J. Math. Crypt. **1** (2007), 373–384

DOI 10.1515 / JMC.2007.019

The average transmission overhead for broadcast encryption

Sarang Aravamuthan and Sachin Lodha

Communicated by Douglas R. Stinson

Abstract. We consider broadcast encryption schemes wherein a center needs to broadcast a secret message to a privileged set of receivers. We prescribe a probability distribution $\mathcal{P}$ on the privileged set. In this setting, the transmission overhead can be viewed as a random variable over $\mathcal{P}$ and we define its expected value as the *average transmission overhead* (or ATO).

Given $\mathcal{P}$, the Shannon's entropy function $H(.)$ provides a lower bound on the average number of bits required to identify every privileged set. This implies a natural lower bound for the ATO in terms of $H(\mathcal{P})$. For session key distribution, we consider the subset cover framework and bound the ATO in terms of the size of the cover. We further specialize our bound to accommodate storage constraints at receivers.

We consider two families of distributions for $\mathcal{P}$ that occur naturally in broadcast networks.

1. Each receiver independently joins the privileged set with probability p.
2. The privileged set is selected uniformly from a collection of subsets of receivers.

We evaluate the ATO of some practical schemes such as the subset difference method, the LSD scheme and the Partition-and-Power scheme under these distributions. Our investigations lead us to conclude that each scheme is inherently tailored to perform optimally for specific distributions.

Keywords. Broadcast encryption, Shannon's entropy.

AMS classification. 68P25, 94A17.

1 Introduction

The average transmission overhead is an important statistical measurement of the effectiveness of communication systems. It finds applications in diverse fields such as video transmission over wireless networks [13], management of replicated data [14] and classification of access networks [16]. The overhead is determined by prescribing a probability distribution on some parameter of the system and using standard methods from probability theory (see [1] for a good introduction) to compute its value.

The setting in a broadcast encryption scheme consists of a broadcaster $\mathcal{B}$ and a collection of receivers $\mathcal{R}$. Periodically, $\mathcal{B}$ will want only some *privileged* subset $R \subseteq \mathcal{R}$ of receivers to avail the broadcast. To ensure this, $\mathcal{B}$ first broadcasts a message $\mathcal{M}$ that only the receivers in R can decipher. The receivers are pre-configured with the right bits of information to achieve this. $\mathcal{M}$ includes a session key that is used for encrypting further broadcasts intended for R. The *transmission overhead* is the length of $\mathcal{M}$, that is, the amount of information (in bits) required to transfer the session key to R. Naturally, this overhead depends on R.

We now impose a probability distribution $\mathcal{P}$ on the privileged set. The distribution

is determined by the particular application the broadcaster intends to deploy. For instance, if the programs belong to the premium pay-per-view category, the privileged sets are more likely to be small and $\mathcal{P}$ would reflect this fact.

We assume that $\mathcal{P}$ holds true for every broadcast and is not affected by past broadcast history. Let us denote the probability of a subset $R \subseteq \mathcal{R}$ being privileged as $\mathcal{P}(R)$. The transmission overhead may then be viewed as a random variable over $\mathcal{P}$ whose expected value we define as the average transmission overhead (or ATO). In other words,

$$\text{ATO} = \sum_{R \subseteq \mathcal{R}} \mathcal{P}(R) \cdot \text{TO}(R),$$

where $\text{TO}(R)$ is the transmission overhead for R.

We observe that the Shannon's entropy function [15]

$$H(\mathcal{P}) = -\sum_{R \subseteq \mathcal{R}} \mathcal{P}(R) \log \mathcal{P}(R)^{1}$$

provides a lower bound on the average number of bits required to identify every privileged set. Since the broadcasts are *independent*, any broadcast encryption technique must provide for enough communication to allow identification of the privileged set. This observation translates into the following fact.

Fact 1.1.

$$\text{ATO} \quad \geq \quad H(\mathcal{P}).$$

In Section 2, this bound is improved to take into account the transfer of session key.

1.1 Related work

To the best of our knowledge, the concept of average transmission overhead has not been studied before in the broadcast encryption context. Broadcast encryption itself is a well-studied topic though; see for instance [7, 3, 8] for an introduction. In [3], a broadcast scheme is proposed using polynomial interpolation and related vector formulation methods. Fiat and Naor [8] consider k-resilient broadcast schemes where coalitions of k users not in the privileged set cannot recover the session key. Blundo et al study the trade-off between communication and storage in an information theoretic setting for unconditionally secure broadcast encryption schemes [4, 5]. Luby and Staddon [11] use a combinatorial model to study this trade-off.

