

# On Dynamic Subset Difference Revocation Scheme

Weifeng Chen<sup>1</sup>, Zihui Ge<sup>2</sup>, Chun Zhang<sup>1</sup>, Jim Kurose<sup>1</sup>, and Don Towsley<sup>1</sup>

<sup>1</sup> Department of Computer Science  
University of Massachusetts at Amherst, MA 01002, USA  
{chenwf, czhang, kurose, towsley}@cs.umass.edu  
<sup>2</sup> AT&T Labs-Research, Florham Park, NJ 07932  
gezihui@research.att.com

**Abstract.** Subset Difference Revocation (SDR) [8] has been proposed to perform group rekeying in a stateless manner. However, statelessness comes at a cost in terms of storage and message overhead when the number of currently active members is much smaller than the number of potential group members [3]. We propose a *dynamic* SDR scheme to address these problems. Rather than maintaining a large static key tree that accommodates all potential group members, we use a smaller dynamic key tree for only currently active members. We dynamically assign current members to the positions in the key tree rather than using fixed pre-assignment. The smaller key tree requires less storage and dynamic assignment achieves a smaller rekeying cost. Our evaluation shows that the dynamic scheme significantly improves the performance of SDR, reducing by half the rekey communication cost in the case that the number of the currently active members is much less than the total number of potential members. Compared to SDR in [8], dynamic SDR does not need to know the maximum number of potential group members in advance, a value that can be difficult to estimate in practice.

**Keywords:** Multicast security, Group rekeying, Subset Difference Revocation

## 1 Introduction

Membership-based applications, such as pay-per-view and specialized information services (e.g., stock price, live news), require that information content be delivered to (and only to) subscribed members. This is typically accomplished by encrypting data using a common *Traffic Encryption Key* (TEK) that is shared by all currently active members. When a member joins the group, the TEK must be changed to ensure that the newly joining member cannot decrypt previous communications (a requirement known as “backward confidentiality”). Similarly, the TEK must be changed when a member leaves the group to ensure that future messages cannot be decrypted by the departing member (a requirement known as “forward confidentiality”). The algorithms that manage the distribution, updating and revocation of the TEK are collectively known as *group key management*

protocols. The IETF MSEC framework suggests using a *Group Controller and Key Server (GCKS)* for rekeying. Generally, the TEK is encrypted using *Key Encryption Keys (KEKs)* and then multicast by the GCKS. Several works have dealt with the group rekeying problem [2,6,7,8,9,10,11]. Subset Difference Revocation (SDR) [8] has been proposed as a “stateless” group rekeying algorithm. By stateless, it is meant that members do not need to keep track of rekeying messages in order to maintain their states. This desirable property comes at a price, however [3] - SDR can require high key storage at both the member and GCKS sides, and can generate a significant amount of messaging traffic during rekeying.

These two problems arise from the fact that SDR maintains a static key tree that is constructed in advance. This tree must be large enough to hold all potential members. Since key storage cost is determined by the size of the key tree [8], key storage costs can thus be large. Further, a member joining the group is attached to a pre-assigned position in the key tree. Assuming that member activity is independent of position, when the number of currently active members is much smaller than the number of potential members, the positions occupied by active members are likely to be *sparse*, i.e., there are many holes (positions not occupied by an active member) among the positions occupied by active members. The sparse distribution of these positions can cause SDR to perform as inefficiently as encrypting the TEK separately for each active member [3]. We will refer to the SDR proposed in [8] as *static SDR* in this paper.

In this paper, we propose a *dynamic SDR* scheme to reduce both the key storage cost and the rekey communication cost of static SDR. Dynamic SDR uses a key tree that is large enough to hold currently active members (as opposed to *all* members, both active and inactive) in order to reduce key storage. This is done by dynamically assigning a joining member a new position in the tree. Ideally, our goal will be to position active members adjacent in the key tree in order to reduce rekeying cost.

This paper has the following contributions. First, we design a new group rekeying algorithm, dynamic SDR. The algorithm is based on Subset Difference and uses a dynamic key tree to reduce both storage costs and messaging overhead. The algorithm is “multicast stateless” (i.e., it does not require nodes to maintain states when receiving the multicast rekeying messages, which distribute only TEKs) and does not require *a priori* knowledge of the number of potential members. Secondly, we propose and evaluate several enhancements on dynamic SDR. Our simulations show that dynamic SDR significantly reduces both the key storage cost and rekey communication cost when the number of currently active members is much less than the number of potential members. Finally, we investigate the tradeoff between the unicast and multicast costs in dynamic SDR.

The rest of the paper is organized as follows: In Section 2, we briefly overview the static Subset Difference Revocation algorithm. The dynamic SDR scheme is described in Section 3. Section 4 presents our evaluation. Section 5 discusses

various properties of dynamic SDR. Related work is given in Section 6. Finally, we conclude the paper in Section 7.

