_2}{(g^x)^{ch_1}}$
3. $(g^y)^{res'} \stackrel{?}{=} \frac{com'_1}{(z/g)^{ch_2}}$

$$4. g^{res'} \stackrel{?}{=} \frac{com'_2}{(g^x)^{ch_2}}$$

This proof comprises 4 commitments, 2 responses, and 2 challenges.

Case II: Now, suppose $Z = g^{xy}g$. Here, the prover needs to provide a real proof for the statement $Z = g^{xy}g$, and a simulated proof for the second statement $Z = g^{xy}$. The prover selects a random value r_1 , and generates commitments $com'_1 = (g^y)^{r_1}, com'_2 = g^{r_1}$. The prover also generates two other commitments: $com_1 = (g^y)^{res}Z^{ch_1}$ and $com_2 = g^{res}(g^x)^{ch_1}$, where $res, ch_1 \in_R \mathbb{Z}_p$. Let ch be the grand challenge generated by the hash function that takes all system parameters and commitments as inputs. The prover computes $ch_2 = ch - ch_1$ and then calculates response as $res' = r_1 - ch_2 * x$.

The verification equations are as below:

1. $(g^y)^{res} \stackrel{?}{=} \frac{com_1}{z^{ch_1}}$
2. $g^{res} \stackrel{?}{=} \frac{com_2}{(g^x)^{ch_1}}$
3. $(g^y)^{res'} \stackrel{?}{=} \frac{com'_1}{(z/g)^{ch_2}}$
4. $g^{res'} \stackrel{?}{=} \frac{com'_2}{(g^x)^{ch_2}}$

This proof comprises 4 commitments, 2 responses, and 2 challenges.

4.5. Verification

The PrivBox system allows consumers to verify the correctness of the information provided by the marketplace or the business entity. The marketplace publishes the aggregated reputation scores on the account area of the retailer’s portal. The consumer simply accesses the encrypted scores and the associated non-interactive zero-knowledge proofs from the bulletin board, to verify the correctness by using the verification equations mentioned in Section 4.4. If all the NIZK proofs are found to be correct then the aggregated score θ_r is computed using the information from the public bulletin board. All verification operations are

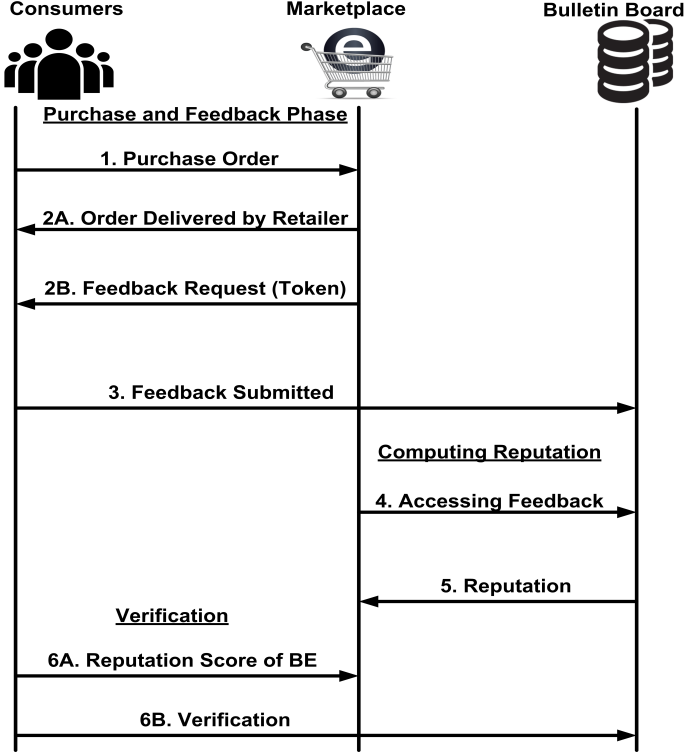


Figure 3: Transaction Workflow in PrivBox System.

performed without actually decrypting the individual responses presented by the consumers, thus preserving each customer's privacy.

4.6. Protocol Workflow

The transaction workflow in the PrivBox system is presented in Figure 3, and each step is described as follows.

1. The buyer/consumer places an order for the product or service from the service provider of an online marketplace (eBay, Amazon, cloud provider).
2. Once the product has been delivered to the buyer, the online marketplace asks the buyer for the feedback about his experience with the retailer. The marketplace asks buyers to provide answers to a set of questions in the form of a survey. The responses are submitted to the bulletin board.
3. Each buyer U_{τ_k} ; $k \in [1, \eta]$ who has received a token from the marketplace then generates the cryptographic parameters i.e. private key $(x_{\tau_k j} \in \mathbb{Z}_p, j \in [1, \eta])$ and the public key $(X_{\tau_k j} = g^{x_{\tau_k j}}, j \in [1, \eta])$. The public key and the NIZK proof $\Pi[x_{\tau_k j} : g, R_{\tau_k j}]$ of knowledge of $x_{\tau_k j} = \log_g R_{\tau_k j}$ are then posted to the bulletin board for all $j \in [1, \eta]$. This proof proves that the buyer U_{τ_k} knows the value of $x_{\tau_k j}$ for all $j \in [1, \eta]$. Finally, the buyer generates the restructured key used for encrypting the feedback score.

4. The buyer submits the feedback in the following form:

$$c_{\tau_k j} = R_{\tau_k j}^{x_{\tau_k j}} g^{v_{\tau_k j}}, j \in [1, \eta]$$

Here $v_{\tau_k j} \in \{0, 1\}$ is the secret feedback of U_{τ_k} for the seller. A value of $v_{\tau_k j}$ equal to 1 represents a positive feedback and a value equal to 0 implies a negative feedback. U_{τ_k} also posts a NIZK proof $\Pi[x_{\tau_k j} : g, X_{\tau_k j}, R_{\tau_k j}, c_{\tau_k j}]$ on the bulletin board. The construction of this NIZK proof is detailed in the section 4.4. This NIZK proof shows that $c_{\tau_k j}$ is either $R_{\tau_k j}^{x_{\tau_k j}}$ or $R_{\tau_k j}^{x_{\tau_k j}} g$ given $g, X_{\tau_k j}$ and $R_{\tau_k j}$, where $R_{\tau_k j} = \prod_{i=1}^{k-1} X_{\tau_i j} / \prod_{i=k+1}^{\eta} X_{\tau_i j}$ for $j \in [1, \eta]$. $R_{\tau_k j}$ is called the restructured key of U_{τ_k} and can be computed by anyone immediately after every buyer has completed step 3.

