Secret Sharing with Public Reconstruction

Amos Beimel and Benny Chor

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 $2n-2$ one times pads to each party. This results in shares whose size is $2n-1$ 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.

Keywords— **Secret Sharing, Public and Private Channels, Unrestricted and One-time Schemes, Space Efficiency, Cryptography.**

I. Introduction

SECRET sharing schemes were introduced by Blakley [1] and Shamir [2], and were the subject of a considerable amount of work, e.g., [3]–[7]. 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 of the parties (a subset containing at least t parties) collects their shares, and uses 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 require that after a reconstruction, only the parties who took part will know the secret. 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 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 an $O(n)$ 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 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 it. 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 V 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 securely. 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 the 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. Next, 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 jointly have enough information to reconstruct the secret. However, they cannot reconstruct the secret over the public channel, because communicating it from members of the authorized set is not possible in a secure way. This means that the authorized sets that can securely reconstruct the secret do not

Research supported by Technion VPR funds. A preliminary version of this paper was presented at CRYPTO '95, Santa Barbara, CA, and appears in *Advances in Cryptography – CRYPTO '95*, D. Coppersmith, Ed., Springer, LNCS 963, pp. 353–366.

Amos Beimel is with the Division of Engineering and Applied Sciences, Harvard University, 40 Oxford st., Cambridge, MA 02138. Email: beimel@deas.harvard.edu. This work was done while the author was a Ph.D. student at the Department of Computer Science of the Technion.

Benny Chor is with the Department of Computer Science, Technion, Haifa 32000, Israel. Email: benny@cs.technion.ac.il.

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 the information of $t-1$ parties) is at least the entropy of the secrets [4]. If the parties start with shares of smaller conditional entropy, then the parties and the dealer cannot increase it by communicating over public channels, even if interaction is allowed [8], [9]. Thus in our model, it is necessary to have secure initial distribution of shares from the dealer to the participants. However, from practical point of view the distribution stage is an off-line process which is typically done upon system initialization (unlike the reconstruction stage). Thus, assuming private initial distribution is reasonable.

Some bibliographical remarks: A similar setting of public interaction was considered for interactive key distribution schemes (e.g., [10]–[12]). Our schemes employ key distribution schemes, though not interactive ones. Another solution for 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 [13], [14]. 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., [7], [15], [16]). A different scenario in which a public broadcast channel is used (in addition to private channels) is to protect against Byzantine parties [17], [18]. Unlike our scenario, in those works the broadcast channel is heard by *all* parties.

The rest of this paper is organized as follows: In Section II we define the model, secret sharing schemes, and key distribution schemes. In Section III we describe the unrestricted schemes, and in Section IV the one time schemes. In Section V we introduce non-reactive, unrestricted schemes. In Section VI we provide the lower bound for unrestricted schemes. Finally, in Section VII we summarize our results and give two numerical examples of the sizes of shares in our various schemes.

II. Definitions

In this section we define our model, secret sharing schemes (traditional and public channels), and key distribution schemes.

A. The Model

We consider a system with n parties, denoted by $\{P_1, P_2, \ldots, P_n\}$. In addition to the parties, there is a dealer in the system, who has a secret input s taken from some finite domain S. A *distribution scheme* is a probabilistic mapping (namely a mapping which depends on the secret and on an independent random input), which the dealer applies to the input, and generates n pieces of information. These pieces of information are called shares, and the i-th pieces 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, subsets of the parties can communicate, according to some pre-defined, possibly randomized, protocol. The parties are *honest*, that is, they follow their protocols. However, a subset, disjoint of the communicating subsets, is *curious* and after the protocol has ended the members of the subset collude and try to gain some partial information about the secret.

Definition 1: Let B be a disjoint coalition (set of parties), and $s \in S$ the input of the dealer. The *view* of B, after a distribution of the input and an execution of a protocol, denoted by $\mathrm{VIEW}_B(s)$, is all the information the coalition gets. This information includes the shares of the parties in the coalition and the messages exchanged by parties over the communication channels to which the coalition has access. (This definition can accommodate different levels of security of the communication medium.) The view of B is a probabilistic function of the input (since the ensemble of shares of the parties is a function of the input).

The coalition B *has no information* on the inputs if for every two values $s_1, s_2 \in S$ and every value v of the view:

$$\Pr[\ \mathrm{VIEW}_B(s_1) = v\ ] \ = \ \Pr[\ \mathrm{VIEW}_B(s_2) = v\ ]\ .$$

The probability is taken over the random input of the dealer, and the random inputs of the members outside the coalition.

Notice that we do not make any assumptions on the probability distribution of the inputs S. This is desirable since the designer of a scheme does not necessarily know the probability distribution of the inputs that will be used (e.g., there might be only two possible inputs). This idea is inspired by the definition of probabilistic encryption [14]. An alternative definition assumes that there is some probability distribution on the inputs, and requires that for every probability distribution the probability of the input, given the view of the coalition, is the a-priori probability of the input. It can be proved that these two definitions are equivalent [19]. The alternative definition can be rephrased to state that the conditional entropy of the inputs given the view of the coalition is equal to the original entropy of the inputs. This is a common definition of secret sharing schemes [4], [20]. We use this latter definition in Section VI. Another comment is that in general the view of the parties in a coalition includes their local random inputs. However, since we consider only disjoint coalitions which only listen and do not participate in any computation, we can ignore their local random inputs.