## 2 Background: Subset Difference Revocation Algorithm

Static SDR [8] is a tree-based group key management protocol. For a finite set of potential members  $\mathcal{N}$  ( $N = |\mathcal{N}|$ ), the GCKS maintains a tree with  $N$  leaves and assigns a fixed position (a leaf in the key tree) to each distinct member<sup>1</sup>. For a node  $i$  in the key tree, let  $S_i$  be the set of the potential members which are descendants of  $i$ . Given two nodes  $i$  and  $j$ , where  $j$  is a descendant of  $i$ , subset  $S_{i,j}$  is defined as  $S_i \setminus S_j$ . Each subset  $S_{i,j}$  is initially assigned a long-lived key  $K_{i,j}$  that is only known by members covered by  $S_{i,j}$ . Let  $\mathcal{M} \subseteq \mathcal{N}$  be the set of members currently active in the group and  $M = |\mathcal{M}|$ . To distribute an updated TEK, the GCKS uses a *Subset-Cover* framework to partition  $\mathcal{M}$  into disjoint subsets. More specifically, the GCKS divides  $\mathcal{M}$  into  $s$  subsets such that  $\bigcup_{d=1}^s S_{i_d,j_d} = \mathcal{M}$  and  $\forall m \in \mathcal{M}$ ,  $m$  is covered by one and exact one resultant subset. The updated TEK is then encrypted using  $K_{i_d,j_d}$ . Consequently, only members in  $\mathcal{M}$  can deduce the new TEK.

SDR is a *stateless algorithm* for the KEKs ( $K_{i,j}$ ) are unchanged after initialization such that each member  $m$  in  $\mathcal{M}$  is able to deduce the KEKs at each rekeying instance based on the *secret information*, denoted as  $I_m$ , received during initialization.  $I_m$  allows  $m$  to deduce the keys of *all* possible subsets to which  $m$  may belong.

Although the total number of possible subsets to which a member  $m$  may belong is  $O(N)$ , the size of  $I_m$  is  $|I_m| = (\log^2 N + \log N)/2 + 1$  [3]. It is important to note that  $I_m$  is the only required information for member  $m$  to participate in the group communication for the statelessness of SDR. Also note that  $I_m$  is determined by the position of  $m$  in the key tree. As a consequence, once a leaf position in the key tree is assigned to a member  $m$ , that position cannot be assigned to any other member even when  $m$  is currently not in the group. That is the reason why the GCKS in static SDR needs to maintain a key tree large enough for  $\mathcal{N}$ . Typically, in static SDR, a *returned* member (a member joining the group again after leaving) is assigned to the position that the member was assigned last time.

The number of the resultant subsets of  $\mathcal{M}$  is determined by the positions of members of  $\mathcal{M}$  in the key tree of size  $N$ . Generally speaking, under the assumption that member activity is independent of position in the key tree, the larger the difference between  $N$  and  $M$ , the more likely that active members are sparsely distributed in the key tree, resulting in  $O(M)$  disjoint subsets. On the other hand, the smaller the difference between  $N$  and  $M$ , the more adjacent the positions of active members in the key tree, resulting in  $O(N - M)$  disjoint subsets. As pointed out by [3], given  $M$ , the expected number of resultant subsets


<sup>1</sup> In the rest of this paper, a leaf node in the key tree and the member assigned to that node are treated indistinguishable when there is no risk of ambiguity.

of  $\mathcal{M}$  is  $\Theta(\min(N - M, M))$ . Moreover, static SDR requires  $O(N \log N)$  storage overhead at the GCKS [3], and  $O(\log^2 N)$  storage overhead at the member side.

### 3 Dynamic SDR

We have seen that static SDR generally requires a very large key tree to accommodate all unique members. As a consequence, currently active members, usually a small fraction of all of the unique members, are likely to be widely dispersed in the key tree space. Both of these factors decrease the performance of SDR. In this section, we propose a *dynamic SDR* approach that addresses these two inefficiencies of static SDR. First, we describe an observation on static SDR, based on which dynamic SDR is proposed.

In the binary key tree of SDR, a node  $i$  has a height of  $h$  if the subtree  $T_i$  rooted at  $i$  has  $2^h$  leaves. A leaf itself has a height of 1. The heights of  $T_i$  and  $S_{i,j}$  are also defined as the height of node  $i$ .



**Fig. 1.** Constructing dynamic SDR key tree

**Proposition 1.** *In the key tree  $T$  of static SDR, if there exists a set of disjoint subtrees  $\{T_i\}$  with the same height  $H$  such that*

- (1) *leaves of  $T = \bigcup_i \{\text{leaves of } T_i\}$ ;*
- (2) *each  $T_i$  has at least one leaf in  $\mathcal{N} \setminus \mathcal{M}$ , i.e., a member currently inactive, then all resultant subsets of  $\mathcal{M}$  have height  $h \leq H$  (Fig. 1(a)).*

Proposition 1 implies that, if such set of  $\{T_i\}$  exists in the key tree of static SDR, all nodes with height  $h > H$  will not contribute to the resultant subsets of  $\mathcal{M}$ . Thus the GCKS only needs to maintain a set of smaller subtrees satisfying Proposition 1 with an appropriate height  $H$  (we will address how to choose  $H$  shortly). The idea of dynamic SDR is to dynamically maintain such set of subtrees.

