Privacy-Preserving Distributed Profile Matching in Proximity-Based Mobile Social Networks

Ming Li, *Member, IEEE,* Shucheng Yu, *Member, IEEE,* Ning Cao, *Member, IEEE,* and Wenjing Lou, *Senior Member, IEEE*

***Abstract*—Making new connections according to personal preferences is a crucial service in mobile social networking, where an initiating user can find matching users within physical proximity of him/her. In existing systems for such services, usually all the users directly publish their complete profiles for others to search. However, in many applications, the users' personal profiles may contain sensitive information that they do not want to make public. In this paper, we propose FindU, a set of privacy-preserving profile matching schemes for proximity-based mobile social networks. In FindU, an initiating user can find from a group of users the one whose profile best matches with his/her; to limit the risk of privacy exposure, only necessary and minimal information about the private attributes of the participating users is exchanged. Two increasing levels of user privacy are defined, with decreasing amounts of revealed profile information. Leveraging secure multi-party computation (SMC) techniques, we propose novel protocols that realize each of the user privacy levels, which can also be personalized by the users. We provide formal security proofs and performance evaluation on our schemes, and show their advantages in both security and efficiency over state-of-the-art schemes.**

***Index Terms*—Mobile social networks, private matching, secure multiparty computation.**

I. Introduction

With the proliferation of mobile devices, mobile social networks (MSNs) are becoming an inseparable part of our lives. Leveraging networked portable devices such as smart phones and PDAs as platforms, MSN not only enables people to use their existing online social networks (OSNs) at anywhere and anytime, but also introduces a myriad of mobility-oriented applications, such as location-based services and augmented reality. Among them, an important service is to make new social connections/friends within physical proximity based on the matching of personal profiles. For example, MagnetU and E-SmallTalker [2] are MSN applications that match one with nearby people for dating or friend-making based on common interests. In such an application, a user only needs to input some (query) attributes in her profile, and the system would automatically find the persons around with similar profiles. The scopes of these applications are very broad, since people can input anything as they want, such as hobbies, phone contacts and places they have been to. The latter can even be used to find "lost connections" and "familiar strangers".

However, such systems also raise a number of privacy concerns. Let us first examine a motivating scenario. In a hospital, patients may include their illness symptoms and medications in their personal profiles in order to find similar patients, for physical or mental support. In this scenario, an initiating user (initiator) may want to find out the patient having the maximum number of identical symptoms with her, while being reluctant to disclose her sensitive illness information to the rest of the users, and the same for the users being matched with. If users' private profiles are directly exchanged with each other, it will facilitate user profiling where those information can be easily collected by a nearby user, either in an active or passive way; and those user information may be exploited in unauthorized ways. For example, a salesman from a pharmacy may submit malicious matching queries to obtain statistics on patients' medications for marketing purposes. To cope with user profiling in MSNs, it is essential to disclose minimal and necessary personal information to as few users as possible.

In fact, the ideal situation is to let the initiator and its best matching user directly and privately find out and connect to each other, without knowing anything about other users' profile attributes, while the rest of the users should also learn nothing about the two user's matching attributes. However, it is challenging to find out the matching users privately while efficiently. One may think of simply turning off the cellphone or input very few attributes, but these would interfere with the system usability. Recently, Yang *et. al*. proposed E-SmallTalker [2], a practical system for matching people's interests before initiating a small-talk. However, E-SmallTalker suffers from the dictionary attack which does not fully protect the non-match attributes between two users. Another difficulty of private matching under a MSN setting is the lack of a centralized authority. Lu *et. al.* [3] proposed a symptom matching scheme for mobile health social networks, assuming the existence of a semi-online central authority.

In this paper, we overcome the above challenges and make the following main contributions.

(1) We formulate the privacy preservation problem of profile matching in MSN. Two levels of privacy are defined along with their threat models, where the higher privacy level leaks

Manuscript received January 30, 2012; revised June 22, 2012; accepted February 27, 2013. The associate editor coordinating the review of this paper and approving it for publication was Y. Guan.

The preliminary version of this paper appeared in IEEE INFOCOM 2011 [1].

M. Li is with the Dept. of CS, Utah State University, Logan, UT 84322 (e-mail: ming.li@usu.edu).

S. Yu is with the Dept. of CS, University of Arkansas at Little Rock (e-mail: sxyu1@ualr.edu).

N. Cao is with Google Inc., Mountain View, CA (e-mail: ncao@ece.wpi.edu).

W. Lou is with the Dept. of CS, Virginia Tech (e-mail: wjlou@vt.edu).

Digital Object Identifier 10.1109/TWC.2013.032513.120149

less profile information to the adversary than the lower level.

(2) We propose two fully distributed privacy-preserving profile matching schemes, one of them being a private set-intersection (PSI) protocol and the other is a private cardinality of set-intersection (PCSI) protocol. However, solutions based on existing PSI schemes are far from efficient. We leverage secure multi-party computation (SMC) based on polynomial secret sharing, and propose several key enhancements to improve the computation and communication efficiency. Also, users can choose personalized privacy levels when running the same matching instance.

(3) We provide formal security proofs and extensive performance evaluation for our schemes. Our two protocols are shown to be secure under the honest-but-curious (HBC) model, with information-theoretic security (for PSI) and standard security (for PCSI), respectively. We also discuss possible extensions to prevent malicious attacks. Meanwhile, they are shown to be more efficient than previous schemes that achieve similar security guarantees under the typical settings of MSN.

II. Problem Definition

A. System Model

Our system consists of N users (parties) denoted as $P_1, ..., P_N$, each possessing a portable device. We denote the initiating party (*initiator*) as P_1. P_1 launches the matching process and its goal is to find one party that best *"matches"* with it, from the rest of the parties $P_2, ..., P_N$ which are called *candidates*. Each party P_i's profile consists of a set of attributes $\mathcal{S}_i$, which can be strings up to a certain length. P_1 defines a matching query to be a subset of $\mathcal{S}_1$, and in the following we use $\mathcal{S}_1$ to denote the query set unless specified. Also, we denote $n = |S_1|$ and $m = |S_i|, i > 1$, assuming each candidate has the same set size for simplicity. Note that, we assume that the system adopts some standard way to describe every attribute, so that two attributes are exactly the same if they are the same semantically.

There could be various definitions of "match". In this paper, we consider a popular similarity criterion, namely the intersection set size $|\mathcal{S}_1 \cap \mathcal{S}_i|$ (also used in [2]). The larger the intersection set size, the higher the similarity between two users' profiles. User P_1 can first find out her similarity with each other users via our protocols, and then will decide whether to connect with a best matching user based on their actual common attributes.

We assume devices communicate through wireless interfaces such as bluetooth or WIFI. For simplicity, we assume every participating device is in the communication range with each other. In addition, we assume that a secure communication channel has been established between each pair of users. We will discuss practical approaches for secure setup in Sec. IV-E.

We *do not* assume the existence of a trusted third party during the protocol run; all parties carry out profile matching in a completely distributed way. They may cooperate with each other, i.e., when P_1 runs the protocol with each P_i, a subset of the rest of parties would help them to compute their results. Note that, providing incentives for the users to cooperate is an important topic, and there are some existing mechanisms [4].

B. Adversary Model

In this paper, we are mainly interested in insiders who are legitimate participators of the matching protocol and try to perform *user profiling*, i.e., obtain as much personal profile information of other nearby users as possible. For example, With a user's attributes, a bad guy could correlate and identify that user via its MAC addresses or public keys. However, we cannot absolutely prevent user profiling, because at least the initiator and its best matching user will mutually learn their intersection set. Thus we focus on minimizing the amount of private information revealed in one protocol run.

The main adversary model considered in this paper is *honest-but-curious* (HBC), i.e., a participant will infer private information from protocol run but honestly follow the protocol. Although we do not specifically address the malicious attacker model [5] where an adversary may arbitrarily deviate from the protocol run, we will discuss how our protocols can be extended to achieve security in that model. The adversary may act alone or several parties may collude. We assume that the size of a coalition is smaller than a threshold t, where t is a parameter. And we assume $N \geq 2t + 1$ (honest majority) in this paper.

