

Efficiently Computing Private Recommendations

Zekeriya Erkin¹

Michael Beye¹

Thijs Veugen^{1,2}

Reginald L. Lagendijk¹

¹Information Security and Privacy Lab, Faculty of EEMCS
Delft University of Technology, 2628 CD, Delft, The Netherlands

²TNO Information and Communication Technology
P.O. Box 5050, 2600 GB Delft, The Netherlands

{z.erkin, m.r.t.beye, p.j.m.veugen, r.l.lagendijk}@tudelft.nl

Abstract

Online recommender systems enable personalized service to users. The underlying collaborative filtering techniques operate on privacy sensitive user data, which could be misused if it is leaked or by the service provider him self. To protect user's privacy, we propose to encrypt the data and generate recommendations by processing them under encryption. Thus, the service provider observes neither user preferences nor recommendations. The proposed method uses homomorphic encryption and secure multiparty computation (MPC) techniques, which introduce a significant overhead in computational complexity. The second contribution of this paper lies in minimizing this overhead by packing data. The improvements are illustrated by a complexity analysis.

Keywords: Recommender systems, user privacy, MPC, homomorphic encryption, data packing.

1 Introduction

In the last decade, we have experienced phenomenal progress in information and communication technologies. Cheaper, more powerful, less power consuming devices and high bandwidth communication lines enabled us to create a new virtual world in which people mimic activities from their daily lives without the limitations imposed by the physical world. As a result, online applications have become very popular for millions of people [1].

Personalization is a common approach to further improve online services and attract more users. Instead of making general suggestions for the users of the system, the system can suggest personalized services targeting only a particular user based on his preferences [2]. Since the personalization of the services offers high profits to the service providers and poses interesting research challenges, research for generating recommendations, also known as collaborative filtering, attracts attention both from academia and industry.

The techniques to generate recommendations for users strongly rely on information gathered from the user. This information can be provided by the user himself as in profiles or the service provider can observe users' actions, such as click logs. On one hand, more user information helps the system to improve the accuracy of the recommendations. On the other hand, the information on the users creates a severe privacy risk since there is no solid guarantee for the service provider not to misuse the users' data. It is often seen that whenever a user enters the system, the service provider claims the ownership of the information provided by the user and authorizes itself to distribute the data to third parties for its own benefits [14].

As an example, consider pay-TV boxes. A small box purchased by the user provides high quality broadcasting with several interesting features like recording programs. Companies in this field also suggest programs and movies that they think their customers may like. In order to make useful recommendations to their customers, the small box observes the user behavior:

it records the programs watched, the duration spent in front of the TV and so on. The information gathered by the box is then sent to a server and processed to deduce meaningful information about users. It is obvious that this system can be used for harming the user's privacy.

In this paper, we propose a cryptographic solution to preserve the privacy of users in a recommender system. In particular, the privacy-sensitive data of the users are kept encrypted and the service provider generates recommendations by processing encrypted data. The cryptographic protocol developed for this purpose is based on homomorphic encryption [3] and secure multiparty computation (MPC) techniques [15]. While the homomorphic property is used for realizing linear operations, protocols based on MPC techniques are developed for non-linear operations (e.g. finding the most similar users). This added privacy comes at a cost: increased computational complexity. Aside from MPC techniques, multiplications in the encrypted domain are expensive, because they involve exponentiations of relatively large numbers. Another issue is data expansion, because (small) plain texts are transformed into cipher texts of a large, fixed length. A main contribution of this paper is to provide a significant reduction of the overhead of working in the encrypted domain, by using data packing. In this way, many smaller plain texts are combined into one large cipher text, thus performing several multiplications by one single exponentiation. As an added benefit, the message expansion is reduced.

After looking at related work in Section 2, we will outline the general workings of a recommender system in Section 3. In Sections 4, the cryptographic primitives used throughout the paper will be introduced. In Sections 5 and 6, we present our privacy-preserving protocol for generating recommendations. A security analysis of the proposed scheme is provided in Section 7. The complexity analysis is given in Section 8. Finally, conclusions will be drawn in Section 9.

