Disenrollment with Perfect Forward Secrecy in Threshold Schemes

Mingyan Li and Radha Poovendran

Abstract

In this paper, we propose a new model for threshold schemes with disenrollment capability (TSDC), to address the scenarios in which the ability of a coalition to construct future shared secrets is prohibited. Compared to existing TSDC models, our model provides forward secrecy by adding a constraint that the broadcast from the dealer is required to activate the reconstruction of any secret. We also present a TSDC model in which the dealer has the enhanced capability of disenrolling any subset of participants, to prevent the otherwise unnecessary rekey of the entire group when a large number of participants are compromised. We establish the lower bounds on the entropy of broadcast messages in both proposed models, as guidelines on constructing broadcast efficient schemes, and present bound achieving schemes.

Keywords: secret sharing, threshold schemes with disenrollment, entropy, broadcast, collusion

I. Introduction

A (t, n) threshold scheme is a technique to split and distribute a secret among a group of n participants (or members) in such a way that any subset of t or more participants can reconstruct the shared secret by pooling the information they have, while any subset of participants of cardinality less than t is unable to recover the secret [1], [2]. The information held by a participant is called a *share*, which is distributed securely by a trusted third party, called *dealer* or *owner*, to the participant at the time of initialization.

A cryptographic key is one of the secrets that can be protected using a threshold scheme. Since keys have finite lifetime, knowing the duration of the key use, one can generate multiple keys with key K_i being used in time period i. Each of these keys is then shared by n members using

Authors are with Department of Electrical Engineering, University of Washington, (e-mail: myli, radha@ee.washington.edu).

Preliminary version of this work was presented at the Selected Areas in Cryptography 2003.

a (t, n) threshold scheme, with only a valid set of t or more members being able to construct the key K_i at the time period i.

It is preferable, that should it be necessary to delete one or more members or if a compromise of all keys up to period i occurs, these events will not lead to a compromise of any keys intended to be used in future time periods. This is called the *perfect forward secrecy* or simply *forward secrecy* property [3]. We intend to study the problem of revoking untrustworthy members in threshold schemes while preserving forward secrecy.

When one member is compromised or deleted, the share held by the member should be treated as compromised and known to the public. Then, any $(t-1)$ shares from the valid group members along with the publicly known share suffice to reconstruct the shared key uniquely. Hence, the effective threshold of the secret sharing scheme reduces to $(t-1)$, and the key used at that time period is protected by the reduced number of shares.

To preserve the same level of security, the same threshold size t needs to be maintained and the shared key needs to be updated. If a secure channel is available all the time, a dealer can choose a new shared secret, construct a $(t, n-1)$ threshold scheme, and deliver new shares securely to the remaining participants. However, an expensive secure channel is normally set up only to distribute initial shares and is no longer available after initialization. Hence, the goal is to develop methods that use a public broadcast channel to update shares in order to maintaining the threshold level t.

The problem of *maintaining the threshold via only broadcast* in case of share disclosure or loss was considered in a seminal paper by Blakley, Blakley, Chan and Massey in [4], and the model presented was called a *threshold scheme with disenrollment capability* (TSDC). Blakley *et al.* formally defined *threshold schemes with L-fold disenrollment capability* as TSDC schemes that have the capability of disenrolling L participants successively, one at a time, without reducing the threshold t. However, their TSDC model addresses an environment where members are deemed benign and no coalition of colluders arise to derive future shared secret or such a collusive coalition is harmless. When such a coalition is undesired, forward secrecy needs to be enforced to ensure no future shared secret can be derived under collusion of members. We present a TSDC model in this paper to address the problem of *disenrolling participants with forward secrecy in threshold schemes* while maintaining the threshold through broadcast.

A TSDC scheme is characterized by the size of a share held by a member and the size of the broadcast the owner has to use when a member is disenrolled. The share size in any threshold scheme is lower bounded by the entropy of the key it is derived from [5]. In [4], Blakley *et al.* established a lower bound on the size of total number of shares in a TSDC with L-fold disenrollment capability.

Broadcast size in a TSDC indicates the communication cost to the dealer under a single member deletion. In [4], the authors conjectured a lower bound on the size of public broadcast required for a TSDC. In [6], Barwick *et al.* presented a revised version of the lower bound on the broadcast size for a TSDC with L-fold disenrollment capability.

Our contributions: We make four contributions in this paper.

- We first show that the original TSDC model [4] does not guarantee forward secrecy under collusion by showing that if any $(t+j), (j > i)$ members collaborate in one TSDC scheme, they can construct keys $\{K_1, K_2, \cdots, K_i, \cdots, K_j\}$.
- We present an enhancement to the TSDC model to incorporate the condition that the broadcast from the dealer is required to construct the key K_i at a time period i by t or more valid members, in order to prevent the exposure of future shared secrets to collusion of participants. The idea of activating the reconstruction of the shared secret by additional information, rather than the collection of shares only, was originally proposed in Simmons' "prepositioned" threshold schemes (PTS) [7]. However, no dynamic disenrollment is involved in prepositioned schemes, and no theoretical study on characterizing such threshold schemes was presented [7].
- We provide an enhancement to the dealer's ability to disenroll multiple members at a time. In the original TSDC [4] with L disenrollment stages, the owner is capable of deleting one member at each stage and hence total L members. However, observing that any subset of members can collaborate and may have to be disenrolled, the owner needs to have the capability to delete any subset of members at a time period, as long as the number of the remaining valid members is greater than t. We propose a new TSDC model that allows deletion of an arbitrary number of members at any time period, and maximum number of disenrolled members can be more than the total number of disenrollment stages.
- We derive the necessary lower bounds on the entropy of broadcast messages in our enhanced models. The bounds serve as guidelines in the search of broadcast efficient TSDC.

