Key Management in Distributed Online Social Networks WoWMoM 2011 – D-SPAN, Lucca



TECHNISCHE UNIVERSITÄT DARMSTADT

Felix Günther, Mark Manulis, Thorsten Strufe TU Darmstadt & CASED







Users and the Web



Modern Web is dominated by social interaction, networking, online communities. Services are offered by users for users – Users are at the heart of the Web.



Social Networks as Global Players



Social networking on the Web enjoys popularity all over the world.



Facebook alone has over 500 Mio. active users, 50% logging on every day.

Current Social Networks are Centralized



All big online social networks today are centralized.



Current Social Networks are Centralized



All big online social networks today are centralized.

Many trust has to be put into the Provider, which can act as a Big Brother.





One approach to avoid a Big Brother is to decentralize storage and control. This however introduces some challenges:



One approach to avoid a Big Brother is to decentralize storage and control. This however introduces some challenges:

availability of profiles



One approach to avoid a Big Brother is to decentralize storage and control. This however introduces some challenges:

 \blacktriangleright availability of profiles \rightarrow replication on friends' systems



One approach to avoid a Big Brother is to decentralize storage and control. This however introduces some challenges:

- \blacktriangleright availability of profiles \rightarrow replication on friends' systems
- privacy/access control



One approach to avoid a Big Brother is to decentralize storage and control. This however introduces some challenges:

- \blacktriangleright availability of profiles \rightarrow replication on friends' systems
- \blacktriangleright privacy/access control \rightarrow encryption of profile data



One approach to avoid a Big Brother is to decentralize storage and control. This however introduces some challenges:

- \blacktriangleright availability of profiles \rightarrow replication on friends' systems
- \blacktriangleright privacy/access control \rightarrow encryption of profile data

Qualitative Requirements for Encryption Schemes

- confidentiality: hide data from unauthorized users
- privacy: hide identities of authorized users



One approach to avoid a Big Brother is to decentralize storage and control. This however introduces some challenges:

- \blacktriangleright availability of profiles \rightarrow replication on friends' systems
- \blacktriangleright privacy/access control \rightarrow encryption of profile data

Qualitative Requirements for Encryption Schemes

- confidentiality: hide data from unauthorized users
- privacy: hide identities of authorized users

Quantitative Requirements for Encryption Schemes

- Iow storage overhead
- little interaction with authorized users
- low ressources requirements for computations

Formal Model for User Profiles and Profile Management



A profile P is modeled as a set of pairs

$$P \stackrel{\text{def}}{=} \left\{ (a, \bar{d}) | a \in \mathcal{I}, \bar{d} \in \{0, 1\}^* \right\}$$

- \mathcal{I} is set of unique attribute indices *a*; \overline{d} is the corresponding value stored in *P*.
- ▶ *P* is public + authenticated by owner U_P , having profile management key *pmk*.
- ▶ U_P given $(a, \bar{d}) \in P$ knows attribute d and group G of authorized users.
- Profile Management Scheme offers: Init(κ), Publish(pmk, P, (a, d), G), Retrieve(rk_U, P, a), Delete(pmk, P, a), ModifyAccess(pmk, P, a, U)



Security Goal: Confidentiality



U_P publishes pairs (a, \bar{d}) in P and gives U retrieval key rk_U for some indices. Confidentiality: Attributes d should remain hidden from unauthorized users.



Security Goal: Confidentiality



U_P publishes pairs (a, \bar{d}) in P and gives U retrieval key rk_U for some indices. Confidentiality: Attributes d should remain hidden from unauthorized users.



Indistinguishability approach:

A without access rights to (a, \overline{d}) should not be able to distinguish which attribute d is encrypted in \overline{d} .

 \ldots even if $\mathcal A$ can access other attributes in the same profile.

Privacy Goal: Unlinkability



Owner U_P knows which users were granted access to which pairs (a, d) in P. Unlinkability: Profiles should hide which users can access which attributes.



Privacy Goal: Unlinkability



Owner U_P knows which users were granted access to which pairs (a, d) in P. Unlinkability: Profiles should hide which users can access which attributes.



Indistinguishability approach:

A without access rights to (a, \bar{d}) should not be able to distinguish whether user A or user B was granted to access *a*.

Privacy Goal: Unlinkability



Owner U_P knows which users were granted access to which pairs (a, d) in P. Unlinkability: Profiles should hide which users can access which attributes.