2 Related Work

In [4], Canny proposes a system where the private user data is encrypted and recommendations are generated by applying an iterative procedure based on the conjugate gradient algorithm. The algorithm computes a characterization matrix of the users in a subspace and generates recommendations by calcu-

lating reprojections in the encrypted domain. Since the algorithm is iterative, it takes many rounds for convergence and in each round users need to participate in an expensive decryption procedure which is based on a threshold scheme where a significant portion of the users are assumed to be online and honest. The output of each iteration, which is the characterization matrix, is available in clear. In [5], Canny proposes a method to protect the privacy of users based on a probabilistic factor analysis model by using a similar approach as in [4].

While Canny works with encrypted user data, Polat and Du suggest to protect the privacy of users by using randomization techniques [12, 13]. In their paper, they blind the user data with a known random distribution assuming that in aggregated data this randomization cancels out and the result is a good estimation of the intended outcome. The success of this method highly depends on the number of users participating in the computation since for the system to work, the number of users need to be vast. This creates a trade-off between accuracy/correctness of the recommendations and the number of users in the system. Moreover, the outcome of the algorithm is also available to the server who may constitute a privacy threat to the users. Finally, the randomization techniques are believed to be highly insecure [16].

3 Generating Recommendations

A centralized system for generating recommendations is a common approach in e-commerce applications. To generate recommendations for user A , the server follows a two-step procedure. In the first step, the server searches for users similar to user A . Each user in the system is represented by a preference vector which is usually composed of ratings for each item within a certain range. Finding similar users is based on computing similarity measures between users' preference vectors. Pearson correlation (Eq. 1) is a common similarity measure for two users with preference vectors $V_A = (v_{(A,0)}, \dots, v_{(A,M-1)})$ and $V_B = (v_{(B,0)}, \dots, v_{(B,M-1)})$, respectively, where M is the number of items and \bar{v} represents the average value of the vector v .

$$\text{sim}_{A,B} = \frac{\sum_{i=0}^{M-1} (v_{(A,i)} - \bar{v}_A) \cdot (v_{(B,i)} - \bar{v}_B)}{\sqrt{\sum_{i=0}^{M-1} (v_{(A,i)} - \bar{v}_A)^2 \cdot \sum_{i=0}^{M-1} (v_{(B,i)} - \bar{v}_B)^2}}. \quad (1)$$

Once the similarity measure for each user is computed, the server proceeds with the second step. In this step, the server chooses those L users with similarity values above a threshold δ and averages their ratings. These average ratings are then presented as *recommendations* to user A .

In e-commerce applications the number of items offered to users is usually in the order of hundreds or thousands. Apart from many smart ways of determining the likes and dislikes of users for the items, we assume the users are asked to rate the items explicitly with integer values in the range of $[0, K]$. This rating matrix is usually highly sparse, meaning that most of the items are not rated. Finding similar users in a sparse dataset can easily lead the server to generate inaccurate recommendations. To cope with this problem, one approach is to introduce a small set of items that is rated by most users. Such a base set can be explicitly given to the users or implicitly chosen by the server from the most commonly rated items. Given such a small set of items that is rated by most users, the server can compute similarities between users more confidently, resulting in more accurate recommendations. Therefore, we assume that the user preference vector V is split into two parts: the first part consists of R elements that are rated by most of the users and the second part contains $M - R$ sparsely rated items that the user would like to get recommendations on [2].

4 Cryptographic Primitives and Security Model

We use encryption to protect user data against the service provider and other users. A special class of cryptosystems, namely homomorphic cryptosystems, allows us to process data in the encrypted form. We choose the Paillier cryptosystem [11] as it is *additively homomorphic* meaning that the product of two encrypted values $[a]$ and $[b]$, (where $[\cdot]$ denotes the encryption function), corresponds to a new encrypted message whose decryption yields the sum of a and b

as $[a] \cdot [b] = [a + b]$. As a consequence of the additive homomorphism, any cipher text $[m]$ raised to the power of a public value c corresponds to the multiplication of m and c in the encrypted domain: $[m]^c = [m \cdot c]$. In addition to the homomorphism property, the Paillier cryptosystem is semantically secure implying that each encryption has a random element that results in different cipher texts for the same plain text.

