

# Cryptographic Treatment of Private User Profiles

Kryptographische Handhabung privater Benutzerprofile  
Bachelor-Thesis von Felix Günther  
November 2010



TECHNISCHE  
UNIVERSITÄT  
DARMSTADT

Fachbereich Informatik  
Peer-to-Peer Networks Group

---

Cryptographic Treatment of Private User Profiles  
Kryptographische Handhabung privater Benutzerprofile

Vorgelegte Bachelor-Thesis von Felix Günther

1. Gutachten: Prof. Dr. Thorsten Strufe
2. Gutachten: Prof. Dr. Mark Manulis

Tag der Einreichung:

---

# Erklärung zur Bachelor-Thesis

Hiermit versichere ich, die vorliegende Bachelor-Thesis ohne Hilfe Dritter nur mit den angegebenen Quellen und Hilfsmitteln angefertigt zu haben. Alle Stellen, die aus Quellen entnommen wurden, sind als solche kenntlich gemacht. Diese Arbeit hat in gleicher oder ähnlicher Form noch keiner Prüfungsbehörde vorgelegen.

Darmstadt, den 15. November 2010

---

(Felix Günther)

---

“The right to be let alone is indeed the beginning of all freedom.”

William O. Douglas (1898 – 1980)  
Associate Justice of the United States Supreme Court

---

## Abstract

The publication of private data in user profiles in a both secure and private way is a rising problem and of special interest in, e.g., online social networks that become more and more popular. Current approaches, especially for decentralized networks, often do not address this issue or impose large storage overhead.

In this work, we present a cryptographic approach to *private profile management* that is seen as a building block for applications in which users maintain their own profiles, publish and retrieve data, and authorize other users to access different portions of data in their profiles. We formalize *confidentiality* and *unlinkability* as two main security and privacy goals for the data which is kept in profiles and users who are authorized to retrieve this data.

Based on the formal model, we present two distinct constructions leveraging different encryption techniques and analyze their characteristics as well as their applicability to real-world online social networks, discussing the imposed overhead in multiple scenarios.

---

## Zusammenfassung

Die Veröffentlichung von privaten Daten in Benutzerprofilen in einer sicheren und privaten Art und Weise, insbesondere im Hinblick auf soziale Netzwerke im Internet (Online Social Networks), ist ein Problem von stark wachsender Bedeutung. Bisherige Ansätze, gerade im Bereich von dezentralen Netzwerken, blenden diese Fragestellung häufig aus oder lösen sie mit unzufriedenstellendem zusätzlichem Speicherbedarf.

In dieser Arbeit präsentieren wir einen kryptographischen Ansatz zur Verwaltung von privaten Benutzerprofilen, der als Baustein in Systemen eingesetzt werden kann, in denen Benutzer eigene Profile verwalten, Daten veröffentlichen und aus Profilen abrufen sowie anderen Benutzern Zugriff auf bestimmte Daten in ihrem eigenen Profil gewähren können. Wir formalisieren in diesem Zusammenhang *Confidentiality* (Vertraulichkeit) und *Unlinkability* (Unverbindbarkeit) als zwei zentrale Ziele im Bezug auf Sicherheit von Daten in Profilen und Privatsphäre von Benutzern, denen der Zugriff auf bestimmte Daten erlaubt ist.

Basierend auf diesem formalen Model präsentieren wir zwei verschiedene Konstruktionen, die unterschiedliche kryptographische Verfahren einsetzen, und analysieren ihre Eigenschaften sowie ihre Verwendbarkeit in realen Online Social Networks im Hinblick auf den durch sie hervorgerufenen, zusätzlichen Ressourcenbedarf.

---

# Contents

<b>1</b>	<b>Introduction</b>	<b>6</b>
1.1	Organization . . . . .	6
<b>2</b>	<b>Related Work</b>	<b>7</b>
<b>3</b>	<b>Private User Profiles: Model and Definitions</b>	<b>8</b>
3.1	Management of User Profiles . . . . .	8
3.1.1	Users . . . . .	8
3.1.2	Profiles . . . . .	8
3.1.3	Profile Management Scheme . . . . .	8
3.2	Adversarial Model . . . . .	9
3.3	Security and Privacy Requirements . . . . .	9
3.3.1	Confidentiality . . . . .	9
3.3.2	Unlinkability . . . . .	10
<b>4</b>	<b>Private Profiles with Shared Keys</b>	<b>11</b>
4.1	Specification of PMS-SK . . . . .	11
4.2	Complexity Analysis . . . . .	11
4.3	Security and Privacy Analysis . . . . .	11
4.4	Further Optimizations . . . . .	12
<b>5</b>	<b>Private Profiles with Broadcast Encryption</b>	<b>13</b>
5.1	Specification of PMS-BE . . . . .	13
5.2	Complexity Analysis . . . . .	14
5.3	Confidentiality of PMS-BE . . . . .	14
5.4	Privacy of PMS-BE . . . . .	15
<b>6</b>	<b>Taking the Model to Reality</b>	<b>18</b>
6.1	Key Overhead in general . . . . .	18
6.2	Key Overhead in real-world Settings . . . . .	19
6.2.1	Facebook . . . . .	19
6.2.2	Twitter . . . . .	19
6.2.3	XING . . . . .	19
6.2.4	Flickr . . . . .	19
6.3	Discussion . . . . .	19
<b>7</b>	<b>Conclusion and Outlook</b>	<b>21</b>

---

# 1 Introduction

Publishing personal profiles and other means of sharing private data are increasingly popular on the web. Online social networks (OSNs) arguably are the most accepted networked service, today. Facebook alone, serving a claimed base of over 500 Million active users<sup>1</sup>, surpassed Google and currently enjoys the highest utilization duration by their users and one of the highest access frequencies of all web sites since January 2010<sup>2</sup>. Its users share 90 pieces of content per month on average, mainly consisting of personally identifiable information. Protecting this data against unauthorized access is of utmost importance, since users store private and sensitive data in their OSN profiles.

The *confidentiality* of published data, meant to be shared with only a chosen group of users, is already important in centralized services. Yet, it becomes even more pressing when establishing decentralized OSNs, which have been proposed recently [9, 4, 15] in an attempt to avoid the centralized control and omnipotent access of commercial service providers. *Unlinkability* is the more subtle requirement of protecting the identity of users who interact with or access certain chunks of published data. It is frequently missed and only few general solutions achieve this privacy goal [2].

A serious corpus of solutions has been proposed to address these issues in the past. Yet, there so far exist no appropriate definitions for secure and private management of user profiles, even from the cryptographic point of view, as we show in Section 2.

This work hence takes a cryptographic approach to address the management of user profiles in a secure and privacy-friendly way. To this end, its goal is to come up with a well-defined model and provably secure solutions, both to ensure the confidentiality of data, which individually is published by users in their social profiles, as well as the privacy of the users who are allowed to access this data. For this purpose we formally define confidentiality and privacy in the given context, first. Our model further addresses several fundamental properties of a profile management scheme (PMS). They comprise the ability of users (profile owners) to publish and remove data and their ability to grant, modify, and revoke access rights to the published data as well as to retrieve data published in other users' profiles. In particular, we consider PMS as an independent building block, without relying on the higher-level application or any other parties to perform these tasks.

After the specification of the model, we describe two provably secure solutions that use different techniques: our first solution, called PMS-SK, combines symmetric encryption with shared keys that are distributed amongst the authorized users. Our second solution, called PMS-BE, involves broadcast encryption techniques to reduce the number of keys that have to be distributed. Both solutions have their advantages and disadvantages with respect to the key overhead they impose and the level of privacy they guarantee, as we show in our subsequent analysis. In particular, PMS-SK provides confidentiality and perfect unlinkability, but imposes an overhead of keys linear in the number of attributes a user is allowed to access, which results in equivalent storage overhead. PMS-BE reduces the key overhead to a constant value at the cost of lower privacy, expressed through the requirement of *anonymity*, which we also model and formally relate to the stronger notion of unlinkability. We further discuss the trade-off between privacy and efficiency by evaluating some complexity characteristics of both approaches and suggest several optimizations that could enhance their performance, while preserving their security and privacy guarantees.

Finally, we give a short idea of how both constructions perform in real-world OSN settings, analyzing the concrete overhead they impose based on current statistics and research. We discuss the results with a view to the advantages and disadvantages of both constructions.

---

## 1.1 Organization

---

The rest of this work is organized as follows. First, we discuss previous cryptographic and non-cryptographic work on private management of user profiles in the following Section. In Section 3, we introduce our formal model of a profile management scheme and define two requirements upon this model: confidentiality and unlinkability. Subsequently, we present our shared-key construction PMS-SK in Section 4, evaluate its complexity, and formally address its security and privacy properties. There, we also present ideas on how to optimize the scheme using further cryptographic techniques. The second construction, PMS-BE, which bases on broadcast encryption, is introduced in Section 5, where we discuss its complexity and relate it to that of PMS-SK. Moreover, we analyze the confidentiality and privacy of the scheme, introduce the weaker privacy notion of anonymity, and show its separation from unlinkability. In Section 6, we analyze the behavior of both constructions in real-world OSNs, discuss the overhead they impose and relate it to advantages and disadvantages of the approaches. Finally, we conclude in Section 7.

---

<sup>1</sup> <http://www.facebook.com/press/info.php?statistics>, October 2010

<sup>2</sup> <http://blog.nielsen.com/nielsenwire/>, October 2010



---

## 2 Related Work

Substantial amount of work has been carried out in the field of secure and private publication of sensitive data in online social networks (OSNs), demonstrating threats and proposing countermeasures. For example, Gross et al. [16] and Zheleva et al. [27] studied how access patterns of users to the information stored in user profiles and membership of users in different social groups can be exploited for the disclosure of private data.

Amongst the non-cryptographic solutions is the approach proposed by Carminati et al. [7, 8], where access to private data is modeled using semantic rules taking into account the depth of social relationships and the amount of trust amongst the users. In addition to being semi-centralized, this approach requires synchronous communication — a significant limitation in our case.

