MITSUBISHI ELECTRIC RESEARCH LABORATORIES
http://www.merl.com

An Attribute-Based Framework for Privacy Preserving Image Querying

Rane, S.; Sun, W.

TR2012-066 September 2012

Abstract

We are specifically concerned with scenarios in which multimedia data is stored once on the server and the same data is queried by multiple parties. We propose a framework for privacy preserving querying, in which encryption is performed only once, and the ciphertexts are stored on a database server. Rather than using public-key homomorphic cryptosystems, the parties querying the database first derive an "attribute" from their query signal. They can decrypt the server's ciphertext only if their attribute satisfies a specified mathematical condition. This query-specific decryption capability makes attribute based cryptography a vital addition to the secure signal processor's toolkit. We give an example of a construction for privacy preserving querying, in which a client can privately retrieve an image from the server if attribute vectors extracted from the server's and client's images are close enough in Euclidean distance.

IEEE International Conference on Image Processing (ICIP)

AN ATTRIBUTE-BASED FRAMEWORK FOR PRIVACY PRESERVING IMAGE QUERYING

Shantanu Rane and Wei Sun

Mitsubishi Electric Research Laboratories, Cambridge, MA.

ABSTRACT

We are specifically concerned with scenarios in which multimedia data is stored once on the server and the same data is queried by multiple parties. We propose a framework for privacy preserving querying, in which encryption is performed only once, and the ciphertexts are stored on a database server. Rather than using public-key homomorphic cryptosystems, the parties querying the database first derive an "attribute" from their query signal. They can decrypt the server's ciphertext only if their attribute satisfies a specified mathematical condition. This query-specific decryption capability makes attribute-based cryptography a vital addition to the secure signal processor's toolkit. We give an example of a construction for privacy preserving querying, in which a client can privately retrieve an image from the server if attribute vectors extracted from the server's and client's images are close enough in Euclidean distance.

Index Terms— Attribute-based encryption, Bilinear maps, Privacy preserving querying

1. INTRODUCTION

Alice is a medical researcher who has found an unusual symptom in a medical image. To further her studies, she wants to query a medical database for images that show similar symptoms. However, she does not wish to share her image with Bob, the database administrator. Moreover, Bob is concerned about the privacy of the patients in the database, and therefore keeps his database anonymized and encrypted at all times. Is it possible for Alice to search through Bob's encrypted database and retrieve images similar to her own, while still satisfying these privacy constraints? Such privacy preserving querying scenarios are expected to become increasingly common with the deployment of cloud-based storage for healthcare data, financial records, census data and other kinds of sensitive information.

The problem of determining whether two images are similar without revealing the images falls under the realm of secure multiparty computation. If the function to be computed – in this case a distance function between two images – can be expressed as an algebraic circuit, there exists a generalized protocol to compute it while satisfying all the privacy requirements [1, 2]. In practice, however, such a generalized protocol is extremely complex in terms of computation and communication overhead. Therefore, it is necessary to develop efficient protocols for commonly used functions, such as Euclidean distance, Hamming distance, cross-correlation, etc.

One line of work that has received increased attention in recent years is the application of public-key homomorphic cryptosystems for computing functions in the encrypted domain. Depending on the encrypted-domain computation that these cryptosystems allow, they can be classified into additively, multiplicatively and doubly homomorphic cryptosystems. See [3, 4, 5] for seminal examples of homomorphic cryptosystems. Such cryptosystems have been used to construct privacy preserving protocols for string comparisons [6], nearest-neighbor clustering [7], face recognition [8, 9], biometric authentication [10, 11] and other applications. Most of these protocols require encryption of the data using the public encryption key of the querying party and operate in two stages (1) Secure computation of the distance or correlation between data entities, and (2) Information retrieval based on a distance criterion. While these protocols are efficient for a single retrieval request, they may not scale for a very large number of users simultaneously querying a database. For example, if a second researcher, Charlie, also wants to retrieve similar images from Bob's database, the entire protocol must be replicated using the encryption/decryption key pair of Charlie. For several simultaneous queries, a large amount of ciphertext is produced, only to be discarded later. We consider an alternative scenario wherein Bob can encrypt his data once and for all, after which, anybody can retrieve images from the server using decryption keys *calculated from their query images*. In this paper, our first goal is to incorporate these desiderata into a new framework for privacy preserving querying of multimedia databases. Our second goal is to show how this framework can be realized using attribute based encryption (ABE).