As a part of a cryptographic protocol introduced in Section 6, we use another additively homomorphic and semantically secure encryption scheme, DGK [7, 6]. The DGK cryptosystem is used to replace the Paillier cryptosystem in a subprotocol, for reasons of efficiency. For the same level of security, DGK has a much smaller message space compared to the Paillier cryptosystem and thus, encryption and decryption operations are more efficient than under Paillier.

We use the semi-honest security model, which assumes that all players follow the protocol steps but are curious and thus keep all messages from previous and current steps to extract more information than they are allowed to have. Our protocol can be adapted to the active attacker model by using the ideas in [10] with additional overhead.

5 Privacy Preserving Recommender System

In this section we propose a protocol based on additively homomorphic encryption schemes and MPC techniques. In particular the service provider, i.e. the server, receives the encrypted preference vector of user A and sends it to the other users in the system who can then compute the similarity value on their own by using the homomorphism property of the encryption scheme. Once the users compute the similarity values, they are sent to the server. After that, the server and user A run a protocol to determine which similarity values are above a threshold δ . The server - being unaware of the number of users with a similarity value above a threshold, and their identities - accumulates the ratings of all users in the encrypted domain. Then, the encrypted sum is sent to user A along with the encrypted number of similarities above the threshold, L . User A decrypts the sum and L and computes the average values, obtaining the recommendations. Each step of the proposed protocol

is detailed in the following sections.

5.1 Key Generation and Preprocessing

Any user in the system who wants to get recommendations generates personal public key pairs for the Paillier and the DGK cryptosystems. We assume that the public keys of the users are available publicly.

Since the Pearson correlation given in (1) for user A and B can be also written as:

$$\begin{aligned} \text{sim}_{A,B} &= \sum_{i=0}^{R-1} C_{A,i} \cdot C_{B,i}, \text{ where} \\ C_{X,i} &= \frac{(v_{(X,i)} - \bar{v}_X)}{\sqrt{\sum_{j=0}^{R-1} (v_{(X,j)} - \bar{v}_X)^2}}. \end{aligned} \quad (2)$$

The terms $C_{A,i}$ and $C_{B,i}$ can be easily computed by users A and B , respectively. Each user computes a vector from which the mean is subtracted and normalized. Since the elements of the vector are real numbers and cryptosystems are only defined on integer values, they are all scaled by a parameter f and rounded to the nearest integer resulting in a new vector $V'_i = (v'_{(i,0)}, \dots, v'_{(i,R-1)})$ whose elements are now k -bit positive integers. Note that the threshold value δ should also be adjusted accordingly.

5.2 Computing Similarity Measures

The similarity value between user A and any other user B is computed over the rating vectors of size R . The elements of the user vector $V'_A = (v'_{(A,0)}, \dots, v'_{(A,R-1)})$ are encrypted individually by using the public key of the user A . Then, the encrypted vector $[V'_A]_{pk_A}$ is sent to the server. The server then sends the encrypted vector to the other users in the system. Any user B who receives the encrypted vector $[V'_A]_{pk_A}$ can compute the encrypted similarity as follows:

$$\begin{aligned} [\text{sim}_{A,B}] &= \left[\sum_{i=0}^{R-1} v'_{(A,i)} \cdot v'_{(B,i)} \right] \\ &= \left[v'_{(A,0)} \cdot v'_{(B,0)} + \dots + v'_{(A,R-1)} \cdot v'_{(B,R-1)} \right] \\ &= \left[v'_{(A,0)} \right]^{v'_{(B,0)}} \cdot \dots \cdot \left[v'_{(A,R-1)} \right]^{v'_{(B,R-1)}} \\ &= \prod_{i=0}^{R-1} \left[v'_{(A,i)} \right]^{v'_{(B,i)}}. \end{aligned} \quad (3)$$

Note that we omit the encryption key pk_A above and in the rest of the paper for the sake of readability. The computed similarity value is then sent back to the server in encrypted form.

5.3 Finding the Most Similar Users

Upon receiving similarity values from users, the server initiates a cryptographic protocol with user A to determine the most similar users whose similarity values are above a public threshold δ . The protocol receives N encrypted similarity values and outputs an encrypted vector $[\Gamma_A] = ([\gamma_{(A,0)}], [\gamma_{(A,1)}], \dots, [\gamma_{(A,N-1)}])$. The elements of this vector $\gamma_{(A,i)}$ are either an encryption of 1, if the similarity value between user A and user i is above the threshold δ , or an encryption of 0, otherwise. The details of this protocol can be found in Section 6.

