Dynamic access policies for unconditionally secure secret sharing schemes

Keith M. Martin
Information Security Group
Royal Holloway, University of London
Egham, Surrey TW20 0EX, U.K.
Email: keith.martin@rhul.ac.uk

***Abstract*— Secret sharing schemes are cryptographic primitives that can be employed in any situation where it is desirable to distribute information relating to secret data amongst a number of entities. Following a secure initialization process a secret sharing scheme normally has a fixed access policy (specified by an access structure). However in dynamic environments it is likely that access policies will be subject to change: for example entities may wish to join the scheme, entities may be expelled from the scheme or security thresholds may change. Rather than undergo a costly re-initialization process, there have been several different proposals for designing secret sharing schemes that permit dynamic access policies. We will review various existing models and proposals, and identify several open problems.**

I. Introduction

Secret sharing schemes are cryptographic primitives that provide techniques for protecting high value pieces of information (or *secrets*) by distributing knowledge related to that information across a number of network entities (often referred to as *participants* or *shareholders*). Participants must co-operate by combining privately held related information (*shares* of the secret) in order for the secret to be reconstructed or activated. Each scheme is associated with a fixed collection of sets of entities (often called the *access structure* or *access policy*), which specifies the sets whose collaboration will result in the secret being recovered. If entities not in the access structure attempt to reconstruct the secret then they should not gain any useful information about the secret itself and we say that a scheme is *perfect* if this holds in an unconditionally secure sense. Threshold schemes are probably the most oft-studied shared control systems. In a (k, n)-*threshold scheme* the access structure consists of all subsets of k entities or more.

Secret sharing schemes (and in particular threshold schemes) have found numerous uses within cryptographic systems. They are used by the banking community to directly protect cryptographic master keys, by Certification Authorities to protect private root certificate keys, in the design of key recovery systems, in the shared generation of other secret cryptographic keys and in multiparty computation. They are particularly useful for designing cryptographic primitives for application in increasingly common distributed environments, such as ad hoc networks.

Since the *initialisation phase* of a secret sharing scheme normally requires the secure distribution of related information across a number of entities, they are relatively complex to set up. As a result most schemes are usually intended to be relatively long lasting. This raises the interesting question of what the cost implications are should the access policy (specified by the access structure) need to be changed (or updated) before the expected expiry of the scheme (during the *running phase*). This might happen for a number of reasons: an entity involved in the system may leave or become corrupted (often referred to as *disenrollment*); a change in the organizational security policy might mandate a strengthening (or even a weakening) of the current access structure; or it may be desired that new entities be added to the existing access structure (*enrollment*).

In all the above cases, one solution would simply be to design a new secret sharing scheme and re-issue share information to all the involved entities in the new system. However this requires that the conditions in place during the initialization phase remain in place during the running phase. In particular, the initialization phase normally involves information being passed from some trusted source (often called the *dealer*) to each participant through secure channels. This is normally an expensive process, both in terms of time and resources. For many applications there is no guarantee during the running phase that such channels either remain in place or can be re-established, especially in the event that an urgent change in access policy is mandated by a security compromise of one of the existing participants. There is thus a need to consider alternative techniques for enabling access policy change that do not necessarily involve the need for secure communication between the dealer and the shareholders. These solutions should ideally be light and flexible, without weakening in any way the security of the original scheme.

In this paper we overview a number of different models and schemes that have been considered for enabling dynamic access policies to be enabled for unconditionally secure secret sharing schemes. In the next section we will introduce terminology and present a taxonomy of models within which this problem can be studied. In the subsequent sections we will review work that has been undertaken on these different models. Throughout the paper we will identify open issues where subsequent research could be done.

II. SECRET SHARING SCHEMES

Secret sharing schemes were first proposed in [3], [18]. Since then a considerable body of literature has built up on them, with a good introduction in [21].

