

Private Web Search with an Expected Constant Round

Myungsun Kim

Abstract—Web searching is becoming an essential activity because it is often the most effective and convenient way of finding information. However, a Web search can be a threat to the privacy of the searcher because the queries may reveal sensitive information about the searcher. Private Web search (PWS) solutions allow users to find information on the Internet while preserving their privacy. Here, privacy means maintaining the confidentiality of the identity of the communicating users. According to their underlying technology, existing PWS solutions can be divided into three types: proxy-based solutions, obfuscation-based solutions, and cryptography-based solutions. Of these, cryptography-based PWS (CB-PWS) systems are particularly interesting because they provide strong privacy guarantees in the cryptographic sense. In this paper, we present a round-efficient CB-PWS protocol that preserves computational efficiency compared to other known CB-PWS systems. Assuming a broadcast channel, our protocol is a *four-round* cryptographic scheme that requires $O(n)$ communication complexity. However, if only point-to-point interaction is available, with the users emulating the broadcast channel, our protocol requires an expected $O(1)$ -round complexity and the same computation and communication overhead. Further analyzing the efficiency of our protocol shows that our proposal requires only $3n$ modular exponentiations for n users. To evaluate the security of our protocol, we demonstrate that our construction is secure in terms of a semi-honest model. We then discuss how to enhance its security to render it secure in the presence of malicious adversaries. We provide a specific protocol for managing users' groups, which is also an advantage over existing systems.

Index Terms—Private Web search, secret sharing, homomorphic encryption, round efficiency



1 INTRODUCTION

Using a private Web search (PWS) prevents Web-search service providers such as Google and Bing from building user profiles while still allowing the users to enjoy full search functionality when performing Web searches. User profiling is usually defined as the process of implicitly developing a user profile from the search-engine queries submitted by the user. The service provider can then use user-profiling information, which might include the users' private interests and preferences, to assign the user to a predefined user class, such as a demographic or taste category, or to capture the online behavior of the user. Although user profiles may enable service providers to offer a better service, clearly this raises *privacy concerns* because sensitive information, such as a user's name and location, can be inferred from search-engine queries. In addition to the query terms themselves, other information, such as the source IP address and timestamp, may reveal sensitive information about the user.

Various approaches have been proposed to address this problem. Balsa et al. [1] identified three types of PWS protocols, in terms of the key techniques used for anonymity. First, if a PWS solution introduces a proxy server to submit query words on behalf of the user, it is called a proxy-based technique (e.g., [2]–[4]). A second group of PWS solutions enables users to submit a collection of queries in such a way that the real query term is buried among many other (fake) terms, thereby obscuring the identity of the real term. This is called obfuscation-based PWS (e.g., [5]–[7]). The third approach relies on cryptographic tools to eliminate

the possibility of linking users to their query terms. This is called cryptography-based PWS (CB-PWS) (e.g., [8]–[11]) and is used in our proposal.

In this paper, we are particularly interested in constructing an efficient CP-PWS protocol. Our choice relates to our technical standpoint on the trade-off between performance and security. Controversially, PWS solutions based on proxy or obfuscation techniques place more importance on performance than on security. Therefore, we can say that their goal is to find a way of enhancing security while not affecting the performance of legacy Web-searching services. On the other hand, CB-PWS solutions pursue strong security in spite of the high possibility of degrading search performance. Accordingly, PWS protocols of this type aim to find methods that minimize the performance degradation.

Our goal in this work is twofold.

First, we construct an efficient CB-PWS protocol with *constant-round* complexity, but with the same computation and communication costs as a standard protocol. When designing an interactive protocol, researchers will always investigate the round complexity of their method, because interaction over a network is usually affected by lagging or network congestion, becoming the most time-consuming aspect of the operation. We have many good examples of this type of secure protocol (e.g., [12]–[15]). Unfortunately, although CB-PWS solutions are typical examples of interaction-intensive protocols, round complexity has not been taken into account as a key efficiency metric to accompany computation and communication complexity. Therefore, it is very important to devise protocols that require a minimal number of rounds to complete. Our challenging goal is to find a CB-PWS solution with $O(1)$ -round complexity for a number of users.

• M. Kim is with the Department of Information Security, The University of Suwon, Hwasong-si, Gyeonggi-do 18323, South Korea. E-mail: msunkim@suwon.ac.kr

Second, we provide our efficient CB-PWS protocol with *simulation-based security*, focusing on the model of *malicious users*. We have observed that existing CB-PWS systems have not undergone rigorous proof of security, although ad-hoc security analysis has been performed in some cases. The only exception is Lindell and Waisbard’s scheme [9], but this has involved only game-based security. We believe that it is important to define precisely, and prove formally, the security that a cryptographic protocol offers, because the history of cryptographic-scheme design gives good reasons to distrust heuristic approaches.

Having stated the technical goals of this work and before giving a conceptual description of our PWS solution, we now summarize our contributions and present a high-level overview of our construction. For readers who would first like to review existing work in the PWS field to better understand where our work fits in, Section 6 surveys the PWS state of the art. However, in explaining our work, we will sometimes also give a brief survey of what is known in the literature about the relevant PWS issues.

1.1 Our Contributions

The technical contributions of this paper can be summarized as:

- We design a constant-round CB-PWS solution with comparable computation and communication complexity to existing systems, assuming the existence of a broadcast channel. Without the broadcast, our protocol expects to involve constant rounds for a number of users while preserving the computation and communication cost. Kang et al. [16] argued that their CB-PWS protocol has a constant number of rounds to complete. However, their work does not consider a network where only point-to-point communication is allowed. Furthermore, their scheme does not provide a specific protocol to set up a group of users, which should be a mandatory part of any CB-PWS protocol. In fact, no CB-PWS solutions other than that of Lindell and Waisbard have handled this issue. These CB-PWS schemes do not provide a way to determine a group leader among n users, but it is well known that group-leader election is a nontrivial issue, particularly with the constant-round restriction. Therefore, it is not clear if Kang et al.’s scheme will complete the protocol activities in a constant number of rounds. Kim and Kim [11] proposed a round-efficient CB-PWS scheme, but it significantly restricts the message-space size (at most 5 bits for a group of 20 users). Therefore, it does not lead to a preferable solution to the problem.
- We provide a so-called formal proof of security, which uses a simulation-based technique. It is clearly preferable to prove security using a standard simulation-based technique with a universally composable model. Differently from [9], we define an ideal functionality to fit the PWS model and show that our protocol can be transformed efficiently into the ideal functionality.

- We provide a specific protocol for efficiently creating a group of n users. In addition, we solve the inefficiency problem in Kang et al.’s work. Even if Kang et al. resolved the incompleteness issues above, their scheme is far from efficient in the sense of computation complexity. This is because the group manager has $O(n^3 \log^2 n)$ computation complexity in the number of users n . Therefore, their scheme can run efficiently only for a very small number of users. (3 ~ 4). We develop a solution to these technically crucial problems without restricting the plaintext length, which indeed is the primary difference from Kang et al.’s results.

1.2 A Key Idea behind Our Scheme

We now describe the properties of our scheme, compared to existing CB-PWS solutions, from a design-philosophy viewpoint.

Techniques common to our protocol and existing protocols are, first, to encrypt users’ query terms via a proper encryption algorithm, and then to rerandomize and mix the ciphertexts to remove linkability between users and their query terms. Figure 1 shows these general protocol actions, together with the number of rounds required to complete subprotocols such as a shuffle protocol.

Protocol flow	Main operations	# of Rounds
$u_1 \quad \dots \quad u_n$ \downarrow $\bar{q}_1 = E(q_1) \quad \dots \quad \bar{q}_n = E(q_n)$	Encrypt	1
\downarrow <div style="border: 1px solid black; padding: 2px; display: inline-block;">Shuffle</div> (π, r_1, \dots, r_n) \downarrow $\{\hat{q}_1, \dots, \hat{q}_n\}$	Remask & Mix	n
\downarrow <div style="border: 1px solid black; padding: 2px; display: inline-block;">Group decryption</div> \downarrow $\{q_1, \dots, q_n\}$	Decrypt	1
\downarrow <div style="border: 1px solid black; padding: 2px; display: inline-block;">Web query</div> \downarrow $\{a_1, \dots, a_n\}$	Broadcast query result	1

Fig. 1. General CB-PWS design framework

Existing works and their basic framework. We are mainly interested in *Shuffle* phase because this step requires $O(n)$ rounds for a number of users n . In existing solutions, each user u_i first computes an encryption of its own query term q_i , denoted by \bar{q}_i , under an encryption algorithm E . Then all users join and run a shuffle protocol, taking as inputs a private permutation π over the set $\{1, 2, \dots, n\}$ and a set of fresh randomizers (r_1, r_2, \dots, r_n) . This shuffle protocol should be performed in a relay manner among all users because the sequential shuffles they employ can only achieve unlinkability in this manner. Therefore, a round complexity of $O(n)$ seems to be unavoidable.

Our idea. We utilize a similar design framework. However, our scheme does not require the users to engage in a sequential shuffle protocol. However, note that this does not mean that our scheme does not use the shuffle protocol at all. Similarly, our scheme also outputs a list of rerandomized and mixed ciphertexts during its execution, but without any interaction with other users, i.e., in a stand-alone mode. To achieve this, instead of using a query term q_i , we encrypt a share of q_i after letting its other shares be distributed to other users. This is the main difference between our protocol and other existing solutions. More specifically, consider a situation before invoking the shuffle protocol in both cases. In existing solutions, all users have the same list of ciphertexts $\{\bar{q}_1, \bar{q}_2, \dots, \bar{q}_n\}$ in clear. In contrast, users in our solution have a different list of ciphertexts at this point, i.e., $\{\bar{q}_{i,1}, \bar{q}_{i,2}, \dots, \bar{q}_{i,n}\}$ by denoting the j -th share of q_i as $q_{i,j}$ ($1 \leq j \leq n$). Following this basic technical property, our protocol does not require any interaction between neighbors and does not incur a number of rounds proportional to the number of users.

However, we do face a new problem. Because all users' query terms have been distributed in the form of shares and then encrypted, after performing group decryption, all shares should be presented to an algorithm that can reconstruct a mixed list of original query terms $\{q_{\pi(1)}, \dots, q_{\pi(n)}\}$. Because we cannot match the decrypted shares to their original messages, a naïve method involving trial and error runs the reconstruction algorithm $O(n^2)$ times. This incurs $O(n^3 \log^2 n)$ computational complexity in total, assuming that fast interpolation can be done by $O(n \log^2 n)$ multiplications. In this work, we develop a lightweight solution to resolve this computational problem at the cost of some slight damage in the plaintext domain.