5.4 Generating Recommendations

After obtaining the vector $[\Gamma_A]$, the server can generate the recommendations for user A . For this purpose, the server sends $[\gamma_{(A,i)}]$ to the i^{th} user in the system. User i , referred to as user B , can raise $[\gamma_{(A,B)}]$ to the power of each rating he has left in his ratings vector to obtain another encrypted vector $[\Phi_{(A,B)}] = ([\phi_{(A,R)}], [\phi_{(A,R+1)}], \dots, [\phi_{(A,M-1)}])$ where $\phi_{(A,j)} = [\gamma_{(A,B)} \cdot v'_{(B,j)}] = [\gamma_{(A,B)}]^{v'_{(B,j)}}$ for $(R \leq j < M)$. Notice that user B does not know the content of $\gamma_{(A,B)}$. The resulting vector $[\Phi_{(A,B)}]$ is either the encrypted rating vector of user B or a vector of encrypted 0's. Vector $[\Phi_{(A,B)}]$ is then sent to the server to be accumulated with other users' vectors.

The above procedure can be improved in order to minimize the computational and communication cost by using data packing. Instead of raising $[\gamma_{(A,B)}]$ to the power of each rating, the ratings can be represented in a compact form and then used as an exponent:

$$v'_{(B,R)} | v'_{(B,R+1)} | \dots | v'_{(B,M-1)}, \quad (4)$$

where $|$ represents the concatenation operation. Assuming that each $v'_{(B,j)}$ is k -bits and N of such vectors are to be accumulated by the server, where N is the number of users participating in the protocol, each compartment should have a bit size of $k + \log(N)$. Thus, packing is achieved by the following formula:

$$v''_B = \sum_{j=0}^{M-R} 2^{j(k+\log(N))} \cdot v'_{(B,j+R)}. \quad (5)$$

By packing values, the communication cost reduces significantly as we obtain a packed value rather than a vector of encrypted vectors. Packing also reduces the number of exponentiations which is a costly operation in the encrypted domain, introducing a gain in computation. However, depending on the message space of the encryption scheme, n , and the number of ratings, $M - R$, it may not be possible to pack all values in one encryption. The number of values that can fit into one encryption is $T = n/(k + \log(N))$. Therefore, we may need $S = \lceil (M - R)/T \rceil$ encryptions.

Once user B packs his ratings to obtain v_B'' , he can compute $[\Phi_{(A,B)}]$ as follows:

$$\begin{aligned} [\Phi_{(A,B)}] &= [\gamma_{(A,B)}]^{v_B''} \\ &= \begin{cases} [v_B''] & \text{if } \gamma_{(A,B)} = 1 \\ [0] & \text{if } \gamma_{(A,B)} = 0 \end{cases}, \end{aligned} \quad (6)$$

and sends $[\Phi_{(A,B)}]$ to the server. Upon receiving $[\Phi_{(A,i)}]$ values from all users, the server accumulates them:

$$[\Phi_A] = \prod_{i=0}^N [\Phi_{(A,i)}] = \left[\sum_{i=0}^N \Phi_{(A,i)} \right]. \quad (7)$$

Notice that the result will be equal to the sum of ratings of the users who have similarity values above threshold δ . The server also accumulates the $[\gamma_{(A,i)}]$ values to obtain the number of users above the threshold also encrypted:

$$[L] = \prod_{i=0}^N [\gamma_{(A,i)}] = \left[\sum_{i=0}^N \gamma_{(A,i)} \right]. \quad (8)$$

These two values, $[\Phi_A]$ and $[L]$ are then sent to user A . After decrypting, user A decomposes Φ_A and divides each extracted value by L , obtaining the average ratings of L users. This concludes our protocol.

An important observation at this point is the value of L . If $L = 0$, the user can notify the server to repeat the second step of the protocol with a new threshold. If $L = 1$, the user obtains exactly the same ratings vector of some user but he does not have the identity of that particular user.