5. The marketplace or the analyst can utilize cryptograms from the bulletin board to compute the aggregated reputation of the retailer or seller on the marketplace $\prod_{i=1}^{\eta} c_{\tau_i j} = g^{\sum_{i=1}^{\eta} v_{\tau_i j}}$. A limited brute force search will yield the value of $S_j = \sum_{i=1}^{\eta} v_{\tau_i j}$ for all $j \in [1, \eta]$.
6. The new buyer gets the aggregate reputation of the seller or the retailers from the seller's portal on the marketplace. He can verify the statistics by making the query to retrieve the public data from the bulletin board. Similarly, the buyer can also verify whether his feedback is included in the bulletin board as intended.

4.7. Allowing Multiple Choices

So far, we have presented the working of the protocol for binary answers: i.e., participants are allowed to choose one choice from the two available choices (zero or one). However, in online marketplaces, consumers are often allowed to provide the feedback over a range of discrete values (e.g 0-5 stars). The scheme can be trivially extended to allow the participant to have more than two options at their disposal. In order to enable multiple options, we need to run c parallel instances of the same scheme, where c is the total number of choices available to the participants. A participant will have to generate c distinct private/public key pairs, each for a single choice. The user will provide c cryptograms, each one for a distinct option. A one-cryptogram will mean that the particular user has selected the option that corresponds to the cryptogram. Similarly, a zero-cryptogram will mean that the participant does not choose the particular option. Hence, each of the c cryptograms will be a zero-cryptogram or a one-cryptogram. In addition to that, exactly one of the c cryptograms should be a one-cryptogram. This is because of the fact that the participant can choose exactly one of the c available options. The participant will need to provide NIZK proofs that their cryptograms meet these two constraints. That is:

- Each of c cryptograms is either a one-cryptogram or a zero-cryptogram

- Exactly one of the c -cryptograms is a one-cryptogram.

Thus, the participant will have to provide $c+1$ distinct NIZK proofs for $c+1$ statements. They together constitute the proof of well-formedness of the cryptograms.

4.8. Public Bulletin Board

A public bulletin board (BB) is used for facilitating the communication between the entities participating in the reputation aggregation process. The use of BB is a realistic choice as it has been used in privacy-preserving recommendation systems [59, 60, 61] and in electronic voting systems [62, 55, 63]. The bulletin board serves as an authenticated public channel with memory. Participating parties have write-only access (append-only), and other parties (not participating in protocol operation, e.g. analyst, marketplace, or third party system) can read the information from the bulletin board. The bulletin board itself does not have the ability to generate the cryptograms; it can only write the committed information from the participants. The bulletin board can also validate the received scores for their correctness before putting them on the bulletin board, and make sure that no entity could delete or change the published data.

The bulletin board holds the following information: the cryptograms of the feedback scores, the zero-knowledge proof to prove that cryptograms are well-formed, the token to prove that the feedback provider has really interacted with the marketplace, and the identity of the retailer. Anyone can access the information from the bulletin board to compute the aggregated reputation of the retailer or seller. All the computation performed by using information from the bulletin board will be publicly verifiable by executing the aggregation process.

5. Security and Privacy Analysis

In this section, we prove that our scheme is secure in the sense that it does not reveal individual ratings of the users; in other words, at the end of protocol operations, the secret vector T_i s is not compromised provided that some conditions are met.

5.1. Verifiability of Output

In this scheme, the aggregated score of the retailer or seller is publicly available on the bulletin board along with non-interactive zero-knowledge proofs. The proofs provide assurance that the provided encrypted ratings are well-formed. Further, NIZK proofs provide public verifiability for the computation of the aggregated score. If a retailer displays the aggregated score on their website then everyone can check its correctness by looking into the bulletin board. As such, no retailer can misrepresent their own aggregated score on their web page without being caught.

5.2. Unlink-ability

In this section, we show that our scheme protects privacy of the participants of the protocol. Our main result is in Lemma 1. This lemma proves that if there are at least a pair of participants, who provided feedback differently (i.e. one of them submits 0 as an input and the other one provides 1 as an input) then the adversary will not be able to breach the privacy of any of the two participants. In other words, if the feedback values of all the honest participants are not the same, then the adversary will not be able to learn the feedbacks of any of the honest participant. Note that, if the aggregate score of all the participants (both honest and colluding) is known by the end of the protocol operations and the adversary knows the inputs of the colluding users, then the adversary can trivially learn the aggregate or sum of all the inputs of honest users. Now, if all the inputs of the honest users are the same, then the sum will simply expose the inputs of all the honest users. This is because, if all the honest users submit 0 as inputs, then the sum of their inputs will be 0, and if they all submit 1 as inputs, then the sum of their inputs will be equal to the number of honest users. So, in order to preserve the privacy of all the honest users, we must have at least a pair of honest users, who provided distinct inputs. More precisely, the adversary learns nothing more than the aggregate of all the inputs of honest users. Lemma 1 proves this fact. We prove Lemma 1 by reducing it to the well-known Decisional Diffie Hellman problem (DDH). Hence, if the DDH problem is intractable in the mathematical group G , the PrivBox protocol is secure.

Assumption 1. DDH assumption: Given g, g^a, g^b , and $\Omega \in \{g^{ab}, g^{ab}g\}$, it is hard to decide whether $\Omega = g^{ab}$ or $\Omega = g^{ab}g$.

Lemma 1. If there exist at least two participants U_{τ_α} and U_{τ_β} , $\alpha, \beta \in [\eta], \alpha \neq \beta$, such that $v_{\tau_\alpha r} + v_{\tau_\beta r} = 1$ for some seller $r \in [n]$, then no adversary can distinguish between following two cases:

1. $v_{\tau_\alpha r} = 1, v_{\tau_\beta r} = 0$
2. $v_{\tau_\alpha r} = 0, v_{\tau_\beta r} = 1$

Proof. We show that an adversary \mathcal{A} who can distinguish between two bulletin boards where the values of $v_{\tau_\alpha r}$ and $v_{\tau_\beta r}$ are interchanged, then \mathcal{A} could be used to construct another adversary \mathcal{B} against the assumption 1. \mathcal{B} works as follows:

It receives as input g^a, g^b and a challenge $\Omega \in \{g^{ab}, g^{ab}g\}$. It allows \mathcal{A} to choose a set of secret keys $\{x_{\tau_i r} : i \in [\eta] \setminus \{\alpha, \beta\}\}$. \mathcal{A} also chooses the set of scores $\{v_{\tau_i r} : i \in [\eta] \setminus \{\alpha, \beta\}\}$. Then \mathcal{A} sets $X_{\tau_i r} = g^{x_{\tau_i r}}$. \mathcal{B} sets $X_{\tau_\alpha r} = g^{x_{\tau_\alpha r}} = g^a$ and $X_{\tau_\beta r} = g^{x_{\tau_\beta r}} = g^b$. Now, \mathcal{A} computes the set of ballots for each participant $U_{\tau_i} : i \in [\eta] \setminus \{\alpha, \beta\}$ as follows:

$c_{\tau_i r} = g^{x_{\tau_i r} y_{\tau_i r}} g^{v_{\tau_i r}} : i \in [\eta] \setminus \{\alpha, \beta\}$. \mathcal{B} sets $c_{\tau_\alpha r} = (g^a)^{z_{\tau_\alpha r}} g / \Omega$ and $c_{\tau_\beta r} = (g^b)^{z_{\tau_\beta r}} * \Omega$. Here, $z_{\tau_\alpha r} = \sum_{i=1}^{\alpha-1} x_{\tau_i r} -$

$\sum_{i=\alpha+1}^{\beta-1} x_{\tau_i r} - \sum_{i=\beta+1}^{\eta} x_{\tau_i r}$ and $z_{\tau_{\beta} r} = \sum_{i=1}^{\alpha-1} x_{\tau_i r}$
 $+ \sum_{i=\alpha+1}^{\beta-1} x_{\tau_i r} - \sum_{i=\beta+1}^{\eta} x_{\tau_i r}$. $z_{\tau_{\alpha} r}$ and $z_{\tau_{\beta} r}$ can be computed by \mathcal{B} with the help of \mathcal{A} , as the values of $x_{\tau_i r}$ are chosen by \mathcal{A} for all $i \in [n] \setminus \{\alpha, \beta\}$. Now note that, if $\Omega = g^{ab}$, then $c_{\tau_{\alpha} r} = R_{\tau_{\alpha} r}^{x_{\tau_{\alpha} r}} g$ and $c_{\tau_{\beta} r} = R_{\tau_{\beta} r}^{x_{\tau_{\beta} r}}$. That is if $\Omega = g^{ab}$, then $v_{\tau_{\alpha} r} = 1$ and $v_{\tau_{\beta} r} = 0$. Alternatively, if $\Omega = g^{ab} g$, then $c_{\tau_{\alpha} r} = R_{\tau_{\alpha} r}^{x_{\tau_{\alpha} r}}$ and $c_{\tau_{\beta} r} = R_{\tau_{\beta} r}^{x_{\tau_{\beta} r}} g$. That is if $\Omega = g^{ab} g$, then $v_{\tau_{\alpha} r} = 0$ and $v_{\tau_{\beta} r} = 1$. If \mathcal{A} can distinguish between these two cases, \mathcal{B} can identify $\Omega \in \{g^{ab}, g^{ab} g\}$ correctly. \square

5.3. Correctness of Protocol

Lemma 2. *In the PrivBox system, participants can learn the correct aggregated reputation of other participants. We prove this under the model where participants are honest in providing their feedback, but they try to learn the feedback score of other participants.*

Proof. For any seller r , feedbacks provided by the participant $U_{\tau_i r}$ is $c_{\tau_i r} = g^{x_{\tau_i r} y_{\tau_i r}} g^{v_{\tau_i r}}$. The product of all cryptograms is given by $\theta_r = \prod_{i=1}^{\eta} c_{\tau_i r} = \prod_{i=1}^{\eta} g^{x_{\tau_i r} y_{\tau_i r}} g^{v_{\tau_i r}} = g^{\sum_{i=1}^{\eta} x_{\tau_i r} y_{\tau_i r} + \sum_{i=1}^{\eta} v_{\tau_i r}}$. We know that $\sum_{i=1}^{\eta} x_{\tau_i r} y_{\tau_i r} = 0$, then $\theta_r = g^{\sum_{i=1}^{\eta} v_{\tau_i r}}$, or $\sum_{i=1}^{\eta} v_{\tau_i r} = \log_g \theta_r$. This proves computation correctness of the protocol. \square

5.4. Attacks on the System

As described in Section 5.2, the scheme is secure as long as for each $r \in [n]$, there are at least two participants P_{α} and P_{β} satisfying $v_{\alpha r} + v_{\beta r} = 1$, that is if P_{α} and P_{β} assigns different scores to P_r . Let us suppose, the attacker has colluded with some c participants to deduce the score $v_{\alpha r}$ assigned to P_r by an honest participant P_{α} . Alternatively, the attacker can also launch a Sybil attack by creating c fake participants to deduce the score $v_{\alpha r}$. Now, if there exists at least one uncompromised participant P_{β} such that either $v_{\alpha r} = 0 \wedge v_{\beta r} = 1$ or $v_{\alpha r} = 1 \wedge v_{\beta r} = 0$, then it would be computationally infeasible for the attacker to find $v_{\alpha r} = 0$ or $v_{\beta r} = 1$. The scheme assumes that all the participants who completed step 2 of the reputation aggregation protocol in Section 4.6, would provide their feedback as in step 3 of the protocol. If some participants abort after step 2, the system cannot compute the aggregated reputation in that iteration. As such, the iteration has to be started afresh. Though, this kind of denial of service attacks cannot compromise the privacy of any participant, an attacker can use this as a means to disrupt the normal flow of the protocol. A participant who has intentionally aborted the protocol in the middle should be excluded from the network in order to ward off possibilities of sabotage in the future.

5.5. Privacy and Integrity Analysis

In this section, we analyze privacy aspects of participants in the reputation aggregation process. At the end of the reputation aggregation process, each participant or the marketplace can only hold the global reputation score of a particular retailer, which cannot be linked to infer the individual feedback of users, neither can be used to infer who voted positively or negatively for the retailer. The published feedback is the valid score of either 0 or 1 in the following format $g^{xy} g^v$ for $v = 0$ or 1. It is indistinguishable from random feedback and the associated 1-out-of-2 ZKP reveals nothing more than the statement: the v is either 0 or 1. The encrypted feedback value and computation on the encrypted data ensure that participants would not learn anything about the feedback value other than the final aggregated reputation score. The scheme is also secure if a number of feedback providers collude with each other to learn the feedback score of some target user. However, the scheme can reveal the feedback value when $n - 1$ participants collude against the single remaining user. Note that, the final aggregate of all feedbacks has to be made public. Hence, it is impossible to ensure the privacy of an honest feedback provider when all other feedback providers collude against her. The protocol assures the maximum possible privacy for any honest feedback providers, which a scheme of this sort can achieve. Moreover, our protocol does not require any centralized trusted third party for the generation of cryptographic parameters and can provide reputation aggregation even if a certain participant is not online at the time of the reputation aggregation.

