Secret Sharing with Public Reconstruction (extended abstract)

Amos Beimel* and Benny Chor **

Department of Computer Science
Technion, Haifa 32000, Israel

Abstract. All known constructions of information theoretic t -out-of- n secret sharing schemes require *secure, private* communication channels among the parties for the reconstruction of the secret. In this work we investigate the cost of performing the reconstruction over *public* communication channels. A naive implementation of this task distributes $O(n)$ one times pads to each party. This results in shares whose size is $O(n)$ times the secret size. We present three implementations of such schemes that are substantially more efficient:

- A scheme enabling multiple reconstructions of the secret by different subsets of parties, with factor $O(n/t)$ increase in the shares' size.
- A one-time scheme, enabling a single reconstruction of the secret, with $O(\log(n/t))$ increase in the shares' size.
- A one-time scheme, enabling a single reconstruction by a set of size *exactly* t , with factor $O(1)$ increase in the shares' size.

We prove that the first implementation is optimal (up to constant factors) by showing a tight $\Omega(n/t)$ lower bound for the increase in the shares' size.

1 Introduction

Secret sharing schemes were introduced by Blakley [7] and Shamir [19], and were the subject of a considerable amount of work, e.g. [18, 15, 16, 4, 20]. In these schemes, a dealer holds a secret piece of information. Upon system initialization, the dealer gives one share of the secret to each of n parties. These shares are distributed privately, and are kept by each party in a secure way. Later on, any authorized subset (a subset containing at least t parties) of the parties collect their shares, and use them to reconstruct the secret. All known schemes that guarantee information theoretic secrecy require the use of secure, private communication channels between the parties that participate in the reconstruction.

The question we raise in this work is whether reconstruction can be done without assuming that the channels are secure, while maintaining the security of the schemes. We consider the worst case scenario: The "bad" parties can overhear any communication, so from their point of view the channels are public. On the

* email: beimel@cs.technion.ac.il, URL: <http://www.cs.technion.ac.il/~beimel>

** email: benny@cs.technion.ac.il, URL: <http://www.cs.technion.ac.il/~benny>

other hand, "good" parties hear only messages sent to them. (In particular, from the point of view of the "good guys", the channels do not carry any of the potential advantages of a broadcast channel.)

The simplest way to implement such public reconstruction securely is to hand to each party upon system initialization, in addition to his original share, $2(n-1)$ one time pads. These pads are used in order to simulate a private channel on a public one. In the private channel scenario, reconstruction is typically done by exchanging shares among parties. To enable such exchange with every other participant, each party will need two pads per participant: one for receiving a share, and one for sending the share. Thus the simple implementation results in $O(n)$ multiplicative factor increase in the size of each share.

We design substantially more efficient schemes of three types. The first type is *unrestricted schemes*. In these schemes, any number of authorized sets (each containing at least t parties) may reconstruct the secret, after communicating on the public channel. Any disjoint coalition of at most $t-1$ parties, does not gain any partial information on the secret, given the coalition's shares and the the communication of the sets that reconstructed the secret. We describe unrestricted schemes in which the size of the shares is $O(n/t)$ times the size of the original secret. We complement this result by proving a tight $\Omega(n/t)$ lower bound on the increase in the shares' size for any unrestricted scheme.

In order to participate in more than one reconstruction, every party that has already reconstructed the secret must store the secret. This is problematic in applications where an adversary might break into the computer of the secret holder. (One of the advantages of traditional secret sharing is that breaking into the computer of a "share holder" does not compromise the secret.) The unrestricted non-reactive schemes of Section 5 solve this problem, but the share size there is n times the secret size.

The second type is *one time schemes*, in which only a single authorized set (containing at least t parties) will reconstruct the secret. It is not known during system initialization which set will reconstruct the secret, and the dealer has to accommodate any possible set. For example, these schemes can be used to enable one time activities like the firing of a ballistic missile or opening of a sealed safe. We describe one-time schemes in which the size of the shares is $O(\log(n/t))$ times the original secret size. It is an open problem if this bound is tight for one-time schemes. Finally, we consider one time schemes where one authorized set of size *exactly* t will reconstruct the secret. Additional parties in supersets with more than t parties may not reconstruct the secret, because communicating it from members of the authorized set over the public channel is not possible in a secure way. This means that the authorized sets that can securely reconstruct the secret do not form a *monotone* access structure. We design such schemes with just $O(1)$ multiplicative increase in the share size (for any threshold t).

In light of our results, one may wonder if the initial distribution of shares can also be done over public channels. By the properties of "regular" schemes, each participant requires a share whose conditional mutual information with the secret (given $t-1$ shares) is at least the entropy of the secrets [15]. This conditional

entropy cannot be increased by communicating over public channels [17, 1]. Thus in our model, it is necessary to have secure initial distribution of shares from the dealer to the participants.