For any given secret sharing scheme we let $\mathcal{P}$ denote the set of participants, S the secret and Γ the monotone collection of sets that defines the access structure. We will try to avoid excessive formalization in this paper but note that as first noted in [12] any unconditionally secure (perfect) secret sharing scheme for Γ can easily be expressed as a probability distribution on a collection of possible sharings of the secret such that for any $A \subseteq \mathcal{P}$,

$$H(S \mid A) = \begin{cases} 0 & \text{if } A \in \Gamma \\ H(S) & \text{if } A \notin \Gamma. \end{cases}$$

Much of the secret sharing literature is devoted to determining lower bounds on the size of shares that are necessary for specific access structures. It is easy to establish that for any participant $P \in \mathcal{P}$ in a perfect secret sharing scheme we must have $H(P) \geq H(S)$ and schemes that meet this bound are often referred to as being *ideal*.

III. DYNAMIC ACCESS POLICIES

There are many different conditions under which it might be possible to implement a secret sharing scheme whose access policy can be changed after initialization. In this section we comment on a number of issues that affect the design of a solution.

A. Management models

In [7] a number of management models were identified under which policy changes could be considered. Table I lists twelve plausible management models, categorized in the following terms:

1) Presence of a dealer D during initialization phase (X = present).
2) Type of communication link in place between dealer D and participants P during initialization phase (S = secure channel, B = public broadcast channel, - = no channel).
3) Type of communication link in place between any two participants (P - P) during initialization phase (S = secure channel, B = public broadcast channel, - = no channel).
4) Presence of a dealer D during running phase (X = present)
5) Type of communication link in place between dealer D and participants P during running phase (S = secure channel, B = public broadcast channel, - = no channel).
6) Type of communication link in place between any two participants (P - P) during running phase (S = secure channel, B = public broadcast channel, - = no channel).

TABLE I
MANAGEMENT MODELS FOR DYNAMIC SECRET SHARING SCHEMES

	Initialization phase			Running phase		
Model	D	D - P	P - P	D	D - P	P - P
1 2 3	X	S	-	X	S	S B -
4 5 6	X	S	-	X	B	S B -
7 8 9	X	S	-	-	-	S B -
10 11 12	-	-	S	-	-	S B -

TABLE II
TYPE CLASSIFICATION FOR DYNAMIC SECRET SHARING SCHEMES

Type	Knowledge prior to initialization phase
A	*No prior knowledge*
B	*It is known that changes will be necessary, but not the nature of these changes.*
C	*Some information is known about the nature of changes that will be necessary.*

Note that not all of these models are suitable for facilitating every type of policy change. For example, in any model where there are no secure channels in place during the running phase (such as Models 5, 6, 8 and 9) it is not possible to add new shareholders to the scheme as there is no means of delivering fresh shares to new shareholders. This type of policy change is only possible in these models if there remains some form of dealer entity who can enrol new shareholders as part of an initialization process.

B. Advance information about changes

In addition we can classify techniques for changing the access policy of a secret sharing scheme in terms of the extent to which a change of access policy needs to be anticipated prior to the initialization phase (and hence possibly built into the design of the scheme). These are identified in Table II.

Note the following:

- Type A techniques are by far the most flexible. These are techniques that are largely independent of the underlying secret sharing scheme and can be applied to a scheme that was designed without dynamic access policy change in mind.
- Type B techniques typically only require the designer to know how many times an access policy is likely to require changing.
- Type C techniques typically apply to schemes that are specially designed to allow particular changes such as schemes that permit disenrollment. In such schemes it is known that disenrollments will be necessary during the running phase, but not which participants will need to be disenrolled.

C. Robustness issues

Most of the schemes with dynamic access policies that have been proposed in the literature are robust against coalitions of unauthorised participants attempting to reconstruct secrets that they are not permitted to (although the precise security models vary). They also typically assume that disenrolled participants can no longer be trusted. Note that under certain conditions this forces the secret being protected to change at the same time as the access policy. For example, a (k, n)-threshold scheme can be changed to a $(k-1, n-1)$-threshold scheme under Model 6 simply by broadcasting the share of one of the participants. In this case the secret remains the same. However if a (k, n)-threshold scheme has to be changed to a $(k, n-1)$-threshold scheme then regardless of the technique used the secret after the change will have to be independent of the original one because there exists a set of participants who were allowed to reconstruct the secret before the change, but not after. More generally, this applies to any access policy change from Γ to Γ', where there exists a set $A \in \Gamma$ but $A \notin \Gamma'$.

