



Secret Sharing Lower Bound: Either Reconstruction is Hard or Shares are Long

Kasper Green Larsen and Mark Simkin^(✉)

Aarhus University, Aarhus, Denmark
{larsen, simkin}@cs.au.dk

Abstract. A secret sharing scheme allows a dealer to distribute shares of a secret among a set of n parties $P = \{p_1, \dots, p_n\}$ such that any authorized subset of parties can reconstruct the secret, yet any unauthorized subset learns nothing about it. The family $\mathcal{A} \subseteq 2^P$ of all authorized subsets is called the access structure. Classic results show that if \mathcal{A} contains precisely all subsets of cardinality at least t , then there exists a secret sharing scheme where the length of the shares is proportional to $\lg n$ bits plus the length of the secret. However, for general access structures, the best known upper bounds have shares of length exponential in n , whereas the strongest lower bound shows that the shares must have length at least $n/\lg n$. Beimel conjectured that the exponential upper bound is tight, but proving it has so far resisted all attempts. In this paper we make progress towards proving the conjecture by showing that there exists an access structure \mathcal{A} , such that any secret sharing scheme for \mathcal{A} must have either exponential share length, or the function used for reconstructing the secret by authorized parties must have an exponentially long description. As an example corollary, we conclude that if one insists that authorized parties can reconstruct the secret via a constant fan-in boolean circuit of size polynomial in the share length, then there exists an access structure that requires a share length that is exponential in n .

1 Introduction

A secret sharing scheme allows a dealer to distribute shares of a secret among a set of parties $P = \{p_1, \dots, p_n\}$ such that any authorized subset $A \subseteq P$ can reconstruct the secret, yet any unauthorized subset learns nothing about it. The family $\mathcal{A} \subseteq 2^{\{p_1, \dots, p_n\}}$ of all authorized subsets is called the access structure. Secret

K. G. Larsen—Supported by a Villum Young Investigator grant and an AUFF starting grant.

M. Simkin—Supported by the European Research Council (ERC) under the European Unions’s Horizon 2020 research and innovation programme under grant agreement No 669255 (MPCPRO), grant agreement No 803096 (SPEC), Danish Independent Research Council under Grant-ID DFF-6108-00169 (FoCC), and the Concordium Blockchain Research Center.

© Springer Nature Switzerland AG 2020

C. Galdi and V. Kolesnikov (Eds.): SCN 2020, LNCS 12238, pp. 566–578, 2020.

https://doi.org/10.1007/978-3-030-57990-6_28

sharing was introduced independently by Shamir [35] and Blakley [11], who presented constructions for threshold access structures that contains all subsets with a cardinality larger than some threshold t . The first construction for general (monotone) access structures was presented by Ito, Saito, and Nishizeki [21].

The main measure of efficiency for secret sharing schemes is the share size. For threshold access structures it is known that Shamir's secret sharing, which has a share size of $\Theta(\lg n)$, is optimal up to additive constants [13]. This stands in stark contrast to the smallest share sizes we can achieve for general monotone access structures. The construction of Ito, Saito, and Nishizeki has a share size of $\mathcal{O}(2^n/\sqrt{n})$ and 29 years later the best known upper bound on the share size, due to Applebaum et al. [3] is still $2^{0.892n}$. A widely believed conjecture suggests that these upper bounds are, up to constants, the best ones one can hope for. More concretely, Beimel conjectured:

Conjecture 1 ([5, 6]). There exists an $\epsilon > 0$ such that for every integer n there exists an access structure with n parties for which every secret sharing scheme distributes shares of length exponential in the number of parties, that is, $2^{\epsilon n}$.

Proving this conjecture is a major open problem in the research area of secret sharing schemes. Karnin, Greene, and Hellman [23] initiated a line of works [12, 14, 17, 18] that proved different lower bounds on the share size using tools from information theory. The best of those lower bounds is due to Csirmaz [17, 18], who uses Shannon information inequalities to prove that there exists an explicit access structure that requires shares of size $\Omega(n/\lg n)$. Csirmaz himself and subsequent works [9, 30] indicate that it is unlikely that one can prove a super-polynomial lower bound on the share size using such information inequalities.

A different line of works focuses on linear secret sharing schemes, where the shared secret is a linear combination of the shares. Many of the existing schemes, e.g. [35], are linear and applications like multiparty computation [10, 16, 34] crucially rely on this property. Karchmer and Wigderson [22] introduce monotone span programs and show that these are closely related to linear secret sharing schemes. Through the lens of monotone span programs, a series of works obtained increasingly stronger lower bounds. Karchmer and Wigderson prove the first super-linear lower bound on the share size. Babai, Gál, and Wigderson [4] prove the first super-polynomial lower bound. Finally, Pitassi and Robere [33] prove an exponential lower bound, however, the gap between the constants in the exponent of the lower and upper bound remain far apart.

Several works consider different flavors of the original secret sharing notion. Beimel and Franklin [7] consider a relaxed security notion of weak privacy, which only requires that any unauthorized subset can not exclude any secret value with certainty. The unauthorized subset can, however, conclude that some secret is more probable than another one. The authors show that this notion is strictly weaker than the original notion of secret sharing by constructing schemes with share sizes that are impossible for secret sharing schemes with perfect privacy. Among other results, the authors construct a weakly-private secret sharing scheme for the threshold access structures, where the share size is independent

of n . The authors conclude that any sensible lower bound proof has to make use of the privacy requirement of secret sharing schemes. Applebaum et al. [1, 2] consider the efficiency of secret sharing schemes for large secrets. The authors show that, for a certain class of access structures, one can construct secret sharing schemes, where the share size does not grow with an increasing number n of parties. Their approach requires the secrets to be exponentially large in n .