C. Design Goals

1) Security Goals: Since the users may have different privacy requirements and it takes different amount of efforts to achieve them, we hereby (informally) define two levels of privacy where the higher level leaks less information to the adversaries. Aformal definition will be given in Sec. V.

Definition 1 (Privacy level 1 (PL-1)): When the protocol ends, P_1 and each candidate $P_i, 2 \leq i \leq N$ mutually learn the intersection set between them: $\mathcal{I}_{1,i} = \mathcal{S}_1 \cap \mathcal{S}_i$. An adversary $\mathcal{A}$ (whose capability is defined in Sec. II-B) should learn nothing beyond what can be derived from the above outputs and its private inputs.

If we assume the adversary has unbounded computing power, PL-1 actually corresponds to unconditional security for all the parties under the HBC model. In PL-1, P_1 can obtain all candidates' intersection sets within one protocol run, which may still reveal much user information to the attacker. Therefore we define privacy level 2 in the following.

Definition 2 (Privacy level 2 (PL-2)): When the protocol ends, P_1 and each candidate $P_i, 2 \leq i \leq N$ mutually learn the size of their intersection set: $m_{1,i} = |\mathcal{S}_1 \cap \mathcal{S}_i|$. The adversary $\mathcal{A}$ should learn nothing beyond what can be derived from the above outputs and its private inputs.

In PL-2, except when $m_{1,i} = |\mathcal{S}_1|$ or $|\mathcal{S}_i|$, P_1 and each P_i both will not learn exactly which attributes are in $\mathcal{I}_{1,i}$. The adversary needs to run the protocol multiple times to obtain the same amount of information with what he can obtain under PL-1 when he assumes the role of P_1.

2) Usability and Efficiency: For profile matching in MSN, it is desirable to involve as few human interactions as possible. In this paper, a human user only needs to explicitly participate in the end of the protocol run, e.g., decide whom to connect to based on the common interests. In addition, the system design should be *lightweight and practical*, i.e., being enough efficient in computation and communication to be used in

MSN. Finally, different users (especially the candidates) shall have the option to flexibly *personalize their privacy levels*.

D. Design Challenges and Related Works

It is very challenging to achieve all the design goals simultaneously, especially if we desire high level of security but are unwilling to pay the cost of high computation and communication overhead. Similar problems to ours can be found in the literature, namely two-party private set intersection (PSI) [5]–[7], private cardinality of set intersection (PCSI) [5], [6], [8]. Our privacy goals can be achieved if given multiple instances of PSI and PCSI, respectively. They are usually tackled with Secure Multi-party Computation (SMC). The general SMC techniques [9] heavily rely on cryptography, and are well-known for their inefficiency. Researchers have proposed various customized solutions for those problems; for example, based on oblivious polynomial evaluation [5], [6], [8], [10], [11] and oblivious pseudo-random functions [7], [12], [13] that are secure in the HBC model. But when applied to the problem presented here, they incur either high computation or communication cost, thus are impractical to be used in MSN.

Concurrently with our work, a secure friend discovery protocol has been proposed in [14]. Different from us, their matching is based on computing the similarity (dot product) between two users' coordinates (which is not as intuitive as the intersection of the profile attributes as ours). In addition, a centralized trusted authority is needed to provide the coordinates. In [15], a private contact discovery protocol is proposed, where contact list manipulation is prevented by distributed certification. However, for general sensitive profile attributes it is difficult to find a distributed certifier in practice, whereas our protocols are not limited in the type of attributes to share with. In [16], privacy-preserving multi-party interest sharing protocols for smartphone applications are proposed. However, their protocols rely on an online semi-trusted server, which may not be available when the users do not have connections to it (e.g., poor signals).

In this paper, we propose two fully-distributed privacy-preserving profile matching protocols, without relying on a client-server relationship nor any central server. We propose novel methods to reduce energy consumption and protocol run time, while achieving reasonable security levels. Specifically, we exploit the homomorphic properties of Shamir secret sharing to compute the intersection between user profiles privately, and due to the smaller computational domain of secret sharing, our protocols achieve higher performance and lower energy consumption for practical parameter settings of an MSN. Such a framework is also applicable to many scenarios beyond the motivating problems in this paper, for example, in patient matching in online healthcare social networks.

III. Technical Preliminaries

Notations. We give the main notations in the following. Note that, unless specified, we denote $[s]_i$ as P_i's $(t, 2t+1)$-share of secret s under Shamir secret sharing (SS) scheme, and when we mention P_i, we refer to $2 \le i \le N$.

Main Notations

N, t :	Number of parties, maximum number of colluders
$[s]_i^{t,w}$:	Party P_i's secret share of s (under (t,w)-SS)
$\mathcal{S}_1$, $\mathcal{S}_i$:	P_1's query attribute set, and P_i's profile attribute set
$x_j, 1 \le j \le n$:	P_1's query set elements, $n = \lvert\mathcal{S}_1\rvert$
$y_{ij}, 1 \le j \le m$:	P_i's profile set elements, $m = \lvert\mathcal{S}_i\rvert, i \in \{2, \cdots, N\}$
$\mathcal{I}_{1,i}$:	Intersection set between P_1 and P_i; $m_{1,i} = \lvert\mathcal{I}_{1,i}\rvert$
$\mathcal{F}_p$:	The finite field used; $\kappa = \log p$: security parameter
$\mathcal{H}()$:	A cryptographic hash function
$\xleftarrow{R}$, $\lvert\rvert$:	Random sampling from a set, concatenation
$\mathcal{P}$, P_1, P_i:	The set of all parties, the initiator and the ith party
$\mathcal{P}_i$, $\mathcal{P}_i'$:	The computing set and reconstruction set for P_i

Preliminaries. **Shamir secret sharing scheme (SS)**. A (t, w)-SS scheme [17] shares secret s among w parties by giving each party P_i the value $[s]_i^{t,w}$, and if any at most t parties collude they cannot gain any information about s.

Secure multiparty computation (SMC) based on SS. For addition, SS is homomorphic: let α and β be two secrets shared using (t, w)-SS, we have $[\alpha + \beta]_i^{t,w} = [\alpha]_i^{t,w} + [\beta]_i^{t,w}$, denoted as SS-Add. However, for secure multiplication, one round of communication is needed and it is required that $w \ge 2t + 1$ [18]. Gennaro *et. al.* proposed the following efficient secure multiplication protocol [19]. Let the inputs of party P_i be $[\alpha]_i^{t,w}$ and $[\beta]_i^{t,w}$; the idea is to first locally multiply these shares, which lie on a $2t$-degree polynomial (but not random). Thus their protocol realizes randomization and degree-reduction in one round by letting each P_i pick a random t-degree polynomial and re-share $[\alpha]_i^{t,w}[\beta]_i^{t,w}$ to others:

Round 1. Each party P_i shares the value $[\alpha]_i^{t,w}[\beta]_i^{t,w}$ by choosing a t-degree random polynomial $h_i(x)$, s. t. $h_i(0) = [\alpha]_i^{t,w}[\beta]_i^{t,w}$. He sends the value $h_i(j)$ to party P_j, $1 \le j \le w$.

Round 2: Every party P_j computes his share of $\alpha\beta$, i.e., the value $H(j) = [\alpha\beta]_j^{t,w}$ under a t-degree random polynomial H, by locally computing the linear combination $H(j) = \sum_{i=1}^{w} \lambda_i h_i(j)$, where $\lambda_1, ..., \lambda_w$ are known constants (Lagrangian coefficients).

This protocol incurs $O(w^2\kappa)$ communication cost in total, and $O(w\kappa)$ for each party. We denote the above protocol as SS-Mul.

Additive homomorphic encryption. An additive homomorphic encryption scheme E allows one to compute $E(m_1 + m_2)$ given $E(m_1)$ and $E(m_2)$, without knowing the underlying plain texts. This is only used in our protocol for PL-2.

IV. Main Design of FindU

In this section, we first outline the idea of FindU, and then present two core designs for the PSI and PCSI protocols. Finally we address practical issues including user discovery.