Some bibliographical remarks: A similar setting of public interaction was considered for interactive key distribution schemes (e.g. [10, 2]). Our schemes employ key distribution schemes, though not interactive ones. Another solution to eliminating the use of secure private channels assumes that the parties have limited computing power. A common assumption is that the parties are probabilistic polynomial-time Turing machines, and the security of the channels is achieved by means of public key cryptography [12, 13]. Public channels have been used in secret sharing (in addition to private channels) in dynamic sharing of secrets. These are schemes where the dealer enables parties to reconstruct different secrets in different time instants (e.g. [20, 6, 9]). A different scenario in which a public broadcast channel is used (in addition to private channels) is to protect against Byzantine parties [3]. Unlike our scenario, in that work the broadcast channel is heard by *all* parties.

The rest of this paper is organized as follows: In Section 2 we define the model, secret sharing schemes, and key distribution schemes. Section 3 describes the unrestricted schemes, and Section 4 the one time schemes. In Section 5 we introduce non-reactive, unrestricted schemes. Finally, Section 6 provides lower bounds for unrestricted schemes.

2 Definitions

In this section we define our model, secret sharing schemes (traditional and public channels), and key distribution schemes. We consider a system with n parties denoted by $\{P_1, P_2, \dots, P_n\}$. In addition to the parties, there is a dealer in the system, who has a secret input s . A *scheme* is a probabilistic mapping, which the dealer applies to the input, and generates n pieces of information. These pieces of information are called shares, and the i -th piece is called the share of P_i . For every i , the dealer gives the i -th share to P_i . The dealer is only active in this initial stage. After the initial stage, the parties can communicate, according to some pre-defined, possibly randomized, protocol. The parties are honest, that is, they follow their protocols. However, they are curious and after the protocol has ended some of them can collude and try to gain some partial information on the secret.

Definition 1. Let B be a (bad) coalition (set of parties). The *view* of B , denoted by VIEW_B , after an execution of a protocol is all the information it has, i.e. the shares of the parties in the coalition, and the messages exchanged by *all* parties over the communication channels. (Here the insecurity of the communication is manifested.) The coalition B has *no information* on a random variable X if for every two possible values x_1, x_2 of X :

$$\Pr[\text{VIEW}_B \mid X = x_1] = \Pr[\text{VIEW}_B \mid X = x_2],$$

where the probability is taken over the random inputs of the dealer, and the random inputs of the members outside the coalition. Notice that we do not make any assumptions on the distribution of X .

We define both traditional secret sharing scheme, i.e with private channels, and secret sharing schemes with public reconstruction.

Definition 2. Let S be a finite set of secrets. A t -out-of- n secret sharing scheme is a scheme, in which the input is a secret taken from S , and which satisfies the following two conditions:

Reconstructability: Any set of parties whose size is at least t can reconstruct the value of the secret after communicating among them self. Any party in the reconstructing set gets the value of the secret with certainty.

Security: Every disjoint coalition B of size at most $t - 1$ has no information on the secret as defined in Definition 1. There are three variates we consider:

1. *Traditional* secret sharing schemes in which the reconstruction takes place via secure, private channels. In this case the view of a disjoint coalition is its shares.
2. *Unrestricted* secret sharing scheme with public reconstruction in which a coalition B can hear all communications that took place. The security is guaranteed even if any collection of sets (maybe even all) will reconstruct the secret using the public channel. In this case the view of a disjoint coalition is its shares and all the communications that took place.
3. *One-time* scheme in which the security is guaranteed only if one set will reconstruct the secret. In this case the view of a disjoint coalition is its shares and a communication of one reconstructing set.

The security should hold for any coalition of at most $t - 1$ parties. As a special case ($B = \emptyset$), a listener who heard all communications but has no shares should gain no partial information about the secret.

Shamir [19] presented a traditional scheme in which the size of the shares is the same as the size of the secrets (for domains of secrets which contain at least $n + 1$ secrets). The domain of shares in Shamir's scheme is the smallest possible, since the size of the share has to be at least as large as the size of the secrets [15]. In traditional secret sharing schemes, while one set reconstructs the secret, no information is leaked to disjoint coalitions (due to the security of the channels). Hence, these schemes are always unrestricted. Furthermore, in traditional schemes, if a set can reconstruct the secret, then every superset of the set can reconstruct the secret. However, secret sharing schemes with public reconstruction do not necessarily have this monotone property. We require that every party of the superset should know the reconstructed secret. However, it is not necessarily possible to "distribute" the secret to members of a superset without leaking information on it to other parties.

We describe unrestricted, non-interactive key distribution schemes. (Formal definition can be found in [10, 2].) These schemes are used in the constructions of the schemes with public reconstruction.

Let b be a positive integers such that $b \leq n - 2$, and K be a set of keys. A $(2, b)$ *key distribution scheme* with n users and domain of keys K is a scheme in which a dealer (who has no input) generates n shares such that the following requirements hold:

Every pair of parties has a key, which is uniformly distributed. Each member of the pair can deterministically reconstruct G 's key from his share. Any "bad" coalition "B" of cardinality at most b gets no information on the key of any disjoint pair. In this case their view is the collection of their pieces.

