Privacy-Preserving Distributed Collaborative Filtering

Antoine Boutet, Davide Frey, Rachid Guerraoui, Arnaud Jégou, Anne-Marie Kermarrec

► To cite this version:

Antoine Boutet, Davide Frey, Rachid Guerraoui, Arnaud Jégou, Anne-Marie Kermarrec. Privacy-Preserving Distributed Collaborative Filtering. Computing, 2016, Special Issue on NETYS 2014, 98 (8), pp.827-846. 10.1007/s00607-015-0451-z . hal-01251314

HAL Id: hal-01251314

<https://hal.inria.fr/hal-01251314>

Submitted on 5 Jan 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Privacy-Preserving Distributed Collaborative Filtering

Antoine Boutet¹, Davide Frey¹, Rachid Guerraoui², Arnaud Jégou¹, and Anne-Marie Kermarrec¹

¹ INRIA Rennes, France `first.last@inria.fr`

² EPFL, Switzerland `first.last@epfl.ch`

Abstract. We propose a new mechanism to preserve privacy while leveraging user profiles in distributed recommender systems. Our mechanism relies on two contributions: *(i)* an original obfuscation scheme, and *(ii)* a randomized dissemination protocol. We show that our obfuscation scheme hides the exact profiles of users without significantly decreasing their utility for recommendation. In addition, we precisely characterize the conditions that make our randomized dissemination protocol differentially private.

We compare our mechanism with a non-private as well as with a fully private alternative. We consider a real dataset from a user survey and report on simulations as well as planetlab experiments. We dissect our results in terms of accuracy and privacy trade-offs, bandwidth consumption, as well as resilience to a censorship attack. In short, our extensive evaluation shows that our twofold mechanism provides a good trade-off between privacy and accuracy, with little overhead and high resilience.

1 Introduction

Collaborative Filtering (CF) leverages interest similarities between users to recommend relevant content [19]. This helps users manage the ever-growing volume of data they are exposed to on the Web [7]. But it also introduces a trade-off between ensuring user privacy and enabling accurate recommendations. Decentralized collaborative filtering partially addresses this trade-off by removing the monopoly of a central entity that could commercially exploit user profiles. However, it introduces new privacy breaches: users may directly access the profiles of other users. Preventing these breaches is the challenge we address in this paper. We do so in the context of a news-oriented decentralized CF system.

We propose a twofold mechanism: *(i)* an obfuscation technique applied to user profiles, and *(ii)* a randomized dissemination protocol satisfying a strong notion of privacy. Each applies to one of the core components of a decentralized user-based CF system: clustering and dissemination. Clustering consists in building an interest-based topology, implicitly connecting users with similar preferences: it computes the similarity between profiles, capturing the opinions of users on the items they have been exposed to. The dissemination protocol propagates the items along the resulting topology.

Our obfuscation scheme prevents user machines from exchanging their exact profiles while constructing the interest-based topology. We compute similarities using coarse-grained obfuscated versions of user profiles that reveal only the least sensitive information. To achieve this, we associate each disseminated item with an *item profile*. This profile aggregates information from the profiles of users that liked an item along its dissemination path. This reflects the interests of the portion of the network the item has traversed, gathering the tastes of a community of users that have liked similar items. We use this information to construct filters that identify the least sensitive parts of user profiles: those that are the most popular among users with similar interests. Albeit lightweight, our *obfuscation* scheme prevents any user from knowing, with certainty, the exact profile of another user. Interestingly, we achieve this without significantly hampering the quality of recommendation: the obfuscated profile reveals enough information to connect users with similar interests.

We also characterize the parameters that make our dissemination protocol differentially private [8]. Differential privacy bounds the probability of the output of an algorithm to be sensitive to the presence of information about a given entity—the interests of a user in our context—in the input data. We obtain differential privacy by introducing randomness in the dissemination of items. This prevents malicious players from guessing the interests of a user from the items she forwards.

We compare our mechanism with a non-private baseline as well as with an alternative solution that applies differential privacy to the entire recommendation process. We consider a real dataset from a user survey and report on simulations as well as planetlab experiments. We dissect our results in terms of accuracy and privacy trade-offs, bandwidth consumption, as well as resilience to a censorship attack. Our extensive evaluation shows that our twofold mechanism provides a good trade-off between privacy and accuracy. For instance, by revealing only the least sensitive 30% of a user profile, and by randomizing dissemination with a probability of 0.3, our solution achieves an F1-Score (trade-off between precision and recall) of 0.58, against a value of 0.59 for a solution that discloses all profiles, and a value of 0.57 for the differentially private alternative in a similar setting. Similarly, malicious users can predict only 26% of the items in a user’s profile with our solution, and as much as 70% when using the differentially private one. In addition, our mechanism is very resilient to censorship attacks, unlike the fully differentially private approach.

2 Setting

We consider a decentralized news-item recommender employing user-based collaborative filtering (CF). Its architecture relies on two components: *user clustering* and *item dissemination*. We aim to protect users from *privacy threats*.