6 Cryptographic Protocol for Finding Similar Users

Finding similar users is based on comparing the similarity value between user A and B , $\text{sim}_{A,B}$, to a public threshold δ . As the similarity value is privacy sensitive and should be kept secret both from the server and the user, we compare it in the encrypted domain. For this purpose, we use a comparison protocol that has been introduced in [8]. The cryptographic protocol in [8] takes two encrypted values, $[a]$ and $[b]$, and outputs the result λ again in the encrypted form: if $a > b$ $[\lambda = 1]$, and $[\lambda = 0]$ otherwise. For the completeness of the paper, we give a brief description of the protocol. More explanation and implementation details on the comparison protocol can be found in [8].

Given the similarity value $\text{sim}_{A,B}$ and public threshold δ , both of which are ℓ bits, the most significant bit of the value $z = 2^\ell + \text{sim}_{A,B} - \delta$ is the outcome of the comparison. However, we need to obtain the most significant bit of z in the encrypted domain. While the encrypted value $[z]$ can be computed by the server, the most significant bit of $[z]$ requires running a protocol between the server and user A who has the decryption key. Note that the similarity value cannot be trusted to the user as it leaks information about other users in the system. Therefore, the server adds a random value r to z : $[c] = [z + r]$ and sends it to user A who then decrypts it. Notice that the most significant bit now can be computed as:

$$\begin{aligned} [\gamma_{(A,i)}] &= \\ [2^{-\ell}(z - (c \bmod 2^\ell - r \bmod 2^\ell) + \alpha \cdot 2^\ell)] &, \end{aligned} \quad (9)$$

where the last term is necessary depending on the relation between c and r . The variable α is a single bit representing whether $c > r$ or not. At this point, we convert the problem of comparing $[\text{sim}_{A,i}]$ and δ to the problem of comparing c and r which are owned by the user and the server respectively.

Comparing c and r requires another cryptographic protocol in which the server and user A evaluate the following formula for each of ℓ bits:

$$[e_i] = \left[1 - c_i + r_i + 3 \sum_{j=i+1}^{\ell-1} c_j \oplus r_j \right], \quad (10)$$

where c_i and r_i are the i^{th} bits of c and r , respectively. The value of e_i can be 0 if and only if $c > r$, when $c_i = 0$, $r_i = 1$ and the upper part of c and r are the same. After these computations, the server sends the randomized and shuffled $[e_i]$ values to the user A . User A decrypts them and checks whether there is a zero among the values e_i . Existence of a 0 value indicates that $r > c$. However, this leaks information about the comparison of $\text{sim}_{A,B}$ and δ thus, the server randomizes the direction of the comparison by replacing $1 - c_i + r_i$ in Eq. 10 with $-1 - c_i + r_i$ at random. User A then returns $[\alpha]$ which is either [1] or [0] depending on the existence of a 0 among the e_i values. The server can correct the direction of the comparison and obtain the $[\gamma_{(A,i)}]$ by replacing α in Eq. 9.

By using this comparison protocol, each similarity value is compared to threshold δ in parallel. The outcomes of the comparisons, $[\Gamma_A] = ([\gamma_{(A,0)}], [\gamma_{(A,1)}], \dots, [\gamma_{(A,N-1)}])$, are then used in the subsequent steps.

7 Security Analysis

Our protocol for generating recommendations can be considered as a secure multi-party computation in the semi-honest model. The parties that participate in the computation are the N users and the server. The private input of user i , $0 \leq i < N$, is his (or her) normalized preference vector $V'_i = (v'_{(i,0)}, \dots, v'_{(i,M-1)})$, together with the decryption key K_i . The output for user A , A being the user that is requesting a recommendation, equals the accumulated recommendation vector Φ_A of length $M - R$ and the number L of similar users, i.e. the number of users for which the similarity value with A exceeds a threshold δ .

In order to show that our protocol privately computes the recommendation vector for user A in the semi-honest model, we have to show that whatever can be computed by any party in the protocol from its view of a protocol execution, can be computed from its input and output (see Definition 7.2.1 from Goldreich [9]). We have basically three different types of parties, namely the server, user A , and any other user $i \neq A$.