Other related work on TSDC includes [8], [9], [10], [11], [12]. In [8], disenrollment of untrustworthy participants in general secret sharing schemes that have different subsets of members allowed to construct a secret was discussed. However, no analytical results on the bounds of share size or broadcast size was established. Blundo *et al.* [9] presented the generalized problem of enabling participants of different sets to reconstruct different secrets at different times via an insecure public channel, and established lower bounds on share size. However, they did not investigate the lower bound on the broadcast size for a TSDC. Charnes *et al.* [10] presented a computationally secure TSDC based on an exponentiation version of Shamir threshold scheme, and the secrecy of the scheme is built on the hardness of the discrete logarithm problem. A TSDC scheme was proposed to realize disenrollment without changing shares in [11], but was later shown to be unable to maintain the threshold in [12].

This paper is organized as follows. In Section II, we review the definition of TSDC [4] and previous results on quantifying share size [4] and broadcast size [6]. We then formulate a TSDC model with forward secrecy in Section III, and we derive lower bounds on the broadcast size in such a model in Section IV. In Section V, we present a TSDC model with enhanced capability of deleting multiple participants simultaneously. Section VI concludes the paper.

II. Preliminaries

For clarity of presentation, the notations used in this paper are listed in Table I in Appendix A.

A. Threshold Schemes

Let K be a shared secret that is a random variable taking values from space $\mathcal{K}$, and S_j be a share held by participant j, that is a random variable taking values from space $\mathcal{S}$, and $j \in \mathcal{N} = \{1, .., n\}$.

Definition 1: A (t, n) threshold scheme is to share a secret K among n participants so that:

1) The secret K is recoverable from at least t shares. That is, for any set of k $(t \leq k \leq n)$ indices $\{l_1, l_2, ..., l_k\} \subset \mathcal{N}$, a collection of k shares $S_{l_1:l_k}$ can reconstruct the secret K, $H(K|S_{l_1:l_k}) = 0$.
2) The secret K remains uncertain with the knowledge of $(t-1)$ or less shares. That is, $H(K|S_{l_1:l_k}) > 0$ for $k < t$.

A (t, n) threshold scheme is called *perfect* in an information theoretic sense if $(t-1)$ or fewer shares reveal absolutely no information on the secret K. That is, $H(K|S_{l_1:l_k}) = H(K)$ for $k < t$. A necessary condition for a perfect threshold scheme is given in [5] as:

$$H(S_j) \geq H(K) \qquad \text{for} \qquad j = 1, ..., n. \tag{1}$$

A perfect threshold scheme is called *ideal* if share size achieves the lower bound in (1), i.e., $H(S_j) = H(K), \forall j \in \mathcal{N}$.

B. Threshold Schemes with Disenrollment Capability (TSDC)

A threshold scheme with L-fold disenrollment capability (TSDC) addresses the problem of maintaining a constant threshold through an insecure broadcast channel while disenrolling an untrustworthy participant at each of L successive updates [4]. Let $i = 1, ..., L$ be the indices of update stages. At the i^{th} update, let K_i denote the shared secret, P_i denote the broadcast message, $V_i(k)$ denote the set of indices of k valid participants, and $\tilde{\mathcal{D}}_i$ denote the set of indices of all disenrolled participants up to and including stage i. Let $d_i \ \in \mathcal{N} \backslash \tilde{\mathcal{D}}_{i-1}$ be the index of the disenrolled participant at the i^{th} update, and $v_l \in \mathcal{N} \backslash \tilde{\mathcal{D}}_i$ for $l = 1, ..., n - i$ be the index of one of the remaining valid participants.

Definition 2: A (t, n) threshold scheme with L-fold disenrollment capability with $n - L \geq t$ is a collection of shares S_j for $j = 1, ..., n$; shared secrets K_i for $i = 0, ..., L$; and public broadcast messages P_i for $i = 1, ..., L$, that satisfy the following conditions:

1) Initially (i.e., at stage $i = 0$), the scheme is a (t, n) threshold scheme that shares the secret K_0 among n participants.
2) At stage i for $i = 1, ..., L$, one additional share S_{d_i} is disenrolled. Any set of k ($t \leq k \leq n-i$) valid shares $S_{V_i(k)}$ plus the broadcast messages $P_1, ..., P_i$ can reconstruct the new secret K, leading to
$$H(K_i | S_{V_i(k)}, P_{1:i}) = 0 \quad \text{for} \quad t \leq k \leq n - i. \tag{2}$$
3) At stage i for $i = 1, ..., L$, given broadcast information $P_1, ..., P_i$ and all disenrolled shares $S_{d_1}, ..., S_{d_i}$, the shared secret K_i remains uncertain if the number of valid shares is less than t. That is,
$$H(K_i | S_{V_i(k)}, S_{d_1:d_i}, P_{1:i}) > 0 \quad \text{for} \quad k < t. \tag{3}$$

A (t, n) threshold scheme with L-fold disenrollment capability is called *perfect* if

$$H(K_i | S_{V_i(k)}, S_{d_1:d_i}, P_{1:i}) = H(K_i) \quad \text{for} \quad k < t. \tag{4}$$

A perfect TSDC is denoted as PTSDC in the rest of the paper.

In order to be able to collectively reconstruct K_i, each participant must have a component in his share corresponding to K_i. Let $S_j^{(i)}$ denote such a component and we call $S_j^{(i)}$ a *subshare* of participant j. Note that S_j is a collection of components to reconstruct $K_0, ..., K_L$, i.e., S_j is the union of all its subshares over L disenrollment stages, denoted by $S_j = \{S_j^{(0)}, S_j^{(1)}, ..., S_j^{(L)}\}$. Without component $S_j^{(i)}$, K_i cannot be reconstructed even if all the rest subshares $S_j^{(-i)} = \{S_j^{(0)}, S_j^{(1)}, ..S_j^{(i-1)}, S_j^{(i+1}, ..., S_j^{(L)}\}$ are collected from more than t valid participants. Hence, subshares satisfy

$$H(K_i|S_{V_i(k)}^{(i)}, P_{1:i}) \quad = \quad 0 \quad \text{for} \quad t \leq k \leq n-i; \tag{5}$$

$$H(K_i|S_{V_i(k)}^{(-i)}, P_{1:i}) \quad = \quad m \quad \text{for} \quad t \leq k \leq n-i; \tag{6}$$

where $S_{V_i(k)}^{(-i)}$ denotes a collection of $S_j^{(-i)}$'s with $j \in V_i(k)$.