User clustering aims at identifying the k nearest neighbors of each user ³. It maintains a dynamic interest-based topology consisting of a directed graph

³ we use the terms ‘node’ and ‘user’ interchangeably to refer to the pair ‘user/machine’

$G(U, E)$, where vertices, $U = u_1, u_2, u_3, \dots, u_n$, correspond to users, and edges, $E = e_1, e_2, e_3, \dots, e_n$, connect users that have the most similar opinions about a set of items $I = i_1, i_2, \dots, i_m$. The system is decentralized: each node records the interests of its associated user, u , in a *user profile*, a vector of tuples recording the opinions of the user on the items she has been exposed to. Each such tuple $P_u = \langle i, v, t \rangle$ consists of an item identifier, i , a score value, v , and a timestamp, t , indicating when the opinion was recorded. Profiles track the interests of users using a sliding window scheme: each node removes from its profile all the tuples that are older than a specified *time window*. This allows the interest-based topology to quickly react to emerging interests while quickly forgetting stale ones. We focus on systems based on binary ratings: a user either *likes* or *dislikes* an item. The interest-based topology exploits two gossip protocols running on each node. The lower-layer random-peer-sampling (RPS) [22] protocol ensures connectivity by maintaining a continuously changing random graph. The upper-layer clustering protocol [23, 5] starts from this random graph and quickly provides each node with its k closest neighbors according to a similarity metric. Several similarity metrics have been proposed [21], we use the Jaccard index in this paper.

Item dissemination exploits the above clustering scheme to drive the dissemination. When a user generates a new item or receives an item she likes, the associated node assumes that this is an interesting item for other users with similar interests. It thus forwards the item to its neighbors in the interest-based topology. If, instead, the user marks an item as *dislike*, the node simply drops it.

Privacy Threats. While decentralization removes the prying eyes of *Big-Brother* companies, it leaves those of curious users who might want to discover the personal tastes of others. In the decentralized item recommender considered, malicious nodes can extract information in two ways: (i) from the profiles they exchange with other nodes (profiles contain information about the interests of users); and (ii) from the predictive nature of the dissemination (a node sends an item only when it likes it). We consider the Honest-But-Curious adversary model [10] where malicious nodes can collude to predict interests from received profiles but cannot cheat in the protocol. In Section 6.6, we also consider attackers modifying their obfuscated profiles to control their location in the interest-based topology (*i.e.* their clustering views).

3 Obfuscation Protocol

Our first contribution is an obfuscation protocol that protects user profiles by (i) aggregating their interests with those of similar users, and (ii) revealing only the least sensitive information to other users. By tuning these two mechanisms, system designers can manage the trade-off between disclosed information and recommendation quality [15]. An excessively obfuscated profile that reveals very

little information is difficult to compromise, but it also provides poor recommendation performance. Conversely, a highly accurate profile yields better recommendations, but does not protect privacy-sensitive information effectively. As we show in Section 6, our obfuscation mechanism provides good recommendation while protecting privacy.

For clarity, this Section describes a simplified version of our obfuscation protocol. Section 4 completes this description with features required by our differentially-private dissemination scheme. Figure 1 gives an overview of the complete protocol.

Fig. 1: *Simplified information flow through the protocol's data structures.*

3.1 Overview

Our protocol relies on random indexing, an incremental dimension reduction technique [24, 13]. To apply it in our context, we associate each item with an *item vector*, a random signature generated by its source node. An *item vector* consists of a sparse d -dimensional bit array. To generate it, the source of an item randomly chooses $b \ll d$ distinct array positions and sets the corresponding bits to 1. It then attaches the *item vector* to the item before disseminating it.

Nodes use item vectors when recording information about items in their obfuscated profiles. Let us consider a node A that receives an item R from another node C as depicted in Figure 1. Node A records whether it likes or dislikes the item in its *private profile*. A node never shares its private profile. It only uses it as a basis to build an *obfuscated profile* whenever it must share interest information with other nodes in the clustering process. Nodes remove the items whose timestamps are outside the latest *time window*. This ensures that all profiles reflect the current interests of the corresponding nodes.

Upon receiving an item R that she likes, user A first updates the item profile of R and then forwards it (Figure 1). To this end, A combines the item vectors of the liked items in its *private profile* and obtains a *compact profile* consisting of a bit map. This dimension reduction introduces some uncertainty because

Algorithm 1: Receiving an item.

```

1 on receive (item  $\langle id^N, t^N \rangle$ , item vector  $S^N$ , item profile  $P^N$ ) do
2   if  $iLike(id^N)$  then
3      $P \leftarrow \langle id^N, t^N, 1, S^N, P^N \rangle$ 
4     buildCompactProfile( $S^N$ )
5     updateItemProfile( $P^N$ )
6     forward( $\langle id^N, t^N \rangle$ ,  $S^N$ ,  $P^N$ )
7   else
8      $P \leftarrow \langle id^N, t^N, 0 \rangle$ 
9   function buildCompactProfile()
10    for all  $\langle id, t, 1, S, P^N \rangle \in P$ 
11    |  $\tilde{P}[i] = S[i]$  OR  $\tilde{P}[i]$ 
12  function updateItemProfile(item vector  $P^N$ )
13    for all  $i \in P^N$ 
14    |  $Sum[i] = Integer(\tilde{P}[i]) + Integer(P^N[i])$ 
15    for all  $i \in$  the  $s$  highest values in  $Sum$ 
16    |  $P^N[i] = 1$ 
17  function forward( $\langle id^R, t^R \rangle$ , item vector  $S^N$ , item profile  $P^N$ )
18    for all  $n \in Neighbors$ 
19    | send  $\langle id^R, t^R \rangle$  with associated  $S^N$  and  $P^N$  to  $n$ 

```

different sets of liked items may result in the same *compact profile* as described in Section 3.2. Then A updates the *item profile* of R : a bitmap that aggregates the compact profiles of the nodes that liked an item. To update it, A combines its own *compact profile* and R 's old *item profile*. This aggregation amplifies the uncertainty that already exists in *compact profiles* and makes R 's *item profile* an obfuscated summary of the interests of the nodes that like R .

