Spectral Approach to the Communication Complexity of Multi-Party Key Agreement

Geoffroy Caillat-Grenier
LIRMM, University of Montpellier, CNRS, Montpellier, France

Andrei Romashchenko
LIRMM, University of Montpellier, CNRS, Montpellier, France

Abstract

We propose a linear algebraic method, rooted in the spectral properties of graphs, that can be used to prove lower bounds in communication complexity. Our proof technique effectively marries spectral bounds with information-theoretic inequalities. The key insight is the observation that, in specific settings, even when data sets X and Y are closely correlated and have high mutual information, the owner of X cannot convey a reasonably short message that maintains substantial mutual information with Y. In essence, from the perspective of the owner of Y, any sufficiently brief message $m = m(X)$ would appear nearly indistinguishable from a random bit sequence.

We employ this argument in several problems of communication complexity. Our main result concerns cryptographic protocols. We establish a lower bound for communication complexity of multi-party secret key agreement with unconditional, i.e., information-theoretic security. Specifically, for one-round protocols (simultaneous messages model) of secret key agreement with three participants we obtain an asymptotically tight lower bound. This bound implies optimality of the previously known *omniscience* communication protocol (this result applies to a non-interactive secret key agreement with three parties and input data sets with an arbitrary symmetric information profile).

We consider communication problems in one-shot scenarios when the parties inputs are not produced by any i.i.d. sources, and there are no ergodicity assumptions on the input data. In this setting, we found it natural to present our results using the framework of Kolmogorov complexity.

2012 ACM Subject Classification Mathematics of computing → Information theory; Theory of computation → Communication complexity; Security and privacy → Information-theoretic techniques; Theory of computation → Expander graphs and randomness extractors

Keywords and phrases communication complexity, Kolmogorov complexity, information-theoretic cryptography, multiparty secret key agreement, expander mixing lemma, information inequalities

Digital Object Identifier 10.4230/LIPIcs.STACS.2024.22

Related Version *Full Version*: https://arxiv.org/abs/2305.01355

Funding supported in part by the ANR project FLITTLA ANR-21-CE48-0023.

Acknowledgements We thank the anonymous referees for valuable and helpful comments.

1 Introduction

Within computer science, a broad range of communication complexity problems has been studied in recent decades. In these problems several (two or more) agents solve together some task (compute a function, search an elements in a set, sample a distribution, and so on) when the input data are distributed among the agents. In different context we may impose different constraints on the class of admissible protocols (protocols can be deterministic or randomized, one-way or interactive, with a one shot of simultaneous messages or with several rounds, etc.). The cost of a communication protocol is the total number of bits that must be exchanged between participants, typically in the worst-case situation.

41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024).
Editors: Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshtanov;
Article No. 22; pp. 22:1–22:19
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In this paper we focus on communication problems with three parties (Alice, Bob, and Charlie), though our techniques can be extended to bigger number of participants. We deal with the situation when the input data accessible to Alice, Bob, and Charlie are correlated. In a popular model *number-on-forehead*, the datasets given to Alice, Bob, and Charlie have large intersections, which is a very particular form of correlation between the data. We study a more general setting (more usual in cryptography and information theory) where the input data sets given to the parties have large mutual information, but it might be impossible to materialize this mutual information as common chunks of bits shared by several parties.

The principal communication problem under consideration is *secret key agreement*: Alice, Bob, and Charlie use the correlation between their input data sets to produce a common secret key. A special feature of this setting is the implicit presence of another participant in the game, Eve (eavesdropper/adversary). The eavesdropper can intercept all messages between Alice, Bob, and Charlie, but this should not give Eve any information about the final result of the protocol – the produced secret key. A secret key agreement (for two or many participants) is one of the basic primitives in cryptography; it can serve as a part of more sophisticated protocols (the produced secret key can be used in a one-time pad encryption or in more complicated cryptographic schemes).

In practice, the most standard and well known method of secret key agreement is the Diffie-Hellman key exchange [8, 22] and its generalizations, see [27]. The security of this protocol is based on assumptions of computational complexity. In particular, the Diffie-Hellman scheme is secure only if the eavesdropper cannot solve efficiently the problem of discrete logarithms. Such an assumption looks plausible for most practical applications. However, theoretical cryptography studies also secret key agreement in information-theoretic settings, where we impose no restrictions on the computational power of the eavesdropper. Besides a natural theoretical interest, such a scheme can be useful as a building block in more complex protocols. In particular, a protocol of information-theoretic secret key agreement (pretty conventional, involving communication and computational tools conceivable in the framework of the classical physics) is an indispensable component of the protocol of quantum key distribution ([4, 7, 16]). Besides quantum cryptography, secret-key agreement based on correlated information appears in various cryptographic schemes connected with noisy data (biometric information, observations of an inherently noisy communication channel or other physical phenomenon, see the discussions in [17, 10]), in the bounded-storage model ([9, 11]), and so on. We refer the reader to the survey [5] for a more detailed discussion.

In the Diffie-Hellman scheme, the parties may start the protocol from zero, holding initially no secret information. In contrast, a secret key agreement with information-theoretic secrecy is impossible if the parties start from scratch. To produce a key that is secret in information-theoretic sense, the participants of the protocol need to be given some input data (inaccessible to the eavesdropper). The pieces of input data provided to the parties must be correlated with each other, and the measure of this correlation determines the optimal size of the common secret key that can be produced.

So far we were very informal and did not specify the mathematical definitions behind the words *secrecy* (of the key) and *correlation* (between parties' inputs). Let us describe the settings of information-theoretic secret key agreement more precisely. This can be done in different mathematical frameworks.

Historically, information-theoretically secure protocols of secret key agreement were introduced in classical information theory, [1, 21]. In this setting, the input data of the parties are produced by correlated random variables. In the settings with two parties it is usually assumed that there is a sequence of i.i.d. pairs of random variables with finite range, (X_i, Y_i), $i = 1, \dots, n$, and Alice and Bob receive the values of $(X_1 \dots X_n)$ and $(Y_1 \dots Y_n)$ respectively. Then Alice and Bob run a communication protocol and try to produce a common value (secret key) W asymptotically independent of the *transcript* (the transcript consist of the messages sent by Alice and Bob to each other). Ahlswede–Csiszar [1] and Maurer [21] found a characterization of the optimal size of W in terms of Shannon's entropy of the input data. They showed that the optimal size of the secret key is asymptotically equal to the mutual information between Alice's and Bob's inputs. A similar characterization of the optimal secret key is known for multi-party protocols, with $k \geq 3$ parties, [6]. The problem of secret key agreement and a related problem of *common randomness generation* were extensively studied in the information theory community and also (in somewhat different settings) in theoretical computer science, see, e.g., [28, 13] and the survey [29].