A different line of works, which is closely related to secret sharing, deals with the conditional disclosure of secrets (CDS) problem [20]. In this setting, n parties have a common secret s , some common randomness r and separate inputs x_i . The goal of the parties is to each send a single message to a referee, who should learn the secret s iff the inputs x_i satisfy some publicly known predicate F , i.e. if $F(x_1, \dots, x_n) = 1$. Beimel et al. [8] show that any predicate F can be realized with communication complexity $\mathcal{O}(2^{n/2})$ and subsequently Liu, Vaikuntanathan and Wee [29] improve this upper bound to $2^{\tilde{\mathcal{O}}(\sqrt{n})}$. Gay, Kerenidis, and Wee [19] prove a lower bound for CDS schemes, which, very roughly speaking, shows that the communication complexity of any CDS is at least as large as the one-way communication complexity of the underlying predicate F . The techniques in this work are similar to some of the techniques in their work and our lower bound proof can be seen as a non-trivial generalization of their initial proof strategy to the case of secret sharing schemes. In contrast to their linear lower bound, our lower bound here is exponential.

Apart from being an interesting primitive on its own, CDS is also the main building block underlying the secret sharing scheme of Liu and Vaikuntanathan [28] described above. All of the CDS schemes mentioned above run in time exponential in n for certain predicates F .

Despite all progress that was made towards understanding the complexity of secret sharing, a lower bound on the share size of secret sharing schemes for general access structures remained out of reach.

1.1 Our Contribution

In this work we make some progress towards proving Beimel's conjecture. Informally, we show that either the total share size or the computational effort for reconstructing the secret has to be exponential in n . A bit more formally, let us consider a secret sharing scheme Σ for some access structure \mathcal{A} that takes a 1-bit secret as input and outputs n shares, which are at most k bits long in total. Let \mathcal{F} be some family of reconstruction functions. We require that for any authorized subset of parties $A \subseteq \mathcal{A}$, there exists at least one function in \mathcal{F} that these parties can use to reconstruct the correct secret with probability at least $3/4$. For any $A \notin \mathcal{A}$, we require that all functions in \mathcal{F} reconstruct the correct secret with probability at most $1/4$. These correctness and privacy requirements are very weak. Neither do we require perfect correctness, nor do we require privacy against an unauthorized set of parties that may use some function outside of \mathcal{F} to reconstruct the secret. Proving a lower bound for such a secret sharing scheme makes our result only stronger, since any lower bound we can prove here

also applies to any secret sharing scheme with better correctness and privacy guarantees. In this work we prove:

Theorem 1 (Informal). *For any secret sharing scheme Σ for general access structures, with domain of secrets $\{0, 1\}$ and total share length k , then there exists an access structure \mathcal{A} such that*

$$\lg(|\mathcal{F}|) \cdot k = \Omega(2^n / \sqrt{n}).$$

Our result does not prove Beimel’s conjecture, but it tells us that any secret sharing scheme for 1-bit secrets for general access structures, which has a reconstruction function whose description is sub-exponentially large in n , must have a share size that is exponential in n . In particular, this holds even if the secret sharing itself runs in time exponential in n .

To get a better feeling of what \mathcal{F} is, one can, for example, imagine it to be the set of all functions from $\{0, 1\}^k \rightarrow \{0, 1\}$ that are computable by a constant fan-in boolean circuit of some size $t(k) \geq k$. Any one circuit can compute exactly one function, there are a constant amount of different gates types, and for any gate with constant fan-in, there are $t(k)^{O(1)}$ choices for the input wires. It follows that there are at most $t(k)^{O(t(k))}$ different reconstruction functions in \mathcal{F} . Now, if for example $t(k) \leq k^c$ for a constant $c \geq 1$ (decoding by a circuit of size polynomial in the secret share length), then our theorem says that there exists an access structure \mathcal{A} for which the share length k must be exponential in n . On the other hand, if k is for example polynomial in n , then our theorem tells us that there exists some access structure \mathcal{A} which requires an exponentially large reconstruction circuit.

We prove Theorem 1 via a counting argument, meaning that we do not explicitly provide an access structure \mathcal{A} that is affected by the lower bound. The high-level idea of our proof is as follows. Assume that there exists some secret sharing scheme $\Sigma_{\mathcal{A}}$ for every access structure \mathcal{A} with the desired correctness and privacy properties and a total share size of $k^{\mathcal{A}} \leq k$. In the first step, we construct a family \mathcal{D} that contains all access structures \mathcal{A} of a certain type and we show that the size of this family is $2^{\Omega(2^n / \sqrt{n})}$. By the pigeon-hole principle, we know that the description of any $\mathcal{A} \in \mathcal{D}$ is at least $\lg |\mathcal{D}| = \Omega(2^n / \sqrt{n})$ bits long. On the other hand, we show that for any $\mathcal{A} \in \mathcal{D}$ one can use $\Sigma_{\mathcal{A}}$ to construct a $\mathcal{O}(\lg |\mathcal{F}| \cdot k)$ -bit long *lossless* encoding from which \mathcal{A} can be uniquely recovered. Combining the two observations directly yields the theorem stated above. The main challenge in realizing this proof idea lies in the construction of an appropriate encoding (and decoding) algorithm with the desired efficiency. Our encoding algorithm proceeds in two steps. First, we exploit the correctness and privacy properties of our secret sharing scheme to construct a randomized lossless encoding algorithm that works well for 99% of the sets A in any given \mathcal{A} and encodes them into $\mathcal{O}(\lg |\mathcal{F}| \cdot k)$ bits. A careful analysis reveals that we can simply write out the remaining 1% of $A \in \mathcal{A}$ as part of the encoding and still obtain a lower bound on $\lg |\mathcal{F}| \cdot k$.

