

Key Bundles and Parcels: Secure Communication in Many Groups

Eunjin Jung, Xiang-Yang Alex Liu, Mohamed G. Gouda
{ejung, alex, gouda}@cs.utexas.edu

Department of Computer Sciences
The University of Texas at Austin

Abstract. We consider a system where each user is in one or more elementary groups. In this system, arbitrary groups of users can be specified using the operations of union, intersection, and complement over the elementary groups in the system. Each elementary group in the system is provided with a security key that is known only to the users in the elementary group and to the system server. Thus, for any user u to securely multicast a data item d to every user in an arbitrary group G , u first forwards d to the system server which encrypts it using the keys of the elementary groups that comprise G before multicasting the encrypted d to every user in G . Every elementary group is also provided with a key tree to ensure that the cost of changing the key of the elementary group, when a user leaves the group, is small. We describe two methods for packing the key trees of elementary groups into key bundles and into key parcels. Packing into key bundles has the advantage of reducing the number of encryptions needed to multicast a data item to the complement of an elementary group. Packing into key parcels has the advantage of reducing the total number of keys in the system. We apply these two methods to a class of synthetic systems: each system has 10000 users and 500 elementary groups, and each user is in 2 elementary groups on average. Simulations of these systems show that our proposals to pack key trees into key bundles and key parcels live up to their promises.

1 Introduction

We consider a system that consists of n users denoted u_i , $0 \leq i < n$. The system users share one security key, called the system key. Each user u_i can use the system key to encrypt any data item before sending it to any subset of the system users, and can use it to decrypt any data item after receiving it from any other system user. (Examples of such systems are secure multicast systems [1], [2], [3], [4], secure peer-to-peer systems [5], and secure wireless networks [6].)

When a user u_i leaves the system, the system key needs to be changed so that u_i can no longer decrypt the encrypted data item exchanged within the system. This requires to add a server S to the system and to provide each system user u_j with an individual key K_j that only user u_j and server S know. When a user u_i leaves the system, server S changes the system key and sends the new key to each user u_j , other than u_i , encrypted using its individual key K_j . The cost of this rekeying scheme, measured by the number of needed encryptions, is $O(n)$, where n is the number of users in the system.

Clearly, this solution does not scale when the number of users become large. More efficient rekeying schemes have been proposed in [7], [8], [9], [10], [11], and [12]. A particular efficient rekeying scheme [3] and [4] is shown to cost merely $O(\log n)$ encryptions. This scheme is extended in [13], [14], and [15], and is shown to be optimal in [16], and has already been accepted as an Internet standard [3].

This scheme is based on a distributed data structure called a key tree. A *key tree* is a directed, incoming, rooted, balanced tree where each node represents a key. The root of the tree represents the system key and each leaf node represents the individual key of a system user. The number of leaf nodes is n , which is the number of users in the system. Each user knows all the keys on the directed path from its individual key to the root of the tree, and the server knows all the keys in the key tree. Thus, in a binary key tree, each user knows $\lceil \log_2 n \rceil + 1$ keys, and the server knows $(2n - 1)$ keys.

An example of a key tree for a system of 8 users is depicted in Figure 1(a). The root of the key tree represents the system key $K_{01234567}$ that is known to all users in the system. Each user also knows all the keys on the directed path from its individual key to the root of the key tree. For example, user u_7 knows all the keys K_7 , K_{67} , K_{4567} , and $K_{01234567}$.