The outline. The structure of this paper is as follows. Section 2 introduces basic definitions and cryptographic primitives: secret sharing and public-key homomorphic encryption. The system model for running our scheme is described in Section 3. Section 4 provides a detailed description of our construction, together with a full description of the performance and security analysis of our protocol. Section 5 deals with a key subprotocol by which a group can be constructed and a group manager can be elected. Section 6 contains a review of the relevant literature and we make some concluding remarks in Section 7.

2 DEFINITIONS AND BASICS

In this section, we review briefly the concepts and notations in cryptographic building blocks. We then give a definition of security for CB-PWS which will be used in proving formally that our proposal is secure in the presence of malicious adversaries.

Mathematical notation. If A is a probabilistic polynomial-time (PPT) machine, we use $a \leftarrow A$ to denote making A produce an output according to its internal randomness.

In particular, if U is a set, then $r \xleftarrow{\$} U$ is used to denote sampling from the uniform distribution on U . Letting $U^t[x]$ be the set of all polynomials of degree $0, \dots, t$ with coefficients from U , we denote by $f \xleftarrow{\$} U^t[x]$ a polynomial chosen independently and uniformly from $U^t[x]$.

A *negligible* function, denoted by $\text{negl}(\lambda)$, is a $\nu(\lambda)$ such that $\nu(\lambda) = o(\lambda^{-\kappa})$ for every fixed constant κ . For $n \in \mathbb{N}$, $[n]$ denotes the set $\{1, \dots, n\}$. For any integer x the length of the binary representation of x is denoted by $|x|$, but when the context is clear we also use $|X|$ to denote the cardinality of a set X .

Let $X = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $Y = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ be ensembles. Two ensembles X and Y are computationally indistinguishable, denoted by $X \stackrel{c}{=} Y$, if for every PPT algorithm D and all $\lambda \in \mathbb{N}$,

$$\left| \mathbb{P}[D(X_\lambda, 1^\lambda) = 1] - \mathbb{P}[D(Y_\lambda, 1^\lambda) = 1] \right| < \text{negl}(\lambda)$$

2.1 Threshold Secret Sharing

We assume that the readers are familiar with the notion of secret sharing and Shamir's implementation.

A secret sharing scheme (e.g., [17], [18]) consists typically of an algorithm for sharing a secret and an algorithm for reconstructing the shared secret. We denote the first as Sh and the second as Rc. Shamir's scheme is based on polynomial interpolation and involves n points on the Cartesian plane. Using these n points, a unique polynomial $f(x)$ over a finite field is guaranteed to exist such that $f(x) = y$ for each of the points given. For a concrete instantiation of Shamir's scheme, we need first to determine an appropriate field F for the subsequent modular arithmetic. For instance, we can take F as $\mathbb{Z}_{\tilde{q}}$ for a prime \tilde{q} . Given $f \xleftarrow{\$} \mathbb{Z}_{\tilde{q}}^{n-1}[x]$ and a secret $q \in \mathbb{Z}_{\tilde{q}}$, we use f whose constant term has been replaced by the q , in running algorithms Sh and Rc subsequently.

However, Shamir's secret sharing is insufficient for settings involving active corruption, where an adversary may corrupt users in an arbitrary way. In particular, ordinary secret sharing is not effective in settings that may require a secret to be secured for a long period of time. Indeed, standard secret sharing schemes are no longer secure after some number t of users have been corrupted. To avoid advanced threats which, given sufficient time, will successfully corrupt sufficient users to break the threshold that guarantees security, a proactive security model can be considered, such as Ostrovsky and Yung [19]. Alternatively, verifiable secret sharing (VSS) could be considered, to prevent a malicious dealer from distributing spurious shares.

Fortunately, there are three considerations that free us from imposing a heavy computation burden on the users to make up for the security weakness of ordinary secret sharing in a stronger security model:

- 1) Because it is not necessary to preserve the confidentiality of users' query terms for a long period, we do not need to consider proactive security.
- 2) Users distribute encryptions of shares of their own query words rather than the shares themselves. The more important thing is that the shares should be encoded into a specific form before being given to an encryption algorithm. Whether the decryption of encrypted share is well-formed can be checked later. If the decryption is not a well-formed share, then the decryption is simply discarded.
- 3) Our technique for encoding secret shares involves both additions and multiplications, but, because our

underlying encryption has only a group homomorphism, we cannot plug our encoding scheme into an existing VSS scheme.

2.2 Threshold Homomorphic Encryption

A public-key encryption scheme $\mathcal{E} = (\text{Kg}, \text{E}, \text{D})$ comprises the following algorithms:

- Kg is a randomized algorithm that takes a security parameter λ as input and outputs a secret key sk and a public key pk . pk defines a plaintext space \mathcal{M}_{pk} and a ciphertext space \mathcal{C}_{pk} .
- E is a randomized algorithm that takes pk and a plaintext $m \in \mathcal{M}_{pk}$ as input and outputs a ciphertext $c \in \mathcal{C}_{pk}$. Note that this process is usually randomized using a randomizer $r \in \mathcal{R}_{pk}$, denoted by $c = \text{E}_{pk}(m; r)$.
- D takes sk and $c \in \mathcal{C}_{pk}$ as input and outputs the plaintext m .

We say that an encryption scheme is *correct* if, for any $(pk, sk) \leftarrow \text{Kg}(1^\lambda)$ and any $m \in \mathcal{M}_{pk}$, $m = \text{D}_{sk}(\text{E}_{pk}(m))$.

We say that a public-key cryptosystem \mathcal{E} is *homomorphic* for the binary relations (\oplus, \otimes) if for all $(pk, sk) \leftarrow \text{Kg}(1^\lambda)$, given \mathcal{M}_{pk} and \mathcal{C}_{pk} , $(\mathcal{M}_{pk}, \oplus)$ forms a group and $(\mathcal{C}_{pk}, \otimes)$ forms a group. Further, for all $c_1, c_2 \in \mathcal{C}_{pk}$, $\text{D}_{sk}(c_1 \otimes c_2) = \text{D}_{sk}(c_1) \oplus \text{D}_{sk}(c_2)$.

Informally, when two ciphertexts are combined in a specific manner, the resulting ciphertext encodes the combination of the underlying plaintexts under a specific group operation, usually multiplication or addition. As a consequence, a cryptosystem's homomorphic property allows us to perform *rerandomization*: given a ciphertext c , anyone can create a different ciphertext \bar{c} that encodes the same plaintext as c . Therefore, given a group homomorphic cryptosystem \mathcal{E} , we can define the rerandomization algorithm as follows:

$$\text{Re}_{pk}(c; \gamma) := c \otimes \text{E}_{pk}(0; \gamma),$$

where 0 is an identity such that $\forall m \in \mathcal{M}_{pk}, m \oplus 0 = m$.

Security for homomorphic encryption. Because group homomorphic encryption (GHE) allows malleability in ciphertexts, we need only discuss security for homomorphic cryptosystems, the so-called *semantic security*.

Semantic security was first defined by Goldwasser and Micali [20]. Intuitively, a cryptosystem is said to be semantically secure if, given a ciphertext c , an adversary cannot determine any property of the underlying plaintext m . Specifically, an adversary cannot extract any semantic information about plaintext m from an encryption of m .

We say that a GHE scheme $\mathcal{E} = (\text{Kg}, \text{E}, \text{D})$ is *semantically secure* if, for all polynomial-time algorithms $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$,

$$\mathbb{P}_{b,r} \left[\begin{array}{l} (pk, sk) \leftarrow \text{Kg}(1^\lambda); \\ (m_0, m_1, s) \leftarrow \mathcal{A}_1(pk); \\ c_b \leftarrow \text{E}_{pk}(m_b; r); \\ b' \leftarrow \mathcal{A}_2(m_0, m_1, c_b, s) \end{array} \middle| b = b' \right] - \frac{1}{2} \leq \text{negl}(\lambda),$$

where $b, b' \in \{0, 1\}$ and s is the state information of \mathcal{A} .

Threshold GHE. We require a threshold group-homomorphic encryption scheme. This property is satisfied by most known homomorphic encryption

schemes, including Goldwasser-Micali [20], El Gamal [21], Paillier [22] and threshold Paillier [23]. As mentioned above, our scheme does not require a particular type of GHE scheme. However, because El Gamal's scheme is somewhat more efficient than Paillier's scheme in the computational sense, we review briefly the El Gamal encryption scheme and its threshold variant.

For large primes p and q such that $q|(p-1)$, let \mathbb{G}_q be the unique subgroup of \mathbb{Z}_p^\times of order q , and let g be a generator of \mathbb{G}_q . Because any element $1 \neq \beta \in \mathbb{G}_q$ generates the group, the discrete logarithm of $\alpha \in \mathbb{G}_q$ with respect to the base β is defined as usual. All computations in the remainder of this paper are modulo p unless otherwise noted.

The standard El Gamal encryption scheme is as follows:

- $\text{Kg}(1^\lambda)$ outputs a group description (\mathbb{G}_q, g, p, q) by taking the security parameter λ , then publicly opening $y = g^x$ and keeping x secret, where $x \xleftarrow{\$} \mathbb{Z}_q^\times$.
- $\text{E}_{pk}(m; r)$ outputs $c = (g^r, m \cdot y^r)$ with $r \xleftarrow{\$} \mathbb{Z}_q^\times$.
- $\text{D}_{sk}(c)$ first parses c into (α, β) and outputs $\beta \cdot \alpha^{-x}$.

We can easily verify that the El Gamal encryption scheme is multiplicatively homomorphic. Its rerandomization algorithm $\text{Re}_{pk}(c; \gamma) := c \otimes \text{E}_{pk}(1_{\mathbb{G}}; \gamma)$ is given by

$$c \otimes \text{E}_{pk}(1_{\mathbb{G}}; \gamma) = (\alpha \cdot g^\gamma, \beta \cdot 1_{\mathbb{G}} \cdot y^\gamma) = (g^{r+\gamma}, m \cdot y^{r+\gamma}),$$

where the \otimes operation means componentwise group multiplication and $1_{\mathbb{G}}$ is the identity element in \mathbb{G}_q .

