

CERIAS Tech Report 2000-15

**Protocols for Secure Remote Database
Access with Approximate Matching**

Wenliang Du, Mikhail J. Atallah
Center for Education and Research in
Information Assurance and Security
Purdue University, West Lafayette, IN 47907

Protocols for Secure Remote Database Access with Approximate Matching *

Wenliang Du

CERIAS

Purdue University

West Lafayette, IN 47907

Email: duw@cerias.purdue.edu

Tel: (765) 496-6765

Mikhail J. Atallah

CERIAS

Purdue University

West Lafayette, IN 47907

Email: mja@cs.purdue.edu

Tel: (765) 494-6017

Abstract

Suppose that Bob has a database D and that Alice wants to perform a search query q on D (e.g., “is q in D ?”). Since Alice is concerned about her privacy, she does not want Bob to know the query q or the response to the query. How could this be done? There are elegant cryptographic techniques for solving this problem under various constraints (such as “Bob should know neither q nor the answer to the query” and “Alice should learn nothing about D other than the answer to the query”), while optimizing various performance criteria (e.g., amount of communication).

We consider the version of this problem where the query is of the type “is q approximately in D ?” for a number of different notions of “approximate”, some of which arise in image processing and template matching, while others are of the string-edit type that arise in biological sequence comparisons. New techniques are needed in this framework of approximate searching, because each notion of “approximate equality” introduces its own set of difficulties; using encryption is more problematic in this framework because the items that are approximately equal cease to be so after encryption or cryptographic hashing. Practical protocols for solving such problems make possible new forms of e-commerce between proprietary database owners and customers who seek to query the database, with privacy.

We first present four secure remote database access models that are used in the e-commerce, each of which has different privacy requirement. We then present our solutions for achieving privacy in each of these four models.

Keywords: Privacy, security, secure multi-party computation, pattern matching, approximate pattern matching.

*Portions of this work were supported by Grant EIA-9903545 from the National Science Foundation, and by sponsors of the Center for Education and Research in Information Assurance and Security.

1 Introduction

Consider the following real-life scenario: Alice thinks that she may have some genetic disease, so she wants to investigate it further. She also knows that Bob has a database containing known DNA patterns about various diseases. After Alice gets a sample of her DNA sequence, she sends it to Bob, who will then tell Alice the diagnosis. However, if Alice is concerned about her privacy, the above process is not acceptable because it does not prevent Bob from knowing Alice’s private information—both the query and the result.

This kind of situation, which is likely to arise as e-commerce develops, motivates the following general problem formulation:

Secure Database Access (SDA) Problem: Alice has a string s , and Bob has a database of strings $T = \{t_1, \dots, t_N\}$; Alice wants to know the result of whether there exists a string t_i in Bob’s database that “matches” s . The “match” could be an exact match or an approximate match. How to design a protocol that can accomplish this task without revealing Alice’s secret s to Bob?

Because of its practical importance and also because not much work has been done for approximate pattern matching in the SDA context, our work particularly focuses on approximate pattern matching.

The exact matching problem has been extensively considered in the literature [19, 6, 17, 16, 20, 22, 21, 13], even though it can theoretically be solved using the general techniques of secure multi-party computation [10]. The motivation for giving specialized solutions to it is that they are more *efficient* than those that follow from the above-mentioned general techniques. This is also our motivation in considering approximate pattern matching even though it too is a special case of the general secure multi-party computation problem. Unlike exact pattern matching that produces “yes” and “no” answers, approximate pattern matching measures the difference between the two targets, and produces a *score* to indicate how different the two targets are. The metrics used to measure the difference usually are heuristic and are application-dependent. For example, in image template matching [14, 18], $\sum_{i=1}^n (a_i - b_i)^2$ and $\sum_{i=1}^n |a_i - b_i|$ are used to measure the difference between two sequences a and b . In DNA sequence matching [15], *edit distance* [2] makes more sense than the above measurements; *edit distance* measures the cost of transforming one given sequence to another given sequence, and its special case, *longest common subsequence* is used to measure how similar two sequences are.

Solving approximate pattern matching problems within the SDA framework is quite a nontrivial task. Consider the $\sum_{i=1}^n |a_i - b_i|$ metric as an example. The known PIR (private information retrieval) techniques [19, 6, 17, 16, 20, 22, 21, 13] can be used by Alice to efficiently access each individual b_i without revealing to Bob anything about which b_i (or even which b) Alice accessed (more on this later), but doing this for each individual b_i and then calculating $\sum_{i=1}^n |a_i - b_i|$ violates the requirement that Alice should know the total score $\sum_{i=1}^n |a_i - b_i|$ *without knowing anything other than that score*, i.e., without learning anything about the individual b_i values. Using a general secure multi-party computation protocol typically does not lead to

an efficient solution. The goals of our research, and the results presented in this paper, are finding efficient ways to do such approximate pattern matchings without disclosing private information.

The actual practice of remote database access does not all fit into the same model we described in the above SDA formulation. For example, in some situations, Bob’s database could be proprietary whereas in some others it could be public (in either case the protocol should reveal nothing to Bob about Alice’s query). The “proprietary” nature of a database might make the solution more difficult because Alice should not be able to know more information than the response to her query. There is also another practical framework, within which Alice uses Bob to store a (suitable disguised) version of her private database (a form of outsourcing) and whose solutions could be much different. Based on these various practical variants of the problem, we have investigated four SDA models, and defined a class of SDA problems for each model according to the metrics we use for approximate pattern matching. Of course the difficulties of the problems are not the same for the different metrics, and so far we have solved a subset of those problems. A summary of our results is listed below (the results are stated more precisely in Section 4, and the models are defined in Section 3 – in the meantime see Figure 1 in that section for a summary of each model).

