THEORY OF COMPUTING, Volume 16 (2), 2020, pp. 1–18
www.theoryofcomputing.org

Threshold Secret Sharing Requires a Linear-Size Alphabet

Andrej Bogdanov* Siyao Guo† Ilan Komargodski‡

Received November 26, 2016; Revised May 3, 2019; Published September 7, 2020

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 1979) 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.

ACM Classification: F.1.3

AMS Classification: 68Q17, 94A60

Key words and phrases: secret sharing, threshold, lower bound

---

A conference version of this paper appeared in the Proceedings of the The 14th Theory of Cryptography Conference (TCC), 2016-B.

*Supported by RGC GRF grants CUHK410113 and CUHK14208215.

†Part of the work done in the Chinese University of Hong Kong supported by RGC GRF grants CUHK410112 and CUHK410113.

‡Supported by an AFOSR grant FA9550-15-1-0262. Part of this work done while visiting CUHK, supported by RGC GRF grant CUHK410113, and while being a Ph.D. student at the Weizmann Institute of Science, 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.

 DOI: 10.4086/toc.2020.v016a002

1 Introduction

In 1979, Shamir [36] and Blakley [11] 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 cryptography and distributed computing; see the extensive survey of Beimel [3] and the recent book of Cramer et al. [17].

Threshold secret sharing was generalized by Ito et al. [26] 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 [36] 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 in 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 [28][1] showed that $\log n$ bits are necessary when t is not too large. Specifically, they proved 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 [12], are useful for various applications (see, e. g., [24, 37, 16, 32]). If $r-s$ is large, they can sometimes be realized with shorter shares than standard threshold schemes (especially in the case of a long secret).

[1] Their result is unpublished and independently obtained (and generalized in various ways) by [15]. The original argument of Kilian an Nisan appears in [15, Appendix A] and was referenced earlier in [4, 2, 5].

[2] Another common definition (see [22, Definition 2.7] and [23, 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.

Generalizing the lower bound of Kilian and Nisan, Cascudo et al. [15] showed that $\log((n-s+1)/(r-s))$-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 [34] 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 [26, 10, 27] 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 super-polynomial size shares.[3] On the other hand, the best known lower bound on share size for sharing an ℓ-bit secret is $\Omega(\ell \cdot n/\log n)$ bits, by Csirmaz [20] (improving on [14]).[4]

Bridging the exponential gap between the upper and the lower bounds is the major open problem in the study of secret sharing schemes. While it is widely believed that the lower bound should be exponential (see, e. g., [2, 3]), no major progress has been made in the last two decades. Moreover, even a non-explicit linear lower bound is not known, that is, whether there *exists* an access structure that requires linear size shares.

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 1.1. *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 Section 5. By combining Theorem 1.1 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

$$\max\{\log(n-t+2), \log(t+1)\} \geq \log\frac{n+3}{2}. \tag{1.1}$$

Theorem 1.1 is a special case of the following theorem, which applies more generally to ramp schemes.

Theorem 1.2. *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.*

[3] One 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] In a follow-up work, Csirmaz [19] proved a lower bound of $\Omega(\ell \cdot n^2/\log n)$ for the total share size of all parties.

By combining Theorem 1.2 with the lower bound of [15], we get that any (s,r,n)-ramp secret sharing scheme must have share size at least

$$\max\left\{\log\frac{n-s+1}{r-s},\log\frac{r+1}{r-s}\right\}\geq\log\frac{n+r-s+2}{2\cdot(r-s)}. \tag{1.2}$$

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 3.1) 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 1.2 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 1.1 is tight with respect to this method:

Theorem 1.3. *For all $1<t<n$ and $0<\theta\leq 1/t$, Alice wins in the game $G(\mathsf{THR}^n_t,\theta)$.*

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 1.4. *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 [20] is a general framework for proving lower bounds on share size in various access structures. 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) and whose constraints are entropy inequalities. A solution to the dual linear program yields an $\Omega(n/\log n)$ lower bound on the *information ratio*, namely the length of the largest share divided by the length of the secret.

Csirmaz's framework does not give any non-trivial lower bounds on share size for a 1-bit secret over threshold access structures. The reason is that for sufficiently long secrets, the information ratio of threshold schemes is 1 in Shamir's construction (see Claim 5.1). Kilian and Nisan's [28] proof is the only known argument for threshold schemes. It does not appear to be useful for any other access structure, including the (t,n)-threshold access structures with t being close to n.