Distributed key generation. Each participant chooses $x_i \xleftarrow{\$} \mathbb{Z}_q^\times$ and publishes $y_i = g^{x_i}$. The public key is $y = \prod_{i=1}^N y_i$, and the secret key is $x = \sum_{i=1}^N x_i$. This requires N multiplications, but their computational cost is negligible compared to exponentiations. Further broadcast round complexity is $O(1)$. We write this algorithm as $\text{TKg}(1^\lambda, N)$.

Distributed decryption. Given an encryption $c = (\alpha, \beta)$, each participant publishes their decryption share α^{x_i} . The plaintext can be derived by computing $\frac{\beta}{\prod_{i=1}^N \alpha^{x_i}}$. As in key generation, decryption can be performed in a constant number of rounds, requiring one exponentiation. We write this algorithm as $\text{TD}_{x_i}(c)$.

2.3 Message Space Compatibility

As described briefly in §1.2, our protocol requires that each share of a query be encrypted by a threshold GHE scheme. Accordingly, all shares from the Sh algorithm of a secret sharing scheme need to be in the plaintext domain of the encryption scheme.

Let F be an underlying field for Shamir's sharing scheme and let \mathcal{M} be the plaintext domain of the GHE scheme. To resolve this message compatibility issue, F needs to be embedded into \mathcal{M} . Specifically, our protocol requires that $|\mathcal{M}| \geq |F| + 2 \lceil \log n \rceil$, for the number of users n . In particular, letting $F = \mathbb{Z}_{\tilde{q}}$ and $\mathcal{M} = \mathbb{G}_q$ for two primes \tilde{q} and q , we will take \tilde{q} and q such that $\lceil \log q \rceil = \lceil \log \tilde{q} \rceil + 2 \lceil \log n \rceil$. We will discuss a technical reason for these parameter selections in Section 4.

2.4 Security Definition

We now present a security definition for our CB-PWS scheme. In principle, we follow the standard definition for secure multiparty computation [24, §7].

Simulation-based security. We define the ideal execution of a function Q on inputs (q_1, \dots, q_n) and security parameter λ as the outputs of the honest users and the adversary \mathcal{A} from the above ideal execution. More specifically, let $Q : (\{0, 1\}^*)^n \rightarrow (\{0, 1\}^*)^n$ be an n -ary functionality and let $Q_i(q_1, \dots, q_n)$ be the i -th element of $Q(q_1, \dots, q_n)$. For an index set $J = \{i_1, \dots, i_t\} \subset [n]$ such that $t < n$, let $I = [n] \setminus J$ and denote by $Q_J(q_1, \dots, q_n)$ the sequence $Q_{i_1}(q_1, \dots, q_n), \dots, Q_{i_t}(q_1, \dots, q_n)$. A pair (J, \mathcal{S}) where \mathcal{S} is a PPT algorithm, represents an ideal-model adversary. The ideal execution of Q under (J, \mathcal{S}) on input $\mathbf{q} = (q_1, \dots, q_n)$, denoted by $\text{IDEAL}_{Q, \mathcal{S}(\cdot), J}(\mathbf{q}, \lambda)$, is defined as the output pair of the honest users and the ideal-model adversary \mathcal{S} from the ideal execution. The output of the honest users is $Q_I(\mathbf{q}')$ and the input of \mathcal{S} is $(q_{j \in J}, J, Q_J(\mathbf{q}'))$, where $\mathbf{q}' := (q'_1, \dots, q'_n)$ such that q'_i is given by \mathcal{S} for $i \in J$ and $q'_i = q_i$ otherwise.

Let Π be an n -party protocol for computing Q , defined as above. In the real model, there is no trusted party and the users communicate directly with each other. The real-model adversary \mathcal{A} controls the corrupted users and therefore sends all messages of its choice in their place. Further, the adversary is not obliged to follow the specifications prescribed in the protocol Π . The real execution of Π on inputs \mathbf{q} and security parameter λ , denoted by $\text{REAL}_{\Pi, \mathcal{A}(\cdot), J}(\mathbf{q}, \lambda)$, is then defined as the output vector of the honest users and the real adversary \mathcal{A} from the real execution of Π .

Following the ideal-vs-real standard simulation technique, we also require that a secure protocol emulates the ideal execution in the real model where a trusted party does not exist. We consider only a *static* adversary, which is not allowed to corrupt a user during protocol execution. Implicitly, we assume that the adversary takes as input some auxiliary information.

We can now describe the security situation more formally.

Definition 1 (Security in the malicious model). Let Q and Π be defined as above. Protocol Π is said to securely compute Q if, for every PPT algorithm \mathcal{A} , there exists a PPT algorithm \mathcal{S} , such that for every $J \subset [n]$,

$$\{\text{IDEAL}_{Q, \mathcal{S}, J}(\mathbf{q}, \lambda)\}_{\mathbf{q}, \lambda} \stackrel{c}{=} \{\text{REAL}_{\Pi, \mathcal{A}, J}(\mathbf{q}, \lambda)\}_{\mathbf{q}, \lambda}$$

Hybrid execution. In our construction, we will use zero-knowledge proof (ZKP) protocols as subroutines. A standard technique for plugging a secure subprotocol Π_G computing a functionality \mathcal{G} into a protocol Π is to adopt a hybrid model. In this model, any execution of the protocol Π calling Π_G as a subprotocol requires that users interact with each other as in the real model, but access the ideal functionality of \mathcal{G} as in the ideal model. Specifically, when user u_i needs to send a message q_i to the trusted party, it begins to execute Π_G on input q_i instead. On the other hand, when the execution of Π_G ends with output a_i , user u_i continues with Π as if a_i were given by the trusted party. Then, by the composition theorem of [25], if Π_G securely

computes \mathcal{G} , then the output distribution of a protocol Π in a hybrid execution with \mathcal{G} is computationally indistinguishable from the output distribution of Π invoking Π_G . Therefore, we need only analyze the security of Π when working with \mathcal{G} in a hybrid model.

The ideal functionality of PWS. We now describe an ideal functionality of PWS, where each user's input is a search keyword q_i . Given a query q_i , when its query result corresponds to a_i , the functionality outputs the union of all a_i results to all users. If no results are found, u_i is given a null string \perp . More formally:

Definition 2. Let q_i be a query word of user u_i and a_i be a corresponding query result (without loss of generality, we assume that all query words have the same size). Then the ideal functionality \mathcal{F}_{PWS} is:

$$(q_1, \dots, q_n) \mapsto (\{a_1, \dots, a_n\}, \dots, \{a_1, \dots, a_n\})$$

In the following sections, we present details of our PWS construction for a protocol realizing \mathcal{F}_{PWS} in the presence of malicious adversaries.

3 SYSTEM MODEL

Participants. We work in a setting which involves three semi-honest entities: (1) the users, (2) the group manager, and (3) the Web search engine. More specifically, The *users* are the individuals who submit query terms to the search engine and who wish to prevent the search engine from building user profiles. We use u to denote a user. The role of the *group manager*, denoted by G , is to group users for execution of the protocol introduced above. The *Web search engine provider*, denoted by W , is the entity that provides a list of best-matching Web pages, usually accompanied by a short summary and/or parts of the document. Note that a search engine has no incentive to protect the users' privacy.

Communication channels. For our communication model, we assume that there is a broadcast channel whereby users send messages to all other users in a single round. However, sometimes the network may not have a broadcast channel, and we will need to emulate a broadcast channel using point-to-point communication.

Restrictions on adversary. As in the security model, we consider that an adversary is not allowed to break current encryption schemes that are computationally secure. Further, we assume that, for correctness, there is at least one honest user in a group. In contrast to [8], we allow collusion between two different participants in the protocol. However, our solution does not consider several desirable properties for efficiency reasons, namely abrupt termination, fairness, and guaranteed output delivery. The first of these properties has been relatively better studied than the remainder. For example, Garg et al. [26] constructed a two-round protocol for general computation. However, for protocols that consider fairness and guaranteed output delivery, our understanding of round complexity is still incomplete.

4 OUR SHUFFLE PROTOCOL FOR WEB QUERIES

In this section, we describe the concrete construction of our CP-PWS system. We begin by describing the model of

execution for our PWS construction and its participants, explaining the concepts of queries, shuffle, and cryptographic keys. We then give the description of our scheme. Finally, we discuss security and efficiency issues.

4.1 Our Basic CB-PWS Protocol

As mentioned above, the key idea behind CB-PWS solutions is for a group to submit a set of search words on behalf of individual users. Following this design philosophy, we describe how our protocol hides the link between users and their query terms. We begin by giving a high-level overview of our protocol.

Overview. Our PWS scheme, denoted by Π_{PWS} , is divided logically into three phases:

- 1) *Setup.* The main goal of the Setup phase is to create a group of users who would like to make searches via the search engine. In addition, all system parameters for Shamir's secret sharing scheme and a public-key encryption scheme will be published to all users in the group.
- 2) *Mixing query.* Upon completing this phase, all users will hold a reencrypted and permuted version of the distributed query terms.
- 3) *Submitting query.* The group manager receives a set of queries, without knowing who submitted which query. It then submits the queries to the search engine. Upon receiving a set of query results from the search engine, it broadcasts the result.

We provide an abstract description of our proposal in Figure 2, which summarizes all interactions between the entities (e.g., users and the group manager). From this, we can estimate approximately the complexity results.

4.1.1 Setup

Let n be the size of the group and let N be the number of group managers. For convenience, we simply assume that the group managers know n and that all users know where the group managers are and how to contact them. Further, all messages are assumed to be encoded automatically into the working group of a given encryption scheme.

Let $\mathcal{E} = (\text{TKg}, \text{E}, \text{TD}, \text{Re})$ be a semantically secure threshold GHE scheme, and let K be a public parameter specifying Shamir's secret sharing scheme. The setup phase then comprises three main activities:

- (1) The N group managers G_1, \dots, G_N collaboratively run the key generation $\text{TKg}(1^\lambda, N)$ and publish the system parameters, including the public key $pk = y$, and with the parameter K , within the restrictions specified in §2.3.
- (2) The N group managers elect a group-manager leader, denoted by $G \in \{G_1, \dots, G_N\}$, and publicize G to users as a representative of the group managers.
- (3) When the leading group manager G receives n requests for a private query, it responds to all n users indicating the group size of n . The group manager then constructs a group $\{u_1, u_2, \dots, u_n\}$ and publishes the group information, including the

group name, a list of participating users, and each user's label.

We defer the details of techniques for choosing the leader of the group managers and for group construction to Section 5. Note that these subprotocols require only a constant number of rounds.