6. Implementation and Evaluation

We have implemented functionalities of the PrivBox system in a Java program using the bouncycastle⁶ cryptographic library. We choose the standard elliptic curve NIST P-256 for the cryptographic setting and SHA-512 for the hashing. The functionality of the bulletin board is implemented as a web server, and the functionality of the feedback provider is implemented as a Java client. We evaluate the performance of various cryptographic operations in terms of computational cost and bandwidth overhead. All experiments were performed on an Intel i-7 3.4GHz system running Windows 10, with 8 GB RAM. The experiments were carried out for the single core.

6.1. Computation Complexity

We present the computation costs in terms of time require for creating the cryptograms (the encrypted feedback and the non-interactive zero-knowledge proof) at the client side, and the time require for aggregating the cryptograms from the bulletin board. Table 5 presents a comparison of the PrivBox system architecture and computational complexity with other reputation systems. The

⁶<https://www.bouncycastle.org/>

Proposal	Architecture	Adversarial Model	Privacy	Complexity	Verifiability
Hassan et al. [24]	Decentralized	Malicious	depends on preselected peers	$O(n)+O(\log n)$	No
Androulaki et al. [64]	Decentralized	Semi-honest	compromised if user colludes	$O(n)$	No
Gudes et al. [22]-1	Decentralized	Semi-honest	depend on witness peers	$O(n^2)+O(N)$	No
Gudes et al. [22]-2	Decentralized	Semi-honest	depend preselected peers	$O(1)$	No
Zhai et al. [18]	Distributed	Honest	depend selected peers	$O(\log n)+O(\log n)$	No
Schaub et al. [47]	Decentralized	Malicious	protects privacy	not provided	No
Bethencourt et al. [36]	Centralized	Malicious	depend trusted party	$O(1)$	No
Stefanos et al. [52]	Decentralized	Semi-honest	protects privacy	not provided	No
Clark et al. [51]	Decentralized	Semi-honest	protects privacy	not provided	No
PrivBox	Decentralized	Malicious/Semi-honest	protects privacy	$O(n)+O(n)$	Yes

Table 5: Comparison of the PrivBox system with other Centralized and Decentralized Reputation Systems. n is the number of users, and N is the number of preselected users for the privacy protection.

Entity	Computational overhead (number of exponentiations)			Communication overhead (bits)		
	Key	Encrypted Feedback	NIZK Proof	Key	Encrypted Feedback	NIZK Proof
User	n	n	$7n$	n	n	$8n$
Aggregator.	—	—	$8n * \eta$	$n * \eta$	$n * \eta$	$8n * \eta$

Table 6: Computation and Communication complexity for n retailers and η number of feedbacks.

Item	Computation Cost	Communication Cost
Setup	—	n
Key	—	$n\eta$
Feedback	—	$n\eta$
NIZKP	$8n\eta$ exponentiations	$8n\eta$

Table 7: Computation and Communication complexity for the Public Verification.

computation and communication complexities of the privacy-preserving reputation system for the client and the analyst are given in Table 6. Table 7 presents the computation time and communication overhead required for verifying the reputation statistics.

Generating the cryptograms of feedback score:

Each participant in the PrivBox system has to compute the cryptograms of the feedback score. An encrypted feedback is of the form $\langle c_{\tau_i r}, X_{\tau_i r} \rangle : i \in [\tau], r \in [n]$. $X_{\tau_i r} = g^{x_{\tau_i r}}$ is the public key and cipher text $c_{\tau_i r} \in \{R_{\tau_i r}^{x_{\tau_i r}}, R_{\tau_i r}^{x_{\tau_i r} g}\}$. The restructured key is $R_{\tau_i r} = \prod_{j=1}^{i-1} X_{\tau_j r} / \prod_{j=i+1}^{\eta} X_{\tau_j r} = \prod_{j=1}^{i-1} g^{x_{\tau_j r}} / \prod_{j=i+1}^{\eta} g^{x_{\tau_j r}} = g^{\sum_{j=1}^{i-1} x_{\tau_j r} - \sum_{j=i+1}^{\eta} x_{\tau_j r}}$. For computing the encrypted feedback for one retailer, the feedback provider needs to compute one exponentiation, and for n retailer the computational cost would be n exponentiations.

Computing NIZK Proof: The feedback provider also has to compute the NIZK proof to prove that his commitment is within the prescribed range without revealing the actual value of the feedback. The computation of the NIZK proof of the public key requires one exponentiation for one retailer. The computation of the NIZK proof of well-formedness of the feedback scores needs 7 exponentiations for one retailer. Hence, for the n retailers, the total computation cost for the NIZK proofs will be equal to $7n$ exponentiations. This is the most expensive operation among all the operations at the client side. Thus, the total time for generating the complete cryptograms for the

n retailers would be $7n$ exponentiations.

Aggregation: The aggregation process consists of three steps: checking the NIZK proof of well-formedness, the multiplication of cryptograms, and the brute force search to get the final aggregated score. The checking of well-formedness requires 8 exponentiations per feedback, thus for the n retailers and η participants, the total number of exponentiations required for checking the well-formedness is $8(n\eta)$. The multiplication of cryptograms is carried out with negligible cost and the brute force search is carried out $O(m)$ where m is the number of patterns to be matched.

6.2. Communication Complexity

The communication overhead depends on the size of data sent by the feedback provider to the bulletin board. The most expensive data unit in the PrivBox protocol is the NIZK proof of well-formedness, which consumes most of the communication bandwidth. An encrypted feedback is of the form $g^{xy}g^v$ where $v \in \{0,1\}$, and the feedback provider has to generate n cryptograms for n retailers. The NIZK proof for one retailer consists of 4 commitments, 2 responses, and 2 challenges. The entire communication cost for n retailers thus will be $8n \times (\text{Bytes required for one commitment})$. For the aggregation, the analyst has to download all the data i.e. $8n\eta$, where η is a number of participants providing the feedback to the bulletin board.




Figure 4: Computation time required for generating the restructured key at the user side.