A. Overview

We present two protocols that aim at realizing one level of privacy requirement each. We start with the basic scheme realizing PSI under PL-1, which is based on secure polynomial evaluation using secret sharing. At a high level, for P_1 and each P_i $(2 \le i \le N)$, their inputs are shared among a subset $\mathcal{P}_i$ of $2t + 1$ parties (the computing set) using $(t, 2t + 1)$-SS, based on which they cooperatively compute shares of the function $F_i(x_j) = R_{ij} \cdot f_i(x_j) + x_j$ for each $1 \le j \le n$,

where $f_i(y)$ is the polynomial representing P_i's set, and R_{ij} is a random number jointly generated by P_1 and P_i but not known to any party. We have $x_j \in \mathcal{I}_{1,i}$ iff. $F_i(x_j) = x_j$. The values of $\{F_i(x_j)\}_{1\le j\le n}$ remain in secret-shared forms between P_1 and P_i before their shares are revealed to each other. To reduce the communication complexity, we propose an enhancement that aggregates multiple multiplication and addition operations into one round during the secure polynomial evaluation computation.

For PL-2, the advanced scheme achieves efficient PCSI. The main idea is that, the parties in $\mathcal{P}_i$ first compute the $(t, 2t+1)$-shares of the function $F_i(x_j) = R_{ij} \cdot f_i(x_j)$, $1 \le j \le n$ securely using the basic scheme, whereas $x_j \in \mathcal{I}_{1,i}$ iff. $R_{ij} \cdot f_i(x_j) = 0$. In order to blind from P_1 the correspondence between its inputs $\{x_j\}$ ($j \in \{1, \cdots, n\}$) and the outputs $F_i(x_{j'})$ ($j' \in \{1, \cdots, n\}$), we employ a blind-and-permute (BP) method. To reduce the number of invocations of the BP protocol, we use share conversion to convert the $(t, t+1)$-shares of $\{F_i(x_j)\}_{1\le j\le n}$ (held by parties in the reconstruction set $\mathcal{P}'_i$) into $(1, 2)$-shares shared between P_1 and P_i, so that only one BP invocation is needed between P_1 and each P_i.

B. The Basic Scheme

We first give two definitions that capture the idea to involve the minimum number of parties during computation.

Definition 3 (Computing set of P_i): A set of $2t+1$ parties $\mathcal{P}_i \subset \mathcal{P}$, who help P_1 and P_i to compute the shares of $F_i(x_j)$, $1 \le j \le n$. $\mathcal{P}_i$ includes P_1 and P_i, and the rest $2t-1$ parties are chosen as $P_{i+1}, P_{i+2}, \cdots$ with indices wrapping around.

Definition 4 (Reconstruction set of P_i): A set of $t+1$ parties $\mathcal{P}'_i \subset \mathcal{P}_i$, who will contribute the shares of $F_i(x_j)$, $1 \le j \le n$ to P_1 and P_i for reconstruction, $\mathcal{P}'_i$ also includes P_1 and P_i, and the rest $t-1$ parties are chosen in the same way as in the computing set.

As input, each party has a set of attributes: P_1 has $\mathcal{S}_1 = \{x_1, x_2, ..., x_n\}$ and P_i has $\mathcal{S}_i = \{y_{i1}, y_{i2}, ..., y_{im}\}$, respectively, where each element is an encoded attribute in $\mathcal{F}_p$. For example, a hash algorithm can be used for encoding. Rather than publishing the sets as they are, each P_i first generates an m-degree polynomial based on $\mathcal{S}_i$ as follows:

$$f_i(y) = (y - y_{i1}) \cdot (y - y_{i2}) \cdots (y - y_{im}) = \sum_{k=0}^{m} a_{ik} y^k, \quad (1)$$

where $\{a_{ik}\}_{0\le k\le m-1}$ are coefficients. We require $a_{im} \equiv 1$ so that P_i cannot give an all-zero polynomial. The function to be computed is: $F_i(x_j) = R_{ij} \cdot f_i(x_j) + x_j$ for each $1 \le j \le n$, where $R_{ij} = r_{ij} r'_{ij}$, r_{ij} and r'_{ij} are random numbers generated by P_1 and P_i, respectively. In this way, if $F_i(x_j) \in \mathcal{S}_i$, $x_j \in \mathcal{I}_{1,i}$ with high probability, and if $F_i(x_j) \notin \mathcal{S}_1$ then $x_j \notin \mathcal{I}_{1,i}$.

The basic scheme consists of three phases, and Fig. 1 describes one run between two parties - P_1 and P_i. The whole protocol between P_1 and $P_2, ..., P_N$ consists of $N-1$ instances of the two-party protocol, which can be parallelized/aggregated to save time (due to space limitations, details are shown in [1]). In the *data share distribution* phase, P_1 shares the 1 to m powers of each of its set elements, while P_i shares its private inputs among P_i's computing set. In addition, P_1 and P_i also share their n random numbers, respectively.

In the *computation* phase, the parties in $\mathcal{P}_i$ participate in secure computation of the shares of $\{F_i(x_j)\}_{1\le j\le n}$. In particular, to evaluate $f_i(x_j)$, a straightforward way is to compute $m-1$ multiplications of $a_{ik}x_j^k, 1 \le k \le m-1$ by invoking the SS-multiplication protocol $m-1$ times. However, this will introduce too much communication cost.

Therefore, we propose to aggregate those multiplications into one round. That is, each party $P_l \in \mathcal{P}_i$ first locally compute a product-sum of shares $z_{ijl} = \sum_{k=1}^{m-1} [a_{ik}]_l [x_j^k]_l$ based on $m-1$ pairs of local shares $\{[a_{ik}]_l, [x_j^k]_l\}_{1\le k\le m-1}$. Then, after computing z_{ijl}, each party $P_l \in \mathcal{P}_i$ proceeds in the same way as in SS-Mul. Specifically, each P_l shares the value z_{ijl} to others by choosing a t-degree random polynomial $h_l(x)$, and then locally computes the same linear combination $(\sum_{k=1}^{2t+1} \lambda_k h_k(l))$ of the received secondary shares to get its own share of the product-sum - $[\sum_{k=1}^{m-1} a_{ik} x_j^k]_l$. We denote this variant of SS-Mul as SS-Mul-Add, whose correctness follows from the homomorphic properties of SS-Add and SS-Mul. Since $F_i(x_j) = r_{ij} r'_{ij} (a_{i0} + \sum_{k=1}^{m-1} a_{ik} x_j^k + x_j^m) + x_j$, P_l's share of $F_i(x_j)$ can then be easily computed by invoking two more SS-Mul.

In the *reconstruction* phase, at least $t+1$ shares of $F_i(x_j)$ are needed to reconstruct $F_i(x_j)$. To this end, the parties reveal their shares to P_1 and P_i, who can obtain $F_i(x_j)$ by polynomial interpolation. P_1 and P_i can test if $F_i(x_j) = x_j, 1 \le j \le n$ and $F_i(x_j) = y_j, 1 \le j \le m$ respectively, to determine their intersection set.

C. The Advanced Scheme

Observe that, in the basic scheme if we set $F_i(x_j) = r_{ij} r'_{ij} f_i(x_j)$, then the result will be 0 if $x_j \in \mathcal{I}_{1,i}$, otherwise a random number. In order to obtain the number of matching attributes ($m_{1,i}$), one way is to employ the equality-test protocol [20] based on SS. However, this method incurs too high communication cost, since even the most efficient (probabilistic) algorithm takes $12k$ invocations of the SS multiplication protocol, where 2^{-k} is the error probability. When $k = 10$, the communicated bits for test one number ($F_i(x_j)$) amounts to $120N^2p$, and there are $(N-1)n$ numbers to test. Considering that in modern smart mobile devices, the wireless transmission is more costly than computation ability, we would like to tradeoff computation for communication efficiency.