4.1.2 Mixing Query

Let \mathcal{M}_{pk} be the plaintext domain of \mathcal{E} and let F_K be the working domain for Shamir's sharing scheme such that $F_K \subset \mathcal{M}_{pk}$. Let $q_i \in F_K$ be a query term from u_i . After receiving both the group information and the system parameters from G as a response to its query request, each user $u_{i \in [n]}$ performs six steps:

- (1) Chooses a set of random coefficients $\{r_{i,k}\}_{k=1}^{n-1}$ such that $r_{i,k} \in F_K$ and $|r_{i,k}| = \lceil \log \tilde{q} \rceil$, and determines

$$R_i(x) = \sum_{k=1}^{n-1} r_{i,k} x^k + q_i \in F_K^{n-1}[x]$$

- (2) Computes the shares of the query term q_i by evaluating R_i at each point $j \in [n]$ and sets $v_{i,j} = R_i(j) \in F_K$ for each $i, j \in [n]$.
- (3) Generates a random integer α_i such that $|\alpha_i| = 2 \lceil \log n \rceil$ and, for all $j \in [n]$, defines the shares by appending it to each of them as $v_{i,j} = v_{i,j} \parallel \alpha_i$. We then have $|v_{i,j}| = \lceil \log \tilde{q} \rceil + 2 \lceil \log n \rceil \leq \lceil \log q \rceil$.
- (4) Computes $\bar{v}_{i,j} = \text{E}_{pk}(v_{i,j})$ for each $j \in [n]$, and broadcasts a list $\langle i, j, \bar{v}_{i,j} \rangle_{j \in [n] \setminus \{i\}}$ to all other users.
- (5) Because u_i can build the array of 4-tuples:

$$\begin{bmatrix} \perp & \cdots & \langle 1, i, \bar{v}_{1,i} \rangle & \cdots & \langle 1, n, \bar{v}_{1,n} \rangle \\ \langle 2, 1, \bar{v}_{2,1} \rangle & \cdots & \langle 2, i, \bar{v}_{2,i} \rangle & \cdots & \langle 2, n, \bar{v}_{2,n} \rangle \\ & & \vdots & & \\ \langle n, 1, \bar{v}_{n,1} \rangle & \cdots & \langle n, i, \bar{v}_{n,i} \rangle & \cdots & \perp \end{bmatrix}$$

where \perp indicates that a 4-tuple with an encrypted share is unknown to the corresponding cell, it sets $\bar{v}_i = (\bar{v}_{1,i}, \bar{v}_{2,i}, \dots, \bar{v}_{n,i})$ and then computes a rerandomized and shuffled version of \bar{v}_i :

$$\hat{v}_i = (\hat{v}_{1,i}, \hat{v}_{2,i}, \dots, \hat{v}_{n,i}),$$

where, for each $\ell \in [n]$, $\hat{v}_{\ell,i} = \text{Re}_{pk}(\bar{v}_{\pi_i(\ell),i}; \gamma_\ell)$ with π_i is a random permutation on $[n]$ and $\gamma_{i,\ell} \in \mathbb{Z}_q^\times$.

- (6) Sends the new list \hat{v}_i to the group manager G .

4.1.3 Submitting Query

The group manager performs the following four steps:

- (1) Constructs an $n \times n$ matrix M by decrypting all of the received ciphertext vectors:

$$M = \begin{bmatrix} v_{\pi_1(1),1} & v_{\pi_1(2),1} & \cdots & v_{\pi_1(n),1} \\ v_{\pi_2(1),2} & v_{\pi_2(2),2} & \cdots & v_{\pi_2(n),2} \\ \vdots & \vdots & \vdots & \vdots \\ v_{\pi_n(1),n} & v_{\pi_n(2),n} & \cdots & v_{\pi_n(n),n} \end{bmatrix}$$

protocol Private Web Search

- Inputs: a list of query terms (q_1, q_2, \dots, q_n) from each user
- Auxiliary inputs: a security parameter and a group size n
- The protocol actions:
 - 1) *Setup*. This step consists of two major tasks as follows:
 - Group setup: A group of n users is created, and a group information is published.
 - Parameter selection: N group managers jointly generate parameters for threshold cryptographic primitives and publicize the parameters to the users.
 - 2) *Mixing query*. In this step, the users perform the following:
 - a) Splitting the query term q_i into n shares using Shamir's secret sharing scheme.
 - b) Encrypting the shares into a list of ciphertexts under the public key of group.
 - c) Broadcasting the list of ciphertexts.
 - d) Re-encrypting and mixing the received list of ciphertexts.
 - e) Sending the updated list of ciphertexts to the group managers.
 - 3) *Submitting query*. In this step, the group managers jointly recover all query terms by applying the Lagrange interpolation to the decrypted lists, and submit the recovered terms to the search engine on behalf of the group users.

Fig. 2. An algorithmic description of our protocol Π_{PWS}

- (2) Parses each $v_{\pi_i(j),i}$ into a share of a query term and a random padding, and collects only the n shares as

$$\mathbf{v}_i = \{v_{i,1}, v_{i,2}, \dots, v_{i,n}\},$$

where the random padding is the same for all shares.

- (3) Reconstructs the polynomial $R_i(x)$ of degree $n - 1$ using the Lagrange interpolant and n points in \mathbf{v}_i , for all $i \in [n]$, and recovers a list of query terms $Q = \{q_1, q_2, \dots, q_n\}$ from the fact that q_i is the constant term of $R_i(x)$ for each $i \in [n]$.
- (4) Submits the set Q to the search engine W and then broadcasts to the users a result set $\{a_1, \dots, a_n\}$ received from W .

In the following subsection, we analyze the performance of our construction in terms of its efficiency. Next, we analyze the security of our protocol by examining its various behaviors.

4.2 Performance

Our protocol Π_{PWS} is compared to other CB-PWS solutions in terms of three efficiency measures: computation, communication, and rounds. For this purpose, we first analyze the performance of our proposal. Then, the schemes proposed by Castellà-Roca et al. [8] and Lindell and Waisbard [9] are compared to our proposal. Because the scheme by Romero-Tris et al. [10] is identical to Castellà-Roca et al.'s scheme except for adding ZKPs and replacing the shuffling by a permutation network, we omit Romero-Tris et al.'s scheme from the comparison. Furthermore, we are unable to provide a fair comparison between the scheme proposed by Kim and Kim [11] and our scheme because their scheme restricts the query term size to $\frac{\log q}{n}$, i.e., $|q| \leq \frac{\log q}{n}$.

4.2.1 Computation & Communication Complexity

We first analyze the computation and communication costs for running the protocol Π_{PWS} . For a fair comparison, we assume that our construction also employs an El Gamal encryption scheme whose working group is \mathbb{G}_q of order q , which is the subgroup of \mathbb{Z}_p^\times for primes q, p . For example, p may be 1024 bits long while q is 512 bits long. We denote by $\exp(\ell)$ a modular exponentiation of an ℓ -bit integer and by $\text{mul}(\ell)$ a modular multiplication of two ℓ -bit integers. We note that, because Lindell and Waisbard's scheme uses a two-layer encryption where Cramer and Shoup's cryptosystem [27] takes as input an El Gamal ciphertext, most of their modular exponentiations should be carried out modulo a 2048-bit integer rather than a 1024-bit integer.

Computations. Table 1 presents a comparison of computations for our scheme and our competitors. For a fair comparison, we did not count the number of cryptographic operations against a malicious adversary, e.g., ZKPs.

TABLE 1
Comparisons of computation costs

	Modular Exponentiations	Modular Multiplications
Ours	$4n \exp(\lceil \log q \rceil)$	$(3n + O(n \log^2 n)) \text{mul}(\lceil \log q \rceil)$
[9]	$(n + 2) \exp(\lceil \log q \rceil) + 11n \exp(2 \lceil \log q \rceil)$	$(n + 1) \text{mul}(\lceil \log q \rceil) + 6n \text{mul}(2 \lceil \log q \rceil)$
[8]	$(3n + 2) \exp(\lceil \log q \rceil)$	$(3n + 1) \text{mul}(\lceil \log q \rceil)$

The 'Mixing Query' step in our scheme requires evaluation of a polynomial at many points. It is well known that a fast evaluation of a polynomial in $F_K[x]$, of degree less than n and for n points in F_K , can be performed using at most $O(n \log^2 n)$ operations in F_K . As a result, the computation complexity of our protocol is in total $O(n)$ modular exponentiations in addition to $O(n \log^2 n)$ modular multiplications.

Communications. In our scheme, each user first needs to send $(n - 1)$ El Gamal ciphertexts and then n El Gamal ciphertexts, where each El Gamal ciphertext comprises two $\lceil \log p \rceil$ -bit integers. Because of the n -layered encryption in the Lindell and Waisbard scheme, the ciphertext size may grow exponentially. To avoid this problem, the authors apply a hybrid encryption technique by introducing a secure block cipher that can process a κ -bit block in one operation. In their scheme, the i -th user outputs an i -layered Cramer-Shoup ciphertext whose size is $(i \cdot 4 \cdot 2^{\lceil \log p \rceil} + \kappa)$ bits. For n users, the transmission size becomes $\frac{n(n+1)}{2} \cdot 4 \cdot 2^{\lceil \log p \rceil} + n\kappa = 4n(n+1)\lceil \log p \rceil + n\kappa$ bits. Because each user holds n of these ciphertexts, the transmission requires $4n^2(n+1)\lceil \log p \rceil + n^2\kappa$ bits in total. Furthermore, each user sends $n\lceil \log p \rceil$ bits for decryption to other users.

TABLE 2
Comparisons of communication costs

	Transmissions (in bits)
Ours	$2n(2n - 1)\lceil \log p \rceil$
[9]	$4n^2(n + 1)\lceil \log p \rceil + n^2\kappa + n^2\lceil \log p \rceil$
[8]	$3n^2\lceil \log p \rceil$

In consequence, our scheme has $O(n^2\lceil \log p \rceil)$ communication complexity, whereas Lindell and Waisbard's scheme has $O(n^3\lceil \log p \rceil)$ complexity.

4.2.2 Round Complexity

We have yet to show that the total number of rounds for running our PWS solution is constant because we do not know the round number required to build a group of n users.