Subset cover schemes were introduced in [12]. The authors describe a *subset difference method* to cover the set of non-revoked users by a collection of disjoint subsets. Each non-revoked user uses the key corresponding to his subset to recover the session key. If r is the number of revoked users and n the total number of users, their scheme

[1] In this paper, all logs are in base 2.

has a transmission overhead of $(2r-1)$ messages, each receiver stores $O(\log^2 n)$ keys and the processing time at a receiver is $O(\log n)$ operations. These bounds are improved by Halevy and Shamir [9] using a *Layered Subset Difference* (LSD) scheme. Specifically, they reduce the storage size to $O(\log^{3/2} n)$ while the processing time remains the same. However, the transmission overhead increases to $4r$ messages.

In [10, 6], the authors develop new methods to reduce transmission overheads in broadcast encryption. The methods are based on the idea of assigning one key per each partition using one-way key chains after partitioning the users. One method adopts skipping chains on partitions containing up to p revoked users and the other adopts cascade chains on partitions with layer structure. The scheme using the former reduces the transmission overhead down to $\frac{r}{p+1}$ messages asymptotically as r grows, and the scheme using the latter keeps the transmission overhead very small when r approaches 0, where r is the number of revoked users. Combining the two schemes, the authors propose a new broadcast encryption scheme with the least transmission overhead. However, these schemes require a large storage at a receiver and this storage may grow exponentially in some cases.

The Partition-and-Power scheme (PaP) was introduced in [2] wherein each user stores m keys and the subset cover is formed by partitioning the set of users into $n/(\log m + 1)$ equally sized sets and taking the union of all subsets of each set. A lower bound for the maximum transmission overhead for a subset cover scheme was determined and the PaP scheme was shown to achieve within a constant factor of this bound.

1.2 Our contribution

In this article, we assume that keys are distributed using the subset cover method. These offer several advantages over other key distribution methods. For instance, the receivers in this scheme are stateless, that is, once configured, the keys don't have to be updated. This scheme is also fully resilient in the sense that the session key cannot be computed by any coalition of non-receivers. The keys may either be explicitly stored as in the complete subtree method [12] and partition-and-power (PaP) scheme [2] or computed through pseudo-random generators as in the subset difference (SD) method [12] or the layered subset difference (LSD) method [9].

Our main result is a generic lower bound for the ATO in terms of the storage at a receiver, the distribution $\mathcal{P}$ and the key length (Theorem 2.3). Using this bound as the benchmark, we evaluate different subset cover schemes (SD, LSD and PaP) under two families of distributions encountered in practice. Our findings show that each scheme is inherently tailored to perform optimally for specific distributions.

Outline of the paper: In the next section, we introduce subset cover schemes and derive the expression for the ATO. We specialize this bound in terms of the storage at receivers when all subset keys are explicitly stored with the receiver. In Section 3, we consider some probability models for the privileged set that are natural and useful in practice, and compute the bounds for ATO under these distributions. In Section 4,

we consider some specific subset cover schemes (such as SD, LSD, PaP schemes) and estimate their ATO for these distributions. We conclude with Section 5.

2 The subset cover framework

We consider subset cover schemes introduced in [12]. The idea is to express the privileged set as a union of a collection of subsets from a cover and encrypt the session key with the subset key corresponding to each subset in the union. These encrypted values are broadcast. Each privileged user either stores the subset keys explicitly or is able to compute them using some pre-stored information. The subset key is then used to recover the session key from the encrypted message.

Let $\mathcal{R} = \{r_1, \ldots, r_n\}$ be a set of n receivers capable of receiving transmissions from a broadcaster $\mathcal{B}$. Let $\mathcal{S}, \mathcal{X} \subseteq 2^{\mathcal{R}}$ be collections of subsets of $\mathcal{R}$ defined in the following manner. $\mathcal{S}$ is the collection of privileged sets. In the most general case, $\mathcal{S}$ would comprise of all subsets of $\mathcal{R}$. $\mathcal{X}$ is defined to be a *cover* for $\mathcal{S}$, that is, every element in $\mathcal{S}$ can be expressed as a union of some elements from $\mathcal{X}$. Formally,

$$\forall R \in \mathcal{S},\ \exists X_R \subseteq \mathcal{X} \text{ such that } \bigcup_{X \in X_R} X = R.$$

Let

$$f_{\mathcal{X}}(R) = \min_{\substack{X_R \subseteq \mathcal{X}, \\ \bigcup_{X \in X_R} X = R}} |X_R|$$

be the minimum number of elements from $\mathcal{X}$ required to cover R.