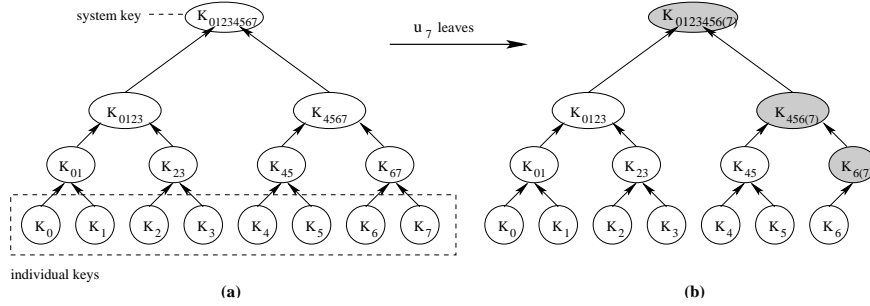


Fig. 1. A binary key tree before and after u_7 leaves

Figure 1(a) and 1(b) illustrates the protocol for updating the system key when user u_7 leaves the system. In this case, the system server S is required to change the keys $K_{01234567}$, K_{4567} , and K_{67} that user u_7 knows. To update these keys, S selects new keys $K_{0123456(7)}$, $K_{456(7)}$, and $K_{6(7)}$, encrypts them, and sends them to the users that need to know them. To ensure that u_7 cannot get a copy of the new keys, S needs to encrypt the new keys using keys that u_7 does not know. Therefore, S encrypts the new $K_{0123456(7)}$ with the old K_{0123} , encrypts the new $K_{0123456(7)}$ and the new $K_{456(7)}$ with the old K_{45} , encrypts the new $K_{0123456(7)}$, the new $K_{456(7)}$, and the new $K_{6(7)}$ with K_6 . Then, S multicasts the encrypted keys to the corresponding holders of these keys. The protocol can be specified as follows.

$$\begin{aligned}
 S &\rightarrow u_0, \dots, u_6 : \{u_0, u_1, u_2, u_3\}, K_{0123} < K_{0123456(7)} | \text{chk} > \\
 S &\rightarrow u_0, \dots, u_6 : \{u_4, u_5\}, K_{45} < K_{0123456(7)} | K_{456(7)} | \text{chk} > \\
 S &\rightarrow u_0, \dots, u_6 : \{u_6\}, K_6 < K_{0123456(7)} | K_{456(7)} | K_{6(7)} | \text{chk} >
 \end{aligned}$$

This protocol consists of three steps. In each step, server S broadcasts a message consisting of two fields to every user in the system. The first field defines the set of the intended ultimate destinations of the message. The second field is an encryption, using an old key, of the concatenation of the new key(s) and a checksum computed over the new key(s). Note that although the broadcast message is sent to every user in the system, only users in the specified destination set have the key used in encrypting the message and so only they can decrypt the message.

The above system architecture is based on the assumption that the system users constitute a single group. In this paper, we extend this architecture to the case where the system users form many groups.

2 Groups and Group Algebra

Assume that the system has m , $m \geq 1$, elementary groups: each elementary group is a distinct subset of the system users and one elementary group has all the system users. Every elementary group has a unique identifier G_j , $0 \leq j \leq m-1$. The identifier for the elementary group that has all users is G_0 . As an example, Figure 2 illustrates a system that has eight users u_0 through u_7 and five elementary groups G_0 , G_1 , G_2 , G_3 , and G_4 .




Fig. 2. A sample system

The system needs to be designed such that any user u_i can securely multicast data items to all users in any elementary group G_j . Moreover, any user u_i can securely multicast data items to all users in any group, where a group is defined recursively according to the following four rules:

- i. Any of the elementary groups G_0, \dots, G_{m-1} is a group.
- ii. The union of any two groups is a group.
- iii. The intersection of any two groups is a group.
- iv. The complement of any group is a group. (Note that the complement of any group G is the set of all users in G_0 that are not in G)

Thus, the set of groups is closed under the three operations of union, intersection, and complement.

Each group can be defined by a group formula that includes the following symbols.

