Electronic Colloquium on Computational Complexity, Report No. 131 (2016)

Threshold Secret Sharing Requires a Linear Size Alphabet

Andrej Bogdanov * Siyao Guo † Ilan Komargodski ‡

August 21, 2016

Abstract

We prove that for every n and $1 < t < n$ any t-out-of-n threshold secret sharing scheme for one-bit secrets requires share size $\log(t+1)$. Our bound is tight when $t = n - 1$ and n is a prime power. In 1990 Kilian and Nisan proved the incomparable bound $\log(n - t + 2)$. Taken together, the two bounds imply that the share size of Shamir's secret sharing scheme (Comm. ACM '79) is optimal up to an additive constant even for one-bit secrets for the whole range of parameters $1 < t < n$.

More generally, we show that for all $1 < s < r < n$, any ramp secret sharing scheme with secrecy threshold s and reconstruction threshold r requires share size $\log((r+1)/(r-s))$.

As part of our analysis we formulate a simple game-theoretic relaxation of secret sharing for arbitrary access structures. We prove the optimality of our analysis for threshold secret sharing with respect to this method and point out a general limitation.

1 Introduction

In 1979, Shamir [Sha79] and Blakley [Bla79] presented a method for sharing a piece of secret information among n parties such that any $1 < t < n$ parties can recover the secret while any $t - 1$ parties learn *nothing* about the secret. These methods are called (t, n)-threshold secret sharing schemes. This sharp threshold between secrecy and reconstruction is fundamental in applications where a group of mutually suspicious individuals with conflicting interests must cooperate. Indeed, threshold secret sharing schemes have found many applications in

*Chinese University of Hong Kong. Email: andrejb@cse.cuhk.edu.hk. Supported by RGC GRF grants CUHK410113 and CUHK14208215.

†New York University. Email: sguo@cims.nyu.edu. Part of the work done in the Chinese University of Hong Kong supported by RGC GRF grants CUHK410112 and CUHK410113.

‡Weizmann Institute of Science, Israel. Email: ilan.komargodski@weizmann.ac.il. Part of this work done while visiting CUHK, supported by RGC GRF grant CUHK410113. Supported in part by a Levzion fellowship, by a grant from the I-CORE Program of the Planning and Budgeting Committee, the Israel Science Foundation, BSF and the Israeli Ministry of Science and Technology.

ISSN 1433-8092

cryptography and distributed computing; see the extensive survey of Beimel [Bei11] and the recent book of Cramer et al. [CDN15].

Threshold secret sharing was generalized by Ito et al. [ISN93] to allow more general structures of subsets to learn the secret, while keeping the secret perfectly hidden from all other subsets. The collection of qualified subsets is called an access structure.

A significant goal in secret sharing is to minimize the share size, namely, the amount of information distributed to the parties. Despite the long history of the subject, there are significant gaps between lower and upper bounds both for general access structures and for the special case of threshold structures.

Threshold access structures. For (t, n)-threshold access structures (denoted by THR_t^n) and a 1-bit secret, Shamir [Sha79] gave a very elegant and efficient scheme: the dealer picks a random polynomial of degree $t-1$ conditioned on setting the free coefficient to be the secret, and gives the i-th party the evaluation of the polynomial at the point i. The computation is done over a field $\mathbb{F}$ of size $q > n$.

The correctness follows because one can recover the unique polynomial from any t points (and thus recover the secret). Security follows by a counting argument showing that given less than t points, all possibilities for the free coefficient are equally likely. The share of each party is an element in the field $\mathbb{F}$ that can be represented using $\log q \approx \log n$ bits (all our logarithms are base 2). The efficiency of this scheme makes it very attractive for applications.

A natural question to ask is whether $\log n$-bit shares are necessary for sharing a 1-bit secret for threshold access structures. Kilian and Nisan [KN90][1] showed that $\log n$ bits are necessary when t is not too large. Specifically, they showed a $\log(n - t + 2)$ lower bound on share size for (t, n)-threshold schemes. For large values of t, especially those close to n, their bound does not rule out schemes with shares much shorter than $\log n$ bits. Their bound leaves open the possibility that, in particular, $(n - 1, n)$-threshold schemes with two-bit shares exist.

Ramp schemes are a generalization of threshold schemes that allow for a gap between the secrecy and reconstruction parameters. In an (s, r, n)-ramp scheme, we require that any subset of at least r parties can recover the secret, while any subset of size at most s cannot learn anything about the secret.[2] When $r = s + 1$, an (s, r, n)-ramp scheme is exactly an (r, n)-threshold scheme. Ramp schemes, defined by Blakley and Meadows [BM84], are useful for various applications (see e.g. [SW99, CC06, MPS11]) since if $r - s$ is large, they can sometimes be realized with *shorter* shares than standard threshold schemes (especially in the case of long secret).

Generalizing the lower bound of Kilian and Nisan, Cascudo et al. [CCX13] showed that

[1] Their result is unpublished and independently obtained (and generalized in various ways) by [CCX13]. The original argument of Kilian an Nisan appears in [CCX13, Appendix A] and was referenced earlier in [BC94, Bei96, BF07].

[2] Another common definition (See [FHKP14, Definition 2.7] and [FMP16, Example 2.11] for examples) for a ramp scheme is where the information about the secret increases with the size of the set. We focus only on the definition in which sets of size below a certain threshold have no information about the secret, while sets of size larger than some threshold can recover it.

$\log((n-s+1)/(r-s))$-bit shares are necessary to realize an (s,r,n)-ramp scheme. When $s = n - O(1)$, however, their share size bound is a constant independent of n. Paterson and Stinson [PS13] showed that this bound is tight for specific small values of s.