Several cryptographic approaches to improve confidentiality and privacy in existing, mostly centralized online social networks have been suggested in the past: Lucas et al. [18] presented *flyByNight*, an application to encrypt sensitive data in Facebook. Tootoonchian [24] proposed a system called *Lockr* to improve privacy in both centralized and decentralized social networks. Yet, both approaches are not able to keep security and/or privacy up under certain attacks: *flyByNight* relies on the Facebook servers as middlemen and thus enables them to introduce malicious code or keys whilst in *Lockr* malicious users are able to reveal relationship keys or disclose relationship metadata for access control, compromising privacy properties of the system. Another cryptographic approach is *Scramble!* [21], a Firefox plugin that uses the OpenPGP standard [5] to encrypt data relying on its public-key infrastructure. Moreover, *Scramble!* tries to achieve recipient anonymity by omitting the public identifiers of recipients in the ciphertext and allows for data storage on third-party systems using “tiny URLs”, thus reducing the size of ciphertexts. Nevertheless, the approach implies linear storage overhead and, as it relies on OpenPGP, is vulnerable to active attacks as shown by Barth et al. [2].

A number of solutions aim at binding the access to private data with some fine-grained access policies. For example, Graffi et al. [14] implemented an approach based on symmetric encryption of profile items with independent shared keys, yet without actually specifying or formally analyzing the desired security and privacy properties. The OSN *Persona*, presented by Baden et al. [1], implements ciphertext-policy attribute-based encryption (CP-ABE) [3] for the enforcement of access rules to the encrypted profile data (e.g., “ ‘neighbor’ AND ‘football fan’ ”). Their approach aims at confidentiality of attributes but does not guarantee privacy.

Recently, Zhu et al. [28] proposed a collaborative framework to enforce access control in OSNs by the use of a new group-oriented convergence cryptosystem. Their scheme is centralized and focuses on the joint publication of data within the communities, less on the individual users and protection of their own profiles and data.

A somewhat more general construct for privacy-preserving distribution of an encrypted content was proposed by Barth et al. [2] using public-key broadcast encryption. Of particular interest is their notion of *recipient privacy*, which is supposed to hide the identities of recipients of the broadcast content and can be applied for the private distribution of shared keys in our PMS-SK approach (cf. Remark 1 in Section 4) at the cost of linear storage overhead in the number of recipients. Their scheme could also be used as a building block for a private profile management scheme that would, however, require linear storage overhead for the distribution of ciphertexts.

---

## 3 Private User Profiles: Model and Definitions

In this Section, we formally model our scheme for user profile management based on previously introduced notions of users, profiles, and the like. We then define an adversarial model which we use to formally specify the security requirements of confidentiality and unlinkability.

---

### 3.1 Management of User Profiles

---

Subsequently, we introduce our notions of users and profiles, followed by the formal definition of our profile management scheme.

---

#### 3.1.1 Users

---

Let  $\mathcal{U}$  denote a set of at most  $N$  users. We do not distinguish between users and their identities but assume that each identity  $U \in \mathcal{U}$  is unique. Furthermore, we assume that users can create authentic and, if necessary, confidential communication channels. This assumption is motivated by the fact that the profile management scheme will likely be deployed as a building block within an application like an online social network, where users typically have other means of authentication. In this way we can focus on the core functionality of the profile management scheme, namely the management of and access to the profile data.

---

#### 3.1.2 Profiles

---

A profile  $P$  is modeled as a set of pairs  $(a, \bar{d}) \in \mathcal{I} \times \{0, 1\}^*$  where  $\mathcal{I} \subseteq \{0, 1\}^*$  is the set of possible attribute indices  $a$  and  $\bar{d}$  are corresponding values stored in  $P$ . We assume that within a profile  $P$  attribute indices are unique. Furthermore, we assume that each profile  $P$  is publicly accessible but is distributed in an authentic manner by its owner  $U_p \in \mathcal{U}$ . Also, every user  $U$  owns at most one profile and the profile owned by  $U$  is denoted by  $P_U$ . The authenticity of profiles means, that their content can only be manipulated by their respective owner who is in possession of the corresponding profile management key  $pmk$ . Since one of the goals will be to ensure confidentiality of attributes, we assume that for each publicly accessible value  $\bar{d}$  there exists the actual attribute  $d$  and that for any pair  $(a, \bar{d}) \in P$  the profile owner  $U_p$  can implicitly retrieve the corresponding  $d$  as well as the group  $\mathcal{G} \subseteq \mathcal{U}$  of users who are currently authorized to access  $d$ . By  $\mathcal{G}_{a,p}^*$  we denote the set of users that have ever been authorized to access the attribute indexed by  $a$  within the profile  $P$  (we assume that  $U_p \in \mathcal{G}_{a,p}^*$  for all attributes in  $P$ ).

---

#### 3.1.3 Profile Management Scheme

---

We are now ready to define the syntax of the profile management scheme.

**Definition 1** (Profile Management Scheme). A profile management scheme PMS consists of the five algorithms `Init`, `Publish`, `Retrieve`, `Delete` and `ModifyAccess` defined as follows:

`Init`( $\kappa$ ): On input the security parameter  $\kappa$ , this probabilistic algorithm initializes the scheme and outputs an empty profile  $P$  together with the private profile management key  $pmk$ . `Init` is executed by the owner  $U_p$ .

`Publish`( $pmk, P, (a, d), \mathcal{G}$ ): On input a profile management key  $pmk$ , a profile  $P$ , a pair  $(a, d) \in \mathcal{I} \times \{0, 1\}^*$  (such that  $a$  is not yet in the profile), and a group of users  $\mathcal{G}$ , this probabilistic algorithm transforms the attribute  $d$  into value  $\bar{d}$ , adds  $(a, \bar{d})$  to  $P$ , and  $\mathcal{G}$  to  $\mathcal{G}_{a,p}^*$ . It outputs the modified  $P$  and a retrieval key  $rk_U$  for each  $U \in \mathcal{G}$  (that may be newly generated or modified). Optionally, it updates  $pmk$ . `Publish` is executed by the owner  $U_p$ .

`Retrieve`( $rk_U, P, a$ ): On input a retrieval key  $rk_U$ , a profile  $P$ , and an attribute index  $a$ , this deterministic algorithm checks whether  $(a, \bar{d}) \in P$ , and either outputs  $d$  or rejects with  $\perp$ . `Retrieve` can be executed by any user  $U \in \mathcal{U}$  being in possession of the key for  $a$  in  $rk_U$ .

**Delete**( $pmk, P, a$ ): On input a profile management key  $pmk$ , a profile  $P$ , and an attribute index  $a$ , this possibly probabilistic algorithm checks whether  $(a, \bar{d}) \in P$ , and if so outputs modified profile  $P = P \setminus (a, \bar{d})$ . Optionally, it updates  $pmk$  and  $rk_U$  of all  $U \in \mathcal{G}$  where  $\mathcal{G}$  denotes the set of users authorized to access the pair with index  $a$  at the end of the execution. **Delete** is executed by the owner  $U_p$ .

**ModifyAccess**( $pmk, P, a, U$ ): On input a profile management key  $pmk$ , a profile  $P$ , an attribute index  $a$ , and some user  $U \in \mathcal{U}$  this probabilistic algorithm checks whether  $(a, \bar{d}) \in P$  for some  $\bar{d}$ , and if so finds the set  $\mathcal{G}$  of users that are authorized to access the attribute  $d$ . The algorithm then proceeds according to the one of the following two cases:

- If  $U \in \mathcal{G}$  then it updates  $\mathcal{G} = \mathcal{G} \setminus \{U\}$  (i.e., user  $U$  is removed from  $\mathcal{G}$ ).
- If  $U \notin \mathcal{G}$  then it updates  $\mathcal{G} = \mathcal{G} \cup \{U\}$  and  $\mathcal{G}_{a,p}^* = \mathcal{G}_{a,p}^* \cup \{U\}$  (i.e.,  $U$  is added to both  $\mathcal{G}$  and  $\mathcal{G}_{a,p}^*$ ).

Finally, the algorithm outputs the modified profile  $P$ . Optionally, it updates  $pmk$  and the retrieval keys  $rk_U$  of all  $U \in \mathcal{G} \cup \{U\}$ . **ModifyAccess** is executed by the owner  $U_p$ .

We remark that some profile management schemes may include an additional algorithm **ModifyAttribute** allowing  $U_p$  to modify the attribute  $d$  behind some pair  $(a, \bar{d}) \in P$  in a more efficient way than by consecutive execution of **Delete** and **Publish**. However, for the security treatment of profile management schemes it is sufficient to consider only **Delete** and **Publish** that make the modification of attributes possible.

---

### 3.2 Adversarial Model

---

In order to define security and privacy of a profile management scheme PMS, we consider a probabilistic polynomial-time (PPT) adversary  $\mathcal{A}$  that knows all users in the system, i.e., the set  $\mathcal{U}$  is assumed to be public, and interacts with them via the following set of queries:

**Corrupt**( $U$ ): This corruption query gives  $\mathcal{A}$  all secret keys known to  $U$ , including the profile management key  $pmk$  and all retrieval keys  $rk_U$  (with which  $U$  can access other users' profiles).

$U$  is added to the set of corrupted users that we denote by  $\mathcal{C} \subseteq \mathcal{U}$ .

**Publish**( $P, (a, d), \mathcal{G}$ ): In response, **Publish**( $pmk, P, (a, d), \mathcal{G}$ ) is executed using  $pmk$  of  $U_p$ .  $\mathcal{A}$  is then given the modified profile  $P$  and all updated keys of corrupted users  $U \in \mathcal{C}$ .

**Retrieve**( $P, a, U$ ): In response, **Retrieve**( $rk_U, P, a$ ) is executed using  $rk_U$  of  $U$  and its output is given back to  $\mathcal{A}$ .

**Delete**( $P, a$ ): In response, **Delete**( $pmk, P, a$ ) is executed using  $pmk$  of  $U_p$ .  $\mathcal{A}$  is then given the modified profile  $P$  and all updated keys belonging to corrupted users  $U \in \mathcal{C}$ .

**ModifyAccess**( $P, a, U$ ): In response, **ModifyAccess**( $pmk, P, a, U$ ) is executed using  $pmk$  of  $U_p$ .  $\mathcal{A}$  is then given the modified profile  $P$  and all updated keys belonging to corrupted users  $U \in \mathcal{C}$ .