#### 3.1 Scheme of Dynamic SDR

In dynamic SDR, the positions of members are not pre-assigned. Instead, the spaces in the key tree are dynamically allocated and reclaimed, adapting to the current set of active members. More specifically, the GCKS dynamically creates

leaves when new member joins, or discards a subtree when all positions of the subtree are inactive. By doing this, the GCKS maintains active members in a dynamic key tree, rather than a large key tree constructed in advance.

For simplicity, we use set of subtrees  $\{T_{i_k}\}$  as allocation units, one can view the subtrees connected to a virtual root  $r$  (Fig. 1(b)). Initially, the GCKS has a single subtree  $T_{i_1}$  connected to the virtual root  $r$ . When a member joins the group, regardless of being a new member or a returned member, the member is assigned to the next available position in the key tree (from left to right) and is *unicast* the secret information associated with the new position. The GCKS thereafter encrypts and *multicasts* the updated TEK to the current members in exactly the same way as in static SDR. If new positions are required, the GCKS creates a new subtree  $T_{i_k}$ . When a member,  $m$ , leaves the group, the position becomes empty and will never be used by any member (even  $m$  itself). And a new TEK is multicast to the members that remain in the group. If all positions of the leftmost subtree become empty, the GCKS discards that subtree. This process is detailed in [4].

The advantages of maintaining such a dynamic-membership key tree are two-fold. First, dynamic SDR may require a much smaller key tree of a size sufficient to accommodate the maximum number of concurrently active members. This helps reduce key storage cost, both at the members and at the GCKS. Second, by assigning members that arrive close in time to positions that are close in the key tree, the GCKS is likely to find a subset that can cover many adjacent members. This implies that the messaging overhead associated with rekeying is also reduced. Dynamic SDR achieves the advantages by introducing additional unicast, corresponding to deliver  $I_m$  to a joining member  $m$ . However, the overall communication cost (in bytes per second), can be reduced by more than 50% in comparison to that of static SDR, as we will see in Section 4.

Since the key tree of dynamic SDR can be extended arbitrarily, dynamic SDR does not require *a priori* knowledge of the size of total member population,  $N$ . This avoids the problem, which exists in static SDR, of estimating  $N$ . Overestimating  $N$  makes the static SDR key tree unnecessarily large, increasing both rekey communication cost and key storage cost, whereas, underestimating  $N$  may introduce the problem of having to reject members when all positions have been assigned.

In dynamic SDR, however, the size,  $L = 2^H$ , of SDR subtree  $T_{i_k}$  should be chosen properly, as we describe next.

### 3.2 Determining $L$

The choosing of  $L$  is a design tradeoff. Based on Proposition 1, one subset covers at most  $L$  members. Therefore, when  $L$  is small, more resultant subsets are required to cover  $\mathcal{M}$ , which consequently increases the multicast cost. When  $L$  is large, we may still encounter the space inefficiency of static SDR that active members disperse in the subtree.



We thus choose  $L$  as a reasonable value of  $2^{\lceil \log E[M] \rceil}$ , where  $E[M]$  is the expected value of the number of concurrent members. Ideally, we want to put the concurrent members in one subtree.

### 3.3 Key Storage of Dynamic SDR

In dynamic SDR, the size of secret information,  $|I_m|$ , is reduced from  $(\log^2 N + \log N)/2 + 1$  to  $(\log^2 L + \log L)/2 + 1$  for the following Proposition.

**Proposition 2.** *In dynamic SDR, if a member,  $m$ , is assigned a position in subtree  $T_{i_k}$ , for any resultant subset  $S_{i,j}$  covering  $m$ ,  $i$  is a descendant of  $i_k$ .*

Although the key storage size required by a member is fixed when  $L$  is chosen, this is not the case for the GCKS, whose key storage is related to the number of subtrees  $T_{i_k}$ . Assuming that the GCKS sequentially assigns the available positions of the key tree from the left to the right to the joining members, we define  $S$  as the distance from the leftmost position occupied by an active member to the first available position at the right side. These  $S$  positions are referred as the *concurrent spaces*, which determine the key storage at the GCKS. To hold  $S$  concurrent spaces, at most  $\lceil S/L \rceil + 1$  subtrees are required. Since the GCKS has  $2L \log L + 1$  key storage for a subtree  $T_{i_k}$  of size  $L$  [3], it follows that the key storage at the GCKS is  $(\lceil S/L \rceil + 1)(2L \log L + 1)$ .

When members arrive according to a Poisson process with rate  $\lambda$  and the time that each active member stays in the group (which is referred as *lifetime*) is exponentially distributed with mean  $1/\mu$ , the expectation of  $S$  can be computed as follows [4], where  $\rho = \lambda/\mu$ :

$$E[S] = \rho e^{-\rho} \sum_{m=0}^{\infty} \frac{\rho^m}{m!} \sum_{i=1}^{m+1} \frac{1}{i} \quad (1)$$

Using Little's law, the average number of the concurrent members  $E[M] = \rho = \lambda/\mu$ , thus  $E[S]$  can be viewed as a function of  $E[M]$ , as shown in Fig. 2(a). The top curve in the figure presents  $E[S]$  as a function of  $E[M]$  according to (1), which shows that  $E[S]$  increases super-linearly with  $E[M]$ .