Consider a $(2, 2t-3)$ key distribution scheme, a coalition B of $t-1$ parties, and a disjoint set G of t parties. It holds that from the point of view of the coalition, the $\binom{2t-3}{2}$ keys of pairs of parties in G are distributed uniformly and independently (for proof see [2]). Blom [8] constructed efficient $(2, b)$ -key distribution schemes. For every prime-power q (where $q \geq n$) he presented a scheme in which the keys are taken from $\text{GF}(q)$ and the shares are taken from $\text{GF}(q)^{b+1}$.

3 Unrestricted Schemes

In this section we construct unrestricted secret sharing schemes with public reconstruction in which the size of the share of every party is $O(n/t)$ times the size of the secret. We first describe a simple scheme in which the size of the shares is $O(n)$ times the size of the secret. In this scheme, the dealer shares the secret using Shamir's secret sharing scheme [19]. The dealer also deals to every pair of parties two random strings whose size is the same as the size of the secret. These two random strings, which we call keys, are given to the two parties of the pair, and will be used as one-time pads. Overall, every party receives $2(n-1)$ keys, each one with the same size as the secret. When the parties of a set of size at least t wish to reconstruct the secret, all the parties "send" their shares to the "leader" of the set, say the party with minimal index in the set. The leader gets at least t shares, which enable him to reconstruct the secret. Then, the leader "sends" the secret to the other parties. The parties use their keys as one time pads to simulate the private channels. Specifically, let P_{i_0} be the party with smallest identity in the set. Every party P_i , holding the share s_i from Shamir's scheme, adds s_i and the first key of the pair $\langle P_{i_0}, P_i \rangle$ and sends this sum on the public channel. The party P_{i_0} can reconstruct all the shares from these messages, and therefore reconstruct the secret. Now, P_{i_0} sends messages, one message to every party in the reconstruction set. For every party P_i , he sums the secret and the second key of the pair $\langle P_{i_0}, P_i \rangle$ and sends this sum on the public channel. Since the one-time pads are independent, coalitions of parties disjoint to the reconstructing set do not gain any information on the shares or the secret. Furthermore, even if many reconstructions took place, this will not leak any information to a disjoint set.

Suppose P_{i_0} is the leader in a set of size at least t . In the previous scheme, during the reconstruction for this set, only the keys that were given to P_{i_0} were used. To improve the space efficiency we will use all the keys of the parties in the reconstructing set. Following [2], we partition the secret into t sub-secrets,

and share each sub-secret using Shamir's scheme. Now we choose t parties of the reconstructing set, and each one will be responsible for reconstructing one sub-secret. Each party will act as the leader in the previous scheme. That is, every leader receives shares only from the other $t-1$ leaders (this is enough!), but sends his sub-secret (after reconstruction) to every member of the reconstructing set. This way we can handle t sub-secrets "at the price of one". The domain of the secrets in the scheme is $\text{GF}(q)^t$, where q is a prime-power such that $q > n$. In the scheme we view the secret as t secrets from $\text{GF}(q)$.

Unrestricted Secret Sharing Scheme

Distribution stage:

Input: t secrets $s_1, s_2, \dots, s_t \in \text{GF}(q)$

Shares:

Share each s_i using Shamir's scheme for every i , where $1 \leq i \leq t$.

Denote the n shares of secret s_i by $s_{i,1}, s_{i,2}, \dots, s_{i,n}$.

For every pair of parties generate 4 independent keys from $\text{GF}(q)$.

Denote the keys of $\langle P_i, P_j \rangle$ by $k_{i,j}^1, k_{i,j}^2, k_{i,j}^3, k_{i,j}^4$.

The share of P_i is $s_{1,i}, \dots, s_{t,i}$ and $k_{i,j}^1, k_{i,j}^2, k_{i,j}^3, k_{i,j}^4$ for $1 \leq j \leq n$.

Reconstruction stage:

A set $G = \{P_{i_1}, \dots, P_{i_\ell}\}$ that wants to reconstruct the secret ($\ell \geq t$).

Every $P_{i_j} \in G$ announces if he has previously reconstructed the secret.

Let P_{i_j} for $1 \leq j \leq t$ be the leaders of G .

Each leader P_{i_j} ($1 \leq j \leq t$) sends (at most) $t-1$ messages to other leaders that have not previously reconstructed the secret:

$$s_{i_j,j'} + k_{i_j,i_{j'}}^1 \text{ to } P_{i_{j'}}, \text{ for } 1 \leq j' < j$$

$$s_{i_j,j'} + k_{i_j,i_{j'}}^2 \text{ to } P_{i_{j'}}, \text{ for } j < j' \leq t$$

Each leader P_{i_j} ($1 \leq j \leq t$) computes s_j from $s_{j,i_1}, \dots, s_{j,i_t}$.