However, if building a group of users can be performed with a constant number of rounds, it is clear that our protocol will have a constant round complexity. The main part of our protocol comprises four rounds: (1) applying Shamir's secret sharing and encryption to a query term, and broadcasting the resulting list, (2) shuffling the resulting set and sending it to the group manager, (3) recovering a set of query terms and submitting them to a search engine, and (4) broadcasting the search results to the users. In contrast, other CB-PWS proposals have $O(n)$ -round complexity.

4.3 Security

We continue by arguing that our protocol Π_{PWS} is secure in terms of the semi-honest model.

Assume that there exists at least one honest group manager in $\{G_1, \dots, G_N\}$. Because all private inputs are encrypted by semantically secure encryption $E_{pk}(\cdot)$, no users can learn nontrivial information about other users' private query values during the protocol Π_{PWS} with any significant probability. Because a random permutation of an honest user is secret to the adversary and, in this security model, the adversary should follow the instructions of the protocol, the protocol leaks no nontrivial information, even for the case of a conspiracy between a group manager and a set of t corrupted users, where $t < n$.

If all group managers are corrupt, then the adversary could identify which query term has been submitted by

which user. This is because each share of a query term carries a random padding, and therefore decrypting an El Gamal ciphertext originating from a specific user would compromise the unlinkability of our protocol.

4.4 Upgrading to Malicious Security

A standard technique for ensuring the security of a cryptographic protocol is to show the achievability of simulation-based security of the protocol. To achieve our security goal against malicious entities, we will need to utilize cryptographic tools as subprotocols. The first tool is a ZKP of knowledge about a discrete logarithm (DL). There have been many efforts to construct a ZKP for the language of the nonzero exponent of g :

$$\mathcal{L}_{\text{zk-DL}} = \{(pk, h) \mid \exists x \text{ s.t. } h = g^x\}$$

We use a ZKP protocol designed for a signature scheme by Schnorr [28], denoted by $\text{zk-DL}_{pk}\{(x) \mid g^x\}$.

The second tool is a ZKP that, given a public ciphertext, a user knows the corresponding plaintext. We use standard techniques to design a ZKP protocol for the language:

$$\mathcal{L}_{\text{zk-PK}} = \{(pk, \bar{v}) \mid \exists v \in \mathbb{G}_q \text{ s.t. } \bar{v} = E_{pk}(v)\},$$

and we denote by $\text{zk-PK}_{pk}\{(v) \mid \bar{v} = E_{pk}(v)\}$ a ZKP from this subroutine. ZKP techniques for proving plaintext knowledge have been well studied elsewhere [29].

The final tool is a verifiable shuffle that enables proof of the correctness of a shuffle, where a shuffle of ciphertexts $\bar{v} = (\bar{v}_1, \dots, \bar{v}_n)$ is a new set of ciphertexts $\hat{v} = (\hat{v}_1, \dots, \hat{v}_n)$ with the same plaintexts in permuted order. We also use a ZKP protocol for the language of the shuffle of $2n$ El Gamal encryptions:

$$\mathcal{L}_{\text{zk-CS}} = \{(pk, \bar{v}, \hat{v}) \mid \exists (\pi, \gamma) \text{ s.t. } \forall i \in [n], \hat{v}_{\pi(i)} = \text{Re}_{pk}(\bar{v}_i; \gamma_i)\},$$

where π is a random permutation on $[n]$ and a random vector $\gamma = (\gamma_1, \dots, \gamma_n) \xleftarrow{\$} (\mathbb{Z}_q^\times)^n$. $\text{zk-CS}_{pk}\{(\pi, \gamma) \mid \hat{v}_{\pi(i)} = \text{Re}_{pk}(\bar{v}_i; \gamma_i) \wedge \bar{v}_i \in \bar{v} \wedge \hat{v}_{\pi(i)} \in \hat{v}\}$ is shorthand notation for a ZKP that a user knows a witness (π, γ) to the correctness of a shuffle. There is a considerable literature on this topic. For example, VS protocols by Groth [30] and Neff [31] are well known. Note that all values not enclosed by $()$ are assumed to be known to the verifier.

4.4.1 Our CB-PWS Protocol for the Malicious case

Using the above subprotocols, we can strengthen the basic protocol described in §4.1 so that the modified protocol is secure against a malicious adversary. We denote this extended version as Π_{PWS}^* . The remainder of this section deals with steps that should be modified for stronger security. To avoid repeating the basic descriptions, we will describe just the required modifications to the three phases:

Setup. We need to change only Step (1). Specifically, each group manager $G_{i \in [N]}$ chooses a random private key $x_i \xleftarrow{\$} \mathbb{Z}_q^\times$ and publishes its public-key share $y_i = g^{x_i}$ along with a ZKP of knowledge of y_i 's discrete logarithm using $\text{zk-DL}_{pk}\{(x_i) \mid g^{x_i}\}$.

Mixing Query. We make no changes in Steps (1) to (3) of this phase.

In Step (4), each user u_i computes $\bar{v}_{i,j}$ with proof of plaintext knowledge:

$$\text{zk-PK}_{pk}\{(v_{i,j})|\bar{v}_{i,j} = E_{pk}(v_{i,j})\},$$

for each $j \in [n]$ and sends it to all other users.

In Step (5), on receiving an encrypted share, each user verifies proofs of plaintext knowledge. If this check fails, the protocol is terminated. If all verifications of proofs from other users are valid, for a ciphertext vector $\bar{v}_i = (\bar{v}_{1,i}, \dots, \bar{v}_{n,i})$, each user computes a rerandomized and shuffled version of this vector, $\hat{v}_i = (\hat{v}_{1,i}, \dots, \hat{v}_{n,i})$ together with proofs of correct shuffle, for each $\ell \in [n]$:

$$\text{zk-CS}_{pk}\{(\pi_i, \gamma_i)|\hat{v}_{\pi_i(\ell),i} = \text{Re}_{pk}(\bar{v}_{\ell,i}; \gamma_{i,\ell})\},$$

using a vector of randomizers $\gamma_i = (\gamma_{i,1}, \dots, \gamma_{i,n})$ and a random permutation π_i of its own choice, where $\bar{v}_{\ell,i}$ is a component of \bar{v}_i and $\hat{v}_{\pi_i(\ell),i}$ is a component of \hat{v}_i . The user then sends it to the group manager G in Step (6).

Submitting Query. On receiving a vector of ciphertexts \hat{v}_i from a user u_i , the group manager verifies the proofs of correct shuffle. If this verification fails, it terminates the protocol. Only after verification from all users, it constructs the matrix M in Step (1).

In Step (2), the group manager receives a vector of n shares $v_i = (v_{i,1}, \dots, v_{i,n})$ such that all $v_{i,j}$'s have the same random padding α_i , without knowing which user submitted the vector v_i .

The remaining two steps are unchanged.

4.4.2 Security Proof for the Malicious Case

Before considering the proof of security, we observe that if all users and a group manager are honest, u_i outputs a set of query results $\{a_1, \dots, a_n\}$ with a probability greater than $\frac{1}{n^4}$. In this case, if a random padding $\alpha_i \in \{0, 1\}^{n^2}$ is unique in an execution of our protocol, then, because of the Lagrange interpolation, a unique polynomial is reconstructed from the shares with the same α_i and a query term q_i is correctly recovered. However, for the case of $\alpha_i = \alpha_{j \neq i}$, the group manager drops both query terms q_i, q_j . Therefore, the query term q_i cannot be recovered with probability $\mathbb{P}[\alpha_i = \alpha_j] \leq \frac{1}{n^4}$. For example, for a modest group size of $n = 30$, the probability is $\frac{1}{30^4} \approx 2^{-20}$.

We now prove that our protocol Π_{PWS}^* satisfies the security definition given in Definition 1.

Theorem 1. Assume that at least one group manager is honest, and that zk-DL, zk-PK, and zk-CS are the ZKPs for $\mathcal{L}_{\text{zk-DL}}, \mathcal{L}_{\text{zk-PK}},$ and $\mathcal{L}_{\text{zk-CS}}$, respectively, and (TKg, E, TD, Re) is the threshold El Gamal encryption scheme with semantic security. Then our protocol Π_{PWS}^* for any coalition J of colluding users such that $|J| < n$ securely computes \mathcal{F}_{PWS} in the presence of malicious adversaries, assuming a broadcast channel.

Proof: We prove the security by constructing an algorithm \mathcal{S} , called a simulator, for an adversary in the ideal model. The simulator works within the ideal model, but it interacts with the corrupted users in J without \mathcal{A} detecting that it is not within the real model. The trusted party takes the input from \mathcal{S} and the honest users, and gives the query result set \mathcal{S} to the honest users. The simulator \mathcal{S} then interacts with the malicious users in J , pretending to be one or

more honest users, enabling them to know the query result set. For the trivial case of $J = \emptyset$, we have demonstrated that the output is correct. Our proof is in a hybrid model where a trusted party runs the ZKPs of knowledge for $\mathcal{L}_{\text{zk-DL}}, \mathcal{L}_{\text{zk-PK}},$ and $\mathcal{L}_{\text{zk-CS}}$.

