Trade-offs Between Communication and Storage in Unconditionally Secure Schemes for Broadcast Encryption and Interactive Key Distribution

Carlo Blundo[1], Luiz A. Frota Mattos[2], and Douglas R. Stinson[3]

[1] Dipartimento di Informatica ed Applicazioni,
Università di Salerno, 84081 Baronissi (SA), Italy

E-mail: carblu@dia.unisa.it
URL: http://www.unisa.it/carblu.dir/

[2] CEPESC/SAE
SAISw Área 5, Quadra 1, Bloco V
70610 Brasília, DF, Brazil

E-mail: frota@antares.linf.unb.br

[3] Department of Computer Science and Engineering
University of Nebraska-Lincoln, Lincoln NE 68588, USA

E-mail: stinson@bibd.unl.edu
URL: http://bibd.unl.edu/~stinson

Abstract. In 1993, Beimel and Chor presented an unconditionally secure interactive protocol which allows a subset of users in a network to establish a common key. This scheme made use of a key predistribution scheme due to Blom.
In this paper, we describe some variations and generalizations of the Beimel-Chor scheme, including broadcast encryption schemes as well as interactive key distribution schemes. Our constructions use the key predistribution scheme of Blundo *et al*, which is a generalization of the Blom scheme. We obtain families of schemes in which the amount of secret information held by the network users can be traded off against the amount of information that needs to be broadcast.
We also discuss lower bounds on the storage and communcation requirements of protocols of these types. Some of our schemes are optimal (or close to optimal) with respect to these bounds.

1 Introduction

When a subset of users in a network wishes to communicate privately in conference, encryption algorithms can be employed to provide security against eavesdropping. If conventional (private-key) cryptography is used, a common key must be shared by the members of the conference, which we call the *privileged subset*.

One solution is to use a *Key Predistribution Scheme* (KPS) [10] in which secret information is given to each user by a trusted authority in such way that specified privileged subsets can compute a secret key. One such scheme

N. Koblitz (Ed.): Advances in Cryptology - CRYPTO '96, LNCS 1109, pp. 387-400, 1996.
© Springer-Verlag Berlin Heidelberg 1996

was described by Blom [3], in which there is a secret key associated with each pair of users. By giving each user $w + 1$ pieces of secret information, any of the resulting secret keys is unconditionally secure against a coalition of size w. A generalization of Blom's method can be found in [4], and a survey of key predistribution schemes can be found in [11].

A different approach is to use a *Broadcast Encryption Scheme* (BES), in which the TA broadcasts an encrypted version of a conference key, whose value can be decrypted only by members of the privileged subset. (Unconditionally secure) broadcast encryption was first introduced by Fiat and Naor [8], and has been further studied in [5, 11].

A third approach is for the members of the privileged set to interactively compute a secret conference key by exchanging messages among themselves. Of course there many such schemes that are computationally secure, going back to the Diffie-Hellman scheme. (One nice conference scheme was recently described by Burmester and Desmedt [6].) Unconditionally secure schemes for conference key agreement have been studied by Beimel and Chor [1, 2] (see also Blundo and Cresti [5]). We will call such a scheme an *Interactive Key Distribution Scheme* (IKDS). An IKDS typically consists of a key predistribution phase (which requires a trusted authority), followed by an interactive communication phase among the conference participants (which does not involve the trusted authority).

In this paper, we present new constructions for BES and IKDS. These schemes are *one-time* schemes in that they can be used for only a single broadcast (in the case of BES) or to establish only one key (in the case of IKDS). Hence, we will use the acronyms OTBES and OTIKDS, respectively. One situation in which this would be very appropriate is if the key that is being established is a (long-term) master key. We should also note, however, that one-time schemes can generally be modified in a straightforward way to obtain τ-time schemes; see [1, 5], for example.

Our new schemes allow a trade-off between communication and storage. In general, a smaller broadcast size can be accomplished if the participants have more secret information, and vice versa. One of the main purposes of this paper is to examine and quantify these trade-offs.

It should be evident that we are interested in *unconditionally secure* schemes in this paper, i.e., schemes that do not depend on any computational assumptions. Although these schemes require a trusted authority, they are nevertheless extremely efficient computationally. There is also some benefit in using an unconditionally secure key distribution mechanism, even if the key is to be subsequently used in a conventional, computationally-secure cryptosystem such as DES.

2 Mathematical Models

2.1 Key Predistribution Schemes

We begin by discussing key predistribution schemes, since we will be using these as building blocks in our consideration of broadcast encryption schemes and interactive key distribution schemes.

Our model consists of a trusted authority (TA) and a set of users $\mathcal{U} = \{1, 2, \ldots, n\}$. We assume that network is a *broadcast channel*, i.e., any information transmitted by the TA (or by a user in the network) will be received by every user. It is assumed to be insecure against passive attacks, i.e., the information that is broadcast can be observed by anyone. However, we will assume that the network is secure against active attacks. (In practice, we could obtain protection against active attacks by using an unconditionally secure authentication code to authenticate all information that is broadcast.)