6.3. Efficiency

We present experimental results of our protocol in terms of generating the restructured key, creating the cryptogram of the feedback values and its associated zero-knowledge proof, checking the well-formedness of NIZK proof, and performing the feedback aggregation. We obtain estimates for the standard NIST Curve P-256 [65]. We performed experiments for 10 times and presented the average of observations. In the first phase, while computing the restructured key, we varied the number of participants who posted their public keys from 100 to 20000. The time required for generating the restructured key increases linearly with the number of feedback providers posting their public keys as shown in Figure 4. Figure 5.A presents the computation time required by a user client to present the cryptogram of his feedback score about the retailers. The number of retailers varies from 100 to 1000, and the number of users is fixed at 10000. The result reveals that the computation time is not very high even for a high number of retailers i.e. 1000. In the experiment, we have not included the time consumed while issuing the token to the user. We consider the time when the user has received the token from the marketplace and is ready to submit the rating score. Each user requires around less than a second to present the feedback to the bulletin board. Figure 5.B represents the computation time required by the analyst for checking the NIZK proof and aggregating the encrypted response from the bulletin board. The expensive operation is NIZK that consumes much of the system resources, i.e. around 700 seconds when the number of feedbacks on the bulletin board is 100K. However, this computation cost can be minimized by using the multiple cores in parallel. On I-7 system with 8 cores, this time reduces to around less than 100 seconds.

In terms of communication cost, the cryptogram of the feedback score and NIZK proof would require the bandwidth of less than one megabyte as shown in Figure 6.A.

Specifically, for 1000 retailers, the bandwidth required is around 800 Kb which is acceptable. The experiments reveal that communication overhead for the user increases linearly with the number of retailers (Figure 6.A), and the communication overhead for the aggregator also increases linearly with the number of cryptograms present on the bulletin board as shown in Figure 6.B.

7. Defenses against other Attacks

In this section, we discuss the system’s defense against other attacks.

7.1. Self-Promotion

In a self-promotion attack, the seller/retailer wants to falsely increase his overall reputation score. In our system, such an attack is possible when the retailer issues a large number of fake tokens to a large number of users that are actually controlled by the retailer. These fake users then submit positive ratings for the retailer. This attack could be mitigated by imposing some cost on the number of tokens the marketplace can issue to the service provider or retailer, or imposing limit on the number of tokens issued to the retailer.

7.2. Reputation Whitewashing

In a whitewashing attack, the non-reputed retailer discards his identity, and re-joins the reputation system with a new identity and develops its reputation from the scratch. The attack is feasible under the conditions when the reputation system does not impose a reasonable cost on the participants of the reputation system. The PrivBox reputation system can be joined and left by any participant. The effect of the whitewashing can be minimized in two ways: first, imposing a certain cost for the new retailer or seller identity, and second, linking the retailer identity to the physical business address or the web-page address.

7.3. Bad Mouthing

In a bad-mouthing attack, retailers or users collude to lower the reputation of a certain retailer. In this attack, the users would internationally issue negative recommendations about the particular retailer, resulting in a negative aggregate reputation of the retailer. The PrivBox system requires the token and well-formedness of feedback to limit the feedback provider to provide a value within the prescribed limits. The token is only issued to a user who has interacted with the retailer. Thus the user who wants to attack the system with bad mouthing has to acquire the token to submit his negative feedback.

8. Other Application Domains

In this section, we discuss other application domains where the PrivBox system can be used as a reputation system.

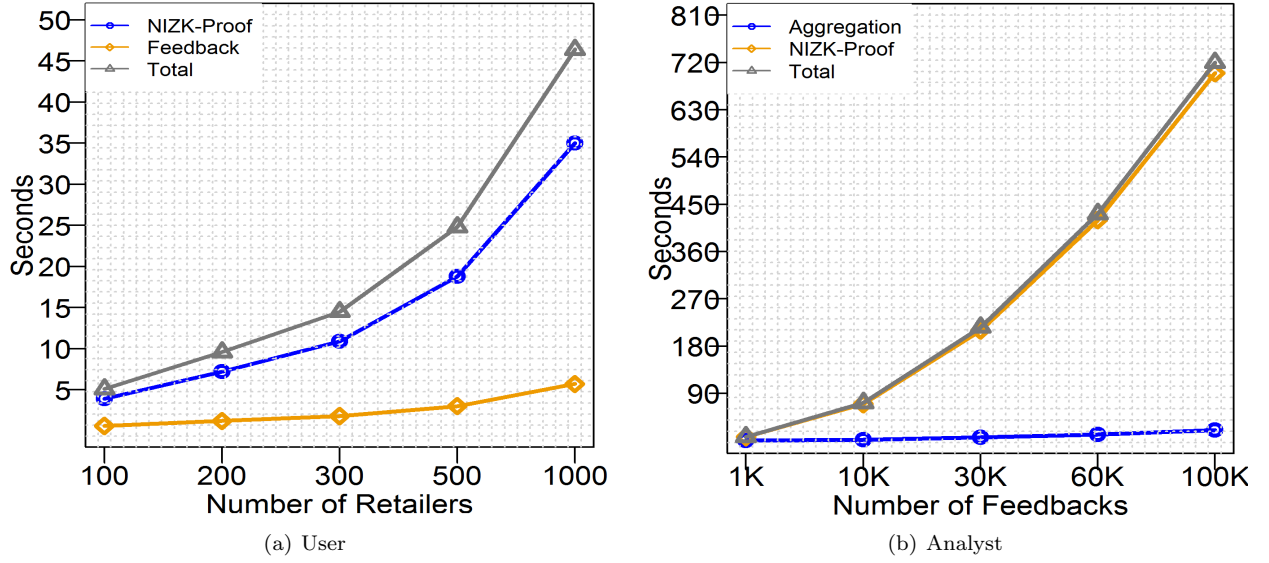


Figure 5: Computation time for the User and Analyst. A) User computation time for encrypted feedback and NIZK proof, and B) Analyst computation time for checking NIZK-proof and aggregating feedback.

8.1. Application to Edge Computing

The Edge computing paradigm allows users to perform computationally expensive tasks near their premises on behalf of the cloud service provider. In an edge computing system, an end user device is not the only one that asks for computational services from the edge node, but it can also act as the service provider by offering a computational platform as the edge node. The outsourcing of sensitive data to malicious edge nodes or offering computational services to malicious nodes could bring catastrophic consequences to edge providers and consumers. Therefore, participants – either the edge devices or the consumers have to assess the trustworthiness of computing platform before providing and accessing services. This is similar to evaluating the trustworthiness of nodes in a P2P network before downloading the content from the specific nodes or peers [66, 67, 68]. In edge computing, the trustworthiness of the edge provider could be evaluated using a service level agreement, but these service level agreements are not considering how consumers of the service see the trustworthiness of the service provider. The reputation system could provide information about the trustworthiness of consumers and edge devices by collecting the feedback from other consumers and service providers.