B. Secret Sharing Schemes

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 distribution scheme, in which the dealer's 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 themselves. Every 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 variants 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 only its shares.
2. *Unrestricted* secret sharing scheme *with public reconstruction* in which a coalition B can overhear all communications taking place. The security is guaranteed even if several sets (maybe even all) 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* secret sharing scheme in which the security is guaranteed only if one set 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. In this case the view of a disjoint coalition is its shares and the communication of one reconstructing set.

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

Shamir [2] presented a traditional secret sharing 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 size 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 [4]. 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, secure traditional 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, one-time secret sharing schemes with public reconstruction do not necessarily have this monotone property. To satisfy monotonicity, it is required 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 to other parties.

C. Key Distribution Schemes

We now define unrestricted, non-communicating key distribution schemes. (Other variations of key distribution schemes can be found in [10], [11].) These key distribution schemes are used in the constructions of our secret sharing schemes.

Definition 3: Let b be a positive integer such that $b \leq n-2$, and let K be a set of keys. A $(2,b)$ *key distribution scheme* with n users and domain of keys K is a distribution scheme in which a dealer (who has only a random input) generates n shares such that the following two requirements hold:

Reconstructability: The shares of every pair of parties G determine a key, which is distributed uniformly over the domain of keys, K. (The key depends on the random input of the dealer.) Each member of the pair G can deterministically reconstruct G's key from his share without any communication.

Security: Let B be any ("bad") coalition of cardinality at most b, and G be a disjoint pair of parties. The coalition B has no information on the key of G, as defined in Definition 1. In this case their view is the collection of their shares.

The security requirement is with respect to the key of a single pair. It does not guarantee that a coalition cannot gain information on relations between different keys (e.g., the coalition might know that the keys of two pairs are equal). However, the security does guarantee some independence between keys. Consider a $(2, 2t-3)$ key distribution scheme, a coalition B of $t-1$ parties, and a disjoint set G of t parties. From the point of view of the coalition, the $\binom{t}{2}$ keys of pairs of parties in G are distributed uniformly and independently (for a proof see [11]). Blom [21] 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 $\mathrm{GF}(q)$ and the shares are taken from $\mathrm{GF}(q)^{b+1}$. Blundo et. al. [10] proved that this is optimal, namely the shares cannot be taken from smaller domains.

III. 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 $2n-1$ times the size of the secret. Our $O(n/t)$ construction can be viewed as an optimization of this simple scheme. In this scheme, the dealer shares the secret using Shamir's secret sharing scheme [2]. 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 in 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 (including his own), 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 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 over the public channel (this is an addition in the appropriate finite field). 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 adds the secret and the second key of the pair $\langle P_{i_0}, P_i \rangle$ and sends this sum over 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 take place, this does 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 [11], 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 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 $\mathrm{GF}(q)^t$, where q is a prime-power such that $q > n$. (We require that $q > n$ since this is the requirement in Shamir's scheme.) In the scheme we view the secret as t sub-secrets from $\mathrm{GF}(q)$. The scheme is presented in Fig. 1.

Unrestricted Secret Sharing Scheme

Distribution stage:
Input: t sub-secrets $s_1, s_2, \ldots, s_t \in \mathrm{GF}(q)$
Shares:
Share each subsecret s_i using Shamir's scheme $(1 \le i \le t)$.
Denote the n shares of sub-secret s_i by $s_{i,1}, s_{i,2}, \ldots, s_{i,n}$.
For every two parties generate 4 independent keys from $\mathrm{GF}(q)$.
Denote the keys of $\langle P_i, P_j \rangle$ by $k^1_{i,j}, k^2_{i,j}, k^3_{i,j}, k^4_{i,j}$.
Share of party P_i**:** $s_{1,i}, \ldots, s_{t,i}$ and $k^1_{i,j}, k^2_{i,j}, k^3_{i,j}, k^4_{i,j}$ for $1 \le j \le n$.

Reconstruction stage:
A set $G = \left\{P_{i_1}, \ldots, P_{i_\ell}\right\}$ wants to reconstruct the secret
Every party in G announces whether he has previously reconstructed the secret.
Let P_{i_j} for $1 \le j \le t$ be the leaders of G.
Each leader P_{i_j} $(1 \le j \le t)$ sends (at most) $t-1$ messages to all other leaders that have not previously reconstructed the secret:

$s_{j',i_j} + k^1_{i_j,i_{j'}}$ to $P_{i_{j'}}$ for $1 \le j' < j$
$s_{j',i_j} + k^2_{i_j,i_{j'}}$ to $P_{i_{j'}}$ for $j < j' \le t$

Each leader P_{i_j} computes s_j from $s_{j,i_1}, \ldots, s_{j,i_t}$.
Each leader P_{i_j} sends a message to every $P_{i_{j'}} \in G$ that has not previously reconstructed the secret:

$s_j + k^3_{i_j,i_{j'}}$ to $P_{i_{j'}}$ for $1 \le j' < j$
$s_j + k^4_{i_j,i_{j'}}$ to $P_{i_{j'}}$ for $j < j' \le \ell$

Each party of G concatenates the sub-secrets $s_1, s_2, \ldots, s_t$ to obtain the secret s.