Proving lower bounds via such encoding arguments has been done quite extensively in the area of data structure lower bounds, see e.g. [15, 24, 25, 31, 32, 36]

and was also used recently to prove optimality of the Johnson-Lindenstrauss lemma in dimensionality reduction [26] and to prove optimality of ORAMs without balls-in-bins assumptions [27].

Remark. At first sight it may seem that the reconstruction function must take the access structure as input to be able to reconstruct. If this was the case, then our lower bound would be meaningless, since we can construct exponentially large access structures. This, however, is *not* the case. Consider the following trivial secret sharing scheme for some bitstring x among parties p_1, \dots, p_n for any access structure \mathcal{A} . For each authorized set $A = \{p_{i_1}, \dots, p_{i_m}\} \in \mathcal{A}$, we pick uniformly random s_{i_1}, \dots, s_{i_m} such that $x = s_{i_1} \oplus \dots \oplus s_{i_m}$ and give (A, s_{i_j}) to p_{i_j} . If a set A wants to reconstruct a secret, they check whether they have a share corresponding to A and xor their corresponding shares together if this is the case. If they do not have a share corresponding to A , they conclude that they are not an authorized set. Note that reconstruction function *only* takes the shares as input and not the access structure itself.

2 Formal Model and Result

In this section, we formally define secret sharing schemes and the precise conditions under which our lower bounds holds. Except for the security requirements, we define a secret sharing scheme precisely as in [6].

Definition 1. Let $\{p_1, \dots, p_n\}$ be a set of parties. A collection $\mathcal{A} \subseteq 2^{\{p_1, \dots, p_n\}}$ is monotone if $B \in \mathcal{A}$ and $B \subseteq C$ imply $C \in \mathcal{A}$. An access structure is a monotone collection $\mathcal{A} \subseteq 2^{\{p_1, \dots, p_n\}}$ of non-empty subsets of $\{p_1, \dots, p_n\}$. Sets in \mathcal{A} are called authorized, and sets not in \mathcal{A} are called unauthorized.

Definition 2. Let $\{p_1, \dots, p_n\}$ be a set of parties. A distribution scheme $\Sigma = (\Pi, \mu)$ with domain of secrets $\{0, 1\}$ is a pair, where μ is a probability distribution on some finite set R called the set of random strings and Π is a mapping from $\{0, 1\} \times R$ to a set of n -tuples $\{0, 1\}^{k_1} \times \dots \times \{0, 1\}^{k_n}$, where $\{0, 1\}^{k_j}$ is called the domain of shares of p_j . A dealer distributes a secret $b \in \{0, 1\}$ according to Σ by first sampling a random string $r \in R$ according to μ , computing a vector of shares $\Pi(b, r) = (s_1, \dots, s_n)$ and privately communicating each share s_j to party p_j . For a set $A \subseteq \{p_1, \dots, p_n\}$, we denote $\Pi(b, r)_A$ as the restriction of $\Pi(b, r)$ to its A -entries.

When designing secret sharing schemes, one would typically consider larger domains of secrets than just a single bit as in Definition 2. In this paper we are proving a lower bound, so focusing on the simplest possible setting of a secret consisting of a single bit only makes our lower bound stronger and the proof simpler. The lower bound we prove in this paper holds for secret sharing schemes that are computationally more efficient when authorized parties reconstruct the secret than when unauthorized parties attempt to. We define this formally in the following:

Definition 3. Let $\{p_1, \dots, p_n\}$ be a set of parties, let $\mathcal{A} \subseteq 2^{\{p_1, \dots, p_n\}}$ be an access structure and $\Sigma = (\Pi, \mu)$ a distribution scheme with domain of secrets $\{0, 1\}$ and domain of shares $\{0, 1\}^{k_1} \times \dots \times \{0, 1\}^{k_n}$. Let \mathcal{F} be a family of functions from $\cup_{i=1}^{\infty} (\{0, 1\}^i \rightarrow \{0, 1\})$ and let \mathcal{U} be the uniform distribution on $\{0, 1\}$. We say that $(\mathcal{F}, \mathcal{A}, \Sigma)$ is an efficient secret sharing scheme if it satisfies the following two conditions:

For any $A \in \mathcal{A}$, there exists a function $f_A \in (\mathcal{F} \cap (\{0, 1\}^{\sum_{j \in A} k_j} \rightarrow \{0, 1\}))$ such that

$$\left| \Pr_{b \sim \mathcal{U}, r \sim \mu} [f_A(\Pi(b, r)_A) = b] - \Pr_{b \sim \mathcal{U}, r \sim \mu} [f_A(\Pi(b, r)_A) \neq b] \right| \geq 3/4.$$

For any $A \notin \mathcal{A}$, it holds for all functions $f \in (\mathcal{F} \cap (\{0, 1\}^{\sum_{j \in A} k_j} \rightarrow \{0, 1\}))$ that

$$\left| \Pr_{b \sim \mathcal{U}, r \sim \mu} [f(\Pi(b, r)_A) = b] - \Pr_{b \sim \mathcal{U}, r \sim \mu} [f(\Pi(b, r)_A) \neq b] \right| \leq 1/4.$$