In this paper we follow the paradigm of building the foundations of cryptography in the framework of algorithmic information theory, as suggested in a general form in [2] and more specifically for secret key agreement in [24, 15]. In this approach, the information-theoretic characteristics of the data are defined not in terms of Shannon's entropy but in terms of Kolmogorov complexity. In this setting, we can talk about properties of *individual* inputs, keys, transcripts, and not about *probability distributions*. We assume that the parties (Alice, Bob, Charlie) are given as inputs binary strings x, y, z respectively, and that the parties know the complexity profile of these strings, i.e., the optimal compression rate of these inputs (precisely or at least approximately, see below). The secrecy of the produced key means that this key must be incompressible, even conditional on the public data including the transcript of the communication protocol. In other words, the mutual information (in the sense of Kolmogorov complexity) between the key and the messages sent via the communication channel (the transcript) must be negligibly small. Practically, this property guarantees that the adversary can crack an encryption scheme based on this key only by the brute-force search, see the discussion in [15].

▶ Remark 1.1. The approach based on Kolmogorov complexity seems more general since we do not need to assume that inputs have any property of stationarity or ergodicity, we do not fix in advance the probability distribution of the pairs of inputs, we do not even assume the existence of such a distribution. However, the frameworks of Shannon and Kolmogorov for the definition of secrecy have similar practical interpretations. Indeed, a distribution W on $\{0,1\}^n$ has a high entropy, i.e., $H(W) \approx n$, if and only if with a high probability W returns an n-bit string with Kolmogorov complexity close to n. For a more detailed discussion of the connection between Shannon's and Kolmogorov's formalism see [14]. The formal statements in Kolmogorov's framework are usually stronger than their homologues in Shannon's framework, and theorems from the former theory in most cases formally imply the corresponding results from the latter theory, see [24]. ⌟

A characterization of the optimal size of the secret key in term of Kolmogorov complexity was suggested in [24]. We begin with the case of two parties, see Theorem 1.2 below. In this theorem, a communication protocol is randomized (we assume that the parties may use a public source of random bits, which is also accessible to the eavesdropper). Let x and y stand for inputs of Alice and Bob, r denote the string of bits produced by a public source of randomness (used by the parties and accessible to the eavesdropper), and t denote the transcript of the protocol.

▶ **Theorem 1.2** ([24]).

(i) *For any numbers $k, \ell \in \mathbb{N}$ and $\epsilon, \delta > 0$ there exist a randomized communication protocols $\pi_{k,\ell,\epsilon,\delta}$ such that on every pair of input strings (x, y) (of length at most n) satisfying*[1] $\mathrm{C}(x) \overset{\delta}{=} k$ *and* $\mathrm{C}(x \mid y) \overset{\delta}{=} \ell$, *Alice and Bob with probability $1 - \epsilon$ both obtain a result $w = w(x, y, r)$ such that*

$$[\textit{length of } w \textit{ in bits}] = \mathrm{C}(x) - \mathrm{C}(x \mid y) - O(\delta) - o(n) \textit{ and } \mathrm{C}(w \mid \langle t, r\rangle) \geq |w| - o(n) \quad (1)$$

(for $n = |x| + |y|$), which means that the size of the produced secret key is asymptotically equal to the mutual information between Alice's and Bob's inputs, and the leakage of information on the key to the eavesdropper (who can access the transcript of the protocol t and public randomness r) is negligibly small.

(ii) *The size of the key in (i) is pretty much optimal: no communication protocol can produce a key w longer than $\mathrm{C}(x) - \mathrm{C}(x \mid y) + O(\delta) + o(n)$ without loosing the property of secrecy $\mathrm{C}(w \mid \langle t, r\rangle) \geq [\textit{length of } w \textit{ in bits}] - o(n)$ (the size of a secret key cannot be made asymptotically greater than the mutual information between Alice's and Bob's inputs).*

▶ Remark 1.3. In Theorem 1.2, the values of k and ℓ are embedded in the communication protocol $\pi_{k,\ell,\epsilon,\delta}$. This means that the parties in some sense "know" (at least approximately) the values of $\mathrm{C}(x)$ and $\mathrm{C}(x \mid y)$. This is similar to the settings of the classical information theory, where the parties "know" the probability distribution on random inputs and can use a suitable protocol. The theorem is nontrivial if the approximation rate $\delta = o(n)$ as $n \to \infty$.

The secrecy of the key is understood in the information-theoretic sense: the last inequality in (1) claims that complexity of the key w conditional one *all data accessible to the adversary* must be (almost) maximal. The theorem can be adapted to a non-uniform setting where the adversary is given an auxiliary inputs s_n. In this case, all terms of Kolmogorov complexity appearing in the theorem should be relativized conditional on s_n. The theorem remains meaningful if the size of s_n is $o(n)$. ⌟

Theorem 1.2 can be extended to the multi-party setting, where $k > 2$ parties are given correlated data and need to agree on common secret key communicating via a public channel. Let us discuss in more detail the version with $k = 3$ participants. We assume now that three parties (Alice, Bob, and Charlie) are involved in the protocol. They are given inputs x, y, z respectively. We assume that all parties have an access to a common source of random bits (we denote by r the bits produced by this source) and exchange messages via a public channel (we use the conventional definition of a multi-party communication protocol with a public source of random bits, see [19]). It is assumed that every message sent by any party reaches every other party (and the eavesdropper). In what follows we consider only triples of inputs (x, y, z) with a "symmetric" complexity profile such that $\mathrm{C}(x) \approx \mathrm{C}(y) \approx \mathrm{C}(z)$ and $\mathrm{C}(x, y) \approx \mathrm{C}(x, z) \approx \mathrm{C}(y, z)$.

▶ **Theorem 1.4** (symmetric version of [24, Theorem 5.11]).

(i) *For any profile $(k_1, k_2, k_3) \in \mathbb{N}^3$ and $\epsilon, \delta > 0$ there exist a randomized communication protocols $\pi_{k_1,k_2,k_3,\epsilon,\delta}$ for three parties such that on every triple of binary input strings (x, y, z) (of length at most n) satisfying*

$$\mathrm{C}(x) \overset{\delta}{=} \mathrm{C}(y) \overset{\delta}{=} \mathrm{C}(z) \overset{\delta}{=} k_1,\ \mathrm{C}(x, y) \overset{\delta}{=} \mathrm{C}(x, z) \overset{\delta}{=} \mathrm{C}(y, z) \overset{\delta}{=} k_2,\ \mathrm{C}(x, y, z) \overset{\delta}{=} k_3 \quad (2)$$

[1] The term $\mathrm{C}(x)$ stands for the plain Kolmogorov complexity of x (optimal compression of x), the term $\mathrm{C}(x|y)$ stands for conditional Kolmogorov complexity of x conditional on y (optimal compression of x given advice y), and the notation $\mathrm{C}(x) \overset{\delta}{=} k$ and $\mathrm{C}(x|y) \overset{\delta}{=} \ell$ means that $|\mathrm{C}(x) - k| \leq \delta$ and $|\mathrm{C}(x|y) - \ell| \leq \delta$.