$E[S]$  also depends on the distribution of members' lifetime. If members' lifetime is deterministic (e.g., members are First-In-First-Out),  $E[S]$  is identical to  $E[M]$ . Generally, the higher variance members' lifetime has, the larger  $E[S]$  is.

A large value  $S$  results in an increased key storage at the GCKS side. Also, currently active members disperse in the key tree as  $S$  increases, incurring more resultant subsets and thus more rekeying messages. As a result, it is desirable to keep  $S$  small, a topic we address next.

### 3.4 Reducing $S$ by Shifting

In this subsection, we propose a simple operation, namely *shifting*, to reduce  $S$ .

We define *shifting* as the operation of detaching the leftmost active member in the key tree and re-attaching the member to the next available position (for



**Fig. 2.** Relationship between  $E[S]$  and  $E[M]$  (a); The shifting scheme (b)

new arrivals) in the key tree. When some holes (i.e., positions with departed members) are generated in the key tree, shifting the leftmost active member may reduce the concurrent spaces,  $S$ , and make active members more adjacent, as illustrated in Fig. 2(b).

When active members are shifted, they are delivered new secret information associated with the new positions by unicast. From the collusion-proof property of static SDR, we can show that shifting does not jeopardize the confidentiality [4]. There are many strategies for shifting. Here, we investigate two approaches, namely *probabilistic shifting* and *threshold-based shifting*.

**Probabilistic shifting.** Probabilistic shifting is defined as follows. When a member joins the group, that member is assigned to the next available position in the key tree. Meanwhile, with a probability  $p$ , the GCKS shifts the leftmost member to the position just to the right of the newly arrived member. Note that the shifting probability  $p$  affects the tradeoff between the shifting cost, associated with unicasting the secret information for the new position to the shifted member, and the value of  $S$ . In practice, the shifting probability  $p$  is a parameter set by the GCKS and can be application specific.

When members join the group according to a Poisson process with rate  $\lambda$  and a member's lifetime is exponentially distributed with mean  $1/\mu$ , we can compute the expected number of the concurrent spaces,  $E(S)$ , as a function of  $p$ , where  $\rho = \lambda/\mu$ :

$$E[S] = \rho e^{-\rho} \sum_{m=0}^{\infty} \frac{\rho^m}{m!} \left( \sum_{i=1}^{m+1} \frac{1}{i + p\rho} + p \sum_{i=1}^m \frac{1}{i + p\rho} \right) \quad (2)$$

Fig. 2(a) presents the result of  $E[S]$  for  $p = 0, 0.1$  and  $0.5$  respectively. We observe that  $E[S]$  decreases as a function of  $p$ . In particular, when the expected group size  $E[M] = 10^5$ , the space-member-ratio,  $E[S]/E[M]$  is 12.1 for dynamic SDR without shifting, and reduces to 2.6 when  $p = 0.1$ , and further reduces to 1.6 when  $p = 0.5$ . Thus, by introducing a small shifting probability  $p$ , dynamic SDR can effectively reduce the average concurrent space  $S$ , therefore reducing the key storage cost at the GCKS and potentially the rekey cost as well.

We next consider a different kind of shifting strategy – threshold-based shifting.

**Threshold-based shifting.** We define the occupancy ratio  $\gamma$  as the number of active group members to the number of concurrent spaces, i.e.,  $\gamma = M/S$ . Informally, the larger the occupancy ratio is, the more likely the members are adjacent to each other in the key tree, and thus can be covered by fewer subsets. To keep the rekey process efficient, the GCKS should keep the occupancy ratio high. A natural way to achieve this is to define a threshold  $\Gamma < 1$ ; when a member leaves the group, the GCKS computes the occupancy ratio  $\gamma$  and compares it to the threshold  $\Gamma$ . If the occupancy ratio falls below the threshold (i.e.,  $\gamma < \Gamma$ ), the GCKS will keep shifting the leftmost member until  $\gamma \geq \Gamma$ . We refer to this strategy as threshold-based shifting.

As with the shifting probability  $p$  in the probabilistic shifting scheme, the threshold  $\Gamma$  is also an application-specific parameter affecting the tradeoff between the shifting cost (unicast) and the rekey cost (multicast), as discussed in Section 4.2.

### 3.5 Block Alignment

So far, we have treated the newly arrived members and the shifted members identically when assigning a member to an available position. However, a member that has been in the group for a long time and has become the leftmost member in the key tree may have very different characteristics in terms of the remaining service time, than that of a member who just joined the group.