The simulator \mathcal{S} proceeds as follows:

- 1) \mathcal{S} performs the *Setup* phase as follows:
 - a) For each honest group manager G , \mathcal{S} chooses a uniformly random $x_G \in \mathbb{Z}_q^\times$. \mathcal{S} sends $y_G = g^{x_G}$ to \mathcal{A} and emulates the ideal functionality of $\mathcal{L}_{\text{zk-DL}}$ by sending x_G to \mathcal{A} .
 - b) For each malicious group manager \tilde{G} , \mathcal{S} receives from \mathcal{A} , $(y_{\tilde{G}}, x_{\tilde{G}})$ for the ideal functionality of $\mathcal{L}_{\text{zk-DL}}$ and records $x_{\tilde{G}}$ only when $y_{\tilde{G}} = g^{x_{\tilde{G}}}$. Otherwise, it aborts, sending an error message to the trusted party for \mathcal{F}_{PWS} .
- 2) Let $I = [n] \setminus J$. For each simulated honest user $u_i \in I$, the simulator \mathcal{S} :
 - a) chooses a random polynomial $R_i(x)$ such that all coefficients are in F_K and its constant term is a query word q_i of its choice.
 - b) chooses random paddings α_i and constructs encryptions of shares with them.
- 3) To perform the *Mixing Query* phase of the protocol, the simulator \mathcal{S} :
 - a) sends the encrypted shares to all malicious users in J and simulates the ideal functionality of $\mathcal{L}_{\text{zk-PK}}$ by sending each share.
 - b) receives from \mathcal{A} a list of encrypted shares for each malicious user $u_{\xi \in J}$.
 - c) receives from \mathcal{A} its input $V_{\xi} = (v_{\xi,j})_{j \in [n] \setminus \{\xi\}}$ for the ideal implementation $\mathcal{L}_{\text{zk-PK}}$ for each malicious user $u_{\xi \in J}$. If \mathcal{A} fails to prove its correctness, then \mathcal{S} sends an error message to the trusted party \mathcal{F}_{PWS} and aborts.
 - d) emulates the ideal functionality of $\mathcal{L}_{\text{zk-CS}}$ as follows: it constructs an $n \times n$ matrix M_i for each honest user u_i in I and sets its i -th column vector to \bar{v}_i . For all $i \in I$, \mathcal{S} chooses a random vector $\gamma_i = (\gamma_{i,1}, \dots, \gamma_{i,n}) \stackrel{\$}{\leftarrow} (\mathbb{Z}_q^\times)^n$ and a random permutation π_i over $[n]$ and computes a vector \hat{v}_i . \mathcal{S} sends \hat{v}_i to \mathcal{A} and emulates the ideal functionality of $\mathcal{L}_{\text{zk-CS}}$ by sending the vector γ_i and the permutation π_i .
 - e) receives from \mathcal{A} a vector of mixed encryptions \hat{v}_{ξ} for each malicious user $u_{\xi \in J}$.
 - f) receives from \mathcal{A} its input $(\pi_{\xi}, \gamma_{\xi})$ for the ideal functionality of $\mathcal{L}_{\text{zk-CS}}$ for each malicious user u_{ξ} . \mathcal{S} communicates with \mathcal{A} to check whether its proof is valid or not; if not, then \mathcal{S} sends an error message to \mathcal{F}_{PWS} and aborts.
- 4) To perform the *Submitting Query* phase, \mathcal{S} emulates a group manager G :
 - a) \mathcal{S} sends each vector $\hat{v}_{i \in I}$ for all honest users and $\hat{v}_{\xi \in J}$ from \mathcal{A} to the group manager.

- b) For each $\xi \in J$, \mathcal{S} extracts $\{v_{j,\xi}\}_{j \in [n]}$ from the encryption vector \hat{v}_ξ using (\cdot, γ_ξ) in Step 2-f.
- c) Using V_ξ in Step 2-c and (π_ξ, \cdot) in Step 2-f, \mathcal{S} constructs a vector $v_\xi = (v_{\xi,j})_{j \in [n]}$ and calculates q_ξ from the vector using Lagrange interpolation.
- d) \mathcal{S} sends a set of query terms $Q = \{q_\xi\}_{\xi \in J} \cup \{q_i\}_{i \in I}$ to the trusted party \mathcal{F}_{PWS} and receives as the answer a set $\{a_1, \dots, a_n\}$.

5) \mathcal{S} outputs whatever \mathcal{A} does.

We can see that the simulator runs in polynomial time, and the malicious users cannot identify that they are communicating with the simulator, which is working with the ideal model, rather than with other users in the real model. The correct answer is learned by all users for both real and ideal models. Therefore, we may conclude the theorem. \square

Limitations. To maintain security against malicious adversaries, we require that all participants prove the correctness of each protocol step. The above subprotocols can be proven correct by using only so-called Σ -protocols, which need just three rounds of interaction [32], [33]. Σ -protocols are not known to be zero-knowledge, but they enable efficient proofs of correctness for the protocol steps. Applying the Fiat-Shamir heuristic [34] enables the obtained proof to be zero-knowledge with a single message, but for the *random oracle* model.

If a broadcast channel is not available, we need to modify Theorem 1. First, the number of corrupted users t should be strictly bound. Because we are working in an authenticated setting, we are allowed to set $t < \frac{n}{2}$. We can then emulate the broadcast channel with *expected constant rounds* using point-to-point interaction. [35] The round overhead of our protocol in this setting will be dominated by this cost.

5 GROUP CONSTRUCTION

We have a variety of approaches to ensuring the establishment of an anonymous channel between users and search engines (e.g., the PIR technique [36] and Chaum's mixnet [37]). In principle, they can be applied to obtain a secure PWS solution. The most convincing reason for their applicability to PWS solutions is that they require users to engage in a group, enabling the users to hide their identity within the group. The larger the group, the greater is the anonymity available to users. Despite the importance of the group setup issue, many previous CB-PWS approaches assume a particular group or do not offer a specific group-setup protocol.

Unlike other approaches in the CB-PWS literature, we pay attention to developing a set of specific protocol activities for users to create a group of a fixed size and to elect a group manager for the group. In the remainder of this section, we present our group-setup protocol and analyze its security for the random oracle model. To start, we summarize the relevant abbreviations and notations in Table 3.

TABLE 3
Notation and abbreviations

n_u	the total number of users in the system
n	the group size
n_g	the number of groups, i.e., $n_g = \lfloor \frac{n_u}{n} \rfloor$
t	the number of corrupted users
$\{0, 1\}^*$	the set of all binary strings
$\{0, 1\}^\alpha$	the set of all binary strings of length α
$s_1 \parallel s_2$	the concatenation of strings $s_1, s_2 \in \{0, 1\}^*$
J	an index set of corrupted users in a group
I	an index set of honest users in a group
\mathcal{G}_j	the j^{th} group of size n in the system

$H_1 : \{0, 1\}^* \rightarrow \{0, 1\}^M$	
$H_2 : \{0, 1\}^{M \cdot (n_u)} \rightarrow \{0, 1\}^{n_u \cdot \log(n_u)}$	
$H_3 : \{0, 1\}^{\log(n_u) + n_u \cdot r} \rightarrow \{0, 1\}^{\log(n_g)}$	
} random oracles	

M : a positive integer of polynomial size for λ	
r : a random string in \mathbb{G}_q	

5.1 Group Setup

As for all other protocols in the literature, we assume the existence of an ideal public bulletin-board (PBB) functionality. In fact, this assumption is equivalent to the assumption of a broadcast channel. Informally the information is intentionally published and recorded permanently on the PBB (e.g., Web pages that serve as a PBB). Figure 3 shows our resulting new protocol for constructing a group of n users.

protocol Group setup

GOAL: n users build a new group \mathcal{G}_j in a 2-pass protocol

RESULT: a group \mathcal{G}_j of honest majority, i.e., $|J| < n/2$

1. *Protocol messages.*

$$u_i \rightarrow \text{PBB: } \langle s_i, x_i = H_1(\text{IP}_i, \text{id}_i, r_i) \rangle$$

$$u_i \leftarrow \text{PBB: } \langle j = H_3(y^{(i)}, s_1, \dots, s_{n_u}) \rangle$$

2. *Protocol actions.* The following steps are performed each time a user u_i requests. Assume that n_u users have been registered in the system.

- (a) u_i computes $x_i = H_1(\text{IP}_i, \text{id}_i, r_i)$ where r_i, s_i are randomizers from \mathbb{G}_q , IP_i is the IP address, and id_i is its ID registered in the system.
- (b) u_i sends (s_i, x_i) to PBB.
- (c) PBB computes $y = H_2(x_1, \dots, x_{n_u})$ and defines $y := y^{(1)} \parallel \dots \parallel y^{(n_u)}$ where $\forall i, |y^{(i)}| = \log n_u$.
- (d) PBB computes $j = H_3(y^{(i)}, s_1, \dots, s_{n_u})$ and assigns to \mathcal{G}_j the user u_i .
- (e) PBB sends the j to u_i .

Fig. 3. Group setup protocol

Our protocol is a variant of Lindell and Waisbard's group-setup protocol in [9, §5.2]. For several technical reasons, we need to develop a modified version rather than use their original protocol.

The primary reason for our modification is that our PWS protocol frequently uses a *broadcast channel* to control the round complexity. Because our group-setup protocol should be run before our main CB-PWS protocol, the group-setup protocol might appear to be run in isolation. However, we need to emulate the broadcast channel in a point-to-point network using a broadcast protocol. If we wish to realize

the broadcast protocol within a fixed number of rounds, we have to strictly restrict the number of corrupted users t .

Classical results (e.g., [38], [39]) show that achieving broadcast among n users incurs $\Omega(t)$ -round complexity only if the number of corrupted users t satisfies $t < n/3$. Fortunately, more recent results based on randomization techniques show that it is achievable in *expected* constant round time for $t < n/2$, i.e., in an honest-majority setting [40]–[42]. Unlike Lindell and Waisbard’s scheme, our protocol that relies on broadcasting does not allow the adversary to have an index set J such that $|J| \geq n/2$. In an authenticated setting like our protocol, restricting to a *computationally bounded* adversary allows a limit of $t < n$. Therefore, because of the different assumptions in the communication model, we need to devise a new group-setup protocol.

A second reason relates to practical matters. To justify our assumption that $|J| < n/2$ for any index set of corrupted users J , we need to guarantee the probability that the adversary can create a set J whose cardinality $|J| \geq n/2$ is very small, i.e., negligible in λ . However, our new protocol should not only consider the strong assumption described above, but also ensure that the impact of our modification on the probability of maliciously grouping should be minimal, even under this assumption. For this reason, a new random oracle H_3 is added to our protocol. Because the output of a random oracle is uniformly distributed, the additional random oracle H_3 makes it harder for the adversary to add members to a group, as it would like.

We now compute the probability that an adversary can maliciously generate a group \mathcal{G}_j so that $|\mathcal{G}_j|=n$, but half of its member users are corrupted (i.e., $n/2$ users are malicious), for some $j \in [n_g]$, where $n_g = \lfloor \frac{n_u}{n} \rfloor$. We mean by a “malicious” group that the group contains at least $\frac{n}{2}$ malicious users, with the number of honest users being less than $\frac{n}{2}$. This is clearly an undesirable case because ambiguity with respect to users’ query terms will drastically increase, particularly when the size of the group is not large.

The following theorem shows that the probability still remains very small with respect to the security parameter λ . We write the probability as $\mathbb{P}[\text{MalGrp}]$ in the theorem.

Lemma 1. Let n_u be the total number of registered users, and let t be the number of corrupted users. Let n be the number of users in a group. For every positive integer $n_u, n, t \in \mathbb{N}$, and for all PPT adversaries \mathcal{A} , the probability of a bad event is given by:

$$\mathbb{P}[\text{MalGrp}] \leq \frac{2^n}{n_g} \cdot \left(\frac{n_u - t}{t} \right)^{n/2} \cdot \left(\frac{t}{n_u} \right)^n,$$

where the number of groups $n_g = \lfloor \frac{n_u}{n} \rfloor$.