For each $X \in \mathcal{X}$, there exists a (unique) *subset key* K_X that is known only to the receivers in X. To distribute a *session key* K to a set of receivers, say $R \in \mathcal{S}$, $\mathcal{B}$ first finds the smallest cover for R by the elements from $\mathcal{X}$. Suppose it is $R = X_1 \cup X_2 \cup \cdots \cup X_f$ where $f = f_{\mathcal{X}}(R)$. Then $\mathcal{B}$ sends the following broadcast

$$\langle [M_{X_1}, \ldots, M_{X_f}, E_{K_{X_1}}(K), \ldots, E_{K_{X_f}}(K)], E_K(M) \rangle, \tag{2.1}$$

where $E_{K_X}(K)$ is the encryption of K under K_X and $E_K(M)$ is the broadcast message M encrypted under K. Here the message M_{X_i} uniquely addresses X_i. Note that the quantity in square brackets in (2.1) is the transmission overhead.

Suppose all keys are t bits in length (so that $|\mathcal{X}| \leq 2^t$) and let $M_{\mathcal{X}}(R) = |M_{X_1}| + \cdots + |M_{X_f}|$. Let $\text{TO}_{\mathcal{X}}(R)$ be the transmission overhead to target R so that

$$\text{TO}_{\mathcal{X}}(R) \quad = \quad M_{\mathcal{X}}(R) + f_{\mathcal{X}}(R)t. \tag{2.2}$$

To define the average transmission overhead, we assume a probability distribution $\mathcal{P}$ on the privileged sets. Assume that a privileged set $R \in \mathcal{S}$, occurs with probability $0 < \mathcal{P}(R) \leq 1$ so that $\sum_{R \in \mathcal{S}} \mathcal{P}(R) = 1$. The average transmission overhead is then the quantity

$$\text{ATO}_{\mathcal{X}} \quad = \quad \sum_{R \in \mathcal{S}} \mathcal{P}(R) \cdot \text{TO}_{\mathcal{X}}(R).$$

We observe that $\text{ATO}_{\mathcal{X}}$ depends only on the cover and not on any specific privileged set. Since R is identified by $M_{\mathcal{X}}(R)$, using (2.2) and Fact 1.1, we get

$$\begin{aligned} \text{ATO}_{\mathcal{X}} &= \sum_{R \in \mathcal{S}} \mathcal{P}(R) \cdot M_{\mathcal{X}}(R) + \sum_{R \in \mathcal{S}} \mathcal{P}(R) \cdot f_{\mathcal{X}}(R)t \\ &\geq H(\mathcal{P}) + \mathrm{a}_{\mathcal{X}} t, \end{aligned} \tag{2.3}$$

where $\mathrm{a}_{\mathcal{X}} = \sum_{R \in \mathcal{S}} \mathcal{P}(R) f_{\mathcal{X}}(R)$ is the *average covering number* of $\mathcal{X}$. Let us now bound $\mathrm{a}_{\mathcal{X}}$ in terms of $H(\mathcal{P})$.

Claim 2.1.

$$\mathrm{a}_{\mathcal{X}} \log \frac{|\mathcal{X}|}{\mathrm{a}_{\mathcal{X}}} \geq H(\mathcal{P}).$$

Proof. There are $|\mathcal{X}|$ different elements of $\mathcal{X}$. Let's suppose that they are represented by numbers 1 to $|\mathcal{X}|$. We denote by N_X the number representing the set $X \in \mathcal{X}$.