The PrivBox system could provide the platform for computing the trustworthiness of edge providers and edge consumers in a decentralized privacy-preserving way. Similar to our e-commerce application, consumers in the edge-computing scenario are the users requesting computational resources from the edge nodes, and edge nodes are the service providers providing computational resources to consumers. The consumer rates the services of an edge node on the scale of 0 or 1 or 0-5 stars for a different set of features (for example, latency, effectiveness, overhead). The PrivBox system can be deployed in an edge computing

scenario in the following way. First, the consumer requests services from the edge node, and upon completion of services, the cloud marketplace handling the edge node would ask the consumer to rate his transaction with the edge node. In this setting, the cloud marketplace sends the questionnaire to the consumers (having 0 and 1 answer), and the consumer replies these queries to the bulletin board similar to the way discussed in Section 4. Once, the feedback responses have been published on the bulletin board, the consumer first checks the reputation of the nearby nodes using information from the bulletin board. The consumer chooses the most trustworthy edge node for the transaction based on the aggregated reputation of nodes and his selection criteria.

8.2. Collaborative Intrusion and Misbehaviour Detection

The traditional standalone intrusion detection systems (IDS) have been mostly deployed within the network. These systems monitor and analyze the Internet traffic logs at the single point. However, malicious users can circumvent the detection system by making stealthy attacks to a large number of operators. The standalone system does not have any information about behavior of the traffic originating from other networks thus does not show effective resistance against sophisticated attackers that make a slow rate or stealthy attack against a large number of networks simultaneously. It is estimated that 20% of the malicious IP sources can attack multiple networks [69]. Naturally, collaboration among network providers would significantly improve the detection of the attacker in a timely manner, show effective and early defense against zero-day and stealthy attackers. However, the collaborative system has the challenge of privacy-preservation as network operators are not willing to share private information of their users, because they are concerned about their user privacy and

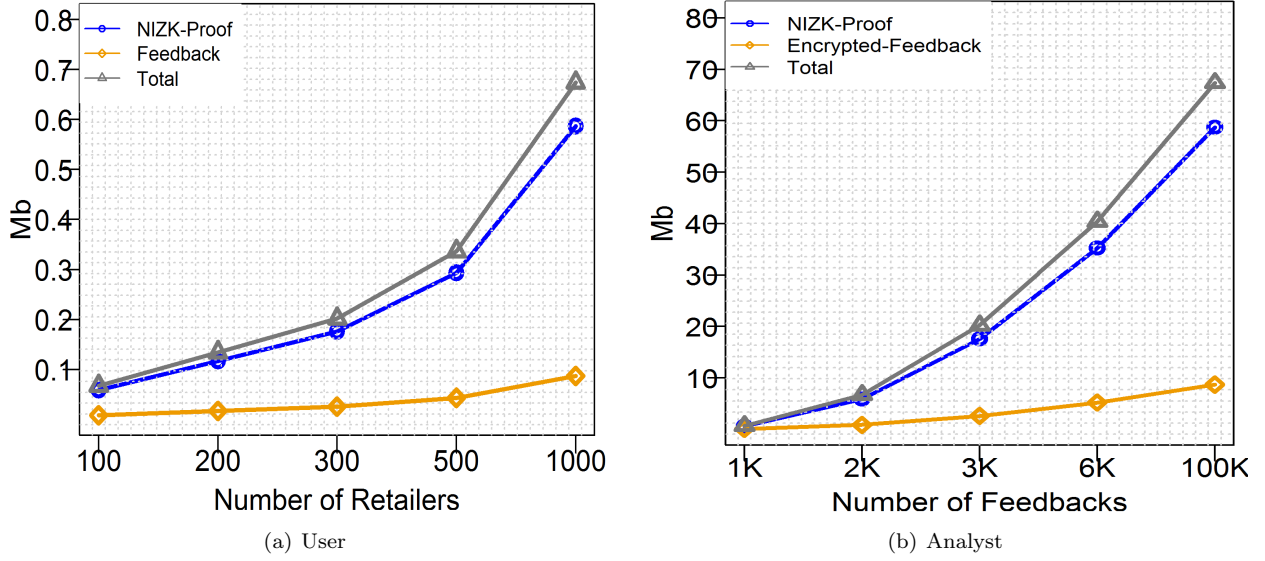


Figure 6: Bandwidth consumption for the user and analyst while generating an encrypted response and NIZK proof.

their own network configurations. However, network operators can be encouraged for the collaboration [70] if the privacy of their exchanged data is guaranteed. The proposed system can be used to ensure the privacy of collaborating operators using the encrypted exchange of feedback score for monitoring the Internet traffic in their networks. The implementation would ensure privacy with small communication and computation overheads. The proposed approach can also be used in a dshield (Internet storm center) setup for aggregating the feedback without anonymizing the identity of feedback providers.

9. Conclusion and Future works

Online marketplaces utilize reputation systems for evaluating the trustworthiness of their retailers based on the feedback submitted by their consumers. The reputation system can assist consumers in the marketplace to decide whether to have a transaction with the retailer or not. Existing reputation systems have serious privacy problems as these systems either require trusted centralized systems or a set of trusted peers to protect the private information of consumers. Further, marketplaces do not have an ability to prove that their reported statistics are correct and the results are not publicly verifiable. In this paper, we have presented a verifiable privacy-preserving reputation protocol that aggregates feedback values provided by consumers of the marketplace in a secure and private way. The protocol is designed to prevent participants from providing out-of-range high or low scores to their retailers. The protocol performs its operations in a decentralized way. The performance of the proposed system has been evaluated through the prototype implementation that demonstrates the effectiveness of the system in terms of computational and bandwidth overheads.

Acknowledgment

The research in this paper was supported by the ERC Starting Grant, No. 306994. We thank the anonymous reviewers for their insightful comments and suggestions.