Indistinguishability approach:

A without access rights to (a, \bar{d}) should not be able to distinguish whether user A or user B was granted to access *a*.

► Formal model + Conf./Unlink. games in Günther et al. FC/RLCPS 2011.

Shared Key (SK) Approach



- intuitive approach: shared secret key for each attribute
- ▶ separate keys $K_a \leftarrow SE.KGen(\kappa)$ for each pair (a, \bar{d}) : $\bar{d} = SE.Enc(K_a, d)$
- revocation: re-encryption with new K_a



Shared Key (SK) Approach



- intuitive approach: shared secret key for each attribute
- ▶ separate keys $K_a \leftarrow SE.KGen(\kappa)$ for each pair (a, \bar{d}) : $\bar{d} = SE.Enc(K_a, d)$
- revocation: re-encryption with new K_a



- provides confidentiality and perfect unlinkability
- each user has to store one key per attribute per profile
- two storage variants: at the users or (encrypted) in the profile
- key updates can be optimized with group key management for K_a (LKH, OFT)

Broadcast Encryption (BE) Approach



- ► each user manages own broadcast group using $(pk, sk) \leftarrow BE.Setup(\kappa, n)$
- each authorized user i receives a single key sk_i per profile
- ▶ for each (a, d) : $(Hdr, K_a) \leftarrow BE.Enc(S, pk)$, authorized users S, $\overline{d} = SE.Enc(K_a, d)$ and finally $\overline{d} = (Hdr, S, \overline{d})$.
- ▶ revocation: re-encryption with new (Hdr, K_a) for the modified set S



Broadcast Encryption (BE) Approach



- ▶ each user manages own broadcast group using $(pk, sk) \leftarrow BE.Setup(\kappa, n)$
- each authorized user i receives a single key sk_i per profile
- ▶ for each (a, d) : $(Hdr, K_a) \leftarrow BE.Enc(S, pk)$, authorized users S, $\overline{d} = SE.Enc(K_a, d)$ and finally $\overline{d} = (Hdr, S, \overline{d})$.
- ▶ revocation: re-encryption with new (*Hdr*, *K_a*) for the modified set *S*



- confidentiality and perfect anonymity (strictly weaker than unlinkability)
- each user has to store one key per profile





- storage requirements
 - at the profile owner (outside the profile)
 - in the profile
 - at the authorized users



- storage requirements
 - at the profile owner (outside the profile)
 - in the profile
 - at the authorized users
- number of encryptions
 - on initialization
 - on user addition
 - on user removal



- storage requirements
 - at the profile owner (outside the profile)
 - in the profile
 - at the authorized users
- number of encryptions
 - on initialization
 - on user addition
 - on user removal
- number of messages
 - on initialization
 - on user addition
 - on user removal



How much storage capacity is needed?



- storage plots for a single attribute published for N users
- linear growth for some property for SK and OFT
- BE only needs constant storage



How many encryptions are needed?



- encryptions needed for a single attribute published for N users
- linear or logarithmic number of encryptions for Shared Key and OFT
- only a single encryption needed for BE, however more expensive



How many messages are sent?



- messages needed for a single attribute published for N users
- no messages needed for Shared Key with profile-side storage and BE
- client-side SK and OFT approaches need linear messaging

Impact on Real-Life Communities



Analysis for Facebook, Twitter, XING, Flickr (based on their own statistics)

community	# contacts	# attributes	# keys SK BE		storage (KB) [*] SK BE	
facebook.	150	180	~27000	332	650	8
twitter	50	180	~9000	232	220	6
XING <mark>*</mark>	168	~36	~8350	220	200	5
flickr	12	200	2000	214	62	5

* 192bit keys (SE and BE)

- SK and BE costs differ by a factor of 10 to 80
 - SK profile-side storage adds factor "#contacts" (→ quadratic, difference 10²−10⁴)
- SK and BE overhead remains below 1 MB which could be acceptable

Impact on Real-Life Communities



Analysis for Facebook, Twitter, XING, Flickr (based on their own statistics)

community	# contacts	# attributes	# ke SK	eys BE	storage SK	e (KB)* BE
facebook.	150	180	~27000	332	650	8
twitter	50	180	~9000	232	220	6
XING <mark>*</mark>	168	~36	~8350	220	200	5
flickr	12	200	2000	214	62	5