- For the Private Information Matching Model, we have a solution to the approximate pattern matching based on the $\sum_{i=1}^n (a_i - b_i)^2$ with $O(n^2 * N)$ communication cost, where n is the length of each pattern and N is the size of the database.
- For the Private Information Matching Model, We also have a solution to the approximate pattern matching based on $\sum_{i=1}^n |a_i - b_i|$ metrics using a Monte Carlo technique; the solution gives an estimate result, and it has $O(n * W * N)$ communication cost, where W is a parameter that affects the accuracy of the estimate.
- For the Private Information Matching Model, if we assume that the alphabet is known to the involved parties and its size is finite, we have a solution to approximate pattern matching based on general $\sum_{i=1}^n f(a_i - b_i)$ metrics, hence the solutions for the special cases of $\sum_{i=1}^n |a_i - b_i|$, $\sum_{i=1}^n (a_i - b_i)^2$, and $\sum_{i=1}^n \delta(a_i, b_i)$ (where $\delta(x, y)$ is 1 if $x = y$ and 0 otherwise). These solutions have $O(m * n * N)$ communication cost, where m is the number of the symbols in the alphabet. In many cases, m is small. For instance, m is four in DNA databases.
- For the Secure Storage Outsourcing Model, we have a solution to approximate pattern matching based on the $\sum_{i=1}^n (a_i - b_i)^2$ metrics. The solution is practical because it has only $O(n)$ communication cost, and $O(n)$ is optimal because that is how long the answer is.
- For the Secure Storage Outsourcing and Computation Model, we also have a solution to approximate pattern matching based on the $\sum_{i=1}^n (a_i - b_i)^2$ metrics. This solution is also practical because of its $O(n^2)$ communication cost.

Motivation

Why do we care about the privacy of a database query? In the example used earlier in this section, if a match is found in the database, Bob immediately knows that Alice has such a disease; even worse, after receiving Alice's DNA sequence, Bob can derive much about Alice from the DNA, such as other health problems that Alice might have. If Bob is not trustworthy, Bob could disclose the information about Alice to other parties, and Alice might have difficulty getting employment, insurance, credit, etc. On the other hand, even if Alice trusts Bob, and Bob has no intention of disclosing Alice's private information, Bob might still prefer that Alice's query be kept private out of liability concerns: If Bob knows Alice's DNA information, and that information is accidentally disclosed (perhaps by a disgruntled employee of Bob's, or after a system break-in), Bob might face an expensive lawsuit from Alice. From this perspective, a trusted Bob will actually prefer not to know either Alice's query or its response.

With the growth of the Internet, more and more e-commerce transactions like the above will take place. There are already DNA pattern databases, public databases about diseases, patent databases, and in the future we may see many more commercial databases and the related database access services, such as fingerprint databases, signature databases, medical record databases, and many more. Privacy will be a major issue in such e-commerce. Assuming the trustworthiness of the service providers, as is done today, is risky; therefore protocols that can support remote access operations while protecting the client's privacy are of growing importance.

One of the fundamental operations behind the queries described in the examples above is pattern matching. Therefore, the basic problem that we face is how to conduct pattern matching operations at the server side while the server has no knowledge of the client's actual query (or the response to it). In some database access situations, exact pattern matching is used, such as query by name, query by social security number, etc. However, in many other situations, exact pattern matching is unrealistic. For instance, in fingerprint matching, even if two fingerprints come from the same finger, they are unlikely to be exactly the same because there is some information loss in the process of deriving an electronic form (usually a complex data structure of features) from a raw fingerprint image. Similarly in voice, face, and DNA matching; in these and many other situations, exact matching is not expected and some form of approximate pattern matching is more useful.

Background Information on Secure Multi-party Computation

The above problem is a special case of the general secure multi-party computation problem [28]. Generally speaking, a multi-party computation problem deals with computing any probabilistic function on any input, in a distributed network where each participant holds one of the inputs, ensuring independence of the inputs, correctness of the computation, and that no more information is revealed to a participant in the computation than can be computed from that participant's input and output [12]. Other examples of such computations include: elections over the Internet, electronic bidding, joint signatures, and joint decryption. The history

of the multi-party computation problem is extensive since it was introduced by Yao [28] and extended by Goldreich, Micali, and Wigderson [23], and by many others: Goldwasser [12] predicts that “the field of multi-party computations is today where public-key cryptography was ten years ago, namely an extremely powerful tool and rich theory whose real-life usage is at this time only beginning but will become in the future an integral part of our computing reality”.

Goldreich states in [10] that the general secure multi-party computation problem is solvable in theory. However, Goldreich also points out that using the solutions derived by these general results for special cases of multi-party computation, are impractical; special solutions should be developed for special cases for efficiency reasons.

One of the well-known special cases of multi-party computation is the Private Information Retrieval (PIR) problem: The problem consists of a client and server. The client needs to get the i th bit of a binary sequence from the server without letting the server know the i ; the server does not want the client to know the binary sequence either. A solution for this problem is not difficult; however an efficient solution, in particular a solution with minimal communication cost, is not easy. Studies [19, 6, 17, 16, 20, 22, 21, 13] have shown that one can design a protocol to solve the PIR problem with much better communication complexity than the theoretical solutions. Pattern matching is another such specific computation, and the recent progress in the PIR problem motivated us to speculate that there exist solutions that are better than the general theoretical one for this particular kind of secure multi-party computation.

Secure Multi-party Protocol v.s. Anonymous Communication Protocol

Anonymous communication protocols [24, 11] were designed to achieve somewhat related goals, so why not use them? Anonymity techniques help to hide the identity of the information sender, rather than the information being sent. For example, when people browse the web, they can use anonymous communication protocols to keep their identities secret, but the web query usually is not secret because the web server has to know the query in order to send a reply back. In situations where the identity of the information sender needs to be protected, anonymous communication protocols are appropriate. However, there are situations where anonymous communication protocols cannot replace secure multi-party computation protocols. First, certain types of information intrinsically reveal the identity of someone related to the information (e.g., social security number). Secondly, in some situations, it is the information itself that needs to be protected, not the identity of the information sender. For instance, if Alice has an invention, she has to search if such an invention is new before she files for a patent. When conducting the query, Alice may want to keep the query private (perhaps to avoid part of her idea being stolen by people who have access to her query); she does not care whether her identity is revealed. Thirdly, in certain situations, one has to be a registered member in order to use the database access service; this makes hiding user’s identity difficult because the user has to register and login first, which might already disclose her identity.

Furthermore, most of the known practical anonymous protocols, such as Crowds [24], Onion routing

[11] and `Anonymizer.com` use one or several *trusted* third parties. In our secure multi-party computation protocols, we do not use a trusted third party; even if a third party is used, we generally assume that the third party is not trusted, and should learn nothing about either Alice’s query, or Bob’s data, or about the response to the query.

Therefore anonymity does not totally solve our problems, and cannot replace secure multi-party computation. Rather, by combining anonymity techniques with secure multi-party computation techniques, one can achieve better overall privacy more efficiently.

2 Related Work

As Goldwasser points out in [12], in the 80’s the focus of research was to show the most general result possible, yielding multi-party protocol solutions for any probabilistic function. Much of the current work is to focus on *efficient* and *non-interactive* solutions to special important problems such as joint-signatures, joint-decryption, and secure and private database access.