Thus, we adopt a *blind-and-permute* (BP) method to obliviously permute P_1's shares of each $F_i(x_j)$, so that the linkage between $F_i(x_j)$ and its corresponding attribute x_j is broken. A BP protocol between two parties A and B where each data item (e.g., $F_i(x_j)$) is additively split between them is described in [21]. The main idea is, A encrypts each of its shares using additive homomorphic encryption and sends to B. B then generates a different random number r_j for each shared item, and randomizes each of A's shares by adding r_j, while subtracting r_j for its own corresponding shares. B permutes the randomized shares using a pseudo-random permutation (PRP), and sends back to A. All the computations are done over the ciphertexts.

P_1's private inputs: $\mathcal{S}_1 = \{x_1, ..., x_n\}$, $\mathcal{I}_{1,i} \leftarrow \emptyset$; P_i's private inputs: $\mathcal{S}_i = \{y_{i1}, ..., y_{im}\}$, $\mathcal{I'}_{1,i} \leftarrow \emptyset$. Public inputs: t, $\mathcal{P}_i, \mathcal{P}'_i, 2 \leq i \leq N$. Output: $\mathcal{I}_{1,i}$ for both P_1 and P_i.

1. Data share distribution:

1) P_1 generates $\{r_{ij}\}_{1\leq j\leq n}$ randomly from $\mathcal{F}_p$. It shares $\{r_{ij}\}_{1\leq j\leq n}$ and $\{x_j, x_j^2, ..., x_j^m\}_{1\leq j\leq n}$ among parties in $\mathcal{P}_i$ using $(t, 2t+1)$-SS.
2) P_i generates $\{r'_{ij}\}_{1\leq j\leq n}$ randomly from $\mathcal{F}_p$, lets $f_i(y) = \prod_{k=1}^{m}(y - y_{ik}) = \sum_{k=0}^{m} a_{ik}y^k$, and shares $\{r'_{ij}\}_{1\leq j\leq n}$, $\{a_{ik}\}_{0\leq k\leq m-1}$ among parties in $\mathcal{P}_i$ using $(t, 2t+1)$-SS.

2. Computation: Each party $P_l \in \mathcal{P}_i \in \mathcal{P}$ executes the following:

1) $\forall 1 \leq j \leq n$, P_l computes the share of $\sum_{k=1}^{m-1}(a_{ik}x_j^k)$ through one invocation of SS-Mul-Add (namely, $[a_{i1}x_j + ... + a_{i,m-1}x_j^{m-1}]_l \leftarrow$ SS-Mul-Add$(\{[a_{ik}]_l, [x_j^k]_l\}_{1\leq k\leq m-1})$). (one communication round)
2) P_l computes $[f_i(x_j)]_l = [a_{i0}]_l + [a_{i1}x_j + ... + a_{i,m-1}x_j^{m-1}]_l + [x_j^m]_l$ via two SS-Add operations (local).
3) P_l computes $[R_{ij}]_l = [r_{ij}r'_{ij}]_l$ and $[F_i(x_j)]_l = [R_{ij}f_i(x_j) + x_j]_l$, for all $1 \leq j \leq n$, through two invocations of SS-Mul (two communication rounds), and one SS-Add (local).

3. Reconstruction:

1) Aggregate shares: each $P_l \in \mathcal{P}'_i$ sends its shares $\{[F_i(x_j)]_l\}_{1\leq j\leq n}$ to P_1 and P_i.
2) Reconstruction: P_1 reconstructs $F_i(x_j), 1 \leq j \leq n$ from their shares, using Lagrangian interpolation. If $F_i(x_j) \in \mathcal{S}_1$, $\mathcal{I}_{1,i} \leftarrow \mathcal{I}_{1,i} \bigcup F_i(x_j)$. The similar is done for P_i.

Fig. 1: The basic scheme, run between the initiator P_1 and each candidate $P_i \in \mathcal{P}, 2 \leq i \leq N$.

However, the BP protocol cannot be applied directly. In our protocol, each $F_i(x_j)$ is polynomially shared among $t+1$ parties in $\mathcal{P}'_i$, where at most t of them may be adversarial. Without loss of generality, assume P_i is the party that generates the random permutation (π). Then in order to randomly permute the shares of all the rest parties in $\mathcal{P}'_i$, there will be at least t invocations of the BP protocol between P_i and other parties in $\mathcal{P}'_i$, which is too computationally expensive.

Hence, in our scheme we propose an improvement that only requires *one* invocation of BP protocol. The idea is to convert the $(t, t+1)$-shares of $F_i(x_j)$ among $P_i \in \mathcal{P}'_i$ into $(1, 2)$-shares shared between P_1 and P_i. The conversion from $(t, t+1)$-share to $(1, t+1)$-share is a standard procedure [22] which involves one round of re-sharing. To transform $(t, t+1)$-shares into $(1, 2)$-shares, we use a trick in which all the parties in $\mathcal{P}'_i$ only re-share their $(t, t+1)$-shares of $F_i(x_j)$ to P_1 and P_i, so that *only P_1 and P_i together can reconstruct s*. Thus the BP protocol is performed only once between them. The parties P_1 and P_i then reconstruction the result $m_{1,i}$ by exchanging their shares. The protocol to compute $m_{1,i}$ is described in Fig. 2.

D. Personalizing Users' Privacy Levels

In our design, a user can choose her privacy level by telling other parties her choice in the beginning. For example, if a candidate P_i chooses PL-1, she can broadcast a message indicating "P_i selects PL-1". Then the parties in P_i's computing set and reconstruction sets will follow the basic scheme to compute the desired results for P_i. However, the initiator should always agree on the privacy level that each candidate proposes, since P_1 is at a position easier to conduct user profiling.

E. Practical Implementation Issues

In practice, proximity-based user discovery and key establishment are two important issues for the usability of our profile matching protocols. We envision that our FindU scheme can be used in mobile devices equipped with short-range wireless interfaces like WiFi or Bluetooth (most of today's smartphones have both interfaces), and operate in the ad-hoc mode. We have done some prior work in practical trust initialization in wireless networks [23], [24]. Here we describe possible setup processes that involve little human effort.

User discovery If using Bluetooth, the existing Service Discovery Protocol (SDP) can be utilized to search for nearby FindU users (range is about 10m). As shown in [2], the SDP protocol can be used to publish/exchange information. We can use it to initiate the protocol (including key establishment). For WiFi, the ad-hoc mode would be sufficient for device discovery.

Key establishment As pairwise keys should be established between all nearby devices, a straightforward design (e.g., in Bluetooth each user needs to manually pair up with all the other users) would require $O(N^2)$ complexity. In order to reduce the complexity and minimize the need of explicit human participation, Diffie-Hellman key exchange (DHKE) can be used. Specifically, each device P_i chooses a random number $X_i \overset{R}{\leftarrow} \mathbb{Z}_p$ and publishes only $Y_i = g^{X_i}$ using broadcast. After receiving all the Y_j's, each P_i can calculate the pairwise key as P_j as $Y_j^{X_i}$. To achieve key authenticity, we can adopt tamper-evident pairing (TEP) [25], in which any modification of key exchange messages between two users by an attacker will be detected. Although TEP was implemented in the two-party setting, the same idea can also be applied to broadcast (each device in the group can negotiate to fix their messaging sequence and timing after device discovery and before pairing). In this method, the communication and time complexity is only $O(N)$. Though DHKE uses two modular exponentiation operations, this overhead can be amortized to multiple profile matching protocol runs. As in TEP, this approach works strictly in-band and only requires each user to press a button once, which is less intrusive. In addition, it is fully distributed and do not need any central server.

Indeed this approach is more favorable for WiFi devices with richer resources; however, for Bluetooth devices it can also be used. We note that although the SDP protocol does not support broadcast, each phone can publish up to 10 numbers with a maximum of 128 bytes each [2]. This would be sufficient for DHKE with 1024- bit group size. The downside is that unicast entails communication complexity of $O(N^2)$; however, given that the number of users within 10m range is usually small, it will not be a big problem.