In a key predistribution scheme, the TA generates and distributes secret information to each user. The information given to user i is denoted by u_i and must be distributed "off-band" (i.e., not using the network) in a secure manner. For $1 \leq i \leq n$, let U_i denote the set of all possible secret values that might be distributed to user i by the TA. For any subset of users $X \subseteq \mathcal{U}$, let U_X denote the cartesian product $U_{i_1} \times \ldots \times U_{i_j}$, where $X = \{i_1, \ldots, i_j\}$ and $i_1 < \ldots < i_j$. We assume that there is a probability distribution on $U_{\mathcal{U}}$, and the TA chooses $u_{\mathcal{U}} \in U_{\mathcal{U}}$ according to this probability distribution. This secret information will enable various *privileged subsets* to compute keys.

Let $2^{\mathcal{U}}$ denote the set of all subsets of users. $\mathcal{P} \subseteq 2^{\mathcal{U}}$ will denote the collection of all privileged subsets to which the TA is distributing keys. $\mathcal{F} \subseteq 2^{\mathcal{U}}$ will denote the collection of all possible coalitions (called *forbidden subsets*) against which each key is to remain secure.

Once the secret information is distributed, each user i in a privileged set P should be able to compute the key k_P associated with P. On the other hand, no forbidden set $F \in \mathcal{F}$ disjoint from P should be able to compute any information about k_P.

The desired properties can be described mathematically using the entropy function (see [7] for basic terminology and results on information theory). We say that the scheme is a $(\mathcal{P}, \mathcal{F})$-**Key Predistribution Scheme** (or $(\mathcal{P}, \mathcal{F})$-KPS) provided the following conditions are satisfied:

(KPS1) Each user i in any privileged set P can compute k_P: $H(K_P|U_i) = 0$ for all $i \in P$, $P \in \mathcal{P}$.

(KPS2) No forbidden subset F disjoint from a privileged subset P has any information on k_P: $H(K_P) = H(K_P|U_F)$ for all $P \in \mathcal{P}$ and $F \in \mathcal{F}$ such that $P \cap F = \emptyset$.

Usually, we will be considering schemes where $\mathcal{P}$ consists of all t-subsets of $\mathcal{U}$, and $\mathcal{F}$ consists of all subsets of $\mathcal{U}$ of size at most w. Such a KPS will be denoted as a (t, w)-KPS.

2.2 Broadcast Encryption Schemes

We now turn to the notion of a one-time broadcast encryption scheme. We describe the model from [11] (a different model, which is not a one-time scheme, is presented in [5]). In our model, there is an initial set-up phase in which the TA distributes secret information to the network users, just as in a key predistribution system. As before, we denote the secret information given to user i by u_i.

At a later time, the TA will want to broadcast a *message* (i.e., a plaintext) to a privileged subset P. We will let $\mathcal{P} \subseteq 2^{\mathcal{U}}$ denote the collection of all privileged subsets to which the TA might want to broadcast a message. The particular privileged subset $P \in \mathcal{P}$ to which the TA will broadcast a message is, in general, not known ahead of time.

The message to be broadcast to P will be denoted as m_P, and is chosen by the TA from a specified set M_P according to a specified probability distribution on M_P. Then the *broadcast* b_P (which is an element of a specified set B_P) is computed as a function of m_P and u_P.

Once b_P is broadcast, each user $i \in P$ should be able to decrypt b_P and obtain m_P. On the other hand, we will desire that the broadcast should be secure against specified coalitions. $\mathcal{F} \subseteq 2^{\mathcal{U}}$ will denote the collection of all possible forbidden subsets against which a broadcast is to remain secure. We require that no $F \in \mathcal{F}$ disjoint from P should be able to compute any information about m_P.

As mentioned above, we discuss the security in terms of a single broadcast, so we call the scheme "one-time". Here is a formal definition. We say that the scheme is a $(\mathcal{P}, \mathcal{F})$-**One-Time Broadcast Encryption Scheme** (or $(\mathcal{P}, \mathcal{F})$-OTBES) provided the following conditions are satisfied:

(OTBES1) Without knowing the broadcast, no subset of users has any information about m_P, even given all the secret information $U_{\mathcal{U}}$: $H(M_P|U_{\mathcal{U}}) = H(M_P)$ for all $P \in \mathcal{P}$.

(OTBES2) The message for a privileged user is uniquely determined by the broadcast and the user's secret information: $H(M_P|U_iB_P) = 0$ for all $i \in P, P \in \mathcal{P}$.

(OTBES3) After receiving the broadcast, no forbidden subset F disjoint from P has any information on m_P: $H(M_P) = H(M_P|U_FB_P)$ for all $P \in \mathcal{P}$ and $F \in \mathcal{F}$ such that $P \cap F = \emptyset$.