General access structures. For most access structures, the best known secret sharing schemes require shares of size $2^{O(n)}$ for sharing a 1-bit secret. Specifically, viewing the access structure as a Boolean indicator function for qualified subsets, the schemes of [ISN93, BL88, KW93] result with shares of size proportional to the DNF/CNF size, monotone formula size, or monotone span program size of the function, respectively. Thus, even for many access structures that can be described by a small monotone uniform circuit, the best schemes have exponential size shares.[3] On the other hand, the best known lower bound on share size for sharing an ℓ-bit secret is $\ell \cdot n/\log n$ bits, by Csirmaz [Csi97] (improving on [CSGV93]).

Bridging the exponential gap between upper and lower bounds is the major open problems in the study of secret sharing schemes. While it is widely believed that the lower bound should be exponential (see e.g. [Bei96, Bei11]), no major progress has been obtained in the last two decades. Moreover, a non-explicit linear lower bound is not known, that is, whether there *exists* an access structure that requires linear size shares.[4]

1.1 Our results

Share size lower bound. We close the gap in share size for threshold secret sharing up to a small additive constant. We assume for simplicity that all parties are given equally long shares.

Theorem 1. *For every $n \in \mathbb{N}$ and $1 < t < n$, any (t,n)-threshold secret sharing scheme for a 1-bit secret requires shares of at least $\log(t+1)$ bits.*

The assumption $1 < t < n$ is necessary, as $(1,n)$-threshold and (n,n)-threshold secret sharing schemes with share size 1 do exist.

Our bound is tight when $t = n-1$ and n is the power of a prime; see Appendix A. By combining Theorem 1 with the lower bound of Kilian and Nisan, we determine the share size of threshold schemes up to a small additive constant. That is, we get that any such scheme requires shares of size

$$\max\{\log(n-t+2), \log(t+1)\} \geq \log \frac{n+3}{2}. \tag{1}$$

Theorem 1 is a special case of the following theorem, which applies more generally to ramp schemes.

[3] One such notable example is the *directed connectivity* access structure: the parties correspond to edge slots in the complete *directed* graph and the qualified subsets are those edges that connect two distinguished nodes s and t.

[4] The usual counting arguments do not work here since one needs to enumerate over the sharing and reconstruction algorithms whose complexity may be larger than the share size.

Theorem 2. *For every $n \in \mathbb{N}$ and $1 \leq s < r < n$, any (s, r, n)-ramp secret sharing scheme for a 1-bit secret requires shares of at least $\log((r+1)/(r-s))$ bits.*

By combining Theorem 2 with the lower bound of [CCX13], we get that any (s, r, n)-ramp secret sharing scheme must have share size at least

$$\max\left\{\log\frac{n-s+1}{r-s}, \log\frac{r+1}{r-s}\right\} \geq \log\frac{n+r-s+2}{2\cdot(r-s)}. \tag{2}$$

Proof technique and limitations. We prove our lower bounds by analyzing a new game-theoretic relaxation of secret sharing. Here, we focus on threshold schemes, although our argument also applies to ramp schemes.

Given an access structure $\mathcal{A}$ and a real-valued parameter $\theta > 0$ we consider the following zero-sum game $G(\mathcal{A}, \theta)$: Alice and Bob pick sets A and B in the access structure $\mathcal{A}$, respectively, and the payoff is $(-\theta)^{|A \setminus B|}$, where $A \setminus B$ denotes set difference. We say Alice wins if she has a strategy with non-negative expected payoff, and Bob wins otherwise.

We show (in Lemma 8) that if Bob wins in the game $G(\mathcal{A}, 1/(q-1))$, then no secret sharing scheme with share size $\log q$ exists. We prove Theorem 2 by constructing such a strategy for Bob.

On the negative side, we show that our analysis is optimal for threshold access structures, so the lower bound in Theorem 1 is tight with respect to this method:

Theorem 3. *For all $1 < t < n$ and $0 < \theta \leq 1/t$, Alice wins in the game $G(\mathsf{THR}_t^n, \theta)$.*

We also show that, for any total access structure $\mathcal{A}$, this method cannot prove a lower bound exceeding $\log|\min \mathcal{A}| \leq \log\binom{n}{\lfloor n/2 \rfloor} = n - \Omega(\log n)$, where $\min \mathcal{A} = \{A \in \mathcal{A} : \forall B \in \mathcal{A}, B \not\subset A\}$ is the set of min-terms in $\mathcal{A}$.

Theorem 4. *For every access structure $\mathcal{A}$ and every $0 < \theta \leq 1/(|\min \mathcal{A}| - 1)$ Alice wins in the game $G(\mathcal{A}, \theta)$.*

1.2 Related work

Known frameworks for proving lower bounds. The method of Csirmaz [Csi97] is one of the only previously known general frameworks for proving lower bounds on share size in various access structures.[5] Csirmaz's framework is a linear programming relaxation whose variables are the entropies of the joint distributions of the shares, one for each subset of the parties. Using several Shannon information inequalities, Csirmaz was able to prove an $n/\log n$ lower bound on the entropy of shares (in a specific access structure) which, in turn, imply the same lower bound on share size (for a 1-bit secret).

We note that Csirmaz's framework does not give any non-trivial lower bounds on share size for sharing a 1-bit secret for the threshold access structure. Indeed, Csirmaz's method

[5] Some lower bounds were proven using other methods such as counting arguments and other tools from information theory.

gives a lower bound on the information ratio of an access structure,[6] namely on the ratio between the size of the shares and the size of the secret, and for threshold schemes this ratio is 1 (using Shamir's scheme for a long enough secret; see Claim 9). Kilian and Nisan's [KN90] proof is the only known argument for threshold schemes and it does not seem to be useful for any other access structure, including the (t, n)-threshold access structures with t being close to n.