Each leader P_{i_j} sends a message to every $P_{i_{j'}} \in G$ that has not previously reconstructed the secret:

$$s_j + k_{i_j,i_{j'}}^3 \text{ to } P_{i_{j'}}, \text{ for } 1 \leq j' < j$$

$$s_j + k_{i_j,i_{j'}}^4 \text{ to } P_{i_{j'}}, \text{ for } j < j' \leq \ell$$

Each party concatenates the sub-secrets $s_1, \dots, s_t$ to obtain the secret.

Figure 1: Unrestricted scheme.

As described, the scheme has two technical points we should elaborate. The first is the fact that in one reconstruction two parties P_i and P_j might need to exchange 4 different messages. This is the reason for giving them 4 common keys. The second difficulty is that in different reconstructions the same party can be responsible for different sub-secrets. This means that P_i will have to send to P_j two different messages, using the same key as a one time pad. This might leak information to disjoint coalitions. Therefore, every party that participated in one reconstruction will remember the secret, and in latter reconstructions will inform other parties (in the clear) that he need not receive new messages. It does not

prevent a leader from sending all messages that he has to send according to the scheme, since these message either depend on his share, or on the secret. (The party need only remember the secret, and *not* the messages that he heard.) Thus, every key is used as a one-time pad at most once (in the first reconstruction that the pair participates together). Therefore, the scheme satisfies the unrestricted security requirement. A detailed code of the scheme appears in Figure 1.

Let us calculate the size of the share of every party in the unrestricted scheme. Each party is given t shares generated by Shamir's scheme for secrets taken from $GF(q)$. The dealer also distributes to each party $4(n-1)$ keys taken from $GF(q)$. Hence, each share contains $(4n+t-4)$ elements from $GF(q)$, compared to t elements from $GF(q)$ for the secret. We summarize these results in the the next theorem.

Theorem 3. *Let q be a prime-power such that $q > n$. The above mentioned scheme is an unrestricted t -out-of- n secret sharing scheme with public reconstruction for secrets taken from $GF(q)^t$. The share of each party is an element of $GF(q)^{4n+t-4}$. So the size of each share is $1 + 4(n-1)/t$ times the size of the secrets.*

4 One-Time Schemes

In the unrestricted scheme, we need totally independent keys in order to guarantee the security of the scheme during repeated reconstructions. In this section we deal with the scenario where the secret is going to be reconstructed only once. For example, to enable the firing of a ballistic missile or opening of a sealed safe. In this case, total independence among the keys is not needed, and weaker independence requirements suffice. Shares can therefore be taken from a smaller sample space, which translates into smaller size shares. Specifically, we use Blom's key distribution scheme [8] for this purpose.

The first scheme we present enables one-time reconstruction of the secret by sets of size *exactly* t . The size of the shares is a constant (less than 10) times the size of the secret, namely only $O(1)$ increase in shares' size. We employ this "exactly t " scheme as a building block for "at least t " schemes. We use $\log(n/t)$ independent instances of "exact schemes" for thresholds $t, 2t, 4t, \dots$ up to n , and an additional instance of size t . Now, given any set G with ℓ parties ($\ell \geq t$), we represent it as a union of subsets (not necessary disjoint) with cardinalities $t, 2t, 4t, \dots$ – at most two sets subsets of cardinality t and at most one subset of cardinality $2^i t$ for each $i \geq 1$. The secret is now separately reconstructed by each subset. Any member of G takes part in at least one of these reconstructions, and thus learns the secret. On the other hand, any disjoint coalition containing at most $t-1$ parties gets no partial information on the secret from any single instance. Due to the independence of the instances, this remain valid with respect to the joint reconstructions. We get a one-time scheme for set of size at least t , with just $O(\log(n/t))$ increase in share size. We now describe in detail the "exact t " scheme. The distribution phase is depicted in Fig. 2.

Distribution in Exactly t -out-of- n one-time scheme

Input: secret $s \in \text{GF}(q)^t$.

Consider the secret as t secrets $s_1, \dots, s_t \in \text{GF}(q)$.

Share each secret s_i using Shamir's t -out-of- n secret sharing scheme.

Let $b = \min\{2t - 3, n - 2\}$.

Generate shares using $(2, b)$ -key distribution scheme with key domain $\text{GF}(q)^4$ (which we consider as 4 keys from $\text{GF}(q)$).

Share of P_j : the j -th share of each s_i ,
and the share of the key distribution scheme.

Figure 2: Exactly t -out-of- n one-time scheme.

The reconstruction is done exactly as in the unrestricted scheme. The security of one reconstruction of a set of exactly t parties follows from the property of $(2, 2t - 3)$ key distribution schemes discussed in Section 2: Given the shares of any disjoint coalition of at most $t - 1$ parties, the keys held by any set of size t are distributed uniformly and independently. Thus, when used as one-time pads, the reconstruction is secure (using the same arguments as in the unrestricted case). This scheme uses t shares of Shamir's t -out-of- n secret sharing scheme with secrets taken from $\text{GF}(q)$. In addition, each party gets a share of a $(2, 2t - 3)$ key distribution scheme with keys taken from $\text{GF}(q)^4$ and with shares taken from $\text{GF}(q)^{4(2t-2)}$. Overall, each share contains $(9t - 8)$ elements from $\text{GF}(q)$ (if $2t > n + 1$, then the shares are even shorter). Recall that the secret is taken from $\text{GF}(q)^t$, and therefore the size of the share is less than 9 times the size of the secrets.