---

### 3.3 Security and Privacy Requirements

---

Subsequently, we define two security requirements for a profile management scheme. Our first requirement, called *confidentiality*, aims to protect attributes  $d$  stored in a profile from unauthorized access. Our second requirement, called *unlinkability* aims at protecting the privacy of users in the following sense: a profile management scheme should hide information on whether a user  $U$  has been authorized to access some attribute in a profile of another user  $U_p$ , and, moreover, it should not leak information whether different attributes within a profile can be accessed by the same user, even if the adversary has access to these attributes as well.

---

#### 3.3.1 Confidentiality

---

We model *confidentiality* in Definition 2 using the standard indistinguishability approach, similar to the one used in definitions of secure encryption, i.e., it should be computationally infeasible for an adversary  $\mathcal{A}$  to decide which attribute  $d$  is referenced by an index  $a$ .

**Definition 2** (Confidentiality). Let PMS be a profile management scheme from Definition 1 and  $\mathcal{A}$  be a PPT adversary interacting with users via queries from Section 3.2 within the following game  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{conf}}$ :

1.  $\text{Init}(\kappa)$  is executed for all users  $U \in \mathcal{U}$ .
2.  $\mathcal{A}$  can execute arbitrary operations and ask queries. At some point it outputs  $(a, d_0), (a, d_1) \in \mathcal{I} \times \{0, 1\}^*$ ,  $\mathcal{G}_t \subset \mathcal{U}$ , and  $U_p \in \mathcal{U} \setminus \mathcal{G}_t$  such that neither  $U_p$  nor any  $U \in \mathcal{G}_t$  is corrupted (i.e.,  $(\{U_p\} \cup \mathcal{G}_t) \cap \mathcal{C} = \emptyset$ ) and  $|d_0| = |d_1|$  (i.e.,  $d_0$  and  $d_1$  have the same length).
3. A bit  $b \in_R \{0, 1\}$  is chosen uniformly,  $\text{Publish}(\text{pmk}, P, (a, d_b), \mathcal{G}_t)$  with  $\text{pmk}$  of  $U_p$  is executed, and the modified  $P$  is given to  $\mathcal{A}$ .
4.  $\mathcal{A}$  can execute arbitrary operations and ask queries. At some point it outputs a bit  $b' \in \{0, 1\}$ .
5.  $\mathcal{A}$  wins, denoted by  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{conf}} = 1$ , if all of the following holds:
  - $b' = b$ .
  - $U_p \notin \mathcal{C}$ .
  - $\mathcal{A}$  did not query  $\text{Retrieve}(P, a, U)$  with  $U \in \mathcal{G}_{a, P}^*$ .
  - $\mathcal{G}_{a, P}^* \cap \mathcal{C} = \emptyset$  (users that have ever been authorized to access the attribute indexed by  $a$  in  $P$  are not corrupted).

The advantage probability of  $\mathcal{A}$  in winning the game  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{conf}}$  is defined as

$$\text{Adv}_{\mathcal{A}, \text{PMS}}^{\text{conf}}(\kappa) := \left| \Pr \left[ \text{Game}_{\mathcal{A}, \text{PMS}}^{\text{conf}} = 1 \right] - \frac{1}{2} \right|$$

We say that PMS provides confidentiality if for any PPT adversary  $\mathcal{A}$  the advantage  $\text{Adv}_{\mathcal{A}, \text{PMS}}^{\text{conf}}(\kappa)$  is negligible.

---

### 3.3.2 Unlinkability

---

We model *unlinkability* of a profile management scheme in Definition 3 using the indistinguishability approach as well, but this time  $\mathcal{A}$  has to decide which user has been authorized to access the attribute, either via  $\text{Publish}$  or  $\text{ModifyAccess}$  queries.

**Definition 3 (Unlinkability).** Let PMS be a profile management scheme from Definition 1 and  $\mathcal{A}$  be a PPT adversary interacting with users via queries from Section 3.2 within the following game  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{unlink}}$ :

1.  $\text{Init}(\kappa)$  is executed for all users  $U \in \mathcal{U}$ .
2.  $\mathcal{A}$  can execute arbitrary operations and ask queries. At some point it outputs  $U_0, U_1, (a, d)$ , and  $U_p$  (owner of some profile  $P$ ).
3. A bit  $b \in_R \{0, 1\}$  is chosen uniformly and
  - if  $(a, *) \notin P$ , then  $\text{Publish}(\text{pmk}, P, (a, d), \{U_b\})$  with  $\text{pmk}$  of  $U_p$  is executed.
  - if  $(a, *) \in P$ , then  $\text{ModifyAccess}(\text{pmk}, P, a, \{U_b\})$  with  $\text{pmk}$  of  $U_p$  is executed.

In both cases  $\mathcal{A}$  is given the modified profile  $P$  and the possibly updated retrieval keys  $rk_U$  for all  $U \in \mathcal{C}$ .
4.  $\mathcal{A}$  can execute arbitrary operations and ask queries. At some point it outputs a bit  $b' \in \{0, 1\}$ .
5.  $\mathcal{A}$  wins, denoted by  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{unlink}} = 1$ , if all of the following holds:
  - $b' = b$ .
  - $\{U_0, U_1, U_p\} \cap \mathcal{C} = \emptyset$ .
  - $\mathcal{A}$  neither queried  $\text{Retrieve}(P, a, U_0)$  nor  $\text{Retrieve}(P, a, U_1)$ .

The advantage probability of  $\mathcal{A}$  in winning the game  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{unlink}}$  is defined as

$$\text{Adv}_{\mathcal{A}, \text{PMS}}^{\text{unlink}}(\kappa) := \left| \Pr \left[ \text{Game}_{\mathcal{A}, \text{PMS}}^{\text{unlink}} = 1 \right] - \frac{1}{2} \right|$$

We say that PMS provides unlinkability if for any PPT adversary  $\mathcal{A}$  the advantage  $\text{Adv}_{\mathcal{A}, \text{PMS}}^{\text{unlink}}(\kappa)$  is negligible.

---

## 4 Private Profiles with Shared Keys

Our first construction, called PMS-SK, is simple and uses *shared keys* to encrypt profile attributes for a group of authorized users. An independent symmetric key  $K_a$  is chosen by the owner of a profile  $P$  for each pair  $(a, d)$  and distributed to the group  $\mathcal{G}$  of users that are authorized to access  $d$ . The key is updated on each modification of  $\mathcal{G}$ . We use a *symmetric encryption scheme*  $SE = (SE.KGen, SE.Enc, SE.Dec)$  (cf. for example [17, p. 60]) for which we assume classical indistinguishability against chosen-plaintext attacks (IND-CPA) and denote by  $\text{Adv}_{\mathcal{A}, SE}^{\text{IND-CPA}}(\kappa)$  the corresponding advantage of the adversary (cf. for example [17, p. 82]).

The distribution of  $K_a$  may be performed in two ways:  $K_a$  can be communicated to the authorized users online (over secure channels) or offline, e.g., by storing  $K_a$  securely (possibly using asymmetric encryption) either within the profile or at some centralized server. Our specification of PMS-SK leaves open how distribution of shared keys is done. In particular, the use of one or another technique may be constrained by the application that will use the scheme.

---

### 4.1 Specification of PMS-SK

---

The construction of PMS-SK is as follows.

**Init**( $\kappa$ ): Output  $P \leftarrow \emptyset$  and  $pmk \leftarrow \emptyset$ .

**Publish**( $pmk, P, (a, d), \mathcal{G}$ ):  $K_a \leftarrow SE.KGen(1^\kappa)$ , add  $(a, SE.Enc(K_a, d))$  to  $P$ ,  $K_a$  to  $rk_U$  for each  $U \in \mathcal{G}$ , and  $K_a$  to  $pmk$ .

**Retrieve**( $rk_U, P, a$ ): Extract  $K_a$  from  $rk_U$ . If  $(a, \bar{d}) \in P$  for some  $\bar{d}$  then output  $SE.Dec(K_a, \bar{d})$ , else  $\perp$ .

**Delete**( $pmk, P, a$ ): Delete  $(a, \bar{d})$  from  $P$ . Delete  $K_a$  from  $pmk$ .

**ModifyAccess**( $pmk, P, a, U$ ): If  $U \in \mathcal{G}$  then remove  $U$  from  $\mathcal{G}$ , otherwise add  $U$  to  $\mathcal{G}$ . Execute **Delete**( $pmk, P, a$ ) followed by **Publish**( $pmk, P, (a, d), \mathcal{G}$ ) where  $d$  is the attribute indexed by  $a$ .

The description of **ModifyAccess** is kept general in the sense that it does not specify how the profile owner  $U_p$  reveals an attribute  $d$  indexed by  $a$ . Our scheme allows for different realizations:  $d$  can be stored by  $U_p$  locally (not as part of  $P$ ) or it can be obtained through decryption of  $\bar{d}$  using  $K_a$  which is part of  $pmk$ .

We assume in our constructions that the uniqueness of indices  $a$  in a profile  $P$  is implicitly ensured or checked by corresponding algorithms.

---

### 4.2 Complexity Analysis

---

PMS-SK requires each profile owner  $U_p$  to store one key per attribute  $(a, \bar{d})$  currently published in  $P$ . Additionally, each user has to store one key per attribute she is allowed to access in any profile. Therefore, assuming the worst case where all users in  $\mathcal{U}$  have profiles containing  $|P|$  attributes that can be accessed by all other users, PMS-SK requires each  $U \in \mathcal{U}$  to store  $N \cdot |P|$  keys from which  $|P|$  keys are stored in its  $pmk$  and  $(N - 1) \cdot |P|$  in the retrieval keys  $rk_U$  for all other users' profiles. For each **Publish** or **ModifyAccess** operation the profile owner needs further to perform one symmetric encryption.

---

### 4.3 Security and Privacy Analysis

---

In this section we prove that PMS-SK ensures confidentiality of attributes and provides unlinkability for the authorized users.