Csirmaz [20] showed that his framework cannot be used to prove a super-linear lower bound on share size for any access structure when the constraints are the Shannon inequalities from information theory. This claim was strengthened by Beimel and Orlov [8], who showed that certain additional "non-Shannon type" information inequalities do not help bypass the linear share size barrier (see [33] for a follow-up).

Linear schemes. A secret sharing scheme is *linear* if the reconstruction procedure is a linear function of the shares (over a finite field). Many known schemes are linear (see [7, 9, 29, 38, 13, 31, 30] for exceptions) and super-polynomial lower bounds for linear schemes were given in [1, 6, 25, 35] via their equivalence to monotone span programs [27].

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 [27]. This was generalized by Cramer et al. [18] to a lower bound as in Equation (1.1).[5] For *linear* (s,r,n)-ramp secret sharing schemes, Cascudo et al. [15] obtained a lower bound as in Equation (1.2). We emphasize that our lower bounds match the lower bounds of [15] but are not restricted to linear (ramp) secret sharing schemes.

2 Access structures and secret sharing

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

Definition 2.1. An *access structure* is a pair $\mathcal{A} = (\mathcal{S},\mathcal{R})$ where

(i) $\mathcal{S}$ and $\mathcal{R}$ are non-empty disjoint subsets of $2^{\mathcal{P}}$,

(ii) $\mathcal{R}$ is upward-closed and $\mathcal{S}$ is downward-closed.

Subsets in $\mathcal{R}$ are called *qualified* and subsets in $\mathcal{S}$ are called *unqualified*.

Remark 2.2. The condition that $\mathcal{R} \neq \emptyset$ is equivalent to saying that $\mathcal{P} \in \mathcal{R}$, i. e., all parties together are qualified (to unlock the secret). The condition $\mathcal{S} \neq \emptyset$ is equivalent to saying that $\emptyset \in \mathcal{S}$, i. e., the secret is not public (the empty set is not qualified).

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 $S \in \mathcal{S}$.

In this paper we are primarily concerned with the following two types of access structures:

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

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

[5] Their proof is based on the duality of linear secret sharing schemes, by which the optimal share sizes of linear (s,r,n)-ramp schemes and linear $(n-r,n-s,n)$-ramp schemes are equal.

A secret sharing scheme involves a dealer who has a secret, a set of n parties, and an 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 infer any information on the secret. We restrict our definition to 1-bit secrets.

Definition 2.3 (Secret sharing)**.** A secret sharing scheme of a 1-bit secret for an access structure $\mathcal{A} = (\mathcal{S}, \mathcal{R})$ over n parties over share alphabet Σ is a pair (p_0, p_1) of probability distributions over Σ^n with the following properties:

Reconstruction: For every $R \in \mathcal{R}$ the marginal distributions[6] 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 [2, Definition 3.6] and in [3, 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 positions with non-zero entries, i. e., $[x] = \{i \colon x_i \neq 0\}$, and $[x]^\complement$ for the complement of this set, i. e., $[x]^\complement = \{i \colon x_i = 0\}$. In this notation, $[x-y]$ is the set of positions where x and y differ and $[x-y]^\complement$ is the set of positions where they agree.

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 2.4. *A secret sharing scheme for a 1-bit secret for an 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 conditions.*

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

Secrecy: *For every $S \in \mathcal{S}$ and every S-junta $\phi_S \colon \mathbb{Z}_q^n \to \mathbb{C}$, we have $\mathbf{E}[f(x)\phi_S(x)] = 0$, 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 $p_1(x) = C \cdot \max\{-f(x), 0\}$, where $1/C = (1/2)\sum_x |f(x)|$. We show reconstruction by contrapositive: If p_0 and p_1 did not have disjoint support

[6] The marginal distribution of p over S is the function $p_S \colon \mathbb{Z}_q^S \to \mathbb{R}$ given by $p_S(z) = \sum_{x \colon x_S = z} p(x)$.