For many applications this is not significant (if reconstruction of the secret is primarily a demonstration of authorization) but this limitation must be kept in mind for applications where the absolute value of the reconstructed secret is important. The only way that this change in secret can be avoided is through a strengthening of the trust model (for example in some applications a disenrolled shareholder might be a device that ceases to exist and so is not a security liability).

Most of the techniques proposed in the literature do however require participants remaining within the scheme to follow protocols honestly and correctly. Improving the robustness of these techniques against malicious or careless participants generally remains an open problem.

D. Measures of efficiency

We also note that several different measurements of efficiency of techniques for access policy change have been proposed. We will comment more precisely on these later but note for now that these tend either to be measures of:

1) The amount of secret information that participants need to store.
2) The amount of secret information that participants need to communicate as part of the policy change.
3) The amount of information that needs to be broadcast on public channels to facilitate a policy change.

IV. Environments with a fully functional dealer

In this section we briefly consider the case where a fully functional dealer remains active during the running phase of the secret sharing scheme. By "fully functional" we mean that the dealer can still conduct all operations that he was able to do during the initialization phase. This corresponds to Models 1, 2 and 3 of Table I.

These models have received almost no attention in the literature because they all permit a trivial Type A solution to the problem through the reissuing of new shares to all participants by means of the existing secure channels open from the dealer to the shareholders.

Before dismissing these models further we note that this is not quite the end of the story. While the trivial Type A solution mentioned above is functional, it is not necessarily the most efficient solution to the problem. As a very basic illustration, assume that a generic (k, n)-threshold scheme was established during the initialization phase. If a change in access policy requires this to be extended to a $(k, n+1)$-threshold scheme through a single enrollment, it is clearly not necessary to reissue $n+1$ new shares. Rather it suffices that one new share of the original scheme is generated by the dealer and sent out to the new shareholder.

While the above case is very clear cut, a more generic problem remains open.

Open Problem 1: Given access structures Γ and Γ', what is the most efficient Type A technique for converting a secret sharing scheme for Γ into a secret sharing scheme for Γ' in models 1, 2 or 3?

V. Environments with a broadcast-capable dealer

The environment that has been most studied for enabling access policy change is when a dealer remains available during the running phase but can only communicate with shareholders through public channel broadcast messages. This type of scenario makes most sense for applications where shareholders are securely initialized at a local trusted centre but during the running phase are remotely located and cannot easily communicate with that centre. Such environments are represented by Models 4, 5 and 6, with Model 6 being the most studied.

A. Two basic techniques

There are two "obvious" techniques that apply in Model 6.

The trivial advanced share technique
This Type C scheme involves issuing each shareholder during the initialization phase with one share for every possible access policy that they might need to change to in the future. The resulting scheme has impractically large secret storage requirements for each shareholder, but has the advantage that the access policy can be switched simply by an instruction to change over to a new set of shares.

The advanced key technique
This type B scheme involves issuing each shareholder during the initialization phase with a share in the original scheme and one unique random shareholder key for every change that will be need in the future. When an access policy change is required the dealer broadcast the new shares, each protected (XORed) with the random shareholder key for that change. This involves storing as many keys as changes required (which might be fairly small) but broadcasting a substantial amount of information at each access policy change.

Most of the work on designing Model 6 schemes has essentially involved examining the extent to which these "obvious" techniques can be bettered. We now review this work.

B. Disenrollment schemes

The first secret sharing schemes with dynamic access policies to be studied were exclusively for disenrollment [4], [15]. Disenrollment involves removing shareholders from a secret sharing scheme without reducing the effective level of security. For a (k, n)-threshold scheme this means changing the access policy to that of a $(k, n-1)$-threshold scheme. A secret sharing scheme with *L-fold disenrollment capability* is one that permits up to L sequential disenrollments. These are Type C schemes, because it is not known exactly which shareholders will require to be disenrolled, but it is known that disenrollments of some shareholders might be necessary.