Consider the case when $R \in \mathcal{S}$ is the privileged set whose minimum cover is $R = X_1 \cup X_2 \cup \ldots \cup X_f$ (where $f = f_{\mathcal{X}}(R)$), and $N_{X_0} = 0 < N_{X_1} < N_{X_2} < \cdots < N_{X_f} \leq |\mathcal{X}|$. Then, $\mathcal{B}$ sends $M_{X_i} = N_{X_i} - N_{X_{i-1}}$ for $1 \leq i \leq f$. Note that

$$N_{X_i} = \sum_{j=1}^{i} M_{X_j},$$

so it is easy for receivers to reconstruct the numbers N_{X_i}s and identify X_is. Here,

$$\begin{aligned} M_{\mathcal{X}}(R) &= |M_{X_1}| + |M_{X_2}| + \cdots + |M_{X_f}| \\ &= \log N_{X_1} + \log (N_{X_2} - N_{X_1}) + \cdots + \log (N_{X_f} - N_{X_{f-1}}) \\ &\leq f \log \frac{|\mathcal{X}|}{f} \quad \left(\text{Since } \prod_{i=1}^{f} r_i \leq \left(\frac{N}{f}\right)^f \text{ when } \sum_{i=1}^{f} r_i \leq N\right) \\ &= f_{\mathcal{X}}(R) \log |\mathcal{X}| - f_{\mathcal{X}}(R) \log f_{\mathcal{X}}(R). \end{aligned}$$

Now,

$$\begin{aligned} H(\mathcal{P}) &\leq \sum_{R \in \mathcal{S}} \mathcal{P}(R) M_{\mathcal{X}}(R) \\ &\leq \sum_{R \in \mathcal{S}} \mathcal{P}(R) f_{\mathcal{X}}(R) \log |\mathcal{X}| - \sum_{R \in \mathcal{S}} \mathcal{P}(R) f_{\mathcal{X}}(R) \log f_{\mathcal{X}}(R) \\ &= \mathrm{a}_{\mathcal{X}} \log |\mathcal{X}| - \sum_{R \in \mathcal{S}} \mathcal{P}(R) f_{\mathcal{X}}(R) \log f_{\mathcal{X}}(R) \\ &\leq \mathrm{a}_{\mathcal{X}} \log |\mathcal{X}| - \mathrm{a}_{\mathcal{X}} \log \mathrm{a}_{\mathcal{X}}. \end{aligned}$$

Note that we used the Jensen's inequality for convex function $\phi = x \log x$ in the final simplification. The Jensen's inequality says that $E[\phi(X)] \geq \phi(E[X])$ when ϕ is convex. □

From Claim 2.1, we get the following bound on $a_{\mathcal{X}}$:

$$a_{\mathcal{X}} \geq \frac{H(\mathcal{P})}{\log|\mathcal{X}| - \log a_{\mathcal{X}}} \geq \frac{H(\mathcal{P})}{\log|\mathcal{X}| - \log H(\mathcal{P}) + \log\log|\mathcal{X}|}.$$

Plugging this bound in equation (2.3), we get the following important lemma.

Lemma 2.2. *Under a probability distribution $\mathcal{P}$ over the privileged sets $\mathcal{S} \subseteq 2^{\mathcal{R}}$, with a subset key of length t bits, the average transmission overhead is bounded by*

$$ATO_{\mathcal{X}} \geq H(\mathcal{P})\left(1 + \frac{t}{\log|\mathcal{X}| - \log H(\mathcal{P}) + \log\log|\mathcal{X}|}\right) \tag{2.4}$$

for broadcast encryption schemes that use subset cover method as outlined by (2.1).

Suppose a receiver can store a maximum of m subset keys and we consider schemes where the subset keys are stored explicitly in the receiver. The cover size then satisfies $|\mathcal{X}| \leq mn$ and (2.4) specializes to the following result.

Theorem 2.3. *Under a probability distribution $\mathcal{P}$ over the privileged sets $\mathcal{S} \subseteq 2^{\mathcal{R}}$, with a storage bound of m keys per receiver and a subset key of length t bits, the average transmission overhead is bounded by*

$$ATO_{\mathcal{X}} \geq H(\mathcal{P})\left(1 + \frac{t}{\log mn - \log H(\mathcal{P}) + \log\log mn}\right) \tag{2.5}$$

for broadcast encryption schemes that use subset cover method as outlined by (2.1) and explicit storage of subset keys at the receiver.

Theorem 2.3 bounds the average transmission overhead in terms of the parameters of the broadcast network, such as the number of receivers, key length, receiver storage and the distribution of the privileged set. Thus a practical scheme can determine its effectiveness by measuring how far it deviates from this bound.

3 Some special probability distributions

We evaluate the lower bound expression in (2.5) for a couple of probability distributions that are natural and useful in practice.