Fig. 1. Unrestricted t-out-of-n secret sharing scheme.

As described, the scheme has two technical points which should be clarified. The first is the fact that in one reconstruction two parties P_i and P_j might need to exchange 4 different messages: Assume that P_i is responsible for reconstructing the sub-secret $s_{i'}$, and P_j is responsible for reconstructing the sub-secret $s_{j'}$. The party P_i has to receive from P_j the share corresponding to the sub-secret $s_{i'}$, and then will send the sub-secret $s_{i'}$. Similarly, P_j has to receive from P_i the share corresponding to the sub-secret $s_{j'}$, and then will send the sub-secret $s_{j'}$. 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. To overcome the problem, every party that participates in one reconstruction will remember the secret, and in later reconstructions will inform other parties (in the clear) that he knows the secret. In such case, other parties will not send him any messages. He will continue to send the messages that he has to send according to the scheme (to "new" parties). 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.

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

Theorem 4: 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 $\mathrm{GF}(q)^t$. The share of each party is an element of $\mathrm{GF}(q)^{4n+t-4}$. So the size of each share is $1 + 4(n-1)/t$ times the size of the secrets.

IV. 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 [21] 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 $1+\log(n/t)$ independent instances of "exact schemes" for thresholds $t, 2t, 4t, \ldots$ 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 $t, 2t, 4t, \ldots$ – at most two 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 remains 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 "exactly t" scheme. The distribution phase is depicted in Fig. 2.

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

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

- Consider the secret as t secrets $s_1, \ldots, s_t \in \mathrm{GF}(q)$.
- 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 a $(2,b)$-key distribution scheme with key domain $\mathrm{GF}(q)^4$ (considered as 4 keys from $\mathrm{GF}(q)$).

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

Fig. 2. One-time for exactly t-out-of-n secret sharing scheme with public reconstruction

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 II-C: 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 $\mathrm{GF}(q)$. In addition, each party gets a share of Blom's $(2, 2t-3)$ key distribution scheme with keys taken from $\mathrm{GF}(q)^4$ – these shares are taken from $\mathrm{GF}(q)^{4(2t-2)}$. Overall, the total share contains $(9t-8)$ elements from $\mathrm{GF}(q)$ (if $2t > n+1$, then the shares are even shorter). Recall that the secret is taken from $\mathrm{GF}(q)^t$, and therefore the size of the share is less than 9 times the size of the secrets.

In this scheme, the domain of secrets has to be $\mathrm{GF}(q)^t$ (for some prime-power q). Restricting the domain of secrets 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 $q = 2^{\lceil m/t \rceil}$, we have $m' = t\lceil m/t \rceil \leq m+t$. If we assume that $m > 9t$ then the size of the share is $9m' \leq 9(m+t) < 10m$

Theorem 5: Let m be a natural number such that $m > \max\{t\log n, 9t\}$. There exists a one-time secret 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.

One-time schemes are a special case of traditional secret sharing schemes even if only sets of size exactly t can securely reconstruct the secret, since every set of at least t parties has enough information to reconstruct the secret on secure private channels. Thus, the size of each share has to be at least the size of the secret [4]. Therefore, our scheme is tight up to a constant factor. We can slightly improve this lower bound, by observing that every one-time exactly t-out-of-n secret sharing scheme with public reconstruction can be used as a one-time communicating $(t, t-1)$ key distribution scheme (for $t \leq n/2$) (for definition of communicating key distribution scheme see [11]). The size of the share in every $(t, t-1)$ key distribution scheme is at least twice the size of the key [11]. Therefore, the size of the share in every one-time secret sharing scheme with public reconstruction is at least twice the size of the secret.

In Fig. 3 we 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 $G = \left\{P_{i_1}, \ldots, P_{i_\ell}\right\}$ that wants to reconstruct the secret (where $\ell \geq t$).
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.

Fig. 3. One-time t-out-of-n secret sharing scheme with public reconstruction

Theorem 6: The scheme of Fig. 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 $\max\{n \log n, 9n\}$, then the size of the shares of every party is less than $10(\log(n/t) + 2)$ times the size of the secrets.

Remark 7: 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. To achieve this goal we use a slightly weaker building block (instead of the exactly t-out-of-n scheme). This building block is a scheme in which exactly $2^i t$ parties can reconstruct the secret, while any coalition of size $t-1$ does not gain any information on the reconstructed secret. Schemes which satisfy this requirement on secure, private channels were presented by Blakley and Meadows [22] (see also [3], [4] and [23]), and are called ramp schemes. We use ramp schemes instead of regular secret sharing schemes to design our building boxes. Assume that sets of ℓ parties should be able to reconstruct the secret, while sets of fewer than t parties should get no information about the secret. The size of the share in such ramp scheme is $1/(\ell - t + 1)$ times the size of the secret, i.e. substantially smaller than traditional secret sharing schemes.