The necessary condition (1) can be extended to subshares as

$$H(S_j^{(i)}) \geq H(K_i) \quad \text{for} \quad j = 1, .., n \quad i = 0, ..., L. \tag{7}$$

C. Bounds of Share and Broadcast Size in TSDC

Blakley *et al.* [4] established a lower bound on the entropy of each share as follows.

Theorem 1: (presented as Theorem 4 in [4]) Let $S_{1:n}, P_{1:L}, K_{0:L}$ form a (t, n) PTSDC and $H(K_i) = m$ for $i = 0, 1, ..., L$. Then,

$$H(S_j) \geq (L+1)m \qquad \text{for} \quad j = 1, 2, ..., n. \tag{8}$$

A PTSDC in which each share achieves its lower bound is called *share minimal* [6], i.e.,

$$H(S_j) = (L+1)m, \quad \text{for} \quad j = 1, ..., n, \tag{9}$$

where $H(K_i) = m$ for $i = 0, ..., L$. For a share minimal PTSDC, we showed each subshare achieves its minimal size [13], i.e., equation (7) holds with equality for $i = 0, 1, ..., L$ and $\forall j \in \mathcal{N}$. Using union bound based arguments in [14], it follows that the individual subshares of a member are independent in a share minimal PTSDC.

In [4], Blakley *et al.* also conjectured on the lower bound of the entropy of broadcast. A modified version of the conjecture was proven by Barwick *et al.* [6] as follows.

Theorem 2: (presented as Theorem 2 in [6]) Let $S_{1:n}, P_{1:L}, K_{0:L}$ form a (t, n) share minimal PTSDC satisfying properties (2), (4), and (9), then

$$\sum_{l=1}^{i} H(P_l) \geq \sum_{l=1}^{i} \min(i, n-i-t+1)m \quad \text{for} \quad i = 1, ..., L. \tag{10}$$

III. Adding Forward Secrecy

In this section, we reexamine the TSDC model and present a scheme from [4] that satisfies Definition 2 but allows any $(t+i)$ colluders to construct any keys $K_1, K_2, \cdots, K_i$.

Scheme 1: (Presented in [4])

- Participant j holds the share $S_j = \{S_j^{(0)}, S_j^{(1)}, ..., S_j^{(L)}\}$, where subshare $S_j^{(i)}$ corresponds to a share of a $(t+i, n)$ ideal PTSDC sharing secret K_i.
- At stage i, the dealer broadcasts $P_i = \{S_{d_1}^{(i)}, S_{d_2}^{(i)}, ..., S_{d_i}^{(i)}\}$.

Scheme 1 satisfies Definition 2 presented in Section II-B and thus is a valid TSDC. From Scheme 1, if $t+i < L$ participants collude and the dealer knows the identities of the colluding participants, then the dealer can revoke every colluding participant[1] by discarding intermediate subshares $S_j^{(l)}$'s for $l = 1, ..., (t+i)$ of *all* the group members, and broadcasting the subshares of the disenrolled members $\{S_{d_1}^{(t+i)}, S_{d_2}^{(t+i)}, ..., S_{d_{t+i}}^{(t+i)}\}$, where d_l's are indices of colluding participants. The new shared secret among the group becomes K_{t+i} and is reconstructible by t shares from remaining valid members of the group.

However, the problem lies in the observation that if $(t+i)$ participants were to collaborate by exchanging their shares, they can construct keys $K_1, ..., K_i$ regardless of the current time period. The dealer has no mechanism to prevent this from happening.

A. Preserving Forward Secrecy

To prevent any colluding parties from obtaining future keys, i.e., to provide forward secrecy, we seek a model in which not only are a threshold number of shares needed, but side channel information from the dealer is also needed to reconstruct the key K_i at time period i. Since the dealer has only the broadcast channel available after the group initialization, we enhance the TSDC model by incorporating the condition that broadcast P_i from the dealer is necessary in reconstructing the secret K_i. The reconstruction of K_i should not be possible without P_i even if all the shares are available. This can be formalized in terms of the mutual information as the following condition to ensure *forward secrecy*:

$$I(K_i; S_{1:n}, P_{1:i-1}) = 0 \quad \text{for} \quad i = 1...L. \tag{11}$$

[1] We assume that the dealer is able to identify all colluding participants, the mechanism of identification of compromised participants is out of the scope of this paper.

Condition (11) states that the mutual information of K_i and all shares S_j for $j = 1...n$ and all previous broadcast message $P_1, ..., P_{i-1}$ is zero. By jointly considering (2) and (11), we note that (11) expresses the importance of broadcast message P_i at stage i. Without the broadcast P_i, the new shared secret K_i cannot be derived even if all shares S_j and all previous broadcast messages $P_1, ..., P_{i-1}$ are known. Hence, forward secrecy is preserved in TSDC.

We now present Scheme 2 that satisfies (11) and ensures forward secrecy by preventing a set of $(t + i)$ colluders from reconstructing keys $K_1, ..., K_i$.

Scheme 2: (A TSDC with forward secrecy)

- The dealer generates the secrets K_i's to be shared, and random strings R_i of length $m = H(K_i)$, computes $(R_i + K_i)$ for $i = 1, ..., L$.
- The dealer generates all subshares $S_j^{(i)}$ with $j = 1, ..., L$ and distributes $S_j = \{S_j^{(0)}, S_j^{(1)}, ..., S_j^{(L)}\}$ to participant j, where $S_j^{(i)}$ is a share of a $(t+i, n)$ ideal perfect threshold scheme sharing $K_i + R_i$.
- At update i, the dealer broadcasts $P_i = \{R_i, S_{d_1}^{(i)}, ..., S_{d_i}^{(i)}\}$. With the broadcast, a set of t shares from valid participants suffices to recover $K_i + R_i$ and thus to decipher K_i using R_i.

Note that Scheme 2 satisfies forward secrecy condition (11) since P_i is needed to remove the effect of R_i from the reconstructed secret $K_i + R_i$. If $(t + i)$ participants collude after initialization, the dealer can send the broadcast message $\{S_{d_1}^{(t+i)}, ..., S_{d_{t+i}}^{(t+i)}, K_{t+i}+R_{t+i}+K_1\}$, disenrolling the $(t+i)$ colluding participants and updating the shared secret to K_1. We note that K_1 is not disclosed to the colluders in Scheme 2 and can still be used as a shared secret.