* 192bit keys (SE and BE)

- SK and BE costs differ by a factor of 10 to 80
 - SK profile-side storage adds factor "#contacts" (→ quadratic, difference 10²−10⁴)
- SK and BE overhead remains below 1 MB which could be acceptable
- generic approaches also applicable to secure established networks

Summary



- modern web dominated by social interaction
- decentralization to avoid omnipotent service provider
- introduces the need for encryption of profile data

Summary



- modern web dominated by social interaction
- decentralization to avoid omnipotent service provider
- introduces the need for encryption of profile data
- we introduced formal model for user profiles and profile management, security goal confidentiality, and privacy goal unlinkability
- two generic encryption approaches: shared key and broadcast encryption
- analysis in real-world settings shows low storage overhead
- practical trade-off between storage overhead and privacy

Summary



- modern web dominated by social interaction
- decentralization to avoid omnipotent service provider
- introduces the need for encryption of profile data
- we introduced formal model for user profiles and profile management, security goal confidentiality, and privacy goal unlinkability
- two generic encryption approaches: shared key and broadcast encryption
- analysis in real-world settings shows low storage overhead
- practical trade-off between storage overhead and privacy
- open question: possible to achieve unlinkability with sub-linear key overhead?
- implementation of plugin for existing OSN is ongoing

Security Goal: Confidentiality Formal Definition



Confidentiality Game (high level):

- 1. Execute $\text{Init}(\kappa)$ for each user U.
- 2. \mathcal{A} interacts with users through queries (incl. Corrupt) until it outputs
 - ► (a, d₀), (a, d₁) two index-attribute pairs
 - G_t group of users
 - U_P profile owner who is not in G_t
- 3. Bit $b \in_R \{0, 1\}^*$. Execute Publish(*pmk*, *P*, (*a*, *d*_b), *G*_t).
- 4. A interacts with users through queries until it outputs some bit b^* .

 \mathcal{A} is successful if:

- ▶ *b* = *b**
- A did not corrupt U_P or any user who was ever authorized to access a
- ► *A* did not retrieve *d_b* trivially via some suitable Retrieve query

Profile Management Scheme is confidential if for all A:

|Pr[successfull attack] - 1/2| is negligible in κ .

Security Goal: Unlinkability Formal Definition



Unlinkability Game (high level):

- 1. Execute $\text{Init}(\kappa)$ for each user U.
- 2. \mathcal{A} interacts with users through queries until it outputs
 - ▶ U_0, U_1 two users, (a, d) index-attribute pair, U_P profile owner
- 3. Bit $b \in_R \{0, 1\}^*$.
 - ▶ If $(a, \cdot) \notin P$: execute Publish $(pmk, P, (a, d_b), \{U_b\})$.
 - If $(a, \cdot) \in P$: execute ModifyAccess (pmk, P, a, U_b) .
- 4. A interacts with users through queries until it outputs some bit b^* .

 \mathcal{A} is successful if:

- ▶ *b* = *b**
- U_P , U_0 , and U_1 are uncorrupted
- ► A did not query Retrieve(P, a, U₀) or Retrieve(P, a, U₁)

Profile Management Scheme is unlinkable if for all A:

|Pr[successfull attack] - 1/2| is negligible in κ .

Security Goal: Anonymity Formal Definition



Anonymity Game (high level):

- 1. Execute $\text{Init}(\kappa)$ for each user U.
- 2. \mathcal{A} interacts with users through queries until it outputs
 - ▶ U_0, U_1 two users, (a, d) index-attribute pair, U_P profile owner
- 3. Bit $b \in_R \{0, 1\}^*$.
 - ▶ If $(a, \cdot) \notin P$: execute Publish $(pmk, P, (a, d_b), \{U_b\})$.
 - If $(a, \cdot) \in P$: execute ModifyAccess (pmk, P, a, U_b) .
- 4. A interacts with users through queries until it outputs some bit b^* .

 \mathcal{A} is successful if:

- ▶ *b* = *b**
- U_P , U_0 , and U_1 are uncorrupted
- \mathcal{A} did not query Retrieve(P, a, U_0) or Retrieve(P, a, U_1)
- ▶ U_0 authorized to access some attribute $\iff U_1$ is also authorized

Profile Management Scheme is unlinkable if for all A:

Pr[successfull attack] – 1/2| is negligible in κ .