For the scheme with public reconstruction, assume that the domain of secrets is $\mathrm{GF}(q)^{\ell(\ell-t+1)}$. That is, the dealer has ℓ secrets, each one taken from $\mathrm{GF}(q)^{\ell-t+1}$. The share of each party is one share of the ramp scheme for ℓ sub-secrets, each share taken from $\mathrm{GF}(q)$. In the reconstruction of the secret by a subset containing exactly ℓ parties, each party will be responsible for one sub-secret. Each pair of parties in this set first exchanges two shares of the ramp scheme. Now each party reconstructs his sub-secret, and every pair of parties exchange two sub-secrets. Therefore, every pair of parties needs two keys of a $(2, \ell + t - 3)$-key distribution scheme with keys from $\mathrm{GF}(q)$ (the domain of shares of the ramp scheme), as well as two keys from a $(2, \ell + t - 3)$-key distribution scheme with keys from $\mathrm{GF}(q)^{\ell-t+1}$ (the domain of sub-secrets). Overall, the share of each party is an element taken from $\mathrm{GF}(q)^{3\ell+2t-4+2(\ell+t-2)(\ell-t+1)}$. That is, the size of the share is $2 + 2t/\ell + O(1/(\ell - t))$ times the size of the secrets. In the t-out-of-n scheme for every set, we use these schemes with $\ell = t, 2t, 4t, \ldots$ – therefore the size of the share is only $2\log(n/t) + O(1)$ times the size of the secret. In this construction we required that the size of the secret m is greater than $n^2 \log n$, this requirement can be weakened to $m \geq n \log^3 n$.

V. 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 improvement (by a factor of 2) over the reactive scheme of Section III 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. For general access structures it 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 $r_1, \ldots, r_n$. The share of P_i is $r_1, \ldots, r_{i-1}, r_i + s, r_{i+1}, \ldots, r_n$. 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 a general secret sharing scheme, first suggested by [24], 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 jointly know anything about the secret. Secret sharing schemes satisfying these requirements can exist only for monotone collections. Indeed, it is known that for every monotone collection there exists a traditional secret-sharing scheme [24]–[26]. 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 structure. The unrestricted, non-reactive, 2-out-of-n scheme can be generalized to an unrestricted, non-reactive scheme realizing the access structure $\mathcal{A}$. Assume there is a traditional secret sharing scheme realizing $\mathcal{A}$ with domain of secrets S and domain of shares U. In our scheme we use the following observation: Denote by $\mathcal{A}_i$ the access structure $\mathcal{A}_i = \{B : B \cup \{P_i\} \in \mathcal{A}\}$. There exists a traditional secret sharing scheme realizing $\mathcal{A}_i$ in which the domain of shares is U (fix some possible share u for P_i, and share the secret using the scheme for $\mathcal{A}$ conditioned on the fact that the share of P_i is u).

We now describe an unrestricted, non-reactive scheme realizing $\mathcal{A}$ with domain of shares $S \times U^{n-1}$. To share a secret s, the dealer chooses n random independent elements from S, the domain of secrets, denoted $r_1, \ldots, r_n$. For every i, the dealer distributes the share $r_i + s$ to P_i, and shares r_i among $\{P_1, \ldots, P_{i-1}, P_{i+1}, \ldots, P_n\}$ using the scheme realizing $\mathcal{A}_i$ with domain of shares U. That is, the share of P_i is $r_i + s$ together with the shares of the $n-1$ schemes realizing $\mathcal{A}_1, \ldots, \mathcal{A}_{i-1}, \mathcal{A}_{i+1}, \ldots, \mathcal{A}_n$ with secrets

$r_1, \ldots, r_{i-1}, r_{i+1}, \ldots, r_n$ respectively. The total share is an element taken from $S \times U^{n-1}$. Now, when a subset B wishes to reconstruct the secret, every $P_i \in B$ sends (in the open) the share of the secret r_j to every $P_j \in B$. Thus, P_i holds $r_i + s$ and hears the shares of $B \setminus \{P_i\}$ from the scheme realizing $\mathcal{A}_i$ with the secret r_i. Since $B \setminus \{P_i\} \in \mathcal{A}_i$, the party P_i can reconstruct r_i, and with $r_i + s$ reconstructs the secret.

We next claim that the reconstructions are secure. That is, every coalition $C \notin \mathcal{A}$ hearing communications during the reconstruction of the secret by all sets in $\mathcal{A}$ that are disjoint to C does not gain information on the secret. The coalition C is not in $\mathcal{A}_i$ for any $P_i \in C$, thus, the shares of C give no information on r_i for $P_i \in C$. Hence, prior to any reconstruction the coalition C has no information on the secret even if it knows r_i for every $P_i \notin C$. But the information that the coalition gains from the communication is at most the r_i's for $P_i \notin C$, and it does not gain any information on the secret. That is, the reconstruction is secure. Thus,

Theorem 8: 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.

We can apply the previous construction for threshold schemes using Shamir's scheme.

Corollary 9: 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 $\mathrm{GF}(q)$. The share of each party is an element of $\mathrm{GF}(q)^n$. So the size of each share is n times the size of the secret.

VI. Lower Bounds for Unrestricted Schemes

In this section we prove an $\Omega(n/t)$ lower bound on the increase in the shares' size for unrestricted t-out-of-n schemes. The specific lower bound that we prove is tight for $t = 2$ (by the non-reactive scheme of Section V). For $t > 2$ our lower bound is tight up to a constant factor (by the reactive scheme of Section III). We first prove an $\Omega(n)$ lower bound on the increase in size of shares for 2-out-of-n schemes. Then, we show that this lower bounds translates into an $\Omega(n/t)$ 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 [27]. We assume an arbitrary probability distribution on the secrets, and we denote the secret by the random variable S.