Inputs: basically the same as those in Fig. 1. In addition, P_1 sets $m_{1,i} = 0$, and has a public key pk_1 and private key sk_1, and pk_1 is known by all others. P_i sets $m'_{1,i} = 0$. Output: intersection set size $m_{1,i}$.
1. Data share distribution: the same as in Fig. 1, except that P_1 first randomly permutes its set S_1.
2. Computation: Basically the same as in Fig. 1, except each $P_l \in \mathcal{P}_i$ computes $[F_i(x_j)]_l = [R_{ij} f_i(x_j)]_l$, $\forall 1 \le j \le n$. For each $P_l \in \mathcal{P}'_i$, its shares are also denoted as $[F_i(x_j)]_l^{t,t+1}$.
3. Reconstruction:
1) Share conversion: Each $P_l \in \mathcal{P}'_i$ first computes and sends $[[F_i(x_j)]_l^{t,t+1}]_1^{1,t+1}$, $[[F_i(x_j)]_l^{t,t+1}]_i^{1,t+1}$ to P_1 and P_i, respectively. Then, P_1 and P_i compute $[F_i(x_j)]_1^{1,2}$ and $[F_i(x_j)]_i^{1,2}, 1 \le j \le n$ using Lagrangian interpolation.
2) Blind and permute P_1's shares: P_1 and P_i involve in one BP-protocol where P_i generates permutation π and $\{r''_j\}_{1 \le j \le n}$, and P_1's shares will be first randomized by $\{r''_j\}_{1 \le j \le n}$ and then permuted by π.
3) Reconstruction: P_1 and P_i exchange their shares of $\{F_i(x_{j'})\}_{j' \in \{1,\cdots,n\}}$. Then, they reconstruct $F_i(x_{j'})$ for each j'. Then both count the number of $F_i(x_{j'}) = 0$, and set this number to $m_{1,i}$.

Fig. 2: The advanced scheme, run between P_1 and P_i.

V. Security Analysis

In this section, we first prove the security of each scheme under the HBC (semi-honest) model, and then discuss extensions to resist active attacks that deviate from the protocol run.

A. Security Definition

We define the security of our schemes based on the standard definition of secure multi-party computation under the HBC model [26] (c.f., Chapter 7), assuming private channels. Loosely speaking, "a multi-party protocol privately computes F, if whatever a set of semi-honest parties can obtain after participating in the protocol could be essentially obtained from the input and output of these very parties". Our protocols (denoted as Π) compute a two-ary functionality F (either the set intersection or cardinality of intersection set) between A (initiator) and B (a responder) with the help of at least $2t-1$ other parties (for simplicity, in our proofs we assume that all N parties $\mathcal{P}$ participate in their computation). That is, $F_(S_A, S_B) = S_A \bigcap S_B = \mathcal{I}_{A,B}$, and $F'_(S_A, S_B) = |S_A \bigcap S_B| = |\mathcal{I}_{A,B}|$ which are deterministic functions. Note that, the input of A is: $\vec{x}_A = S_A = \{x_1, ..., x_n\}$ while for B, the input is $\vec{x}_B = S_B = \{y_1, ..., y_m\}$. However, parties in $\mathcal{P} \backslash \{A, B\}$ do not have inputs nor yield any output. Thus, our protocol is a special case of multi-party computation. The view of the i-th party during an execution of Π on input $\bar{S} = \{S_A, S_B\}$ is defined as: $\text{VIEW}_i^{\Pi}(\bar{S}) \triangleq (S_i, r, m_1, ..., m_l)$ if $i = A$ or B, and $\text{VIEW}_i^{\Pi}(\bar{S}) \triangleq (r, m_1, ..., m_l)$ otherwise, where m_i represents the i-th message it has received and r stands for the internal randomness. Now we give formal definitions for privacy.

Definition 5 (t-privacy with respect to semi-honest behavior): We say that Π t-privately computes deterministic function F if there exist a P.P.T. algorithm, denoted as $\mathcal{A}$, such that for every coalition of semi-honest parties $\Gamma = \{i_1, ..., i_t\} \subset \mathcal{P}$ with size at most t, it holds that

$$\{\mathcal{A}(\Gamma, (\vec{x}_{i_1}, ..., \vec{x}_{i_t}), F_\Gamma(\vec{x}))\}_{\vec{x} \in \{0,1\}^*} \equiv \{\text{VIEW}_\Gamma^{\Pi}(\vec{x})\}_{\vec{x} \in \{0,1\}^*} \quad (2)$$

where $\vec{x}$ is the vector of all the inputs. In the middle when we write "$\equiv$", we mean perfect indistinguishability (correspondingly to computationally unbounded adversaries); while "$\stackrel{c}{\equiv}$" refers to computational indistinguishability (for computationally bounded adversaries).

In what follows, our proofs are based on simulation in the "real-world"/"ideal-world" paradigm.

B. Security Proof of the Basic Scheme Under the HBC Model

Theorem 1: The basic protocol (in Fig. 1) t-privately computes the set intersection between two parties A and B (in the semi-honest model), against computationally unbounded adversaries.

Due to space limitations, the full proof is provided in our technical report [27].

Remark on security parameters. Due to unconditional security, it is enough as long as the field $\mathcal{F}_p$ can represent all the attributes in the attribute dictionary. Therefore, assume there are 1×10^6 attributes, we choose $\Bbbk = \log p = 24$.

C. Security Proof of the Advanced Scheme Under the HBC Model

Theorem 2: Assuming a secure pseudorandom permutation and semantically secure additive homomorphic encryption scheme, the advanced protocol (in Fig. 2) t-privately computes the cardinality of set intersection between two parties A and B (in the semi-honest model), against computationally bounded adversaries.

The proof is provided in our technical report [27].

D. Discussion: Preventing Malicious Attacks

Our protocols in this paper are only proven secure in the HBC model; it would be interesting to make it secure under the stronger malicious model, i.e., to prevent an adversary from arbitrarily deviating from a protocol run. In the conference version of this paper [1], we showed that with an additional commitment round before final reconstruction (which adds little additional overhead), a specific type of "set inflation attack" can be easily prevented where a malicious user influences the final output in her favorable way by changing her shares after seeing others'.

Further, we observe it is possible that our protocols can be extended to the malicious model with some additional cost[1], following the same method in [28]. The idea is to use verifiable secret sharing (VSS) [18], [19], which is secure under the malicious model with $t < N/3$ malicious parties. Asharov and Lindell proved in [28] that (c.f. Sec. 7), with a VSS scheme and a secure multiplication protocol F_{mul} that are both secure under the malicious model, an arbitrary multi-party function F can be computed securely in the malicious

[1]Theoretically, a general SMC protocol that is secure in the semi-honest model can be converted to one secure in the malicious model, using a "compiler" [26], which, however, bears too much cost.

TABLE I: Comparison of complexity ($q = 1024$, $\kappa = 24$)

	Party	Basic scheme	PSI [7]	Advanced scheme	PCSI [5]
Computation	P_1	$6nNt + mnN(1 + t\log N/\kappa)$ mul_1	$1.5nN$ mul_2	$3nN$ exp_3	$(2n + mN)$ exp_3
	P_i	$2mnt + 2(m + 6nt)t^2\log N/\kappa$ mul_1	$(m + n)$ exp_2	n exp_3	$m\log(\log n)$ exp_3
Comm. (TX)	P_1	$(mnN + 8nNt)\kappa$	$2nNq$	$(mnN + 8nNt)\kappa + 2nNq$	$2nq$
	P_i	$2t(m + 6nt)\kappa$	$(n + m)q$	$2t(m + 6nt)\kappa + 2nq$	$2mq$
Comm. (RX)	P_1	$[N(m + n) + 8nNt]\kappa$	$(n + m)Nq$	$N[m + (t + 3N)n]\kappa + 2nNq$	$2mNq$
	P_i	$[m(n + 2t) + 12nt^2]\kappa$	$2nq$	$[m(n + 2t) + 12nt^2]\kappa + 2nq$	$2nq$
Comm. total	All	$[mnN + tN(8n + 2m + 12nt)]\kappa$	$N(3n + m)q$	$[mnN + tN(8n + 2m + 12nt)]\kappa + 4nNq$	$2(n + mN)q$

model. In addition they proposed a construction of F_{mul} that is based on VSS. They represent F using an arithmetic circuit with three kinds of gates: addition, multiplication-by-constant, and multiplication; secure computation of F replaces each of those gate by its secret sharing counterpart, namely SS-Add, SS-Mul-by-Const, SS-Mul. As long as these components are secure, they can be composed into a secure protocol according to the composition theorem [29].