- G_0 through G_{m-1}

- \vee for union
- \wedge for intersection
- \neg for complement

Group formulae can be manipulated using the well-known laws of algebra: associativity, commutativity, distribution, De-Morgan's, and so on. For example, the group formula

$$G_1 \vee \neg(\neg G_2 \wedge G_1)$$

can be manipulated as follows:

$$\begin{aligned}
& G_1 \vee \neg(\neg G_2 \wedge G_1) \\
&= \{\text{by De Morgan's}\} G_1 \vee (\neg\neg G_2 \vee \neg G_1) \\
&= \{\text{by associativity of } \vee\} G_1 \vee \neg\neg G_2 \vee \neg G_1 \\
&= \{\text{by definition of complement}\} G_1 \vee G_2 \vee \neg G_1 \\
&= \{\text{by commutativity of } \vee\} G_1 \vee \neg G_1 \vee G_2 \\
&= \{\text{by definition of complement}\} G_0 \vee G_2 \\
&= \{\text{by definition of } \vee\} G_0
\end{aligned}$$

From this formula manipulation, it follows that the group defined by the formula $G_1 \vee \neg(\neg G_2 \wedge G_1)$ is the set of all system users. Thus, for a user u_i to securely multicast a data item d to every user in the group $G_1 \vee \neg(\neg G_2 \wedge G_1)$, it is sufficient for u_i to securely broadcast d to every user in the system.

In the rest of this paper, we consider solutions for the following problem. How to design the system so that any system user u_i can securely multicast data items to any group G in the system. Any reasonable solution for this problem needs to take into account that the users can leave any elementary group in the system or leave the system altogether, and these activities may require to change the security keys associated with the elementary groups from which users leave. In particular, the solution should utilize key trees, discussed in Section 1, that can reduce the cost of changing the security keys from $O(n)$ to $O(\log n)$, where n is the total number of users in the system.

The above problem has many applications. As a first example, consider a music file sharing system that has four elementary groups: Rock, Jazz, Blues, and Do-Not-Disturb. A user u_i in this system may wish to securely distribute a song of Louis Armstrong to all interested users. In this case, user u_i securely multicasts the song to all users in the group, $\text{Jazz} \wedge \neg \text{Do-Not-Disturb}$.

As a second example, consider a student registration system in some university. This system has m elementary groups G_0 through G_{m-1} , where each G_i is a list of the students registered in one course section. A professor who is teaching three sections G_5 , G_6 , G_7 of the same course, may wish to securely multicast any information related to the course to all the students in the group $G_5 \vee G_6 \vee G_7$.

3 Key Bundles

The above problem suggests the following simple solution (which we show below that it is ineffective). First, assign to each elementary group G_j a security key to be shared

by all the users of G_j . Second, assign to the complement $\neg G_j$ of each elementary group G_j a security key to be shared by every member of this complement. Third, provide a key tree for each elementary group and another key tree for its complement. Note that the two key trees provided for an elementary group and its complement span all the users in the system. Thus, these two trees can be combined into one *complete* key tree that spans all system users in the system. Figure 3 shows the four complete key trees that are provided for the four elementary groups and their complements in the system in Figure 2.