Alice, Bob, and Charlie can agree with probability $1-\epsilon$ *on a key* $w = w(x, y, z, r)$ *such that*

$$[\textit{length of } w \textit{ in bits}] = \tfrac{I(x:y|z)+I(x:z|y)+I(y:z|x)}{2} + I(x:y:z) \; - O(\delta) - o(n), \quad (3)$$

$$\mathrm{C}(w \mid \langle t, r\rangle) \geq |w| - o(n), \quad (4)$$

where r *is the bit string produced by the public source of randomness, and* t *is the communication transcript (concatenation of the messages sent by Alice, Bob, and Charlie).*

(ii) *The size of the key in (i) is asymptotically optimal, i.e., no communication protocol can give a key* w *asymptotically longer than*

$$\frac{1}{2}\left(I(x:y \mid z) + I(x:z \mid y) + I(y:z \mid x)\right) + I(x:y:z) + O(\delta) + o(n) \quad (5)$$

without loosing the property of secrecy (4)*.*

▶ Remark 1.5. The general version of [24, Theorem 5.11] applies to a triple of inputs with arbitrary (possibly non-symmetric) complexity profile. In the general case, the characterization of the optimal size of the secret key is more involved than (3), see [24]. We discuss only symmetric complexity profiles in order to avoid cumbersome formulas and focus on the most essential combinatorial ideas behind the proofs.

Ineq. (4) means that the eavesdropper (who can access the communication transcript t and the public randomness r) gets no information on the produced secret key. Similarly to Theorem 1.2, this secrecy condition remains meaningful even if the adversary is a non-uniform agent having advice s_n of size $o(n)$. ⌟

The known proofs of the positive parts of Theorem 1.2 and Theorem 1.4 (the existence of protocols) are quite explicit and constructive: we know specific communication protocols that allow to produce a secret key of the optimal size. More specifically, the proofs suggested in [24] provide a protocol for Theorem 1.2(i) with communication complexity

$$\min\{\mathrm{C}(x \mid y), \mathrm{C}(y \mid x)\} + O(\delta) + O(\log n) \quad (6)$$

and a protocol[2] for Theorem 1.4(i) with communication complexity

$$\mathrm{C}(x, y, z) - \frac{1}{2}\big(I(x:y \mid z) + I(x:z \mid y) + I(y:z \mid x)\big) - I(x:y:z) + O(\delta) + O(\log n). \quad (7)$$

The communication complexity (6) from Theorem 1.2(i) is known to by asymptotically optimal, see [15]. In this paper we study the communication complexity of the problem from Theorem 1.4. In fact, (7) is *not* optimal for general communication protocols; however, we show that this communication complexity is asymptotically optimal in the class of protocols with *simultaneous messages*, i.e., in the model where Alice, Bob, and Charlie send their messages in parallel, receive the messages sent by their vis-a-vis, and compute the result (secret key) without any further interaction.

[2] The scheme proposed in [24] is the so called *omniscience* protocol. In this protocol, all parties send simultaneously their messages (random hash-values of the inputs) so that each of them learns completely the entire triple of inputs (x, y, z) (this explains the term *omniscience*). The total length of the sent messages is less than $\mathrm{C}(x, y, z)$, so an eavesdropper can learn only a partial information on the inputs. The gap between the total complexity of $\mathrm{C}(x, y, z)$ and the divulged information is used to produce a secret key.

▶ **Theorem 1.6** (main result). *In the setting of Theorem 1.4, communication complexity of a protocol with simultaneous messages (the total number of bits sent by Alice, Bob, and Charlie) for triples of inputs (x, y, z) with a symmetric complexity profile (2)) cannot be smaller than*

$$\mathrm{C}(x,y,z) - \frac{1}{2}\big(I(x:y \mid z) + I(x:z \mid y) + I(y:z \mid x)\big) - I(x:y:z) - O(\delta) - O(\log n). \quad (8)$$

Communication complexity (8) is not optimal for general (multi-round) communication protocols of secret key agreement, see Proposition 8.1.

The proof of our main result combines information-theoretic techniques and spectral bounds for graphs (the expander mixing lemma). Spectral bounds *per se* are not new in communication complexity (see, e.g., the usage of Lindsey's lemma in [3]). Information-theoretic methods are also pretty common in this area. But the combination of these two techniques seems to be less standard. The key step of the proof is the observation that in some setting, when parties hold correlated data sets, for each of them it is hard to send a message that has non-negligible mutual information with the partners' data. In other words, a "too short" message sent by Alice would have zero mutual information with the data (y, z) given to Bob and Charlie. For secret key agreement protocols, this observation implies that the messages of every party inevitably have to be quite long. A similar argument can be used in problems that are not connected with cryptography, see Theorem 5.1.

▶ Remark 1.7. It is instructive to compare our work with [15], where similar questions were addressed in the setting of two parties. Our work was motivated by the observation that the argument from [15] fails in the setting with $k > 2$ parties. In fact, the multi-party setting is qualitatively different. This becomes clear when we consider secret key agreement with a sub-optimal size of the key. The technique of [15] (in the setting with *two* parties) implies that communication complexity of the secret key agreement cannot be reduced even if Alice and Bob agree on a key of a pretty small size, see the "threshold phenomenon" discussed in [15]. Apparently, this phenomenon does not occur in the multi-party setting.