Given the above considerations, we further propose an enhancement to the dynamic SDR with shifting scheme by allocating different blocks with size  $B$  in the key tree for new members and shifted members. Detailed procedure of block alignment is left in [4] for page limitation.  $B$  is an application-specific parameter similar to  $p$  and  $\Gamma$ . We also evaluate the effects of such block alignment through simulation in Section 4.2.

## 4 Evaluation

In this section, we evaluate the performance of the dynamic SDR through simulation. We will first describe the simulation model and the performance metrics, then present the results.

### 4.1 Simulation Model and Performance Metrics

We assume a fixed  $N = 2^{17}$  to which a GCKS provides key management, i.e., the number of potential members is  $2^{17}$ . Each member independently decides to join or leave the group. We approximate the members' arrival by a fixed-rate Poisson process and assume that the lifetimes are iid random variables.



We evaluate different lifetime distribution functions, which include exponential, lognormal and uniform distribution.

In the following, we will only present results for immediate-rekey scheme since the performance of SDR is insensitive to the rekey period [3].

The performance metrics of interest are the *key storage cost* (both at the member side and at the GCKS side) and *rekey communication cost*. The key storage cost is measured as the key storage described in the previous section. More specifically, we assume to use 3DES for encryption and each label has 128-bit of storage.

The rekey communication cost is measured as the number of unit-size rekey messages (assuming one message contains one 128-bit key) per unit time, which is further divided into *multicast cost* and *unicast cost*.

The multicast cost equals to the minimum number of subsets used to cover the active members in the key tree. Since the GCKS is performing immediate rekeying, and when the system is in steady state, the rate at which members depart the group should equal to the rate that members join the group, we can compute the overall multicast cost  $C_M$  as  $C_M = 2\lambda\overline{N_{SD}}$ , where  $\lambda$  is the arrival rate and  $\overline{N_{SD}}$  is the average number of subsets that the GCKS uses for one TEK update.

The unicast cost includes the messages for delivering the secret information to a joining member or a shifted member. For static SDR scheme, the secret information is delivered to a user when that user joins the group for the first time. Since a member's position in the key tree is fixed, no additional unicast costs are incurred when the member returns to the group. To favor static SDR, we assume that the system has been running for long enough so that each member has received the secret information for its position. Thus, we count the unicast cost for static SDR as zero.

For dynamic SDR schemes, the unicast cost is computed as  $C_U = (\lambda + v)N_K$ , where  $v$  is the rate that members are shifted and  $N_K$  is the key storage at the member side and equals to  $(\lceil \log E[M] \rceil^2 + \lceil \log E[M] \rceil)/2 + 1$  based on the analysis result in Section 3.

The overall rekey communication cost of dynamic SDR is the weighted sum of the multicast cost and the unicast cost. In the rest of the evaluation, we treat the cost of unicasting a message the same as that of multicasting, even though the unicast cost should be much lower than the multicast cost with respect to the number of links that the message travels. Some more discussion on the relative weight of the unicast cost and multicast cost is included in Section 4.2.

In summary, the overall rekey communication cost of static SDR is

$$C_s = 2\lambda\overline{N_{SD}^s} \quad (3)$$

where  $\overline{N_{SD}^s}$  is the average number of subsets using static SDR for one TEK update.

The overall rekey communication cost of dynamic SDR is

$$C_d = 2\lambda\overline{N_{SD}^d} + (\lambda + v)N_K \quad (4)$$


where  $\overline{N_{SD}^d}$  is the average number of subsets using dynamic SDR for one TEK update.

## 4.2 Simulation Results

**Effects of group size.** We first compare the performance of static SDR and dynamic SDR for different group sizes.

With a fixed total population ( $N = 2^{17} = 131072$ ) and a fixed mean lifetime ( $1/\mu = 100$ ), we use different value of the arrival rate ( $\lambda$ ) to vary the expected group size ( $E[M] = \lambda/\mu$ ) from 100 to  $1.3 \times 10^5$ . For now, we only consider the dynamic SDR scheme without shifting.

Fig. 3(a) compares the key storage of static SDR and dynamic SDR for different group size with exponentially-distributed lifetimes. The top two curves represent the key storage cost at GCKS and the bottom two curves represent the key storage cost at member side. Since static SDR maintains a fixed key tree whose size is determined by the total member population, its key storage costs, at both GCKS and member side, are invariant with different group size. For dynamic SDR, however, the key storage increases when the expected group size increases. Compared to static SDR, the member-side key storage cost of dynamic SDR is consistently lower, since  $M$  is always smaller than  $N$ . However the GCKS key storage cost of dynamic SDR begins to exceed that of static SDR as  $E[M]$  increases. This is because the expected size of concurrent spaces  $E[S]$ , which includes both active members in the key tree and departed members in between, becomes larger than  $N$  when  $E[M]$  is significant ( $> 10\%$ ) compared to  $N$ .