As with KPS, we will be considering OTBES where privileged sets have size t and forbidden sets have size (at most) w; the notation (t, w)-OTBES will be used to describe this scenario.

In Section 6, we will describe our model for interactive key distribution schemes.

3 The Blundo *et al* KPS and its Properties

We will be using the key predistribution scheme of Blom [3], and the generalization of Blundo *et al* [4]. The Blundo *et al* scheme is a (t, w)-KPS (and the Blom

scheme is the special case $t = 2$). Let p be a prime such that $p \geq n$ (the number of users). The TA chooses n distinct random numbers $s_i \in \mathbf{Z}_p$, and gives s_i to user i ($1 \leq i \leq n$). These values s_i do not need to be secret.

Next, the TA the constructs a random symmetric polynomial in t variables with coefficients from $\mathbf{Z}_p$, in which the degree of any variable is at most w:

$$f(x_1, \ldots, x_t) = \sum_{i_1=0}^{w} \cdots \sum_{i_t=0}^{w} a_{i_1,\ldots,i_t} {x_1}^{i_1} \cdots {x_t}^{i_t}.$$

Then, for $1 \leq i \leq n$, the TA computes a polynomial g_i in the $t-1$ variables $x_2, \ldots, x_t$ by setting $x_1 = s_i$ in $f(x_1, \ldots, x_t)$. The coefficients of g_i comprise the secret information which is given to user i. The key associated with the t-subset $P = \{i_1, \ldots i_t\}$ is

$$k_P = f(s_{i_1}, \ldots, s_{i_t}) \bmod p.$$

Each user $i_j \in P$ can compute

$$k_P = g_{i_j}(s_{i_1}, \ldots, s_{i_{j-1}}, s_{i_{j+1}}, \ldots, s_{i_t}) \bmod p.$$

It can be shown that no subset of w users disjoint from P can compute any information about k_P (see [4]). In fact, given the secret information held by a coalition F of size t, every possible ℓ-tuple of keys held by the ℓ-subsets of P occurs with equal probability $p^{-\alpha}$, where $\alpha = \binom{t}{\ell}$.

Lemma 1. *Suppose we have a Blundo et al $(t + w - \ell, \ell)$-KPS, where $\ell \leq t$. Let P and F two sets of t and w users, respectively, such that $P \cap F = \emptyset$. Let $\alpha = \binom{t}{\ell}$ and let $Y_1, \ldots, Y_\alpha$ be all the subsets of P of cardinality ℓ. Let $k_1, \ldots, k_\alpha$ be arbitrary elements of $\mathbf{Z}_p$, and let u_F be the secret information given to the subset F. Then*

$$p(K_{Y_1} = k_1, \ldots, K_{Y_\alpha} = k_\alpha) = p(K_{Y_1} = k_1, \ldots, K_{Y_\alpha} = k_\alpha | U_F = u_F) = \frac{1}{p^\alpha}.$$

Remark: In the case $\ell = 2$, this is the result stated in [1, Lemma 9].

4 One-time Broadcast Encryption Schemes

4.1 A Construction using Resolvable Designs

In this section, we present a (t, w)-OTBES based on the combinatorial structures called "resolvable designs" (for formal definitions and the results in design theory that we use, we refer to [9]).

Suppose that $\ell \geq 1$ is any integer such that $t \equiv 0 \bmod \ell$. (The "best" choice for ℓ for a given parameter situation will be discussed in Section 5.2.) The set-up phase consists of the TA distributing secret information corresponding to a Blundo *et al* $(\ell, t + w - \ell)$-KPS implemented over $\mathbf{Z}_p$, p prime. For an ℓ-subset of users A, we denote by k_A the key associated with the subset A.

Now suppose that the TA wishes to broadcast a message to a privileged set P of cardinality t. Consider the collection of all ℓ-subsets of P, which we call *blocks*. (This collection is sometimes called the *complete ℓ-uniform hypergraph* on P.) By a famous theorem of Baranyai (see, for example, [9]), this set of $\binom{t}{\ell}$ blocks can be partitioned into $r = \binom{t-1}{\ell-1}$ *parallel classes*, each of which consists of of t/ℓ blocks. (In other words, the hypergraph is *resolvable*.) We will denote these parallel classes by $C_1, \ldots, C_r$, and we will denote the blocks in C_i by $B_{i,j}$, $1 \leq i \leq r$, $1 \leq j \leq t/\ell$. (The resolution of the hypergraph into parallel classes is public knowledge.)