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 $R \subseteq [x-y]^{\complement} \in \mathcal{R}$ (by the monotonicity of $\mathcal{R}$), 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 the positions 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. $\square$

By Lemma 2.4, ruling out the existence of a secret sharing scheme is equivalent to refuting the solvability of a real-valued constraint satisfaction problem with quadratic constraints. The class of refutations that we investigate are sum-of-squares inequalities of the form $\sum_i (\sum_x \rho_i(x) f(x))^2 < 0$, i. e., dual solutions to the canonical semidefinite relaxation of this problem.

Owing to the symmetries inherent in secret sharing, the sum-of-squares certificate of interest is a positive *linear* combination of the squared magnitudes of the Fourier coefficients of f. Our lower bound can be expressed as a dual certificate for a linear program with variables $|\hat{f}(a)|^2, a \in \mathbb{Z}_q^n$. This linear program has the form of a zero-sum game, which we describe next.

3 A zero-sum game and proof of Theorem 1.2

Given an 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. Each party chooses its action independently of the other. The payoff of the game is $(-\theta)^{|A \setminus B|}$. We say that Alice wins if she has a strategy with non-negative expected payoff (over her randomness) against all strategies of Bob. Analogously, we say that Bob wins if he has a strategy with negative expected payoff. By von Neumann's minimax theorem the game has a unique winner.

Lemma 3.1. *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 3.1 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 $\chi_a(x) = \omega^{\langle a,x \rangle}$, $\omega = \mathrm{e}^{2\pi i/q}$. 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 3.1. 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 2.4, $\mathbf{E}[f(x)\cdot\overline{\chi_a(x)}]=0$ 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 2.4 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

$$f(x)\cdot\overline{f(x-z)}=f(x)\cdot f(x-z)\geq 0. \tag{3.1}$$

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

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

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 $\mathbf{E}[\omega^{a_iz_i}]$ 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

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

Therefore, $\prod_{i\in[a]}\mathbf{E}_z[\omega^{a_iz_i}]=(-1/(q-1))^{|[a]\setminus B|}$, and by Equation (3.1)

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

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

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

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

Proof of Theorem 1.2. It is sufficient to prove Theorem 1.2 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 3.1 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 $A \nsubseteq B$ and then $|A \setminus B| = 1$ (since B includes all parties except one). Thus, for every $A \notin \mathcal{S}$, the expected payoff is

$$\begin{aligned}\mathbf{E}_B\left[(-\theta)^{|A\setminus B|}\right] &= 1 \cdot \mathbf{Pr}_B[A \subseteq B] - \theta \cdot \mathbf{Pr}_B[A \nsubseteq B] \\ &= 1 \cdot \frac{r+1-|A|}{r+1} - \theta \cdot \frac{|A|}{r+1} \\ &\leq \frac{r-s}{r+1} - \theta \cdot \frac{s+1}{r+1}, \end{aligned} \tag{3.2}$$

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 1.2 directly from Lemma 3.1 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 (3.2) 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.

Extension to unequal shares. Theorem 1.1 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

$$\frac{1}{q_1} + \cdots + \frac{1}{q_{t+1}} \leq 1. \tag{3.3}$$

In particular, inequality (3.3) implies that the *average* share size must be at least $\log(t+1)$. Kilian and Nisan [28] prove the same for $(n-t+1,n)$-threshold access structures.

The proof of inequality (3.3) is a direct extension of the proof of Theorem 1.1. We describe the differences and main steps. Given an access structure $\mathcal{A} = (\mathcal{S}, \mathcal{R})$ and real positive numbers $\theta_1, \ldots, \theta_n$, we define the following zero-sum game $G(\mathcal{A}, \theta_1, \ldots, \theta_n)$ 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 $\prod_{i \in A\setminus B}(-\theta_i)$. We can obtain the following generalized version of Lemma 3.1.

Lemma 3.2. *If there exists a secret sharing scheme for $\mathcal{A}$ with alphabet size $q_i \in \mathbb{N}$ for the i-th party, then Alice wins in the game $G(\mathcal{A}, 1/(q_1-1), \ldots, 1/(q_n-1))$.*

The generalized lemma can be proved via Fourier analysis over the product group $\mathbb{Z}_{q_1} \times \cdots \times \mathbb{Z}_{q_n}$. We omit this proof.