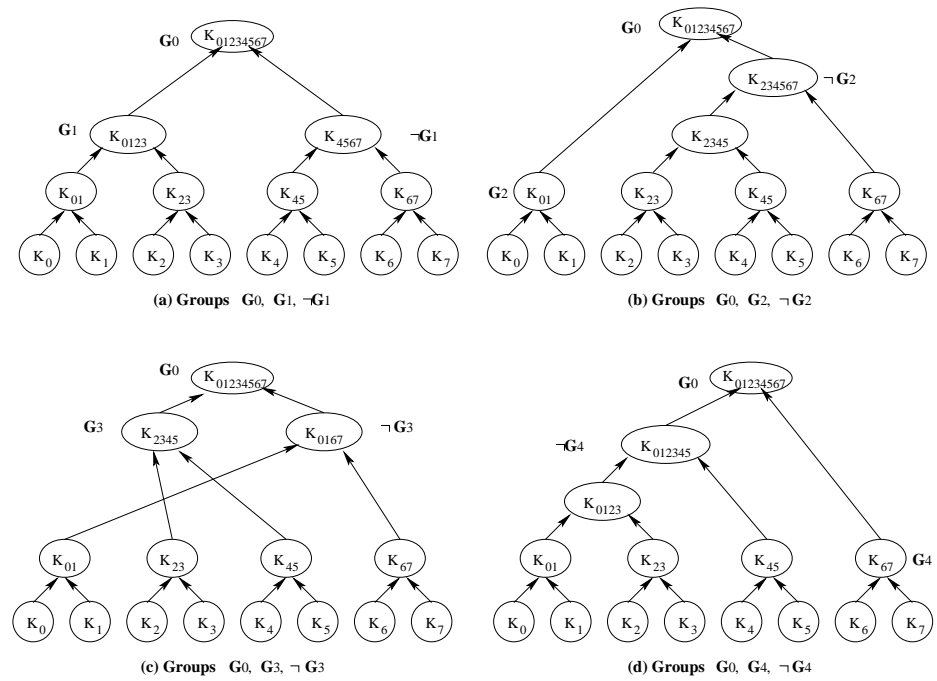


Fig. 3. The complete key trees for the elementary groups and their complements

From Figure 3(a), the key for the elementary group G_1 is K_{0123} and the key for its complement is K_{4567} . From Figure 3(b), the key for the elementary group G_2 is K_{01} and the key for its complement is K_{234567} . From Figure 3(c), the key for the elementary group G_3 is K_{2345} , and the key for its complement is K_{0167} . From Figure 3(d), the key for the elementary group G_4 is K_{67} , and the key for its complement is K_{012345} .

Note that these complete trees have the same key for group G_0 , and the same individual key for each user. Nevertheless, the total number of distinct keys in these complete trees is 19, which is relatively large for this rather simple system. In general, this method requires $O(mn)$ keys, where m is the number of elementary groups and n is the number of users in the system.

To reduce the total number of needed keys, several elementary groups can be added to the same complete key tree, provided that these elementary groups are “nonconflicting”. This idea suggests the following three definitions of nonconflicting elementary groups, bundles, and bundle covers.

Two elementary groups are *nonconflicting* if and only if either their intersection is empty or one of them is a subset of the other. In the system example in Figure 2, the three elementary groups G_0 , G_1 and G_2 are nonconflicting since G_1 is a subset of G_0 , and G_2 is a subset of G_1 . On the other hand, the two elementary groups G_1 and G_3 are conflicting, because they share two users u_2 and u_3 and neither group is a subset of the other.

A *bundle* of a system is a maximal set of nonconflicting elementary groups of the system. In the system example in Figure 2, the four elementary groups G_0 , G_1 , G_2 , G_4 constitute one bundle B_0 , and the four elementary groups G_0 , G_2 , G_3 , G_4 constitute a second bundle B_1 .

A *bundle cover* of a system is a set $\{B_0, \dots, B_{m-1}\}$ of system bundles such that the following two conditions hold:

- i. *Completeness*: Each elementary group of the system appears in some bundle B_i in the bundle cover.
- ii. *Compactness*: Each bundle B_i has at least one elementary group that does not appear in any other bundle B_j in the bundle cover.

Note that the set $\{B_0, B_1\}$, where $B_0 = \{G_0, G_1, G_2, G_4\}$ and $B_1 = \{G_0, G_2, G_3, G_4\}$, is a bundle cover for the system in Figure 2.