In a conventional cryptosystem, when Alice wants to transmit a message securely to Bob, she must encrypt it either with a symmetric key known to her and Bob, or with Bob's public key. Instead, in an ABE system such as [12, 13], Alice obtains some public encryption parameters from a Key Authority and generates a ciphertext that contains two entities: the encryption of the message m and a so-called attribute vector $\mathbf{x}$. The encryption can only be reversed by a decryption key that satisfies a mathematical condition on the attribute $\mathbf{x}$ of Alice, and the attribute $\mathbf{y}$ (say) of Bob. In order to perform decryption, Bob applies to the Key Authority for a decryption key which is a function of his attribute vector $\mathbf{y}$. For e.g., Bob can decrypt m if and only if $\mathbf{x} \perp \mathbf{y}$, or equivalently $\mathbf{x}^T\mathbf{y} = 0$ [13]. To realize the privacy-preserving querying framework in this paper, we present an ABE construction in which decryption is conditional on the Euclidean distance between attribute vectors.

The remainder of this paper is organized as follows: Section 2 presents the proposed framework for privacy preserving querying. Section 3 reviews bilinear mappings and some mathematical background needed to construct an attribute based cryptosystem that can be exploited to realize the proposed framework. In Section 4, we briefly present one such attribute based cryptosystem in which decryption is possible only if the squared Euclidean distance, i.e., the squared ℓ_2 distance between the attributes of the encryptor and decryptor are below a threshold. Section 5 discusses the challenges involved in this line of research and concludes the paper.

2. ATTRIBUTE-BASED QUERYING FRAMEWORK

We now describe a framework for media querying that does not require encryption to be repeated while servicing a request from a new

client. We describe below the three stages in this framework, viz., generation of encrypted content, storage and retrieval.

2.1. Generation of Encrypted Content

Let $\mathbf{M}^{(i)}, i \in \{1, 2, ..., m\}$ represent the images to be included in the database. For each i, the administrator, Bob generates a secret key L_i for a symmetric cryptosystem of his choice, for e.g., AES [14]. Using the secret key for the symmetric cryptosystem, he generates the ciphertext $S(\mathbf{M}^{(i)}, L_i)$. After this, Bob extracts from each image $\mathbf{M}^{(i)}$, an attribute vector $\mathbf{x}^{(i)}$. The attribute vector can be any efficient representation of the image; candidates for the attribute vector $\mathbf{x}^{(i)}$ are an image digest, a feature vector, a robust hash, a locality-sensitive hash [15], an image fingerprint, or in the degenerate case, the vectorized image matrix $\mathbf{M}^{(i)}$ itself. Then, Bob generates a vector of public encryption parameters $\mathbf{W}$, and computes the ciphertext $C(\mathbf{x}^{(i)}, L_i, \mathbf{W})$. A concrete example of the public encryption parameters and the function $C(\cdot, \cdot, \cdot)$ will be given in Section 4; for now, it should be noted that the ciphertext hides both the attribute vector $\mathbf{x}^{(i)}$ and the symmetric key L_i used to encrypt the image $\mathbf{M}^{(i)}$. The encryption process is depicted in Fig. 1.

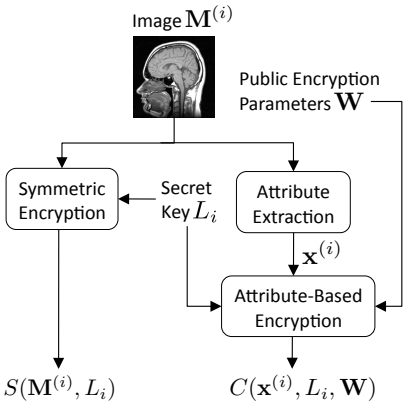

Fig. 1. Bob generates a symmetric encryption of each image, and a ciphertext based on an attribute vector extracted from each image.