As in the proof of Theorem 1.1, it is sufficient to establish inequality (3.3) in the special case $t = n-1$. In this case, Bob plays set $B = \{1,\dots,n\} \setminus \{i\}$ 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. Specifically, 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). Let $C = \sum_{i=1}^n (1-1/q_i)$. Thus, for every $A \notin \mathcal{S}$, the expected payoff is

$$\begin{aligned}
\mathbf{E}_B\Big[\prod_{i\in A\setminus B} -\theta_i\Big] &= \mathbf{Pr}_B[A \subseteq B]\cdot 1 - \sum_{i\in A} \mathbf{Pr}_B[A\setminus B = \{i\}]\cdot(-\theta_i)\\
&= \sum_{i\notin A} \frac{1-1/q_i}{C}\cdot 1 - \sum_{i\in A}\frac{1-1/q_i}{C}\cdot\theta_i\\
&= \sum_{i\notin A} \frac{1-1/q_i}{C} - \sum_{i\in A}\frac{1/q_i}{C}\\
&= 1-\frac{|A|}{C}.
\end{aligned}$$

By assumption $|A| \geq n-1$. So if $\sum_{i=1}^n 1/q_i > 1$ this expression is less than zero, i. e., Bob wins.

4 Limitations of the game relaxation

In the case of threshold access structures Theorem 1.2 shows that Bob has a winning strategy in the game $G(\mathsf{THR}_t^n, \theta)$ whenever $\theta > 1/t$. We now prove Theorem 1.3, which states that our analysis is optimal: There exists a winning strategy for Alice when $\theta \leq 1/t$.

We also prove Theorem 1.4: 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 1.4 is simpler, we present that one first. We remark that Theorem 1.4 can be generalized to any access structure $(\mathcal{S},\mathcal{R})$ by replacing $\mathcal{A}$ with $\mathcal{R}$ in the proof.

Proof of Theorem 1.4. 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

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

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

Proof of Theorem 1.3. Let $a_0,\dots,a_n$ 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, $\mathbf{E}_A[(-\theta)^{|A \setminus B|}]$ is proportional to

$$\sum_A a_{|A|} \prod_{j \in A} \theta_j = \sum_{s=0}^{n} a_s w_s \quad \text{where} \quad w_s = \sum_{A:\ |A|=s} \prod\nolimits_{j \in A} \theta_j.$$

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 equal 1 and the other $n-|B|$ equal $-\theta$, it follows that

$$g_0(z) = (1+z)^{|B|} \cdot (1-\theta z)^{n-|B|}. \tag{4.1}$$

The numbers $a_0, \ldots, a_n$ (as defined in the beginning of the proof) are defined by an order-t homogeneous linear recurrence 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

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

As g_0 is the generating function of the w_s, $g_0(z) = \sum_{s=0}^{n} w_s \cdot z^s$. 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, so 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 On the tightness of Theorem 1.2

We show that Theorem 1.2 is tight when $t = n-1$ and n is a prime power. This result is known (see, e. g., [17, Theorem 11.13]) and we give it here for completeness.

Claim 5.1. *For every prime power n there exists a $(n-1)$-out-of-n secret sharing scheme for $\lfloor \log n \rfloor$-bit secrets with $\lceil \log n \rceil$-bit shares.*

Claim 5.1 follows by an optimization of Shamir's secret sharing scheme, where we place the secret as the coefficient of the highest-degree term (instead of as the constant term). We give the construction and sketch the proof of correctness.

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 $\mathbb{F}_n$. 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 $x_1, \ldots, x_{n-2} \in \mathbb{F}_n$, 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 to its values $r(x_1), \ldots, r(x_{n-2})$ is a full-rank Vandermonde matrix.

6 Concluding remarks

Relation to the lower bound of Kilian and Nisan. By Theorem 1.3 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 [28] is incomparable with ours, their analysis cannot be cast in terms of a winning strategy in our game. The analysis of Kilian and Nisan relies on the collision probabilities $\mathbf{Pr}[X_A = Y_A]$, where X and Y are independent samples of the distributions p_0 and p_1, respectively, and X_A, Y_A are their projections on A. In the Fourier basis these probabilities can be written as

$$\mathbf{Pr}[X_A = Y_A] = \sum_{a\colon [a] \subseteq A} \hat{p}_0(a) \cdot \overline{\hat{p}_1(a)}.$$