Proof: Let MalGrp denote the event that a group is constructed by the adversary in such a way that the group

contains $\frac{n}{2}$ corrupted users. We then have:

$$\begin{aligned} \mathbb{P}[\text{MalGrp}] &= \frac{1}{\binom{n_g}{1}} \cdot \frac{\binom{t}{n/2} \binom{n_u-t}{n/2}}{\binom{n_u}{n}} = \frac{1}{n_g} \cdot \frac{\binom{t}{\tilde{n}} \binom{n_u-t}{\tilde{n}}}{\binom{n_u}{n}} \\ &= \frac{1}{n_g} \cdot \frac{t!}{(t-\tilde{n})!\tilde{n}!} \cdot \frac{(n_u-t)!}{(n_u-t-\tilde{n})!\tilde{n}!} \cdot \frac{(n_u-n)!n!}{n_u!} \\ &= \frac{1}{n_g} \cdot \binom{n}{\tilde{n}} \cdot \frac{\prod_{i=0}^{\tilde{n}-1} (t-i) \prod_{i=0}^{\tilde{n}-1} (n_u-t-i)}{\prod_{i=0}^{n-1} (n_u-i)} \\ &= \frac{1}{n_g} \cdot \binom{n}{\tilde{n}} \cdot \frac{\prod_{i=0}^{\tilde{n}-1} (n_u-t-i)}{\prod_{i=\tilde{n}}^{n-1} (t-i)} \cdot \prod_{i=0}^{n-1} \binom{t-i}{n_u-i} \\ &\leq \frac{2^n}{n_g} \left(\frac{n_u-t}{t} \right)^{\tilde{n}} \left(\frac{t}{n_u} \right)^n \end{aligned}$$

The last inequality is derived from $\binom{k}{\varepsilon k} \leq 2^{kH(\varepsilon)}$, where $H(\varepsilon) = -\varepsilon \lg \varepsilon - (1-\varepsilon) \lg(1-\varepsilon)$ for $0 \leq \varepsilon \leq 1$, and we assume that $0 \lg 0 = 0$, for convenience. This completes the proof of the theorem. \square

For example, consider the case that $n_u = 10^6, t = 10^3$, and $n = 30$. For this case, we have the probability of a malicious group being smaller than $\frac{30 \cdot 10^9}{10^6} \cdot \left(\frac{10^6-10^3}{10^3} \right)^{15} \cdot \left(\frac{10^3}{10^6} \right)^{30}$, which approximates to 10^{-40} . It could be argued that this probability is insufficient to protect the users’ privacy. However, because the adversary would need to create a malicious group within seconds, we guess that the probability would be meaningful in practice.

5.2 Group Manager

The architectural view of our PWS protocol, described conceptually in Section 1.2, is depicted as in Figure 4. The figure can be considered as a snapshot following construction of a group \mathcal{G}_j of size n using the group-setup protocol specified in the previous section.

Because the problem we handled in the previous section was the matter of an algorithm to create a group of proper size, it is sufficient to find an algorithm and to analyze it. However, there remain two crucial issues to be addressed. One issue is the maintenance of a set of group managers, and the second is the election of a representative among them.

Concerning maintenance, our idea is to employ the functionality of a PBB. In a cryptographic sense, a PBB is equivalent to a broadcast channel, which has already been discussed in Section 5.1. Specifically, the PBB is used by users to announce their messages. That is, a message can be posted by any user and read by any other one. By saying that a message is “published”, we mean that the message appears on the PBB. Moreover, the published message can not be deleted or modified once posted. Next, a natural scenario would be that perhaps a dozen users would undertake voluntarily the role of group manager. They could construct a pool of group managers using the group construction protocol shown in Figure 3.




Fig. 4. Overall communication architecture with grouping

The second issue is the election of a group manager G_i from the pool $\{G_1, \dots, G_N\}$. Leader election plays an important role in the design of fault-tolerant applications. There is an extensive literature about this issue (e.g., [43]–[45]) and refer to reference [46] for a detailed survey of recent results.

Intuitively, a leader can operate as a central coordinator who enforces consistent behavior among users. However, in our setting, the leader of group managers does not matter if there exists at least one honest group manager. In other words, a malicious group manager would be allowed to be the leader. Therefore, we can opt for an auto-election mechanism to choose the group leader after the pool containing honest group managers has been constructed.

For the sake of security, we can use the cryptographic leader-election protocol designed by Katz and Koo [35]. Their scheme is built on a moderated VSS, which adds an extra entity, called a moderator, to a standard VSS. Their moderated VSS scheme in an authenticated setting (i.e., assuming a public-key infrastructure) requires only a constant number of rounds, tolerating $t_G < N/2$ malicious group managers [35, Corollary 2]. Specifically, the moderated VSS protocol requires thirteen rounds of interaction. In addition, they provide a two-phase constant-round protocol for leader election in the same setting, where the round complexity of each phase is determined by the underlying moderated VSS [35, Corollary 4].

Much simpler leader election protocols could be employed for efficiency reasons. For example, assuming a trusted leader, we could adapt an efficient variant of Toueg’s constant-round protocol [47].

5.3 Other Considerations

We now identify issues that may arise in deploying our protocol in a real scenario, discuss the privacy that our protocol covers, and address some other issues.

Group size. As discussed in Lemma 1, our main protocol provides privacy that depends on the group size n . That is, the larger the group, the stronger the privacy that users can obtain. However, considering the quality of the Web search service, we need to limit the size of groups according to the computing power of users’ machines and the acceptable level of latency. Our example implementation, which used the El Gamal encryption scheme over 1024-bit p and 512-bit q , showed that, for $34 \leq n \leq 36$, a group manager required approximately 1.02 seconds to submit a list of queries. We conducted this experiment on an iMac with a 3.4 GHz Core i5 CPU and 16 GB memory.

Static and Dynamic Groups. The discussion of this matter depends on whether a user or a group manager joins a group. For users, it might appear that frequently changing groups would achieve higher privacy. However, if a unique user identity such as the IP address is repeatedly used in different groups, the adversary might detect this and have a better chance of designating the user. Therefore, changing groups, in itself, could lead to loss of privacy.

The primary reason for maintaining multiple group managers is to distribute the key pair in the El Gamal encryption scheme. No subsets of group managers whose size is less than a fixed threshold can then learn user plaintexts. Both cases seem to imply that any solution involving dynamic groups might be unnecessary in our setting.

6 LITERATURE REVIEW

Balsa et al. divide existing PWS protocols into three classes according to the underlying technique used to obtain anonymous channels [1]. That is, if a PWS solution employs a proxy server to submit their query words on behalf of users, it is considered a proxy-based technique. However, the server machine can easily become a target of attackers and has single-point-of-the-failure limitations. A second class of PWS solutions enables users to submit a set of queries so that the real query is well buried within the set. Balsa et al. calls this obfuscation-based PWS. Its technical core is the method that embeds a real query into the set and the indistinguishability level achieved by the method. The third class, to which our PWS solution belongs, eliminates the possibility of linkability between users and their query terms by using cryptographic tools. We call these schemes cryptography-based PWS (CB-PWS), following Balsa et al.’s naming convention.

Proxy-based PWS. The first approach is through the use of an anonymous proxy (e.g., [2]–[4]). Users can expect that anonymizers will prohibit the creation of user profiles through query unlinkability. There are several options in this category, from simple mechanisms achieving a low level of anonymity in Web searches to more reliable but more complicated systems based on onion routing, [48] such as the Tor network. [49] However, the effectiveness of simple solutions is clearly limited. In addition, as highlighted by [8], Tor is not always easy to install and configure. Furthermore, it is well known that the HTTP requests over Tor can become very slow. [4] For example, it takes 10 seconds on average to submit a query to Google even when using paths of length 2 (the default length is 3).

Obfuscation-based PWS. Another approach to providing privacy during Web search is based on a query obfuscation technique (e.g., [5]–[7], [50]). The class of solutions using query obfuscation involves blending the real queries into a stream of fake queries so that Web search engines cannot create an accurate profile. From a privacy point of view, these obfuscation-based solutions have a critical drawback, namely, that automated queries have features that are different from real queries entered by a user, such as randomness. The authors in [51] demonstrated a classifier implementation that can distinguish real queries from fake queries generated by TrackMeNot, [50] with a mean misclassification rate of only approximately 0.02%.

Cryptography-based PWS. The third class of solutions involves using cryptographic algorithms such as public-key encryption and shuffle. One of the main advantages of CB-PWS systems over others is that they provide strong privacy guarantees. In addition, they are not affected by the misclassification issue and are generally faster than anonymizer-based solutions. To our knowledge, the known solutions can be found in [8]–[11]. Without loss of generality, because we may consider Romero-Tris et al.’s scheme as a malicious variant of Castellà-Roca et al.’s scheme, the differences between the two schemes do not affect the round complexity.

These known solutions utilize the basic idea that, after joining a small group, each user encrypts the search query and sends it to other group members. Then, according to a predefined order, each user provides a shuffled list of encrypted queries to his neighbor. The last user broadcasts the final shuffled version. After group decryption, each user obtains a set of queries, but cannot know who submitted which query. As a result, Web search engines cannot build accurate user profiles.

We note that the only approach that comes close to achieving our requirements in the restricted setting is the work by Kim et al. [11]. The authors proposed a round-efficient CB-PWS scheme based on the notion of decomposable encryption. However, this approach significantly restricts the length of plaintexts (e.g., to 3 or 4 bits) to be encrypted, which does not lead to practical solutions to the problem. We have provided a detailed evaluation and analysis of existing CB-PWS solutions (see Section 4.2).

To our knowledge, the known CB-PWS constructions either require $O(n)$ rounds, where n is the number of users [8]–[10] or, by significantly restricting the length of messages to be encrypted, do not lead to practical solutions. [11] Moreover, the solution in [11] requires Web search engines to implement and run the protocols. However, search-engine service providers have no incentive to implement costly protocols that they cannot profit from. As yet, there are no known constructions of practical constant-round CB-PWS systems.

7 SUMMARY AND ONGOING WORK

Web searches have been shown to be sensitive in many cases. Any information leaked from search histories can endanger user privacy. Search histories may contain health-related data and other personal information, including, but not restricted to, political or religious views and sexual orientation data. For example, Google provides signed-in users with personalized search results based on their search and navigation histories. Furthermore, users typing search queries in the Web interface are prompted with suggestions derived from their search history. To this end, Google tracks all Web searches performed by a signed-in user as well as the target Web pages clicked from the search result page.