Scheme 2 can be regarded as an enhanced version of Scheme 1 [4], with additional forward secrecy. Unlike Scheme 1, shared secret and shares are encrypted by a one-time pad, and one-time pad keys R_i's are released only via broadcast at disenrollment. Hence, the future key exposure due to collusion, exhibited in Scheme 1, is prevented by enforcing broadcast from the dealer in the reconstruction of shared keys in Scheme 2. The forward secrecy is achieved at the expense of increasing the size of broadcast message by that of a shared secret.

IV. Lower Bounds on Broadcast Entropy

To establish lower bounds on the entropy of broadcast in a TSDC satisfying (2), (4) and (11), we consider two cases, (i) no constraints on the share size; (ii) the size of each share achieves its lower bound (9), i.e., share minimal PTSDC.

Theorem 3: Let $S_{1:n}, P_{1:L}, K_{0:L}$ form a (t, n) TSDC satisfying properties (2), (4) and (11), and $H(K_i) = m$ for $i = 0, 1, ..., L$, then

$$H(P_i) \geq H(K_i) = m \qquad i = 1, ..., L. \tag{12}$$

Proof:

$$\begin{aligned} H(P_i) &\geq I(P_i; K_i | S_{1:n}, P_{1:i-1}) \\ &\overset{(a)}{=} H(K_i | S_{1:n}, P_{1:i-1}) - H(K_i | S_{1:n}, P_{1:i-1}, P_i) \\ &= H(K_i) - I(K_i; S_{1:n}, P_{1:i-1}) \\ &\overset{(b)}{=} H(K_i) = m. \end{aligned}$$

The second term of (a) is zero due to (2) and equation (b) holds because of forward secrecy condition (11). ■

Now we consider share minimal perfect threshold schemes with L-fold disenrollment, i.e., the case in which $H(S_j) = (L + 1)m$ for $j = 1, .., n$. It will be proven for this case,

$$H(P_i) \geq \min(i + 1, n - i - t + 1)m \qquad i = 1, ..., L, \tag{13}$$

if all previous broadcast P_l's satisfy lower bounds $\min(l + 1, n - l - t + 1)m$ for $l = 1, .., i - 1$.

In order to establish the bound (13), we first prove some lemmas as follows.

Lemma 1: In a $(1, n)$ share minimal PTSDC with forward secrecy (11), any $i + 1$ subshares $S^{(i)}_{l_1}, ..., S^{(i)}_{l_{i+1}}$ are independent.

Proof: Let v_l denote the index of one valid participant at stage i. A $(1, n)$ PTSDC satisfies the following two equations in terms of subshares.

$$\begin{aligned} H(K_i | S^{(i)}_{d_1:d_i}, P_{1:i}) &= H(K_i) = m \qquad (14) \\ H(K_i | S^{(i)}_{v_l}, P_{1:i}) &= 0, \qquad (15) \end{aligned}$$

where (14) and (15) are obtained from (4) and (5), respectively.

Substituting $X = K_i$, $Y = S^{(i)}_{v_l}$, $Z = S^{(i)}_{d_1:d_i}$ and $W = P_{1:i}$ into Lemma 5 (see Appendix B), we establish from (15), (14), and $H(S^{(i)}_j) = H(K_i)$ that

$$I(S^{(i)}_{v_l}; S^{(i)}_{d_1}, ..., S^{(i)}_{d_i}) = 0. \tag{16}$$

Assume that there is one set of $i+1$ subshares $\{S_{l_1}^{(i)}, ..., S_{l_{i+1}}^{(i)}\}$, which are not independent. That is, in $\{S_{l_1}^{(i)}, ..., S_{l_{i+1}}^{(i)}\}$, there is at least one subshare that is not independent of the rest i subshares. Without loss of generality, we assume that $S_{l_1}^{(i)}$ is dependent on $S_{l_2}^{(i)}, ..., S_{l_{i+1}}^{(i)}$,

$$I(S_{l_1}^{(i)}; S_{l_2}^{(i)}, ..., S_{l_{i+1}}^{(i)}) > 0 \tag{17}$$

For a valid $(1, n)$ threshold scheme, the dealer should be able to disenroll any subset of i participants at stage i. Hence, the indices of the deleted members can be random. One possible value of these indices is given by $\{d_1, d_2, \cdots, d_i\} = \{l_2, ..., l_{i+1}\}$. All remaining indices correspond to the valid members. Therefore, we note that the equation (17) can be written as:

$$I(S_{v_l}^{(i)}; S_{d_1}^{(i)}, ..., S_{d_i}^{(i)}) > 0, \tag{18}$$

leading to a contradiction with (16). ■

Lemma 2: In a $(1, n)$ share minimal PTSDC with forward secrecy,

$$H(P_i) \geq \min(i+1, n-i)m, \tag{19}$$

if all previous broadcast messages meet their lower bound on entropy, i.e., $H(P_w) = \min(w + 1, n - w)$ for $w = 0, ..., i-1$. When $H(P_i)$ achieves its minimum at $(i+1, n-i)m$, then $I(P_i; S_j^{(l)}) = 0$ for $l = i+1, ..., L$ and $j = 1, ..., n$, i.e., P_i is independent of subshares used to reconstruct future shared secrets.

Proof: The proof is built on the fact $H(S_{V_i(u)}^{(i)} | K_i, P_{1:i}) = 0$ with $u = \min(i+1, n-i)$, which is proved in Appendix C. We will prove the lower bound of $H(P_i)$ by induction.