Thus, to deal with multi-party setting, we have to revise significantly the techniques from [15]. We have to change both the *information-theoretic* and *algebraic* components of the proof. The most important new components appear in the information-theoretic part of the proof. In particular, we need to use the exact expression (5) for the size of the secret key. The key new element of our argument is the observation *Alice must send a message having large mutual information with Bob's and Charlie's inputs, and the cost of this task can be high*, see the discussion in Sections 3.3-3.4 (this idea was irrelevant in the setting with two parties). In the algebraic component of the proof, we adapt the definition of a spectral expander to hypergraphs and then construct a hypergraph with the required properties (see Definition 3.5 and Section 3.5). Only the bridge between the algebraic and the information-theoretic components is pretty much the same as in [15] (in the proof of Theorem 4.1 we use an argument very similar to [15, Lemma 6]). ⌟

The rest of the paper is organized as follows. In Section 2 we recall several standard definitions and introduce the notation. In Section 3 we explain informally the scheme of our argument. In Section 4 we prove the main technical tool of this paper, Theorem 4.1 (which claims that in some setting, it is hard to send a message that has non-negligible mutual information with the partners' data). In Section 5 we illustrate the application of our technique with a simple example that is not related to cryptography. In Section 6 we prove Theorem 1.6 for a restricted ("the most important") class of complexity profiles. In Section 7 we extend this result and prove Theorem 1.6 for all (symmetric) complexity profiles. We conclude with a discussion of limitations of our technique and open problems.

2 Preliminaries and Notation

2.1 General notation

For a binary string x we denote its length $|x|$. For a finite set S we denote its cardinality $\#S$.

We manipulate with equalities and inequalities for Kolmogorov complexity. Since many of them hold up to a logarithmic term, we use the notation $A \overset{\text{lg}}{=} B$, $A \overset{\text{lg}}{\leq} B$, and $A \overset{\text{lg}}{\geq} B$ for $|A - B| = O(\log n)$, $A \leq B + O(\log n)$, and $B \leq A + O(\log n)$ respectively, where n is clear from the context (n is usually the length of the strings involved in the inequality).

$\mathbb{F}_q$ denotes the field of q elements (usually $q = 2^n$). A k-dimensional vector over $\mathbb{F}_q$ is a k-tuple $(x_1, \ldots, x_k) \in \mathbb{F}_q^k$. We say that two vectors $(x_1, \ldots, x_k)$ and $(y_1, \ldots, y_k)$ in $\mathbb{F}_q^k$ are orthogonal to each other if $x_1y_1 + \ldots + x_ky_k = 0$ (the addition and multiplication are computed in the field $\mathbb{F}_q$). A vector is called self-orthogonal if it is orthogonal to itself. In a k-dimensional space over the field of characteristic 2 there are 2^{k-1} self-orthogonal vectors $(x_1, \ldots, x_k)$ and they form a linear subspace of co-dimension 1 (a vector is self-orthogonal iff $x_1 + \ldots + x_k = 0$). A *direction* in $\mathbb{F}_q^k$ is an equivalence class of non-zero vectors over $\mathbb{F}_q$ that are proportional to each other (a direction can be understood as a point in the projective space of dimension $k - 1$).

$\mathrm{C}(x)$ stands for Kolmogorov complexity of x (the length of the shortest program[3] producing x) and $\mathrm{C}(x \mid y)$ (the length of the shortest program producing x given input y) stands for Kolmogorov complexity of x given y. Respectively, $\mathrm{I}(x : y)$ and $\mathrm{I}(x : y \mid z)$ denote the mutual information between x and y and the conditional information between x and y given z. We use the notation $\mathrm{I}(x : y : z) := \mathrm{I}(x : y) - \mathrm{I}(x : y \mid z)$. For a tuple of strings $(x_1, \ldots, x_n)$ its *complexity profile* is the vector consisting of the complexity values $\mathrm{C}(x_{i_1}, \ldots, x_{i_s})$ (for all $2^n - 1$ sub-tuples $1 \leq i_1 < \ldots < i_s \leq n$). Kolmogorov complexity can be relativized: $\mathrm{C}^{\mathcal{O}}(x)$ and $\mathrm{C}^{\mathcal{O}}(x \mid y)$ stand for Kolmogorov complexity of x (conditional on y) assuming that the universal decompressor can access oracle $\mathcal{O}$. If the oracle is a finite string s, then $\mathrm{C}^{\mathcal{O}}(x) = \mathrm{C}(x \mid s) + O(1)$. For more detail on the basic facts about Kolmogorov complexity, see the full paper. A comprehensive introduction in the theory of Kolmogorov complexity can be found in [20] and [26].

2.2 Communication complexity

We use the conventional notion of a communication protocol for two or three parties, see for detailed definitions [19]. We discuss *deterministic protocols* and *randomized protocols with a public source of random bits* (see the full version of the paper for more detail).

In general, a communication protocol may consist of several rounds, when each next message of every party depends on the previously sent messages. In the *simultaneously messages* model there is no interaction: all parties send in parallel their messages that depend only on their own input data (and the random bits), and then compute the final result.

We will assume that the communication protocol has a "uniform" description. More technically, we assume that for n-bit inputs (the full description of such a protocol) has an efficient description of size $O(\log n)$. For such a protocol we do not loose much security even if the description of the protocol is available to the eavesdropper. Thus, we cannot "cheat" by embedding in the structure of the protocol any secret information hidden from the adversary.

[3] In an optimal programming language, see [20, 26] for more detail.

2.3 Reminder of the spectral graph technique

Let $G = (L \cup R, E)$ be a bi-regular bipartite graph where each vertex in L has degree D_L, each vertex in R has degree D_R, and each edge $e \in E$ connects a vertex from L with a vertex from R (observe that $\#E = \#L \cdot D_L = \#R \cdot D_R$). The adjacency matrix of such a graph is a zero-one matrix $M = \begin{pmatrix} 0 & A \\ A^\top & 0 \end{pmatrix}$ where A is a matrix of dimension $(\#L) \times (\#R)$ ($A_{xy} = 1$ if and only if there is an edge between the x-th vertex in L and the y-th vertex in R). Let $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_N$ be the eigenvalues of M, where $N = \#L + \#R$ is the total number of vertices. Since M is symmetric, all λ_i are real numbers. It is well known that for a bipartite graph the spectrum is symmetric, i.e., $\lambda_i = -\lambda_{N-i+1}$ for each i, and $\lambda_1 = -\lambda_N = \sqrt{D_L D_R}$ (see, e.g., [12]). The graphs with a large gap between λ_1 and λ_2 have the property of good *mixing*, see [25].

▶ **Lemma 2.1** (Expander Mixing Lemma for bipartite graphs, see [12]). *Let $G = (L \cup R, E)$ be a regular bipartite graph where each vertex in L has degree D_L and each vertex in R has degree D_R. Then for each $A \subseteq L$ and $B \subseteq R$ we have $\left| E(A, B) - \frac{D_L \cdot \#A \cdot \#B}{\#R} \right| \leq \lambda_2 \sqrt{\#A \cdot \#B}$, where λ_2 is the second largest eigenvalue of the adjacency matrix of G and $E(A, B)$ is the number of edges between A and B.*

▶ **Corollary 2.2.** *Let $G = (L \cup R, E)$ be a graph from Lemma 2.1. Then for $A \subseteq L$ and $B \subseteq R$ such that $\#A \cdot \#B \geq \left(\frac{\lambda_2 \#R}{D_L}\right)^2$ we have $E(A, B) = O\left(\frac{D_L \cdot \#A \cdot \#B}{\#R}\right)$.*

3 Main technical tools and the scheme of the proof

In this section we sketch the main ideas used in the proof of our principal result (Theorem 1.6).