Before sharing interest information with other nodes, A must build its *obfuscated profile*. First, it creates a *filter profile* that aggregates the information contained in the *item profiles* of the items it liked. Then, it uses this filter to identify the bits from its *compact profile* that will appear in its *obfuscated profile*. The *filter profile* allows A to select the bit positions that are most popular among the nodes that liked the same items as it did. This has two advantages. First, using the most popular bits makes A 's *obfuscated profile* likely to overlap with those of similar nodes. Second, these bits carry less information than less popular ones, which makes them preferable in terms of privacy.

3.2 Profile Updates

Private Profile A node updates its private profile whenever it generates a new item or receives an item it likes (lines 3 and 8 in Algorithm 1). In either case,

the node inserts a new tuple into its private profile. This tuple contains the item identifier, its timestamp (indicating when the item was generated) and a score value (1 if the node liked the item, 0 otherwise). For liked items, the tuple also contains two additional fields: the item vector, and the item profile upon receipt.

Compact Profile. Unlike private profiles, which contain item identifiers and their associated scores, the compact profile stores liked items in the form of a d -dimensional bit array. As shown in Figure 1, and on lines 14 of Algorithm 1 and 24 of Algorithm 2, a node uses the compact profile both to update the item profile of an item it likes and to compute its obfuscated profile when exchanging clustering information with other nodes. In each of these two cases, the node computes a fresh compact profile as the bitwise OR of the item vectors of all the liked items in its private profile (line 11 of Algorithm 1).

This on demand computation allows the compact profile to take into account only the items associated with the current *time window*. It is in fact impossible to remove an item from an existing compact profile. The reason is that compact profile provides a first basic form of obfuscation of the interests of a user through bit collisions: a bit with value 1 in the compact profile of a node may in fact result from any of the liked items whose vectors have the corresponding bit set.

Compact profiles bring two clear benefits. First, the presence of bit collisions makes it harder for attackers to identify the items in a given profile. Second, the fixed and small size of bit vectors limits the size of the messages exchanged by the nodes in the system. As evaluated, in Section 6.7, this drastically reduces the bandwidth cost of our protocol.

Item Profile. A node never reveals its compact profile. Instead, it injects part of it in the item profiles of the items it likes. Consequently, the item profile of an item aggregates the interests of the users that liked the item along its dissemination path. A parameter s controls how much information from the compact profile nodes include in the item profile.

Let n be a node that liked an item R . When receiving R for the first time, n computes its compact profile as described above. Then, n builds an integer vector as the bit-by-bit sum of the item profile and its own compact profile (line 14 in Algorithm 1). Each entry in this vector has a value in $\{0, 1, 2\}$: node n chooses the s vector positions with the highest values, breaking ties randomly, and creates a fresh profile for item R by setting the corresponding bits to 1 and the remaining ones to 0. Finally, when n generates the profile for a new item, (line 16 in Algorithm 1), it simply sets to 1 the values of s bits from those that are set in its compact profile. This update process ensures that each item profile always contains s bits with value 1.

Filter Profile. Nodes compute their filter profiles whenever they need to exchange clustering information with other nodes (line 22 in Algorithm 2). Unlike the other profiles associated with nodes, this profile consists of a vector of integer values and does not represent the interests of a user. Rather it captures the interests of the community of users that have liked similar items. A node computes the

Algorithm 2: Building obfuscated profile.

```

20 on demand do
21   Algorithm1.buildCompactProfile()
22   buildFilterProfile()
23   for all  $i \in$  the  $s$  highest values in  $F$ 
24   |  $P^*[i] = \tilde{P}[i]$ 
25 function buildFilterProfile()
26   for all  $\langle id, t, 1, S, P^N \rangle \in P$  in the current time window
27   |  $F[i] = F[i] + Integer(P^N[i])$ 

```

value at each position in its filter profile by summing the values of the bits in the corresponding position in the profiles of the items it liked (line 27 in Algorithm 2) in the latest *time window*. This causes the filter profile to record the popularity of each bit within a community of nodes that liked similar items.

Obfuscated Profiles. As shown in Figure 1, a node computes its obfuscated profile whenever it needs to exchange it with other nodes as part of the clustering protocol. As shown in Figure 1, it achieves this by filtering the contents of its compact profile using its filter profile: this yields a bit vector that captures the most popular bits in the node’s community and thus hides its most specific and unique tastes. The fine-grained information contained in the node’s private and compact profiles remains instead secret throughout the system’s operation.