Csirmaz [Csi97] showed that his framework cannot be used to show a super-linear lower bound on share size for any access structure. This claim was strengthened by Beimel and Orlov [BO11] who showed that certain additional "non-Shannon type" information inequalities cannot bypass the linear share size barrier (see [MPY16] for a follow-up).

Linear schemes. A secret sharing scheme is *linear* if the reconstruction procedure is a linear function of the shares (over some abelian group). Most previously known schemes are linear (see [BI01, BIVW16, KNY16] for exceptions) and super-polynomial lower bounds for linear schemes were given in [BGW99, BGP97, Gál01] via its equivalence to monotone span programs [KW93]. In a very recent work, Cook et al. [CPRR16] gave the first exponential lower bound for linear secret sharing schemes by giving an exponential lower bound for monotone span programs.

For *linear* $(2, n)$-threshold secret sharing schemes for a 1-bit secret, a $\log n$ lower bound on share size was proven by Karchmer and Wigderson [KW93]. This was generalized by Cramer et al. [CFS05] (via a duality argument) to get a lower bound as in Equaiton (1). For *linear* (s, r, n)-ramp secret sharing schemes, Cramer et al. obtained a lower bound as in Equation (2). We emphasize that our lower bounds match the lower bounds of [CFS05] but are not restricted to linear (ramp) secret sharing schemes.

2 Access Structures and Secret Sharing

Let $\mathcal{P} \triangleq \{1, \ldots, n\}$ be a set of n parties. A collection of subsets $\mathcal{A} \subseteq 2^{\mathcal{P}}$ is monotone (upward-closed) if for every $B \in \mathcal{A}$ and $B \subseteq C$ it holds that $C \in \mathcal{A}$. The collection is *anti-monotone* if for every $B \in \mathcal{A}$ and $C \subseteq B$ it holds that $C \in \mathcal{A}$.

Definition 5. *A (partial) access structure $\mathcal{A} = (\mathcal{S}, \mathcal{R})$ is a pair of non-empty disjoint collections of subsets $\mathcal{R}$ and $\mathcal{S}$ of $2^{\mathcal{P}}$ such that $\mathcal{R}$ is monotone and $\mathcal{S}$ is anti-monotone. Subsets in $\mathcal{R}$ are called qualified and subsets in $\mathcal{S}$ are called unqualified.*

The access structure is *total* if $\mathcal{R}$ and $\mathcal{S}$ form a partition of $2^{\mathcal{P}}$. If $\mathcal{A} = (\mathcal{S}, \mathcal{R})$ is total we write $R \in \mathcal{A}$ for $R \in \mathcal{R}$ and $S \notin \mathcal{A}$ for $S \in \mathcal{S}$. Our work is mostly about the following two types of access structures:

- The *threshold access structure* THR_t^n is a total access structure over n parties in which any t parties can reconstruct and secrecy is guaranteed against any subset of $t - 1$

[6] We thank a reviewer for pointing this out.

parties:

$$\mathcal{S} = \{S\colon |S| \leq t-1\} \qquad \mathcal{R} = \{R\colon |R| \geq t\}.$$

- More generally, in the *ramp access structure* $\mathsf{RAMP}^n_{s,r}$, any r parties can reconstruct and secrecy is guaranteed against any s parties:

$$\mathcal{S} = \{S\colon |S| \leq s\} \qquad \mathcal{R} = \{R\colon |R| \geq r\}.$$

A secret sharing scheme involves a dealer who has a secret, a set of n parties, and a partial access structure $\mathcal{A} = (\mathcal{S}, \mathcal{R})$. A secret sharing scheme for $\mathcal{A} = (\mathcal{S}, \mathcal{R})$ is a method by which the dealer distributes shares to the parties such that any subset in $\mathcal{R}$ can reconstruct the secret from its shares, while any subset in $\mathcal{S}$ cannot reveal any information on the secret. We restrict our definition to 1-bit secrets.

Definition 6 (Secret sharing)**.** *A secret sharing scheme of a 1-bit secret for a partial access structure $\mathcal{A} = (\mathcal{S}, \mathcal{R})$ over n parties over share alphabet Σ is a pair of probability distributions p_0 and p_1 over Σ^n with the following properties:*

Reconstruction: *For every $R \in \mathcal{R}$ the marginal distributions[7] of p_0 and p_1 on the set R are disjoint.*

Secrecy: *For every $S \in \mathcal{S}$ the marginal distributions of p_0 and p_1 on the set S are identical.*

An implementation of a secret sharing scheme consists of a sharing algorithm that samples the shares from the probability distribution p_0 or p_1 depending on the value of the secret and of a reconstruction algorithm that recovers the secret from the joint values of the shares of any qualified subsets of parties. The disjointness requirement ensures that recovery by qualified subsets of parties is possible with probability 1. The secrecy requirement ensures that unqualified subsets of parties can extract no information about the secret. Thus, our definition is equivalent to the ones given, for example, in [Bei96, Definition 3.6] and in [Bei11, Definitions 2 and 3].

An alternative formulation of secret sharing. Here is an equivalent formulation of secret sharing. For $x \in \mathbb{Z}_q^n$, we use $[x]$ to denote the set of non-zero entries of x, namely $[x] = \{i\colon x_i \neq 0\}$, and $[x]^{\complement}$ for the complementary set of zero entries. In this notation, $[x-y]$ is the set of coordinates that x and y differ on and $[x-y]^{\complement}$ is the set of coordinates that they agree on. A function $\phi_S\colon \mathbb{Z}_q^n \to \mathbb{C}$ is an *S-junta* if the value $\phi_S(x_1, \ldots, x_n)$ is determined by the inputs $x_i\colon i \in S$.