3.1 Independent and identically distributed case

Consider the case where each receiver, independently, joins the privileged set with probability $0 \leq p \leq 1$. Therefore, the probability of $R \subseteq \mathcal{R}$ being privileged is

$$\mathcal{P}(R) = p^{|R|}(1-p)^{n-|R|}.$$

For such a distribution $\mathcal{P}_p$ over $2^{\mathcal{R}}$, its Shannon entropy is

$$H(\mathcal{P}_p) \quad = \quad nH_2(p),$$

where $H_2(p) = -p\log p - (1-p)\log(1-p)$ is the standard binary entropy function. This follows from the fact that the entropy of the joint distribution of independent events is the sum of the individual entropies of each distribution.

Thus, equation (2.5) translates into

$$\text{ATO}_{\mathcal{X}} \quad \geq \quad n\left(1 + \frac{t}{\log m + \log\log mn}\right) H_2(p). \tag{3.1}$$

This expression is symmetric around $p = 1/2$. It reaches its maximum value $n(1 + \frac{t}{\log m+\log\log mn})$ when $p = 1/2$ which corresponds to the case where every subset of $\mathcal{R}$ is equally likely to be a privileged set. It is interesting to note that this bound is comparable with the lower bound of [2] for the maximum transmission overhead in broadcast communication.

Note that equation (3.1) shows that the average transmission overhead is high for mid-range p. In fact, broadcast schemes addressing such distributions are hard to implement since all subsets of $\mathcal{R}$ are more or less equally likely to be privileged.

3.2 Equi-probable case

Next, we assume a uniform distribution over $\mathcal{S} \subseteq 2^{\mathcal{R}}$, that is, all privileged sets occur with the same probability (of $1/|\mathcal{S}|$). Let us label this probability distribution as $\mathcal{P}_{\mathcal{S}}$. In this case, its Shannon entropy is

$$\begin{aligned} H(\mathcal{P}_{\mathcal{S}}) \quad &= \quad -\sum_{R\in\mathcal{S}} \mathcal{P}(R)\log\mathcal{P}(R) = \sum_{R\in\mathcal{S}} \frac{\log|\mathcal{S}|}{|\mathcal{S}|} \\ &= \quad \log|\mathcal{S}|. \end{aligned}$$

Thus, equation (2.5) translates into

$$\text{ATO}_{\mathcal{X}} \quad \geq \quad \log|\mathcal{S}| \cdot \left(1 + \frac{t}{\log mn - \log\log|\mathcal{S}| + \log\log mn}\right). \tag{3.2}$$

When $\mathcal{S} = 2^{\mathcal{R}}$, $\mathcal{P}_{\mathcal{S}}$ is the uniform distribution over all subsets of $\mathcal{R}$, and the expression in (3.2) reaches its maximum of $n(1 + \frac{t}{\log m+\log\log mn})$ (as in the IID case when $p = 1/2$). The scenario of practical interest is the case where $\mathcal{S} = \{R \subseteq \mathcal{R} : l \leq |R| \leq u\}$, that is, privileged sets are restricted to a size in the interval $[l, u]$. Then,

$$|\mathcal{S}| = \binom{n}{l} + \cdots + \binom{n}{u} \geq \begin{cases} \binom{n}{n/2} & \text{if } l \leq n/2 \leq u \\ (u-l+1)\binom{n}{\frac{l+u}{2}} & \text{otherwise.} \end{cases}$$

This bounds $|\mathcal{S}|$ from below and hence the ATO in equation (3.2).

4 Analysis of some subset cover schemes

We consider the SD, LSD and the PaP schemes and evaluate their ATO values for distributions $\mathcal{P}_p$ and $\mathcal{P}_{\mathcal{S}}$ studied in Section 3.

4.1 The subset difference method

In the SD method [12], the storage per receiver is $m = O(\log^2 n)$ keys. The cover $\mathcal{X}$ is chosen in such a way that if R is a privileged set of size $n - r$, then $f_{\mathcal{X}}(R) \approx 1.25r$.

4.1.1 Independent and identically distributed case

The average covering number under distribution $\mathcal{P}_p$ is then

$$a_{\mathcal{X}} \approx \sum_{r=0}^{n} \binom{n}{r} 1.25rp^{n-r}(1-p)^r = 1.25n(1-p).$$

Since identification of the privileged set in this scheme requires a negligible number of bits compared to the key distribution, the ATO value is

$$\text{ATO}_{\mathcal{X}} \approx 1.25nt(1-p).$$

Consider the ratio of $\text{ATO}_{\mathcal{X}}$ to the minimum ATO. Using $m = \log^2 n$ and the lower bound provided by equation (3.1), we get

$$\begin{aligned} \frac{\text{ATO}_{\mathcal{X}}}{\text{min ATO}} &= O\left(\frac{nt(1-p)\log\log n}{nH_2(p)(t+\log\log n)}\right) \\ &= \frac{(1-p)}{H_2(p)} \cdot O(\log\log n). \end{aligned} \tag{4.1}$$

Note that we have used $2^t \geq |\mathcal{X}| \geq n$ while deriving (4.1). The ratio in (4.1) decreases with p. Note that the average size of a privileged set is np. Thus this scheme is effective when the privileged set is likely to be large.