In this previous scheme the domain of secrets has to be $\text{GF}(q)^t$ (for some prime-power q). Restricting the domain of the secret to such cardinality can cause problems when we employ simultaneously many schemes with the same secret but with different thresholds. To overcome this, given any domain of secrets we consider a slightly bigger domain whose size (which can depend on the threshold) is of the desired form. That is, given a secret of size m which is at least $t \log n$, we choose a prime power q such that $m \leq t \log q$, and use the previous scheme with secrets of size $m' = t \log q$. Choosing the smallest prime-power satisfying these conditions, we have $m' \leq m + t \leq 2m$. Thus,

Theorem 4. *Let m be a natural number such that $m > 9t \log n$. There exists a one-time sharing scheme with public reconstruction for exactly t -out-of- n , in which the size of the secret equals m , and the size of the share of each party is less than 10 times the size of the secrets.*

We next describe the one time scheme in which every set of *at least* t parties can securely reconstruct the secret.

One-time Secret Sharing Scheme

Distribution stage:

Input: secret s of size m

Share the secret s using two independent copies of a one time exactly t -out-of- n secret sharing schemes.

For every i , $1 \leq i < \log(n/t)$:

Share s with an exactly $2^i t$ -out-of- n one time secret sharing scheme.

Reconstruction stage:

A set $G = \{P_{i_1}, \dots, P_{i_\ell}\}$ that wants to reconstruct the secret ($\ell \geq t$).

Cover the set G by (possibly intersecting) sets of size $2^i t$

(at most one set for every $i > 1$, and at most 2 sets of size t).

Each set of size $2^i t$ independently reconstructs the secret using the shares of the exactly $2^i t$ -out-of- n secret sharing scheme.

Figure 3: t -out-of- n one-time scheme for every set.

Theorem 5. *The scheme of Figure 3 is a one-time t -out-of- n secret sharing scheme with public reconstruction in which every set of parties of size at least t can securely reconstruct the secret. If the size of the secrets m is larger than $9n \log n$, then the size of the shares of every party is less than $10(\log(n/t) + 1)$ times the size of the secrets.*

Remark If we require that the size of the secret m is greater than $n^2 \log n$, then we can construct a scheme in which the size of the shares is only $2 \log(n/t) + O(1)$ times the size of the secret, i.e a smaller leading constant.

5 Unrestricted Non-Reactive Schemes

A secret sharing scheme with public reconstruction is called *non-reactive* if the messages sent by each party depend only on his share (and not on messages received during the reconstruction). Non-reactive schemes are simpler to implement, as they require less synchronization. Therefore, they are desirable from practical point of view. In this section we present non-reactive, unrestricted t -out-of- n schemes. The size of the shares in these schemes is n times the size of the secret. This represents a slight improvements (by a factor of 2) over the reactive scheme of Section 3 for $t = 2$, but is strictly less efficient (in terms of share size) for $t \geq 5$. We extend these schemes to general access structures. The size of the share in our public reconstruction schemes is n times the size of the share in the original scheme. This is typically not a significant increase, as the best schemes for most access structures to date require shares whose size is exponential in n .

We first present a simple, non-reactive, 2-out-of- n secret sharing scheme. Let $s \in \mathcal{Z}_m$ be the secret which the dealer wants to share. The dealer chooses

n independent random elements from $\mathcal{Z}_m$, denoted $r_1, \dots, r_n$. The share of P_i is $r_1, \dots, r_{i-1}, r_i + s, r_{i+1}, \dots, r_n$. Each share is uniformly distributed in $\mathcal{Z}_m^n$, regardless of the secret. Hence, prior to any reconstruction every party has no information on the secret (as defined in Definition 1). To reconstruct the secret, P_i sends the message r_j , and P_j sends the message r_i . Now, P_i , who holds $r_i + s$, hears the message r_i , so he can reconstruct the secret. Every third party hears messages that he already knows, and gains no information on the secret. That is, the reconstruction is secure. The size of the shares in this scheme is n times the size of the secret. During the reconstruction in this scheme every party is deterministic and sends only one message that depends only on its share.