For intuition on Definition 3, consider as an example instantiating \mathcal{F} to be the set that contains for each i , the set of all functions from $\{0, 1\}^i \rightarrow \{0, 1\}$ that are computable by a constant fan-in boolean circuit of size $t(i) \leq i^c$ for a constant $c > 1$, i.e. \mathcal{F} contains functions computable by polynomially sized circuits. With this choice of \mathcal{F} , consider an access structure \mathcal{A} . A distribution scheme Σ gives an efficient secret sharing scheme $(\mathcal{F}, \mathcal{A}, \Sigma)$ precisely if any authorized set of parties $A \in \mathcal{A}$ can recover the secret using *some* constant fan-in boolean circuit with size polynomial in the share length, whereas no unauthorized set of parties can recover the secret using *any* constant fan-in boolean circuit with size polynomial in the share length. We can thus think of \mathcal{F} as defining the computational resources with which authorized parties can recover the secret, but unauthorized parties cannot. For ease of notation, define $\mathcal{F}_{\leq k}$ as

$$\mathcal{F}_{\leq k} := \mathcal{F} \cap (\cup_{i=1}^k (\{0, 1\}^i \rightarrow \{0, 1\}))$$

and define

$$\mathcal{F}_{=k} := \mathcal{F} \cap (\{0, 1\}^k \rightarrow \{0, 1\}).$$

in the remainder of the paper.

Discussion 1. When designing secret sharing schemes, one would typically insist that authorized parties can reconstruct the secret with probability $1 - \text{negl}(n)$. Similarly, one would insist that unauthorized parties cannot reconstruct the secret except with probability $\text{negl}(n)$. Since we are proving a lower bound, using the constants $3/4$ and $1/4$ in Definition 3 only makes our results stronger.

Discussion 2. One could consider allowing randomization in the algorithms used for reconstructing the secret, both for the authorized and unauthorized parties. That is, a natural extension of Definition 3 would say that *there exists a distribution γ_A over functions in*

$$\left(\mathcal{F} \cap \left(\{0, 1\}^{\sum_{j \in A} k_j} \rightarrow \{0, 1\} \right) \right)$$

such that $\Pr_{b \sim \mathcal{U}, r \sim \mu, f_A \sim \gamma_A} [\dots]$. We remark that the definition would be equivalent to Definition 3 since one can always fix the randomness in f_A to achieve the same guarantees (equivalent to one direction of Yao’s minimax principle).

Discussion 3. Our definition may seem superficially similar to the definition of weakly-private secret sharing schemes by Beimel and Franklin [7]. Their definition states that any unauthorized set cannot exclude any potential secret with probability 1. It does, however, allow the adversary to guess the secret correctly with a probability that is arbitrarily close to 1. In contrast to their definition, ours is strictly stronger, since it requires a sharp upper bound on the probability that an unqualified set of parties guesses the correct secret.

We are ready to present our main theorem in its full generality:

Theorem 2. *Let \mathcal{F} be a family of functions from $\cup_{i=1}^\infty (\{0, 1\}^i \rightarrow \{0, 1\})$ and let $\{p_1, \dots, p_n\}$ be a set of parties. There exists an access structure $\mathcal{A} \subseteq 2^{\{p_1, \dots, p_n\}}$ such that any efficient secret sharing scheme $(\mathcal{F}, \mathcal{A}, \Sigma)$ with domain of secrets $\{0, 1\}$ and domain of shares $\{0, 1\}^{k_1} \times \dots \times \{0, 1\}^{k_n}$ with $k = \sum_j k_j$, satisfies*

$$\lg(|\mathcal{F}_{\leq k}|) \cdot k = \Omega(2^n / \sqrt{n}).$$

To appreciate Theorem 2, consider instantiating \mathcal{F} to be the set that contains for each i , the set of all functions from $\{0, 1\}^i \rightarrow \{0, 1\}$ that are computable by a constant fan-in boolean circuit of size $t(i)$ (with $t(i) \geq i$). A simple counting argument shows that $|\mathcal{F}_{\leq k}| \leq t(k)^{O(t(k))}$ (A circuit computes only one function and there are $t(k)^{O(1)}$ choices for the input wires to each gate, there are $O(1)$ choices for the function computed by each gate, and there are $t(k)$ gates). Theorem 2 thus gives us that there must exist an access structure \mathcal{A} such that any efficient secret sharing scheme $(\mathcal{F}, \mathcal{A}, \Sigma)$ with domain of shares $\{0, 1\}^{k_1} \times \dots \times \{0, 1\}^{k_n}$ with $k = \sum_j k_j$ must satisfy $t(k) \lg(t(k))k = \Omega(2^n / n^{1/2})$. If we plug in polynomially sized constant fan-in boolean circuits, i.e. $t(i) \leq i^c$ for a constant $c \geq 1$, this gives us that $k^{c+2} = \Omega(2^n / n^{1/2}) \Rightarrow k = 2^{\Omega(n)}$, i.e. any secret sharing scheme for \mathcal{A} must have shares with exponential length if decoding can be done by constant fan-in boolean circuits with size polynomial in the share length. Moreover, the lower bound holds *even if we only require* that unauthorized parties cannot reconstruct the secret using a polynomially sized constant fan-in boolean circuit (polynomial in the length of the shares). Notice that since this is a lower bound, it only makes the result stronger than if we e.g. required that a computationally unbounded set of parties cannot reconstruct the secret. We can also deduce from Theorem 2 that the size of the decoding circuit must be exponential in n , regardless of the share length.