As shown on line 21 and line 22 of Algorithm 2, a node n computes its obfuscated profile by first generating its compact and filter profiles as described above. Then it selects the s positions that have the highest values in the filter profile, breaking ties randomly, and sets the corresponding bits in the obfuscated profile to the values they have in its compact profile. It then sets all the remaining bits in the obfuscated profile to 0.

The resulting profile has s bits (set at 0 or 1) that reflect the node’s compact profile and provide a coarse-grained digest of user interests. Through the value of s , the system designer can control the amount of information that can filter from the compact to the obfuscated profile, and can therefore tune the trade-off between privacy and recommendation quality. It is important to note that the positions of the bits whose value is 1 in the obfuscated profile depend on the filter profile and thus do not suffice to identify the item vectors that contributed to the corresponding compact profile. This prevents isolated attackers from precisely understanding which news items the node liked as shown in Section 6.5.

4 Randomized Dissemination

An attacker can discover the opinions of a user by observing the items she forwards (Section 2). We address this vulnerability through our second contribution: a differentially-private randomized dissemination protocol.

The key idea of our protocol is to randomize the forwarding decision: a node that likes an item drops it with probability pf , while a node that does not like it

Fig. 2: Complete information flow through the protocol's data structures.

forwards it with the same pf . This prevents an attacker from acquiring certainties about a user's interests by observing which items she forwards. However, the attacker could still learn something from the content of the associated *item profiles* (modification of the item profile only when the user likes it). To ensure that the whole dissemination protocol does not expose any non-differentially-private information, we therefore randomize not only forwarding actions, but also the *item profiles* associated with forwarded items. This requires us to modify the protocol described in Section 3 as follows.

First, we introduce a new field in the *private profile*: the *randomized decision*. In addition to record whether the node liked or disliked an item, we use this new field to store the corresponding forwarding decision taken as a result of the randomization process (1 for forward and 0 for drop).

We then introduce a new *randomized compact profile* (as shown in Figure 2). The node fills this profile analogously to the compact profile but it uses the *randomized decision* instead of its actual opinion on the item. The node iterates through all the items for which the randomized decision is 1 and integrates their signatures into the *randomized compact profile* using the same operations described for the non-randomized one.

Finally, the node updates the *item profile* of an item when it decides to forward it as a result of randomization, regardless of whether it likes it or not. Moreover, the node performs this update as described in Section 3.2 except that the node uses its *randomized compact profile* instead of its *compact profile*.

Nodes still use their non-randomized *compact profile* when choosing their neighbors. In this case, they compare their *compact profile* with the *obfuscated profiles* of candidate neighbors. However, the above modifications guarantee that the actual content of the *compact profile* never leaks during dissemination. This guarantees that our dissemination protocol is *differentially private* [8].

A randomized algorithm $\mathcal{A}$ is ϵ -differentially private if it produces approximately the same output when applied to two neighboring datasets (*i.e.* which differ on a single element). In the context of dissemination, the datasets that need to be randomized are vectors of user opinions. Given two neighboring vectors of opinions (*i.e.* differing on a single opinion) $o1 \in \mathcal{D}^n$ and $o2 \in \mathcal{D}^n$, we define differential privacy as follows.

Differential privacy [9] A randomized function $\mathcal{F} : \mathcal{D}^n \rightarrow \mathcal{D}^n$ is ϵ -differentially private, if for any pair of neighboring opinion vectors $\mathbf{o1}, \mathbf{o2} \in \mathcal{D}^n$ and for all $\mathbf{t} \in \mathcal{D}^n$:

$$\Pr[\mathcal{F}(\mathbf{o1}) = \mathbf{t}] \leq e^\epsilon \cdot \Pr[\mathcal{F}(\mathbf{o2}) = \mathbf{t}]$$

This probability is taken over the randomness of $\mathcal{F}$, while e is the base of the natural logarithm.

In the case of our algorithm, we toss a coin each time the user expresses her opinion about an item in order to decide whether the item should be forwarded. This scheme is known as randomized response [25]: instead of randomizing the output of a function f , we randomize each of its inputs independently. Because these inputs as well as the output values are binary $\in \{0, 1\}$, we can rewrite the above equation as follows.

$$\Pr[f(o) = b] \leq e^\epsilon \cdot \Pr[f(1 - o) = b]$$

Our randomization function f flips the opinion o and produces the output $1 - o$ with probability pf . In order to achieve ϵ -differential privacy the value of pf must be such that:

$$1/(e^\epsilon + 1) \leq pf \leq 1/2$$

For space reasons, we omit the details of the reasoning leading to this result, as well as the proof of the equivalence between randomized response and Definition 4. Nonetheless they are similar to those in [4].

This algorithm bounds the amount of information an observer gets when receiving an item from a user. Instead of knowing with certainty that the user liked the item, the observer knows that the user liked it with probability $1 - pf$. However, this does not make our solution fully differentially private, but only the dissemination component. In addition, it can only ensures ϵ -differential privacy when a user expresses her opinion about an item she received, not when she generates a new one. In the latter case, the user always forwards the item.