**Fig. 3.** Key storage (a) and rekey communication cost (b) of static SDR and dynamic SDR for different group size

We next compare the rekey communication cost of static SDR,  $C_s$ , and dynamic SDR,  $C_d$ , as described in (3) and (4) respectively (Fig. 3(b)). We observe that the rekey communication cost of static SDR increases as  $E[M]$  increase, reaching a maximum when  $E[M]$  is about  $N/2$ , and then starts to decrease when  $E[M]$  gets close to  $N$ . This behavior matches well the reasoning in Section 2.

Furthermore, we observe that when  $E[M] < 2 \times 10^4$ , the rekey communication cost  $C_d$  is much lower, nearly half, compared to  $C_s$ . Only when  $E[M]$  is greater than  $N/2$ , does  $C_s$  outperform  $C_d$ .

We have seen the benefit of dynamic SDR when  $M \ll N$ . In fact, many practical applications have this property. For example, the MBone STS-71 session has  $M \approx 360$  while having  $N \approx 4000$  [1]. Another example is in pay-per-view service: the number of people watching a movie at the same time,  $M$ , is usually orders of magnitude smaller than the total number of people having cable TV,  $N$ .


In the next, we will focus on the scenario where  $M \ll N$  and evaluate the impact of shifting and block alignment.

**Table 1.** Parameters of the simulation configurations


| Config. ID             | Poisson arrival | Distr. of lifetime | Mean of lifetime | Var. of lifetime  |
|------------------------|-----------------|--------------------|------------------|-------------------|
| <i>EXP</i>             | $\lambda = 10$  | exponential        | 100              | $10^4$            |
| <i>UNI</i>             | $\lambda = 10$  | uniform            | 100              | $3.3 \times 10^3$ |
| <i>LOG<sub>l</sub></i> | $\lambda = 10$  | lognormal          | 100              | $10^3$            |
| <i>LOG<sub>m</sub></i> | $\lambda = 10$  | lognormal          | 100              | $10^4$            |
| <i>LOG<sub>h</sub></i> | $\lambda = 10$  | lognormal          | 100              | $10^5$            |

**Impact of shifting.** In this subsection, we study the performance of dynamic SDR with shifting. We consider the case where  $E[M] = 1000 (\ll N)$ . Table 1 shows the five different configurations that we use to obtain the simulation results. We simulate probabilistic shifting and threshold-based shifting strategies for each configuration.

Figure 4(a) shows the communication cost,  $C_d$ , of dynamic SDR with probabilistic shifting for the five different configurations. We observe that, without shifting ( $p = 0$ ),  $C_d$  tends to be higher when members lifetime is of high variance.



(a) Probabilistic shifting



(b) Threshold-based shifting

**Fig. 4.** Rekey communication cost of dynamic SDR combined shifting ( $E[M] = 10^3$ )

With shifting ( $p > 0$ ), we find two different kinds of behaviors among the five configurations. For configurations *EXP*, *LOG<sub>m</sub>* and *LOG<sub>h</sub>*, which have high

variance lifetime, increasing the shifting probability generally reduces  $C_d$ . This is because, with no shifting,  $S$  is large; while introducing shifting, although unicast cost is increased,  $S$  can be substantially reduced, which results in a reduced multicast cost.

For low variance configurations,  $UNI$  and  $LOG_l$ , when increasing the shifting rate, the communication cost  $C_d$  does not necessarily decrease. In these cases,  $S$  is close to  $M$  without shifting. Shifting cannot reduce  $S$  much. On the contrary, shifting might affect the distribution of members in the key tree, which may increase the number of subsets needed to cover the active members. As a result, the choice of optimal value of  $p$  depends on the lifetime distribution.

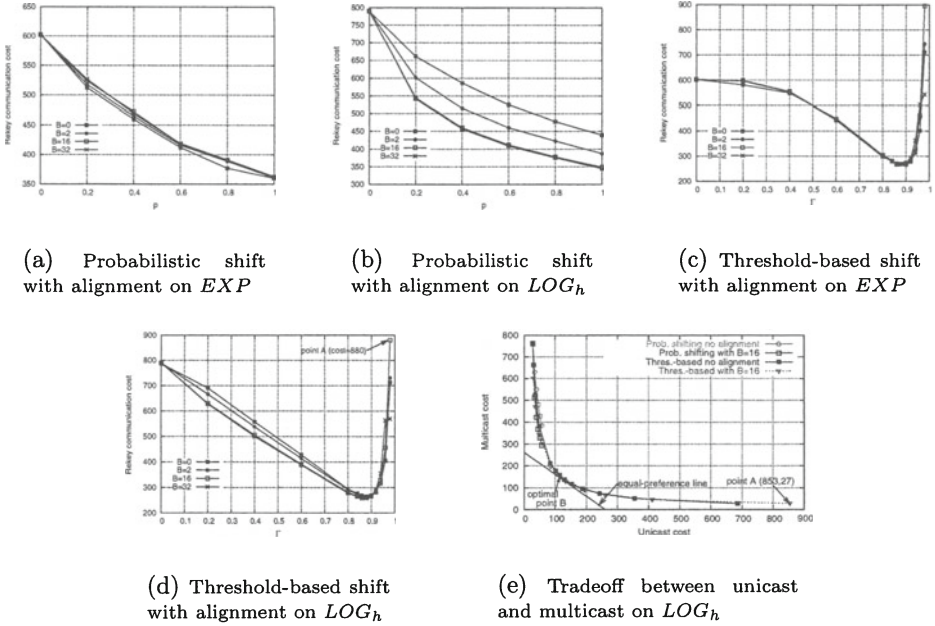
Fig. 4(b) shows the evaluation of threshold-based shifting scheme applied to the five configurations in Table 1. In the figure, it appears that, all configurations have a local minimum rekey communication cost when  $\Gamma \approx 0.9$ . When  $\Gamma$  goes beyond 0.9, there is a dramatic increase of the rekey cost. This is because, for a threshold  $\Gamma > 0.9$ , even though multicast cost may be reduced, unicast cost associated with frequent shifting becomes so high that the overall rekey cost is dramatically increased. Except for  $LOG_l$ , the local minimum rekey cost is also the global minimum. For configuration  $LOG_l$ , although the global minimum cost is achieved when  $\Gamma = 0$ , the local minimum rekey cost at  $\Gamma \approx 0.9$  is very close to the global minimum. Thus choosing a proper value of the threshold parameter  $\Gamma$  is not as sensitive to members' lifetime distribution as in probabilistic shifting.