In this work, we present a constant-round CB-PWS protocol for protecting users’ privacy when a broadcast channel is available. Our solution can be deployed easily in current systems because it does not require any changes on the service-provider side. However, further work is required:

- We are attempting to provide a full implementation of our solution to demonstrate its practicability.
- We should improve the performance of the group-manager aspect, particularly when the group membership is dynamic.

REFERENCES

- [1] E. Balsa, C. Troncoso, and C. Díaz, “OB-PWS: Obfuscation-based private web search,” in *IEEE Security & Privacy*, 2012, pp. 491–505. pages 1, 12
- [2] Anonymizer, “Anonymizer,” <http://www.anonymizer.com>, 2014. pages 1, 12
- [3] Scroogle, “Scroogle,” <http://scroogle.org>, 2014. pages 1, 12
- [4] F. Saint-Jean, A. Johnson, D. Boneh, and J. Feigenbaum, “Private web search,” in *WPES*, P. Ning and T. Yu, Eds., 2007, pp. 84–90. pages 1, 12
- [5] Y. Elovici, B. Shapira, and A. Meshiach, “Cluster-analysis attack against a private web solution (PRAW),” *Online Information Review*, vol. 30, no. 6, pp. 624–643, 2006. pages 1, 12
- [6] J. Domingo-Ferrer, A. Solanas, and J. Castellà-Roca, “ $h(k)$ -private information retrieval from privacy-uncooperative queryable databases,” *Online Information Review*, vol. 33, no. 4, pp. 720–744, 2009. pages 1, 12
- [7] D. Rebollo-Monedero and J. Forné, “Optimized query forgery for private information retrieval,” *IEEE Transactions on Information Theory*, vol. 56, no. 9, pp. 4631–4642, 2010. pages 1, 12
- [8] J. Castellà-Roca, A. Viejo, and J. Herrera-Joancomartí, “Preserving user’s privacy in web search engines,” *Computer Communications*, vol. 32, no. 13–14, pp. 1541–1551, 2009. pages 1, 5, 7, 8, 12, 13
- [9] Y. Lindell and E. Waisbard, “Private web search with malicious adversaries,” in *Privacy Enhancing Technologies*, ser. LNCS 6205, M. Atallah and N. Hopper, Eds., 2010, pp. 220–235. pages 1, 2, 7, 8, 10, 13
- [10] C. Romero-Tris, J. Castellà-Roca, and A. Viejo, “Multi-party private web search with untrusted partners,” in *SecureComm*, M. Rajarajan, F. Piper, H. Wang, and G. Kesidis, Eds., 2011, pp. 261–280. pages 1, 7, 13
- [11] M. Kim and J. Kim, “Privacy-preserving web search,” in *ICUFN*, 2012, pp. 480–481. pages 1, 2, 7, 13
- [12] J. Bar-Ilan and D. Beaver, “Non-cryptographic fault-tolerant computing in constant number of rounds of interaction,” in *PODC*, 1989, pp. 201–209. pages 1
- [13] D. Beaver, S. Micali, and P. Rogaway, “The round complexity of secure protocols,” in *STOC*, 1990, pp. 503–513. pages 1
- [14] C. Cachin, J. Camenisch, J. Kilian, and J. Müller, “One-round secure computation and secure autonomous mobile agents,” in *ICALP*, ser. LNCS 1853, U. Montanari, J. Rolim, and E. Welzl, Eds., 2000, pp. 512–523. pages 1
- [15] Y. Ishai and E. Kushilevitz, “Randomizing polynomials: A new representation with applications to round-efficient secure computation,” in *FOCS*, 2000, pp. 294–304. pages 1
- [16] B. Kang, S. C. Goh, and M. Kim, “Private web search with constant round efficiency,” in *ICISSP*, 2015, pp. 205–212. pages 2
- [17] A. Shamir, “How to share a secret,” *Communications of the ACM*, vol. 22, no. 11, pp. 612–613, 1979. pages 3
- [18] G. Blakley, “Safeguarding cryptographic keys,” *American Federation of Information Processing Societies Proceedings*, vol. 48, pp. 313–317, 1979. pages 3
- [19] R. Ostrovsky and M. Yung, “How to withstand mobile virus attacks,” in *PODC*, 1991, pp. 51–59. pages 3
- [20] S. Goldwasser and S. Micali, “Probabilistic encryption,” *J. Comput. Syst. Sci.*, 1984. pages 4
- [21] T. El Gamal, “A public key cryptosystem and a signature scheme based on discrete logarithms,” in *Advances in Cryptology-Crypto*, ser. LNCS 196, G. R. Blakley and D. Chaum, Eds., 1984, pp. 10–18. pages 4
- [22] P. Paillier, “Public-key cryptosystems based on composite degree residuosity classes,” in *Advances in Cryptology-Eurocrypt*, ser. LNCS 1592, J. Stern, Ed., 1999, pp. 223–238. pages 4
- [23] P.-A. Fouque, G. Poupard, and J. Stern, “Sharing decryption in the context of voting of lotteries,” in *Financial Cryptography*, ser. LNCS 1962, Y. Frankel, Ed., 2000, pp. 90–104. pages 4
- [24] O. Goldreich, *The foundations of cryptography: Volume 2—Basic Applications*. Cambridge University Press, 2004. pages 5

- [25] R. Canetti, "Security and composition of multiparty cryptographic protocols," *Journal of Cryptology*, vol. 13, no. 1, pp. 143–202, 2000. pages 5
- [26] S. Garg, C. Gentry, S. Halevi, and M. Raykova, "Two-round secure MPC from indistinguishability obfuscation," in *TCC*, 2014, pp. 74–94. pages 5
- [27] R. Cramer and V. Shoup, "A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack," in *Advances in Cryptology-Crypto*, ser. LNCS 1462, H. Krawczyk, Ed., 1998, pp. 13–25. pages 7
- [28] C.-P. Schnorr, "Efficient identification and signatures for smart cards," in *Advances in Cryptology-Crypto*, ser. LNCS 435, G. Brassard, Ed., 1989, pp. 239–252. pages 8
- [29] R. Cramer, I. Damgård, and J. B. Nielsen, "Multiparty computation from threshold homomorphic encryption," in *Advances in Cryptology-Eurocrypt*, ser. LNCS 2045, B. Pfitzmann, Ed., 2001, pp. 280–299. pages 8
- [30] J. Groth, "A verifiable secret shuffle of homomorphic encryptions," *J. of Cryptology*, vol. 23, pp. 546–579, 2010. pages 8
- [31] C. Neff, "A verifiable secret shuffle and its application to e-voting," in *ACM Conference on Computer and Communications Security*, 2001, pp. 116–125. pages 8
- [32] I. Damgård, "On Σ -protocols," Lecture Notes, Computer Science Department of University of Aarhus, 2002. pages 10
- [33] R. Cramer, I. Damgård, and B. Schoenmakers, "Proofs of partial knowledge and simplified design of witness hiding protocols," in *Advances in Cryptology-Crypto*, ser. LNCS 839, Y. Desmedt, Ed., 1994, pp. 174–187. pages 10
- [34] A. Fiat and A. Shamir, "How to prove yourself: Practical solutions to identification and signature problems," in *Advances in Cryptology-Crypto*, ser. LNCS 263, A. Odlyzko, Ed., 1986, pp. 186–194. pages 10
- [35] J. Katz and C.-Y. Koo, "On expected constant-round protocols for byzantine agreement," in *Advances in Cryptology-Crypto*, 2006, pp. 445–462. pages 10, 12
- [36] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan, "Private information retrieval," in *FOCS*, 1995, pp. 41–50. pages 10
- [37] D. Chaum, "Untraceable electronic mail, return addresses, and digital pseudonyms," *Communications of the ACM*, vol. 24, no. 2, 1981. pages 10
- [38] M. Pease, R. Shostak, and L. Lamport, "Reaching agreement in the presence of faults," *J. ACM*, vol. 27, no. 2, pp. 228–234, 1980. pages 11
- [39] M. Fischer and N. Lynch, "A lower bound for the time to assure interactive consistency," *Inf. Process. Lett.*, vol. 14, no. 4, pp. 183–186, 1982. pages 11
- [40] P. Feldman and S. Micali, "An optimal probabilistic protocol for synchronous Byzantine agreement," *SIAM J. Comput.*, vol. 26, no. 4, pp. 873–933, 1997. pages 11
- [41] M. Fitzi and J. A. Garay, "Efficient player-optimal protocols for strong and differential consensus," in *PODC*, 2003, pp. 211–220. pages 11
- [42] J. Katz and C.-Y. Koo, "Round-efficient secure computation in point-to-point networks," in *Advances in Cryptology-Eurocrypt*, ser. LNCS 4515, M. Naor, Ed., 2007, pp. 311–328. pages 11
- [43] G. Singh, "Leader election in complete networks," *SIAM Journal on Computing*, vol. 26, no. 3, pp. 772–785, 1997. pages 12
- [44] N. Malpani, J. Welch, and N. Vaidya, "Leader election algorithms for mobile ad hoc networks," in *DIAL-M*, 2000, pp. 96–103. pages 12
- [45] N. Schiper and S. Toueg, "A robust and lightweight stable leader election service for dynamic systems," in *DSN*, 2008, pp. 207–216. pages 12
- [46] A. Iványi, "Leader election in synchronous networks," *Acta Univ. Sapientiae, Mathematica*, vol. 5, no. 1, pp. 54–82, 2013. pages 12
- [47] S. Toueg, "Randomized byzantine agreements," in *PODC*, 1984, pp. 163–178. pages 12
- [48] M. Reed, P. Syverson, and D. Goldschlag, "Anonymous connections and onion routing," *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 4, pp. 482–494, 1998. pages 12
- [49] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *USENIX Security Symposium*, M. Blaze, Ed., 2004, pp. 303–320. pages 12
- [50] TracMeNot, "TracMeNot, <http://mrl.nyu.edu/dhowe/trackmenot>," 2014. pages 12
- [51] S. T. Peddinti and N. Saxena, "On the privacy of web search based on query obfuscation: A case study of TrackMeNot," in *Privacy*
- Enhancing Technologies*, ser. LNCS 6205, M. Atallah and N. Hopper, Eds., 2010, pp. 19–37. pages 12