At stage $k = 1$, there are $n-1$ valid participants, $u = \min(k+1, n-k) = \min(2, n-1)$

$$\begin{aligned} H(P_1) &\geq I(P_i; S_{V_i(u)}^{(1)} | K_1) \\ &= H(S_{V_i(u)}^{(1)} | K_1) - H(S_{V_i(u)}^{(1)} | K_1, P_1) \\ &\overset{(a)}{=} H(S_{V_i(u)}^{(1)}, K_1) - H(K_1) = H(S_{V_i(u)}^{(1)}) + H(K_1 | S_{V_i(u)}^{(1)}) - H(K_1) \\ &\overset{(b)}{=} H(S_{V_i(u)}^{(1)}) + H(K_1) - H(K_1) \\ &\overset{(c)}{=} \sum_{l=1}^{u} H(S_{v_l}^{(1)}) = um = \min(2, n-1)m \end{aligned}$$

Equation (a) holds because $H(S_{V_i(u)}^{(1)} | K_1, P_1) = 0$. Without P_1, $H(K_1 | S_{V_i(u)}^{(1)}) = H(K_1)$ and hence (b) holds. Equation (c) holds due to the independence of $S_{v_l}^{(i)}$'s for $l = 1, ..., i+1$ shown in Lemma 1.

Now we show if $H(P_1) = \min(2, n-1)m$, then $I(P_1; S_j^{(l)}) = 0$ for $l = 2, ..., L$ and $j = 1, ..., n$.

$$\begin{aligned}
H(P_1) &\geq I(P_1; S_j^{(l)}, K_1, S_{V_i(u)}^{(1)}) \\
&\geq I(P_1; S_j^{(l)}) + H(S_{V_i(u)}^{(1)} | K_1, S_j^{(l)}) - H(S_{V_i(u)}^{(1)} | K_1, P_1, S_j^{(l)}) \\
&= I(P_1; S_j^{(l)}) + H(S_{V_i(u)}^{(1)} | S_j^{(l)}) + H(K_1 | S_j^{(l)}) - H(K_1) \\
&\overset{(d)}{=} I(P_1; S_j^{(l)}) + H(S_{V_i(u)}^{(1)}) \\
&= I(P_1; S_j^{(l)}) + \min(2, n-1)m
\end{aligned}$$

Equation (d) holds because of independence of subshares of one participant and forward secrecy condition (11). From $H(P_1) \geq I(P_1; S_j^{(l)}) + \min(2, n-1)m$, a necessary condition for $H(P_1)$ to achieve its lower bound is $I(P_1; S_j^{(l)}) = 0$ for $l = 2, ..., L$.

Assume Lemma 2 is true for stage $k = i-1$, i.e., $H(P_{i-1}) \geq \min(i, n-i+1)m$ if all previous broadcast messages reach their lower bound on entropy, i.e., $H(P_w) = \min(w+1, n-w)$ for $w = 0, ..., i-2$, and $I(P_{i-1}; S_j^{(l)}) = 0$ for $l = i, ..., L$.

When $k = i$, $u = \min(k+1, n-i) = \min(i+1, n-i)$

$$\begin{aligned}
H(P_i) &\geq I(P_i; S_{V_i(u)}^{(i)} | K_i, P_{1:i-1}) \\
&= H(S_{V_i(u)}^{(i)} | K_i, P_{1:i-1}) - H(S_{V_i(u)}^{(i)} | K_i, P_{1:i}) \\
&= H(S_{V_i(u)}^{(i)}, K_i | P_{1:i-1}) - H(K_i | P_{1:i-1}) \\
&= H(S_{V_i(u)}^{(i)} | P_{1:i-1}) + H(K_i | S_{V_i(u)}^{(i)}, P_{1:i-1}) - H(K_i | P_{1:i-1}) \\
&= H(S_{V_i(u)}^{(i)}) = um = \min(i+1, n-i)m
\end{aligned}$$

The proof of $I(P_i; S_j^{(l)}) = 0$ for $l = i+1, ..., L$ when $H(P_i)$ achieves its minimum at $(i+1, n-i)m$ is similar to the base case ($k = 1$) and thus is omitted. ■

Now we can establish the lower bound of $H(P_i)$ for a share minimal PTSDC.

Theorem 4: Let $S_{1:n}, P_{1:L}, K_{0:L}$ form a (t, n) TSDC satisfying properties (2), (4), (9) and (11), then

$$H(P_i) \geq \min(i+1, n-i-t+1)m \qquad i = 1, ..., L. \tag{20}$$

if all previous broadcast messages P_l's for $l = 1, .., i-1$ achieve their lower bounds as $\min(l+1, n-l-t+1)m$.

Proof: As shown in [6], if $\{S_{1:n}, P_{1:L}, K_{0:L}\}$ form a (t, n) share minimal PTSDC, then

$$\{S_{l_1:l_{n-t+1}}|S_{V_i(t-1)}, P_{1:L}|S_{V_i(t-1)}, K_{0:L}|S_{V_i(t-1)}\}$$

form a $(1, n-t+1)$ share minimal PTSDC, where $(\cdot|S_{V_i(t-1)})$ represents the conditioning on $S_{V_i(t-1)}$ and $\{l_1, ..., l_{n-t+1}\} = \mathcal{N}\backslash\{V_i(t-1)\}$. For the $(1, n-t+1)$ TSDC, we have $H(P_i|S_{V_i(k)}) \geq \min(i+1, n-t+1-i)m$ from Lemma 2, if all previous broadcast messages meet their lower bound on their entropy.

Since $H(P_i) \geq H(P_i|S_{V_i(k)})$, then (20) holds. ■

Comparing Theorem 2 and Theorem 4, we notice that when adding forward secrecy condition (11) to prevent unwanted key disclosure, the lower bound on broadcast size (20) is different from (10) as expected.

To construct a scheme that achieves the lower bound on broadcast size (20), we will first propose Scheme 3, and then show that by combining Scheme 3 and Scheme 2, we can obtain Scheme 4 that *achieves the lower bound on the share size and the broadcast size*.