**Theorem 1** (Confidentiality of PMS-SK). *If  $SE$  is IND-CPA secure, then PMS-SK provides confidentiality from Definition 2, and*

$$\text{Adv}_{\mathcal{A}, \text{PMS-SK}}^{\text{conf}}(\kappa) \leq (1 + q) \cdot \text{Adv}_{\mathcal{A}^*, SE}^{\text{IND-CPA}}(\kappa)$$

with  $q$  being the number of invoked **ModifyAccess** operations per attribute.

*Proof.* Assume a PPT adversary  $\mathcal{A}$  against the confidentiality of PMS-SK. We construct a PPT adversary  $\mathcal{A}^*$  against IND-CPA security of  $SE$  which simulates the execution of PMS-SK operations and interacts with  $\mathcal{A}$  as specified in  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{conf}}$ :

After initializing PMS-SK for all users in  $\mathcal{U}$ ,  $\mathcal{A}^*$  responds to the queries issued by  $\mathcal{A}$ , acting on behalf of the respective profile owners as specified for the scheme (e.g., choosing attribute keys  $K_a$ , changing profiles, etc.).

At some point in time,  $\mathcal{A}$  outputs its challenge  $(a, d_0), (a, d_1), \mathcal{G}_t, U_P$ .  $\mathcal{A}^*$  then forwards  $(m_0, m_1) = (d_0, d_1)$  as its own IND-CPA challenge and obtains the ciphertext  $c$  (for  $m_b$ ).  $\mathcal{A}^*$  picks  $i \in [1, q + 1]$  at random. If  $i = 1$  then  $\mathcal{A}^*$  adds  $(a, c)$  to  $P$  as part of corresponding Publish operation and outputs  $P$  to  $\mathcal{A}$ . Otherwise,  $(a, c)$  is added to  $P$  by  $\mathcal{A}^*$  as part of the  $i^{\text{th}}$  ModifyAccess operation on  $a$ . In all other operations on  $a$  the IND-CPA adversary  $\mathcal{A}^*$  proceeds according to the specification of PMS-SK by choosing corresponding keys  $K_a$  on its own.

$\mathcal{A}^*$  continues responding to the queries issued by  $\mathcal{A}$ . At some point in time,  $\mathcal{A}$  outputs a bit  $b'$ , which  $\mathcal{A}^*$  forwards as its own bit  $b'$  to the IND-CPA challenger. Assuming that  $\mathcal{A}$  breaks the confidentiality of PMS-SK the bit  $b'$  forwarded by  $\mathcal{A}^*$  is equal to the bit  $b$  chosen by the IND-CPA challenger with probability  $\frac{1}{1+q}$ , which is due to the independence of keys  $K_a$  chosen by  $\mathcal{A}^*$  in all but the  $i^{\text{th}}$  operation on  $a$  and the successful guess of  $\mathcal{A}^*$  with regard to  $i$ . In this way we obtain the desired upper-bound for  $\text{Adv}_{\mathcal{A}, \text{PMS-SK}}^{\text{conf}}(\kappa)$  which is negligible assuming the IND-CPA security of  $SE$ .  $\square$

**Theorem 2** (Unlinkability of PMS-SK). *PMS-SK provides perfect unlinkability as defined in Definition 3, i.e.,  $\text{Adv}_{\mathcal{A}, \text{PMS-SK}}^{\text{unlink}}(\kappa) = 0$ .*

*Proof.* The attribute keys  $K_a$  are statistically independent of the identities of users in  $\mathcal{G}$  who have been authorized to access the attribute indexed by  $a$ . Therefore,  $\mathcal{A}$  cannot win in  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{unlink}}$  better than by a random guess, i.e., with probability  $\frac{1}{2}$ .  $\square$

**Remark 1.** *The perfect unlinkability property of our PMS-SK construction proven in the above theorem should be enjoyed with caution when it comes to the deployment of the scheme in practice. The reason is that PMS-SK does not specify how shared keys are distributed, leaving this to the application that will use the scheme. One approach to distribute keys in a privacy-preserving manner is given by Barth, Boneh, and Waters [2] and the CCA recipient privacy of their scheme, which however comes with storage overhead linear in the number of recipients and may be undesirable when encrypting small-sized attributes in user profiles. In any case it is clear that the distribution process will eventually have impact on the unlinkability property of the scheme, maybe to the point of ruling out its perfectness.*

---

## 4.4 Further Optimizations

---

Regardless of the question, whether shared keys  $K_a$  are distributed by the application in an online or an offline fashion, there is a way to further optimize and improve the actual management of these keys. In our specification of PMS-SK, these keys are currently chosen fresh for each modification of the authorized group  $\mathcal{G}$ . However, by using *group key management schemes* that allow efficient update of group keys such as LKH [26, 25] or OFT [6, 22, 19] with all the resulting efficiency differences, the overhead for the distribution can be further reduced.

Another optimization concerns generation of shared keys  $K_a$  in case a profile owner  $U_P$  does not wish to store corresponding attributes  $d$  (outside of the profile). Instead of storing linear (in the number of attributes in  $P$ ) many shared keys in  $\text{pmk}$ , the profile owner can derive each  $K_a$  using some pseudorandom function [13]  $f_s(a, i)$  where  $s$  is a seed used for all attributes,  $a$  is the unique attribute index, and  $i$  is a counter that is updated on each execution of ModifyAccess on  $a$  to account for possible repetitions of the authorized group  $\mathcal{G}$  over the life time of the profile. This optimization allows to trade in the storage costs for  $\text{pmk}$  for the computation overhead for deriving  $K_a$ .

We do not analyze the efficiency effects of the proposed optimizations in detail here, as the construction based on broadcast encryption presented in the next section has only a constant overhead of retrieval keys.

## 5 Private Profiles with Broadcast Encryption

Our second generic construction of a profile management scheme, called PMS-BE, is based on an adaptively secure (identity-based) broadcast encryption scheme, e.g. [10], whose syntax and requirements we recall in the following.

**Definition 4** (Broadcast Encryption Scheme [10]). A broadcast encryption scheme  $BE = (BE.Setup, BE.KGen, BE.Enc, BE.Dec)$  consists of the following algorithms:

$BE.Setup(\kappa, n, \ell)$ : On input the security parameter  $\kappa$ , the number of receivers  $n$ , and the maximal size  $\ell \leq n$  of the recipient group, this probabilistic algorithm outputs a public/secret key pair  $\langle PK, SK \rangle$ .

$BE.KGen(i, SK)$ : On input an index  $i \in \{1, \dots, n\}$  and the secret key  $SK$ , this probabilistic algorithm outputs a private (user) key  $sk_i$ .

$BE.Enc(S, PK)$ : On input a subset  $S \subseteq \{1, \dots, n\}$  with  $|S| \leq \ell$  and a public key  $PK$ , this probabilistic algorithm outputs a pair  $\langle Hdr, K \rangle$  where  $Hdr$  is called the header and  $K \in \mathcal{K}$  is a message encryption key.

$BE.Dec(S, i, sk_i, Hdr, PK)$ : On input a subset  $S \subseteq \{1, \dots, n\}$  with  $|S| \leq \ell$ , an index  $i \in \{1, \dots, n\}$ , a private key  $sk_i$ , a header  $Hdr$ , and the public key  $PK$ , this deterministic algorithm outputs the message encryption key  $K \in \mathcal{K}$ .

Correctness of  $BE$  requires that for all  $S \subseteq \{1, \dots, n\}$  and all  $i \in S$ , if  $\langle PK, SK \rangle \leftarrow_R BE.Setup(\kappa, n, \ell)$ ,  $sk_i \leftarrow_R BE.KGen(i, SK)$ , and  $\langle Hdr, K \rangle \leftarrow_R BE.Enc(S, PK)$ , then  $BE.Dec(S, i, sk_i, Hdr, PK) = K$ .

The adaptive security of  $BE$  against chosen plaintext attacks as defined in [10] can be extended to chosen-ciphertext attacks as follows.

**Definition 5** (Adaptive CCA-Security of  $BE$ ). Let  $BE$  be a broadcast encryption scheme from Definition 4 and  $\mathcal{A}$  be a PPT adversary in the following game, denoted  $\text{Game}_{\mathcal{A}, BE, n, \ell}^{\text{ad-CCA}}(\kappa)$ :

1.  $\langle PK, SK \rangle \leftarrow_R BE.Setup(\kappa, n, \ell)$ .  $\mathcal{A}$  is given  $PK$  (together with  $n$  and  $\ell$ ).
2.  $\mathcal{A}$  adaptively issues private key queries  $BE.KGen(i)$  for  $i \in \{1, \dots, n\}$  and obtains corresponding  $sk_i$ . In addition,  $\mathcal{A}$  is allowed to query  $BE.Dec(S, i, Hdr, PK)$  to obtain message encryption keys  $K$ .
3.  $\mathcal{A}$  outputs a challenge set of indices  $S^*$ , such that no  $BE.KGen(i)$  with  $i \in S^*$  was asked. Let  $\langle Hdr^*, K_0 \rangle \leftarrow_R BE.Enc(S^*, PK)$  and  $K_1 \in_R \mathcal{K}$ . A bit  $b \in_R \{0, 1\}$  is chosen uniformly and  $\mathcal{A}$  is given  $(Hdr^*, K^*)$  with  $K^* = K_b$ .
4.  $\mathcal{A}$  is allowed to query  $BE.Dec(S, i, Hdr, PK)$ , except on inputs of the form  $\langle S^*, i, Hdr^*, PK \rangle$ ,  $i \in S^*$ .
5.  $\mathcal{A}$  outputs a bit  $b' \in \{0, 1\}$  and wins the game, denoted  $\text{Game}_{\mathcal{A}, BE, n, \ell}^{\text{ad-CCA}}(\kappa) = 1$ , if  $b' = b$ .

We define  $\mathcal{A}$ 's advantage against the adaptive CCA-security of  $BE$  as

$$\text{Adv}_{\mathcal{A}, BE, n, \ell}^{\text{ad-CCA}}(\kappa) = \left| \Pr \left[ \text{Game}_{\mathcal{A}, BE, n, \ell}^{\text{ad-CCA}}(\kappa) = 1 \right] - \frac{1}{2} \right|$$

We say that  $BE$  is adaptively CCA-secure if for all PPT adversaries  $\mathcal{A}$  the advantage  $\text{Adv}_{\mathcal{A}, BE, n, \ell}^{\text{ad-CCA}}(\kappa)$  is negligible.