Another interesting instantiation of Theorem 2 is to let \mathcal{F} consist of all functions computable by a Turing machine with at most 10^6 states and alphabet $\{0, 1\}$ (or some other constant number of states). Then $|\mathcal{F} \cap (\{0, 1\}^i \rightarrow \{0, 1\})| = O(1)$ and the lower bound says that there exists an access structure \mathcal{A} for which any efficient secret sharing scheme $(\mathcal{F}, \mathcal{A}, \Sigma)$ must satisfy $k^2 = \Omega(2^n/\sqrt{n}) \Rightarrow k = 2^{\Omega(n)}$, i.e. shares must have exponential length if the secret can be reconstructed by authorized parties using a Turing machine with at most 10^6 states and binary alphabet. The lower bound holds as long as we require that unauthorized parties cannot recover the secret using a Turing machine with at most 10^6 states and alphabet $\{0, 1\}$.

An even more exotic instantiation of Theorem 2 follows by letting \mathcal{F} contain, for every i , the set of functions from $\{0, 1\}^i \rightarrow \{0, 1\}$ that are computable by a C-program with up to t ASCII characters. A counting argument shows that $|\mathcal{F}_{\leq k}| \leq k2^{O(t)}$ (there are $2^{O(t)}$ sequences of t ASCII characters, and any program computes at most one function from $\{0, 1\}^i \rightarrow \{0, 1\}$) and we conclude that it must be the case that there exists an access structure \mathcal{A} such that any efficient secret sharing $(\mathcal{F}, \mathcal{A}, \Sigma)$ must have $(t + \lg k) \cdot k = \Omega(2^n/\sqrt{n})$. This means that either the length of the C-program has to grow exponentially with the number of parties n , or the length of the shares has to grow exponentially with n . Thus if we insist on short shares, then the C-programs for reconstructing the secret have to be extremely non-uniform, and if we insist on reconstructing secrets using C-programs of any constant length t independent of n , then the shares must have exponential length. This lower bound holds as long as we require that unauthorized parties cannot recover the secret via a C-program of length t or less.

Finally, if one insist that authorized parties can *efficiently* reconstruct the secret via a C-program of length at most t ASCII characters, then the previous lower bound is strengthened. That is, we can now let \mathcal{F} contain, for every i , the set of functions from $\{0, 1\}^i \rightarrow \{0, 1\}$ that are computable by a C-program with up to t ASCII characters that terminates in at most h steps. If we insist that authorized parties can reconstruct the secret by running such a C-program, then the lower bound $(t + \lg k) \cdot k = \Omega(2^n/\sqrt{n})$ holds even if we only require that unauthorized parties cannot reconstruct the secret via a C-program of length t and running time at most h steps.

3 Lower Bound Proof

To prove Theorem 2, let $\{p_1, \dots, p_n\}$ be a set of parties and let \mathcal{F} be a family of functions from $\cup_{i=1}^\infty (\{0, 1\}^i \rightarrow \{0, 1\})$. Assume that there is a parameter k such that it holds for all access structures $\mathcal{A} \subseteq 2^{\{p_1, \dots, p_n\}}$, that there exists an efficient secret sharing scheme $(\mathcal{F}, \mathcal{A}, (\Pi_{\mathcal{A}}, \mu_{\mathcal{A}}))$ with domain of secrets $\{0, 1\}$ and domain of shares $\{0, 1\}^{k_1^{\mathcal{A}}} \times \dots \times \{0, 1\}^{k_n^{\mathcal{A}}}$ with $\sum_j k_j^{\mathcal{A}} = k^{\mathcal{A}} \leq k$.

We will prove a lower bound on $\lg(|\mathcal{F}_{\leq k}|) \cdot k$ via a counting argument. The high level intuition is that two distinct access structures \mathcal{A}_1 and \mathcal{A}_2 must be different either in terms of the shares they use, or in terms of the procedures used for

reconstructing the secrets. Since there are overwhelmingly many distinct access structures, this gives a lower bound on either the share length (a lower bound on k), or on the descriptive size of the procedures used for reconstructing secrets (a lower bound on $\lg(|\mathcal{F}_{\leq k}|)$).

More formally, let \mathcal{D} be the family containing all access structures $\mathcal{A} \subseteq 2^{\{p_1, \dots, p_n\}}$ such that \mathcal{A} contains no sets A of cardinality less than $\lfloor n/2 \rfloor$ and \mathcal{A} contains all sets A of cardinality more than $\lfloor n/2 \rfloor$. We claim that $|\mathcal{D}| = 2^{\binom{n}{\lfloor n/2 \rfloor}} = 2^{\Omega(2^n / \sqrt{n})}$. To see this, observe that \mathcal{A} is monotone for any choice of subsets with cardinality $\lfloor n/2 \rfloor$ that we might include in it. Since there are $\binom{n}{\lfloor n/2 \rfloor}$ subsets of cardinality $\lfloor n/2 \rfloor$, we conclude that there are $2^{\binom{n}{\lfloor n/2 \rfloor}}$ ways of choosing which subsets to include in \mathcal{A} .