Lemma 7. *A secret sharing scheme of a 1-bit secret for a partial access structure $\mathcal{A} = (\mathcal{S}, \mathcal{R})$ over share alphabet $\mathbb{Z}_q$ exists if and only if there exists a function $f\colon \mathbb{Z}_q^n \to \mathbb{R}$ that is not identically zero satisfying the following properties:*

[7] Given two random variables X and Y whose joint distribution is known, the marginal distribution of X is the probability distribution of X averaging over all possible values of Y. Namely, it is $\mathbf{Pr}[X = x] = \sum_y \mathbf{Pr}[X = x, Y = y]$.

Reconstruction: *For all $x, y \in \mathbb{Z}_q^n$ such that $[x-y]^{\complement} \in \mathcal{R}$, $f(x) \cdot f(y) \geq 0$.*

Secrecy: *For every $S \in \mathcal{S}$ and every S-junta $\phi_S \colon \mathbb{Z}_q^n \to \mathbb{C}$, $\mathbf{E}[f(x)\phi_S(x)] = 0$, where the expectation is over the uniform probability distribution of $x \in \mathbb{Z}_q^n$.*

Proof. For a secret sharing scheme p_0, p_1, we set $f(x) = p_0(x) - p_1(x)$. The functions p_0 and p_1 have disjoint support (otherwise even reconstruction by all parties is impossible) so f cannot be identically zero. The reconstruction implies that if $[x-y]^{\complement} \in \mathcal{R}$, then at least one of p_0 and p_1 must assign zero probability to both x and y, so $f(x) \cdot f(y)$ equals either $p_0(x) \cdot p_0(y)$ or $(-p_1(x)) \cdot (-p_1(y))$. In either case $f(x) \cdot f(y) \geq 0$. For secrecy, since p_0 and p_1 have the same marginals on $S \in \mathcal{S}$, $\mathbf{E}[p_0(x)\phi_S(x)] = \mathbf{E}[p_1(x)\phi_S(x)]$ so $\mathbf{E}[f(x)\phi_S(x)] = 0$.

In the other direction, let $p_0(x) = C \cdot \max\{f(x), 0\}$ and let $p_1(x) = C \cdot \max\{-f(x), 0\}$ for a suitable scaling constant $C > 0$ that makes p_0 and p_1 be valid probability distributions (it exists since f is nonzero). We show reconstruction by contrapositive: If p_0 and p_1 did not have disjoint support on some set $R \in \mathcal{R}$, there would exist $x, y \in \mathbb{Z}_q^n$ such that $p_0(x) > 0$, $p_1(y) > 0$, and $[x-y]^{\complement} = R$, implying $f(x) > 0$, $f(y) < 0$, and therefore $f(x) \cdot f(y) < 0$. For secrecy, by construction we have $f = (p_0 - p_1)/C$, so $\mathbf{E}[p_0(x)\phi_S(x)] = \mathbf{E}[p_1(x)\phi_S(x)]$ for every test function ϕ_S that only depends on coordinates in $S \in \mathcal{S}$. Since no ϕ_S can distinguish between p_0 and p_1 on S, the statistical distance between the marginal distribution of p_0 and p_1 on S is zero, so the two are identical. □

3 A Zero-Sum Game and Proof of Theorem 2

Given a partial access structure $\mathcal{A} = (\mathcal{S}, \mathcal{R})$ and a real parameter $\theta > 0$ we define the following zero-sum game $G(\mathcal{A}, \theta)$ between Alice and Bob. The actions are a set $A \notin \mathcal{S}$ for Alice and a set $B \in \mathcal{R}$ for Bob. The payoff of the game is $(-\theta)^{|A \setminus B|}$. We say Alice wins if she has a strategy with non-negative expected payoff and we say Bob wins if he has a strategy with negative expected payoff (the expectations are over the randomness of Alice and Bob, respecively). By von Neumann's minimax theorem the game has a unique winner.

Lemma 8. *If there exists a secret sharing scheme for $\mathcal{A}$ with alphabet size $q \in \mathbb{N}$, then Alice wins in the game $G(\mathcal{A}, 1/(q-1))$.*

Our proof of Lemma 8 uses Fourier analysis, which we briefly recall here. The characters of the group $\mathbb{Z}_q^n$ are the complex-valued functions $\chi_a \colon \mathbb{Z}_q^n \to \mathbb{C}$, where a ranges over $\mathbb{Z}_q^n$, defined as $\chi_a(x) = \omega^{\langle a, x \rangle}$, $\omega = e^{2\pi i/q}$. The characters are an orthonormal basis with respect to the inner product $\langle f, g \rangle = \mathbf{E}_x[f(x) \cdot \overline{g(x)}]$ with x chosen uniformly from $\mathbb{Z}_q^n$. The characters inherit the group structure: $\chi_a \cdot \chi_b = \chi_{a+b}$ and $\chi_a^{-1} = \overline{\chi_a} = \chi_{-a}$. Every function $f \colon \mathbb{Z}_q^n \to \mathbb{C}$ can then be uniquely written as a linear combination $f = \sum_{a \in \mathbb{Z}_q^n} \hat{f}(a) \cdot \chi_a$ with the Fourier coefficients $\hat{f}(a)$ given by $\hat{f}(a) = \langle f, \chi_a \rangle = \mathbf{E}_x[f(x) \cdot \overline{\chi_a(x)}]$.

Proof of Lemma 8. We show that Alice has a winning strategy. That is, we show that Alice has a strategy such that for every possible action of Bob, the expected payoff of the game is non-negative.