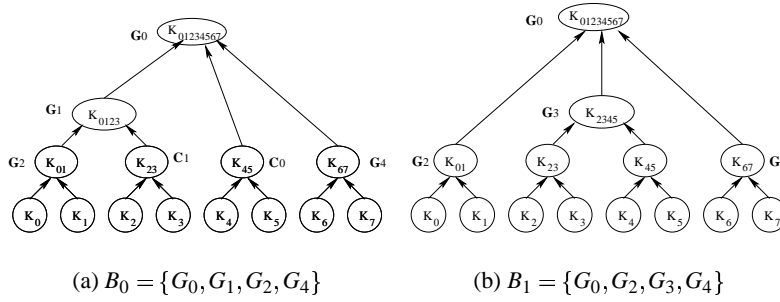


Fig. 4. The complete key trees for the two bundles B_0 and B_1

The security keys for the elementary groups in a bundle can be arranged in a complete key tree. For example, Figure 4(a) shows the complete key tree for B_0 . In this tree, the key for group G_0 is $K_{01234567}$, the key for group G_1 is K_{0123} , the key for G_2 is K_{01} , and the key for G_4 is K_{67} . Note that users u_4 and u_5 in G_0 do not belong to any other elementary group in the bundle, and so they are viewed as forming a complement group C_0 whose key is K_{45} . We refer to a complete key tree that corresponds to a bundle as a *key bundle*.

Figure 4(b) shows the complete key bundle for B_1 . Note that in this bundle every user in G_0 is also in another elementary group. Thus, the resulting complete key tree does not have a complement group as in the former key tree in Figure 4(a).

Comparing the two key bundles in Figures 4(a) and 4(b), one observes that each of the elementary groups G_0 , G_2 , and G_4 appear in both key bundles because none of them conflict with any elementary group or any group in the system. One also observes that each of these groups has the same group key in both key bundles, and that the individual key of each user is the same in both key bundles. Note that these key bundles have only 15 distinct keys compared with the 19 distinct keys in the four complete trees in Figure 3. This represents more than 20% reduction in the total number of keys in the system.

The system server S knows the two key bundles in Figure 4, and each user u_i knows only the keys that exist on the paths from its individual key K_i to the key of group G_0 . Thus, each user u_i needs to collaborate with the system server S in order to securely multicast data items to any elementary group or any group that can be defined by intersection, union, and complement of elementary groups. This point is illustrated by the following four examples.

For the first example, assume that user u_0 wants to securely multicast a data item d to every user in group G_4 . In this case, user u_0 can execute the following protocol.

$$\begin{aligned} u_0 \rightarrow S & \quad : K_0 < d|G_4|chk > \\ S \rightarrow u_0, \dots, u_7 & : G_4, \quad K_{67} < d|u_0|chk > \end{aligned}$$

This protocol consists of two steps. In the first step, user u_0 sends a message $K_0 < d|G_4|chk >$ to server S . This message consists of three concatenated fields, namely the data item d , its intended destination G_4 , and the checksum chk ; the message is encrypted by the individual key K_0 of user u_0 . In the second step, server S multicasts the message $G_4, K_{67} < d|u_0|chk >$ where the second field consists of the data item d , the message source u_0 , and the checksum chk and is encrypted with the group key of G_4 .

For the second example, assume user u_1 wants to securely multicast a data item d to the users in either group G_1 or G_3 , namely the users in the union of G_1 and G_3 . In this case, user u_1 can execute the following protocol.

$$\begin{aligned} u_1 \rightarrow S & \quad : K_1 < d|G_1 \vee G_3|chk > \\ S \rightarrow u_0, \dots, u_7 & : G_1 \vee G_3, \quad K_{0123} < d|u_1|chk >, \quad K_{2345} < d|u_1|chk > \end{aligned}$$

In the second step of this protocol, server S multicasts the message $G_1 \vee G_3, K_{0123} < d|u_1|chk >, K_{2345} < d|u_1|chk >$ to the two groups G_1 and G_3 . The users in group G_1 can get d by using the group key K_{0123} to decrypt $K_{0123} < d|u_1|chk >$ and the users in group G_3 can get d by using the group key K_{2345} to decrypt $K_{2345} < d|u_1|chk >$. Note that if it is u_2 who wants to send d to $G_1 \vee G_3$, then since u_2 belongs to both G_1 and G_3 , u_2 already knows both K_{0123} and K_{2345} . Therefore, u_2 can send the encrypted d directly to the users in G_1 and G_3 as follows:

$$u_2 \rightarrow u_0, \dots, u_7 : G_1 \vee G_3, \quad K_{0123} < d|u_2|chk >, \quad K_{2345} < d|u_2|chk >$$