We will show that we can encode any $\mathcal{A} \in \mathcal{D}$ into

$$\lambda = O(\lg(|\mathcal{F}_{\leq k}|) \cdot k) + 0.1 \cdot \binom{n}{\lfloor n/2 \rfloor}$$

bits and still uniquely recover \mathcal{A} from the encoding alone. The encoding procedure thus defines an injective mapping from \mathcal{D} to $\{0, 1\}^\lambda$. By the pigeon-hole principle, this implies that

$$\begin{aligned} \lambda &\geq \lg |\mathcal{D}| \Rightarrow \\ O(\lg(|\mathcal{F}_{\leq k}|) \cdot k) &\geq 0.9 \cdot \binom{n}{\lfloor n/2 \rfloor} \Rightarrow \\ \lg(|\mathcal{F}_{\leq k}|) \cdot k &= \Omega(2^n / \sqrt{n}). \end{aligned}$$

We are now ready to describe our encoding and decoding procedures.

Encoding. Let $\mathcal{A} \in \mathcal{D}$. Our procedure for uniquely encoding \mathcal{A} is as follows:

- For $i = 1, \dots, T$ for a parameter T to be fixed, consider sampling $b_i \sim \mathcal{U}$ as a uniform random bit, and sample $r_i \sim \mu_{\mathcal{A}}$. Let $A \subseteq \{p_1, \dots, p_n\}$ be an arbitrary set of cardinality $\lfloor n/2 \rfloor$ and define $k_A^{\mathcal{A}} = \sum_{j \in A} k_j^{\mathcal{A}}$. By Definition 3, it holds that:

- If $A \in \mathcal{A}$, then there exists a function $f_A \in \mathcal{F}_{=k_A^{\mathcal{A}}}$ such that

$$\left| \Pr_{b_i, r_i} [f_A(\Pi_{\mathcal{A}}(b_i, r_i)_A) = b_i] - \Pr_{b_i, r_i} [f_A(\Pi_{\mathcal{A}}(b_i, r_i)_A) \neq b_i] \right| \geq 3/4.$$

- If $A \notin \mathcal{A}$, then for all functions $f \in \mathcal{F}_{=k_A^{\mathcal{A}}}$, it holds that

$$\left| \Pr_{b_i, r_i} [f(\Pi_{\mathcal{A}}(b_i, r_i)_A) = b_i] - \Pr_{b_i, r_i} [f(\Pi_{\mathcal{A}}(b_i, r_i)_A) \neq b_i] \right| \leq 1/4.$$

We use this observation as follows: We set $T = c \lg |\mathcal{F}_{\leq k}|$ for a sufficiently large constant $c > 1$. If $A \notin \mathcal{A}$, then since $|\mathcal{F}_{=k_A^{\mathcal{A}}}| \leq |\mathcal{F}_{\leq k}|$, we can use a

Chernoff bound and a union bound over all $f \in \mathcal{F}_{=k^A}$ to conclude that with probability at least 99/100, it holds simultaneously for all $f \in \mathcal{F}_{=k^A}$ that

$$|\{i : f(\Pi_{\mathcal{A}}(b_i, r_i)_A) = b_i\}| - |\{i : f(\Pi_{\mathcal{A}}(b_i, r_i)_A) \neq b_i\}| < T/3.$$

At the same time, if $A \in \mathcal{A}$ and we have $T = c \lg |\mathcal{F}_{\leq k}|$, then with overwhelming probability, we will have that there exists at least one function $f \in \mathcal{F}_{=k^A}$ such that

$$|\{i : f(\Pi_{\mathcal{A}}(b_i, r_i)_A) = b_i\}| - |\{i : f(\Pi_{\mathcal{A}}(b_i, r_i)_A) \neq b_i\}| > T/3.$$

Thus intuitively, the variables b_1, \dots, b_T and r_1, \dots, r_T reveal whether A is in \mathcal{A} or not, i.e. they carry information about \mathcal{A} . We exploit this as follows: Let $\chi_{\mathcal{A}}$ be the random variable taking the value 1 if the test

$$\begin{aligned} \exists f \in \mathcal{F}_{=k^A} : \\ |\{i : f(\Pi_{\mathcal{A}}(b_i, r_i)_A) = b_i\}| - |\{i : f(\Pi_{\mathcal{A}}(b_i, r_i)_A) \neq b_i\}| > T/3? \end{aligned}$$

correctly predicts whether $A \in \mathcal{A}$. Then $\Pr[\chi_{\mathcal{A}} = 1] \geq 99/100$. Let \mathcal{S} be the family of all subsets of $\{p_1, \dots, p_n\}$ that have cardinality $\lfloor n/2 \rfloor$. It follows by linearity of expectation that $\mathbb{E}[\sum_{A \in \mathcal{S}} \chi_A] \geq 99|\mathcal{S}|/100$. This means that there must exist a choice values $\hat{b}_1, \dots, \hat{b}_T$ and $\hat{r}_1, \dots, \hat{r}_T$ such that the test $\exists f \in \mathcal{F}_{=k^A} : |\{i : f(\Pi_{\mathcal{A}}(\hat{b}_i, \hat{r}_i)_A) = \hat{b}_i\}| - |\{i : f(\Pi_{\mathcal{A}}(\hat{b}_i, \hat{r}_i)_A) \neq \hat{b}_i\}| > T/3?$ correctly predicts whether $A \in \mathcal{A}$ for at least $99|\mathcal{S}|/100$ sets $A \in \mathcal{S}$. Fix such values.