### 5.1 Specification of PMS-BE

The main idea behind PMS-BE is that each profile owner  $U_p$  manages independently its own instance of the  $BE$  scheme in the following way:  $U_p$  assigns fresh indices  $i$ , which we call *pseudonyms*, to the users from  $\mathcal{U}$  (upon their first admission to  $P$ ) and gives them corresponding private (user) keys  $sk_i$ . In order to publish an attribute  $d$  for some authorized group  $\mathcal{G}$ , the owner encrypts  $d$  using the  $BE$  scheme and the set of indices assigned to the users in  $\mathcal{G}$ . This process allows for very efficient modification of the authorized group  $\mathcal{G}$ : In order to admit or remove a member  $U$  with regard to  $d$  the profile owner simply adjusts  $\mathcal{G}$  and re-encrypts  $d$ . In particular, there is no need to distribute new decryption keys. However, this flexibility comes at the price of a somewhat weaker privacy, since  $BE$  schemes include indices  $i$  into ciphertext headers, which in turn allows for linkability of an authorized user  $U$  across multiple attributes within  $P$ . Yet, the use of pseudonyms still allows us to show that PMS-BE satisfies the weaker requirement of *anonymity*, which we introduce and formally relate to unlinkability in Section 5.4.

**Init**( $\kappa$ ): Execute  $\langle PK, SK \rangle \leftarrow BE.Setup(\kappa, n, \ell)$  with  $n = \ell = N$ <sup>1</sup>. Output  $P \leftarrow \emptyset$  and  $pmk \leftarrow \{PK, SK\}$ . Additionally,  $PK$  is made public.

**Publish**( $pmk, P, (a, d), \mathcal{G}$ ): For every  $U \in \mathcal{G}$  without pseudonym for  $P$  pick an unused pseudonym  $i$  at random from  $[1, n]$ , extract  $sk_i \leftarrow BE.KGen(i, SK)$ , and define new  $rk_U \leftarrow \langle i, sk_i \rangle$ . For every  $U \in \mathcal{G}$  add the corresponding pseudonyms to the set  $S$ . Compute  $\langle Hdr, K_a \rangle \leftarrow BE.Enc(S, PK)$ ,  $\hat{d} \leftarrow SE.Enc(K_a, d)$ , and  $\bar{d} \leftarrow \langle Hdr, S, \hat{d} \rangle$ . Add  $(a, \bar{d})$  to  $P$  and  $K_a$  to  $pmk$ . Output  $P$ , all new  $rk_U$ , and  $pmk$ .

**Retrieve**( $rk_U, P, a$ ): Extract  $(a, \bar{d})$  from  $P$ . Parse  $\bar{d}$  as  $\langle Hdr, S, \hat{d} \rangle$ . Extract  $\langle i, sk_i \rangle$  from  $rk_U$ . Set  $K_a \leftarrow BE.Dec(S, i, sk_i, Hdr, PK)$  and output  $SE.Dec(K_a, \hat{d})$ .

**Delete**( $pmk, P, a$ ): Delete  $(a, \bar{d})$  from  $P$ . Delete  $K_a$  from  $pmk$ .

**ModifyAccess**( $pmk, P, a, U$ ): If  $U \in \mathcal{G}$  remove  $U$  from  $\mathcal{G}$ ; otherwise add  $U$  to  $\mathcal{G}$ . Execute **Delete**( $pmk, P, a$ ) followed by **Publish**( $pmk, P, (a, d), \mathcal{G}$ ), where  $d$  is the attribute indexed by  $a$ .

---

## 5.2 Complexity Analysis

---

PMS-BE requires each profile owner  $U_p$  to store one key per index-attribute pair  $(a, \bar{d})$  currently published in  $P$  as well as the key pair  $\langle PK, SK \rangle$ . For each profile containing at least one attribute a user is allowed to access, this user has to store its secret key  $\langle i, sk_i \rangle$  contained in  $rk_U$ . Assuming the worst case where all users in  $\mathcal{U}$  have profiles containing  $|P|$  attributes that can be accessed by all other users, PMS-BE requires each  $U \in \mathcal{U}$  to store  $|P| + N + 1$  keys from which  $|P| + 2$  keys are stored in  $pmk$  and  $N - 1$  secret keys  $\langle i, sk_i \rangle$  are stored in the retrieval keys  $rk_U$  of all others users' profiles. For each **Publish** or **ModifyAccess** operation the profile owner needs further to perform one broadcast encryption  $BE.Enc$  and one symmetric encryption.

The storage overhead may be reduced by omitting the storage of attribute keys  $K_a$  in  $pmk$  as the profile owner is able to reconstruct  $K_a$  by executing  $sk_i \leftarrow BE.KGen(i, SK)$  for any index  $i$  in the set of authorized indices  $S$  for  $a$ . With the authorized user's secret key  $sk_i$ , the profile owner is able to execute  $BE.Dec$ , receiving  $K_a$ . That way, the total number of stored keys is reduced by  $|P|$  to  $N + 1$ , traded in for a higher computation overhead when executing **ModifyAccess**.

Obviously, the main advantage of the PMS-BE construction over the PMS-SK approach is the number of keys that have to be stored in  $rk_U$ , which is only one in PMS-BE whereas PMS-SK imposes key overhead linear in the number of attributes. However, this efficiency benefit comes at the cost of a weaker privacy, as we discuss below.

---

## 5.3 Confidentiality of PMS-BE

---

We first analyze the confidentiality property of the PMS-BE scheme.

**Theorem 3** (Confidentiality of PMS-BE). *If  $SE$  is IND-CPA secure and  $BE$  is adaptively CCA-secure, PMS-BE provides confidentiality from Definition 2, and*

$$\text{Adv}_{\mathcal{A}, \text{PMS-BE}}^{\text{conf}}(\kappa) \leq (1 + q) \cdot \left( \text{Adv}_{\mathcal{B}_1, SE}^{\text{IND-CPA}}(\kappa) + N \cdot \text{Adv}_{\mathcal{B}_2, BE, n, \ell}^{\text{ad-CCA}}(\kappa) \right)$$

with  $q$  being the number of invoked **ModifyAccess** operations per attribute.

*Proof.* Assume an PPT adversary  $\mathcal{A}$  against the confidentiality of PMS-BE. We show how to construct an adversary  $\mathcal{B}_1$  against the IND-CPA security of  $SE$  and an adversary  $\mathcal{B}_2$  against adaptive CCA-security of  $BE$ .

*Construction of  $\mathcal{B}_1$ .* This is similar to the adversary  $\mathcal{A}^*$  against IND-CPA of  $SE$  in the proof of Theorem 1. While responding to queries issued by  $\mathcal{A}$  in  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{conf}}$  on behalf of the respective profile owners (by executing all operations honestly as specified for PMS-BE),  $\mathcal{B}_1$  upon receiving the challenge  $((a, d_0), (a, d_1), \mathcal{G}_t, U_p)$  from  $\mathcal{A}$ , forwards  $(m_0, m_1) = (d_0, d_1)$  as its own challenge to the IND-CPA challenger and obtains a ciphertext  $c$ . It then picks  $j \in [1, q + 1]$  at random, and adds  $(a, \bar{d})$  as  $\langle Hdr, S, c \rangle$  to the challenge profile  $P$  either as part of the **Publish** (if  $j = 1$ ) or as part of the  $j^{\text{th}}$  **ModifyAccess** operation on  $a$  (if  $j \neq 1$ ). At the end,  $\mathcal{B}_1$  forwards the output bit of  $\mathcal{A}$  as its own bit in the IND-CPA game.

*Construction of  $\mathcal{B}_2$ .*  $\mathcal{B}_2$  receives  $PK$  as part of  $\text{Game}_{\mathcal{B}_2, BE, n, \ell}^{\text{ad-CCA}}(\kappa)$ .  $\mathcal{B}_2$  now picks a user  $U_p \in \mathcal{U}$  at random, i.e., it guesses the profile owner that  $\mathcal{A}$  will use as target in its attack. For all other users  $\mathcal{B}_2$  generates all PMS-BE parameters and performs operations on their behalf honestly by itself. Only for the profile  $P$  of selected  $U_p$  it responds to the queries of  $\mathcal{A}$  using own oracle access to the underlying  $BE$  scheme, i.e.,

---

<sup>1</sup> We use this upper bound for simplicity here. One may cut down both on  $n$  and  $\ell$  to improve the efficiency of  $BE$ .



- $\mathcal{B}_2$  answers  $\text{Publish}(P, (a, d), \mathcal{G})$  queries using its own  $BE.Enc(S, PK)$  oracle with the set of pseudonyms  $S$  which correspond to the users in  $\mathcal{G}$ . From the obtained  $(Hdr, K_a)$  it then uses  $K_a$  to encrypt  $d$  as specified in  $\text{Publish}$  operation. Note that the assignment of pseudonyms  $i$  to users  $U$  is done by  $\mathcal{B}_2$  and that  $\mathcal{B}_2$  does not know the private keys  $sk_i$ .
- $\mathcal{B}_2$  answers  $\text{Retrieve}(P, a, U)$  queries using its own oracle  $BE.Dec(S, i, Hdr, PK)$  oracle where  $i$  is the pseudonym of  $U$  within the profile  $P$ , and  $S$  and  $Hdr$  is part of the ciphertext  $\bar{d}$  indexed with  $a$ .
- $\mathcal{B}_2$  answers  $\text{Corrupt}(U)$  queries by identifying all indices  $i$  that  $U$  has in all other profiles  $P' \neq P$  and returning corresponding  $sk_i$  to  $\mathcal{A}$ . For the profile  $P$  it queries the oracle  $BE.KGen(i)$  to obtain  $sk_i$  of  $U$  in that profile and gives  $sk_i$  to  $\mathcal{A}$ . In addition it gives  $\mathcal{A}$  all retrieval keys  $rk_U$  of  $U$ .
- $\mathcal{B}_2$  answers  $\text{ModifyAccess}(P, a, U)$  queries similarly to  $\text{Publish}$  queries above using its  $BE.Enc(S, PK)$  oracle.