Among various multi-party computation problems, the Private Information Retrieval (PIR) problem has been widely studied; it is also the problem most related to what we present in this paper (although here we use none of the elegant techniques for PIR that are found in the literature for reasons we explained earlier in this paper). The PIR problem consists of devising a protocol involving a user and a database server, each having a secret input. The database’s secret input is called the *data string*, an N -bit string $B = b_1 b_2 \dots b_N$. The user’s secret input is an integer i between 1 and n . The protocol should enable the user to learn b_i in a communication-efficient way and at the same time hide i from the database.

The trivial solution is having the database send an encryption of the entire string B to the user. However, this solution is not efficient because of its $O(N)$ communication complexity. Much work has been done to reduce the communication complexity [19, 6, 17, 16, 20, 22, 21, 13]. Our work is motivated by this framework, including its emphasis on reducing communication complexity.

Chor *et al.* point out that a major drawback of all known PIR schemes is the assumption that the user knows the *physical address* of the sought item [9]. In the current database query scenario, the user typically holds a keyword and the database internally converts this keyword into a physical address. To solve this problem, Chor *et al.* propose a scheme to privately access data by keywords [9]. The difference between the problem studied in Chor’s paper and the problems in our paper is that we extend the problem to cover approximate pattern matching.

Song *et al.* propose a scheme to conduct searches on encrypted data [27]. The problem is that Alice has a database, and she has to store the database in a server controlled by Bob; how could Alice query her database without letting Bob know the contents of the database or the query? This problem is different from the PIR problem because Alice now knows all the inputs in this problem, whereas in the PIR problem Alice does not know Bob’s input. Here we primarily focus on extending the problem to also cover approximate

pattern matching.

There is much work on other types of secure multi-party computation problems, such as threshold cryptography [8], private bidding [5] and secret-ballot elections [4]. Although they are different from our problems, we believe that the techniques they use are also useful in solving our problems.

Multi-party protocols use a rich body of tools and sub-protocols, some of which were developed for particular applications, while others were developed for general cryptographic settings. These include zero-knowledge proofs, probabilistic encryption, oblivious transfer, various distributed commitment schemes [25], computing with shares of a secret [26], and instance hiding schemes [7, 1].

3 Framework

3.1 Models

Remote database access has many variants. In some e-commerce models, Bob’s database is private while in some other models, it is public. In the latter case, there is no requirement to keep the database secret from Alice; however, the privacy of Alice’s query still needs to be preserved. In other e-commerce models, Bob hosts Alice’s (encrypted/disguised) database while supporting queries from Alice and other customers, in which case Bob should know neither the database nor the queries.

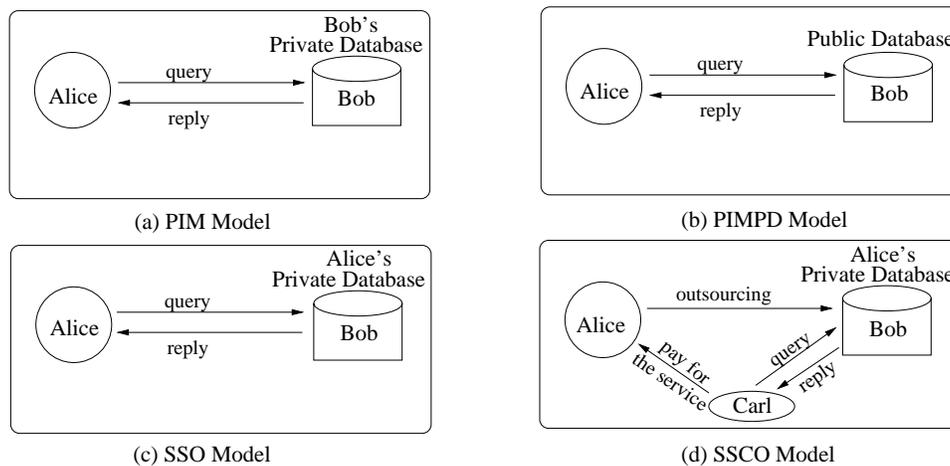


Figure 1: Models

From the various ways that remote database access is conducted, we distinguish four different e-commerce models, all of which require customers’ privacy:

- PIM: Private Information Matching Model (Figure 1.a)
- PIMPD: Private Information Matching from Public Database Model (Figure 1.b).
- SSO: Secure Storage Outsourcing Model (Figure 1.c).

- SSCO: Secure Storage and Computing Outsourcing Model (Figure 1.d).

For the sake of convenience, we will use $Match()$ to represent the pattern matching function, which includes both exact pattern matching and approximate pattern matching.

Private Information Matching Problem (PIM)

Alice has a string x , and Bob has a database of strings $T = \{t_1, \dots, t_N\}$; Alice wants to know the result of $Match(x, T)$. Because of the privacy concern, Alice does not want Bob to know the query x or the result; Bob does not want Alice to know any string in the database except for what can be derived from the reply. Furthermore, Bob wants to make money from providing such a service, therefore Alice should not be able to conduct the querying by herself; in other words, every time Alice wants to perform such a query, she has to contact Bob, otherwise she cannot get the correct answer.

Private Information Matching from Public Database Problem (PIMPD)

Bob has a database of strings $T = \{t_1, \dots, t_N\}$, whose contents are public knowledge. Alice has a query x , and she wants to know the result of $Match(x, T)$. However, because of the privacy concern, Alice does not want to disclose her query x to Bob.

This problem is different from the PIM problem: in the PIM problem, Bob does not allow Alice to know any information about the database except for what can be derived from the reply. In the PIMPD problem, since the database contains only public knowledge, there is no need to prevent Bob from letting Alice know more about the contents of the database than the strict answer to her query (although Bob's doing so may result in unnecessary communication).

Secure Storage Outsourcing Problem (SSO)

Alice has a database of strings $T = \{t_1, \dots, t_N\}$, but she does not have enough storage for the large database, so she outsources her database (suitably disguised—more on this later) to Bob, who provides enough storage for Alice. Furthermore, from time to time, Alice needs to query her database and retrieves the information that matches her query, i.e., Alice wants to know $Match(x, T)$ for her query x . For the sake of privacy, Alice wants to keep the contents of both the database and the query secret from Bob.

Secure Storage and Computing Outsourcing Problem (SSCO)

The SSCO problem is an extension of the SSO problem. While the database is exclusively queried by Alice only in the SSO problem, in the SSCO model the database will also be queried by other clients of Alice. More specifically, in the SSCO model, Alice outsources her database to Bob, and she wants the database to be available to anyone who is willing to pay her for the database access service. When a client accesses the database, neither Alice nor Bob should know the contents of the query. Moreover, Alice wants to charge the

clients for each query they have submitted, so the client should not be able to get the correct query result if Alice is not aware of the query's existence.