3.1 Setting the parameters

Let us assume that $\delta = O(\log n)$, i.e., all parties of the protocol "know" the complexity profile of the triple of inputs (x, y, z) up to an additive logarithmic term[4]. This assumption does not affects significantly the argument, but it helps to avoid minor technical details and makes the explanation more transparent. To simplify the notation, in this section we discuss only triples of inputs with the profile

$$\begin{array}{l} \mathrm{C}(x) \overset{\lg}{=} \mathrm{C}(y) \overset{\lg}{=} \mathrm{C}(z) \overset{\lg}{=} kn, \ \ \mathrm{C}(x,y) \overset{\lg}{=} \mathrm{C}(x,z) \overset{\lg}{=} \mathrm{C}(y,z) \overset{\lg}{=} (2k-1)n, \\ \mathrm{C}(x,y,z) \overset{\lg}{=} (3k-3)n \end{array} \tag{9}$$

In this setting, Theorem 1.4 gives the optimal size of a secret key

$$\frac{1}{2}\big(\mathrm{I}(x : y \mid z) + I(x : z \mid y) + I(y : z \mid x)\big) + \mathrm{I}(x : y : z) \overset{\lg}{=} 1.5n. \tag{10}$$

Our aim is to bound communication complexity for inputs with this complexity profile:

▶ **Theorem 3.1** (special case of Theorem 1.6). *In the setting of Theorem 1.4, communication complexity of a protocol with simultaneous messages for some triples of inputs (x, y, z) with complexity profile (9) cannot be smaller than $(3k - 4.5)n$, which matches Eq. (8).*

[4] A logarithmic error term is, in some sense, the finest meaningful precision for Kolmogorov complexity. All our arguments can be repeated *mutatis mutandis* for any coarser precision δ such that $\log n \ll \delta(n) \ll n$.

3.2 Preliminary consideration: the need for hard inputs

The optimal size of the secret key in Theorem 1.2 and Theorem 1.4 depends only on the complexity profile of (x, y, z) and not on the combinatorial structure of the input. The situation with communication complexity (the number of bits sent by the parties) is different: it may vary significantly for different tuples of inputs with the same complexity profile. When we talk about the communication complexity of a protocol, we mean the worst-case complexity, i.e., the maximal number of sent bits among all admissible inputs. To prove a lower bound for the worst-case communication complexity, we need to provide a triple of inputs for which the parties have to send long messages. We provide a class of inputs that are guaranteed to be "hard" (for all valid protocol, for most triples of inputs from this class, communication complexity is high).

3.3 First step of the argument: conditional on Charlie's message, the mutual information between Alice's and Bob's inputs must increase

We begin with an observation that might seem to have nothing to do with communication complexity. We recall the lower bound for the size of the secret key (that applies to protocols with any communication complexity). In [24] (see Theorem 1.2(ii)) it is shown that two parties, Alice and Bob, can agree on secret key of complexity k *only if* the mutual information between Alice's input x and Bob's input y is greater than k. The proof of this statement can be easily adapted to the following slightly more general setting:

▶ **Lemma 3.2.** *Assume that there is a publicly available information s (accessible to Alice, Bob, and the eavesdropper); besides this, Alice is given a private input x and Bob is given a private input y. Then, by communication via a public channel accessible to the eavesdropper, Alice and Bob cannot agree on a secret key of complexity greater than $I(x : y \mid s)$.*

We apply this proposition to a protocol with three parties. Let t_C denote the concatenation of the messages sent by Charlie. This is a piece of publicly available information (accessible to Alice, Bob, and the eavesdropper). Due to Lemma 3.2, Alice and Bob cannot agree on a secret key with Kolmogorov complexity greater than $I(x : y \mid t_C)$ (at this point we ignore whether Charlie can learn the same key or not). Hence, in the settings (9), a secret key of size (10) can be produced only if $I(x : y \mid t_C) \geq^{\lg} 1.5n$. Observe that in the setting (9) the mutual information between x and y is equal to n. This means that the mutual information between Alice's and Bob's inputs *conditional on Charlie's message*, i.e., $I(x : y \mid t_C)$, is bigger than the unconditional mutual information between Alice's and Bob's inputs, i.e., $I(x : y)$. A pretty standard information-theoretic argument implies that the gap between $I(x : y)$ and $I(x : y \mid t_C)$ is not greater than the mutual information between $\langle x, y\rangle$ and t_C, and we conclude that $I(x, y : t_C) \geq^{\lg} n/2$. In other words, Charlie must send a message t_C that has $\geq n/2$ bits of mutual information with the pair of inputs of Alice and Bob. A similar argument implies that Alice must send a message t_A such that $I(y, z : t_A) \geq^{\lg} n/2$ and Bob must send a message t_B such that $I(x, z : t_B) \geq^{\lg} n/2$.

This part of the argument is based on Lemma 3.2, which re-employs an argument from [24] in a pretty direct way. So at this stage we need no substantially new ideas.

3.4 Second step of the argument: it may be difficult for Alice to send a message increasing the mutual information between Bob's and Charlie's inputs

We have shown above that in the setting (9) Alice, Bob, and Charlie can agree on a secret key of optimal size only if each of them sends a messages that contains $\geq^{\mathrm{lg}} n/2$ bits of mutual information with the inputs of two other parties

We are going to show that this may require sending *very long* messages (much longer than $n/2$ bits). This part of the argument is the main technical contribution of our paper. To explain this idea, we make a digression and discuss a similar problem in simpler settings.

Digression: how to say something that the interlocutor already knows. Let us consider randomized communication protocols with two participants playing non-symmetric roles. We call the participants Speaker and Listener and assume that Speaker holds an input string a and Listener holds another input string b. This is a one-way protocol: Speaker sends a message to Listener in one round, without any feedback. The aim of Speaker is to send to Listener a message that is *not completely unpredictable* from the point of view of Listener. More precisely, Speaker's message must have positive (and non-negligible) mutual information with Listener's input b. We start with a simple example when the task of Speaker is trivial.