Scheme 3:

- On initialization, the dealer randomly chooses strings $R_j^{(i)}$ of length m, and distributes $S_j = \{S_j^{(0)}, R_j^{(1)}, ..., R_j^{(L)}\}$, where $S_j^{(0)}$ is a share of a (t, n) threshold scheme and $R_j^{(i)}$ is used as an one-time pad key for participant j at stage i.
- At update stage i, given the shared secret K_i, the dealer computes the unique polynomial of degree $t-1$, denoted as $f_i(\cdot)$, passing t points $\{v_l, R_{v_l}^{(i)})\}$ for $l = 1, ..., t$. Let $b_t^{(i)} = K_i - f_i(0)$ and $b_l^{(i)} = f_i(v_l) + K_i - f_i(0) - R_l^{(i)}$ for $l = t+1, ..., |V_i|$, the broadcast message is $p_i = \{b_t^{(i)}, b_{t+1}^{(i)}, ..., b_{|V_i|}^{(i)}\}$. For participant v_l with $l = 1, ..., |V_i|$, new subshare $S_{v_l}^{(i)}$ is obtained as

$$S_{v_l}^{(i)} = \begin{cases} R_l^{(i)} + b_t^{(i)} & \text{for } l = 1, ..., t \\ R_l^{(i)} + b_l^{(i)} & \text{for } l = t+1, ..., |V_i|. \end{cases}$$

It follows $S_{v_l}^{(i)} = f_i(v_l) + K_i - f_i(0)$ for $l = 1, ..., |V_i|$. Therefore, t new subshares are sufficient to recover K_i.

In Scheme 3, the share size is of length $(L+1)m$ and the broadcast size at the i^{th} update is $(|V_i| - t + 1)m = (n - i - t + 1)m$ bits.

Scheme 2 has the property of increasing the broadcast size with the disenrollment stage, while Scheme 3 has the property of decreasing the broadcast size. The appropriate combination of these two schemes will result in a bound achieving scheme in a constructive manner, thus realizing the proven lower bounds. We now present a scheme that achieves both the lower bound on share size (8) and the bound on broadcast size (20).

Scheme 4: (A bound achieving TSDC with forward secrecy)

- The share held by participant j after initialization is

$$S_j = \{S_j^{(0)}, ..., S_j^{(l)}, R_j^{(l+1)} ..., R_j^{(L)}\},$$

where $l = \lfloor \frac{n-t}{2} \rfloor$ and $S_j^{(i)}$ is a share of a $(t+i, n)$ share minimal perfect threshold scheme sharing $K_i + R_i$, and $R_j^{(i)}$ is a random string of length m.

- At the i^{th} disenrollment, the dealer broadcasts

$$P_i = \begin{cases} \{R_i, S_{d_1}^{(i)}, ..., S_{d_i}^{(i)}\} & \text{for } i = 1, ..., \lfloor \frac{n-t}{2} \rfloor \\ \{b_t^{(i)}, b_{t+1}^{(i)}, ..., b_{|V_i|}^{(i)}\} & \text{for } i = \lfloor \frac{n-t}{2} \rfloor + 1, ..., L, \end{cases}$$

where $b_t^{(i)}, b_{t+1}^{(i)}, ..., b_{|V_i|}^{(i)}$ are calculated according to Scheme 3.

It can be verified that Scheme 4 is a share minimal PTSDC with forward secrecy that achieves both lower bounds on share size (8) and broadcast size (20).

When comparing Scheme 4 with the bound (10) achieving scheme (Example 4) in [6], the major difference is that when $L < \lfloor \frac{n-t}{2} \rfloor$, Scheme 4 reduces to Scheme 2, while Example 4 reduces to Scheme 1, where forward secrecy is not considered. Hence, the broadcast size of Example 4 is smaller than Scheme 4 by the entropy of a shared secret when $L < \lfloor \frac{n-t}{2} \rfloor$.

We note however, that in the original TSDC, if $(t + L)$ members collaborate, the entire set of keys will be exposed and the group has to be rekeyed. The use of forward secrecy condition (11) only ensures that the dealer plays a role in the disenrollment and it does not prevent the rekey of the group. For example, Scheme 2 ensures that as long as the number of colluders is less than L, the dealer can reach the remaining valid members using a broadcast channel and update shared keys. However, when the number of colluders exceed $(L - 1)$, the dealer has no mechanism to selectively disenroll or reach the remaining valid members to update the keys. Hence, the rekey of entire group is inevitable. This is not a satisfactory condition. To enable the dealer reach any valid subset as long as the subset is of the size t or more, we propose an enhanced TSDC that allows the dealer to simultaneously revoke multiple members.

V. Simultaneous Disenrollment of Multiple Participants

In this section, we present a disenrollment model in which the dealer has the capability to disenroll *any arbitrary number* of participants at each stage, as long as the number of remaining participants is larger than t. We call the capability of disenrolling more than one participant at each stage *enhanced disenrollment capability*, and use D_i and $\tilde{\mathcal{D}}_i$ to denote the set of indices of all disenrolled participants at stage i for $i = 0, 1, ..., L$, and up to stage i, respectively, i.e., $\tilde{\mathcal{D}}_i = \{D_1, ..., D_i\}$. The enhanced disenrollment capability can be expressed as:

$$H(K_i|S_{V_i(k)}, S_{\tilde{\mathcal{D}}_i}, P_{1:i}) = H(K_i) \quad \text{for} \quad k < t. \tag{21}$$

To have L stages in such a disenrollment model, condition $|\tilde{\mathcal{D}}_L| \leq n - t$ should hold, so that there are at least t valid participants after L deletions, i.e., $|V_L| \geq t$.

Now we derive a lower bound on the entropy of broadcast in a TSDC satisfying (2), (11) and (21). If there is no constraint on the share size, it is proven in Theorem 3 that $H(P_i) \geq H(K_i)$. Since only conditions (2) and (11) are used in the proof, the conclusion holds for a TSDC with enhanced disenrollment capability (21). Now we consider the case when the size of each share achieves its lower bound (9), i.e., share minimal PTSDC.

To establish the lower bound on broadcast entropy, we first introduce two lemmas.

Lemma 3: In a TSDC satisfying (2), (9), (11) and (21), n subshares $S_1^{(i)}, ..., S_n^{(i)}$ are independent, where n is the total number of initial participants.
The proof is similar to that of Lemma 1, and thus is omitted due to space limitation.

Lemma 4: In a TSDC satisfying (2), (9), (11) and (21),

$$H(S_{v_l}^{(i)}|P_i, S_{V_i(k)}^{(i)}) = \begin{cases} m & \text{for} \quad k \leq t-1 \\ 0 & \text{for} \quad k \geq t \end{cases} \tag{22}$$

if each broadcast message achieves its lower bound.