Therefore, in our situation, as long as we design a secure SS-based Mul-Add protocol that is secure in the malicious model, our PSI and PCSI protocols can be extended to be secure under the malicious model using VSS. Because the Mul-Add aggregates multiple multiplications into one round, we believe the same method to construct a secure F_{mul} still applies. This will be left as an interesting future work.

VI. Performance Evaluation

In this section, we analyze the complexity of our proposed schemes[2], carry out an extensive simulation study of the protocols' efficiency, and compare them with several state-of-the-art schemes in terms of security and efficiency.

A. Complexity Analysis

The computational cost is evaluated using the number of modular multiplication and exponentiation operations, while the communication cost is calculated in terms of number of transmitted/received bits. Let $\kappa = 24$, and assume one κ-bit modular multiplication (denoted as mul_1) takes κ^2 bit operations. Thus, to share a secret $s \in \mathcal{F}_p$ among N parties using (t, N)-SS, it takes $\kappa Nt\log N$ bit operations [22], which is $\frac{tN\log N}{\kappa}$ multiplications. For 1024-bit and 2048-bit modular multiplications (denoted as mul_2, mul_3) and exponentiations (denoted as exp_2, exp_3), since there exist optimization algorithms, we will use existing benchmark test results instead. Note that we neglect modular additions, and we make the following assumptions to simplify the calculations: $m \gg N$, $q \gg \kappa$ and $m, n, N, t \gg 1$. For the details of the complexity analysis, please refer to our technical report [27].

The complexity results are summarized in Table I, and we compare our basic and advanced schemes with existing schemes: FC10 PSI [7] and FNP PCSI (under HBC model) [5], respectively. It can be seen that, the basic scheme's total computation complexity is much smaller than that of FC10's since $q \gg \kappa$, while that of the advanced scheme's is smaller than FNP's when n is relatively small w.r.t. m. Although the total communication costs may seem large in our schemes, they are on the same orders with the compared schemes in terms of m, n and N. The effect of the $O(t^2)$ factor is moderate unless t scales linearly with N, as we will see in the simulation results.

[2]To reflect the reality, the full matching protocols between the initiator and $N-1$ other users are considered, which are aggregations of $N-1$ individual protocol runs between P_1 and each P_i [1].

B. Simulation Study

1) Methodology: We implement our proposed schemes and two previous schemes, FC10 and FNP, in NS-2 [30]. We simulate the protocols' communications and computations by telling the simulator the sizes and number of packets each party should send, fill each packet with dummy contents, and estimate the latency of each computation. Note that, in each round/step, we exploit the opportunities to aggregate messages sent to the same party into one packet as much as possible, so as to reduce the number of packets sent. In addition, we only simulate one full protocol run, as the time variance is very small due to deterministic scheduling.

For simulation settings, we assume the use of WiFi as the wireless interface which operates in the ad-hoc mode, and IEEE 802.11a is used for MAC and PHY layer. Nodes (parties) are randomly distributed in a $50m \times 50m$ area. The transmission range is set to 200m, such that all nodes are within reach of each other. Two-ray-ground propagation model is assumed, and the wireless channels are reliable.

Evaluation metrics. To evaluate the efficiency of the protocols, we adopt energy consumption and total run time as two unified metrics. Both of them factor into the effects of both computation and communication time, and are closely relate to the user experience; the energy consumption is also affected by the total run time.

We estimate the computation time for primitive operations based on the existing benchmark test results [31]. Assuming a 400MHz CPU, the times for the primitive crypto operations are (seconds): 1.4×10^{-6} for mul_1, 8×10^{-5} for mul_2, 0.04 for exp_2, 2.4×10^{-4} for mul_3, and 0.25 for exp_3.

Our energy consumption model is based on [32]: $E = n_tE_t + n_rE_r$ [3], where n_t and n_r are sent and received in bytes, $E_r \approx 6.7\mu J$ is the receiving energy per byte, and $E_t \approx 4.8\mu J$ is transmitting energy per byte. Also, if the connection time exceeds 15 seconds, it is more efficient to shut down WIFI radio [33]. Thus, in simulation we employ the following strategy to save energy: for each party, whenever it estimates its computing time in one step to be longer than 15s, it shuts down its own WIFI radio and also indicates others to close theirs for a known time period, and reopens it when the computation is done. Now, the model can be described in Eq. (2), where n_{shut} is the # of times WIFI radio has shut down, T_{shut} is the total radio shut down time, P_{comp}, T_{comp} are the CPU's power consumption and time

[3]We omit the initial connection establishment energy since it is common in all schemes.

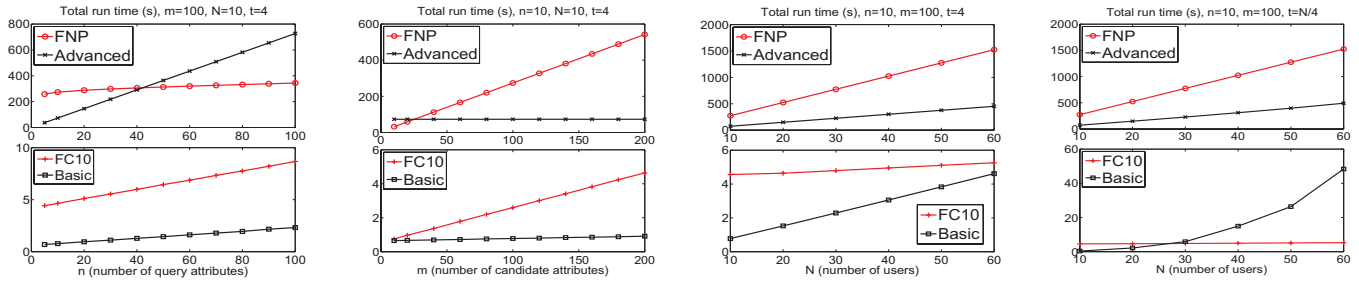

(a) Change n. (b) Change m. (c) Change N, fix $t = 4$. (d) Change N, change $t = \lfloor N/4 \rfloor$.

Fig. 3: Total protocol run time. (Y-axis: protocol run time (s).)

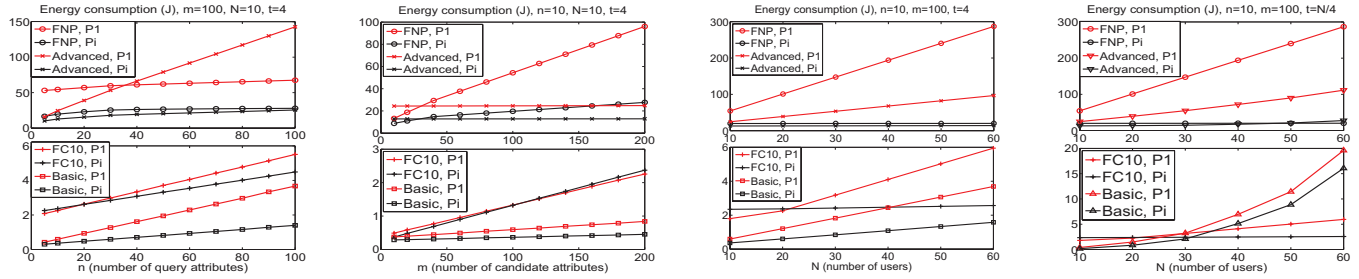

(a) Change n. (b) Change m. (c) Change N, fix $t = 4$. (d) Change N, change $t = \lfloor N/4 \rfloor$.

Fig. 4: Energy consumption for P_1 and P_i. (Y-axis: energy consumption (J).)

$$E = \begin{cases} 10^{-6}(6.7n_t + 4.8n_r) + P_{comp}T_{comp} + 0.3167T_{total} \quad (J), & \text{If computations in every step take time less than } 15s \\ 10^{-6}(6.7n_t + 4.8n_r) + P_{comp}T_{comp} + 0.3167(T_{total} - T_{shut}) + 5n_{shut} \quad (J), & \text{Otherwise.} \end{cases} \tag{3}$$

spent for computation, and T_{total} stands for the total protocol run time. For a smart phone with 400MHz CPU, we choose $P_{comp} \approx 0.18$w [34].