The intuition behind the proof is that P_i has to expose $H(S)$ *"new"* bits of his share in *each* reconstruction, and P_i can participate in $n-1$ reconstructions. After all $n-1$ reconstructions, the uncertainty of the share of P_i has to remain at least $H(S)$, as an outsider who listened to all reconstructions still has $H(S)$ uncertainty on the secret. Thus, 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_j the messages exchanged when P_1 and P_j reconstruct the secret (all these are random variables). We denote $C = C_3 \ldots C_n$, the concatenation of all messages exchanged between P_1 and the parties $P_3, \ldots, P_n$. Recall that the communication C_2, together with P_2's share S_2, enables 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 10: $H(C_2|S_2C) \geq H(S)$.

Proof: Since P_2 can reconstruct the secret S, given his share S_2 and the messages C_2 exchanged between P_1 and P_2, the conditional entropy $H(S|C_2S_2)$ 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 $I(C_2;S|S_2C)$ of the message C_2 and the secret S, given the share S_2 and C. We have

$$\begin{aligned} & H(C_2|S_2C) - H(C_2|SS_2C) \\ = \;& I(C_2;S|S_2C) \\ = \;& H(S|S_2C) - H(S|C_2S_2C) \\ = \;& H(S) \; . \end{aligned}$$

Since the entropy is non-negative, $H(C_2|S_2C) \geq H(S)$. ∎

The next claim is the heart of the proof of the lower bound. It states that the mutual information between S_1 and C_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_2. Hence, C_2 must contain $H(S)$ bits of information originating from the share S_1. Claim 11 is stated for deterministic parties – the outgoing messages are determined by the given share and previous incoming messages. An analogous statement is proved in Section VI-A for randomized parties, whose outgoing messages could in addition depend on random local inputs.

Claim 11: For deterministic reconstruction protocols we have

$$I(C_2;S_1|C) = H(S_1|C) - H(S_1|C_2C) \geq H(S) \; .$$

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 can 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, $C_2 = A_1B_1 \ldots A_kB_k$. The message A_i is determined by the share S_1 and previous messages, that is, $H(A_i|S_1A_1B_1\ldots A_{i-1}B_{i-1})=0$. The following inequality holds for any deterministic communication protocol:

$$\begin{aligned} H(C_2|S_1C) &= H(A_1B_1\ldots A_kB_k|S_1C) \\ &= \sum_{i=1}^{k} H(A_i|S_1CA_1B_1\ldots A_{i-1}B_{i-1}) \\ &\quad +\sum_{i=1}^{k} H(B_i|S_1CA_1B_1\ldots A_{i-1}B_{i-1}A_i) \\ &= \sum_{i=1}^{k} H(B_i|S_1CA_1B_1\ldots A_{i-1}B_{i-1}A_i) \\ &\le \sum_{i=1}^{k} H(B_i|CA_1B_1\ldots A_{i-1}B_{i-1}A_i)\ . \end{aligned}$$

Similarly,

$$H(C_2|S_2C) \le \sum_{i=1}^{k} H(A_i|CA_1B_1\ldots A_{i-1}B_{i-1})\ .$$

Combing the two inequalities

$$\begin{aligned} & H(C_2|S_1C)+H(C_2|S_2C) \\ &\le \sum_{i=1}^{k} H(B_i|CA_1B_1\ldots A_{i-1}B_{i-1}A_i) \\ &\quad +\sum_{i=1}^{k} H(A_i|CA_1B_1\ldots A_{i-1}B_{i-1}) \\ &= H(A_1B_1\ldots A_kB_k|C) \\ &= H(C_2|C)\ . \end{aligned}$$

This inequality, together with Claim 10, implies

$$\begin{aligned} I(C_2;S_1|C) &= H(C_2|C)-H(C_2|S_1C) \\ &\ge H(C_2|S_2C) \\ &\ge H(S) \end{aligned}$$

■

We are now ready to prove our lower bound for $t=2$.

Claim 12: In any unrestricted 2-out-of-n secret sharing scheme with public reconstruction, the share of each participant, S_i, satisfies

$$H(S_i)\ge n\cdot H(S)\ .$$

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

$$H(S|C_2C_3\ldots C_n)=H(S)\ .$$

On the other hand, given P_1's share, this communication determines the secret, so

$$H(S|S_1C_2C_3\ldots C_n)=0\ .$$

Therefore,

$$\begin{aligned} H(S) &= H(S|C_2C_3\ldots C_n)-H(S|S_1C_2C_3\ldots C_n) \\ &= I(S;S_1|C_2C_3\ldots C_n) \\ &= H(S_1|C_2C_3\ldots C_n)-H(S_1|SC_2C_3\ldots C_n)\ , \end{aligned}$$

and in particular

$$H(S_1|C_2C_3\ldots C_n)\ge H(S)\ .$$

Claim 11 (or Claim 17 for the case of randomized protocols) states that

$$H(S_1|C_3\ldots C_n)-H(S_1|C_2C_3\ldots C_n)\ge H(S)\ .$$

Similarly it holds that

$$\begin{aligned} H(S_1|C_4\ldots C_n)-H(S_1|C_3C_4\ldots C_n) &\ge H(S)\ . \\ &\vdots \\ H(S_1)-H(S_1|C_n) &\ge H(S)\ . \end{aligned}$$