At some point in time,  $\mathcal{A}$  outputs  $(a, d_0), (a, d_1), \mathcal{G}_t, U_p$ .  $U_p$  matches the target user guessed by  $\mathcal{B}_2$  with probability  $\frac{1}{N}$ . Then,  $\mathcal{B}_2$  assigns to each  $U \in \mathcal{G}_t$  an unused pseudonym  $i$  (unless  $U$  already has a pseudonym for  $P$ ). Next,  $\mathcal{B}_2$  picks  $j \in [1, q + 1]$  at random.

If  $j = 1$ ,  $\mathcal{B}_2$  queries its  $BE.KGen(i)$  oracle with all so-far unassigned pseudonyms  $i$  to obtain private keys  $sk_i$ . (Note that  $\mathcal{B}_2$  does not obtain private keys of users in  $\mathcal{G}_t$  but can still answer  $\text{Corrupt}$  queries of  $\mathcal{A}$  regarding other users later.) Finally,  $\mathcal{B}_2$  outputs set  $S^*$  containing pseudonyms of users in  $\mathcal{G}_t$  as part of  $\text{Game}_{\mathcal{B}_2, BE, n, \ell}^{\text{ad-CCA}}(\kappa)$ , receives the challenge  $(Hdr^*, K^*)$ , and uses  $K^*$  to encrypt the attribute  $d_b$  for a random bit  $b$  of its choice, i.e.  $\mathcal{B}_2$  adds  $(a, \bar{d})$  where  $\bar{d} = \langle Hdr^*, S^*, SE.Enc(K^*, d_b) \rangle$  to  $P$ .

However, if  $j \neq 1$  then  $\mathcal{B}_2$  replies to the challenge by computing  $(a, \bar{d})$  still with the help of its  $BE.Enc$  oracle. Later, in response to the  $j^{\text{th}}$   $\text{ModifyAccess}(P, a, U)$  query of  $\mathcal{A}$  that leads to the  $j^{\text{th}}$  update of the target group  $\mathcal{G}_t$ , it queries its  $BE.KGen(i)$  oracle with all so-far unassigned pseudonyms  $i$  to obtain private keys  $sk_i$  and outputs set  $S^*$  containing pseudonyms of users in the updated  $\mathcal{G}_t$  as part of  $\text{Game}_{\mathcal{B}_2, BE, n, \ell}^{\text{ad-CCA}}(\kappa)$  to receive the challenge  $(Hdr^*, K^*)$ , and adds  $(a, \bar{d})$  where  $\bar{d} = \langle Hdr^*, S^*, SE.Enc(K^*, d_b) \rangle$  to  $P$ . Note that the updated group  $\mathcal{G}_t$  never included corrupted users.

In both cases  $\mathcal{B}_2$  continues answering queries of  $\mathcal{A}$  as described in the beginning until  $\mathcal{A}$  outputs its bit  $b'$ . If  $b' = b$ , meaning that  $\mathcal{A}$  was successful, then  $\mathcal{B}_2$  outputs 0 in  $\text{Game}_{\mathcal{B}_2, BE, n, \ell}^{\text{ad-CCA}}(\kappa)$ , indicating that the received key  $K^*$  was real; otherwise  $\mathcal{B}_2$  outputs 1. If  $\mathcal{A}$  wins then  $\mathcal{B}_2$  wins with probability  $\frac{1}{N(1+q)}$  to account for the correct guess of the profile owner  $U_p$  and  $j$ .

By combining the advantages of  $\mathcal{B}_1$  and  $\mathcal{B}_2$  that both use the confidentiality adversary  $\mathcal{A}$  in their respective games we can give an upper-bound  $\text{Adv}_{\mathcal{A}, \text{PMS-BE}}^{\text{conf}}(\kappa)$ , which is negligible assuming the security of  $SE$  and  $BE$ .  $\square$

## 5.4 Privacy of PMS-BE

Our PMS-BE construction does *not* provide unlinkability as defined in Definition 3 since the indices of users are linkable across different published attributes. Consider the following construction of a PPT adversary  $\mathcal{A}$  against the unlinkability of PMS-BE:

After initialization of PMS-BE,  $\mathcal{A}$  outputs two arbitrary users  $U_0, U_1$ , some pair  $(a, d)$  and a profile owner  $U_p$ . Then,  $\mathcal{A}$  executes  $\text{Publish}(P, (a', d'), \{U_0\})$  for an arbitrary pair  $(a', d')$  and extracts the two pairs  $(a, \bar{d})$  and  $(a', \bar{d}')$  from  $P$  (parsing  $\bar{d}$  as  $\langle Hdr, S, \hat{d} \rangle$  and  $\bar{d}'$  as  $\langle Hdr', S', \hat{d}' \rangle$ ). If  $S = S'$ ,  $\mathcal{A}$  outputs 0, otherwise 1.

$\mathcal{A}$  always wins, thus

$$\text{Adv}_{\mathcal{A}, \text{PMS-BE}}^{\text{unlink}}(\kappa) = \left| \Pr \left[ \text{Game}_{\mathcal{A}, \text{PMS-BE}}^{\text{unlink}} = 1 \right] - \frac{1}{2} \right| = \frac{1}{2}$$

which is non-negligible.

Since PMS-BE has simpler management and distribution of retrieval keys it would be nice to see whether it can satisfy some weaker, yet still meaningful privacy property. It turns out that PMS-BE is still able to provide *anonymity* of users that are members of different authorized groups  $\mathcal{G}$  within the same profile, even in the presence of an adversary in these groups.

We formalize anonymity in Definition 6 based on the following intuition: An adversary shall not be able to decide the identity of some user  $U_b \in \{U_0, U_1\}$  in the setting where the adversary is restricted to publish attributes or modify access to them either by simultaneously including both  $U_0$  and  $U_1$  into the authorized group or none of them. This definition rules out linkability of users based on their pseudonyms, while keeping all other privacy properties of the unlinkability definition.

**Definition 6** (Anonymity). *Let PMS be a profile management scheme from Definition 1 and  $\mathcal{A}$  be a PPT adversary interacting with users via queries from Section 3.2 within the following game  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{anon}}$ :*

1.  $\text{Init}(\kappa)$  is executed for all users  $U \in \mathcal{U}$ .
2.  $\mathcal{A}$  can execute arbitrary operations and ask queries. At some point it outputs  $U_0, U_1, (a, d)$ , and  $U_p$  (owner of some profile  $P$ ).
3. A bit  $b \in_{\mathcal{R}} \{0, 1\}$  is chosen uniformly and:
  - If  $(a, *) \notin P$  then  $\text{Publish}(pmk, P, (a, d), \{U_b\})$  with  $pmk$  of  $U_p$  is executed.
  - If  $(a, *) \in P$  then  $\text{ModifyAccess}(pmk, P, a, \{U_b\})$  with  $pmk$  of  $U_p$  is executed.

In both cases  $\mathcal{A}$  is given the modified profile  $P$  and the possibly updated retrieval keys  $rk_U$  for all  $U \in \mathcal{C}$ .
4.  $\mathcal{A}$  can execute arbitrary operations and ask queries. At some point it outputs a bit  $b' \in \{0, 1\}$ .
5.  $\mathcal{A}$  wins, denoted by  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{anon}} = 1$ , if all of the following holds:
  - $b' = b$ .
  - $\{U_0, U_1, U_p\} \cap \mathcal{C} = \emptyset$ .
  - $\mathcal{A}$  neither queried  $\text{Retrieve}(P, a, U_0)$  nor  $\text{Retrieve}(P, a, U_1)$ .
  - $\text{Publish}(P, (\hat{a}, *), \mathcal{G})$  has not been executed with  $U_0 \in \mathcal{G}$  but  $U_1 \notin \mathcal{G}$  or with  $U_1 \in \mathcal{G}$  but  $U_0 \notin \mathcal{G}$  for any  $\hat{a}$ .
  - $\text{ModifyAccess}(P, \hat{a}, U)$  has not been executed with  $U = U_0$  or  $U = U_1$  for any  $\hat{a}$ .

The advantage probability of  $\mathcal{A}$  in winning the game  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{anon}}$  is defined as

$$\text{Adv}_{\mathcal{A}, \text{PMS}}^{\text{anon}}(\kappa) := \left| \Pr \left[ \text{Game}_{\mathcal{A}, \text{PMS}}^{\text{anon}} = 1 \right] - \frac{1}{2} \right|$$

We say that PMS provides anonymity if for any PPT adversary  $\mathcal{A}$  the advantage  $\text{Adv}_{\mathcal{A}, \text{PMS}}^{\text{anon}}(\kappa)$  is negligible.

We prove that PMS-BE provides perfect anonymity using similar arguments as we used for the perfect unlinkability of the PMS-SK scheme. Nevertheless, we observe that our discussion in Remark 1 regarding the potential loss of perfectness for the unlinkability of PMS-SK when deployed in the concrete application applies to PMS-BE as well, due to the distribution of private user keys  $sk_i$ .

**Theorem 4** (Anonymity of PMS-BE). *The PMS-BE scheme provides perfect anonymity as defined in Definition 6, i.e.,  $\text{Adv}_{\mathcal{A}, \text{PMS-BE}}^{\text{anon}}(\kappa) = 0$ .*

*Proof.* The pseudonym  $i$  of a user  $U \in \mathcal{U}$  for a profile  $P$  is chosen uniformly by the owner  $U_p$  from  $[1, N]$  and remains independent of the user identity  $U$  and other pseudonyms that  $U$  may have in other profiles. As long as the adversary  $\mathcal{A}$  does not corrupt  $U_0, U_1$ , or the profile owner  $U_p$  (which is prohibited by Definition 6) and since the bit  $b$  chosen in  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{anon}}$  is statistically independent of the identities  $U_0$  and  $U_1$ ,  $\mathcal{A}$  cannot infer any information about  $U_b$  by invoking publishing and modification operations due to the final two restrictions of  $\text{Game}_{\mathcal{A}, \text{PMS}}^{\text{anon}}$ . Therefore,  $\mathcal{A}$  cannot break the anonymity of PMS-BE better than by a random guess, i.e.,  $\Pr \left[ \text{Game}_{\mathcal{A}, \text{PMS}}^{\text{anon}} = 1 \right] = \frac{1}{2}$ .  $\square$

We complete our security analysis by showing that the unlinkability requirement for a profile management scheme is strictly stronger than the anonymity requirement. Our separation result holds unconditionally in that it preserves the properties of the starting scheme.