We identify the alphabet with the elements of the group $\mathbb{Z}_q$. Let $f\colon \mathbb{Z}_q^n \to \mathbb{R}$ be the function $f(x) = p_0(x) - p_1(x)$. Alice plays set A with probability proportional to $\sum_{a\colon [a]=A} |\hat{f}(a)|^2$. By the secrecy part of Lemma 7, $\mathbf{E}[f(x) \cdot \overline{\chi_a(x)}] = 0$ whenever $[a] \in \mathcal{S}$, so Alice's strategy is indeed supported on sets outside $\mathcal{S}$.

Now let B be an arbitrary set in $\mathcal{R}$. By the reconstruction part of Lemma 7 and the fact that f is real-valued, for every $x \in \mathbb{Z}_n^q$ and every $z \in \mathbb{Z}_n^q$ such that $[z]^{\complement} = B$, we have that

$$f(x) \cdot \overline{f(x-z)} = f(x) \cdot f(x-z) \geq 0. \tag{3}$$

Let x be uniform in $\mathbb{Z}_q^n$ and z be uniform in $\mathbb{Z}_q^n$ conditioned on $[z]^{\complement} = B$. Averaging over this distribution, we have

$$\begin{aligned}
\mathbf{E}_{x,z}[f(x) \cdot \overline{f(x-z)}] &= \sum_{a,b \in \mathbb{Z}_q^n} \hat{f}(a) \cdot \overline{\hat{f}(b)} \cdot \mathbf{E}_{x,z}[\chi_a(x) \cdot \overline{\chi_b(x-z)}] \\
&= \sum_a |\hat{f}(a)|^2 \cdot \mathbf{E}_z[\chi_a(z)] \\
&= \sum_a |\hat{f}(a)|^2 \cdot \prod_{i \in [a]} \mathbf{E}_z[\omega^{a_i z_i}],
\end{aligned}$$

where the first equality follows by writing $f(x)$ and $\overline{f(x-z)}$ using their Fourier representation and using linearity of expectation, the second equality follows since x and z are independent and since $\mathbf{E}_x[\chi_a(x) \cdot \overline{\chi_b(x)}] = 0$ for $a \neq b$, and the last equality follows since z is chosen from a product distribution.

The expression $\mathbf{E}[\omega^{a_i z_i}]$ evaluates to one when i is in B (since z_i is fixed to zero). Otherwise, z_i is uniformly distributed over the set $\mathbb{Z}_q \setminus \{0\}$ and

$$\mathbf{E}_z[\omega^{a_i z_i}] = \frac{1}{q-1} \sum_{z_i \in \mathbb{Z}_q \setminus \{0\}} \omega^{a_i z_i} = \frac{1}{q-1} \Big(\sum\nolimits_{z_i \in \mathbb{Z}_q} \omega^{a_i z_i} - 1 \Big) = -\frac{1}{q-1}.$$

Therefore, $\prod_{i \in [a]} \mathbf{E}_z[\omega^{a_i z_i}] = (-1/(q-1))^{|[a] \setminus B|}$, and by Equation (3)

$$\sum_a |\hat{f}(a)|^2 \cdot \left(\frac{-1}{q-1} \right)^{|[a] \setminus B|} \geq 0.$$

Grouping all a's for which $[a] = A$, we get that

$$\sum_A \Big(\sum\nolimits_{a\colon [a]=A} |\hat{f}(a)|^2 \Big) \cdot \Big(-\frac{1}{q-1} \Big)^{|A \setminus B|} \geq 0 \qquad \text{for all } B \in \mathcal{R}.$$

Therefore, Alice's strategy has non-negative expected payoff with respect to every possible action of Bob. ∎

Proof of Theorem 2. It is sufficient to prove Theorem 2 in the case $n = r + 1$: If a secret sharing scheme for $\mathsf{RAMP}^n_{s,r}$ existed, then a secret sharing for $\mathsf{RAMP}^{r+1}_{s,r}$ over the same alphabet can be obtained by discarding the remaining $n - r - 1$ parties and their shares.

We now give a winning strategy for Bob in the game $G(\mathsf{RAMP}^{r+1}_{s,r}, \theta)$ for any $\theta > (r - s)/(s + 1)$. By Lemma 8 it then follows that no secret sharing scheme over an alphabet of size $(r + 1)/(r - s)$ exists.

Bob's strategy is to uniformly choose a set B of size r (which is in $\mathcal{R}$). Then for every set $A \notin \mathcal{S}$, either $A \subseteq B$ and then $|A \setminus B| = 0$, or $A \not\subseteq B$ and then $|A \setminus B| = 1$ (since B includes all parties except one). Thus, for every $A \notin \mathcal{S}$, the expected payoff is

$$\begin{aligned}\mathbf{E}_B\left[(-\theta)^{|A\setminus B|}\right] &= 1 \cdot \mathbf{Pr}_B[A \subseteq B] - \theta \cdot \mathbf{Pr}_B[A \not\subseteq B] \\ &= 1 \cdot \frac{r + 1 - |A|}{r + 1} - \theta \cdot \frac{|A|}{r + 1} \\ &\leq \frac{r - s}{r + 1} - \theta \cdot \frac{s + 1}{r + 1}, \end{aligned} \tag{4}$$

where the inequality follows since $|A| \geq s + 1$. If $\theta > (r - s)/(s + 1)$ this expression is less than zero, i.e., Bob wins. ∎

It is also possible to deduce Theorem 2 directly from Lemma 8 by showing the existence of a winning strategy for Bob in the game $G(\mathsf{RAMP}^n_{s,r}, \theta)$ whenever $\theta > (r-s)/(s+1)$ (rather than for $G(\mathsf{RAMP}^{r+1}_{s,r}, \theta)$, as we did above). Let R be a random subset of $r+1$ parties. Bob's strategy has the form $B = B_0 \cup B_1$, where B_0 is a uniformly random subset of R of size r and B_1 is a random subset of $R^{\complement}$ obtained by including each element independently with probability $p = \theta/(1 + \theta)$. The value of p is chosen so that a random variable that equals 1 with probability p and $-\theta$ with probability $1 - p$ is unbiased.