Since Bob can pretend to be a client, the solutions of the SSCO problem should be secure even if Bob can collude against Alice with any client. However, the SSO problem does not have such a concern because the only client is Alice herself.

3.2 Formalized Problems

For each model, there is a family of problems. We will use the following notations to represent each specific problem:

- *M/Exact*: Exact Pattern Matching problem in model *M*.
- *M/Approx*: Approximate Pattern Matching problem in model *M*.
 - *M/Approx/f*: use $\sum_{k=1}^n f(a_k, b_k)$ metrics to measure the distance between two strings, where *f* is a general function.
 - *M/Approx/δ*: use $\sum_{k=1}^n \delta(a_k, b_k)$ metrics to measure the distance between two strings, where δ is the Kronecker symbol: $\delta(x, y) = 0$ if and only if $x = y$ and 1 otherwise.
 - *M/Approx/Abs*: use $\sum_{k=1}^n |a_k - b_k|$ metrics to measure the distance between two strings.
 - *M/Approx/Squ*: use $\sum_{k=1}^n (a_k - b_k)^2$ metrics to measure the distance between two strings.
 - *M/Approx/Edit*:
 - * *M/Approx/Edit/String*: use the string editing criterion to measure the distance between two strings.
 - * *M/Approx/Edit/Tree*: use the tree editing criterion to measure the distance between two trees.

The *M/Exact* problem has been studied extensively in certain model, such as PIM and SSO, but the *M/Approx* problem has not. Our results deal mostly with the *M/Approx* problem.

4 Our Results

4.1 PIM/Approx

Except for the research on the general secure multi-party computation problem, this specific problem has not been studied in the literature. Unless otherwise specified, we assume alphabet used in the following solution to be predefined and its size to be finite. This assumption is quite reasonable in many situations; for instance, DNA sequences use a fixed alphabet of four symbols. Under this assumption, we can solve the

PIM/Approx/ f problem. However, because the way to calculate *edit distance* cannot be represented in the form $\sum_{k=1}^n f(a_k, b_k)$, the PIM/Approx/Edit problem is not a special case of the PIM/Approx/ f problem. We also have a solution for PIM/Approx/Edit/String problem, but because of its complexity and space limitation, we will leave the solution to the journal version of this paper.

In some other situations, the above finite alphabet assumption does not apply. For instance, fingerprint, image and voice patterns use real numbers instead of characters from a known finite alphabet. The above-mentioned solution for the PIM/Approx/ f problem cannot be used anymore, however by exploiting the mathematical property of $\sum_{i=1}^n (a_i - b_i)^2$, we have come up with a solution for the PIM/Approx/Squ problem for infinite alphabet after introducing an *untrusted* third party who does not know the inputs from either of the two parties and learns nothing about them (or about the query, or the answer to it). We also have a solution to the PIM/Approx/Abs problem using a Monte Carlo technique. All of these are given below.

4.1.1 PIM/Approx/Squ Protocol

Suppose that Bob has a database $T = \{t_1, \dots, t_N\}$, and assume the length of each element is n ; Alice wants to know the $t_i \in T$ that most closely matches a query $x = x_1 \dots x_n$ based on the PIM/Approx/Squ metrics. The requirement is that Bob should not know x or the result, and Alice should not be able to learn more information than the reply from Bob.

We propose a protocol to compute the matching score using a untrusted third party, Ursula. Our assumption here is that Ursula cannot conspire with either Alice or Bob. However, the third party is not considered as trusted; therefore, Ursula should not be able to deduce either x or T , or the final matching score s . This protocol works for both finite and infinite alphabet.

Let $\vec{x} = (-2x_1, \dots, -2x_n, 1, R_A, 1)$, where R_A is a random number generated by Alice; for each $t_i = y_{i,1} \dots y_{i,n}$, let $\vec{z}_i = (y_{i,1}, \dots, y_{i,n}, \sum_{k=1}^n y_{i,k}^2 - R_i, 1, R_i)$, where R_A and R_i are random numbers. Observe that:

$$\sum_{k=1}^n (x_k - y_{i,k})^2 = \vec{x} \cdot \vec{z}_i^T + \left(\sum_{k=1}^n x_k^2 - R_A \right).$$

Since $(\sum_{k=1}^n x_k^2 - R_A)$ is a constant, we can use $\vec{x} \cdot \vec{z}_i^T$ instead of $\sum_{k=1}^n (x_k - y_{i,k})^2$ to find the closest match. After we get the closest match, Alice can calculate the actual score because she knows both $\sum_{k=1}^n x_k^2$ and R_A .

Protocol

1. Alice generates a random number R_A , and constructs the vector $\vec{x} = (-2x_1, \dots, -2x_n, 1, R_A, 1)$.
2. Alice generates an $(n + 3) \times (n + 3)$ matrix M , where

$$\mathbf{M} = \begin{pmatrix} \vec{x} \\ R \end{pmatrix}$$

where R is a matrix of size $(n + 2) \times (n + 3)$, each element of which is a random number.

3. Alice generates a random invertible matrix Q of size $(n + 3) \times (n + 3)$. We will use vector \vec{q}_k to represents the k th row of Q^{-1} .
4. Alice sends the result of $Q * M$ to to Bob.
5. Alice sends \vec{q}_1 to Ursula.
6. For each $t_i \in T$, repeat the next two sub-steps, in which $t_i = y_{i,1} \dots y_{i,n}$.
 - (a) Bob constructs $\vec{z}_i = (y_{i,1}, \dots, y_{i,n}, \sum_{k=1}^n y_{i,k}^2 - R_i, 1, R_i)$, and calculates $(QM)\vec{z}_i^T$, then sends the result to Ursula.
 - (b) Ursula calculates $v_i = \vec{q}_1 \cdot (QM\vec{z}_i^T) = \vec{x} \cdot \vec{z}_i^T$.
7. Ursula computes $score' = \min_{i=1}^N v_i$, and sends the result $score'$ to Alice.
8. Alice computes $score = score' + \sum_{k=1}^n x_k^2 - R_A$, which is the closest match between x and the any $t_i \in T$.

The purpose of R_A is to prevent Ursula from knowing the actual score, and the purpose of Q and R is to disguise the query x . Alice does not need to put x in the first row of M , instead, she can put it in any row of M , and then sends to Ursula the corresponding row of Q^{-1} ; only Alice knows which row of M is vector x . The communication cost of the above protocol is $O(n^2 * N)$.