Our game-theoretic formulation cannot express such quantities. It is, however, possible to unify our analysis with that of Kilian and Nisan by a common linear program. One approach is to refine the canonical linear programming formulation of the zero-sum game $G(\mathcal{A}, \theta)$, which has a variable $v(A)$ representing $\sum_{[a]=A} |\hat{f}(A)|^2$ for every subset A of $[n]$. In the refined formulation, each $v(A)$ is split into $v_{00}(A) + 2v_{01}(A) + v_{11}(A)$, where $v_{ij}(A)$ represents the (real-valued) expression $\sum_{[a]=A} \hat{p}_i(a) \cdot \overline{\hat{p}_j(a)}$. These variables satisfy $v_{ii}(A) \geq 0$ (non-negativity), $v_{00}(A) = v_{01}(A) = v_{11}(A)$ for $A \in \mathcal{S}$ (secrecy), and

$$\sum_A v_{ij}(A)(-\theta)^{|A \setminus B|} \begin{cases} = 0, & \text{if } B \in \mathcal{R} \text{ and } ij = 01, \\ \geq 0, & \text{otherwise.} \end{cases} \qquad \text{(reconstruction)}$$

The last quantity represents the probability of the event $\{b\colon X_b = Y_b\} = B$ up to scaling, where X and Y are independent samples of p_i and p_j, respectively.

This linear program captures both our proof technique and that of Kilian and Nisan and is therefore infeasible with respect to THR_t^n when $q < \max\{n-t+2, t+1\}$. Our computer experiments show that this bound is tight for all $n \le 7$, so this particular formulation does not appear to be advantageous for further improving the share size lower bounds on THR_t^n obtained in this paper.

Comparison with Csirmaz's method. A common feature of Csirmaz's method and ours is that the secrecy and reconstruction axioms are formulated as linear constraints on functionals of random variables determined by subsets of the shares. Csirmaz's relevant functional is the Shannon entropy of the corresponding subset. In contrast, under a suitable change of basis, our functionals capture the collision (Rényi) entropies of shares of 0 and/or 1 restricted to the relevant subset of the parties.

Csirmaz's method more generally gives a lower bound on the *information ratio* of schemes with arbitrary secret alphabet, while our formulation only applies to binary secrets. In this setting the two methods are incomparable. On the one hand, our main result shows that Csirmaz's method is less powerful than ours for threshold access structures. On the other hand, the alphabet size lower bound of 4 for the tree access structure $\{01, 12, 03, 34, 05\}$ obtained by Csirmaz and Tardos [21] cannot be matched by our game analysis.

On general access structures. 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 1.1 shows that a lower bound of $\log(n-1)$ is attainable, while Theorem 1.4 shows that a lower bound of $\log\binom{n}{\lfloor n/2 \rfloor}$ cannot be proved by our method.

The best possible bound is the logarithm of

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

where the minimum is taken over all access structures $\mathcal{A}$ on n parties. As indirect evidence that $\log\binom{n}{\lfloor n/2 \rfloor}$ might be far from tight, 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 at most $cn+1$, where $c = 1/\ln(1+\sqrt{2}) \approx 1.135$, when Bob plays tit-for-tat.

Claim 6.1. *If $A, B \subseteq [n]$ are independent and identically distributed then* $\mathbf{E}[(-1/cn)^{|A \triangle B|}] > 0$.

Proof. We first show that $\mathbf{Pr}[|A \triangle B| = s] \le \binom{n}{s} \mathbf{Pr}[A = B]$. Let p be the probability mass function of the sets.

$$\begin{aligned}
\mathbf{Pr}[|A \triangle B| = s] &= \sum_{A,B \colon |A \triangle B| = s} p(A)p(B) = \sum_{A,S \colon |S|=s} p(A)p(A \triangle S) \\
&\le \sqrt{\sum_{A,S \colon |S|=s} p(A)^2} \cdot \sqrt{\sum_{A,S \colon |S|=s} p(A \triangle S)^2} = \binom{n}{s} \sum_A p(A)^2,
\end{aligned}$$

where the inequality is Cauchy-Schwarz. The expectation of interest is then lower bounded by