4.1.2 Equi-probable case

The average covering number under distribution $\mathcal{P}_{\mathcal{S}}$ is

$$a_{\mathcal{X}} = \sum_{R \in \mathcal{S}} f_{\mathcal{X}}(R)\mathcal{P}(R) \approx \frac{1}{|S|} \sum_{R \in \mathcal{S}} 1.25r = 1.25z$$

where $z = \frac{1}{|\mathcal{S}|} \sum_{R \in \mathcal{S}} r$. Note that $(n - z)$ is the average size of a privileged set. Assuming again that the identification of the privileged set requires a negligible number of bits, the ATO is

$$\text{ATO}_{\mathcal{X}} \approx 1.25tz.$$

Using $m = \log^2 n$ and the lower bound for minimum ATO given by equation (3.2), we get

$$\begin{aligned}\frac{\text{ATO}_{\mathcal{X}}}{\text{min ATO}} &\approx \frac{1.25tz(\log mn - \log\log|\mathcal{S}| + \log\log mn)}{\log|\mathcal{S}| \cdot (\log mn - \log\log|\mathcal{S}| + \log\log mn + t)} \\ &\approx \frac{1.25tz(\log n - \log\log|\mathcal{S}| + \log\log n)}{\log|\mathcal{S}| \cdot (\log n - \log\log|\mathcal{S}| + \log\log n + t)} \\ &= z \cdot O(\frac{\log n - \log\log|\mathcal{S}| + \log\log n}{\log|\mathcal{S}|}) \end{aligned} \tag{4.2}$$

where we have used $t \geq \log n \geq \log m$ and $n \geq |\log \mathcal{S}|$. We now analyze this ratio for two important scenarios.

(1) $\mathcal{S} = 2^{\mathcal{R}}$**:** In this case, the expression (4.2) is $\Theta(\log\log n)$, that is, similar to the expression (4.1) for $p = 1/2$. Note that the average size of a privileged set is $\approx n/2$ here.

(2) $\mathcal{S} = \{R \mid R \subseteq \mathcal{R} \bigwedge |R| = n - r\}$**:** In this case, $|\mathcal{S}| = \binom{n}{r}$ and $z = r$.

- For $r \ll n$ (or when r is a constant), the expression (4.2) is $\Theta(1)$, that is, the performance of the SD scheme is within a constant factor of the minimum on the ATO measure.
- When r is large, say $r = cn$ (where $0 < c < 1$ is a constant), the expression (4.2) is $\frac{c}{H_2(c)} \cdot O(\log\log n)$, that is, similar to the expression (4.1). It increases with c. Thus this scheme is effective when the privileged sets are large in size.

4.2 The layered subset difference method

The LSD scheme [9] essentially improves the SD scheme of [12]. Specifically it reduces the storage per receiver to $m = \log^{3/2} n$ keys while the processing time remains the same. The subset cover $\mathcal{X}$ is little bigger - it satisfies $f_{\mathcal{X}}(R) \approx 2r$ where $|R| = n - r$.

Since m is poly-logarithmic (that is, $O(\log^c n)$) and $f_{\mathcal{X}}(R) = O(r)$ for both the SD and the LSD schemes, their ATO analyses are very similar. Therefore, we choose not to repeat it for the LSD scheme. We conclude that, like the SD scheme, the LSD scheme is again effective when the privileged set is likely to be large.

4.3 The PaP scheme

Let m be the number of subset keys that can be stored at a receiver. The PaP scheme [2] partitions $\mathcal{R}$ into groups of size $g := \log m + 1$. The subset cover $\mathcal{X}$ is the union of the power set of each of these groups. Thus, $|\mathcal{X}| = 2^g n/g \approx 2mn/\log m$. Each receiver stores all the keys corresponding to the subsets it belongs to in $\mathcal{X}$. To address

a privileged set R, R is expressed as a disjoint union of at most $f = n/g$ sets from $\mathcal{X}$. Thus for any distribution, the average cover number is always less than n/g.