Summing these n inequalities, we conclude that

$$H(S_1)\ge n\cdot H(S)\ .$$

■

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 13: 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 12 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 $H(S_1)\le\log|S_1|$, 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. ■

A. Lower Bound for Probabilistic Parties

In the proof of Claim 11 we assumed that the parties are deterministic during the reconstruction of the secret. In this section we prove the same claim without this assumption. Recall that S_1 is the share of P_1, C_2 is the communication generated in the reconstruction of the secret by

P_1 and P_2, and C is the communication in previous reconstructions. We prove that the mutual information between S_1 and C_2 given C is at least the entropy of the secret, even if the parties may toss coins during the reconstruction. That is, party P_1 has an independent local random string, denoted R_1, and the messages he generates are a deterministic function of his share, his random input, and previous messages. As the claim concerns the share of P_1, we can assume that other parties in the system are deterministic (the dealer can supply a random string to the other parties as part of their shares). Since R_1 is independent of the shares and the secret, the mutual information between R_1 and the shares and the secret is zero, i.e. $I(R_1; SS_1S_2 \ldots S_n) = 0$. We first quote a general claim from [8] and [9, Lemma 2.2] which states that the mutual information between the inputs of the parties can only decrease by communicating over public channels.

Claim 14: **[9], [8]** Let X, Y be random variables, held by parties P_1, P_2 respectively. The parties P_1 and P_2 communicate on a channel according to some protocol. Denote the communication by the random variable M. Then, $I(\ X\ ;\ Y\ ) \geq I(\ X\ ;\ Y\ \mid M)$.

We next prove that the mutual information between R_1 and the shares of other parties, given S_1 and a communication M is zero (where M is any prefix of CC_2). The claim can be proven directly by induction on the number of messages sent in M (this is the way that Claim 14 is proven). We avoid this induction as we show that our claim can be formulated as a special case of Claim 14.

Claim 15: Let M be a prefix of the communication exchanged between the parties in the system. Then, $H(R_1 \mid SS_2S_1M) = H(R_1 \mid S_1M)$.

Proof: Consider a scenario in which one party holds R_1, and a second party holds the secret S and all the shares – $S_1, S_2, \ldots, S_n$. They communicate via a public channel and the first message is sent by the second party and equals S_1. From now on, the first player can generate the messages of P_1 and the second player can generate all other messages. Thus, the two parties can continue to communicate and generate M. By Claim 14, communicating S_1 and M can only decrease the mutual information, that is,

$$I(R_1; SS_1S_2 \ldots S_n | S_1M) \leq I(R_1; SS_1S_2 \ldots S_n) = 0\ .$$

Recall that,

$$I(R_1; SS_1S_2 | S_1M) \leq I(R_1; SS_1S_2 \ldots S_n | S_1M).$$

Since the mutual information is non-negative,

$$H(R_1|S_1M) - H(R_1|SS_1S_2M) = I(R_1; SS_1S_2|S_1M) = 0\ .$$

■

We restrict our discussion to protocols with an absolute bound k on the number of rounds in each communication C_j. (The case where the protocol terminates after a finite number of rounds with probability 1 can be handled similarly.) Denote by A_i the i-th message sent by P_1 to P_2, and similarly, B_i to be the i-th message sent by P_2 to P_1. I.e, $C_2 = A_1B_1 \ldots A_kB_k$. We next prove that the dependence of A_i on S_1 and the previous messages is greater than the dependence of A_i on S_2, the secret S, and the previous messages. Formally,

Claim 16: $H(A_i|SS_2CA_1B_1 \ldots A_{i-1}B_{i-1}) \geq H(A_i|S_1CA_1B_1 \ldots A_{i-1}B_{i-1})$.

Proof: On one hand, since the entropy is non-negative

$$\begin{aligned} & I(A_i; R_1|SS_1S_2CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad = H(A_i|SS_1S_2CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad\quad - H(A_i|R_1SS_1S_2CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad \leq H(A_i|SS_1S_2CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad \leq H(A_i|SS_2CA_1B_1 \ldots A_{i-1}B_{i-1})\ . \end{aligned}$$

On the other hand, using Claim 15 we can write

$$\begin{aligned} & I(A_i; R_1|SS_1S_2CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad = H(R_1|SS_1S_2CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad\quad - H(R_1|A_iSS_1S_2CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad = H(R_1|S_1CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad\quad - H(R_1|A_iS_1CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad = H(A_i|S_1CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad\quad - H(A_i|R_1S_1CA_1B_1 \ldots A_{i-1}B_{i-1})\ . \end{aligned}$$

Thus,

$$\begin{aligned} & H(A_i|SS_2CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad \geq H(A_i|S_1CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad\quad - H(A_i|R_1S_1CA_1B_1 \ldots A_{i-1}B_{i-1}) \end{aligned} \tag{1}$$

Since the message A_i is completely determined by the random input R_1, the share S_1 and the previous messages $C, A_1, B_1, \ldots, A_{i-1}, B_{i-1}$, it holds that

$$H(A_i|R_1S_1CA_1B_1 \ldots A_{i-1}B_{i-1}) = 0.$$

Combining this equality and Inequality (1), we get

$$\begin{aligned} & H(A_i|SS_2CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad \geq H(A_i|S_1CA_1B_1 \ldots A_{i-1}B_{i-1})\ . \end{aligned}$$

■

We now prove the analogue of Claim 11, without the restriction to deterministic reconstruction.