In general secret sharing schemes scenario, first suggested by [14], we are given a collection $\mathcal{A}$ of sets of parties called an *access structure*. We require that every set in $\mathcal{A}$ can reconstruct the secret, while every set not in $\mathcal{A}$ does not know anything about the secret. It follows that secret sharing schemes can exist only for monotone collections. Indeed, it is known that for every monotone collection there exists a traditional secret-sharing scheme [14, 5, 21]). However, the size of the shares in these schemes is typically exponential in the number of parties (i.e., of size $m2^{\Theta(n)}$ where n is the number of parties in the system and m is the size of the secret). Let $\mathcal{A}$ be any monotone access structures. The unrestricted, non-reactive, 2-out-of- n scheme can be generalized to an unrestricted, non-reactive scheme realizing the access structure $\mathcal{A}$, with the following properties:

Theorem 6. *Assume there exists a (traditional) secret sharing scheme realizing $\mathcal{A}$ with domain of secrets S and domain of shares U . Then there exists an unrestricted, non-reactive secret sharing scheme realizing $\mathcal{A}$ with public reconstruction for secrets taken from S . The share of each party is an element of $S \times U^{n-1}$. So the size of each share is at most n times the size of the shares in the original scheme.*

Corollary 7. *Let q be a prime-power such that $q > n$. There exists an unrestricted, non-reactive t -out-of- n secret sharing scheme with public reconstruction for secrets taken from $GF(q)$. The share of each party is an element of $GF(q)^n$. So the size of each share is n times the size of the secret.*

For most known schemes [5, 21], it is possible to design unrestricted *reactive* schemes with just an *additive* factor of n times the *secret* size (in these schemes it suffices for a party to send a message of the size of the secret, instead of his entire share). This is typically much better, as the shares tend to be much larger than the secret for general access structures. Additional details will be given in the final version.

6 Lower Bounds for Unrestricted Schemes

In this section we prove a $\Omega(n/t)$ lower bound on the increase in the shares' size for unrestricted t -out-of- n schemes. For $t = 2$ this lower bound is tight by the non-reactive scheme of Section 5. For $t > 2$ this lower bound is tight up to

a constant factor (by the reactive scheme of Section 3). We first prove a $\Omega(n)$ lower bound on increase in size of shares for 2-out-of- n schemes. Then, we show that this lower bound translates into $\Omega(n/t)$ increase for t -out-of- n schemes.

We start with the lower bound for $t = 2$. The proof uses entropy and mutual information. For definitions of these information theoretic terms, the reader can refer to [11]. We assume an arbitrary probability distribution on the secret. The intuition behind the proof is that P_i has to expose $H(S)$ extra bits of his share in each reconstruction. Finally, after all reconstructions, the uncertainty of S_1 has to remain at least $H(S)$, as an outsider who listened to all reconstructions still has $H(S)$ uncertainty on the secret. Since, P_1 participates in $n - 1$ reconstructions, the original entropy of the share has to be at least $n \cdot H(S)$.

Without loss of generality, we prove the claim for P_1 . To prove the lower bound on P_1 's share, we only use the requirement that P_1 can reconstruct the secret together with every other P_j (we do not care if other pairs can or cannot reconstruct the secret). We start with some notation. Denote by S_i the share given to P_i in the initial distribution phase, and by $C_{i,j}$ the messages exchanged between P_i and P_j (all these are random variables). We denote $C = C_{1,3} \dots C_{1,n}$, the concatenation of all messages exchanged between P_1 and $P_3, \dots, P_n$. Recall that the communication $C_{1,2}$, together with P_2 's share S_2 , enable P_2 to reconstruct the secret S . On the other hand, the communication C and S_2 give no information (to P_2) about the secret. These facts will imply the next claim.

Claim 8. $H(C_{1,2}|S_2C) \geq H(S)$.

Proof. Since P_2 can reconstruct the secret S , given his share S_2 and the messages $C_{1,2}$ exchanged between P_1 and P_2 , the conditional entropy $H(S|C_{1,2}S_2)$ equals 0. On the other hand, P_2 gets no information about the secret S from his own share S_2 and all messages C exchanged between P_1 and the other $n - 2$ parties. Therefore the conditional entropy $H(S|S_2C)$ equals $H(S)$. Now, consider the conditional mutual information $I(C_{1,2}; S|S_2C)$ of the message $C_{1,2}$ and the secret S , given the share S_2 and C . We have

$$\begin{aligned} H(C_{1,2}|S_2C) - H(C_{1,2}|SS_2C) &= I(C_{1,2}; S|S_2C) \\ &= H(S|S_2C) - H(S|C_{1,2}S_2C) = H(S) \end{aligned}$$

which implies $H(C_{1,2}|S_2C) \geq H(S)$. □

The next claim is the heart of the proof of the lower bound. It states that the mutual information between S_1 and $C_{1,2}$ given the "other" communication C is at least $H(S)$. Intuitively, since P_2 does not know the secret prior to the reconstruction, and knows it after the reconstruction, P_2 has to receive $H(S)$ bits of information which could only originate in S_1 and passed through the communication $C_{1,2}$. Hence, $C_{1,2}$ must contain $H(S)$ bits of information originating from the share S_1 . Claim 9 is stated for deterministic parties – the outgoing messages are determined by the given share and previous incoming messages. An analogous statement, for randomized protocols, will be included in the final version of

this paper. In randomized, outgoing messages can also depend on random local inputs.