4.1.2 PIM/Approx/Abs Protocol

First, we will present a Monte Carlo technique for Alice and Bob to calculate $|x_k - y_k|$, and then use it as a building block to compute $\sum_{k=1}^n |x_k - y_k|$. The protocol involves an untrusted third party, Ursula, who learns nothing except $|x_k - y_k| + R_k$, where R_k is a random number unknown to her. The protocol works for both finite and infinite alphabet. Assume that $0 < x_k \leq U$ and $0 < y_k \leq U$ for some number U . The protocol for $|x_k - y_k|$ is (in what follows W is a parameter that affects the accuracy of the estimate, and $counter = 0$ initially):

1. Alice generates a random number R_k , and then generates $W - R_k$ random numbers uniformly over $(0..U)$.

2. Alice randomly replaces half of these $W - R_k$ numbers with their negative values.
3. Alice inserts R_k zeroes into random positions of these $W - R_k$ numbers, resulting in a sequence S of W numbers.
4. Alice then sends S to Bob.
5. For each number s from S , if $s = 0$, Alice sends 1 to Ursula; if $s > 0$, Alice sends 1 to Ursula if $|s| \geq x_k$, sends 0 otherwise; if $s < 0$, Alice sends 0 to Ursula if $|s| \geq x_k$, sends 1 otherwise.
6. For each number s from S , if $s = 0$, Bob sends 0 to Ursula; if $s > 0$, Bob sends 1 to Ursula if $|s| \geq y_k$, sends 0 otherwise; if $s < 0$, Bob sends 0 to Ursula if $|s| \geq y_k$, sends 1 otherwise.
7. Ursula increases *counter* by 1 if the values she receives from Alice and Bob are different.
8. Ursula computes $score = counter * \frac{U}{W}$, which is shown earlier to be an unbiased estimate of $|x_k - y_k| + R_k * \frac{U}{W}$.

Because of R_k , Ursula does not know the actual distance between x_k and y_k , and because of the negative numbers among those W random numbers, Ursula can not figure out whether $x_k > y_k$ or $x_k < y_k$. Therefore Ursula knows nothing about x_k and y_k .

Now, let us see how to use the above protocol to compute $\sum_{k=1}^n |x_k - y_{i,k}|$, where $x = x_1 \dots x_n$ and $t_i = y_{i,1} \dots y_{i,n}$:

1. Alice generates a random number R .
2. For each $t_i \in T$, suppose $t_i = y_{i,1} \dots y_{i,n}$ and repeat the next three sub-steps:
 - (a) *counter* = 0.
 - (b) For each $k = 1, \dots, n$, Alice, Bob and Ursula use the above protocol to compute $|x_k - y_{i,k}|$. The random numbers $R_{i,1}, \dots, R_{i,n}$ used in the above protocol are generated by Alice, such that $\sum_{k=1}^n R_{i,k} = R$.
 - (c) Ursula computes $score_i = counter * \frac{U}{W}$, which is an unbiased estimate of $\sum_{k=1}^n |x_k - y_{i,k}| + \sum_{k=1}^n R_{i,k} * \frac{U}{W} = \sum_{k=1}^n |x_k - y_{i,k}| + R * \frac{U}{W}$.
3. Ursula computes $score' = \min_{i=1}^N score_i$, and sends $score'$ to Alice.
4. Alice computes $score = score' - R * \frac{U}{W}$ and gets the closest match between x and any $t_i \in T$.

The communication complexity is $O(n * W * N)$. The analysis of the variance will given in the full version of this paper.

4.1.3 PIM/Approx/ f protocol

If the alphabet is predefined and its size is finite, we can solve a general problem—computing $f(x_k, y_k)$. However, we cannot directly use this protocol n times to compute $\sum_{k=1}^n f(x_k, y_k)$ because that would reveal each individual result of $f(x_k, y_k)$. We will present the protocol for computing $f(x_k, y_k)$ here, and then in the following sub-section, we will discuss how to use it as a building block to compute $\sum_{k=1}^n f(x_k, y_k)$ without revealing any individual $f(x_k, y_k)$.

Suppose Alice has an input x_k ; Bob has an input y_k ; Alice wants to know the result of $f(x_k, y_k)$ without revealing x_k and the result to Bob, and Bob does not want to reveal its y_k to Alice. If Alice can derive y_k from $f(x_k, y_k)$, that is beyond the scope of this problem. We present a solution to this problem. Later we will use this solution as a building block to construct solutions to other problems.

f -function Protocol

We assume the encryption methods used below are commutative.

1. Bob computes $f(\alpha_i, y_k)$ for each $\alpha_i \in X$, where X is the finite (known) alphabet. Let m be the size of X .
2. Bob chooses a secret key k , computes $E_k(f(\alpha_i, y_k))$ for each $\alpha_i \in X$, and sends to Alice the m results.
3. Alice chooses one from $E_k(f(\alpha_i, y_k))$, $i = 1 \dots m$, such that $\alpha_i = x_k$. This can be done because Bob sent the m encrypted results in order.
4. Alice chooses a secret key k' , computes $E_{k'}(E_k(f(x_k, y_k)))$, and sends it back to Bob.
5. Because of the commutative properties of $E_{k'}$ and E_k , $E_{k'}(E_k(f(x_k, y_k)))$ is equivalent to $E_k(E_{k'}(f(x_k, y_k)))$, which could be decrypted to $E_{k'}(f(x_k, y_k))$ by Bob. Bob sends the result $E_{k'}(f(x_k, y_k))$ to Alice.
6. Alice gets $f(x_k, y_k)$ by decrypting $E_{k'}(f(x_k, y_k))$.

The technique used above is similar to the standard oblivious transfer protocol; it protects the privacy of the inputs from both parties without introducing a third-party. The communication cost is $O(m)$, where m is the size of the alphabet.

PIM/Approx/ f Protocol

First, let us see how to securely compute $\sum_{k=1}^n f(x_k, y_k)$. As we discussed above, we cannot run the above f -function protocol n times to get $\sum_{k=1}^n f(x_k, y_k)$. In the following protocol, we will use a disguise technique to hide each individual result of $f(x_k, y_k)$.