$$\begin{aligned}\mathbf{E}\left[(-1/cn)^{|A\triangle B|}\right] &\geq \mathbf{Pr}[A=B] - \sum_{s \text{ odd}} \left(\frac{1}{cn}\right)^s \mathbf{Pr}[|A\triangle B| = s] \\ &\geq \mathbf{Pr}[A=B] - \sum_{s \text{ odd}} \left(\frac{1}{cn}\right)^s \binom{n}{s} \mathbf{Pr}[A=B] \\ &> \mathbf{Pr}[A=B] \cdot \left(1 - \frac{c^{-1}}{1!} - \frac{c^{-3}}{3!} - \frac{c^{-5}}{5!} - \cdots\right) \\ &= \mathbf{Pr}[A=B] \cdot \left(1 - \frac{\mathrm{e}^{1/c} - \mathrm{e}^{-1/c}}{2}\right),\end{aligned}$$

which is nonnegative by our choice of c. □

Acknowledgments.

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

References

[1] LÁSZLÓ BABAI, ANNA GÁL, AND AVI WIGDERSON: Superpolynomial lower bounds for monotone span programs. *Combinatorica*, 19(3):301–319, 1999. [doi:10.1007/s004930050058] 5

[2] AMOS BEIMEL: *Secure Schemes for Secret Sharing and Key Distribution*. Ph. D. thesis, Technion – Israel Institute of Technology, 1996. Accessible in PS on the author's home page and in PDF at DPHU.org. 2, 3, 6, 16

[3] AMOS BEIMEL: Secret-sharing schemes: A survey. In *Coding and Cryptology – 3rd Internat. Workshop, IWCC'11*, volume 6639, pp. 11–46. Springer, 2011. [doi:10.1007/978-3-642-20901-7_2] 2, 3, 6

[4] AMOS BEIMEL AND BENNY CHOR: Universally ideal secret-sharing schemes. *IEEE Trans. Inform. Theory*, 40(3):786–794, 1994. Preliminary version in CRYPTO'92. [doi:10.1109/18.335890] 2, 16

[5] AMOS BEIMEL AND MATTHEW K. FRANKLIN: Weakly-private secret sharing schemes. In *4th Theory of Cryptography Conf., TCC'07*, volume 4392, pp. 253–272, 2007. [doi:10.1007/978-3-540-70936-7_14] 2, 16

[6] AMOS BEIMEL, ANNA GÁL, AND MIKE PATERSON: Lower bounds for monotone span programs. *Comput. Complexity*, 6(1):29–45, 1996. Preliminary version in FOCS'95. [doi:10.1007/BF01202040] 5

[7] AMOS BEIMEL AND YUVAL ISHAI: On the power of nonlinear secrect-sharing. In *Proc. 16th IEEE Conf. on Computational Complexity (CCC'01)*, pp. 188–202. IEEE Comp. Soc., 2001. [doi:10.1109/CCC.2001.933886] 5

[8] AMOS BEIMEL AND ILAN ORLOV: Secret sharing and non-Shannon information inequalities. *IEEE Trans. Inform. Theory*, 57(9):5634–5649, 2011. Preliminary version in TCC'09. [doi:10.1109/TIT.2011.2162183] 4

[9] AMOS BEIMEL AND ENAV WEINREB: Separating the power of monotone span programs over different fields. *SIAM J. Comput.*, 34(5):1196–1215, 2005. Preliminary version in FOCS'03. [doi:10.1137/S0097539704444038] 5

[10] JOSH COHEN BENALOH AND JERRY LEICHTER: Generalized secret sharing and monotone functions. In *Proc. of 8th Internat. Cryptology Conf. (CRYPTO'88)*, pp. 27–35. Springer, 1988. [doi:10.1007/0-387-34799-2_3] 3

[11] GEORGE ROBERT BLAKLEY: Safeguarding cryptographic keys. In *Proc. AFIPS Nat. Computer Conf.*, volume 1, pp. 313–317. IEEE Comp. Soc., 1979. [doi:10.1109/AFIPS.1979.98] 2

[12] GEORGE ROBERT BLAKLEY AND CATHERINE A. MEADOWS: Security of ramp schemes. In *Proc. of 4th Internat. Cryptology Conf. (CRYPTO'84)*, pp. 242–268. Springer, 1984. [doi:10.1007/3-540-39568-7_20] 2