2) Simulation results: We show the simulation results in Fig. 3 and Fig. 4; we compare the basic scheme with the FC10 scheme [7] (the RSA exponents/modulus are both 1024-bits), and the advanced scheme with the FNP scheme [5]. We also assume the use of Paillier's crytosystem for FNP, with 2048 bit modulus. Note that, we use $\mathbb{k} = 24$ bits for our schemes.

In Fig. 3 (a) and Fig. 4 (a), we fix the network size $N = 10$, maximum number of colluders $t = 4$, number of profile attributes $m = 100$, and change number of query attributes n. It can be seen that, the basic scheme takes less than 1 second when $n < 100$. Both the total run time and energy consumption of it is much less than that of FC10's in this case and they all increase linearly with n. For the advanced scheme, the time and energy are smaller than those of FNP's when $n < 40$. This is mainly because the computation in the advanced scheme scales as $O(nN)$, which is $O(n + mN)$ for the FNP scheme. Nevertheless, for the profile matching application, in reality it is often the case that n is small.

On the other hand, from Fig. 3 (b) and Fig. 4 (b) in which n is fixed to 10 but m changes from 10 to 200, it can be see that both the basic and the advanced schemes are more efficient than their counterparts, more importantly our proposed schemes are hardly affected by m. The above shows that our schemes are more practical when the number of profile attributes is large, while the number of query attributes is small.

Next, we change the number of parties. We first fix the maximum number of colluders to 4. From Fig. 3 (c) and Fig. 4 (c), one can observe that the basic scheme's costs increase linearly with N. This is because its run time and energy costs are both dominated by communication which is linear in N when t is fixed, since the computations are quite fast. The FC10 scheme is much more computationally intensive under this case. For the advanced scheme and FNP, their costs are both dominated by computation rather than communication, yet the advanced scheme performs much better due to $n < m$. P_i has almost constant energy consumption except in the basic scheme, since in those three schemes P_i's computation cost is mainly affected by m, but not N.

Finally, we scale t with N ($t = \lfloor N/4 \rfloor$) in Fig. 3 (d) and Fig. 4 (d). Interestingly, the advanced scheme is still much more efficient than the FNP scheme, and it exhibits a super-linear effect ($O(t^3)$ in overall communication) only when $N > 50$ for P_i's energy consumption. Meanwhile, the basic scheme suffers from this effect earlier than the advanced scheme (better than FC10 when $N < 30$), since it is dominated by communication.

The above demonstrates the advantage of our proposed schemes when n and N are both relatively small (in the order of tens), which is usual in mobile social networks. Note that, in all the compared protocols, it is always the case that P_1 uses more energy than P_i. Through using the energy-saving strategy, for the advanced scheme the parties' energy consumption seldom exceed 100J (equivalent to purely using WIFI for 5min), while that of the basic scheme is below 10J.

TABLE II: Comparison of security, and computation/communication efficiency under typical parameters (n=10, m=100, N=10)

Schemes		Basic	Advanced	PSI [7]	PSI [35]	PCSI [5]	PSI [11]	PCSI [11]
Adversary Model		HBC	HBC	HBC	Malicious	HBC	HBC	HBC
Resist active attacks		No*	No*	No	Yes	No	No	No
Computation	P_1	1.6×10^4 mul_1	300 exp_3	250 mul_2	$4 \times 10^4 \mathsf{mul}_2$	1020 exp_3	6×10^4 mul_1	4.4×10^4 mul_1
	P_i	9.5×10^3 mul_1	20 exp_3	110 exp_2	$2.96 \times 10^4 \mathsf{mul}_2$	173 exp_3	4.5×10^4 mul_1	3.5×10^4 mul_1
Communication (KB) Sent/Received	P_1	40/13	76/39	26/141	37/127	2.6/256	280/253	528/528
	P_i	8.2/11	11/13.6	14/2.6	14/4.1	25.6/2.6	200/203	422/422
Total comp. time (s)		0.023	80	4.42	5.57	298	0.14	0.093
Total sent bytes (KB)		114	175	166	164	259	2080	4326
Note: No* means can be easily extended to prevent certain malicious attacks.								

C. Further Comparison with Previous Works

We now compare our schemes with more existing schemes, in terms of security and efficiency. The computation/communication complexities are calculated by aggregating $N-1$ protocol runs between P_1 and each P_i, with P_1 being the initiator/client and each P_i being a responder/server.

First, the PSI scheme in [35] achieves standard security under the malicious adversary model. Its computation complexity is $400n(N-1)\mathsf{mul}_2$ operations for P_1 and $(560n+240m)\mathsf{mul}_2$ for P_i. For sent bytes, those are $(3n+2)(N-1)q$ for P_1 and $(2n+m+1)q$ for P_i. While they are linear with n, m, the constant factors are much higher than our basic scheme under typical MSN scenarios.

The schemes in [11] (CANS'09) represent a category with unconditional security. Their PSI and PCSI schemes are also based on secret sharing. For the PSI scheme, we modify it to fit our problem setting (by adding r_{ij} and r'_{ij}), which differs from our basic scheme in that: (1) it does not aggregate the multiplications in the computation phase; (2) it does not prevent set inflation attack. We denote the computation complexity of P_1 in our basic scheme as $\mathrm{Comp}_{P_1}(\mathrm{Basic})$, the number of bits sent by P_1 in our basic scheme as $\mathrm{Comm}_{P_1}(\mathrm{Basic})$, while $\mathrm{Comm}'_{P_1}(\mathrm{Basic})$ stands for the number of bits received by P_1 in our basic scheme. The computation complexity for the PSI scheme in [11] is $\mathrm{Comp}_{P_1}(\mathrm{Basic}) + 2mnNt^2\log N/\mathbb{k}$ mul_1 operations for P_1, and $\mathrm{Comp}_{P_i}(\mathrm{Basic}) + 4mnt^3\log N/\mathbb{k}$ mul_1 for P_i. The bytes sent by P_1 and P_i are: $\mathrm{Comm}_{P_1}(\mathrm{Basic}) + 2mnNt\mathbb{k}$, and $\mathrm{Comm}'_{P_1}(\mathrm{Basic}) + 4mnt^2\mathbb{k}$, respectively. For the PCSI scheme in [11], computation complexity for P_1 and P_i are: $2mnNt^2\log N/\mathbb{k}$ mul_1 and $4mnt^3\log N/\mathbb{k}$ mul_1, respectively, while the bytes sent are $(2mnNt + 240nNt)\mathbb{k}$ and $(4mnt^2 + 480nt^2)\mathbb{k}$, respectively.

The security and efficiency comparisons are summarized in Table II. For efficiency comparison, we set the parameters to be typical values: $n = 10$, $m = 100$, $N = 10$ and $t = 4$, and numerically evaluate the computation and communication costs. Note that, we also count the offline computations since those also consume energy. From Table II, it can be seen that our basic scheme outperforms all the other PSI schemes, while the advanced scheme is slower in computationally than the PCSI scheme in [11]. However, the latter is expensive in communication, and does not protect against active attacks.

VII. Conclusion

In this paper, we for the first time formalize the problem of privacy-preserving distributed profile matching in MSNs, and propose two concrete schemes that achieve increasing levels of user privacy preservation. Towards designing lightweight protocols, we utilize Shamir secret sharing as the main secure computation technique, while we propose additional enhancements to lower the proposed schemes' communication costs. Through extensive security analysis and simulation study, we show that 1) our schemes are proven secure under the HBC model, and can be easily extended to prevent certain active attacks; 2) our schemes are much more efficient than state-of-the-art ones in MSNs where the network size is in the order of tens, and when the number of query attributes is smaller than the number of profile attributes.

Acknowledgment

This work was supported in part by the US National Science Foundation under grants CNS-1156311, CNS-1156318 and CNS-1218085. The authors would like to thank Prof. William Martin for helpful discussions at the early stage of the protocols. We also thank the anonymous reviewers for their helpful comments.