2.2. Storage of Encrypted Content

Having generated encrypted content corresponding to the images $\mathbf{M}^{(i)}$, Bob constructs a database as shown in Fig. 2. Observe that, in order to recover any image $\mathbf{M}^{(j)}, j \in \{1, 2, ..., m\}$, a client – or an adversary – needs to obtain the corresponding symmetric key L_j, but this key is hidden in the ciphertext $C(\mathbf{x}^{(j)}, L_j, \mathbf{W})$. For computationally bounded adversary, all usable information about the image, embodied by the image $\mathbf{M}^{(j)}$ and its attribute $\mathbf{x}^{(j)}$ are hidden in the ciphertexts $S(\mathbf{M}^{(j)}, L_j)$ and $C(\mathbf{x}^{(j)}, L_j, \mathbf{W})$ respectively.

2.3. Retrieval of Encrypted Content

The retrieval process is depicted in Fig. 3. Suppose that a client, Alice, has a query image $\mathbf{Q}$ from which she has extracted an attribute vector $\mathbf{y}$. For simplicity, we may consider that the attribute extraction algorithms used by Alice and Bob are identical, though this is not a binding requirement. The essential requirement is that Alice should be able to retrieve an image $\mathbf{M}^{(j)}, j \in \{1, 2, ..., m\}$ from Bob if and only if her own attribute $\mathbf{y}$ satisfies a specific mathematical condition with respect to the attribute $\mathbf{x}^{(j)}$ of image $\mathbf{M}^{(j)}$. Therefore, the ciphertext $C(\mathbf{x}^{(j)}, L_j, \mathbf{W})$ is designed such that it can be

Index	Data	Attribute
1	$S(\mathbf{M}^{(1)}, L_1)$	$C(\mathbf{x}^{(1)}, L_1, \mathbf{W})$
⋮	⋮	⋮
j	$S(\mathbf{M}^{(j)}, L_j)$	$C(\mathbf{x}^{(j)}, L_j, \mathbf{W})$
⋮	⋮	⋮
m	$S(\mathbf{M}^{(m)}, L_m)$	$C(\mathbf{x}^{(m)}, L_m, \mathbf{W})$

Fig. 2. The database server stores all the ciphertexts, two for each image.

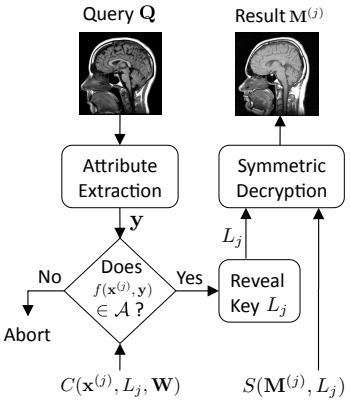

Fig. 3. Alice can decrypt the key L_j only if her attribute $\mathbf{y}$ satisfies a mathematical condition along with the attribute of the j^{th} image. This key enables her to retrieve the j^{th} image.

decrypted by Alice only if a specified function $f(\mathbf{x}^{(j)}, \mathbf{y})$ takes a value in a permissible set $\mathcal{A}$. Decryption reveals the key L_j. To ensure privacy of the database entries, the condition $f(\mathbf{x}^{(j)}, \mathbf{y}) \in \mathcal{A}$ is checked without revealing the attribute $\mathbf{x}^{(j)}$. Note again, that we are only describing a desired mode of operation here; an example of how this can be accomplished by means of attribute based cryptography, is described in Sections 3 and 4. Once L_j is revealed to Alice, she retrieves $S(\mathbf{M}^{(j)}, L_j)$ from Bob by means of Oblivious Transfer (OT) [1], and retrieves the image $\mathbf{M}^{(j)}$. By construction, Alice does not discover any other images in Bob's database, while OT ensures that Bob does not discover the index of the retrieved image.

Note that public key encryption causes significant ciphertext expansion and thus a corresponding increase in storage, computation and communication overhead. That is why we propose to use public key encryption only for the (preferably low-dimensional) attribute vector $\mathbf{x}^{(j)}$, evaluate the querying criteria only based on the attribute vector, and if these criteria are satisfied, transfer the encrypted image $\mathbf{M}^{(i)}$ to Alice. The symmetric encryption of the actual data, $S(\mathbf{M}^{(j)}, L_j)$ is efficient in terms of storage and protocol overhead as it does not cause ciphertext expansion. Such a strategy of searching for matches in a lower dimensional space, while encrypting the media file separately has also been applied in media retrieval schemes based on order-preserving encryption [16].