[13] ANDREJ BOGDANOV, YUVAL ISHAI, EMANUELE VIOLA, AND CHRISTOPHER WILLIAMSON: Bounded indistinguishability and the complexity of recovering secrets. In *Proc. of 36th Internat. Cryptology Conf. (CRYPTO'16)*, pp. 593–618. Springer, 2016. [doi:10.1007/978-3-662-53015-3_21] 5

[14] RENATO M. CAPOCELLI, ALFREDO DE SANTIS, LUISA GARGANO, AND UGO VACCARO: On the size of shares for secret sharing schemes. *J. Cryptology*, 6(3):157–167, 1993. Preliminary version in CRYPTO'91. [doi:10.1007/BF00198463] 3

[15] IGNACIO CASCUDO PUEYO, RONALD CRAMER, AND CHAOPING XING: Bounds on the threshold gap in secret sharing and its applications. *IEEE Trans. Inform. Theory*, 59(9):5600–5612, 2013. [doi:10.1109/TIT.2013.2264504] 2, 3, 4, 5, 16

[16] HAO CHEN AND RONALD CRAMER: Algebraic geometric secret sharing schemes and secure multi-party computations over small fields. In *Proc. of 26th Internat. Cryptology Conf. (CRYPTO'06)*, pp. 521–536. Springer, 2006. [doi:10.1007/11818175_31] 2

[17] RONALD CRAMER, IVAN BJERRE DAMGÅRD, AND JESPER BUUS NIELSEN: *Secure Multiparty Computation and Secret Sharing*. Cambridge Univ. Press, 2015. [doi:10.1017/CBO9781107337756] 2, 12

[18] RONALD CRAMER AND SERGE FEHR: Optimal black-box secret sharing over arbitrary Abelian groups. In *Proc. 22nd Internat. Cryptology Conf. (CRYPTO'02)*, pp. 272–287. Springer, 2002. [doi:10.1007/3-540-45708-9_18] 5

[19] LÁSZLÓ CSIRMAZ: The dealer's random bits in perfect secret sharing schemes. *Studia Scientiarum Mathematicarum Hungarica*, 32(3–4):429–438, 1996. Available at Rényi Institute. 3

[20] LÁSZLÓ CSIRMAZ: The size of a share must be large. *J. Cryptology*, 10(4):223–231, 1997. Preliminary version in EUROCRYPT'94. [doi:10.1007/s001459900029] 3, 4

[21] LÁSZLÓ CSIRMAZ AND GÁBOR TARDOS: Optimal information rate of secret sharing schemes on trees. *IEEE Trans. Inform. Theory*, 59(4):2527–2530, 2013. [doi:10.1109/TIT.2012.2236958] 13

[22] ORIOL FARRÀS, TORBEN BRANDT HANSEN, TARIK KACED, AND CARLES PADRÓ: Optimal non-perfect uniform secret sharing schemes. In *Proc. of 34th Internat. Cryptology Conf. (CRYPTO'14)*, pp. 217–234. Springer, 2014. [doi:10.1007/978-3-662-44381-1_13] 2

[23] ORIOL FARRÀS, SEBASTIÀ MARTÍN MOLLEVÍ, AND CARLES PADRÓ: A note on non-perfect secret sharing. IACR Cryptology ePrint Archive, Report 2016/348, 2016. 2

[24] MATTHEW K. FRANKLIN AND MOTI YUNG: Communication complexity of secure computation (extended abstract). In *Proc. 24th STOC*, pp. 699–710. ACM Press, 1992. [doi:10.1145/129712.129780] 2

[25] ANNA GÁL: A characterization of span program size and improved lower bounds for monotone span programs. *Comput. Complexity*, 10(4):277–296, 2001. Preliminary version in STOC'98. [doi:10.1007/s000370100001] 5

[26] MITSURU ITO, AKIRA SAITO, AND TAKAO NISHIZEKI: Multiple assignment scheme for sharing secret. *J. Cryptology*, 6(1):15–20, 1993. [doi:10.1007/BF02620229] 2, 3

[27] MAURICIO KARCHMER AND AVI WIGDERSON: On span programs. In *Proc. 8th Structure in Complexity Theory Conf. (SCT'93)*, pp. 102–111. IEEE Comp. Soc., 1993. [doi:10.1109/SCT.1993.336536] 3, 5

[28] JOE KILIAN AND NOAM NISAN: Unpublished. Referenced in [4, 2, 5, 15], 1990. 2, 4, 9, 12