For each $t_i = y_{i,1} \dots y_{i,n}$, and for each $k = 1, \dots, n$, let $f_{i,k}(x_k, y_{i,k}) = f(x_k, y_{i,k}) + R_{i,k}$, where $R_{i,k}$ is a random number, the following protocol shows how A and B calculate $\sum_{k=1}^n f(x_k - y_{i,k})$:

1. Bob generates a random number R then sends R to Alice.
2. For each $t_i = y_{i,1}, \dots, y_{i,n}$, repeat the next five sub-steps:
 - (a) Bob constructs $f_{i,k}(x_k, y_{i,k}) = f(x_k, y_{i,k}) + R_{i,k}$ for $k = 1, \dots, n$, where $R_{i,1}, \dots, R_{i,n}$ are n random numbers.
 - (b) Alice and Bob use the f -function protocol to compute $f_{i,k}(x_k, y_{i,k})$, for each $k = 1, \dots, n$.
 - (c) Alice sends $\sum_{k=1}^n f_{i,k}(x_k, y_{i,k})$ to Ursula.
 - (d) Bob sends $\sum_{k=1}^n R_{i,k} - R$ to Ursula.
 - (e) Ursula computes $score_i = \sum_{k=1}^n f_{i,k}(x_k, y_{i,k}) - (\sum_{k=1}^n R_{i,k} - R) = \sum_{k=1}^n f(x_k, y_{i,k}) + R$.
3. Ursula computes $score' = \min_{i=1}^N score_i$, and sends $score'$ to Alice.
4. Alice compute $score = score' - R$, thus getting the actual distance between x and the closest t_i in the database T .

Although Alice knows each individual $f_{i,k}(x_k, y_{i,k})$, she does not know the actual value of $f(x_k, y_{i,k})$ because of $R_{i,k}$. Similarly, because of R , Ursula does not know the actual score of the closest match. The communication cost of the protocol is $O(m * n * N)$, where m is the size of the alphabet, n is the length of each pattern, and N is the size of the database. In many cases, m is quite small. For instance, m is four in DNA databases.

Because $|x_k - y_k|$, $(x_k - y_k)^2$ and $\delta(x_k, y_k)$ functions are special cases of $f(x_k, y_k)$, PIM/Approx/(Abs, Squ, δ) problems can all be solved using the above protocol.

4.2 PIMPD/Approx

The only difference between the PIM model and the PIMPD model is that, in the latter, Bob does not need to keep the database secret from Alice. Therefore, All solutions in the PIM model can be applied to the PIMPD model as well. Whether the “public” feature of the database can result in more efficient solutions is an interesting question. Although we do not yet have an answer to it, we observed the following:

Theorem 1. *There is no secure two-party non-interactive solution for the PIMPD/Approx problem.*

Proof. A two-party non-interactive protocol means Bob, by himself, is able to find the item in the database that has minimal distance from the query.

Assume there is a two-party non-interactive protocol A which solves any of the PIMPD/Approx problems, in another words, given an encrypted/disguised form (q') of a query q , and the database T that Bob

knows, Bob can find the item in the database that has minimal distance from q as follows. We use $A(T, q)$ to represent the algorithm on input T and q .

Since Bob can use any database he wants, he can use a database like this: $T = \{\text{"axxxxxx"}, \text{"bxxxxxx"}, \dots, \text{"zxxxxxx"}\}$, supposing that the alphabet is a set from 'a' to 'z'. After applying $A(T, q)$, Bob will get one that has the minimal distance from q . For instance, if "mxxxxxx" is the result, Bob knows that 'm' is the first character in q . Since A is a non-interactive protocol, Bob can reuse it on another database constructed for the purpose of exposing the second character in q ; he can keep doing this and figure out the rest of the characters in q .

Therefore, if such a protocol existed, the query q would not be kept secret from Bob. \square

The theorem does not rule out the existence of an efficient interactive protocol or a multi-party protocol.

4.3 SSO/Approx

In this model, Bob is a service provider who provides storage and database query services to Alice. According to Alice's privacy requirement, Bob should know nothing about the database that he stores for Alice, nor should he know the query. So Bob has to conduct a database query based on the encrypted or disguised data of Alice.

The requirement that Bob should not know the query result, as in the PIM and PIMPD problem, is not needed anymore in the SSO problem. The reason is that Bob does not know the contents of the database, he does not even know what the database is for, so knowing whether Alice's query is in the database does not disclose any secret information to Bob.

Intuitively, it can look like that the SSO/Approx problem might be more difficult than the PIM/Approx problem because Bob at least knows the contents of the database in the PIM/Approx problem whereas he knows nothing about the database in the SSO/Approx problem. But knowing the contents of the database has a disadvantage, in that Bob cannot know an intermediate result because he knows one of the inputs (the database); if he also knows an intermediate result, he might be able to figure out the other input (query) of the computation. However, in the SSO/Approx problem, Bob knows nothing about the database, so it is safe for him to know intermediate results without exposing the query information.

Whether Bob can know intermediate results is a critical issue to reduce the communication complexity. If he knows intermediate results to some extent, he can conduct the comparison operation to find the minimal or maximal score; otherwise, he has to turn to Alice in order to find the minimal or maximal score, which results in high communication cost in the PIM problem.

The SSO/Approx problem is similar to secure outsourcing of scientific computations problems studied by Atallah *et al.* [3]. The difference is that in secure outsourcing problems, inputs are provided by Alice every time a computation is conducted in Bob's side; therefore, Alice can encrypt/disguise the inputs differently in different rounds of the computation. However, in the SSO problem, one of the inputs (the database)

is encrypted/disguised only once, and this same input is used in all rounds of computations; this makes the problem more difficult.

So far, we have a solution only for SSO/Approx/Squ problem. The solution works for both infinite and finite alphabet.

4.3.1 SSO/Approx/Squ Protocol

Suppose that Alice wants to outsource her database $T = \{t_1, \dots, t_N\}$ to Bob, and wants to know if query string $x = x_1 \dots x_n$ matches any pattern t_i in the database T .

The straightforward solution would be to let Bob send the whole database back to Alice, and let Alice conduct the query by herself. Although this solution satisfies the privacy requirement, much better communication complexity can be achieved. Another intuitive question would be whether Bob can conduct the matching independently after Alice sends him the relevant information about the query. If the answer is true, Bob should be able to find the item t_i that has the closest match to the query x . In another words, if $t_i = y_1 \dots y_n$ and $score_i = \sum_{k=1}^n (x_k - y_k)^2$, then Bob should be able to find the minimum value of $score_i$. However, because of the privacy requirement, Bob is not allowed to know the actual query x , nor is he allowed to know the content of the database, so how does he compute the distance $score_i$ between x and each of the element t_i in the database?