Claim 17: $I(C_2; S_1|C) \geq H(S)$.

Proof: Expressing $I(C_2; S_1|C)$ as a sum and using the fact that mutual information is non-negative, we get

$$\begin{aligned} I(C_2; S_1|C) &= \sum_{i=1}^{k} I(A_i; S_1|CA_1B_1 \ldots A_{i-1}B_{i-1}) \\ & \quad + \sum_{i=1}^{k} I(B_i; S_1|CA_1B_1 \ldots A_{i-1}B_{i-1}A_i) \\ & \geq \sum_{i=1}^{k} I(A_i; S_1|CA_1B_1 \ldots A_{i-1}B_{i-1}) \end{aligned}$$

$$
\begin{aligned}
&= \sum_{i=1}^{k} H(A_i|CA_1B_1\dots A_{i-1}B_{i-1}) \\
&\quad - \sum_{i=1}^{k} H(A_i|S_1CA_1B_1\dots A_{i-1}B_{i-1}) \\
&\geq \sum_{i=1}^{k} H(A_i|S_2CA_1B_1\dots A_{i-1}B_{i-1}) \\
&\quad - \sum_{i=1}^{k} H(A_i|S_1CA_1B_1\dots A_{i-1}B_{i-1})\ .
\end{aligned}
$$

This inequality together with Claim 16 imply

$$
\begin{aligned}
I(C_2;S_1|C) &\geq \sum_{i=1}^{k} H(A_i|S_2CA_1B_1\dots A_{i-1}B_{i-1}) \\
&\quad - \sum_{i=1}^{k} H(A_i|SS_2CA_1B_1\dots A_{i-1}B_{i-1}) \\
&= \sum_{i=1}^{k} I(A_i;S|S_2CA_1B_1\dots A_{i-1}B_{i-1}) \\
&= \sum_{i=1}^{k} H(S|S_2CA_1B_1\dots A_{i-1}B_{i-1}) \\
&\quad - \sum_{i=1}^{k} H(S|S_2CA_1B_1\dots A_{i-1}B_{i-1}A_i)\ .
\end{aligned}
$$

Since the message B_i is determined by S_2 and previous messages,

$$
\begin{aligned}
&H(S|S_2CA_1B_1\dots A_{i-1}B_{i-1}A_i) \\
&\quad = H(S|S_2CA_1B_1\dots A_{i-1}B_{i-1}A_iB_i)\ .
\end{aligned}
$$

Therefore the "internal" summands in the last sum cancel each other, and we are left with

$$I(C_2;S_1|C) \geq H(S|S_2C) - H(S|S_2CC_2)\ . \tag{2}$$

Recall that the communication C and the share S_2 give no information (to P_2) about the secret, i.e. $H(S|S_2C) = H(S)$. On the other hand, P_2 holding S_2 and knowing C_2, can reconstruct the secret, i.e. $H(S|S_2CC_2) = 0$. Therefore, $H(S|S_2C) - H(S|S_2CC_2) = H(S)$. Together with Inequality (2), we get $I(C_2;S_1|C) \geq H(S)$, as claimed. ■

Claim 17 implies that Theorem 13 holds also in the scenario in which the parties can toss coins during the reconstructions.

VII. Conclusions

In this work we investigated the cost of performing the reconstruction over *public* communication channels. In Fig. 4 we summarize our results for the various schemes. We denote by m the size of the secret, and the sizes of the shares are multiples of m (e.g., $10m$). We also give the minimum size of secrets for which this share size applies. In Fig. 5 we give two examples of the sizes of shares in the various schemes. In both examples we consider a system with 1024 parties.

Acknowledgments

We thank Carlo Blundo, Ehud Hausman, and Hugo Krawczyk for helpful discussions, and Arie Freund for comments on earlier versions of this paper.

	share size	min. secret size
Naive	$(2n-1)m$	$\log n$
Non-reactive	nm	$\log n$
Unrestricted	$(1+4(n-1)/t)m$	$t \log n$
One-time	$10(\log(n/t)+2)m$	$\max\{n\log n, 9n\}$
One-time, exactly t	$10m$	$\max\{t\log n, 9t\}$

Fig. 4. Summary of the complexity of our algorithms for t-out-of-n schemes.

	share size	minimum secret size
$n = 1024$; $t = 128$		
Naive scheme	$2047m$	10 bits
Unrestricted scheme	$33m$	160 bytes
One-time at least t scheme	$40m$	1280 bytes
One-time exactly t scheme	$10m$	160 bytes
$n = 1024$; $t = 8$		
Naive scheme	$2047m$	10 bits
Unrestricted scheme	$513m$	10 bytes
One-time at least t scheme	$80m$	1280 bytes
One-time exactly t scheme	$10m$	90 bytes

Fig. 5. Numerical Examples.

References

[1] G. R. Blakley, "Safeguarding cryptographic keys," in *Proc. AFIPS 1979 NCC, vol. 48*, June 1979, pp. 313–317.

[2] A. Shamir, "How to share a secret," *Communications of the ACM*, vol. 22, pp. 612–613, 1979.