For the third example, assume that user u_4 wants to send a data item d to all the users in the intersection of G_1 and G_3 . In this case, user u_4 can execute the following

protocol.

$$\begin{aligned} u_4 \rightarrow S & : K_4 \langle d|G_1 \wedge G_3|chk \rangle \\ S \rightarrow u_0, \dots, u_7 & : G_1 \wedge G_3, \quad K_{0123} \langle K_{2345} \langle d|u_4|chk \rangle \rangle \end{aligned}$$

In the second step of this protocol, server S multicasts a message $G_1 \wedge G_3, K_{0123} \langle K_{2345} \langle d|u_4|chk \rangle \rangle$ to the group $G_1 \wedge G_3$. Here the concatenation of d, u_4 and chk is encrypted by both the group key of G_1 , which is K_{0123} , and the group key of G_3 , which is K_{2345} . The encrypted message can only be decrypted by the users that are in both G_1 and G_3 because only these users know the two group keys K_{0123} and K_{2345} .

For the fourth example, assume that user u_5 wants to send a data item d to all the users in the complement of group G_1 . In this case, user u_5 executes the following protocol.

$$\begin{aligned} u_5 \rightarrow S & : K_5 \langle d|\neg G_1|chk \rangle \\ S \rightarrow u_0, \dots, u_7 & : C_0 \vee G_4, \quad K_{45} \langle d|u_5|chk \rangle, \quad K_{67} \langle d|u_5|chk \rangle \end{aligned}$$

After server S receives this message, it translates $\neg G_1$ to $C_0 \vee G_4$ then multicasts the message $C_0 \vee G_4, K_{45} \langle d|u_5|chk \rangle, K_{67} \langle d|u_5|chk \rangle$. The users in group G_c can get d using the group key K_{45} , and the users in group G_4 can get d using the group key K_{67} .

The algorithm for the construction of a bundle cover is described in detail in the full version of this paper[17].

4 Key Parcels

A bundle is defined as a maximal set of nonconflicting elementary groups in the system. From this definition the elementary group G_0 is in every bundle since it does not conflict with any other elementary group in the system. Thus, every key bundle is a complete key tree.

This feature of bundle maximality has one advantage and one disadvantage. The advantage is that the complement of any elementary group in a bundle B_j can be expressed as the union of some other elementary groups in B_j . Thus, securely multicasting a data item to the complement of any elementary group can be carried out efficiently. The disadvantage is that the number of keys needed in each key bundle is relatively large, and so the total number of keys in the system is relatively large.

Clearly, the disadvantage of bundle maximality outweighs its advantage in systems where users never need to securely multicast data items to the complements of elementary groups. Therefore, in these systems, we use ‘‘parcels’’, which are not maximal, instead of bundles, which are maximal. The definitions of parcels and parcel covers are given next.

A *parcel* of a system is a set of nonconflicting elementary groups of the system.

A *parcel cover* of a system is a sequence of parcels (P_0, \dots, P_{s-1}) such that the following two conditions hold:

- i. *Completeness*: Each elementary group of the system appears in some parcel P_i in the parcel cover.

- ii. *Compactness*: Each elementary group in each parcel P_i conflicts with at least one elementary group in each of the preceding parcels P_0, \dots, P_{i-1} in the parcel cover.




Fig. 5. The complete key trees for the parcel cover

As an example, a parcel cover for the system in Figure 2 is (P_0, P_1) , where $P_0 = \{G_0, G_1, G_2, G_4\}$ and $P_1 = \{G_3\}$. Figure 6 is a parcel cover (P_0, P_1) for the system in Figure 2.