There is a key $k_{B_{i,j}}$ associated with every block $B_{i,j}$. The message to be broadcast to P will be an element of $(\mathbf{Z}_p)^r$, say

$$m_P = (m_1, \ldots, m_r).$$

The TA encrypts each m_i using the t/ℓ keys $k_{B_{i,j}}$, by defining

$$b_{i,j} = k_{B_{i,j}} + m_i \bmod p,$$

$1 \leq i \leq r$, $1 \leq j \leq t/\ell$. Then the broadcast b_P is

$$b_P = (b_{1,1}, \ldots, b_{1,t/\ell}, b_{2,1}, \ldots, b_{2,t/\ell}, \ldots, b_{r,1}, \ldots, b_{r,t/\ell}).$$

Let's see how any user $h \in P$ can decrypt the broadcast. For each i, $1 \leq i \leq r$, there is a block $B_{i,h_i} \in C_i$ such that $h \in B_{i,h_i}$. Thus h can compute all the keys $k_{B_{i,h_i}}$, $1 \leq i \leq r$. Then it is a simple matter for h to compute

$$m_i = b_{i,h_i} - k_{B_{i,h_i}} \bmod p,$$

$1 \leq i \leq r$.

Here is a small example to illustrate.

Example 1. Assume that $t = 6$ and $\ell = 2$. Let $P = \{1,2,3,4,5,6\}$ The $\binom{6}{2} = 15$ pairs of users in P can be partitioned into 5 disjoint parallel classes, as follows:

$C_1 = \{\{5,6\},\{1,4\},\{2,3\}\}, C_2 = \{\{5,1\},\{2,6\},\{3,4\}\}, C_3 = \{\{5,2\},\{3,1\},\{4,6\}\}$,
$C_4 = \{\{5,3\},\{1,6\},\{2,4\}\}, C_5 = \{\{5,4\},\{3,6\},\{1,2\}\}$.

Suppose the TA wants to broadcast the message $m_P = (m_1, m_2, m_3, m_4, m_5) \in (\mathbf{Z}_p)^5$ to the users in P. The broadcast b_P is the concatenation of the following 15 values:

$$\begin{array}{c} k_{\{5,6\}} + m_1,\ k_{\{1,4\}} + m_1,\ k_{\{2,3\}} + m_1, \\ k_{\{1,5\}} + m_2,\ k_{\{2,6\}} + m_2,\ k_{\{3,4\}} + m_2, \\ k_{\{2,5\}} + m_3,\ k_{\{1,3\}} + m_3,\ k_{\{4,6\}} + m_3, \\ k_{\{3,5\}} + m_4,\ k_{\{1,6\}} + m_4,\ k_{\{2,4\}} + m_4, \\ k_{\{4,5\}} + m_5,\ k_{\{3,6\}} + m_5,\ k_{\{1,2\}} + m_5, \end{array}$$

where all addition is modulo p.

Let us see how user 3 will decrypt the broadcast. This user knows the five keys $k_{\{1,3\}}$, $k_{\{2,3\}}$, $k_{\{3,4\}}$, $k_{\{3,5\}}$ and $k_{\{3,6\}}$, and hence he or she can perform the following calculations:

$m_1 = b_{1,3} - k_{\{2,3\}} \bmod p$, $m_2 = b_{2,3} - k_{\{3,4\}} \bmod p$, $m_3 = b_{3,2} - k_{\{1,3\}} \bmod p$,
$m_4 = b_{4,1} - k_{\{3,5\}} \bmod p$, $m_5 = b_{5,2} - k_{\{3,6\}} \bmod p$.

4.2 Security of the Scheme

We now briefly discuss the security of the scheme. It is intuitively clear that a coalition of w users disjoint form a privileged set P has no information about m_P after observation of the broadcast, even if they pool all their secret information. This is because of the property, which we stated in Lemma 1, that the $\binom{t}{\ell}$ keys k_A appear to them to be independent random elements of $\mathbf{Z}_p$. Each of these keys is used to encrypt one element of $\mathbf{Z}_p$, and thus these keys function as a big one-time pad. A formal proof can be given by modifying that given in [2] in a straightforward way.

5 Entropy Bounds for OTBES

5.1 Lower Bounds on Entropy

We are interested in the amount of secret information held by the users in an OTBES, as well as the size of the broadcast. As mentioned earlier, there exists a trade-off between these quantities: one can achieve a small broadcast if the amount of secret information is large, and vice versa. In this section, we present a bound that quantifies this trade-off. More specifically, Theorem 2 provides a lower bound on the information held by w users together with the broadcast size (in the case $t \geq w+1$). Because of space limitations, the proof will not be given here, but it can be found in the long version of this paper, available from the authors.

The following theorem holds for arbitrary entropies on the message spaces M_P, but for clarity we will state it for the simpler case where $H(M_P) = H(M_{P'})$ for all $P, P' \in \mathcal{P}$. We denote this common entropy by $H(M)$.

Theorem 2. *Suppose we have a (t,w)-OTBES, where $t \geq w+1$. Then, for any $P = \{i_1, \ldots, i_l\} \subseteq \mathcal{U}$, it holds that*

$$H(B_P) + \sum_{j=1}^{w} H(U_{i_j}) \geq (2w+1)H(M).$$

5.2 The Entropy of our Schemes

We measure the efficiency of our constructions by considering the amount of secret information stored by each user as compared to the information content of the broadcast; and the size of the broadcast as compared to its information content.