**Enhancement with block alignment.** We next evaluate the performance of block alignment as an enhancement to probabilistic shifting and threshold-based shifting for the five simulation configurations. Here we present the results of configuration  $EXP$  and  $LOG_h$ .

Fig. 5 plots the rekey communication cost when applying block alignment with  $B = 0, 2, 16$  and  $32$  for probabilistic shifting (5(a) and 5(b)) and threshold-based shifting (5(c) and 5(d)). We observe that, for configuration  $EXP$ , introducing block alignment (with various  $B$ ) does not have much improvement on reducing rekey communication cost (sometimes is even worse). This is due to the memoryless property of exponential distribution. However, for configuration  $LOG_h$ , the improvement of block alignment is significant. Furthermore, for this particular group size, we find that increasing block size  $B$  beyond 16 does not provide much additional improvement.

From the figure, one can also see that, compared to threshold-based shifting, improvement of using block alignment is more evident in probabilistic shifting.

**Tradeoff between unicast cost and multicast cost.** As described in Section 3, dynamic SDR reduces multicast costs by introducing additional unicast, by which the secret information is delivered to shifted or returned members. Since the different strategies proposed in this paper have different parameters ( $p$  or  $\Gamma$ ) to configure, each of which reflects the tradeoff between unicast cost and multicast cost. To compare different schemes, we study the tradeoff graph of these proposed schemes as shown in Fig. 5(e). In the tradeoff graph, a point



**Fig. 5.** Costs and tradeoff of dynamic SDR using shifting and alignment ( $E[M] = 10^3$ )

on a curve denotes the multicast and unicast cost for the corresponding strategy with a particular parameter. For example, point A in Fig. 5(e) is associated with unicast cost of 853 and multicast cost of 27, denoting the total cost of 880 for the threshold-based shifting combined alignment of  $B=16$  with parameter  $\Gamma = 0.98$  (point A in Fig. 5(d)).

From the figure, we observe that reducing multicast cost comes at a cost of increasing unicast cost, and vice versa. The relative weight of unicast cost and multicast cost affects the choice of the optimal schemes and the operating parameters. If we treat unicast cost as expensive as multicast cost, in the tradeoff graph, all points on a line with a slope -1 are equally preferable. While points on a line close to point (0,0) are preferred over points on lines far away. In this sense, the threshold-based shifting combined alignment with block size  $B=16$  offers the best tradeoff among the algorithms considered, achieving an optimal value with multicast cost of 146 and unicast cost of 114 (point B in Fig. 5(e)). In general, if the relative weight of unicast cost and multicast cost is  $w$ , the equal-preference lines will have slope  $-w$  in the tradeoff graph. In this case, the best approach and the optimal parameters may be different.

## 5 Discussion

**Security.** In this paper, we use three criteria to measure security: forward confidentiality, backward confidentiality and collusion problem. As we know, static



SDR maintains forward/backward confidentiality, and has no collusion problem. We find that those properties hold for dynamic SDR. Scratch proof is given out in [4].

**Multicast Stateless.** Among all group key management algorithms, static SDR belongs to the so-called stateless algorithms in that members do not need to keep track of history of rekeying. In dynamic SDR, KEKs (keys of resultant subsets) are long-lived and can be computed based on the secret information securely unicast to members. Multicast messages in dynamic SDR distribute only TEKs. The multicast messages are not required to be reliably delivered to members so as to maintain their states correctly. However, in dynamic SDR, the GCKS is required to reliably unicast a member the secret information for the new position when the member joins or shifts. For this reason, we classify dynamic SDR as a multicast stateless algorithm.