Let A, where $|A| \geq s + 1$, be any action of Alice. For a fixed choice of R, if $A \setminus R$ is nonempty, by our choice of probability p the expected payoff is zero. Otherwise, A is a subset of R, and by Equation (4) the expected payoff is at most $-(s + 1) \cdot \theta + (r - s) < 0$. Since the event $A \subseteq R$ has positive probability the expected payoff is negative and Bob wins.

4 Limitations of the Game Relaxation

In the case of threshold access structures Theorem 2 shows that Bob has a winning strategy in the game $G(\mathsf{THR}^n_t, \theta)$ whenever $\theta > 1/t$. We now prove Theorem 3, which states that our analysis is optimal: There exists a winning strategy for Alice when $\theta \leq 1/t$.

We also prove Theorem 4: For every total access structure $\mathcal{A}$ over n parties, Alice has a winning strategy in $G(\mathcal{A}, \theta)$ for every $\theta \leq 1/(|\mathcal{A}| - 1)$. As the proof of Theorem 4 is simpler we present that one first. We remark Theorem 4 can be generalized to any partial access structure $(\mathcal{S}, \mathcal{R})$ by replacing $\mathcal{A}$ by $\mathcal{R}$ in the proof.

Proof of Theorem 4. Alice's strategy is uniformly random over all minterms $A \in \min \mathcal{A}$.

Then, for every $B \in \mathcal{A}$ and $\theta < 1$, it holds that

$$\begin{aligned}\mathbf{E}_A[(-\theta)^{|A\setminus B|}] =& \mathbf{E}_A[(-\theta)^{|A\setminus B|} \mid A \subseteq B] \cdot \mathbf{Pr}_A[A \subseteq B] + \\ & \mathbf{E}_A[(-\theta)^{|A\setminus B|} \mid A \not\subseteq B] \cdot \mathbf{Pr}_A[A \not\subseteq B] \\ \geq& 1 \cdot \mathbf{Pr}_A[A \subseteq B] - \theta \cdot \mathbf{Pr}_A[A \not\subseteq B] \\ =& (1+\theta) \cdot \mathbf{Pr}_A[A \subseteq B] - \theta \\ \geq& (1+\theta) \cdot \frac{1}{|\min \mathcal{A}|} - \theta.\end{aligned}$$

This is non-negative when $\theta \leq 1/(|\min \mathcal{A}| - 1)$. ∎

Proof of Theorem 3. Let $a_0, \ldots, a_n$ be the following sequence of integers:

$$a_0 = \cdots = a_{t-1} = 0, \quad a_t = 1, \quad a_s = k_t \cdot a_{s-1} + \cdots + k_0 \cdot a_{s-t-1}$$

for $t+1 \leq s \leq n$, where k_j is the coefficient of x^j in the formal expansion of $(x+1)^t \cdot (1/\theta - x)$. By expanding this expression according to the Binomial formula, we see that the numbers $k_0, \ldots, k_t$ are non-negative when $\theta \leq 1/t$ because

$$k_j = \binom{t}{j}\left(\frac{1}{\theta} - \frac{j}{t-j+1}\right) \geq 0$$

for all $0 \leq j \leq t$. Therefore a_s is also non-negative for all s.

Alice plays set A with probability proportional to the number $a_{|A|}$. We will prove that this is a winning strategy for Alice. When $B = \{1, \ldots, n\}$, then $\mathbf{E}_A[(-\theta)^{|A\setminus B|}] = 1$ and Alice wins. Now let $B \subseteq \{1, \ldots, n\}$ be any set such that $t \leq |B| < n$. Let

$$\theta_j = \begin{cases} 1, & \text{if } j \in B, \\ -\theta, & \text{if } j \notin B. \end{cases}$$

Then,

$$\mathbf{E}_A[(-\theta)^{|A\setminus B|}] \propto \sum_A a_{|A|} \prod_{j \in A} \theta_j = \sum_{s=0}^{n} a_s w_s \quad \text{where} \quad w_s = \sum_{A:\ |A|=s} \prod\nolimits_{j \in A} \theta_j.$$

The number w_s can be represented as the coefficient of z^s in the formal expansion of $g_0(z) = \prod_{j=1}^{n}(1+\theta_j z)$. Since exactly $|B|$ of the θ_j's equal 1 and the other $n - |B|$ equal $-\theta$, it follows that

$$g_0(z) = (1+z)^{|B|} \cdot (1-\theta z)^{n-|B|}. \tag{5}$$

The numbers $a_0, \ldots, a_n$ (as defined in the beginning of the proof) are defined by an order t homogeneous linear degree relation with constant coefficients whose characteristic equation

is $(x+1)^t \cdot (1/\theta - x) = 0$. This equation has roots -1 (with multiplicity t) and $1/\theta$ (with multiplicity 1). Therefore,

$$a_s = C \cdot \theta^{-s} + \sum_{i=0}^{t-1} c_i \cdot s^i \cdot (-1)^s$$

where $c_0, \ldots, c_{t-1}$ and C are constants determined by the initial conditions on $a_0, \ldots, a_t$. We can now write

$$\sum_{s=0}^{n} a_s \cdot w_s = C \cdot \sum_{s=0}^{n} w_s \cdot \theta^{-s} + \sum_{i=0}^{t-1} c_i \cdot \sum_{s=0}^{n} w_s \cdot s^i \cdot (-1)^s.$$