**Theorem 5** (Unlinkability  $\Rightarrow$  Anonymity). *Let PMS be a profile management scheme from Definition 1 providing unlinkability from Definition 3. Then PMS also provides anonymity from Definition 6.*

*Proof.* The construction of an unlinkability adversary  $\mathcal{A}^*$  that is given black-box access to an anonymity adversary  $\mathcal{A}$  is straightforward.  $\mathcal{A}^*$  answers all queries of  $\mathcal{A}$  by relaying them as queries to its own challenger and outputs whatever  $\mathcal{A}$  outputs. If  $\mathcal{A}$  wins then so does  $\mathcal{A}^*$ , i.e.,  $\text{Adv}_{\mathcal{A}, \text{PMS}}^{\text{unlink}}(\kappa) \geq \text{Adv}_{\mathcal{A}, \text{PMS}}^{\text{anon}}(\kappa)$ .  $\square$

**Theorem 6** (Unlinkability  $\not\Leftarrow$  Anonymity). *Let PMS be a profile management scheme from Definition 1. Then there exists a profile management scheme  $\text{PMS}'$  such that for any adversary  $\mathcal{A}'$  against anonymity of  $\text{PMS}'$  there exists an adversary  $\mathcal{A}$  against anonymity of PMS. Furthermore, there exists an adversary  $\mathcal{B}$  against the unlinkability of  $\text{PMS}'$ .*

*Proof.* Take any PMS scheme which provides anonymity. PMS can be modified to  $\text{PMS}'$  as follows:

- 
- Use identic algorithms  $\text{PMS}'.\text{Init} \equiv \text{PMS}.\text{Init}$  and  $\text{PMS}'.\text{Delete} \equiv \text{PMS}.\text{Delete}$ .
  - The algorithms  $\text{PMS}'.\text{Publish}$  and  $\text{PMS}.\text{Publish}$  as well as  $\text{PMS}'.\text{ModifyAccess}$  and  $\text{PMS}.\text{ModifyAccess}$  differ in that  $\text{PMS}'$  operations update a profile  $P$  with pairs of the form  $(a, \bar{d}') \in P$  with  $\bar{d}' = \bar{d} \parallel S$ , where  $a$  and  $\bar{d}$  are computed using corresponding PMS operations, and the appended set  $S$  contains profile-specific indices  $i$  of users  $U$  from the authorized group  $\mathcal{G}$  (for the attribute  $d$ ), such that unique  $i$  is chosen at random from  $[1, N]$  within these operations for each  $U \in \mathcal{U}$  upon the first admission of  $U$  to  $P$ .
  - $\text{PMS}'.\text{Retrieve}$  and  $\text{PMS}.\text{Retrieve}$  differ in that  $\text{PMS}'.\text{Retrieve}$  first computes  $\bar{d}$  by removing  $S$  from  $\bar{d}'$  and then invoking  $\text{PMS}.\text{Retrieve}$  on the corresponding pair  $(a, \bar{d})$ .

The  $\text{PMS}'$  scheme remains anonymous, as the additional indices in  $S$  do not leak any information about the corresponding identities of users  $\mathcal{G}$ . The  $\text{PMS}'$  scheme clearly does not provide unlinkability since the indices of users are linkable across different published attributes within the profile.  $\square$

## 6 Taking the Model to Reality

In the following, we will provide a short idea of how the presented constructions of profile management schemes behave in settings of different online social networks. Therefore, we first recall and adjust the general key overhead of both presented schemes and subsequently discuss the overhead in specific settings.

### 6.1 Key Overhead in general

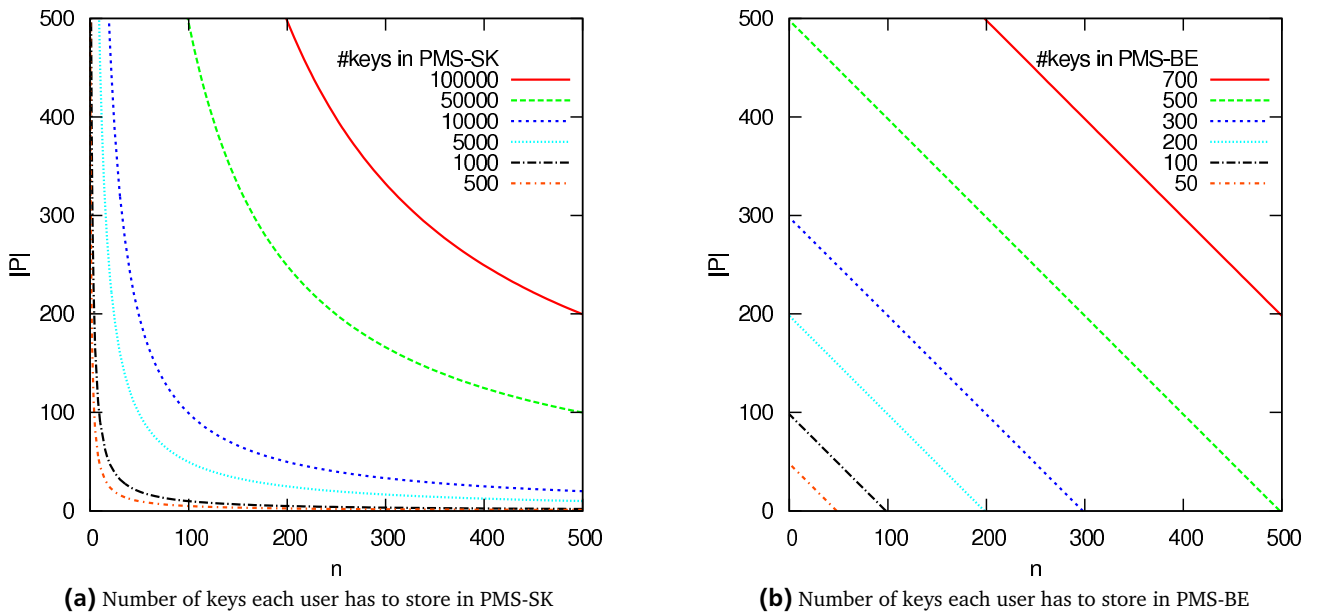
Recalling the complexity analyses of PMS-SK in Section 4.2 and PMS-BE in Section 5.2, each user has to store  $N \cdot |P|$  keys in the PMS-SK construction, whereas in the PMS-BE construction, each user has to store  $N + |P| + 1$  keys ( $N$  being the total number of users and  $|P|$  the number of attributes stored in each profile), assuming the worst case, where each user stores  $|P|$  attributes in her profile and allows every other user in the system to access all attributes.

Obviously, this worst case is very unlikely in real settings. Both constructions however behave quite similar in a setting where each user has on average  $n$  contacts and shares on average  $|P|$  attributes with all of them:

- PMS-SK requires each profile owner to store one key per attribute currently published in her profile and each user to store one (retrieval) key per attribute she is allowed to access in any profile. We thus have  $|P|$  keys of the own profile and  $|P|$  keys for each of the  $n$  contacts, so each user has to store  $(n + 1) \cdot |P|$  keys.
- PMS-BE requires each user to store one key per attribute currently published in her profile and the key pair  $\langle PK, SK \rangle$ . Additionally, one (retrieval) key for each profile containing at least one attribute a user is allowed to access has to be stored by this user. Thus, each user has to store  $n + |P| + 2$  keys in this scheme.

Figure 6.1 shows the number of keys each user has to store in both constructions, depending on  $n$  and  $|P|$ . The number of keys in PMS-SK exceeds the number of keys in PMS-BE if  $N > 1$  or  $P > 1$  and  $(n, |P|) \neq (2, 2)$ . Although this is obviously the case in most applications, it is an interesting question how both constructions perform in different settings of common online social networks.

When analyzing the key overhead of both constructions, it has to be kept in mind that the smaller number of keys that have to be stored in PMS-BE comes at the cost of a lower privacy, as discussed in Section 5.



**Figure 6.1:** Plots of the number of keys each user has to store in PMS-SK respectively PMS-BE, depending on the average number of contacts  $n$  and the average number of attributes per profile  $|P|$ .

---

## 6.2 Key Overhead in real-world Settings

---

We will now discuss the key overhead of both proposed profile management schemes in real-world online social networks of different kind, namely Facebook<sup>1</sup>, Twitter<sup>2</sup>, XING<sup>3</sup>, and Flickr<sup>4</sup>.

---

### 6.2.1 Facebook

---

Being a very general platform for social networking, Facebook users share data with a quite high amount of contacts. Facebook's own statistics<sup>5</sup> state an average of 130 contacts per user, Golder et al. [12] found a mean of about 180 contacts per user in their analysis. According to Facebook's statistics, about 500 million active users share more than 30 billion pieces of content (e.g., web links, blog posts, photo albums, etc.) each month. Assuming a rather short lifetime of only three month per item, each user stores on average about 180 pieces of content, which would be attributes in our profile management scheme.

Obviously, the key overhead is much lower when applying PMS-BE to Facebook: Assuming an average of 150 contacts and 180 attributes, 332 keys have to be stored by each user in this scheme in contrast to over 27,000 keys in PMS-SK. Taking a key length of 192 bits for a symmetric or broadcast encryption key as a basis, this results in a storage overhead of about 8 Kilobytes (for PMS-BE) compared to about 650 Kilobytes (for PMS-SK).

---

### 6.2.2 Twitter

---

In contrast to Facebook, users of the microblogging service Twitter have on average approximately 50 contacts ("followers") and publish about 60 attributes ("tweets") per month.<sup>6</sup> Assuming again an attribute lifetime of three months, the number of stored keys per user is 232 in PMS-BE and over 9,000 in PMS-SK, resulting in about 6 Kilobytes respectively 220 Kilobytes storage overhead.