References

- [1] "Number of digital buyers worldwide from 2014 to 2021 (in billions)," 2018. [Online]. Available: <https://goo.gl/ch44pW>
- [2] "Sales on online marketplaces cross \$1 trillion in 2016," 2017. [Online]. Available: <https://goo.gl/NvSVjv>
- [3] "Global ecommerce markets will reach \$4 trillion by 2020." 2017. [Online]. Available: <https://goo.gl/SXswbY>
- [4] "Overview of cloud market in 2017 and beyond," 2017. [Online]. Available: <https://goo.gl/zGx4zt>
- [5] H. Zhang, Y. Wang, and X. Zhang, "The approaches to contextual transaction trust computation in e-commerce environments," *Security and Communication Networks*, vol. 7, no. 9.
- [6] H. Zhang, Y. Wang, X. Zhang, and E.-P. Lim, "Reputationpro: The efficient approaches to contextual transaction trust computation in e-commerce environments," *ACM Transactions on Web*, vol. 9, no. 1, pp. 2:1–2:49, 2015.
- [7] "E-commerce fraud rates spike 33% in 2016," 2017. [Online]. Available: <https://goo.gl/YjFgnk>
- [8] "Fraud facts 2017—recognize, reject, report fraud," 2017. [Online]. Available: <http://www.competitionbureau.gc.ca>
- [9] "Online bargain hunters buy even if they think it may be a scam," 2017. [Online]. Available: <https://goo.gl/81PLgw>
- [10] "Watch out for online marketplace scams," 2018. [Online]. Available: <https://goo.gl/bStXXd>
- [11] L. De Alfaro, A. Kulshreshtha, I. Pye, and B. T. Adler, "Reputation systems for open collaboration," *ACM Communication*, vol. 54, no. 8, pp. 81–87, 2011.
- [12] A. Jøsang, R. Ismail, and C. Boyd, "A survey of trust and reputation systems for online service provision," *Elsevier Decision Support Systems*, vol. 43, no. 2, pp. 618–644, 2007.
- [13] F. Hendrikx, K. Bubendorfer, and R. Chard, "Reputation systems: A survey and taxonomy," *Journal of Parallel and Distributed Computing*, vol. 75, pp. 184 – 197, 2015.
- [14] S. Bag, M. A. Azad, and F. Hao, "A privacy-aware decentralized and personalized reputation system," *Computers & Security*, vol. 77, pp. 514 – 530, 2018.

- [15] T. Minkus and K. W. Ross, "I know what you're buying: Privacy breaches on ebay," in *Proceedings of 14th International Symposium Privacy Enhancing Technologies*, 2014, pp. 164–183.
- [16] P. Resnick and R. Zeckhauser, "Trust among strangers in internet transactions: Empirical analysis of eBay's reputation system," in *The Economics of the Internet and E-Commerce*, 2002, vol. 11, pp. 127 – 157.
- [17] S. Clauß, S. Schiffner, and F. Kerschbaum, "K-anonymous reputation," in *Proceedings of the 8th ACM SIGSAC Symposium on Information, Computer and Communications Security*, 2013, pp. 359–368.
- [18] E. Zhai, D. I. Wolinsky, R. Chen, E. Syta, C. Teng, and B. Ford, "Anonrep: Towards tracking-resistant anonymous reputation," in *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation*, 2016, pp. 583–596.
- [19] J. Blömer, J. Juhnke, and C. Kolb, "Anonymous and publicly linkable reputation systems," in *Proceedings of 19th Financial Cryptography and Data Security*, R. Böhme and T. Okamoto, Eds., pp. 478–488.
- [20] A. Narayanan and V. Shmatikov, "Robust de-anonymization of large sparse datasets," in *Proceedings of 29th IEEE Symposium on Security and Privacy (sp 2008)*, 2008, pp. 111–125.
- [21] —, "De-anonymizing social networks," in *Proceedings of the 30th IEEE Symposium on Security and Privacy*, 2009, pp. 173–187.
- [22] E. Gudes, N. Gal-Oz, and A. Grubshtein, "Methods for computing trust and reputation while preserving privacy," in *Proceedings of 23rd Annual IFIP WG 11.3 Working Conference Data and Applications Security*, 2009, pp. 291–298.
- [23] N. Gal-Oz, E. Gudes, and D. Hendler, "A robust and knot-aware trust-based reputation model," in *Proceedings of IFIPTM Conferences on Privacy, Trust Management and Security*, 2008, pp. 167–182.
- [24] O. Hasan, L. Brunie, E. Bertino, and N. Shang, "A decentralized privacy preserving reputation protocol for the malicious adversarial model," *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 6, pp. 949–962, 2013.
- [25] O. Hasan, L. Brunie, and E. Bertino, "Preserving privacy of feedback providers in decentralized reputation systems," *Computer & Secur.*, vol. 31, no. 7, Oct. 2012.
- [26] E. Vasilomanolakis, S. Karuppayah, M. Mühlhäuser, and M. Fischer, "Taxonomy and survey of collaborative intrusion detection," *ACM Computing Survey*, vol. 47, no. 4, pp. 55:1–55:33, May 2015.
- [27] M. Gil Pérez, F. Gómez Mármol, G. Martínez Pérez, and A. F. Skarmeta Gómez, "Repcidn: A reputation-based collaborative intrusion detection network to lessen the impact of malicious alarms," *Journal of Network and Systems Management*, vol. 21, no. 1, pp. 128–167, Mar 2013.
- [28] M. A. Azad and S. Bag, "Decentralized privacy-aware collaborative filtering of smart spammers in a telecommunication network," in *Proceedings of the Symposium on Applied Computing*, 2017, pp. 1711–1717.
- [29] M. A. Azad, S. Bag, S. Tabassum, and F. Hao, "privy: Privacy preserving collaboration across multiple service providers to combat telecoms spam," *IEEE Transactions on Emerging Topics in Computing*, vol. PP, no. 99, pp. 1–1, 2017.
- [30] A. Ramachandran, N. Feamster, and S. Vempala, "Filtering spam with behavioral blacklisting," in *Proceedings of the 14th ACM Conference on Computer and Communications Security*, pp. 342–351.
- [31] M. A. Azad and R. Morla, "Rapid detection of spammers through collaborative information sharing across multiple service providers," *Future Generation Computer Systems*, 2018.
- [32] I. E. Akkus, R. Chen, M. Hardt, P. Francis, and J. Gehrke, "Non-tracking web analytics," in *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, ser. CCS '12, 2012, pp. 687–698.
- [33] U. Erlingsson, V. Pihur, and A. Korolova, "Rappor: Randomized aggregatable privacy-preserving ordinal response," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 2014, pp. 1054–1067.
- [34] L. Liu and M. Munro, "Systematic analysis of centralized online reputation systems," *Elsevier Decision Support Systems*, vol. 52, no. 2, pp. 438 – 449, 2012.
- [35] M. Kinateter and S. Pearson, "A privacy-enhanced peer-to-peer reputation system," in *Proceedings of International Conference on Electronic Commerce and Web Technologies*, 2003, pp. 206–215.
- [36] J. Bethencourt, E. Shi, and D. Song, "Signatures of reputation: Towards trust without identity," in *Proceedings of the 14th International Conference on Financial Cryptography and Data Security*, 2010, pp. 400–407.
- [37] S. Schiffner, S. Clauß, and S. Steinbrecher, "Privacy and liveness for reputation systems," in *Proceedings of 6th European Workshop, EuroPKI 2009*, 2010, pp. 209–224.
- [38] M. A. Azad, S. Bag, and F. Hao, "M2m-rep: Reputation of machines in the internet of things," in *Proceedings of the 12th International Conference on Availability, Reliability and Security*, 2017, pp. 28:1–28:7.
- [39] K. Gai, M. Qiu, Z. Ming, H. Zhao, and L. Qiu, "Spoofing-jamming attack strategy using optimal power distributions in wireless smart grid networks," *IEEE Transactions on Smart Grid*, vol. 8, no. 5, pp. 2431–2439, 2017.
- [40] M. Qiu, K. Gai, and Z. Xiong, "Privacy-preserving wireless communications using bipartite matching in social big data," *Elsevier Future Generation Computer Systems*, 2017.
- [41] "Facebook suspends controversial data firm cambridge analytica." 2018. [Online]. Available: <https://goo.gl/FHDLGi>
- [42] "Your private medical data is for sale – and it's driving a business worth billion," 2017. [Online]. Available: <https://goo.gl/ovPLg9>
- [43] A. Singh and L. Liu, "Trustme: Anonymous management of trust relationships in decentralized p2p systems," in *Proceedings of the 3rd International Conference on Peer-to-Peer Computing*, ser. P2P '03, 2003.
- [44] L. Yibin, G. Keke, Q. Longfei, Q. Meikang, and Z. Hui, "Intelligent cryptography approach for secure distributed big data storage in cloud computing," *Elsevier Information Sciences*, vol. 387, pp. 103 – 115, 2017.
- [45] K. Gai and M. Qiu, "Blend arithmetic operations on tensor-based fully homomorphic encryption over real numbers," *IEEE Transactions on Industrial Informatics*, vol. PP, no. 99, pp. 1–1, 2017.
- [46] C. Gentry, "A fully homomorphic encryption scheme," Ph.D. dissertation, Stanford University, 2009, crypto.stanford.edu/craig.
- [47] A. Schaub, R. Bazin, O. Hasan, and L. Brunie, "A trustless privacy-preserving reputation system," in *Proceedings of 31st IFIP TC 11 International Conference ICT Systems Security and Privacy Protection*, 2016, pp. 398–411.
- [48] E. Pavlov, J. S. Rosenschein, and Z. Topol, "Supporting privacy in decentralized additive reputation systems," in *Proceedings of Second International Conference Trust Management*, C. Jensen, S. Poslad, and T. Dimitrakos, Eds., Berlin, Heidelberg, 2004, pp. 108–119.
- [49] N. Rishab and R. Karthik, "Fuzzy privacy preserving peer-to-peer reputation management," *IACR Cryptology ePrint Archive*, 2009.
- [50] K. Soska, A. Kwon, N. Christin, and S. Devadas, "Beaver: A decentralized anonymous marketplace with secure reputation," *IACR Cryptology ePrint Archive*, 2016.
- [51] M. R. Clark, K. Stewart, and K. M. Hopkinson, "Dynamic, privacy-preserving decentralized reputation systems," *IEEE Transactions on Mobile Computing*, vol. 16, no. 9, pp. 2506–2517, 2017.
- [52] O. Stefanos, T. Litos, and D. Zindros, "Trust is risk: A decentralized financial trust platform," *Proceedings of Financial Cryptography and Data Security 2017*, 2017.
- [53] A. Post, V. Shah, and A. Mislove, "Bazaar: Strengthening user reputations in online marketplaces," in *Proceedings of the 8th*