5 Experimental setup

We implemented and extensively evaluated our approach using a real dataset from a user survey. We also compare our solution with a baseline solution with no privacy mechanism, where profiles are exchanged in clear, and a solution that applies a differentially private mechanism both when generating the profiles that users exchange and upon dissemination. We refer to our solution as OPRD (Obfuscation Profile and Randomized Dissemination) in the following.

5.1 Dataset

To evaluate our approach against a real dataset, we conducted a survey on 200 news items involving 120 colleagues and relatives. We selected news items

randomly from a set of RSS feeds illustrating various topics (culture, politics, people, sports,...). We exposed this list to our test users and gathered their opinions (like/dislike) on each news item. This provided us with a small but *real* dataset of users exposed to exactly the same news items. To scale out our system, we generated 4 instances of each user and news item in the experiments. While this may introduce a bias, this affects accuracy of both our mechanisms and the solutions we compare against.

5.2 Alternatives

We compare our approach with the two following alternatives.

Cleartext profile (CT). This baseline approach implements the decentralized CF solution presented in Section 2 where user profiles are exchanged in clear during the clustering process. This solution does not provide any privacy mechanism.

Differentially private approach (2-DP). This alternative, denoted by 2-DP in the following, applies randomization both when generating user profiles and during dissemination. Every time a user expresses an opinion about an item, the algorithm inverses it with probability pd : this results in a differentially private clustering protocol and a differentially private dissemination protocol. The latter is similar to our randomized dissemination. However, unlike our solution, 2-DP also applies randomness when generating user profiles. When a user dislikes an item, 2-DP considers this item as liked with a probability pd , thus integrating it in the profile of the user and disseminating it to her neighbors. Conversely, when a user likes an item, 2-DP considers it as disliked with probability pd . In this case, it silently drops it without including it in the user’s profile.

2-DP builds user profiles that are structurally similar to our compact profiles. However, they gather the item vectors of the items identified as liked after the randomization of user opinions. This extends the privacy guarantee associated with our dissemination protocol to the profiles of users. This represents a contribution in its own right. For space reasons, we do not include the associated proof. However, it follows a similar intuition than the one presented in Section 4.

As user profiles change over time and are impacted by the dissemination of items, applying a randomization function on cleartext profiles as in [4] is not enough. Iteratively probing the profiles of a user and analyzing the dissemination process could be enough to weaken the privacy guarantee. Instead, 2-DP does not randomize profiles, but it randomizes the opinion of a user on the items she is exposed to. Moreover, it does so independently of the user’s opinion on other items.

2-DP uses the output of its randomization function to build user profiles and drive the dissemination. In particular, users use the resulting randomized profiles to compute their clustering views. We show in Section 6.4 that this introduces a weakness in the context of the decentralized CF scheme considered in this paper. Moreover, section 6.6 shows that 2-DP remains more vulnerable to censorship attacks than our solution.

5.3 Evaluation metrics

Accuracy. We evaluate accuracy along the traditional metrics used in information-retrieval systems: *recall* and *precision*. Both measures are in $[0, 1]$. A recall of 1 means that all interested users have received the item. Yet, a trivial way to ensure a recall of 1 is to send all news items to all users, potentially generating spam. Precision precisely captures the level of spam: a precision of 1 means that all news items reach only users that are interested in them. The F1-Score captures the trade-off between these two metrics and is defined as the harmonic mean of precision and recall [21].

Overhead. We evaluate the overhead of the system in terms of the network traffic it generates. For simulations, we compute the total number of sent messages. For our implementation, we instead measure the average consumed bandwidth. A key parameter that determines network traffic is the *fanout* of the dissemination protocol, *i.e.* the number of neighbors from the interest-based overlay to which nodes forward each item.

Privacy. We define privacy as the ability of a system to hide the profile of a user from other users. We measure it by means of two metrics. The first evaluates to what extent the obfuscated profile is close to the real one by measuring the similarity between the two. We consider the Jaccard index [21] to measure the similarity between a compact profile and the corresponding obfuscated one. The second measures the fraction of items present in a compact profile out of those that can be predicted by analyzing the presence of item vectors in the corresponding obfuscated profile. As item vectors are public, a malicious user can leverage them to guess the contents of the obfuscated profiles of other users, thereby inferring their interests.

6 Performance evaluation

In this section, we evaluate the ability of our solution to achieve efficient information dissemination while protecting the profiles of its users. First, we show that compacting user profiles, filtering sensitive information, and randomizing dissemination do not significantly affect the accuracy of dissemination when compared to CT, yielding slightly better results than 2-DP. Then we analyze the trade-off between accuracy and privacy and show the clear advantage of our solution in protecting user profiles in the context of a censorship attack. Finally, we show the benefits of our solution in term of network cost. We conducted an extensive evaluation through simulations, and through a real implementation deployed on PanelLab. In both cases, we randomly select the source of each item among all users. We refer to our solution as OPRD (Obfuscation Profile and Randomized Dissemination) in the following.

Fig. 3: Impact of compacting the profiles (various b -to- d ratios)

6.1 Compacting profiles

As explained in Section 3.2, our solution associates each item with a (sparse) item vector containing b 1's out of d possible positions. When a user likes an item, we add the corresponding item vector to her compact profile by performing a bitwise OR with the current profile. The ratio between b and d affects the probability of having two items sharing bits at 1 in their vectors, which in turn affects the accuracy of the similarity computation between users. Figure 3 evaluates its effect on performance.