The view of the server of a protocol execution consists of all public parameters, its randomly generated variables, and a number of received messages

encrypted with the public key of A . The public parameters are the public keys of all users, the threshold δ , the number of (most rated) items M and R , the size ℓ of the similarity values, the size n of the encryption scheme, the size k of one normalized rating, and the number S of encryptions needed to present the accumulated recommendation vector. The randomly generated variables of the server are the $N - 1$ variables (r) that are used in the comparison protocol to determine whether the similarity value with user i , ($i \neq A$), exceeds the threshold δ or not. The comparison protocol, which we use as a subprotocol, has been proven secure in [7, 6]. Our encryption system Paillier, with a message space size of n and the decryption key K_A , is semantically secure [11]. Therefore, it easily follows that anything that the server could compute from its view, can be computed from the public parameters alone.

The view of user A consists of all public parameters, its private input and output, and all messages received by the server. All messages, different from the output of A , that are received by the server, are related to the $N - 1$ executions of the comparison protocol to compare the similarity value $\text{sim}_{A,i}$ with public threshold δ for each $i \neq A$. Since the comparison protocol is known to be secure, it easily follows that anything user A could compute from its view, can be computed from the public parameters and its private input and output alone.

The view of user i , ($i \neq A$), consists of all public parameters, its private input, and all messages received by the server. The only message that user i received from the server is the normalized preference vector $[V'_A]$, encrypted by the public key of user A . Since Paillier is semantically secure, it easily follows that anything user i could compute from its view, can be computed from the public parameters and its private input alone.

8 Complexity Analysis

The performance of our protocol is mainly determined by the interaction among the server, and user A , who asks for recommendation, and other users in the system. In our construction, the server participates in the computation and relays messages among users. User A , on the other hand, only participates in the protocol in two stages: 1) when he asks for a recommendation and uploads his encrypted data and 2) when he receives the encrypted recommendation. Other users help the

Table 1: Computational complexity.

	Server		User A		User B	
	Paillier	DGK	Paillier	DGK	Paillier	DGK
Encryption	$\mathcal{O}(N)$	$\mathcal{O}(N\ell)$	$\mathcal{O}(R)$	$\mathcal{O}(\ell)$	-	-
Decryption	-	-	$\mathcal{O}(1)$	$\mathcal{O}(\ell)$	-	-
Multiplication	$\mathcal{O}(NS)$	$\mathcal{O}(N\ell^2)$	-	-	$\mathcal{O}(R)$	-
Exponentiation	-	$\mathcal{O}(N\ell)$	-	-	$\mathcal{O}(R + S)$	-

server with the recommendation generation.

Recall that R is the number of heavily rated items, N is the number of users, ℓ is the length of the similarity values and threshold δ , S is the number of encryptions required (with packing), and T is the number of values that fit into one encryption.

8.1 Round Complexity

Our protocol consists of 5 rounds. The data transfer from users to the server in the initialization stage is 0.5 round. To determine the similar users and generate the recommendation, the server needs 4 rounds of interaction. Notice that to obtain $[\Gamma_A]$ in the comparison protocol, all encrypted values are compared to a public value δ , and all comparisons can be done in parallel. In the last stage, the server sends the recommendation to user A which requires another 0.5 round. This gives $\mathcal{O}(1)$ rounds.

8.2 Communication Complexity

The amount of data transferred during the protocol is primarily influenced by the size of the encrypted data. For user A , the amount of encrypted data to be transferred is $\mathcal{O}(R + N\ell)$. The server, on the other hand, has to receive and send $\mathcal{O}(N(R + S + \ell))$ encrypted data which is heavily influenced by the data transmission during the comparison of N similarity values. Other users in the system need to receive and send data $\mathcal{O}(R + S)$. As mentioned in Section 5.4, by packing multiple values into one single encryption, we can reduce the data expansion by a factor of T , significantly reducing our communication overhead.

8.3 Computational Complexity

The computational complexity depends strongly on the cost of operations in the encrypted domain, which can be categorized into four classes: encryptions, decryptions, multiplications and exponentiations. In Table 1, we provide the number of each operation in the Paillier and the DGK cryptosystems. One exception is for the decryption operation in DGK, which is actually a *zero-check* which is a fast and less expensive operation compared to original decryption. Again, data packing allows us to operate on T rating values with one single exponentiation. This reduces the number of exponentiations by a factor of T .