The security keys for the elementary groups in a parcel can be arranged in a key tree that is not necessarily a complete tree. Figure 6(a) shows the key tree for parcel P_0 consisting of the elementary groups $G_0, G_1, G_2,$ and G_4 . Figure 6(b) shows the key tree for parcel P_1 consisting of the elementary group G_3 . Note that the key tree for parcel P_1 is not a complete tree. We refer to a key tree that corresponds to a parcel as a *key parcel*.

The algorithm to construct a parcel cover is described in detail in the full version of this paper[17].

5 Simulation Results

In this section, we present the results of simulations that we carried out to demonstrate the feasibility of key bundles and key parcels. In our simulation, we used a class of synthetic systems with the following properties:

- i. The number of users in each system varies from 1000 to 10000.
- ii. Each system has 500 elementary groups.
- iii. In each system, a user joins 2 elementary groups on average.

Each system is simulated 100 times and the averages of the following four items are computed over the 100 simulation runs for each system: the number of bundles or parcels in the system cover, the total number of keys in the system, the number of keys per user, and the number of encryptions needed to multicast a data item to the users in the complement of an elementary group. The results of these simulations are shown in Figures 6 through 9.

As shown in Figure 6, the number of bundles in a bundle cover more or less equals the number of parcels in a parcel cover. Note that this number increases logarithmically as the number of users in the simulated system grows.

Figure 7 shows that the number of keys in systems that use key bundles and the number of keys in systems that use key parcels grow linearly as the number of users in the system increases. However, the number of keys in the case of key bundles grows much faster in the case of key parcels. This is because each key bundle is a complete key tree while each key parcel is not necessarily complete.




Fig. 6. Number of bundles or parcels




Fig. 7. Number of keys in the system

As shown in Figure 8, the number of keys that each user needs to store increases as a logarithm function with the number of users in the system.

Figure 9 shows that when Key Bundle approach is used, the average number of encryptions performed for a complement of an elementary group decreases as the number of users increases (the actual number is from around 500 to 400). As the number of users increases, the probability of two groups' conflicting increases. Therefore, the average number of groups can be put in a bundle decreases. Since we use the keys of other groups in the same bundle of a complement of an elementary group, the number of encryptions decreases as the number of users increases.

The number of encryptions for a complement performed in Key Parcel approach is constantly larger than that in Key Bundle approach, and the difference becomes greater as the number of users grows. It is because that the number of encryptions for a complement $\neg G$ for a random elementary group G in key parcels is $O(n)$, since the server needs to use the individual keys for the users that are not in the parcel which G belongs to.

6 Conclusion

We consider a system where each user is in one or more elementary groups. In this system, arbitrary groups of users can be specified using the operations of union, intersection, and complement over the elementary groups in the system. Every elementary




Fig. 8. Number of keys per user




Fig. 9. Number of encryptions per complement

group is provided with a key tree to reduce the cost of changing the key of the elementary group. We propose two methods for packing the key trees of elementary groups into key bundles and into key parcels.

Packing into key bundles has the advantage of reducing number of encryptions needed to multicast a data item to the complement of an elementary group. Packing into key parcels has the advantage of reducing the total number of keys in the system. We apply these two methods to a class of synthetic systems: each system has from 1000 to 10000 users and 500 elementary groups, and a user in each system is in 2 elementary groups on average. Simulations of these systems show that our proposal to pack key trees into key bundles and key parcels provides a reasonable performance to be used. The number of keys stored per user in the case of key bundles is 12 for 10000 user system, while that in the case of key parcels is 5. Instead, the number of encryptions needed for a complement in the case of key bundles is far less than that in the case of key parcels by the magnitude of $O(\frac{n}{k})$.

As a future work, we would like to find a hybrid between these two methods, which needs less number of keys in the system than in the case of key bundles and at the same time supports complement of an elementary group with less number of encryptions than in the case of key parcels.