In [4] it was shown that for a (k, n)-threshold scheme with L-fold disenrollment it is necessary that $H(P) \geq (L+1)H(S)$ and two optimal schemes were provided. One was the advanced key technique and the other was a type of advanced share technique that works as follows:

1) Issue each shareholder with a share from each of L independent threshold schemes corresponding to thresholds (k, n), $(k+1, n)$, ..., $(k+L, n)$, respectively.
2) Initially the effective threshold scheme is the (k, n)-threshold scheme corresponding to the first set of shares.
3) After the ith disenrollment, shareholders are instructed to move on to their $(i+1)$th share (which corresponds to the $(k+i, n)$-threshold scheme) and the dealer broadcasts the values of the i shares in this scheme corresponding to each of the disenrolled shareholders. This effectively realizes a $(k, n-i)$-threshold scheme on the remaining shareholders, as required.

The above scheme is a variant of a more general scheme in [15] for providing L-fold disenrollment for arbitrary access policies. While these schemes are optimal with respect to secret storage, they are not optimal with respect to the amount of information that the dealer needs to broadcast. In [4] a conjecture was made about this minimal broadcast size, however this was disproved and corrected in [2], where the correct bound was established and a more complex construction provided that is optimal both with respect to secret storage and broadcast information.

More recently this problem has been studied in [13], where schemes that offer a degree of forward secrecy under a more powerful notion of collusion security are proposed.

C. Making general changes to the access policy

The problem of making more general changes to an access policy in Model 6 Type C environments has been considered in [5], [1]. Let $\mathcal{U}$ be a family of access structures. We say that a secret sharing scheme for access structure Γ is *changeable to* $\mathcal{U}$ if it is possible to change the access structure from Γ to any access structure in $\mathcal{U}$ by means of a public channel broadcast from the dealer.

It would seem intuitive that the larger $\mathcal{U}$ becomes, the more challenging it becomes to design such a scheme. In [5] a closely related formal model for so-called *dynamic secret sharing* was proposed and lower bounds on the secret storage for such schemes were computed. Interestingly it was shown that in fact the amount of secret storage was more closely influenced by the number of required changes to the original access structure than to the set $\mathcal{U}$.

In [1] this problem was investigated more closely for the case of threshold schemes that require just one change. Bounds were proven for the minimal amount of broadcast information that needs to be used for threshold schemes with minimal secret storage that are changeable to $\mathcal{U}$ for the entire range of possible sets $\mathcal{U}$ (where the only restriction was that $\mathcal{U}$ defined a family of threshold access structures). Optimal schemes were also constructed that meet both these bounds.

D. Open problems

A number of open problems emerge from the work already conducted on Model 6.

Open Problem 2: Can the results of [1] for bounding the broadcast size for enabling one change to threshold schemes be extended to multiple changes for arbitrary access policies?

Open Problem 3: The bounds in [1] provide lower bounds on the broadcast information for schemes with minimal secret storage. However the trivial advanced key share technique demonstrates (in extremis) that if the secret storage is not minimal then the broadcast information can be reduced. Are there schemes that offer sensible tradeoffs? In other words, if the secret storage is not minimal (but still "small") can we design schemes with "small" broadcast sizes?

Open Problem 4: The work of [13] has begun to consider a strengthening of the trust assumptions for Model 6 schemes. Can schemes within this model be made more robust against insiders failing to follow the access policy change protocols correctly?

Open Problem 5: We note that all the schemes described in this section, with the exception of the advanced key technique, are Type C solutions. In [22] an intriguing Type B threshold scheme was proposed that was designed to be changeable to all other types of threshold access policy. Unfortunately this was shown to be flawed in [16], however the interesting question remains as to whether it is possible to design other Type B schemes within Model 6.

VI. Environments with an inactive dealer

Perhaps the most intriguing models within which to try to design schemes with dynamic access policies are Models 7, 8 and 9. In these environments the trusted dealer who initialized the scheme is no longer present in any capacity.

A. Redistributions

The most flexible model within which to design schemes permitting access policy change without dealer assistance is Model 7, where the shareholders can communicate amongst themselves using secure channels. It is possible to replace the dealer functionality by collectively conducting secure operations between the shareholders. Such schemes are often refereed to as *redistributions* because they essentially involve the secret information that was distributed amongst the original shareholders being redistributed to the new shareholders in a way that permits a new access policy to protect it. This is best illustrated by the following simple example.