8.4 Optimizations

In order to improve the performance of the system, we may consider a few optimizations. Firstly, once the similarity value between two users is computed, it can be stored by the server for future use. As long as the rating vectors of size R that are used for similarity computation do not change, the similarity value will remain the same. This eliminates most of the expensive encryption, exponentiation and multiplication operations in the encrypted domain. Secondly, in our analysis we assume that the similarity values are computed for all users in the system. In applications with millions of users, this approach can be reconsidered in several ways. A smaller set of users can be selected at random, or a group of users who are trusted by user A , also known as social trust network, can be chosen for the similarity computation. Thirdly, encryption can be optimized for run time efficiency. The random values required for Paillier and DGK cryptosystems can be generated in advance or in the idle time of the processors, resulting a substantial gain in efficiency as suggested in [8].

9 Conclusion

In this paper we proposed a cryptographic approach for generating recommendations to the users within online applications. The proposed method is constructed by homomorphic encryption schemes and MPC techniques. This makes our proposal provably secure and not reliant on the number of users in the system, as opposed to randomization techniques [12, 13].

The inevitable overhead introduced by working in the encrypted domain is reduced significantly by packing data (as well as using the DGK cryptosystem), as shown in our complexity analysis. Unfortunately, we cannot compare our result with previously proposed systems due to space restrictions. However, we conclude that our proposal is based on a realistic scenario and the required technology is not overly demanding compared to cryptographic tools like thresholding schemes, as used in other approaches as in [4].

References

- [1] Internet usage statistics. <http://www.internetworldstats.com/stats.htm>, 2009.
- [2] G. Adomavicius and A. Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Trans. on Knowl. and Data Eng.*, 17(6):734–749, 2005.
- [3] N. Ahituv, Y. Lapid, and S. Neumann. Processing encrypted data. *Commun. ACM*, 30(9):777–780, 1987.
- [4] J. F. Canny. Collaborative filtering with privacy. In *IEEE Symposium on Security and Privacy*, pages 45–57, 2002.
- [5] J. F. Canny. Collaborative filtering with privacy via factor analysis. In *SIGIR*, pages 238–245, New York, NY, USA, 2002. ACM Press.
- [6] I. Damgård, M. Geisler, and M. Krøigaard. Efficient and Secure Comparison for On-Line Auctions. In J. Pieprzyk, H. Ghodosi, and E. Dawson, editors, *Australasian Conference on Information Security and Privacy*, volume 4586 of *LNCS*, pages 416–430. Springer, July 2–4, 2007.
- [7] I. Damgård and M. Jurik. A Generalization, a Simplification and some Applications of Paillier’s Probabilistic Public-Key System. Technical report, Department of Computer Science, University of Aarhus, 2000.
- [8] Z. Erkin, M. Franz, J. Guajardo, S. Katzenbeisser, R. L. Lagendijk, and T. Toft. Privacy-preserving face recognition. In *Proceedings of the Privacy Enhancing Technologies Symposium*, pages 235–253, Seattle, USA, 2009.
- [9] O. Goldreich. *Foundations of Cryptography. Basic Applications*, volume 2. Cambridge University Press, first edition, May 2004.
- [10] O. Goldreich, S. Micali, and A. Wigderson. How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority. In *ACM Symposium on Theory of Computing*, pages 218–229. ACM, May 25–27, 1987.
- [11] P. Paillier. Public-Key Cryptosystems Based on Composite Degree Residuosity Classes. In J. Stern, editor, *Advances in Cryptology*, volume 1592 of *LNCS*, pages 223–238. Springer, May 2–6, 1999.
- [12] H. Polat and W. Du. Privacy-preserving collaborative filtering using randomized perturbation techniques. In *ICDM*, pages 625–628, 2003.
- [13] H. Polat and W. Du. SVD-based collaborative filtering with privacy. In *Proceedings of the 2005 ACM symposium on Applied computing*, pages 791–795, New York, NY, USA, 2005. ACM Press.
- [14] Shopzilla, Inc. Privacy policy, 2009. <http://www.bizrate.com/content/privacy.html>.
- [15] A. C.-C. Yao. Protocols for Secure Computations (Extended Abstract). In *Annual Symposium on Foundations of Computer Science*, pages 160–164. IEEE, November 3–5, 1982.
- [16] S. Zhang, J. Ford, and F. Makedon. Deriving private information from randomly perturbed ratings. In *ICDM*, pages 59–69, 2006.