Hence, we consider the following quantities:

$$r_i = \frac{H(U_i)}{H(M)},$$

$1 \leq i \leq n$, and

$$r_{B_P} = \frac{H(B_P)}{H(M)}.$$

(Note that these quantities are the reciprocals of *information rates* as defined in [11].)

It is easy to see that $r_i \geq 1$ provided that i is a member of at least one privileged set. Also, $r_{B_P} \geq 1$. It is easy to construct schemes in which $r_i = 1$ for $1 \leq i \leq n$, and it is also trivial to construct schemes in which $r_{B_P} = 1$ (see [11], for example). However, as mentioned previously, the results of Section 5.1 show that there is a trade-off between r_i and r_{B_P}: if one is "small" then the other must be "large". In particular, if we rephrase Theorem 2 in terms of the notation defined above, we have the following result.

Theorem 3. *In any (t, w)-OTBES with $t \geq w + 1$,*

$$r_{B_P} + \sum_{j=1}^{w} r_{i_j} \geq (2w + 1).$$

We look now at the construction for (t, w)-OTBES that we presented in Section 4.1. We assume that a message is a random element of $(\mathbf{Z}_p)^r$, where $r = \binom{t-1}{\ell-1}$. Performing some simple arithmetic, we have the following result.

Theorem 4. *Let ℓ be an integer such that $1 \leq \ell \leq t$. Then there exists a $(t, w)-OTBES$ with*

$$r_i = \frac{\binom{t+w-1}{\ell-1}}{\binom{t-1}{\ell-1}},$$

$1 \leq i \leq n$, and

$$r_{B_P} = \frac{t}{\ell}.$$

We have noted that we are free to choose ℓ however we wish. If we wanted to minimize r_i, we would choose $\ell = 1$. The resulting (trivial) scheme has $r_i = 1$, which is optimal. If we wanted to minimize r_{B_P}, we would choose $\ell = t$, yielding $r_{B_P} = 1$, also optimal. But if we wanted to see how close we can get to attaining the bound of Theorem 3, we should choose some intermediate value of ℓ.

As an example, suppose we consider the case $w = 1$. We know from Theorem 3 that $r_i + r_{B_P} \geq 3$ in any $(t, 1)$-OTBES. Our construction allows us to construct a scheme in which

$$r_i + r_{B_P} = \frac{\binom{t}{\ell-1} + \binom{t}{\ell}}{\binom{t-1}{\ell-1}} = \frac{t(t+1)}{\ell(t-\ell+1)}.$$

We minimize $r_i + r_{B_P}$ by choosing ℓ as follows:

$$\ell = \begin{cases} \frac{t+1}{2} & \text{if } t \text{ is odd} \\ \frac{t}{2} & \text{if } t \text{ is even.} \end{cases}$$

These choices for ℓ yield the following:

$$r_i + r_{B_P} \leq \begin{cases} \frac{4t}{t+1} & \text{if } t \text{ is odd} \\ \frac{4(t+1)}{t+2} & \text{if } t \text{ is even.} \end{cases}$$

In particular, for $t = 2, 3$, we get a $(t, 1)$-OTBES with $r_i + r_{B_P} = 3$, which is optimal. For larger t, we always have $r_i + r_{B_P} < 4$, which is not too far away from the lower bound $r_i + r_{B_P} \geq 3$.

6 The Model for Interactive Key Distribution Schemes

We now turn our attention to one-time interactive key distribution schemes. We will use the model described by Beimel and Chor in [1, 2]. In this model, there is an initial distribution of secret information by the TA, followed by a sequence of messages broadcast by the members of a privileged set, P. At the end of the protocol, every member of P should be able to compute the same key, m_P, while no coalition F disjoint from P should have any information about m_P.

In general, the messages exchanged among the members in P can depend on previous messages, and there may be several rounds of communication. In this paper, we confine our attention to schemes of a very special type, termed "non-reactive" by Beimel and Chor. In a non-reactive scheme, every member $i \in P$ independently chooses a value m_i, and uses his or her secret information, u_i, to compute an encrypted version of m_i, denoted by b_i, which is then broadcast.

Thus a non-reactive one-time key distribution scheme can be thought of as several independent executions of a one-time broadcast encryption scheme, with each privileged user in turn broadcasting a message that can be decrypted only by the other privileged users.

The key m_P for the privileged set P consists of the concatenation of all the values m_i, $i \in P$:

$$m_P = (m_{i_1}, \ldots, m_{i_t}),$$

where $P = \{i_1, \ldots, i_t\}$, $i_1 < \ldots < i_t$. We define the broadcast in an analogous way:

$$b_P = (b_{i_1}, \ldots, b_{i_t}).$$

Here are the formal definitions for a $(\mathcal{P}, \mathcal{F})$-**Non-Reactive One-Time Key Distribution Scheme** (or $(\mathcal{P}, \mathcal{F})$-NROTKDS):

(NROTKDS1) The value of the key is independent of all the secret information $U_\mathcal{U}$:

$$H(M_P|U_\mathcal{U}) = H(M_P)$$

for all $P \in \mathcal{P}$.