▶ **Example 3.3.** Let Speaker is given a string $a = uv$ and Listener is given a string $b = uw$, where u, v, and w are independent incompressible strings of length n, i.e., $\mathrm{C}(uvw) \overset{\mathrm{lg}}{=} \mathrm{C}(u) + \mathrm{C}(v) + \mathrm{C}(w) \overset{\mathrm{lg}}{=} 3n$. Observe that

$$\mathrm{C}(a) \overset{\mathrm{lg}}{=} 2n,\ \mathrm{C}(b) \overset{\mathrm{lg}}{=} 2n,\ \mathrm{I}(a:b) \overset{\mathrm{lg}}{=} n \tag{11}$$

In this setting, if Speaker wants to communicate a message of length n with a *high* mutual information with Listener's y, she may send a part of u, which is know to both participants of the protocol. On the other hand, if Speaker wants to communicate a message with a *low* mutual information with Listener's b, this is also possible: Speaker may send a part of v, which is know to Speaker but not to the Listener. ⌟

▶ **Example 3.4.** Now we consider a pair (a, b) with the same complexity profile as in Example 3.3 but with a different combinatorial structure. Let a be a line in the projective plane over the finite field $\mathbb{F}_{2^n}$ and b be a point in the same projective plane incident to a, and the pair (a, b) have the maximal possible complexity (among all incident pairs (line, point) in the plane). For these a and b we have the same complexity profile (11). Indeed, we need two elements of the field ($2n$ bits of information) to specify a line or a point, but we need only one element of the field (n bits of information) to specify a point when a line is known. However, the combinatorial properties of this pair are very different from the properties of the pair in Example 3.3.

If Speaker is given a and Listener is given b as above, then Speaker cannot send a *reasonably short* message having non-negligible mutual information with Listener's input b. In fact, if Speaker wants to send to Listener a message $m = m(a)$ having δ bits of mutual information with b, then the size of m must be at least $n + \delta$. In particular, if the message m is shorter than n, then it cannot contain any information on b, see Section 4. ⌟

Example 3.4 is an instance of a much more general phenomenon. Let us have a bipartite graph $G = (V_L, V_R, E)$, where the set of vertices is $V_L \cup V_R$ and the set of edges is $E \subset V_L \times V_R$. We assume that the graph is bi-regular, i.e., all vertices in V_L have the same degree D_L and all vertices in V_R have the same degree D_R (we always assume that $D_L \geq D_R$). We say that

G is a *spectral expander*[5] if the second eigenvalue of its adjacency matrix $\lambda_2 = O(\sqrt{D_L})$. Let $(x, y) \in E$ be a "typical" edge of this graph (in the sense that its Kolmogorov complexity is close to the maximum possible value), and let x and y be the inputs given to Alice and Bob respectively. Then we have a property similar to Example 3.3: if Alice wants to send a message having δ bits of mutual information with Bob's data y, she must send a message of size at least $\log D_R + \delta$. We prove this fact using the Expander Mixing Lemma. (Example 3.4 corresponds to the graph $G = (V_L, V_R, E)$ where V_L consists of all lines in the plane, V_R consists of all points in the plane, and E is the set of all pairs of incident lines and points; it is known that this graph is a spectral expander.) [End of **Digression**.]

Now we generalize the observations from the *Digression* above and explain the main idea of the proof of Theorem 3.1. We need the following extension of the notion of expander.

▶ **Definition 3.5.** *Let $G = (V_1, V_2, V_3, H)$ be a hypergraph where the set of vertices consists of three disjoint parts V_1, V_2, V_3 of the same cardinality, and the set of hyperedges is a set $H \subset V_1 \times V_2 \times V_3$. We consider three bipartite graphs G_1, G_2, G_3 associated with hypergraph G: each G_i is a bipartite graph $(V_i, V_{j\ell}, E_i)$ (here $j = i + 1 \mod 3$ and $\ell = i + 2 \mod 3$), where $V_{j\ell}$ is the sets of $\langle y, z\rangle \in V_j \times V_\ell$ that are connected in G and $(x, \langle y, z\rangle) \in E_i$ if and only if the triple $\{x, y, z\}$ corresponds to a hyperedge in H. The hypergraph is called* tri-expander *if the graphs G_1, G_2, G_3 are bi-regular spectral expanders*[6].

We show that the communication is costly for a triple of inputs (x, y, z) that is a hyperedge in a tri-expander. To this end, we combine the idea from section 3.3 with an argument similar to the observation sketched in the *Digression*: each party must send a message having non-negligible mutual information with two other inputs (an information-theoretic argument) but this is only possible when each of the messages is very long (due to the spectral bound and the expander mixing lemma).

3.5 Construction of a tri-expander

To conclude the proof of the main result it remains to show that there exists a tri-expander with suitable parameters:

▶ **Proposition 3.6.** *For all integer numbers $k \geq 0$ and $n \geq 1$ there exists a tri-expander $G = (V_1, V_2, V_3, H)$ such that $\#V_1 = \#V_2 = \#V_3 = \Theta(2^{kn})$, $\#H = \Theta(2^{kn} \cdot 2^{(k-1)n} \cdot 2^{(k-2)n})$, and for all $i \neq j$, for every $x \in V_i$ there exists $\Theta(2^{(k-1)n})$ vertices $y \in V_j$ such that x and y are adjacent in the hypergraph.*

Proof. We construct such a tri-expander explicitly. We fix the finite field $\mathbb{F}_{2^n}$ with $q = 2^n$ elements, the $(k+2)$-dimensional space $\mathcal{L}$ over this field, and the subspace $\mathcal{L}_{so} \subset \mathcal{L}$ that consists of self-orthogonal vectors. Observe that $\#\mathcal{L}_{so} = \#\mathcal{L}/q = q^{k+1}$ (a subspace of co-dimension 1 in $\mathcal{L}$). Let V denote the space of all *directions* in $\mathcal{L}_{so}$ except for the direction $(1, \ldots, 1)$ (which is self-orthogonal for even k). Observe that $\#V = \Theta(q^k)$.

We let $V_1 = V_2 = V_3 = V$ and define H as the set of all triple $(x, y, z) \in V^3$ such that x, y, z are *distinct and pairwise orthogonal* directions in $\mathcal{L}_{so}$.

For every vector $x \in \mathcal{L}_{so}$, the condition of being orthogonal to x determines in $\mathcal{L}_{so}$ a subspace of co-dimension 1; this subspace consists of q^k vectors (including x itself as it is self-orthogonal) and, respectively, $(q^k - 1)/(q - 1)$ directions (again, including the direction

[5] We use the term *expander* without assuming that the degree of a graph is constant.