3. BILINEAR GROUPS & SECURITY ASSUMPTIONS

We review the mathematical properties of bilinear groups of composite order, particularly when the group order N is a product of three primes [13]. Let $N = pqr$, where p, q, r are three distinct prime numbers. Let $\mathbb{G}$ and $\mathbb{G}_T$ be cyclic groups of order N. Then,

the mapping $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ is a non-degenerate bilinear map if it satisfies (1) $e(x^\alpha, y^\beta) = e(x, y)^{\alpha\beta}$ for $x, y \in \mathbb{G}$ and $\alpha, \beta \in \mathbb{Z}$, and (2) If g is a generator of $\mathbb{G}$, then $e(g, g)$ is a generator of $\mathbb{G}_T$.

Now, consider the cyclic groups $\mathbb{G}_p, \mathbb{G}_q$ and $\mathbb{G}_r$ with orders p, q and r respectively and generators g_p, g_q and g_r respectively. Then $\mathbb{G} = \mathbb{G}_p \times \mathbb{G}_q \times \mathbb{G}_r$ and any element $x \in \mathbb{G}$ can be represented as $x = g_p^\alpha g_q^\beta g_r^\gamma$, where $\alpha, \beta, \gamma \in \mathbb{Z}$. The bilinear map $e(\cdot, \cdot)$ has the following four well-known properties:

$$e(g_p^\alpha, g_q^\beta) = 1, \qquad e(g_p^\alpha g_q^{\beta'}, g_q^\beta) = e(g_q^{\beta'}, g_q^\beta)$$
$$e(g_p^\alpha, g_p^{\beta'} g_p^\beta) = e(g_p^\alpha, g_p^{\beta'}) \cdot e(g_p^\alpha, g_p^\beta)$$
$$e(g_p^\alpha g_q^\beta, g_p^{\alpha'} g_q^{\beta'}) = e(g_p^\alpha, g_p^{\alpha'}) \cdot e(g_q^\beta, g_q^{\beta'})$$

Proving these properties involves a straightforward application of the definition of the bilinear mapping given above, and the properties of multiplicative cyclic groups, which allow us to substitute $g_p \equiv g^{qr}$, $g_q \equiv g^{pr}$, and $g_r \equiv g^{pq}$, where g is a generator of $\mathbb{G}$.

The security of attribute-based cryptosystems employing bilinear groups of composite order reduces to solving the following two problems that are regarded as computationally intractable:

1. **Subgroup Decision Problem:** It is computationally hard to distinguish elements of the subgroup $\mathbb{G}_p \times \mathbb{G}_q$ from an element of the group $\mathbb{G}$ defined above. In other words, it is computationally hard to determine whether an element is drawn from a uniform distribution on $\mathbb{G}$, or from a uniform distribution on the subgroup $\mathbb{G}_p \times \mathbb{G}_q$.
2. **Pairing Diffie-Hellman Problem:** Consider a bilinear map $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$. Choose $\bar{g}$ as one element from the set $\{g_p, g_q, g_r\}$. Suppose that $e(\bar{g}, \bar{g})^v$ is given and an integer u is chosen at random. Then, it is computationally hard to distinguish $e(\bar{g}, \bar{g})^{uv} \in \mathbb{G}_T$ from a randomly chosen element of $\mathbb{G}_T$. Another way of stating this is that, given $e(\bar{g}, \bar{g})^v$, it is computationally hard to obtain v.

Both these assumptions are related to the computational intractability of finding non-trivial prime factors of N. For a detailed discussion of proving the security of an ABE system using the above hardness assumptions, please refer to [13].

4. DISTANCE ATTRIBUTE-BASED CRYPTOSYSTEM

We now show an example construction of an attribute-based cryptosystem to realize the querying framework of Section 2. A sketch of the security proof of this construction appears in [17] and a detailed proof is deferred to a later work. Alice wants to retrieve images from a database administered by Bob. As shown in Fig. 1, the images $\mathbf{M}^{(j)}$ are encrypted using a traditional symmetric cryptosystem. The corresponding attribute vectors $\mathbf{x}_j$ are hidden in a ciphertext along with the secret keys L_j using an attribute-based cryptosystem. Fig. 3 shows that, in order to retrieve $\mathbf{M}^{(j)}$, it is sufficient for Alice to discover the key L_j of the symmetric cryptosystem. Below, we present an attribute-based cryptosystem which ensures that Alice will discover L_j only if the squared ℓ_2 distance between the attributes $\mathbf{x}^{(j)}, j \in \{1, 2, ..., m\}$ of Bob and $\mathbf{y}$ of Alice is less than a threshold τ. Both Alice and Bob are assumed to be honest but curious, i.e., they follow the rules of the protocol below, but may exploit all their available information to discover each other's data at any step in the protocol.