4.3.1 Independent and identically distributed case

Suppose $R \subseteq \mathcal{R}$ is privileged with distribution $\mathcal{P}_p$. If $f_{\mathcal{X}}(R) = j$, then R contains at least one element from j groups and none from the remaining groups. The average cover number is given by

$$\begin{aligned} a_{\mathcal{X}} &= \sum_{j=1}^{f} j \binom{f}{j} (1-p)^{g(f-j)} \left(\sum_{i=1}^{g} \binom{g}{i} p^i (1-p)^{g-i} \right)^j \\ &= f(1-(1-p)^g) = \frac{n}{g}(1-(1-p)^g) \end{aligned}$$

after routine simplifications. As one would expect, $a_{\mathcal{X}}$ increases with p. Now, the ATO is given by

$$\text{ATO}_{\mathcal{X}} = a_{\mathcal{X}}(\log|\mathcal{X}| + t) = \frac{n(1-(1-p)^g)(\log \frac{n}{g} + g + t)}{g}.$$

Consider the ratio of $\text{ATO}_{\mathcal{X}}$ to the minimum ATO. Using the lower bound provided by equation (3.1), we get

$$\begin{aligned} \frac{\text{ATO}_{\mathcal{X}}}{\text{min ATO}} &\leq \frac{n(1-(1-p)^g)(\log \frac{n}{g} + g + t)}{g} \cdot \frac{\log m + \log\log mn}{nH_2(p)(t + \log m + \log\log mn)} \\ &\leq \frac{2}{H_2(p)} \cdot (1 + \frac{\log\log mn}{\log m}). \end{aligned} \tag{4.3}$$

Note that we have used $g = \log 2m \geq 2$, and $2^t \geq |\mathcal{X}| \geq n$ while deriving (4.3). There are two distinct terms in (4.3): the first term $\frac{2}{H_2(p)}$ underlines the probabilistic setup whereas the second term $(1 + \frac{\log\log mn}{\log m})$ represents the storage constraint. We now analyze this ratio for different values of p.

(1) $p = c/n$ (or $p = 1 - c/n$) where c is a constant: In this case, the expression (4.3) is $\Theta(n)$. Thus, the performance of PaP scheme could be arbitrarily bad.

(2) $p = 1/2$ (or any constant): In this case, the expression (4.3) is $\Theta(1 + \frac{\log\log mn}{\log m})$. If m is fixed, and n is varied, then here again the performance of PaP scheme could be arbitrarily bad.

If $mn < 2^m$ (which is a likely scenario in practice), then the expression (4.3) is $O(1)$. Thus, the ATO of PaP scheme is within a constant factor of the optimal. For example, if $n = 2^{20}, m = 2^{10}$, and $p = 1/2$, then the performance of PaP scheme is within a factor of 3 of the minimum on the ATO measure.

4.3.2 Equi-probable case

Since $\mathcal{S}$ can be any sub-collection, the best we can do is bound the average cover number from above by $f = n/g$. Thus, the ATO of PaP scheme is

$$\text{ATO}_{\mathcal{X}} \leq 2tf = \frac{2tn}{\log 2m}.$$

Comparing with the lower bound given by (3.2), we observe that when $\log |\mathcal{S}| \approx n$, the ATO is within a $\Theta(1 + \frac{\log \log mn}{\log m})$ factor of the minimum. When $|\mathcal{S}|$ is small, the PaP scheme is not effective.

5 Conclusion

The ATO is a statistical measure of the efficacy of a transmission scheme and finds applications in different domains [13, 14, 16]. Thus studying this concept in the field of broadcast encryption is a natural extension. The ATO can be used as a yardstick to measure the effectiveness of a broadcast scheme and to compare different schemes under uniform conditions such as same storage at receivers, same key length and identical distributions.

In this article, we studied the subset cover framework and derived a lower bound for the ATO under a given distribution on the privileged set. Our analysis used the entropy of the distribution to bound the number of bits required to identify any privileged set. We specialized the lower bound when the storage on a receiver was constrained. We considered two probability distributions and evaluated the ATO of some subset cover schemes under these distributions. Our studies show that each scheme is naturally designed to give optimal performance for specific distributions.