2. Write down $\lg k$ bits specifying k^A , followed by k bits specifying k_1^A, \dots, k_n^A (this can be done by writing a length k bit string, where positions $\sum_{i=1}^j k_i^A$ are set to 1 for all $j = 1, \dots, n$). Then write down the bits $\hat{b}_1, \dots, \hat{b}_T$ and $\Pi_{\mathcal{A}}(\hat{b}_1, \hat{r}_1), \dots, \Pi_{\mathcal{A}}(\hat{b}_T, \hat{r}_T)$ for a total of at most $\lg k + k + T(1 + k)$ bits.
3. Let $\bar{\mathcal{S}}$ be the subset of sets from \mathcal{S} where the prediction is incorrect. Encode $\bar{\mathcal{S}}$ as a subset of \mathcal{S} using $\lg \binom{n}{\lfloor n/2 \rfloor} \leq n$ bits to specify $|\bar{\mathcal{S}}|$ and $\lg \binom{|\mathcal{S}|}{|\bar{\mathcal{S}}|} \leq |\bar{\mathcal{S}}| \lg(e|\mathcal{S}|/|\bar{\mathcal{S}}|) \leq (|\mathcal{S}|/100) \lg(100e) < 0.1 \cdot \binom{n}{\lfloor n/2 \rfloor}$ bits to specify the subset.

Next we argue how to recover \mathcal{A} from the above encoding:

Decoding.

1. Read the first $\lg k + k$ bits to recover k^A and k_1^A, \dots, k_n^A . Then use the following $T(k + 1)$ bits to recover $\hat{b}_1, \dots, \hat{b}_T$ and $\Pi_{\mathcal{A}}(\hat{b}_1, \hat{r}_1), \dots, \Pi_{\mathcal{A}}(\hat{b}_T, \hat{r}_T)$.
2. For each $A \in \mathcal{S}$, iterate over all $f \in \mathcal{F}_{=k^A}$ and compute the value

$$\Delta_f := \left| |\{i : f(\Pi_{\mathcal{A}}(\hat{b}_i, \hat{r}_i)_A) = \hat{b}_i\}| - |\{i : f(\Pi_{\mathcal{A}}(\hat{b}_i, \hat{r}_i)_A) \neq \hat{b}_i\}| \right|.$$

Observe that the decoder can extract $\Pi_{\mathcal{A}}(\hat{b}_i, \hat{r}_i)_A$ from $\Pi_{\mathcal{A}}(\hat{b}_i, \hat{r}_i)$ since the decoder knows k_1^A, \dots, k_n^A . Thus the decoder can indeed compute Δ_f . If there is at least one f with $\Delta_f \geq T/3$, we initially predict that $A \in \mathcal{A}$ and otherwise, we predict that $A \notin \mathcal{A}$. These predictions are correct, except for $A \in \bar{\mathcal{S}}$.

3. Finally we read the last part of the encoding to determine which sets A that were predicted incorrectly in step 2. Together with the correct predictions from step 2., this recovers \mathcal{A} .

Analysis. Finally we derive the lower bound. We have just argued that we can give a unique encoding of each $\mathcal{A} \in \mathcal{D}$, hence the length of the encoding must be at least $\lg |\mathcal{D}| = \binom{n}{\lfloor n/2 \rfloor}$ bits. But the above encoding uses at most:

$$\lg k + k + T(1 + k) + n + 0.1 \cdot \binom{n}{\lfloor n/2 \rfloor}$$

bits. Thus we must have

$$\begin{aligned} \lg k + k + T(1 + k) + n + 0.1 \cdot \binom{n}{\lfloor n/2 \rfloor} &\geq \binom{n}{\lfloor n/2 \rfloor} \Rightarrow \\ Tk &= \Omega\left(\binom{n}{\lfloor n/2 \rfloor}\right) = \Omega(2^n / \sqrt{n}). \end{aligned}$$

But $T = c \lg |\mathcal{F}_{\leq k}|$ and we conclude:

$$\lg |\mathcal{F}_{\leq k}| \cdot k = \Omega(2^n / \sqrt{n}).$$

References

1. Applebaum, B., Arkis, B.: On the power of amortization in secret sharing: d -uniform secret sharing and CDS with constant information rate. In: Beimel, A., Dziembowski, S. (eds.) TCC 2018. LNCS, vol. 11239, pp. 317–344. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03807-6_12
2. Applebaum, B., Arkis, B., Raykov, P., Vasudevan, P.N.: Conditional disclosure of secrets: amplification, closure, amortization, lower-bounds, and separations. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10401, pp. 727–757. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63688-7_24
3. Applebaum, B., Beimel, A., Farràs, O., Nir, O., Peter, N.: Secret-sharing schemes for general and uniform access structures. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019. LNCS, vol. 11478, pp. 441–471. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17659-4_15
4. Babai, L., Gál, A., Wigderson, A.: Superpolynomial lower bounds for monotone span programs. *Combinatorica* **19**(3), 301–319 (1999)
5. Beimel, A.: Secure schemes for secret sharing and key distribution. Technion-Israel Institute of technology, Faculty of computer science (1996)
6. Beimel, A.: Secret-sharing schemes: a survey. In: Chee, Y.M., et al. (eds.) IWCC 2011. LNCS, vol. 6639, pp. 11–46. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20901-7_2
7. Beimel, A., Franklin, M.: Weakly-private secret sharing schemes. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 253–272. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-70936-7_14
8. Beimel, A., Ishai, Y., Kumaresan, R., Kushilevitz, E.: On the cryptographic complexity of the worst functions. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 317–342. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54242-8_14