Figure 3a shows the values of the F1-Score depending on network traffic for various values of the b -to- d ratio. The points in each curve correspond to a range of fanout values, the fanout being the number of neighbors to which a user forwards an item she likes: the larger the fanout the higher the load on the network. Figure 3b shows instead the corresponding precision-recall curve. Again, each curve reflects a range of fanout values: the larger the fanout, the higher the recall, and the lower the precision.

Interestingly, the larger the b -to- d ratio, the bigger the difference between our solution and CT. With a low b -to- d ratio, it is unlikely for any two item vectors to contain common bits at 1. As a result, the performance of our solution closely mimics that of CT. When the b -to- d ratio increases, the number of collisions between item vectors—cases in which two distinct item vectors have common bits at 1—also increases. This has two interesting effects on performance.

The first is that the F1-Score increases faster with the fanout and thus with the number of messages: the $b = 10\%$ curve climbs to an F1-Score of 0.4 with less than 400k messages. The curve on Figure 3b shows that this results from a higher recall for corresponding precision values (bump in the $b = 10\%$ curve). The high probability of collisions between item vectors results in some user profiles being similar even though they do not contain many common items. This leads to a topology in which users are less clearly clustered, and in which the items can be disseminated more easily, which explains the high recall value.

The second effect is that the maximum F1-Score attained by the protocol with a large b -to- d ratio (to the right of Figure 3a) stabilizes at lower values. Figure 3b clarifies that this results from a lower maximum recall, as indicated

by the left endpoints of the curves corresponding to high values of b . The artificial similarities caused by a large b —advantageous with small fanout values (small number of messages)—also create false clusters that ultimately inhibit the dissemination of items to large populations of users. This effect is even more prominent with values of b that set a vast majority of the bits in compact profiles to 1 (not shown in the plot).

In the following, we set d to 500 and b to 5 for our evaluations. The values assigned to b and d should be computed depending on the expected number of items per user profile. Explanations about the computation of these values are outside of the scope of this paper, but are similar to those that relate the number of hash functions and the size of a bloom filter [20].

6.2 Filtering sensitive information

In our solution, the size of the filter defines how much information from the compact profile appears in the obfuscated profile. The larger the filter, the more the revealed information. Figure 4a depicts the F1-Score as a function of the number of messages. The performance increases with the size of the filter. Figure 4b shows that this variation comes from the fact that precision strongly decreases when the filter size decreases. The important aspect is that both plots highlight that a filter of 200 bits (*e.g.* 40% of the compact profile) achieves performance values similar to those of a system using full profiles.

Fig. 4: Impact of filtering sensitive information (various filter sizes, fs)

Fig. 5: Impact of obfuscating profiles and randomizing dissemination ($fs = 200$)

6.3 Randomizing the dissemination

We now evaluate the impact of randomizing the dissemination process in addition to the obfuscation protocol evaluated above (the previous results were obtained without randomization). Figure 5a shows the F1-Score for our solution using a filter size of 200 and several values for pf . Performance decreases slightly as we increase the amount of randomness (for clarity, we only show $pf = 0$ and $pf = 0.5$, the other curves being in between). Figure 5b shows that increasing pf results mostly in a decrease in precision.

Fig. 6: Impact of the randomization for 2-DP

6.4 Evaluating 2-DP

In this section, we evaluate the 2-DP alternative defined in Section 5.2. 2-DP reverses the opinions of users with a probability, pd , that affects both the construction of user profiles and the dissemination process. This differs from our solution in which only the dissemination is randomized.

Figure 6a shows the F1-Score of 2-DP versus network traffic for various values of pd . Performance strongly increases at low fanout values for $dp = 0.1$, but decreases for larger values. A small amount of randomness proves beneficial and allows the protocol to disseminate items more effectively with a low fanout. This effect, however, disappears when the number of messages increases at high fanouts. Too much randomness, on the other hand, causes a drastic decrease in the F1-Score. Figure 6b shows that randomness induces an increase in recall with respect to CT and a decrease in precision. The former dominates with low values of pd while the latter dominates for high values.

Figure 7 compares the F1-Score of OPRD using a filter of size of 200 and a pf value of 0.3, with that of CT and 2-DP using a pd of 0.3. We observe that above $2M$ messages, our solution provides slightly better F1-Score values than 2-DP. Overall, however, the best performances of the two approaches are comparable. In the following, we show that this is not the case for their ability to protect user profiles.

Fig. 7: OPRD vs 2-DP: F1-Score vs number of messages

Fig. 8: Randomness vs performance and level of privacy

6.5 Privacy versus accuracy

We evaluate the trade-off between privacy, measured as the ability to conceal the exact profiles of users, and accuracy for both OPRD and 2-DP. OPRD controls this trade-off with two parameters: the size of the filter, and the probability pf . 2-DP controls this trade-off by tuning the probability pd to switch the opinion of the user, impacting both profile generation and the dissemination process.