[29] ILAN KOMARGODSKI, MONI NAOR, AND EYLON YOGEV: How to share a secret, infinitely. *IEEE Trans. Inform. Theory*, 64(6):4179–4190, 2018. Preliminary version in TCC'16. [doi:10.1109/TIT.2017.2779121] 5

[30] TIANREN LIU AND VINOD VAIKUNTANATHAN: Breaking the circuit-size barrier in secret sharing. In *Proc. 50th STOC*, pp. 699–708. ACM Press, 2018. [doi:10.1145/3188745.3188936] 5

[31] TIANREN LIU, VINOD VAIKUNTANATHAN, AND HOETECK WEE: Conditional disclosure of secrets via non-linear reconstruction. In *Proc. of 37th Internat. Cryptology Conf. (CRYPTO'17)*, pp. 758–790, 2017. [doi:10.1007/978-3-319-63688-7_25] 5

[32] 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*, 3(2):65–86, 2011. [doi:10.1007/s12095-010-0039-6] 2

[33] SEBASTIÀ MARTÍN MOLLEVÍ, CARLES PADRÓ, AND AN YANG: Secret sharing, rank inequalities, and information inequalities. *IEEE Trans. Inform. Theory*, 62(1):599–609, 2016. Preliminary version in CRYPTO'13. [doi:10.1109/TIT.2015.2500232] 4

[34] MAURA B. PATERSON AND DOUGLAS R. STINSON: A simple combinatorial treatment of constructions and threshold gaps of ramp schemes. *Cryptography and Communications*, 5(4):229–240, 2013. [doi:10.1007/s12095-013-0082-1] 3

[35] TONIANN PITASSI AND ROBERT ROBERE: Lifting Nullstellensatz to monotone span programs over any field. In *Proc. 50th STOC*, pp. 1207–1219. ACM Press, 2018. [doi:10.1145/3188745.3188914, ECCC:TR17-165] 5

[36] ADI SHAMIR: How to share a secret. *Comm. ACM*, 22(11):612–613, 1979. [doi:10.1145/359168.359176] 2

[37] DOUGLAS R. STINSON AND RUIZHONG WEI: An application of ramp schemes to broadcast encryption. *Info. Processing Lett.*, 69(3):131–135, 1999. [doi:10.1016/S0020-0190(98)00204-X] 2

[38] VINOD VAIKUNTANATHAN AND PRASHANT NALINI VASUDEVAN: Secret sharing and statistical zero knowledge. In *Proc. 21st Internat. Conf. on Theory and Appl. of Cryptology and Inform. Security (ASIACRYPT'15)*, pp. 656–680. Springer, 2015. [doi:10.1007/978-3-662-48797-6_27] 5

AUTHORS

Andrej Bogdanov
Associate professor
Department of Computer Science and Engineering
The Chinese University of Hong Kong
Hong Kong
andrejb@cse.cuhk.edu.hk
http://www.cse.cuhk.edu.hk/~andrejb/

Siyao Guo
Assistant professor
Department of Computer Science
NYU Shanghai
China
siyao.guo@nyu.edu
https://sites.google.com/site/siyaoguo/

Ilan Komargodski
Assistant professor and Scientist
Hebrew University of Jerusalem and NTT Research
Jerusalem, Israel and Palo Alto, CA, USA
ilan.komargodski@mail.huji.ac.il
http://ilan.tech.cornell.edu

ABOUT THE AUTHORS

ANDREJ BOGDANOV is an associate professor of Computer Science at the Chinese University of Hong Kong. He is on short leave from the Iyengar Yoga Institute of Hong Kong after a disastrous attempt at utthita pada sirsasana. He was advised by Luca Trevisan who taught him the fine art of enjoying xiaolongbao.

SIYAO GUO is an assistant professor of Computer Science at NYU Shanghai. She received her Ph. D. from the Chinese University of Hong Kong, advised by Andrej Bogdanov, who taught her how to eat xiaolongbao properly.

ILAN KOMARGODSKI is an assistant professor of Computer Science at the Hebrew University of Jerusalem and a Scientist at NTT Research. During his Ph. D. at the Weizmann Institute of Science, advised by Moni Naor, he visited Andrej Bogdanov and Siyao Guo who introduced to him xiaolongbao and other delicious treats like Fourier analysis.