## 6 Related Work

Most scalable centralized key-management algorithms make use of a tree structure to manage members. These algorithms could be broadly divided into *stateful* algorithms and *stateless* algorithms depending on whether members need to track the communication history to participate in the group communication. In stateful algorithms ([2,6,10,11]), the active members are leaves of the tree. While in stateless algorithms ([8]), the potential members are leaves of the tree.

LKH is a stateful algorithm. In an LKH tree, there is a leaf node corresponding to each active member. There is a key associated with each node in the tree, and each member holds a copy of every key on the path from its corresponding leaf node to the root of the tree. Hence, the key corresponding to the root node is shared by all members, and serves as the TEK.

Static SDR is a stateless algorithm. In a static SDR tree, there is a leaf node corresponding to each potential member. Subsets are defined through the tree, and each member holds subset keys for all subsets to which it belongs.

The performance of key-management algorithms is mostly determined by the positions of concurrent members in the tree. [5,12] propose methods to improve the performance of LKH by adjusting the positions of members dynamically so as to balance the key tree and reduce the overall height. Our work aims to improve the performance of SDR using the similar methodology – dynamically adjusting the positions of members.

A performance comparison between static SDR and LKH is given in [3]. Both the key storage and the rekey communication cost are compared in different scenarios, e.g. immediate rekeying, periodical batch rekeying and membership batch rekeying.

## 7 Conclusion

Static Subset Difference Revocation (SDR) is the current state of the art in stateless group rekeying algorithms. However, it works inefficiently when the number of the active members in the group is much less than the number of potential members, which is the case in many practical applications.

In this paper, we have proposed a group rekeying algorithm, dynamic SDR, which still keeps multicast stateless without the requirement of estimating the number of all potential members. By dynamically constructing the key tree, dynamic SDR uses a smaller key tree sufficiently large for the currently active members rather than the potential members. The smaller key tree reduces both the key storage cost and rekey communication cost compared to static SDR. We also introduce some enhancements to further improve the performance of dynamic SDR. Our evaluation shows that dynamic SDR significantly improves the performance of static SDR, reducing by half the rekey communication cost in the case that the number of the currently active members is much less than the total number of potential members. Also, compared to static SDR, dynamic SDR does not need to know the maximum number of potential group members in advance, a value that can be difficult to estimate in practice.

**Acknowledgments.** This research has been supported in part by the NSF under grant awards UF-EIES-0205003-UMA and EIA-0080119. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. The authors would also like to thank the anonymous reviewers for their helpful comments.

## References

1. Almeroth, K.C., Ammar, M.H.: Collecting and Modeling the Join/Leave Behavior of Multicast Group Members in the MBone. In: Proc. of the Symposium on High Performance Distributed Computing. Syracuse, NY (1996) 209–216
2. Chang, I., Engel, R., Kandlur, D., Pendarakis, D., Saha, D.: Key Management for Secure Internet Multicast using Boolean Function Minimization Techniques. In: Proc. IEEE INFOCOM 1999, Vol. 2. New York, NY (1999) 689–698
3. Chen, W., Dondeti, L.R.: Performance Comparison of Stateful and Stateless Group Rekeying Algorithms. In: Proc. Fourth International Workshop on Networked Group Communication. Boston, MA (2002)
4. Chen, W., Ge, Z., Zhang, C., Kurose, J., Towsley, D.: On Dynamic Subset Difference Revocation Scheme. Technical Report UM-CS-2003-22, University of Massachusetts at Amherst. (2003)
5. Dondeti, L.R., Mukherjee, S., Samal, A.: DISEC: A Distributed Framework for Scalable Secure Many-to-Many Communication. In: Proc. Fifth IEEE Symposium on Computers and Communications, Antibes-Juan les Pins, France (2000)
6. McGrew, D.A., Sherman, A.T.: Key Establishment in Large Dynamic Groups Using One-Way Function Trees. Technical Report No. 0755, TIS Labs at Network Associates, Inc., Glenwood, MD (1998)

7. Mittra, S.: Iolus: A Framework for Scalable Secure Multicasting. In: Proc. ACM SIGCOMM 1997, Cannes, France (1997) 277–288
8. Naor, D., Naor, M., Lotspiech, J.: Revocation and Tracing Schemes for Stateless Receivers. In: Proc. CRYPTO 2001, Lecture Notes in Computer Scienc, Vol. 2139, (2001) 41–62
9. Setia, S., Koussiah, S., Jajodia, S., Harder, E.: Kronos: A Scalable Rekeying Approach for Secure Multicast. In: Proc. IEEE Symposium on Security and Privacy, Berkeley, CA (2000)
10. Wallner, D., Harder, E., Agee, R.: Key Management for Multicast: Issues and Architectures. IETF Informational RFC, (1999)
11. Wong, C.K., Gouda, M.G., Lam, S.S.: Secure Group Communications Using Key Graphs. In: Proc. ACM SIGCOMM 1998, Vancouver, Canada (1998) 68–79
12. Yang, Y., Li, X., Zhang, X., Lam, S.: Reliable Group Rekeying: A Performance Analysis. In: Proc. ACM SIGCOMM 2001, San Diego, CA (2001) 27–38