[6] The usage of the expander mixing lemma for a tri-expander seems to be similar but not literally equivalent to the hypergraph generalization of the expander mixing lemma from [18].

collinear with x). If we have two non-collinear vectors $x, y \in \mathcal{L}_{so}$, then the condition of being orthogonal to x and y determines in $\mathcal{L}_{so}$ a subspace of co-dimension 2; this subspace consists of q^{k-1} vectors (including x and y), which corresponds to $(q^{k-1}-1)/(q-1) = \Theta(q^{k-2})$ directions (once again, including the directions collinear with x and with y).

Thus, we have $\Theta(q^k)$ individual vertices, $\Theta(q^k \cdot q^{k-1})$ pairs of adjacent vertices, and $\Theta(q^k \cdot q^{k-1} \cdot q^{k-2})$ adjacent triples (hyperedges). It remains to compute the eigenvalues of the associated bipartite graphs.

▶ **Lemma 3.7.** *The hypergraph $G = (V_1, V_2, V_3, H)$ defined above is a tri-expander.*

(We give a proof in the full version of the paper.) ◀

▶ Remark 3.8. A standard counting shows that for most hyperedges (x, y, z) in the graph from Proposition 3.6 we have $\mathrm{C}(x) \overset{\lg}{=} \log \Theta(q^k) \overset{\lg}{=} kn$, $\mathrm{C}(x,y) \overset{\lg}{=} \log \Theta(q^k \cdot q^{k-1}) \overset{\lg}{=} (2k-1)n$, $\mathrm{C}(x,y,z) \overset{\lg}{=} \log \Theta(q^k \cdot q^{k-1} \cdot q^{k-2}) \overset{\lg}{=} (3k-3)n$, and we get the profile (9). ⌟

4 When it is hard to say anything that the interlocutor already knows

In this section we explain our main technical tool. We consider randomized communication protocols with two participants, Speaker and Listener. We assume that Speaker holds an input string a and Listener holds another input string b; we assume also that the complexity profile of the pairs (a, b) is known to all parties. The aim of Speaker in this protocol is to send to Listener a message that has non-negligible mutual information with Listener's input b, as we discussed in Section 3.

▶ **Theorem 4.1.** *Let $G = (V_L, V_R, E)$ be a bipartite spectral expander such that $N = \#V_L$, $M = \#V_R$, and (D_L, D_R) are the degrees of the edges in V_L and V_R respectively. Let $(a, b) \in E$ be a "typical" edge in the graph, i.e., $\mathrm{C}(a, b) \overset{\lg}{=} \log \#E$, and $\mathrm{C}(m \mid a) \overset{\lg}{=} 0$. Then $\mathrm{I}(m : b) \leq^{\lg} \max\{0, \mathrm{C}(m) - \mathrm{C}(a \mid b)\}$. In particular, if the length of m is less than $\mathrm{C}(a \mid b)$, then $\mathrm{I}(m : b) \overset{\lg}{=} 0$.*

▶ Remark 4.2. The statement of Theorem 4.1 remain valid if we relativize all terms of Kolmogorov complexity in this statement conditional on a string r such that $\mathrm{I}(r : (a, b)) \overset{\lg}{=} 0$. In what follows we present the proof without r. But every step of this argument trivially relativizes conditional on r assuming that $\mathrm{C}(a, b \mid r) \overset{\lg}{=} \mathrm{C}(a, b) \overset{\lg}{=} \log \#E$. ⌟

Proof of Theorem 4.1. Let $n_a := \log N$, $n_b = \log M$, $n'_a = \log D_R$, $n'_b = \log D_L$, and $n_{ab} := n_a - n'_a$. Using this notation, we have

$$\mathrm{C}(a) \overset{\lg}{=} n_a,\ \ \mathrm{C}(b) \overset{\lg}{=} n_b, \mathrm{C}(a \mid b) \overset{\lg}{=} n'_a,\ \ \mathrm{C}(b) \overset{\lg}{=} n_b,\ \ \mathrm{C}(b \mid a) \overset{\lg}{=} n'_b,\ \ \mathrm{I}(a : b) \overset{\lg}{=} n_{ab}.$$

Since Speaker computes the message m given the input data a, we have $\mathrm{C}(m \mid a) \overset{\lg}{=} 0$. We denote $\alpha := \mathrm{I}(m : a \mid b)$ and $\beta := \mathrm{I}(m : a : b)$. It is easy to verify that $\mathrm{C}(m) = \alpha + \beta$.

Case 1. Assume that $\mathrm{C}(m) \leq n'_a - 2 \cdot \mathsf{const} \cdot \log n$ for some $\mathsf{const} > 0$ (a constant to be specified later). In this case, to prove the theorem, we need to show that $\mathrm{I}(m : b) \overset{\lg}{=} 0$. In our notation this is equivalent to $\beta \overset{\lg}{=} 0$. More technically, we are going to show that

$$\beta \leq \mathsf{const} \cdot \log n. \tag{12}$$

For the sake of contradiction we assume that (12) is false. It is enough to consider the case when β is *somewhat large* but not *too large*, i.e., just slightly above the threshold (12). Indeed, any communication protocol violating (12) can be converted in a different protocols with the same or a smaller value of α and with $\beta = \mathsf{const} \cdot \log n + O(1)$. To this end, we observe that by discarding a few last bits of Speaker's message m we make the protocol only simpler. So, we may replace the initial message m with the shortest prefix of the initial message that still violates (12). Thus, in what follows, we assume w.l.o.g. that

$$\mathsf{const} \cdot \log n < \beta \leq \mathsf{const} \cdot \log n + O(1).$$

Let us define $A := \{a' \ : \ \mathrm{C}(a' \mid m) \leq \mathrm{C}(a \mid m)\}$ and $B := \{b' \ : \ \mathrm{C}(b' \mid m) \leq \mathrm{C}(b \mid m)\}$. We use the following standard claim:

▷ Claim. For A and B defined above we have $\#A = 2^{\mathrm{C}(a|m) \pm O(\log n)} = 2^{n_a - \alpha - \beta \pm O(\log n)}$ *and* $\#B = 2^{\mathrm{C}(b|m) \pm O(\log n)} = 2^{n_b - \beta \pm O(\log n)}$ (see, e.g. [24, Claim 4.7]).