Figure 8a compares their recommendation performance by measuring the F1-Score values for various filter sizes. The x -axis represents the evolution of the probabilities pf , for our solution, and pd , for 2-DP. We show that the F1-Score of 2-DP decreases faster than ours. The F1-Score of 2-DP with a pd of at least 0.2 is smaller than that of our solution with a filter size greater than 100. In addition, revealing the most popular 10% of the compact profile ($fs = 50$) yields similar performance as 2-DP with $pd \geq 0.3$.

Figure 8b measures the level of privacy as the overlap rate (computed with the Jaccard index) between the compact profile and the obfuscated profile: lower overlap rate implies more privacy. As our randomized dissemination protocol hardly impacts the obfuscated profile, our results are almost independent of pf . 2-DP sees instead its similarity decrease with increasing pd . With $pd = 0.3$, 2-DP yields an overlap rate of about 0.55 with an F1-Score (from Figure 8a) of 0.55. Our approach, on the other hand yields the same overlap rate with a filter size between $150 < fs < 200$, which corresponds to an F1-Score value of about 0.57.

Figure 9, instead, assesses privacy by measuring if the items in a user's real profile can be predicted by an attacker that analyzes the user's public profile. Note that in 2-DP, the real profile is the one that would exist without random

Fig. 9: *Profile prediction*

perturbations. We evaluate this aspect by measuring the recall and the precision of predictions. Prediction recall measures the fraction of correctly predicted items out of those in the compact profile. Prediction precision measures the fraction of correct predictions out of all the prediction attempts. For our solution, in Figure 9a, we use a $pf = 0.2$ to control the randomized dissemination, and vary the filter size. For 2-DP (Figure 9b), we instead vary pd .

The plots show that while our approach is subject to fairly precise predictions, these cover only a small fraction of the compact profile with reasonable values of fs . With $fs = 200$, the prediction recall is of about 30%. In contrast, 2-DP exposes a higher number of items from the compact profile. With $pd = 0.2$ the prediction recall is 0.8 with a prediction precision of 0.6. The curves for prediction effectiveness, computed as the harmonic mean of recall and precision, further highlight our approach's ability to strike an advantageous balance between privacy and recommendation performance.

The two plots also show the average popularity of the predicted items. We observe that when the filter size decreases, the correctly predicted items are among the most popular ones, which are arguably the least private.

Finally, we also observe that the compact profile itself provides a small protection to the prediction of items due to its inherent collision rate. With a filter of size 500 (*e.g.* with no difference between the compact and the public profile), the error rate is equal to 0.15.

6.6 Resilience to a censorship attack

We illustrate the resilience of our obfuscation protocol against censorship by implementing a simple eclipse attack [18]. A coalition of censors mirrors the (obfuscated) profile of a target node in order to populate its clustering view. This in turn isolates it from the remaining nodes since its only neighbors are all censors. If the user profiles are exposed in clear, the profile of the censors matches exactly that of the target node: this gives censors a very high probability to enter its view. Once the censors have fully populated the target node's view, they simply intercept all the messages sent by the target node, preventing their dissemination. We evaluate the efficiency of this attack with two metrics: the poisoning rate of the target's clustering view by attackers; and the fraction of honest nodes (*e.g.* not censors) reachable by the target when it sends an item.

Fig. 10: *Resilience to censorship*

We ran this attack for each user in the dataset. The x -axis represents the users in the experiment sorted by their sensitivity to the attack. Figure 10a and Figure 10b depict the results obtained with a cluster size of 50, and 50 censors (we observe similar results independently of the cluster size). In addition, this experiment uses a filter of 125 and $pf = 0.2$ for our solution, and $pd = 0.2$ for 2-DP. We can clearly see that 2-DP is not effective in preventing censorship attacks: only 150 nodes have a poisoning rate lower than 1. This is because 2-DP computes similarities using the randomized compact profile, which it also shares with other users. Therefore 2-DP exhibits exactly the same vulnerability as CT. The censors can trivially match the profile of the target node.

Our approach is more resilient to this censorship attack. It is difficult for censors to intercept all messages sent by the target and only a third of the nodes have a fully poisoned clustering view. The obfuscated profile only reveals the least sensitive information to other nodes: censors only mirror a coarse-grained sub part of the target node’s profile. Consequently, their profiles are more likely to resemble those of users with correlated interests than to match the target profile. Figure 8b confirms this observation by showing the overlap between obfuscated and compact profiles. The resilience of OPRD is driven by the size of the obfuscation filter, the smaller the filter, the more resilient the protocol.

6.7 Bandwidth consumption

We also conducted experiments using our prototype with 215 users running on approximately 110 PlanetLab nodes in order to evaluate the reduction of network cost resulting from the compactness of our profiles. The results in terms of F1-Score, recall, and precision closely mimic those obtained with our simulations and are therefore omitted. Table 1 shows the bandwidth cost of our protocols in terms of bandwidth: our obfuscation protocol is effective in reducing the bandwidth consumption of decentralized collaborative filtering. The cost associated with our obfuscated solution is about one third of that of the solution based on cleartext profiles.