Claim 9. *For deterministic reconstruction protocols we have*

$$I(C_{1,2}; S_1|C) = H(S_1|C) - H(S_1|C_{1,2}C) \geq H(S) .$$

Proof. Since P_1 and P_2 are deterministic, and their domain of shares is finite, there is a bound k on the maximum number of communication rounds which take place during the reconstruction of the secret. Denote by A_i the i -th message sent by P_1 to P_2 , and similarly, let B_i be the i -th message sent by P_2 to P_1 . Then, without loss of generality, $C_{1,2} = A_1 B_1 \dots A_k B_k$. The message A_i is determined by the share S_1 and previous messages, that is, $H(A_i|S_1 A_1 B_1 \dots A_{i-1} B_{i-1}) = 0$. The following inequality holds for any deterministic communication protocol:

$$\begin{aligned} H(C_{1,2}|S_1 C) &= H(A_1 B_1 \dots A_k B_k|S_1 C) \\ &= \sum_{i=1}^k (H(A_i|S_1 C A_1 B_1 \dots A_{i-1} B_{i-1}) + H(B_i|S_1 C A_1 B_1 \dots A_{i-1} B_{i-1} A_i)) \\ &= \sum_{i=1}^k H(B_i|S_1 C A_1 B_1 \dots A_{i-1} B_{i-1} A_i) \\ &\leq \sum_{i=1}^k H(B_i|C A_1 B_1 \dots A_{i-1} B_{i-1} A_i) . \end{aligned}$$

Similarly, $H(C_{1,2}|S_2 C) \leq \sum_{i=1}^k H(A_i|C A_1 B_1 \dots A_{i-1} B_{i-1})$. Combining the two inequalities

$$\begin{aligned} H(C_{1,2}|S_1 C) + H(C_{1,2}|S_2 C) &\leq \sum_{i=1}^k H(B_i|C A_1 B_1 \dots A_{i-1} B_{i-1} A_i) \\ &\quad + \sum_{i=1}^k H(A_i|C A_1 B_1 \dots A_{i-1} B_{i-1}) \\ &= H(A_1 B_1 \dots A_k B_k|C) = H(C_{1,2}|C) . \end{aligned}$$

This inequality, together with Claim 8, implies

$$I(C_{1,2}; S_1|C) = H(C_{1,2}|C) - H(C_{1,2}|S_1 C) \geq H(C_{1,2}|S_2 C) \geq H(S) \quad \square$$

Claim 10. *In any unrestricted 2-out-of- n secret sharing scheme with public reconstruction, the share of each participant, S_i , satisfies $H(S_i) \geq n \cdot H(S)$.*

Proof. We first note that by Definition 2 a listener, who overhears all communication involving P_1 , gets no information on the secret. Therefore,

$$H(S|C_{1,2}C_{1,3} \dots C_{1,n}) = H(S) .$$

On the other hand, given P_1 's share, this communication determines the secret, so $H(S|S_1C_{1,2}C_{1,3}\dots C_{1,n}) = 0$. Therefore,

$$\begin{aligned} H(S) &= H(S|C_{1,2}C_{1,3}\dots C_{1,n}) - H(S|S_1C_{1,2}C_{1,3}\dots C_{1,n}) \\ &= I(S; S_1|C_{1,2}C_{1,3}\dots C_{1,n}) \\ &= H(S_1|C_{1,2}C_{1,3}\dots C_{1,n}) - H(S_1|SC_{1,2}C_{1,3}\dots C_{1,n}), \end{aligned}$$

and in particular $H(S_1|C_{1,2}C_{1,3}\dots C_{1,n}) \geq H(S)$. Claim 9 (or analog claim for the case of randomized protocols, which will appear in the final version) states that

$$H(S_1|C_{1,3}\dots C_{1,n}) - H(S_1|C_{1,2}C_{1,3}\dots C_{1,n}) \geq H(S).$$

Similarly it holds that

$$H(S_1|C_{1,4}\dots C_{1,n}) - H(S_1|C_{1,3}C_{1,4}\dots C_{1,n}) \geq H(S).$$

$$\vdots$$

$$H(S_1) - H(S_1|C_{1,n}) \geq H(S).$$

Summing these n inequalities, we conclude that $H(S_1) \geq n \cdot H(S)$. $\square$

We next show that this lower bounds on increase in size of shares for 2-out-of- n schemes translates into $\Omega(n/t)$ increase for t -out-of- n schemes.