Other distributions for the privileged set are possible. For instance, we have assumed that every receiver is privileged with the same probability p. A natural extension is to consider the scenario where different groups of receivers are privileged with different probabilities. The most general case would prescribe a probability for each privileged set. However, such distributions appear harder to analyze.

References

[1] N. Alon and J. Spencer, *The Probabilistic Method*. John Wiley & Sons, Inc., 2000.

[2] S. Aravamuthan and S. Lodha, *An Optimal Subset Cover for Broadcast Encryption*. In Proceedings of the 6th International Conference on Cryptology in India (INDOCRYPT), Lecture Notes in Computer Science 3797, pp. 221–231. Springer, Berlin, New York, 2005.

[3] S. Berkovits, *How to Broadcast a Secret*. Advances in Cryptology – Eurocrypt 1991, Lecture Notes in Computer Science 547, pp. 536–541. Springer, Berlin, New York, 1991.

[4] C. Blundo and A. Cresti, *Space Requirements for Broadcast Encryption*. Advances in Cryptology – Eurocrypt 1994, Lecture Notes in Computer Science 950, pp. 287–298. Springer, Berlin, New York, 1994.

[5] C. Blundo, L. A. Frota Mattos, and D. R. Stinson, *Tradeoffs Between Communication and Storage in Unconditionally Secure Schemes for Broadcast Encryption and Interactive Key Distribution*. Advances in Cryptology – CRYPTO 1996, Lecture Notes in Computer Science 1109, pp. 387–400. Springer, Berlin, New York, 1996.

[6] J. H. Cheon, N. S. Jho, M. H. Kim, and E. S. Yoo, *Skipping Cascade, and Combined Chain Schemes for Broadcast Encrypyion*, Cryptology ePrint Archive (2005). Available at http://eprint.iacr.org/2005/136.

[7] G. H. Chiou and W. T. Chen, *Secure Broadcasting Using the Secure Lock*, IEEE Transactions on Software Engineering SE-15 (1989), pp. 929–934.

[8] A. Fiat and M. Naor, *Broadcast Encryption*. Advances in Cryptology – CRYPTO 1993, Lecture Notes in Computer Science 773, pp. 480–491. Springer, Berlin, New York, 1993.

[9] D. Halevy and A. Shamir, *The LSD Broadcast Encryption Scheme*. Advances in Cryptology – CRYPTO 2002, Lecture Notes in Computer Science 2442, pp. 47–60. Springer, Berlin, New York, 2002.

[10] N. S. Jho, J. Y. Hwang, J. H. Cheon, M. H. Kim, D. H. Lee, and E. S. Yoo, *One-way Chain Based Broadcast Encrypyion Schemes*. Advances in Cryptology – EUROCRYPT 2005, Lecture Notes in Computer Science 3494, pp. 559–574. Springer, Berlin, New York, 2005.

[11] M. Luby and J. Staddon, *Combinatorial Bounds for Broadcast Encryption*. Advances in Cryptology – EUROCRYPT 1998, Lecture Notes in Computer Science 1403, pp. 512–526. Springer, Berlin, New York, 1998.

[12] D. Naor, M. Naor, and J. Lotspiech, *Revocation and Tracing Schemes for Stateless Receivers*. Advances in Cryptology – CRYPTO 2001, Lecture Notes in Computer Science 2139, pp. 41–62. Springer, Berlin, New York, 2001.

[13] D. Qiao and K. G. Shin, *A Two–Step Adaptive Error Recovery Scheme for Video Transmission over Wireless Networks*. In Proceedings of IEEE INFOCOM, pp. 1698–1704, 2000.

[14] D. Saha, S. Rangarajan, and S. Tripathi, *An Analysis of the Average Message Overhead in Replica Control Protocols*, IEEE Transactions on Parallel and Distributed Systems 7 (1996), pp. 1026–1034.

[15] C. E. Shannon, *A Mathematical Theory of Communication*, Bell Systems Technical Journal 27 (1948), pp. 379–423.

[16] W. Wei, B. Wang, C. Zhang, J. Kurose, and D. Towsley, *Classification of Access Network Types: Ethernet, Wireless LAN, ADSL, Cable Modem or Dialup?*. In Proceedings of IEEE INFOCOM, pp. 1060–1071, 2005.

Received 29 January, 2007; revised 10 July, 2007

Author information

Sarang Aravamuthan, Google, India.
Email: sarang@google.com

Sachin Lodha, Tata Consultancy Services, India.
Email: sachin.lodha@tcs.com