Setup: Bob generates large primes p, q, r and two cyclic groups $\mathbb{G}$ and $\mathbb{G}_T$ of order $N = pqr$. As above, there are cyclic groups $\mathbb{G}_p$, $\mathbb{G}_q$ and $\mathbb{G}_r$ with orders p, q and r respectively and generators g_p, g_q and g_r respectively. Let $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ be a non-degenerate bilinear map[1]. Bob randomly chooses $a \in \mathbb{G}_p$ and $c \in \mathbb{G}_r$, and outputs public parameters $\mathbf{W} = (N, g_p, g_r, g_q c, e(g_p, a))$, and retains a private master key (p, q, r, g_q, a, c). Alice and Bob publicly agree on a distance threshold τ.

Encryption: For every $j \in \{1, 2, ..., m\}$, Bob possesses an integer attribute vector $\mathbf{x}^{(j)} = (x_1^{(j)}, x_2^{(j)}, ..., x_n^{(j)})$ to be hidden in the ciphertext. He randomly chooses $\delta, \gamma \in \mathbb{Z}$ and $s, s_i \in \mathbb{G}_r, i = 1, 2, ..., n$ and computes the ciphertext as shown in (1). Since δ and γ are different for each encryption, the ciphertext is semantically secure.

Decryption Key Generation: Now, Alice needs to calculate a decryption key based on an integer attribute vector, $\mathbf{y} = (y_1, y_2, ..., y_n)$ extracted from her query image $\mathbf{Q}$. However, she must hide both $\mathbf{y}$ and the decryption key from Bob. So, she randomly chooses an integer vector $\mathbf{z} = (z_1, z_2, ..., z_n)$ and sends $\mathbf{z}+\mathbf{y}$ and $\sum_{i=1}^n z_i$ to Bob. Bob randomly chooses integers $\alpha, \beta, \alpha_i, i = 1, 2, ..., n$ and $\sigma_t, \rho_t, t = 0, 1, ..., \tau$, and generates a "pre-decryption key" for Alice. This is given by (2) and the following relations:

$$K_i' = g_p^{\alpha_i} g_q^{\beta(y_i + z_i)} \text{ for } i = 1, 2, ..., n$$
$$K_t'^{(3)} = g_p^\alpha g_q^{\beta(\sum_{i=1}^n (y_i+z_i)^2 - t - \sigma_t)} \text{ for } t = 0, 1, ..., \tau$$

Using the pre-decryption key and her knowledge of $\mathbf{y}$ and $\mathbf{z}$, Alice obtains the decryption key, given by (2) and the following relations:

$$K_i = K_i' K_0^{-z_i} \text{ for } i = 1, 2, ..., n \tag{4}$$
$$K_t^{(3)} = K_t'^{(3)} K_0^{-\sum_{i=1}^n (z_i^2 + 2 z_i y_i)} \text{ for } t = 0, 1, ..., \tau \tag{5}$$

Decryption: Given the attribute $\mathbf{y}$, Alice evaluates the expression (3) repeatedly for $t = 0, 1, ...\tau$ for each $j \in \{1, 2, ..., m\}$. Space constraints preclude us from writing the intermediate steps, but by repeatedly applying the four properties of the bilinear map provided in Section 3, readers can verify that the right hand side of (3) simplifies to:

$$D_{\ell_2} = L_j \cdot e(g_q, g_q)^{\gamma\beta(\|\mathbf{x}^{(j)} - \mathbf{y}\|_2^2 - t))} \tag{6}$$