(NROTKDS2) The value m_i chosen by each privileged user is independent of the all the secret information $U_\mathcal{U}$ and independent of all the information broadcast by the other privileged users:

$$H(M_i|U_\mathcal{U} B_{P\setminus\{i\}}) = H(M_i)$$

for all $i \in P$, $P \in \mathcal{P}$.

(NROTKDS3) The information broadcast by each privileged user $i \in P$ is a function of his or her secret information and the value chosen for m_i:

$$H(B_i|U_iM_i) = 0$$

for all $i \in P$, $P \in \mathcal{P}$.

(NROTKDS4) The key can be computed by a privileged user from the broadcast and the user's secret information:

$$H(M_P|U_iB_P) = 0$$

for all $i \in P, P \in \mathcal{P}$.

(NROTKDS5) After receiving the broadcast, no forbidden subset F disjoint from P has any information on m_P:

$$H(M_P) = H(M_P|U_FB_P)$$

for all $P \in \mathcal{P}$ and $F \in \mathcal{F}$ such that $P \cap F = \emptyset$.

7 Constructions for Key Distribution Schemes

In this section, we present a (t, w)-NROTKDS which uses (t, w)-OTBES constructed from resolvable designs in Section 4.1. Our construction will contain the Beimel-Chor scheme as a special case (namely, when $\ell = 2$).

Suppose that $\ell \geq 2$ is any integer such that $t \equiv 1 \bmod (\ell - 1)$. The set-up phase consists of the TA distributing secret information corresponding to a Blundo *et al* $(\ell, t + w - \ell)$-KPS implemented over $(\mathbf{Z}_p)^\ell$, p prime. For an ℓ-subset of users A, we denote by k_A the key associated with the subset A. We will think of k_A as being made up of ℓ independent keys over $\mathbf{Z}_p$, which we will denote by $k_{A,1}, \ldots, k_{A,\ell}$.

Suppose that the privileged set P wishes to interactively construct a common key. Each user $h \in P$ will perform the following steps.

1. h chooses a random value $m^h = (m_1^h, \ldots, m_r^h) \in (\mathbf{Z}_p)^r$, where $r = \binom{t-2}{\ell-2}$.
2. Since $t \equiv 1 \bmod (\ell - 1)$, the complete $(\ell - 1)$-uniform hypergraph on $P \backslash \{h\}$ can be partitioned into r parallel classes, each of which consists of of $s = (t-1)/(\ell-1)$ blocks. Denote these parallel classes by $C_1^h, \ldots, C_r^h$, and denote the blocks in C_i^h by $B_{i,j}^h$, $1 \leq i \leq r$, $1 \leq j \leq s$.
3. For each block $B_{i,j}^h$, denote

$$B_{i,j}^h \cup \{h\} = \{x_1, \ldots, x_\ell\},$$

where $x_1 < \ldots < x_\ell$, and let $\alpha_{i,j}^h$ denote the index such that $x_{\alpha_{i,j}^h} = h$.

4. User h encrypts each m_i^h using the s keys $k_{B_{i,j}^h, \alpha_{i,j}^h}$, by defining

$$b_{i,j}^h = k_{B_{i,j}^h, \alpha_{i,j}^h} + m_i^h \bmod p,$$

$1 \leq i \leq r$, $1 \leq j \leq s$.

5. User h broadcasts the vector

$$b^h = (b^h_{1,1}, \ldots, b^h_{1,s}, \ldots, b^h_{r,1}, \ldots, b^h_{r,s}).$$

Probably the only aspect of the scheme that requires explanation are the values $\alpha^h_{i,j}$. Their function is to ensure that every $k_{A,j}$ is used to encrypt exactly one of the m^h_i's. The proof that every privileged user can compute m_P, and that no forbidden set of w users can compute any information about m_P, are essentially the same as those used in studying the OTBES scheme in Section 4.1.

We now present an example to illustrate the protocol.

Example 2. Suppose that $t = 5$ and $\ell = 3$. Note that $5 \equiv 1 \bmod 2$. Suppose that the privileged set $P = \{1, 2, 3, 4, 5\}$.