USENIX Conference on Networked Systems Design and Implementation, 2011, pp. 183–196.

- [54] L. Page, S. Brin, R. Motwani, and T. Winograd, “The pagerank citation ranking: Bringing order to the web.” Stanford InfoLab, Technical Report 1999-66, November 1999. [Online]. Available: <http://ilpubs.stanford.edu:8090/422/>
- [55] F. Hao, P. Y. A. Ryan, and P. Zielinski, “Anonymous voting by two-round public discussion,” *IET Information Security*, vol. 4, no. 2, pp. 62–67, June 2010.
- [56] P. Resnick and R. Zeckhauser, “Trust among strangers in internet transactions: Empirical analysis of eBay’s reputation system,” 2002, vol. 11.
- [57] B. E. Commerce, A. Jøsang, and R. Ismail, “The beta reputation system,” in *In Proceedings of the 15th Bled Electronic Commerce Conference*, 2002.
- [58] A. Fiat and A. Shamir, “How to prove yourself: Practical solutions to identification and signature problems,” in *Proceedings of Advances in Cryptology — CRYPTO’ 86*, 1987, pp. 186–194.
- [59] F. Kerschbaum, “A verifiable, centralized, coercion-free reputation system,” in *Proceedings of the 8th ACM Workshop on Privacy in the Electronic Society*, 2009, pp. 61–70.
- [60] L. Melis, G. Danezis, and E. Cristofaro, “Efficient private statistics with succinct sketches,” in *Proceedings of 23rd Network and Distributed System Security Symposium (NDSS)*, 2015.
- [61] J. Canny, “Collaborative filtering with privacy,” in *Proceedings 2002 IEEE Symposium on Security and Privacy*, 2002, pp. 45–57.
- [62] P. Y. A. Ryan, D. Bismark, J. Heather, S. Schneider, and Z. Xia, “Prêt à voter: A voter-verifiable voting system,” *Transaction on Information Security*, vol. 4, no. 4, pp. 662–673, Dec. 2009.
- [63] B. Adida, “Helios: Web-based open-audit voting,” in *Proceedings of the 17th Conference on Security Symposium*, 2008, pp. 335–348.
- [64] E. Androulaki, S. G. Choi, S. M. Bellovin, and T. Malkin, “Reputation systems for anonymous networks,” in *Proceedings of the 8th International Symposium on Privacy Enhancing Technologies*, ser. PETS ’08, 2008, pp. 202–218.
- [65] “Digital signature standard (dss), u.s. department of commerce/national institute of standards and technology.” 2017. [Online]. Available: <https://goo.gl/iiUpFb>
- [66] X. Li and L. Ling, “Peertrust: supporting reputation-based trust for peer-to-peer electronic communities,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 16, no. 7, pp. 843–857, 2004.
- [67] R. Zhou and K. Hwang, “Powertrust: A robust and scalable reputation system for trusted peer-to-peer computing,” *IEEE Transactions on Parallel and Distributed Systems.*, vol. 18, no. 4, pp. 460–473, 2007.
- [68] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, “The eigentrust algorithm for reputation management in p2p networks,” in *Proceedings of the 12th International Conference on World Wide Web*, 2003, pp. 640–651.
- [69] S. Katti, B. Krishnamurthy, and D. Katabi, “Collaborating against common enemies,” in *Proceedings of the 5th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC ’05, 2005, pp. 34–34.
- [70] Y. Chen, K. Hwang, and W. S. Ku., “Collaborative detection of ddos attacks over multiple network domains,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 18, no. 12, pp. 1649–1662, 2007.