Thus, the secret key L_j is unmasked by Alice if and only if $\| \mathbf{x}^{(j)} - \mathbf{y} \|_2^2 = t$. Otherwise, the result D_{ℓ_2} is just some element of $\mathbb{G}_T$. Since Alice does not discover $\mathbf{x}^{(j)}$ at any stage, an important question here is: How does Alice know that she has actually decrypted L_j? One way to resolve this is as follows: Before computing $C(\mathbf{x}^{(j)}, L_j, \mathbf{W})$, Bob left-shifts the digits of L_j, and appends a *publicly known* pattern of digits. Since the ciphertext field is so large, it is extremely unlikely that evaluating (6) will return a value containing the embedded digit pattern for the case $\| \mathbf{x}^{(j)} - \mathbf{y} \|_2^2 \neq t$. If Alice discovers the public embedded pattern in the digits of D_{ℓ_2}, she declares that decryption was successful for some $t \leq \tau$, removes the embedded pattern, and recovers the actual secret key L_j. If she does not discover the embedded pattern in D_{ℓ_2} for any $t \in [0, \tau]$, then decryption is deemed unsuccessful based on the ℓ_2 distance condition on the attributes.

[1]There exist algorithms based on elliptic curves to generate groups of composite order and bilinear mappings using these groups. Examples include the Weil pairing and Tate pairing [18].

$$C(\mathbf{x}^{(j)}, L_j, \mathbf{W}) = (A_0, \{A_i\}_{i=1}^n, \bar{A}, B, H_j) \equiv \left((g_q c)^{\gamma}, \{(g_q c)^{-\gamma x_i^{(j)}} s_i\}_{i=1}^n, (g_q c)^{-\gamma \sum_{i=1}^n (x_i^{(j)})^2} s, g_p^{\delta}, L_j \cdot e(g_p, a)^{\delta} \right) \quad (1)$$

$$(K_0, \{K_t^{(1)}, K_t^{(2)}\}_{t=0}^{\tau}) \equiv \left(g_p^{\alpha} g_q^{\beta}, \{a^{-1} g_p^{\alpha - \rho_t - 2\sum_{i=1}^n \alpha_i + 2\alpha \sum_{i=1}^n z_i}, g_p^{\rho_t} g_q^{\beta \sigma_t}\}_{t=0}^{\tau} \right) \quad (2)$$

$$D_{\ell_2} = H_j \cdot e(B, K_t^{(1)}) \cdot e(A_0, K_t^{(3)}) \cdot e(A_0 B, K_t^{(2)}) \cdot \left(e(\bar{A}B, K_0)\right)^{-1} \cdot \left(\prod_{i=1}^n e(A_i B, K_i) \right)^2 \quad (3)$$

5. DISCUSSION

Compared with previous approaches, the proposed attribute-based framework has several salient features: (1) Unlike most protocols based on homomorphic cryptosystems, encryption is performed only once using the *server's* public key, not with the public key of each querying client. Thus, the protocol scales for a large number of simultaneous or sequential database queries. (2) No additional parties are introduced, so the issue of collusion does not arise. (3) In ABE schemes in the literature, a powerful Key Authority generates public encryption parameters and knows the attribute vectors held by all decryptors. In contrast, in the proposed scheme, the public encryption parameters are generated by the server, Bob, and additive secret sharing is used to hide Alice's attribute vector from Bob. (4) Once the encrypted database is created by Bob, data retrieval proceeds in just two major steps (A) Alice downloads the ciphertexts $C(\mathbf{x}^{(i)}, L_i, \mathbf{W})$, and successfully recovers secret keys $L_j, j \in \mathcal{S}_A \subseteq \{1, 2, ..., m\}$ for which the querying condition is satisfied, and (B) Alice obtains ciphertexts $S(\mathbf{M}^{(j)}, L_j)$ via an oblivious transfer protocol with Bob, and then decrypts the images $\mathbf{M}^{(j)}$.

The construction of Section 4 has one significant limitation: Alice must carry out $O(\tau)$ decryptions per database entry, one decryption for each $t = 0, 1, ..., \tau$. This is a consequence of the ciphertext construction. Specifically, the secret key L_j in (1) is unmasked if the ℓ_2 distance equals t for some $t \in \{0, 1, ..., \tau\}$. This places a limitation on the value of τ or more generally on the range $[\tau_{\min}, \tau_{\max}]$ of distances that can be tested. Alice's computational overhead would be smaller if the distance threshold condition, $\| \mathbf{x}^{(j)} - \mathbf{y} \|_2^2 \lesseqgtr \tau$, could be tested over all $t \in [0, \tau]$ with only one decryption. Incorporating a more efficient way to test the threshold condition on the attributes is an interesting avenue for future research.