Fanout	2	4	6	8	10	15	20
CT	1.8	3.0	4.4	6.5	8.2	12	14
OPRD	0.8	1.1	1.5	1.7	2.7	2.8	4.1

Table 1: *Bandwidth usage in kbps per node in PlanetLab*

7 Related work

Privacy is important in many applications. Several approaches [2, 16, 17] use randomized masking distortion techniques to preserve the privacy of sensitive data. However, [12] shows that the predictable structure in the spectral domain of the random distortion can seriously compromise privacy. In the same vein, [14] shows that the variances of the random noises have an important impact on the possibility to filter noise from the original data. In our solution, instead of adding perturbation to user profiles, we exchange with other users a coarse-grain version of this profile only revealing its least sensitive information. The perturbation applied on the item profile is not random and depends on the interest of users. This makes it harder to separate privacy sensitive information from the introduced distortion.

Some authors [1] designed a statistical measure of privacy based on differential entropy. However, it is difficult to evaluate its meaning and its impact on sensitive data. Differential privacy was considered in [8, 11]. In a distributed settings, [4] proposed a differentially private protocol to measure the similarity between peers. While this solution works well with static profiles, its differential privacy is not preserved when profiles are dynamic as in recommendation systems. In addition, still in the context of recommendation systems, [15] highlights the trade-off between privacy and accuracy.

Other approaches [6] exploit homomorphic encryption in a P2P environment to secure multi-party computation techniques. Similarly, [3] proposes an architecture for privacy preserving CF by replacing the single server providing the service with a coalition of trusted servers.

8 Concluding Remarks

The motivation of this work is to make distributed CF resilient to privacy and censorship attacks without jeopardizing the quality of recommendation. We proposed a mechanism that relies on two components: (i) an obfuscation scheme revealing only the least sensitive information in the profiles of users, and (ii) a randomization-based dissemination protocol ensuring differential privacy during the dissemination. We showed the viability of our mechanism by comparing it with a non-private and a fully (differentially) private alternative. However, many questions remain open. In particular, evaluating the fundamental trade-offs between privacy, resilience to censorship, and recommendation quality constitutes an interesting research direction.

References

1. D. Agrawal and C. C. Aggarwal. On the design and quantification of privacy preserving data mining algorithms. In *PODS*, 2001.
2. R. Agrawal and R. Srikant. Privacy-preserving data mining. In *SIGMOD*, 2000.
3. W. Ahmad and A. Khokhar. An architecture for privacy preserving collaborative filtering on web portals. In *IAS*, 2007.
4. M. Alaggan, S. Gambs, and A.-M. Kermarrec. BLIP: Non-interactive Differentially-Private Similarity Computation on Bloom Filters. In *SSS*, 2012.
5. A. Boutet, D. Frey, R. Guerraoui, A. Jégou, and A.-M. Kermarrec. WhatsUp Decentralized Instant News Recommender. In *IPDPS*, 2013.
6. J. Canny. Collaborative filtering with privacy via factor analysis. In *SIGIR*, 2002.
7. A. S. Das, M. Datar, A. Garg, and S. Rajaram. Google news personalization: scalable online collaborative filtering. In *WWW*, 2007.
8. C. Dwork. Differential privacy: a survey of results. In *TAMC*, 2008.
9. C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*. Springer, 2006.
10. O. Goldreich. Cryptography and cryptographic protocols. *Distrib. Comput.*, 2003.
11. A. Haeberlen, B. C. Pierce, and A. Narayan. Differential privacy under fire. In *SEC*, 2011.
12. Z. Huang, W. Du, and B. Chen. Deriving private information from randomized data. In *SIGMOD*, 2005.
13. P. Kanerva, J. Kristoferson, and A. Holst. Random indexing of text samples for latent semantic analysis. In *CCSS*, 2000.
14. H. Kargupta, S. Datta, Q. Wang, and K. Sivakumar. On the privacy preserving properties of random data perturbation techniques. In *ICDM*, 2003.
15. A. Machanavajjhala, A. Korolova, and A. D. Sarma. Personalized social recommendations: accurate or private. *VLDB*, 2011.
16. H. Polat and W. Du. Privacy-preserving collaborative filtering using randomized perturbation techniques. In *ICDM*, 2003.
17. H. Polat and W. Du. Svd-based collaborative filtering with privacy. In *SAC*, 2005.
18. A. Singh, M. Castro, P. Druschel, and A. Rowstron. Defending against eclipse attacks on overlay networks. In *SIGOPS*, 2004.
19. X. Su and T. M. Khoshgoftaar. A survey of collaborative filtering techniques. *Advances in Artificial Intelligence*, 2009.
20. Sasu Tarkoma, Christian Esteve Rothenberg, and Eemil Lagerspetz. Theory and practice of bloom filters for distributed systems. *IEEE Communications Surveys and Tutorials*, pages 131–155, 2012.
21. C. J. van Rijsbergen. *Information retrieval*. Butterworth, 1979.
22. S. Voulgaris, D. Gavidia, and M. v. Steen. Cyclon: inexpensive membership management for unstructured p2p overlays. *Journal of Network and Systems Management*, 2005.
23. S. Voulgaris and M. v. Steen. Epidemic-style management of semantic overlays for content-based searching. In *Euro-Par*, 2005.
24. M. Wan, A. Jönsson, C. Wang, L. Li, and Y. Yang. A random indexing approach for web user clustering and web prefetching. In *PAKDD*, 2012.
25. Stanley L. Warner. Randomized response: a survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, March, 1965.