Recall that g_0 is the generating function of w_s which means that $g_0(z) = \sum_{s=0}^{n} w_s \cdot z^s$. So, the term $\sum_{s=0}^{n} w_s \cdot \theta^{-s}$ equals $g_0(1/\theta) = 0$. To finish the proof, we show that $\sum_{s=0}^{n} w_s \cdot s^i \cdot (-1)^s = 0$ for all $i \leq t-1$ (this implies that Alice's strategy has a 0 payoff, which means that she wins the game). Let $g_i(z) = z \cdot g'_{i-1}(z)$ for $1 \leq i \leq t-1$ where g'_{i-1} is the derivative of g_{i-1}. On the one hand, since -1 is a root of g_0 of multiplicity t, $g_i(-1) = 0$ for all $i \leq t-1$. On the other hand, $g_i(z)$ has the formal expansion $\sum_{s=0}^{n} w_s \cdot s^i \cdot z^s$. Therefore, $\sum_{s=0}^{n} w_s \cdot s^i \cdot (-1)^s$ must equal zero. ∎

5 Concluding Remarks

Theorem 1 requires that the shares given to all parties have the same length. Its proof extends easily to yield the following generalization: For every n, every $1 < t < n$, and every (t, n)-threshold secret sharing scheme in which party i receives a $\log q_i$-bit share and $q_1 \leq q_2 \leq \cdots \leq q_n$ it must hold that

$$\frac{1}{q_1} + \cdots + \frac{1}{q_{t+1}} \leq 1. \tag{6}$$

In particular, inequality (6) implies that the *average* share size must be at least $\log(t+1)$. We sketch the Proof in Appendix B. Kilian and Nisan [KN90] prove the same for $(n-t+1, n)$-threshold access structures.

By Theorem 3 our analysis of threshold secret sharing is tight within the game-theoretic relaxation that we introduce here. As the lower bound of Kilian and Nisan [KN90] is incomparable with ours, their analysis cannot be cast in terms of a winning strategy in our game. It is, however, possible to capture both our analysis and that of Kilian and Nisan by a single *linear program*. We performed computer experiments to investigate the feasibility of one such family of linear programs, but were unable to obtain better lower bounds on share size.

We do not know what is the best possible lower bound on share size that our method can give among all access structures on n parties. Theorem 1 shows a lower bound of $\log(n-1)$

is attainable, while Theorem 4 shows that a lower bound of $\log \binom{n}{\lfloor n/2 \rfloor}$ cannot be proved. The best possible bound is the logarithm of

$$b_n = \min_{\mathcal{A}} \max\{q\text{: Bob wins in } G(\mathcal{A}, 1/(q-1))\},$$

where the minimum is taken over all access structures $\mathcal{A}$ on n parties. We can prove that if the payoff function is replaced by $(-\theta)^{|A \triangle B|}$, where $\triangle$ is symmetric set difference, then the quantity analogous to b_n is upper bounded by $O(n^2)$.

Acknowledgments.

We thank Moni Naor for telling us about the work of Kilian and Nisan. We thank the anonymous reviewers for their useful advice.

A On the Tightness of Theorem 2

We show that Theorem 2 is tight when $t = n - 1$ and n is the power of a prime. This result is known (see e.g. [CDN15, Theorem 11.13]) and we give it here for completeness.

Claim 9. *For every power of a prime n there exists a $(n-1)$-out-of-n secret sharing scheme for $\log n$-bit secrets with $\log n$-bit shares.*

Claim 9 follows by a small optimization of Shamir's secret sharing scheme. We give the construction and sketch the correctness proof.

To share a secret $s \in \mathbb{F}_n$, let $p(x) = sx^{n-2} + r(x)$, where r is a random polynomial of degree $n-3$ and all algebra is over the finite field $\mathbb{F}_n$. The shares are the n values $p(x)$ as x ranges over $\mathbb{F}_q$. Reconstruction is immediate as the polynomial p can be interpolated from any $n-1$ of its values.

For secrecy, we show for any $s \in \mathbb{F}_n$ and distinct elements $x_1, \ldots, x_{n-2} \in \mathbb{F}_n$, the vector $(p(x_1), \ldots, p(x_{n-2}))$ is uniformly random in $\mathbb{F}_n^{n-2}$. Since $p(x) = sx^{n-2} + r(x)$ it suffices to show that $(r(x_1), \ldots, r(x_{n-2}))$ is uniformly random. This is true because the evaluation map that takes the coefficients of r into its values $r(x_1), \ldots, r(x_{n-2})$ is a full-rank Vandermonde matrix.

B Proof Sketch of Inequality (6)

The proof of inequality (6) is a direct extension of the proof of Theorem 1. We describe the differences. The payoff function in the game G in Lemma 8 becomes $\prod_{i \in A \setminus B} -1/(q_i - 1)$. The generalized lemma can be proved via Fourier analysis over the product group $\mathbb{Z}_{q_1} \times \cdots \times \mathbb{Z}_{q_n}$.

As in the proof of Theorem 1, it is sufficient to establish inequality (6) in the special case $t = n - 1$. Bob then plays set $B = \{1, \ldots, n\} \setminus \{i\}$ with probability proportional to $1 - 1/q_i$. It can be verified that when $\sum_{i=1}^{n} 1/q_i > 1$ this is a winning strategy for Bob.