---

### 6.2.3 XING

---

XING is an online social network with currently over 10 million members<sup>7</sup> that focuses business contacts. It comes along with a set of 36 default attributes that is extended by on average 13.4 attributes individual for each user [23]. Furthermore, each user has on average 168 contacts. Thus, the overhead imposed by PMS-BE is about 220 keys (5 Kilobytes) and the one imposed by PMS-SK is about 8350 keys (200 Kilobytes).

---

### 6.2.4 Flickr

---

Flickr, an online community for image and video hosting, has a very low average of only 12 contacts ("friends") per user according to a study of Mislove et al. [20] in 2007. Assuming the limit of 200 images for a Free Account<sup>8</sup> as average number of attributes per profile, the number of keys each Flickr user has to store would be 214 in the PMS-BE construction and 2600 in PMS-SK, which yields a storage overhead of about 5 Kilobytes respectively 62 Kilobytes.

---

## 6.3 Discussion

---

As already shown in the analysis of both constructions, they differ notably in the complexity of the imposed overhead: Whilst PMS-SK requires about  $n \cdot |P|$  keys to be stored by each user, PMS-BE reduces this to about  $n + |P|$ . However, since  $n$  and  $|P|$  are relatively small values — at least concerning the online social networks we examined in this section —, the absolute difference between the storage overhead of both approaches is not very high. Although they differ relatively by a factor of 10 to 100, the absolute storage overhead is always below 1 Megabyte.

---

<sup>1</sup> <http://www.facebook.com>, October 2010

<sup>2</sup> <http://twitter.com>, October 2010

<sup>3</sup> <http://www.xing.com>, October 2010

<sup>4</sup> <http://www.flickr.com>, October 2010

<sup>5</sup> <http://www.facebook.com/press/info.php?statistics>, October 2010

<sup>6</sup> <http://www.website-monitoring.com/blog/2010/05/04/twitter-facts-and-figures-history-statistics/>, October 2010

<sup>7</sup> <http://corporate.xing.com/english/press/press-releases/details/article/press-releasebrxing-passes-10-million-member-ma/572/0d089bab50da33acf74565b3eb5711c0/>, October 2010

<sup>8</sup> <http://www.flickr.com/help/limits/>, October 2010

---

Thus, the two constructions allow for a trade-off between minimization of the storage overhead and maximization of privacy (as PMS-BE only provides anonymity, but not unlinkability). Applied to large profiles that may occur, e.g., in Facebook, the difference in storage overhead however increases rapidly. A user profile with 300 contacts<sup>9</sup> and 2,000 attributes leads to a key overhead of about 15 Megabytes in PMS-SK compared to 55 Kilobytes in PMS-BE.

An interesting question remaining for future work is how concrete implementations of both schemes behave both regarding the imposed storage overhead and especially the overhead in computation. The latter may be influenced heavily by the distribution of retrieval keys to the users allowed to access some attribute, which, based on the underlying encryption scheme, is significantly more complex in PMS-SK than in PMS-BE.

---

<sup>9</sup> More than 10% of the Facebook users have more than 300 contacts according to Gjoka et al. [11].

---

## 7 Conclusion and Outlook

Privacy preserving publication of personal data in user profiles is a valuable building block that can be used to bootstrap various collaborative and social data sharing applications. So far, security and privacy of user profiles have been addressed in a rather informal way, resulting in several proprietary implementations with unclear requirements.

In this work, we gave a rigorous security model for private user profiles, capturing the confidentiality of profile data and privacy of users that are allowed to retrieve this data. Our model allows for the construction of private profile management schemes, independently of the social application that will use them. It aims at local schemes, which can be used both in centralized and distributed environments.

Furthermore, we gave two concrete constructions of profile management schemes, one using symmetric and the other using broadcast encryption techniques. Our analysis showed that these constructions differ in the privacy guarantees they provide. While one offers strong privacy with the notion of unlinkability, the other represents a trade-off for better efficiency and key management, yet consequently offers only anonymity of the users.

Beyond that, we provided an insight into how the proposed schemes may behave in real-world settings, analyzing and discussing the imposed overhead. Interestingly, even though the broadcast encryption approach performs better in orders of magnitude, the absolute storage overhead of both schemes is quite low in multiple average settings.

An interesting open question remaining from this work is, whether unlinkability of users being authorized to access different attributes within a profile can be achieved with a sub-linear overhead.

---

## Bibliography

- [1] R. Baden, A. Bender, N. Spring, B. Bhattacharjee, and D. Starin. Persona: an online social network with user-defined privacy. In *ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM 2009)*, pages 135–146, 2009.
- [2] A. Barth, D. Boneh, and B. Waters. Privacy in encrypted content distribution using private broadcast encryption. In *10th International Conference on Financial Cryptography and Data Security (FC 2006)*, pages 52–64, 2006.
- [3] J. Bethencourt, A. Sahai, and B. Waters. Ciphertext-policy attribute-based encryption. In *2007 IEEE Symposium on Security and Privacy (S&P 2007)*, pages 321–334, 2007.
- [4] S. Buchegger, D. Schiöberg, L.-H. Vu, and A. Datta. PeerSoN: P2P social networking: early experiences and insights. In *Proceedings of the Second ACM EuroSys Workshop on Social Network Systems (SNS 2009)*, pages 46–52, 2009.
- [5] J. Callas, L. Donnerhacke, H. Finney, D. Shaw, and R. Thayer. OpenPGP Message Format. RFC 4880 (Informational), 2007.
- [6] R. Canetti, J. A. Garay, G. Itkis, D. Micciancio, M. Naor, and B. Pinkas. Multicast security: A taxonomy and some efficient constructions. In *18th IEEE Conference on Computer Communications (INFOCOM 1999)*, pages 708–716, 1999.
- [7] B. Carminati, E. Ferrari, and A. Perego. Rule-based access control for social networks. In *On the Move to Meaningful Internet Systems 2006: OTM 2006 Workshops (2)*, pages 1734–1744, 2006.
- [8] B. Carminati, E. Ferrari, and A. Perego. Enforcing access control in web-based social networks. *ACM Transactions on Information and System Security*, 13(1), 2009.
- [9] L. A. Cutillo, R. Molva, and T. Strufe. Safebook: A privacy-preserving online social network leveraging on real-life trust. *"IEEE Communications Magazine"*, Vol 47, Issue 12, *Consumer Communications and Networking Series*, 2009.
- [10] C. Gentry and B. Waters. Adaptive security in broadcast encryption systems (with short ciphertexts). In *28th Annual International Conference on the Theory and Application of Cryptographic Techniques (EUROCRYPT 2009)*, pages 171–188, 2009.
- [11] M. Gjoka, M. Kurant, C. T. Butts, and A. Markopoulou. Walking in Facebook: A Case Study of Unbiased Sampling of OSNs. In *29th IEEE Conference on Computer Communications (INFOCOM 2010)*, pages 2498–2506, 2010.
- [12] S. A. Golder, D. M. Wilkinson, and B. A. Huberman. Rhythms of social interaction: Messaging within a massive online network. In *Communities and Technologies 2007*, pages 41–66. 2007.
- [13] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, 1986.
- [14] K. Graffi, P. Mukherjee, B. Menges, D. Hartung, A. Kovacevic, and R. Steinmetz. Practical security in p2p-based social networks. In *34th Annual IEEE Conference on Local Computer Networks (LCN 2009)*, pages 269–272, 2009.
- [15] K. Graffi, S. Podrajanski, P. Mukherjee, A. Kovacevic, and R. Steinmetz. A distributed platform for multimedia communities. In *10th IEEE International Symposium on Multimedia (ISM 2008)*, pages 208–213, 2008.
- [16] R. Gross and A. Acquisti. Information revelation and privacy in online social networks. In *2005 ACM Workshop on Privacy in the Electronic Society (WPES 2005)*, pages 71–80, 2005.
- [17] J. Katz and Y. Lindell. *Introduction to Modern Cryptography (Chapman & Hall/Crc Cryptography and Network Security Series)*. Chapman & Hall/CRC, 2007.
- [18] M. M. Lucas and N. Borisov. flyByNight: mitigating the privacy risks of social networking. In *5th Symposium on Usable Privacy and Security (SOUPS 2009)*, 2009.



- 
- [19] D. A. McGrew and A. T. Sherman. Key establishment in large dynamic groups using one-way function trees. Technical Report No. 0755, TIS Labs at Network Associates, Inc., Glenwood, MD, May 1998.
- [20] A. Mislove, M. Marcon, P. K. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *7th ACM SIGCOMM Conference on Internet Measurement 2007*, pages 29–42, 2007.
- [21] PrimeLife. Scramble! <http://www.primelife.eu/results/opensource/65-scramble>, September 2010.
- [22] A. T. Sherman and D. A. McGrew. Key establishment in large dynamic groups using one-way function trees. *IEEE Transactions on Software Engineering*, 29(5):444–458, 2003.
- [23] T. Strufe. Profile popularity in a business-oriented online social network. In *Proceedings of the Third ACM EuroSys Workshop on Social Network Systems (SNS 2010)*, pages 1–6, 2010.
- [24] A. Tootoonchian, S. Saroiu, Y. Ganjali, and A. Wolman. Lockr: better privacy for social networks. In *2009 ACM International Conference on Emerging Networking Experiments and Technologies (CoNEXT 2009)*, pages 169–180, 2009.
- [25] D. Wallner, E. Harder, and R. Agee. Key Management for Multicast: Issues and Architectures. RFC 2627 (Informational), 1999.
- [26] C. K. Wong, M. G. Gouda, and S. S. Lam. Secure group communications using key graphs. In *ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM 1998)*, pages 68–79, 1998.
- [27] E. Zheleva and L. Getoor. To join or not to join: the illusion of privacy in social networks with mixed public and private user profiles. In *18th International Conference on World Wide Web (WWW 2009)*, pages 531–540, 2009.
- [28] Y. Zhu, Z. Hu, H. Wang, H. Hu, and G.-J. Ahn. A collaborative framework for privacy protection in online social networks. Cryptology ePrint Archive, Report 2010/491, 2010. <http://eprint.iacr.org/>.