We are also interested in conducting a case study of these methods in a real world application. The case study includes to define appropriate scopes of elementary groups and to maintain key bundles or key parcels accordingly. As a typical application of secure group communication, a knowledge sharing system can take advantage of these methods.

As described in Section 3, if the sender of message m does not know the keys required to encrypt m appropriately, the system server has to encrypt m and multicast. This requirement for the server's help may cause performance bottleneck at the server. To reduce the workload of the system server, future work will investigate how multiple servers may be placed and coordinated to work in a distributed or a hierarchical manner.

References

1. Gong, L.: Enclaves: Enabling secure collaboration over the internet. *IEEE Journal of Selected Areas in Communications* **15** (1997) 567–575
2. Mitra, S.: Iolus: a framework for scalable secure multicasting. In: the Proceedings of the ACM SIGCOMM '97, ACM Press (1997) 277–288
3. Wallner, D.M., Harder, E.J., Agee, R.C.: Key management for multicast: Issues and architectures. RFC 2627 (1999)
4. Wong, C.K., Gouda, M., Lam, S.S.: Secure group communications using key graphs. *IEEE/ACM Transactions on Networking (TON)* **8** (2000) 16–30
5. Steiner, M., Tsudik, G., Waidner, M.: Key agreement in dynamic peer groups. *IEEE Transactions on Parallel and Distributed Systems* **11** (2000) 769–780
6. Gong, L., Shacham, N.: Multicast security and its extension to a mobile environment. *Wireless Networks* **1** (1995) 281–295
7. Ballardie, A.: Scalable multicast key distribution. RFC 1949 (1996)
8. Chang, I., Engel, R., Kandlur, D.D., Pendarakis, D.E., Saha, D.: Key management for secure internet multicast using boolean function minimization techniques. In: the Proceedings of IEEE Infocom 1999. Volume 2. (1999) 689–698
9. Naor, D., Naor, M., Lotspiech, J.: Revocation and tracing schemes for stateless receivers. In: Advances in Cryptology - CRYPTO 2001. LNCS 2139, Springer-Verlag (2001) 41–62
10. Rodeh, O., Birman, K., Dolev, D.: The architecture and performance of security protocols in the ensemble group communication system: Using diamonds to guard the castle. *ACM Transactions on Information and System Security (TISSEC)* **4** (2001) 289–319
11. Setia, S., Koussih, S., Jajodia, S., Harder, E.: Kronos: A scalable group re-keying approach for secure multicast. In: the Proceedings of IEEE Symposium on Security and Privacy. (2000)
12. Waldvogel, M., Caronni, G., Sun, D., Weiler, N., Plattner, B.: The versakey framework: Versatile group key management. *IEEE Journal on Selected Areas in Communications* **17** (1999) 1614–1631
13. Gouda, M.G., Huang, C.T., Elnozahy, E.: Key trees and the security of the interval multicast. In: the Proceedings of the 22 nd International Conference on Distributed Computing Systems. (2002) 467–468
14. Li, X.S., Yang, Y.R., Gouda, M.G., Lam, S.S.: Batch rekeying for secure group communications. In: the Proceedings of the 10th international World Wide Web conference on World Wide Web, ACM Press (2001) 525–534
15. Yang, Y.R., Li, X.S., Zhang, X.B., Lam, S.S.: Reliable group rekeying: a performance analysis. In: the Proceedings of the 2001 conference on applications, technologies, architectures, and protocols for computer communications, ACM Press (2001) 27–38
16. Snoeyink, J., Suri, S., Varghese, G.: A lower bound for multicast key distribution. In: the Proceedings of IEEE Infocom 2001. (2001) 667–675
17. Jung, E., Liu, X.Y.A., Gouda, M.G.: Key bundles and parcels: Secure communication in many groups. Technical Report TR-03-21, Dept. of Computer Sciences, the University of Texas at Austin (2003)