For each user $i \in P$, we partition the 2-subsets of $P\backslash\{i\}$ into $r = 3$ disjoint parallel classes. This can be done as follows:

$$\begin{array}{lll}
C^1_1 = \{\{2,3\},\{4,5\}\} & C^1_2 = \{\{2,4\},\{3,5\}\} & C^1_3 = \{\{2,5\},\{3,4\}\} \\
C^2_1 = \{\{1,3\},\{4,5\}\} & C^2_2 = \{\{1,4\},\{3,5\}\} & C^2_3 = \{\{1,5\},\{3,4\}\} \\
C^3_1 = \{\{1,2\},\{4,5\}\} & C^3_2 = \{\{1,4\},\{2,5\}\} & C^3_3 = \{\{1,5\},\{2,4\}\} \\
C^4_1 = \{\{1,2\},\{3,5\}\} & C^4_2 = \{\{1,3\},\{2,5\}\} & C^4_3 = \{\{1,5\},\{2,3\}\} \\
C^5_1 = \{\{1,2\},\{3,4\}\} & C^5_2 = \{\{1,3\},\{2,4\}\} & C^5_3 = \{\{1,4\},\{2,3\}\}.
\end{array}$$

Let's look at the computations to be performed by the users in P. First, each user h picks three random values (his or her part of the key) $m^h_1, m^h_2, m^h_3 \in \mathbf{Z}_p$. Next, he or she computes the relevant α values. These are as follows:

$$\begin{array}{llllll}
\alpha^1_{1,1} = 1 & \alpha^1_{1,2} = 1 & \alpha^1_{2,1} = 1 & \alpha^1_{2,2} = 1 & \alpha^1_{3,1} = 1 & \alpha^1_{3,2} = 1 \\
\alpha^2_{1,1} = 2 & \alpha^2_{1,2} = 1 & \alpha^2_{2,1} = 2 & \alpha^2_{2,2} = 1 & \alpha^2_{3,1} = 2 & \alpha^2_{3,2} = 1 \\
\alpha^3_{1,1} = 3 & \alpha^3_{1,2} = 1 & \alpha^3_{2,1} = 2 & \alpha^3_{2,2} = 2 & \alpha^3_{3,1} = 2 & \alpha^3_{3,2} = 2 \\
\alpha^4_{1,1} = 3 & \alpha^4_{1,2} = 2 & \alpha^4_{2,1} = 3 & \alpha^4_{2,2} = 2 & \alpha^4_{3,1} = 2 & \alpha^4_{3,2} = 3 \\
\alpha^5_{1,1} = 3 & \alpha^5_{1,2} = 3 & \alpha^5_{2,1} = 3 & \alpha^5_{2,2} = 3 & \alpha^5_{3,1} = 3 & \alpha^5_{3,2} = 3
\end{array}$$

This determines the values that are broadcast:

$$\begin{aligned}
b^1 &= (m^1_1 + k_{\{1,2,3\},1}, m^1_1 + k_{\{1,4,5\},1}, m^1_2 + k_{\{1,2,4\},1} \\
&\quad m^1_2 + k_{\{1,3,5\},1}, m^1_3 + k_{\{1,2,5\},1}, m^1_3 + k_{\{1,3,4\},1}) \\
b^2 &= (m^2_1 + k_{\{1,2,3\},2}, m^2_1 + k_{\{2,4,5\},1}, m^2_2 + k_{\{1,2,4\},2} \\
&\quad m^2_2 + k_{\{2,3,5\},1}, m^2_3 + k_{\{1,2,5\},2}, m^2_3 + k_{\{2,3,4\},1}) \\
b^3 &= (m^3_1 + k_{\{1,2,3\},3}, m^3_1 + k_{\{3,4,5\},1}, m^3_2 + k_{\{1,3,4\},2} \\
&\quad m^3_2 + k_{\{2,3,5\},2}, m^3_3 + k_{\{1,3,5\},2}, m^3_3 + k_{\{2,3,4\},2}) \\
b^4 &= (m^4_1 + k_{\{1,2,4\},3}, m^4_1 + k_{\{3,4,5\},2}, m^4_2 + k_{\{1,3,4\},3} \\
&\quad m^4_2 + k_{\{2,4,5\},2}, m^4_3 + k_{\{1,4,5\},2}, m^4_3 + k_{\{2,3,4\},3}) \\
b^5 &= (m^5_1 + k_{\{1,2,5\},3}, m^5_1 + k_{\{3,4,5\},3}, m^5_2 + k_{\{1,3,5\},3} \\
&\quad m^5_2 + k_{\{2,4,5\},3}, m^5_3 + k_{\{1,4,5\},3}, m^5_3 + k_{\{2,3,5\},3})
\end{aligned}$$

Now let us see how user 4 will compute the key. This user knows the values of $k_{\{1,2,4\}}$, $k_{\{1,3,4\}}$, $k_{\{1,4,5\}}$, $k_{\{2,3,4\}}$, $k_{\{2,4,5\}}$ and $k_{\{3,4,5\}}$. These can be used to compute values m_i^h as follows:

$k_{\{1,2,4\}}$ can be used to compute m_2^1, m_2^2,
$k_{\{1,3,4\}}$ can be used to compute m_3^1, m_2^3,
$k_{\{1,4,5\}}$ can be used to compute m_1^1, m_3^5,
$k_{\{2,3,4\}}$ can be used to compute m_3^2, m_3^3,
$k_{\{2,4,5\}}$ can be used to compute m_1^2, m_2^5, and
$k_{\{3,4,5\}}$ can be used to compute m_1^3, m_1^5.