The idea behind our solution is based on the fact that $\vec{x} \cdot \vec{z}^T = (\vec{x}Q^{-1}) \cdot (Q\vec{z}^T)$, where Q is an invertible matrix. Alice can store $Q\vec{z}^T$ instead of \vec{z}^T at Bob's site, and keeps Q secret from Bob. She will send $\vec{x}Q^{-1}$ to Bob each time she wants to send a query x ; therefore Bob can compute $\vec{x} \cdot \vec{z}^T$ without even knowing \vec{x} and \vec{z} . If we can use $\vec{x} \cdot \vec{z}^T$ to represent the $\sum_{k=1}^n (x_k - y_k)^2$, we can make it possible for Bob to conduct the approximate pattern matching.

For each $t_i = y_{i,1} \dots y_{i,n}$ in the database T , let $\vec{t}_i = (\sum_{k=1}^n y_{i,k}^2 + R - R_i, y_{i,1}, \dots, y_{i,n}, 1, R_i)$, and let $\vec{x} = (1, -2x_1, \dots, -2x_n, R_A, 1)$, where R, R_A and R_i are random numbers. We will have $\vec{x} \cdot \vec{t}_i^T = \sum_{k=1}^n y_{i,k}^2 - 2 \sum_{k=1}^n x_k y_{i,k} + R + R_A$, and therefore $score_i = \sum_{k=1}^n (x_k - y_{i,k})^2 = \vec{x} \cdot \vec{t}_i^T + (\sum_{k=1}^n x_k^2 - R - R_A)$. Since $(\sum_{k=1}^n x_k^2 - R - R_A)$ is a constant, it does not affect the final result if we only want to find the t_i that produces the minimum $score_i$. Therefore, Bob can use $\vec{x} \cdot \vec{t}_i^T$ to compute the closest match.

Before outsourcing the database to Bob, Alice randomly chooses a secret $(n+3) \times (n+3)$ invertible matrix Q , and computes $\vec{z}_i = Q\vec{t}_i^T$, then sends $T' = \{\vec{z}_1, \dots, \vec{z}_N\}$ to Bob.

Protocol

1. For any query string $x = x_1 \dots x_n$, Alice generates a random number R_A , and constructs a vector $\vec{x} = (1, -2x_1, \dots, -2x_n, R_A, 1)$, then sends $\vec{x}Q^{-1}$ to Bob.
2. Bob computes $score'_i = \vec{x} \cdot \vec{z}_i^T$, for $i = 1, \dots, N$.
3. Bob computes $\min_{i=1}^N score'_i$, and gets the corresponding i .

4. Bob returns \vec{z}_i to Alice.
5. Alice computes $Q^{-1}\vec{z}_i$ and gets t_i , which is the closest match of her query.

Because Alice and Bob are involved in only one round of communication, the communication cost is $O(n)$, which is optimal because that is how long the answer is.

Notice that we have introduced random numbers R, R_A, R_i for $i = 1 \dots N$. The purpose of R is to prevent Bob from knowing the actual distance between x and the items in the database; the purpose of R_A is to prevent Bob from knowing the relationship between two different queries; the purpose of R_i is to prevent Bob from knowing the relationship among items in the database. Without R_i , two similar items in the database T would still be similar to each other in the disguised database T' ; adding a different random number to each different item will make this similarity disappear.

4.4 SSCO/Approx

This model poses more challenges than the SSO model because Bob could now collude against Alice with a client, or he can even become a client. Therefore, one of the threats would be whether Bob is able to compromise the privacy of the database by conducting a small number of queries and deriving the way the database is encrypted or disguised. A secure protocol should resist this type of active attack. We have an solution for the SSCO/Approx/Squ problem that works for both infinite and finite alphabet.

4.4.1 SSCO/Approx/Squ protocol

One of the difference between the SSCO/Approx problem and the SSO/Approx problem is who sends the query. In the SSO/Approx/Squ protocol, Alice transforms the query x to a vector $\vec{x}Q^{-1}$, and sends the vector to Bob; in the SSCO/Approx/Squ protocol, the client Carl will send the query. Because Carl does not know Q , he cannot construct $\vec{x}Q^{-1}$ by himself. If Carl can get the result of $\vec{x}Q^{-1}$ securely, namely without disclosing \vec{x} to Alice and without knowing Q of course, we will have a solution. Because $Q^{-1} = (\vec{q}_1^T, \dots, \vec{q}_m^T)$, computing $\vec{x}Q^{-1}$ securely is basically a task of computing $\vec{x} \cdot \vec{q}_k^T$ for $k = 1..m$, which can be solved using the same technique as that used in solving PIM/Approx/Squ problem.

Therefore, by modifying step 2 of the SSO/Approx/Squ protocol slightly, and also by using a form of “ $R_\alpha * (score + R_A)$ ”, instead of the form of “ $score + R_A$ ” as is used in SSO/Approx/Squ protocol, we have a SSCO/Approx/Squ protocol as the following:

Let $T = \{t_1, \dots, t_N\}$ be the database Alice wants to outsource to Bob, and assume the length of each element is n . Alice generates N random numbers R_1, \dots, R_N . For each $t_i = y_{i,1}, \dots, y_{i,n}$, let $\vec{t}_i = (\sum_{k=1}^n y_{i,k}^2 + R - R_i, y_{i,1}, \dots, y_{i,n}, 1, 1, R_i)$; let $\vec{z}_i = Q\vec{t}_i^T$, where Q is a randomly generated $(n+4) \times (n+4)$ matrix.

In what follows, we assume that Alice outsourced the database $T' = \{\vec{z}_1, \dots, \vec{z}_N\}$ to Bob.

Protocol

1. Whenever a client Carl wants to conduct a search on query $x = x_1 \dots x_n$, he generates a random number R_C .
2. Alice generates random numbers R_A and R_α .
3. Carl and Alice jointly compute $\vec{q} = R_\alpha \vec{x} Q^{-1}$, where $\vec{x} = (1, -2x_1, \dots, -2x_n, R_C, R_A, 1)$. The computation does not reveal Alice's secret Q , R_A or R_α to Carl, nor does it reveal Carl's private query x or R_C to Alice.
4. Carl then sends the vector \vec{q} to Bob.
5. Bob computes $score_i = \vec{q} \cdot \vec{z}_i^T = R_\alpha (\sum_{k=1}^n y_{i,k}^2 - 2 \sum_{k=1}^n x_k y_{i,k} + R_C + R_A)$
6. Bob returns to Alice $score' = \min_{i=1}^N score_i$.
7. Alice computes $score'' = \frac{score'}{R_\alpha} - R_A = \sum_{k=1}^n y_{i,k}^2 - 2 \sum_{k=1}^n x_k y_{i,k} + R_C$ and sends it to Carl.
8. Carl computes $score = score'' + \sum_{k=1}^n x_k^2 - R_C$ and gets the final score.