References

[BC94] Amos Beimel and Benny Chor. Universally ideal secret-sharing schemes. *IEEE Transactions on Information Theory*, 40(3):786–794, 1994.

[Bei96] Amos Beimel. *Secure Schemes for Secret Sharing and Key Distribution*. PhD thesis, Technion - Israel Institute of Technology, 1996.

[Bei11] Amos Beimel. Secret-sharing schemes: A survey. In *Coding and Cryptology - 3rd International Workshop, IWCC*, volume 6639, pages 11–46, 2011.

[BF07] Amos Beimel and Matthew K. Franklin. Weakly-private secret sharing schemes. In *4th Theory of Cryptography Conference, TCC*, volume 4392, pages 253–272, 2007.

[BGP97] Amos Beimel, Anna Gál, and Mike Paterson. Lower bounds for monotone span programs. *Computational Complexity*, 6(1):29–45, 1997.

[BGW99] László Babai, Anna Gál, and Avi Wigderson. Superpolynomial lower bounds for monotone span programs. *Combinatorica*, 19(3):301–319, 1999.

[BI01] Amos Beimel and Yuval Ishai. On the power of nonlinear secrect-sharing. In *16th Annual IEEE Conference on Computational Complexity, CCC*, pages 188–202, 2001.

[BIVW16] Andrej Bogdanov, Yuval Ishai, Emanuele Viola, and Christopher Williamson. Bounded indistinguishability and the complexity of recovering secrets. In *Advances in Cryptology - CRYPTO*, pages 593–618, 2016.

[BL88] Josh Cohen Benaloh and Jerry Leichter. Generalized secret sharing and monotone functions. In *8th Annual International Cryptology Conference, CRYPTO*, pages 27–35, 1988.

[Bla79] George R. Blakley. Safeguarding cryptographic keys. *Proceedings of the AFIPS National Computer Conference*, 22:313–317, 1979.

[BM84] G. R. Blakley and Catherine A. Meadows. Security of ramp schemes. In *Advances in Cryptology - CRYPTO*, pages 242–268, 1984.

[BO11] Amos Beimel and Ilan Orlov. Secret sharing and non-shannon information inequalities. *IEEE Transactions on Information Theory*, 57(9):5634–5649, 2011.

[CC06] Hao Chen and Ronald Cramer. Algebraic geometric secret sharing schemes and secure multi-party computations over small fields. In *Advances in Cryptology - CRYPTO*, pages 521–536. Springer, 2006.

[CCX13] Ignacio Cascudo Pueyo, Ronald Cramer, and Chaoping Xing. Bounds on the threshold gap in secret sharing and its applications. *IEEE Transactions on Information Theory*, 59(9):5600–5612, 2013.

[CDN15] Ronald Cramer, Ivan Damgård, and Jesper Buus Nielsen. *Secure Multiparty Computation and Secret Sharing*. Cambridge University Press, 2015.

[CFS05] Ronald Cramer, Serge Fehr, and Martijn Stam. Black-box secret sharing from primitive sets in algebraic number fields. In *Advances in Cryptology - CRYPTO*, pages 344–360, 2005.

[CPRR16] Stephen A. Cook, Toniann Pitassi, Robert Robere, and Benjamin Rossman. Exponential lower bounds for monotone span programs. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:64, 2016.

[CSGV93] Renato M. Capocelli, Alfredo De Santis, Luisa Gargano, and Ugo Vaccaro. On the size of shares for secret sharing schemes. *Journal of Cryptology*, 6(3):157–167, 1993.

[Csi97] László Csirmaz. The size of a share must be large. *Journal of Cryptology*, 10(4):223–231, 1997.

[FHKP14] Oriol Farràs, Torben Brandt Hansen, Tarik Kaced, and Carles Padró. Optimal non-perfect uniform secret sharing schemes. In *Advances in Cryptology - CRYPTO*, pages 217–234, 2014.

[FMP16] Oriol Farràs, Sebastià Martín Molleví, and Carles Padró. A note on non-perfect secret sharing. *IACR Cryptology ePrint Archive*, page 348, 2016.

[Gál01] Anna Gál. A characterization of span program size and improved lower bounds for monotone span programs. *Computational Complexity*, 10(4):277–296, 2001.

[ISN93] Mitsuru Ito, Akira Saito, and Takao Nishizeki. Multiple assignment scheme for sharing secret. *Journal of Cryptology*, 6(1):15–20, 1993.

[KN90] Joe Kilian and Noam Nisan. Unpublished. Referenced in [BC94, Bei96, BF07, CCX13], 1990.

[KNY16] Ilan Komargodski, Moni Naor, and Eylon Yogev. How to share a secret, infinitely. *IACR Cryptology ePrint Archive*, 2016:194, 2016.

[KW93] Mauricio Karchmer and Avi Wigderson. On span programs. In *8th Annual Structure in Complexity Theory Conference*, pages 102–111, 1993.

[MPS11] Keith M. Martin, Maura B. Paterson, and Douglas R. Stinson. Error decodable secret sharing and one-round perfectly secure message transmission for general adversary structures. *Cryptography and Communications*, pages 65–86, 2011.

[MPY16] Sebastià Martín Molleví, Carles Padró, and An Yang. Secret sharing, rank inequalities, and information inequalities. *IEEE Trans. Information Theory*, 62(1):599–609, 2016.

[PS13] Maura B. Paterson and Douglas R. Stinson. A simple combinatorial treatment of constructions and threshold gaps of ramp schemes. *Cryptography and Communications*, pages 229–240, 2013.

[Sha79] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.

[SW99] Douglas R. Stinson and Ruizhong Wei. An application of ramp schemes to broadcast encryption. *Inf. Process. Lett.*, pages 131–135, 1999.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de