References

[1] M. Li, N. Cao, S. Yu, and W. Lou, "Findu: privacy-preserving personal profile matching in mobile social networks," in *Proc. 2011 IEEE INFOCOM*, pp. 1–9.

[2] Z. Yang, B. Zhang, J. Dai, A. Champion, D. Xuan, and D. Li, "E-smalltalker: a distributed mobile system for social networking in physical proximity," in *2010 IEEE ICDCS*.

[3] R. Lu, X. Lin, X. Liang, and X. Shen, "A secure handshake scheme with symptoms-matching for mhealthcare social network," *2010 Mobile Netw. Applications*, pp. 1–12.

[4] C. Zhang, X. Zhu, Y. Song, and Y. Fang, "C4: a new paradigm for providing incentives in multi-hop wireless networks," in *Proc. 2011 IEEE INFOCOM*, pp. 918–926.

[5] M. Freedman, K. Nissim, and B. Pinkas, "Efficient private matching and set intersection," in *Proc. 2004 EUROCRYPT*, pp. 1–19.

[6] Q. Ye, H. Wang, and J. Pieprzyk, "Distributed private matching and set operations," in *Proc. 2008 ISPEC*, pp. 347–360.

[7] E. D. Cristofaro and G. Tsudik, "Practical private set intersection protocols with linear complexity," in *2010 Financial Cryptography and Data Security*.

[8] L. Kissner and D. Song, "Privacy-preserving set operations," in *Proc. 2005 CRYPTO*, pp. 241–257.

[9] A. C. Yao, "Protocols for secure computations," in *Proc. 1982 SFCS*, pp. 160–164.

[10] D. Dachman-Soled, T. Malkin, M. Raykova, and M. Yung, "Efficient robust private set intersection," in *Proc. 2009 ACNS*, pp. 125–142.

[11] G. S. Narayanan, T. Aishwarya, A. Agrawal, A. Patra, A. Choudhary, and C. P. Rangan, "Multi party distributed private matching, set disjointness and cardinality of set intersection with information theoretic security," in *Proc. 2009 CANS*, pp. 21–40.

[12] C. Hazay and Y. Lindell, "Efficient protocols for set intersection and pattern matching with security against malicious and covert adversaries," in *Proc. 2008 TCC*, pp. 155–175.

[13] S. Jarecki and X. Liu, "Efficient oblivious pseudorandom function with applications to adaptive OT and secure computation of set intersection," in *Proc. 2009 TCC*, pp. 577–594.

[14] W. Dong, V. Dave, L. Qiu, and Y. Zhang, "Secure friend discovery in mobile social networks," in *Proc. 2011 IEEE INFOCOM*, pp. 1–9.

[15] E. De Cristofaro, M. Manulis, and B. Poettering, "Private discovery of common social contacts," in *Proc. 2011 Applied Cryptography and Network Security*, pp. 147–165.
[16] E. De Cristofaro, A. Durussel, and I. Aad, "Reclaiming privacy for smartphone applications," in *Proc. 2011 IEEE International Conf. Pervasive Comput. Commun.*, pp. 84–92.
[17] A. Shamir, "How to share a secret," *Commun. ACM*, vol. 22, no. 11, pp. 612–613, 1979.
[18] M. Ben-Or, S. Goldwasser, and A. Wigderson, "Completeness theorems for non-cryptographic fault-tolerant distributed computation."
[19] R. Gennaro, M. O. Rabin, and T. Rabin, "Simplified vss and fast-track multiparty computations with applications to threshold cryptography," in *Proc. 1998 ACM PODC*, pp. 101–111.
[20] T. Nishide and K. Ohta, "Multiparty computation for interval, equality, and comparison without bit-decomposition protocol," in *Proc. 2007 PKC*, pp. 343–360.
[21] Y. Qi and M. J. Atallah, "Efficient privacy-preserving k-nearest neighbor search," in *Proc. 2008 IEEE ICDCS*, pp. 311–319.
[22] E. Kiltz, "Unconditionally secure constant round multi-party computation for equality, comparison, bits and exponentiation," in *2005 TCC*.
[23] M. Li, S. Yu, J. D. Guttman, W. Lou, and K. Ren, "Secure ad-hoc trust initialization and key management in wireless body area networks," *ACM Trans. Sensor Netw.*, 2012.
[24] Y. Hou, M. Li, and J. D. Guttman, "Chorus: scalable in-band trust establishment for multiple constrained devices over the insecure wireless channel," in *2013 ACM WiSec*.
[25] S. Gollakota, N. Ahmed, N. Zeldovich, and D. Katabi, "Secure in-band wireless pairing," in *Proc. 2011 USENIX Conf. Security*, pp. 16–16.
[26] O. Goldreich, *Foundations of Cryptography: Volume 2, Basic Applications*. Cambridge University Press, 2009.
[27] M. Li, S. Yu, N. Cao, and W. Lou, "Privacy-preserving distributed profile matching in proximity-based mobile social networks," technical Report, 2012. Available: http://digital.cs.usu.edu/~mingli/papers/Findu_techrep.pdf.
[28] G. Asharov and Y. Lindell, "A full proof of the BGW protocol for perfectly-secure multiparty computation," *2011 CRYPTO*.
[29] R. Canetti, "Security and composition of multiparty cryptographic protocols," *J. Cryptology*, vol. 13, pp. 143–202.
[30] "Ns2." Available: http://www.isi.edu/nsnam/ns.
[31] S. Bhatt, R. Sion, and B. Carbunar, "A personal mobile DRM manager for smartphones," *Comput. Security*, vol. 28, no. 6, pp. 327–340, 2009.
[32] A. Rahmati and L. Zhong, "Context-for-wireless: context-sensitive energy-efficient wireless data transfer," in *Proc. 2007 MobiSys*, pp. 165–178.
[33] R. Balani, "Energy consumption analysis for bluetooth, WiFi and cellular networks," Tech. Rep., Dec. 2007, pp. 1–6.
[34] A. Carroll and G. Heiser, "An analysis of power consumption in a smartphone," in *2010 Usenix Tech. Conf.*
[35] G. T. E. De Cristofaro and J. Kim, "Linear-complexity private set intersection protocols secure in malicious model." in *2010 Asiacrypt*.

Ming Li (S'08 - M'11) received his Ph.D. in Electrical and Computer Engineering from Worcester Polytechnic Institute, M.E and B.E in Electronic and Information Engineering from Beihang University in China. He joined the Computer Science Department at Utah State University as an assistant professor in 2011. His research interests are in the general areas of cyber security and privacy, with current emphases on data security and privacy in cloud computing, security in wireless networks and cyber-physical systems. He is a member of IEEE and ACM.

Shucheng Yu (S'07-M'10) received his Ph.D. in Electrical and Computer Engineering from Worcester Polytechnic Institute, a MS in Computer Science from Tsinghua University and a BS in Computer Science from Nanjing University of Post & Telecommunication in China. He joined the Computer Science department at the University of Arkansas at Little Rock as an assistant professor in 2010. His research interests are in the general areas of Network Security and Applied Cryptography. His current research interests include Secure Data Services in Cloud Computing, Attribute-Based Cryptography, and Security and Privacy Protection in Cyber Physical Systems. He is a member of IEEE and ACM.

Ning Cao (S'08-M'12) received his BE and ME degrees in Computer Science from Xi'an Jiaotong University in China, and his Ph.D. degree in Electrical and Computer Engineering from Worcester Polytechnic Institute. He joined the Research and System Infrastructure at Google Inc. in 2012. His research interests are in the areas of security, privacy, and reliability in Cloud Computing, with current focus on search and storage. He is a member of IEEE and a member of ACM.

Wenjing Lou (S'01-M'03-SM'08) is an associate professor at Virginia Polytechnic Institute and State University. Prior to joining Virginia Tech in 2011, she was on the faculty of Worcester Polytechnic Institute from 2003 to 2011. She received her Ph.D. in Electrical and Computer Engineering at the University of Florida in 2003. Her current research interests are in cyber security, with emphases on wireless network security and data security and privacy in cloud computing. She was a recipient of the U.S. National Science Foundation CAREER award in 2008.