Suppose that we wish to change the access policy of an existing secret sharing scheme for Γ (defined on set $\mathcal{P}$) to one for access structure Γ' (for simplicity defined on a disjoint set $\mathcal{P}'$):

1) Each shareholder P in $\mathcal{P}$ generates a personal secret sharing scheme for Γ' whose secret corresponds to P's share in the original scheme.

2) Each shareholder securely transfers the shares of their personal scheme to the corresponding shareholders of the new scheme (in fact it suffices that only a set of shareholders $A \in \Gamma$ need do this).

It is straightforward to see that the end result of this process is a secret sharing scheme for Γ'. As first noted in [9], this protocol can be refined if the underlying secret sharing schemes are *linear* because the new shareholders can then combine all the received shares, thus reducing their storage requirements. Similar ideas have also been applied in other areas of threshold cryptography, notably for *proactive secret sharing* [10] and shared generation of RSA keys [6].

The redistribution problem was formally modelled in [17], where various classes of scheme were defined based on the functionality of the original shareholders following the redistribution. In [17] bounds were computed on the number of secure links required for a successful redistribution and on the amount of exchanged information, and optimal schemes were exhibited.

It is worth noting that redistribution schemes are almost all Type A, although some efficient solutions require the original scheme to be linear (which is a natural property of many efficient secret sharing schemes and so is not regarded as being particularly restrictive).

B. Techniques based on subshares

Once we remove the existence of secure channels between shareholders the design of schemes with dynamic access policies becomes much more challenging. In [14] some Type C (k, n)-threshold schemes were proposed for Model 9. These rely on carefully designed schemes whose original shares have a special structure that allows subshares to be defined that can be used to define (k', n)-threshold schemes for $k' > k$. These schemes are very specialised and the technique does not appear to easily generalise.

C. Towards Type A broadcast solutions?

A intriguing approach based on lattice reduction has recently been proposed in [19], [20] for increasing the threshold of a pre-existing threshold scheme (the former for the shamir threshold scheme and the latter for a threshold scheme based on the Chinese Remainder Theorem). This is the first attempt to design Type A schemes within Model 9. It seems quite counterintuitive that a technique could be found that converts a pre-existing (k, n)-threshold scheme into a (k', n)-threshold scheme for $k' > k$ without any secure links between participants or secure links to a dealer. The problem with these techniques however are that they are complex and appear to only offer probabilistic levels of security for the new schemes. In this sense they do not entirely belong in this review article as the new schemes do not offer unconditional security.

D. Open problems

The following problems for environments with inactive dealers remain open:

Open Problem 6: How efficiently can redistribution protocols be made more robust against insiders failing to behave honestly?

Open Problem 7: Can the lattice-reduction techniques of [19], [20] be improved to deliver a technique that provides guaranteed levels of security for the new schemes and do these techniques extend to chages other than threshold increases?

Perhaps more generally:

Open Problem 8: Is it possible to design any practical schemes in Model 8 or Model 9 for any of the Types A, B or C?

VII. Dealer-free environments

Lastly, we note that Models 10, 11 and 12 all correspond to the special case of *dealer-free* secret sharing schemes (sometimes referred to as *MTA-free* secret sharing schemes [11]). To the best of our knowledge the problem of changing the access policy for such schemes has not been examined. However dealer-free schemes have themselves be used as primitives to facilitate redistributions in Model 7 and so it is likely that these protocols can be adapted fairly easily for Model 10 environments.

Open Problem 9: A thorough investigation of access policy change within Models 10, 11 and 12 remains requires investigation.

VIII. Conclusions

In this paper we have provided a brief overview of the body of work that has been conducted on secret sharing schemes that permit access policy changes. We have shown that techniques for facilitating this vary greatly depending on the environment in place during the running phase of the scheme. In addition we have identified a number of open problems where further research work could be undertaken.

References

[1] S.G. Barwick, W.-A. Jackson, and K.M. Martin, "Updating the parameters of a threshold scheme by minimal broadcast," *IEEE Trans. Inf. Theory*, vol. 51, no. 2, pp.620–633, 2005.