9. Beimel, A., Orlov, I.: Secret sharing and non-Shannon information inequalities. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 539–557. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-00457-5_32
10. Ben-Or, M., Goldwasser, S., Wigderson, A.: Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In: 20th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, 2–4 May 1988, pp. 1–10. ACM Press (1988)
11. Blakley, G.R.: Safeguarding cryptographic keys, pp. 313–317. AFIPS Press (1979)
12. Blundo, C., De Santis, A., Gargano, L., Vaccaro, U.: On the information rate of secret sharing schemes. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 148–167. Springer, Heidelberg (1993). https://doi.org/10.1007/3-540-48071-4_11
13. Bogdanov, A., Guo, S., Komargodski, I.: Threshold secret sharing requires a linear size alphabet. In: Hirt, M., Smith, A. (eds.) TCC 2016. LNCS, vol. 9986, pp. 471–484. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53644-5_18
14. Capocelli, R.M., De Santis, A., Gargano, L., Vaccaro, U.: On the size of shares for secret sharing schemes. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 101–113. Springer, Heidelberg (1992). https://doi.org/10.1007/3-540-46766-1_7
15. Chakraborty, D., Kamra, L., Larsen, K.G.: Tight cell probe bounds for succinct Boolean matrix-vector multiplication. In: Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, 25–29 June 2018, pp. 1297–1306 (2018)
16. Chaum, D., Crépeau, C., Damgård, I.: Multiparty unconditionally secure protocols (extended abstract). In: 20th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, 2–4 May 1988, pp. 11–19. ACM Press (1988)
17. Csirmaz, L.: The size of a share must be large. In: De Santis, A. (ed.) EUROCRYPT 1994. LNCS, vol. 950, pp. 13–22. Springer, Heidelberg (1995). <https://doi.org/10.1007/BFb0053420>
18. Csirmaz, L.: The dealer’s random bits in perfect secret sharing schemes. *Studia Scientiarum Mathematicarum Hungarica* **32**(3), 429–438 (1996)
19. Gay, R., Kerenidis, I., Wee, H.: Communication complexity of conditional disclosure of secrets and attribute-based encryption. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9216, pp. 485–502. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48000-7_24
20. Gertner, Y., Ishai, Y., Kushilevitz, E., Malkin, T.: Protecting data privacy in private information retrieval schemes. In: 30th Annual ACM Symposium on Theory of Computing, Dallas, TX, USA, 23–26 May 1998, pp. 151–160. ACM Press (1988)
21. Ito, M., Saito, A., Nishizeki, T.: Secret sharing scheme realizing general access structure. *Electron. Commun. Jpn. (Part III Fundam. Electron. Sci.)* **72**(9), 56–64 (1989)
22. Karchmer, M., Wigderson, A.: On span programs. In: Proceedings of the Eighth Annual Structure in Complexity Theory Conference 1993, pp. 102–111. IEEE (1993)
23. Karnin, E., Greene, J., Hellman, M.: On secret sharing systems. *IEEE Trans. Inf. Theory* **29**(1), 35–41 (1983)
24. Larsen, K.G.: The cell probe complexity of dynamic range counting. In: Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, 19–22 May 2012, pp. 85–94 (2012)
25. Larsen, K.G.: Higher cell probe lower bounds for evaluating polynomials. In: 53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, 20–23 October 2012, pp. 293–301 (2012)

26. Larsen, K.G., Nelson, J.: Optimality of the Johnson-Lindenstrauss lemma. In: 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, 15–17 October 2017, pp. 633–638 (2017)
27. Larsen, K.G., Nielsen, J.B.: Yes, there is an oblivious RAM lower bound!. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10992, pp. 523–542. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96881-0_18
28. Liu, T., Vaikuntanathan, V.: Breaking the circuit-size barrier in secret sharing. In: Diakonikolas, I., Kempe, D., Henzinger, M. (eds.) 50th Annual ACM Symposium on Theory of Computing, Los Angeles, CA, USA, 25–29 June 2018, pp. 699–708. ACM Press (2018)
29. Liu, T., Vaikuntanathan, V., Wee, H.: Conditional disclosure of secrets via non-linear reconstruction. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10401, pp. 758–790. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63688-7_25
30. Martín, S., Padró, C., Yang, A.: Secret sharing, rank inequalities and information inequalities. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8043, pp. 277–288. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_16
31. Pătraşcu, M., Demaine, E.D.: Logarithmic lower bounds in the cell-probe model. *SIAM J. Comput.* **35**(4), 932–963 (2006)
32. Pătraşcu, M., Viola, E.: Cell-probe lower bounds for succinct partial sums. In: Proceedings of the 21st ACM/SIAM Symposium on Discrete Algorithms (SODA), pp. 117–122 (2010)
33. Pitassi, T., Robere, R.: Lifting Nullstellensatz to monotone span programs over any field. In: Diakonikolas, I., Kempe, D., Henzinger, M. (eds.) 50th Annual ACM Symposium on Theory of Computing, Los Angeles, CA, USA, 25–29 June 2018, pp. 1207–1219. ACM Press (2018)
34. Rabin, T., Ben-Or, M.: Verifiable secret sharing and multiparty protocols with honest majority (extended abstract). In: 21st Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, 15–17 May 1989, pp. 73–85. ACM Press (1989)
35. Shamir, A.: How to share a secret. *Commun. ACM* **22**(11), 612–613 (1979)
36. Verbin, E., Zhang, Q.: The limits of buffering: a tight lower bound for dynamic membership in the external memory model. *SIAM J. Comput.* **42**(1), 212–229 (2013)