Theorem 11. *In every unrestricted t -out-of- n secret sharing scheme with public reconstruction the size of the shares of every party is at least $\lfloor 1 + (n-1)/(t-1) \rfloor$ times the size of the secrets.*

Proof. Consider any t -out-of- n scheme. Denote the party whose share is shortest by P_1 . We construct an unrestricted 2-out-of- $(\lfloor 1 + (n-1)/(t-1) \rfloor)$ scheme in which the entropy of S_1 – the share of P_1 – is the same. Hence, by Claim 10 its entropy is at least $(\lfloor 1 + (n-1)/(t-1) \rfloor)H(S)$. Since the scheme is secure whatever the distribution on the secrets is, we can assume uniform distribution on the secrets. In this case $H(S) = \log|S|$, which is the size of the secret. Since $H(S_1) \leq \log|S_1|$, the size of the share of P_1 is at least $\lfloor 1 + (n-1)/(t-1) \rfloor$ times the size of the secrets.

The construction is simple: the dealer gives P_1 the share of P_1 in the original scheme, and every other party gets shares of $t-1$ disjoint parties. Since every party has at most $t-1$ shares, he does not gain any information on the secret even after hearing communications. On the other hand, every 2 parties have at least t shares, therefore they can communicate on a public channel, and securely reconstruct the secret. $\square$

Acknowledgments We would like to thank Carlo Blundo and Hugo Krawczyk for helpful discussions, and Eyal Kushilevitz for helping us spell Hugo's name.

References

1. R. Ahlswede and I. Csiszar. Common Randomness in Information Theory and Cryptography – Part I: Secret Sharing. *IEEE IT*, 39(4):1121–1132, July 1993.
2. A. Beimel and B. Chor. Interaction in Key Distribution Schemes. To appear in *IEEE IT*. An extended abstract appears in *Crypto '93*, Springer-Verlag, LNCS 773, D.R. Stinson, ed. pp. 444–455.
3. M. Ben-Or and T. Rabin. Verifiable Secret sharing and Multiparty Protocols with Honest Majority. In *Proceeding 21th STOC*, pages 73–85. 1989.
4. J. Benaloh. Secret Sharing Homomorphisms: Keeping Shares of a Secret Secret. In A. M. Odlyzko, editor, *CRYPTO '86*, volume 263 of LNCS, pages 251–260.
5. J. Benaloh and J. Leichter. Generalized Secret Sharing and Monotone Functions. In S. Goldwasser, ed., *CRYPTO '88*, volume 403 of LNCS, pages 27–35.
6. B. Blakley, G. R. Blakley, A. H. Chan, and J. Massey. Threshold Schemes with Disenrollment. In E. F. Brickell, ed., *CRYPTO '92*, vol. 740 of LNCS, pp. 540–548.
7. G. R. Blakley. Safeguarding Cryptographic Keys. In *Proc. AFIPS 1979 NCC*, vol. 48, pages 313–317, June 1979.
8. R. Blom. An Optimal Class of Symmetric Key Generation Systems. In T. Beth, N. Cot, and I. Ingemarsson, editors, *Eurocrypt 84*, volume 209 of LNCS, pages 335–338. Springer-Verlag.
9. C. Blundo, A. Cresti, A. De Santis, and U. Vaccaro. Fully Dynamic Secret Sharing Schemes. In D. R. Stinson, ed., *CRYPTO '93*, vol. 773 of LNCS, pages 110–125.
10. C. Blundo, A. De Santis, A. Herzberg, S. Kuttan, U. Vaccaro, and M. Yung. Perfectly-Secure Key Distribution for Dynamic Conferences. In E. F. Brickell, editor, *CRYPTO '92*, volume 740 of LNCS, pages 471–486. Springer-Verlag, 1993.
11. T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, 1991.
12. W. Diffie and M. E. Hellman. New Directions in Cryptography. *IEEE IT*, 22(6):644–654, 1976.
13. S. Goldwasser and S. Micali. Probabilistic Encryption. *JCSS*, 28(21):270–299, 1984.
14. M. Ito, A. Saito, and T. Nishizeki. Secret Sharing Schemes Realizing General Access Structure. In *Proc. IEEE Globecom 87*, pages 99–102, 1987.
15. E. D. Karnin, J. W. Greene, and M. E. Hellman. On Secret Sharing Systems. *IEEE IT*, 29(1):35–41, 1983.
16. S. C. Kothari. Generalized Linear Threshold Scheme. In G. R. Blakley and D. Chaum, editors, *CRYPTO '84*, volume 196 of LNCS, pages 231–241. Springer-Verlag, 1985.
17. U. M. Maurer. Secret Key Agreement by Public Discussion from Common Information. *IEEE IT*, 39(3):733–742, May 1993.
18. R. J. McEliece and D. V. Sarwate. On Sharing Secrets and Reed-Solomon Codes. *Communications of the ACM*, 24:583–584, September 1981.
19. A. Shamir. How to Share a Secret. *Communications of the ACM*, 22:612–613, 1979.
20. G. J. Simmons. An Introduction to Shared Secret and/or Shared Control and their Application. In G. J. Simmons, editor, *Contemporary Cryptology*, pages 441–497. IEEE Press, 1991.
21. G. J. Simmons, W. Jackson, and K. M. Martin. The Geometry of Shared Secret Schemes. *Bulletin of the ICA*, 1:71–88, 1991.