[3] R. J. McEliece and D. V. Sarwate, "On sharing secrets and Reed-Solomon codes," *Communications of the ACM*, vol. 24, pp. 583–584, September 1981.

[4] E. D. Karnin, J. W. Greene, and M. E. Hellman, "On secret sharing systems," *IEEE Trans. on Information Theory*, vol. 29, no. 1, pp. 35–41, 1983.

[5] S. C. Kothari, "Generalized linear threshold scheme," in *Advances in Cryptology - CRYPTO '84*, G. R. Blakley and D. Chaum, Eds. 1985, vol. 196 of *Lecture Notes in Computer Science*, pp. 231–241, Springer-Verlag.

[6] J. Benaloh, "Secret sharing homomorphisms: Keeping shares of a secret secret," in *Advances in Cryptology - CRYPTO '86*, A. M. Odlyzko, Ed. 1987, vol. 263 of *Lecture Notes in Computer Science*, pp. 251–260, Springer-Verlag.

[7] G. J. Simmons, "An introduction to shared secret and/or shared control and their application," in *Contemporary Cryptology, The Science of Information Integrity*, G. J. Simmons, Ed., pp. 441–497. IEEE Press, 1992.

[8] U. M. Maurer, "Secret Key Agreement by Public Discussion from Common Information," *IEEE Trans. on Information Theory*, vol. 39, no. 3, pp. 733–742, 1993.

[9] R. Ahlswede and I. Csiszár, "Common randomness in information theory and cryptography – part I: Secret sharing," *IEEE Trans. on Information Theory*, vol. 39, no. 4, pp. 1121–1132, 1993.

[10] 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*, E. F. Brickell, Ed. 1993, vol. 740 of *Lecture Notes in Computer Science*, pp. 471–486, Springer-Verlag.

[11] A. Beimel and B. Chor, "Communication in key distribution schemes," *IEEE Trans. on Information Theory*, vol. 42, no. 1, pp. 19–28, 1996.

[12] C. Blundo and A. Cresti, "Space requirement for broadcast encryption," in *Advances in Cryptology - EuroCRYPT '94*, A. De Santis, Ed. 1995, vol. 950 of *Lecture Notes in Computer Science*, pp. 287–298, Springer-Verlag.

[13] W. Diffie and M. E. Hellman, "New directions in cryptography," *IEEE Trans. on Information Theory*, vol. 22, no. 6, pp. 644–654, 1976.

[14] S. Goldwasser and S. Micali, "Probabilistic encryption," *J. of Computer and System Sciences*, vol. 28, no. 21, pp. 270–299, 1984.

[15] B. Blakley, G. R. Blakley, A. H. Chan, and J. Massey, "Threshold schemes with disenrollment," in *Advances in Cryptology - CRYPTO '92*, E. F. Brickell, Ed. 1993, vol. 740 of *Lecture Notes in Computer Science*, pp. 540–548, Springer-Verlag.

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

[17] T. Rabin and M. Ben-Or, "Verifiable secret sharing and multiparty protocols with honest majority," in *Proc. of the 21st Annu. ACM Symp. on the Theory of Computing*, 1989, pp. 73–85.

[18] T. Rabin, "Robust sharing of secrets when the dealer is honest or faulty," *J. of the ACM*, vol. 41, no. 6, pp. 1089–1109, 1994.

[19] E. F. Brickell and D. R. Stinson, "Some improved bounds on the information rate of perfect secret sharing schemes," *J. of Cryptology*, vol. 5, no. 3, pp. 153–166, 1992.

[20] R. M. Capocelli, A. De Santis, L. Gargano, and U. Vaccaro, "On the size of shares for secret sharing schemes," *J. of Cryptology*, vol. 6, no. 3, pp. 157–168, 1993.

[21] R. Blom, "An optimal class of symmetric key generation systems," in *Advances in Cryptology – EuroCRYPT '84*, T. Beth, N. Cot, and I. Ingemarsson, Eds. 1985, vol. 209 of *Lecture Notes in Computer Science*, pp. 335–338, Springer-Verlag.

[22] G. R. Blakley and C. Meadows, "The security of ramp schemes," in *Advances in Cryptology - CRYPTO '84*, G. R. Blakley and D. Chaum, Eds. 1985, vol. 196 of *Lecture Notes in Computer Science*, pp. 242–268, Springer-Verlag.

[23] M. K. Franklin and M. Yung, "Communication complexity of secure computation," in *Proc. of the 24th Annu. ACM Symp. on the Theory of Computing*, 1992, pp. 699–710.

[24] M. Ito, A. Saito, and T. Nishizeki, "Secret sharing schemes realizing general access structure," in *Proc. IEEE Global Telecommunication Conf., Globecom 87*, 1987, pp. 99–102, Journal version: "Multiple assignment scheme for sharing secret". *J. of Cryptology*, vol. 6, no. 1, pp. 15–20, 1993.

[25] J. Benaloh and J. Leichter, "Generalized secret sharing and monotone functions," in *Advances in Cryptology - CRYPTO '88*, S. Goldwasser, Ed. 1990, vol. 403 of *Lecture Notes in Computer Science*, pp. 27–35, Springer-Verlag.

[26] G. J. Simmons, W. Jackson, and K. M. Martin, "The geometry of shared secret schemes," *Bulletin of the ICA*, vol. 1, pp. 71–88, 1991.

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