Because of R_C , Alice cannot figure out the actual score for this query, and because of R_A and R_α , Carl cannot figure out the actual score between his query and other items in the database (except for the matched one) even if Carl can collude with Bob. The communication cost of the protocol is $O(n^2)$, most of which is contributed by the computation of $R_\alpha \vec{x} Q^{-1}$ in step 3.

5 Conclusion and Future Work

We have developed four models for secure remote database access, and presented a class of problems and solutions for these models. For some problems, such as SSO/Approx/Squ and SSCO/Approx/Squ problems, our solutions are practical, and they only need $O(n)$ and $O(n^2)$ communication cost, respectively; while for PIM/Approx and PIMPD/Approx problems, our results are still at the theoretical stage because of their high communication cost. Improving the communication cost for those solutions is one avenue for future work; another avenue is the non-sequential pattern matching: the pattern matching problems that we have discussed only involve patterns of simple sequential structure; in many applications, patterns have a branching structure, such as a tree or a DAG. The M /Approx/Edit/Tree problem in our model is one of the examples. Developing a secure protocol to deal with this type of query is a challenging problem.

References

- [1] M. Abadi and J. Feigenbaum. Secure circuit evaluation: a protocol based on hiding information from an oracle". *Journal of Cryptology*, 2:1–12, 1990.

- [2] A. Apostolico and Z. Galil, editors. *Pattern Matching Algorithms*. Oxford University Press, 1997.
- [3] M. Atallah and J. Rice. Secure outsourcing of scientific computations. Technical Report COAST TR 98-15, Department of Computer Science, Purdue University, 1998.
- [4] J. Benaloh and M. Yung. Distributing the power of a government to enhance the privacy of voters. In *Proceedings of the fifth annual ACM symposium on Principles of distributed computing*, pages 52–62, Calgary, Alta, Canada, August 11 - 13 1986.
- [5] C. Cachin. Efficient private bidding and auctions with an oblivious third party. In *Proceedings of the 6th ACM conference on Computer and communications security*, pages 120–127, Singapore, November 1-4 1999.
- [6] B. Chor and N. Gilboa. Computationally private information retrieval (extended abstract). In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, El Paso, TX USA, May 4-6 1997.
- [7] M. Abadi, J. Feigenbaum and J. Kilian. On hiding information from an oracle. *Journal of Computer and System Sciences*, 39:21–50, 1989.
- [8] P. Gemmell. An introduction to threshold cryptography. In *CryptoBytes*, volume 2. RSA Laboratories, 1997.
- [9] B. Chor, N. Gilboa and M. Naor. Private information retrieval by keywords. Technical Report TR CS0917, Department of Computer Science, Technion, 1997.
- [10] O. Goldreich. Secure multi-party computation (working draft). Available from <http://www.wisdom.weizmann.ac.il/home/oded/public.html/foc.html>, 1998.
- [11] P. F. Syverson, D. M. Goldschlag and M. G. Reed. Anonymous connections and onion routing. In *Proceedings of 1997 IEEE Symposium on Security and Privacy*, Oakland, California, USA, May 5-7 1997.
- [12] S. Goldwasser. Multi-party computations: Past and present. In *Proceedings of the sixteenth annual ACM symposium on Principles of distributed computing*, Santa Barbara, CA USA, August 21-24 1997.
- [13] Y. Gertner, S. Goldwasser and T. Malkin. A random server model for private information retrieval. In *2nd International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM '98)*, 1998.
- [14] R. Gonzalezi and R. Woods. *Digital Image Processing*. Addison-Wesley, Reading, MA, 1992.
- [15] D. Gusfield. *Algorithms on Strings, Trees, and Sequences: Computer Science and Computational Biology*. Cambridge University Press, 1997.
- [16] G. Di-Crescenzo, Y. Ishai and R. Ostrovsky. Universal service-providers for database private information retrieval. In *Proceedings of the 17th Annual ACM Symposium on Principles of Distributed Computing*, September 21 1998.
- [17] Y. Ishai and E. Kushilevitz. Improved upper bounds on information-theoretic private information retrieval (extended abstract). In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, Atlanta, GA USA, May 1-4 1999.
- [18] A. Jain. *Fundamentals of Digital Image Processing*. Prentice Hall, Englewood Cliffs, NJ, 1989.
- [19] B. Chor, O. Goldreich, E. Kushilevitz and M. Sudan. Private information retrieval. In *Proceedings of IEEE Symposium on Foundations of Computer Science*, Milwaukee, WI USA, October 23-25 1995.
- [20] E. Kushilevitz and R. Ostrovsky. Replication is not needed: Single database, computationally-private information retrieval. In *Proceedings of the 38th annual IEEE computer society conference on Foundation of Computer Science*, Miami Beach, Florida USA, October 20-22 1997.
- [21] Y. Gertner, Y. Ishai, E. Kushilevitz and T. Malkin. Protecting data privacy in private information retrieval schemes. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, Dallas, TX USA, May 24-26 1998.

- [22] C. Cachin, S. Micali and M. Stadler. Computationally private information retrieval with polylogarithmic communication. *Advances in Cryptology: EUROCRYPT '99, Lecture Notes in Computer Science*, 1592:402–414, 1999.
- [23] O. Goldreich, S. Micali and A. Wigderson. How to play any mental game. In *Proceedings of the 19th annual ACM symposium on Theory of computing*, pages 218–229, 1987.
- [24] M. K. Reiter and A. D. Rubin. Crowds: anonymity for web transaction. *ACM Transactions on Information and System Security*, 1(1):Pages 66–92, 1998.
- [25] B. Schneier. *Applied Cryptography: Protocols, Algorithms, and Source Code in C*. John Wiley & Sons, Inc., 1996.
- [26] A. Shamir. How to share a secret. *Communication of the ACM*, 22(11):612–613, 1979.
- [27] D. Song, D. Wagner and A. Perrig. Practical techniques for searches on encrypted data. In *Proceedings of 2000 IEEE Symposium on Security and Privacy*, Oakland, California, USA, May 14-17 2000.
- [28] A. Yao. Protocols for secure computations. In *Proceedings of the 23rd Annual IEEE Symposium on Foundations of Computer Science*, 1982.