[2] S.G. Barwick, W.-A. Jackson, K.M. Martin and P.R. Wild, "Size of broadcast in threshold schemes with disenrollment," *Information Security and Privacy, Lecture Notes in Comput. Sci.*, 2384, pp.71–88, 2002.

[3] G.R. Blakley, "Safeguarding cryptographic keys," *Proceedings of AFIPS 1979 National Computer Conference*, 48, pp.313–317, 1979.

[4] B. Blakley, G.R. Blakley, A. Chan and J. Massey, "Threshold schemes with disenrollment," *Adv. in Cryptology - CRYPTO'92, Lecture Notes in Comput. Sci.*, 740, pp.540–548, 1993.

[5] C. Blundo, A. Cresti, A. De Santis and U. Vaccaro, "Fully dynamic secret sharing schemes," *Adv. in Cryptology - CRYPTO'93, Lecture Notes in Comput. Sci.*, 773, pp.110–125, 1994.

[6] D. Boneh and M. Franklin, "Efficient generation of shared RSA keys," *Adv. in Cryptology - CRYPTO'97, Lecture Notes in Comput. Sci.*, 1294, pp.425–439, 1997.

[7] G. Chaddoud and K.M. Martin, "A framework for dynamically updating shared control schemes," unpublished.

[8] L. Chen, D. Gollmann and C.J. Mitchell, "Key escrow in mutually mistrusting domains," *Security Protocols, Lecture Notes in Comput. Sci.*, 1189, pp.139–153, 1997.

[9] Y. Desmedt and S. Jajodia, "Redistributing secret shares to new access structures and its applications," unpublished.

[10] A. Herzberg, S. Jarecki, H. Krawczyk and M. Yung, "Proactive secret sharing or: how to cope with perpetual leakage," *Adv. in Cryptology - CRYPTO'95, Lecture Notes in Comput. Sci.*, 963, pp.339–352, 1995.

[11] W.-A. Jackson, K.M. Martin and C.M. O'Keefe, "Mutually Trusted Authority free secret sharing schemes," *Journal of Cryptology*, vol. 10, pp.261–289, 1997.

[12] E. Karnin, J. Greene and M. Hellman, "On secret sharing systems," *IEEE Trans. Inf. Theory*, vol. 29, pp.35–41, 1983.

[13] M. Li and R. Poovendran, "Broadcast-enforced disnrollment in threshold schemes," *SAC 2003, Lecture Notes in Comput. Sci.*, 3006, pp.101–116, 2004.

[14] K.M. Martin, J. Pieprzyk, R. Safavi-Naini and H. Wang, "Changing thresholds in the absence of secure channels," *Australian Computer Journal*, 31, pp.34–43, 1999.

[15] K.M. Martin, "Untrustworthy participants in secret sharing schemes," *Cryptography and Coding III*, Oxford University Press, Oxford, pp.255–264, 1993.

[16] K.M. Martin and J. Nakahara Jr, "Weaknesses of protocols for updating the parameters of an established threshold scheme," *IEE Proc. Comput. Digit. Tech.*, vol. 148, no. 1, pp. 45–48, 2001.

[17] K.M. Martin, R. Safavi-Naini and H. Wang, "Bounds and techniques for efficient redistribution of secret shares to new access structures," *The Computer Journal*, vol.42, no.8, pp.638–649, 1999.

[18] A. Shamir, "How to share a secret," *Comm. ACM*, 22(11), pp.612–613, 1979.

[19] R. Steinfeld, J. Pieprzyk and H. Wang, "Lattice-based threshold-changeability for standard shamir secret-sharing schemes," *Adv. in Cryptology - ASIACRYPT 2004, Lecture Notes in Comput. Sci.*, 3329, pp.170–186, 2004.

[20] R. Steinfeld, J. Pieprzyk and H. Wang, "Dealer-free threshold changeability for standard CRT secret-sharing schemes," *Finite fields and their applications*, in press.

[21] D.R. Stinson, "An explication of secret sharing schemes," *Des. Codes Cryptogr.*, 2, pp.357–390, 1992.

[22] Y. Tamura, M. Tada and E. Okamoto, "Update of access structure in shamir's (k, n) threshold scheme," *Proc. 1999 Symposium on Cryptography and Information Security*, Kobe, Japan, January 1999.