Finally, user 4 knows the values m_1^4, m_2^4 and m_3^4 since he or she chose them. So user 4 can determine all 15 components of the key.

8 The Entropy of our Key Distribution Schemes

We look now at our construction for (t, w)-NROTKDS. The key is a random element of $(\mathbf{Z}_p)^{tr}$, where $r = \binom{t-2}{\ell-2}$. Simple calculations yield the following.

Theorem 5. *Let ℓ be an integer such that $1 \leq \ell \leq t$. Then there exists a (t, w)−NROTKDS with*

$$r_i = \frac{\ell\binom{t+w-1}{\ell-1}}{t\binom{t-2}{\ell-2}},$$

$1 \leq i \leq n$, and $r_{B_P} = (t-1)/(\ell-1)$.

We mentioned already that the Beimel-Chor construction is the special case $\ell = 2$. In this case, we get

$$r_i = 2 + \frac{2(w-1)}{t}.$$

For any values of t and w, it is always the case in the Beimel-Chor scheme that $r_i \geq 2$. We observe that by using larger values of ℓ, we can sometimes obtain values of r_i very close to 1 (it is easy to see that $r_i \geq 1$ in any scheme, so the values are close to optimal).

As an illustration of a class of examples where this can be done, let's consider the case $w = 1$ in more detail. When $w = 1$, we see that

$$r_i = r_i(\ell) = \frac{\ell(t-1)}{(\ell-1)(t-\ell+1)}.$$

Elementary algebra shows that $r_i(\ell+1) \geq r_i(\ell)$ if and only if $t \leq \ell^2 + \ell - 1$. In the case where $t = \ell^2 + \ell - 1$ for some integer ℓ, the following is obtained.

Theorem 6. *For any $\ell \geq 2$, there exists an $(\ell^2 + \ell - 1, 1)$−NROTKDS in which $r_i = 1 + \frac{2}{\ell}$, for $1 \leq i \leq n$.*

9 Comments

The long version of this paper (available from the authors) contains all omitted proofs. In addition, it contains an alternate construction for OTBES and IKDS, based on polynomial interpolation, that does not have any congruential condition on the parameter ℓ. (The communication and storage requirements of the alternate construction are identical to the construction we give here, but the alternate construction is more complicated computationally.)

Acknowledgments

This research was done while the second author was visiting the Computer Science and Engineering Department at the University of Nebraska. He would like to thank the department for their hospitality.

Research of C. Blundo is partially supported by Italian Ministry of University and Scientific Research (M.U.R.S.T.) in the framework of the project: "Algoritmi, Modelli di Calcolo e Strutture Informative" and by National Council of Research (C.N.R.); research of Luiz Frota Mattos is partially supported by the National Council of Scientific and Technological Development (CNPq), Brazil, under grant 260030/94-5; and research of D. R. Stinson is supported by NSF grant CCR–9402141.

References

1. A. Beimel and B. Chor. Interaction in Key Distribution Schemes. In "Advances in Cryptology — CRYPTO 93", *Lecture Notes in Computer Science* **773** (1994), 444–457.
2. A. Beimel and B. Chor. Communication in Key Distribution Schemes. *IEEE Transactions on Information Theory* **42:1** (1996), 19–28. [This is a revised and expanded journal version of [1].]
3. R. Blom. An Optimal Class of Symmetric Key Generation Systems. In "Advances in Cryptology — EUROCRYPT '84", *Lecture Notes in Computer Science* **209** (1984), 335–338.
4. C. Blundo, A. De Santis, A. Herzberg, S. Kutten, U. Vaccaro and M. Yung. Perfectly Secure Key Distribution for Dynamic Conferences. In "Advances in Cryptology — CRYPTO '92", *Lecture Notes in Computer Science* **740** (1993), 471–486.
5. C. Blundo and A. Cresti. Space Requirements for Broadcast Encryption. In "Advances in Cryptology — EUROCRYPT '94", *Lecture Notes in Computer Science* **950** (1995), 287–298.
6. M. Burmester and Y. Desmedt. A Secure and Efficient Conference Key Distribution System. In "Advances in Cryptology — EUROCRYPT '94", *Lecture Notes in Computer Science* **950** (1995), 275–286.
7. T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, 1991.
8. A. Fiat and M. Naor. Broadcast Encryption. In "Advances in Cryptology — CRYPTO '93", *Lecture Notes in Computer Science* **773** (1994), 480–491.

9. J. H. van Lint and R. M. Wilson. *A Course in Combinatorics*. Cambridge University Press, 1992.
10. T. Matsumoto and H. Imai. On a Key Predistribution System — A Practical Solution to the Key Distribution Problem. In "Advances in Cryptology — CRYPTO '87", *Lecture Notes in Computer Science* **293**, (1988), 185–193.
11. D. R. Stinson. On Some Methods for Unconditionally Secure Key Distribution and Broadcast Encryption. Preprint, (1995).