6. REFERENCES

[1] R. Cramer, "Introduction to secure computation," *Lectures on Data Security - Modern Cryptology in Theory and Practice*, vol. 1561, pp. 16–62, March 1999.

[2] A. C-C. Yao, "How to Generate and Exchange Secrets," in *Proc. 27th Annual Symposium on Foundations of Computer Science (SFCS)*, Washington, DC, USA, 1986, pp. 162–167, IEEE Computer Society.

[3] P. Paillier, "Public-Key Cryptosystems Based on Composite Degree Residuosity Classes," in *Advances in Cryptology, EUROCRYPT*. 1999, vol. 1592, pp. 233–238, Springer-Verlag, LNCS.

[4] T. El Gamal, "A Public Key Cryptosystem and A Signature Scheme Based on Discrete Logarithms," *IEEE Transactions on Information Theory*, vol. 4, pp. 469–472, Jul. 1985.

[5] C. Gentry, "Computing arbitrary functions of encrypted data," *Communications of the ACM*, vol. 53, no. 3, pp. 97–105, March 2010.

[6] M. Atallah and J. Li, "Secure outsourcing of sequence comparisons," *International Journal of Information Security*, vol. 4, no. 4, pp. 277–287, Oct. 2005.

[7] M. Shaneck, Y. Kim, and V. Kumar, "Privacy preserving nearest neighbor search," in *Sixth IEEE Intl. Conf. Data Mining - Workshops*, Washington, DC, USA, 2006, pp. 541–545.

[8] Z. Erkin, M. Franz, J. Guajardo, S. Katzenbeisser, R.L. Lagendijk, and T. Toft, "Privacy-preserving face recognition," in *International Symposium on Privacy Enhancing Technologies (PET)*, Seattle, WA, August 2009, pp. 235–253.

[9] A. Sadeghi, T. Schneider, and I. Wehrenberg, "Efficient privacy-preserving face recognition," in *International Conference on Information Security and Cryptology (ICISC '09)*, Seoul, Korea, December 2009.

[10] M. Blanton and P. Gasti, "Secure and efficient protocols for iris and fingerprint identification," in *European Symposium on Research in Computer Security (ESORICS)*, Leuven, Belgium, September 2011.

[11] M.Barni, T.Bianchi, D.Catalano, M.Di Raimondo, R.Donida Labati, P.Failla, D.Fiore, R.Lazzeretti, V.Piuri, F.Scotti, and A.Piva, "Privacy-preserving fingercode authentication," in *ACM Workshop on Multimedia and Security (MMSEC2010)*, Rome, Italy, Sept 2010.

[12] A. Sahai and B. Waters, "Fuzzy Identity Based Encryption," in *EUROCRYPT*, Aarhus, Denmark, May 2005, pp. 457–473.

[13] J. Katz, A. Sahai, and B. Waters, "Predicate Encryption Supporting Disjunction, Polynomial Equations, & Inner Products," in *EUROCRYPT*, Istanbul, Turkey, April 2008, pp. 146–162.

[14] C. Paar and J. Pelzl, *Understanding Cryptography: A Textbook for Students and Practitioners*, Springer-Verlag New York Inc, 2010.

[15] A. Andoni and P. Indyk, "Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions," *Communications of the ACM*, vol. 51, no. 1, pp. 117–122, Jan. 2008.

[16] W. Lu, A. Swaminathan, A. L. Varna, and M. Wu, "Enabling search over encrypted multimedia databases," in *Media Forensics and Security Conference (SPIE IS&T 2009)*, San Jose, CA, January 2009.

[17] W. Sun and S. Rane, "A distance-sensitive attribute based cryptosystem for privacy-preserving querying," in *International Conference Multimedia and Expo (ICME 2012),* To Appear., Melbourne, Australia, July 2012.

[18] J. Silverman, *The Arithmetic of Elliptic Curves*, Graduate Texts in Mathematics, Springer-Verlag, 1986.