Proof: We identify four cases in (22), (a) $k < t-1$, (b) $k = t-1$, (c) $k = t$, and (d) $k > t$. Before proving (22) for these four cases, we make the following observation.

To have a TSDC in which each broadcast message is of its minimum possible size, only P_i out of $\{P_1, ..., P_i\}$ contributes to the reconstruction of K_i, i.e., $H(K_i|S_{V_i(k)}, P_i) = 0$ for $k \geq t$. If P_w with $w < i$ also helps in the reconstruction of K_i, then the redundancy in P_w regarding K_i

can be removed to have a shorter P_w without affecting the reconstruction of K_w. However, this contradicts with the fact that each broadcast message achieves its lower bound.

Now we consider the four cases one by one to prove (22).

Case (a): $k = t-1$.

$$\begin{aligned} H(S_{v_l}^{(i)}|P_i, S_{V_i(t-1)}^{(i)}) &\geq I(S_{v_l}^{(i)}; K_i|P_i, S_{V_i(t-1)}^{(i)}) \\ &= H(K_i|P_i, S_{V_i(t-1)}^{(i)}) - H(K_i|S_{v_l}^{(i)}, S_{V_i(t-1)}^{(i)}, P_i) \\ &= m - 0 = m \end{aligned}$$

On the other hand, $H(S_{v_l}^{(i)}|P_i, S_{V_i(k)}^{(i)}) \leq H(S_{v_l}^{(i)}) = m$.

Case (b): $k < t-1$.

$$H(S_{v_l}^{(i)}) = m \geq H(S_{v_l}^{(i)}|P_i, S_{V_i(k)}^{(i)}) \geq H(S_{v_l}^{(i)}|P_i, S_{V_i(t-1)}^{(i)}) = m$$

Both the steps follow from the facts that the conditioning reduces entropy and $k < t-1$.

Case (c): $k = t$. Since $H(K_i|S_{V_i(t)}^{(i)}, P_i) = 0$,

$$H(S_{v_l}^{(i)}|P_i, S_{V_i(t)}^{(i)}) = H(S_{v_l}^{(i)}|P_i, S_{V_i(t)}^{(i)}, K_i) \leq H(S_{v_l}^{(i)}|P_i, S_{V_i(t-1)}^{(i)}, K).$$

Let $X = K_i$, $Y = S_{v_l}^{(i)}$, and $Z = \{P_i, S_{V_i(t-1)}^{(i)}\}$, from Lemma 6 in Appendix B, we have $H(S_{v_l}^{(i)}|P_i, S_{V_i(t-1)}^{(i)}, K_i) = 0$.

Case (d): $k > t$. We have $0 \leq H(S_{v_l}^{(i)}|P_i, S_{V_i(k)}^{(i)}) \leq H(S_{v_l}^{(i)}|P_i, S_{V_i(t)}^{(i)}) = 0$.

Therefore, (22) holds for all four cases. ■

Theorem 5: Let $S_{1:n}, P_{1:L}, K_{0:L}$ form a (t, n) TSDC satisfying properties (2), (9), (11) and (21), and $|V_L| \geq t$, then

$$H(P_i) \geq (|V_i| - t + 1)m \qquad i = 1, ..., L. \tag{23}$$

If all previous broadcasts achieve their lower bounds, i.e., $H(P_w) = (|V_w| - t + 1)m$ for $1 \leq w < i$, where V_i denotes the set of indices of all valid participants at stage i.

Proof:

$$\begin{aligned} H(P_i) &\geq I(P_i; S_{V_i}^{(i)}, K_i) \\ &= I(P_i; S_{V_i}^{(i)}) + I(P_i; K_i|S_{V_i}^{(i)}) \\ &= \sum_{w=1}^{|V_i|} H(S_{v_w}^{(i)}|S_{v_{1:v_{w-1}}}^{(i)}) - H(S_{v_w}^{(i)}|S_{v_{1:v_{w-1}}}^{(i)}, P_i) + H(K_i|S_{V_i}^{(i)}) - H(K_i|S_{V_i}^{(i)}, P_i) \\ &\overset{(a)}{=} |V_i|m - tm + m - 0 = (|V_i| - t + 1)m \end{aligned}$$

Note that (a) holds due to Lemma 3, Lemma 4 and conditions (2) and (11). ■

Comparing Theorem 4 and Theorem 5, the lower bound on broadcast in a disenrollment model of revoking multiple participants simultaneously is higher than that of deleting one at each stage. It indicates that larger broadcast size is traded for enhanced disenrollment capability.

Using Scheme 3, the dealer is capable of revoking multiple members simultaneously, and the broadcast size achieves the lower bound (23). Unlike Scheme 2 where participants store the shares corresponding to shared secrets, participants are predistributed random string $R_j^{(i)}$'s used for the decryption of new subshares from broadcast in Scheme 3. The dealer can disenroll multiple participants by not including their updated subshares in the subsequent broadcast message. Contrary to Scheme 2, the disenrollment of multiple members can be realized without abandoning any components of predistributed shares. It follows that enhanced disenrollment capability can lead to storage reduction for the disenrollment of multiple participants simultaneously, at the expense of larger broadcast size.

VI. Conclusions

We proposed a TSDC model with forward secrecy, in order to prevent a collusion of participants in TSDC from extracting any future key. We have shown that when all the keys to be protected are of equal length, the lower bound on the cost of providing forward secrecy by enforcing broadcast from the dealer is equal to the entropy of a key used in the system. We also presented an enhancement to TSDC models that allows the dealer to simultaneously disenroll multiple members and demonstrated that the enhanced disenrollment capability is achieved at the expense of increased broadcast size.

References

[1] A. Shamir, "How to share a secret," *Communications of the ACM*, Vol. 22, No. 11, pp.612–613, November 1979.

[2] G. R. Blakley, "Safeguarding cryptographic keys," *Proceedings of AFIPS 1979 National Computer Conference*, 1979, pp. 313–317.

[3] C. K. Wong, M. Gouda, and S. Lam, "Secure group communications using key graphs," *IEEE/ACM Transactions on Networking*, Vol. 8, No. 1, pp. 16–30, Feburary 2000.

[4] B. Blakley, G. Blakley, A. Chan, and J. Massey, "Threshold schemes with disenrollment," *Advances in Cryptology - CRYPTO '92. 12th Annual International Cryptology Conference Proceedings*, pp. 540–548, 1993.

[5] E. Karnin, J. Greene, and M. Hellman, "On secret sharing systems," *IEEE Transactions on Information Theory*, Vol. IT-29, No. 1, pp. 35–41, January 1983.

[6] S. Barwick, W.-A. Jackson, K. Martin, and P. Wild, "Size of broadcast in threshold schemes with disenrollment," *Information Security and Privacy. 7th Australian Conference, ACISP 2002. Proceedings (Lecture Notes in Computer Science Vol. 2384)*, pp. 71–88, 2002.

[7] G. J. Simmons, "How to (really) share a secret," *Advances in Cryptology - CRYPTO '88. 8th Annual International Cryptology Conference Proceedings*, pp. 390–448, 1990.

[8] K. M. Martin, "Untrustworthy participants in perfect secret sharing schemes," *Cryptography and Coding III*, M. J. Ganley, Ed., Vol. 5033. Oxford University Press, 1993, pp. 255–264.

[9] C. Blundo, A. Cresti, A. De Santis, and U. Vaccaro, "Fully dynamic secret sharing schemes," *Theoretical Computer Science*, Vol. 165, No. 2, pp. 407–440, October 1996.

[10] C. Charnes, J. Pieprzyk, and R. Safavi-Naini, "Conditionally secure secret sharing schemes with disenrollment capability," *2nd ACM Conference on Computer and Communications Security*, pp. 89–95, 1994.

[11] U. Tamura, M. Tada, and E. Okamoto, "Update of access structure in Shamire's (k, n) threshold scheme," *Proceedings of 1999 Symposium on Cryptography and Information Security*, 1999, pp. 26–29.

[12] K. M. Martin and J. N. Jr., "Weakness of protocols for updating the parameters of an established threshold scheme," *IEE Proceedings Computers and Digital Techniques*, Vol. 148, No. 1, pp. 45–48, 2001.

[13] M. Li and R. Poovendran, "Broadcast enforced threshold schemes with disenrollment," University of Washington, Technical Report, 2003.

[14] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons Inc, 1991.

Appendix A: Table of notation

TABLE I

$H(\cdot)$	Shannon entropy [14]	i	index for update stages
$I(\cdot)$	mutual Information [14]	j	index for participants
t	threshold	K_i	secret to be shared at stage i
n	total number of participants	S_j	share held by participant j
L	maximum allowable disenrollments	P_i	broadcast message at stage i
m	entropy of shared secrets, w.l.o.g, we assume $H(K_i) = m$ for all i		
$\mathcal{N}$	set of all the indices of n participants, i.e., $\mathcal{N} = \{1, ..., n\}$		
d_i	index of the disenrolled participant at stage i if only one disenrolled each time		
D_i	set of indices of all disenrolled participants at stage i		
$\tilde{\mathcal{D}}_i$	set of indices of all disenrolled participants up to stage i, i.e., $\tilde{\mathcal{D}}_i = \{D_1, ..., D_i\}$		
V_i	set of indices of all valid participants at stage i, $V_i \cup \tilde{\mathcal{D}}_i = N$		
$V_i(k)$	set of indices of k valid participants at stage i, $V_i(k) \subseteq V_i$		
$v_{i,l}$	index of a valid participant at stage i, or simply v_l when no confusion is caused		
$S_j^{(i)}$	subshare of participant j corresponding to the shared secret K_i		
R	random string used to hide a shared secret or a share		
$\lvert X \rvert$	cardinality of X, for example, $\lvert V_i \rvert + \lvert \tilde{\mathcal{D}}_i \rvert = n$		
$X_{a:b}$	set $\{X_a, X_{a+1}, ...X_b\}$ for $a < b$		

Appendix B: useful Lemmas

The proofs of following lemmas are in our technical report [13].

Lemma 5: Let X, Y, Z and W be random variables. Given $H(X|Y, W) = 0$, $H(X|Z, W) = H(X)$, and $H(X) = H(Y)$, then $I(Y; Z) = 0$, where $H(\cdot)$ denotes Shannon entropy and $I(\cdot)$ denotes mutual information.

Lemma 6: Let X, Y and Z be random variables. Given $H(X|Y, Z) = 0$ and $H(X|Z) = H(Y|Z)$, then $H(Y|X, Z) = 0$.

Appendix C: Proof of Claim 1

Claim 1: In a $(1, n)$ share minimal PTSDC satisfying (11), $H(S^{(i)}_{V_i(u)}|K_i, P_{1:i}) = 0$, where $u = \min(i + 1, n - i)$.

Proof: Let S_{v_l} denote one valid share in the set $\{S_{v_1}, ..., S_{v_k}\}$.

$$\begin{aligned} H(S^{(i)}_{v_l}|P_{1:i}) &\geq I(S^{(i)}_{v_l}; K_i|P_{1:i}) \\ &= H(K_i|S^{(i)}_{v_l}) - H(K_i|S^{(i)}_{v_l}, P_{1:i}) \\ &= H(K_i) = m \end{aligned}$$

Since $H(S^{(i)}_{v_l}|P_{1:i}) \leq H(S^{(i)}_{v_l}) = m$, we have $H(S_{v_l}|P_{1:i}) = m$. From (14), we obtain $H(K_i|P_{1:i}) = H(K_i) = m$. By letting $X = K_i$, $Y = S^{(i)}_{v_l}$ and $Z = P_{1:i}$ and applying Lemma 6, we obtain $H(S^{(i)}_{v_l}|K_i, P_{1:i}) = 0$.

$$0 \leq H(S^{(i)}_{V_i(u)}|K_i, P_{1:i}) \leq \sum_{l=1}^{u} H(S^{(i)}_{v_l}|K_i, P_{1:i}) = 0 \tag{24}$$

Therefore, $H(S^{(i)}_{V_i(u)}|K_i, P_{1:i}) = 0$.

■