From the claim we obtain $\#A \cdot \#B = 2^{n_a - \alpha - \beta + n_b - \beta \pm O(\log n)} = 2^{n_a + n_b - \mathrm{C}(m) - \beta \pm O(\log n)}$. Since $\mathrm{C}(m) \leq n'_a - 2 \cdot \mathsf{const} \log n$ and $\beta < \mathsf{const} \log n + O(1)$, we conclude

$$\begin{aligned} n_a + n_b - \mathrm{C}(m) - \beta n \pm O(\log n) \quad &\geq \quad n_a + n_b - (n'_a - 2 \cdot \mathsf{const} \cdot \log n) \\ &\qquad\qquad - \mathsf{const} \cdot \log n - O(\log n) \\ &\geq \quad n_{ab} + n_b + \mathsf{const} \cdot \log n - O(\log n) \geq n_{ab} + n_b. \end{aligned}$$

(To get the last inequality, we should choose the value of const in (12) so that $\mathsf{const} \cdot \log n$ majorizes the term $O(\log n)$ in the inequality above.) Thus, $\#A \cdot \#B \geq 2^{n_{ab} + n_b} = \frac{M^2}{D_L}$.

With the Corollary 2.2 we obtain $E(A, B) = O\left(\frac{D_L \cdot \#A \cdot \#B}{M}\right) = O\left(\frac{\#A \cdot \#B}{M / D_L}\right)$. Now observe that given m and the numbers $\mathrm{C}(a \mid m)$ and $\mathrm{C}(b \mid m)$ we can enumerate the sets A and B and, therefore, we can describe (a, b) by the index of this edge in the list of all edges between A and B. The size of such an index is $\log E(A, B)$. Hence,

$$\begin{aligned} \mathrm{C}(a, b \mid m) \leq^{\mathrm{lg}} \log E(A, B) \quad &\leq^{\mathrm{lg}} \quad (n_a + n_b - \mathrm{C}(m) - \beta) - (n_b - n'_b) \\ &= \quad n_a + n'_b - \mathrm{C}(m) - \beta = \mathrm{C}(a, b) - \mathrm{C}(m) - \beta, \end{aligned}$$

and $\mathrm{C}(a, b) \leq^{\mathrm{lg}} \mathrm{C}(m) + \mathrm{C}(a, b \mid m) \leq^{\mathrm{lg}} \mathrm{C}(a, b) - \beta$. The terms $O(\log n)$ hidden in the notation $\leq^{\mathrm{lg}}$ and $\overset{\mathrm{lg}}{=}$ in this inequality do not depend on β. Thus, we get a contradiction if the constant in (12) is chosen large enough.

Case 2. Now we assume that $\mathrm{C}(m) = n'_a + \delta$ for an arbitrary δ. Denote by m' the prefix of m of length $(n'_a - \mathsf{const} \log n)$ and by m'' the suffix of m of length $(\delta + \mathsf{const} \log n)$. We know from Case 1 that $\mathrm{I}(m' : b) \overset{\mathrm{lg}}{=} 0$. It remains to apply the chain rule,

$$\mathrm{I}(m : b) \overset{\mathrm{lg}}{=} \mathrm{I}(m' : b) + \mathrm{I}(m'' : b \mid m') \overset{\mathrm{lg}}{=} \mathrm{I}(m'' : b \mid m') \leq^{\mathrm{lg}} |m''| \overset{\mathrm{lg}}{=} \delta.$$

and the theorem is proven. ◀

▶ **Corollary 4.3.** *Let $G = (V_L, V_R, E)$ be a bipartite spectral expander such that $N = \#V_L$, $M = \#V_R$, and (D_L, D_R) are the degrees of the edges in V_L and V_R respectively.*

(a) *We assume that Speaker and Listener are given, respectively, a and b that are ends of a typical edge $(a, b) \in E$ in the graph. We consider a one-round communication protocol where Speaker sends to Listener a message $m = m(a)$. Then $\mathrm{I}(m : b) \leq^{\mathrm{lg}} \max\{0, \mathrm{C}(m) - \mathrm{C}(a \mid b)\}$. In particular, if the length of m is less than $\mathrm{C}(a \mid b)$, then $\mathrm{I}(m : b) \overset{\mathrm{lg}}{=} 0$.*

(b) *A similar statement is true if Speaker and Listener are given instead of a and b some inputs a' and b' such that $\mathrm{C}(a' \mid a) \overset{\mathrm{lg}}{=} 0$ and $\mathrm{C}(b' \mid b) \overset{\mathrm{lg}}{=} 0$ (e.g., if Speaker is given a function of a vertex $a \in V_L$ and Listener is given a function of a vertex $b \in V_R$).*

5 Protocols with simultaneous messages : a warm-up example

In this section we use Theorem 4.1 from the previous section to prove a lower bound for communication complexity of the following problem. Alice and Bob hold, respectively, lines a and b in a plane (intersecting at one point c). They send to Charlie (in parallel, without interacting with each other) some messages so that Charlie can reconstruct the intersection point. We argue that the trivial protocol (where Alice and Bob send the full information on their lines) is essentially optimal.

▶ **Theorem 5.1.** *Let Alice and Bob be given lines in the projective plane over the finite field $\mathbb{F}_{2^n}$ (we denote them a and b respectively), and it is known that the lines intersect at point c. Another participant of the protocol Charlie has no input information. Alice and Bob (without a communication with each other) send to Charlie messages m_A and m_B so that Charlie can find c. For every communication protocol for this problem, for some a, b we have $|m_A| + |m_B| \geq^{\mathrm{lg}} 4n$, which means essentially that in the worst case Alice and Bob must send to Charlie all their data (for a typical pair of lines we have $\mathrm{C}(a) + \mathrm{C}(b) \overset{\mathrm{lg}}{=} 4n$).*

In the setting of Theorem 5.1, the inputs of Alice and Bob contain n bits of the mutual information with c, so an easy lower bound for the communication complexity is $n + n = 2n$. However, due to the spectral properties of graphs implicitly present in this construction, the true communication complexity of this problem is twice bigger.

Sketch of the proof (see the full version of the paper for the details). In this sketch we ignore the public randomness and explain the argument for deterministic protocols. A generalization for protocols with public randomness is pretty straightforward, see full version of the paper.

Let (a, b) be a pair of lines in a projective plane over $\mathbb{F}_{2^n}$ intersecting at a point c, such that $\mathrm{C}(a, b) \overset{\mathrm{lg}}{=} \mathrm{C}(a) + \mathrm{C}(b) \overset{\mathrm{lg}}{=} 4n$ (which is the case for most pairs of lines in the plane). Observe that $\mathrm{I}(a : c) \overset{\mathrm{lg}}{=} n$ and $\mathrm{I}(b : c) \overset{\mathrm{lg}}{=} n$. It follows that for the messages $m_A = m_A(a)$ and $m_B = m_B(b)$ we have $\mathrm{I}(m_A : c) \leq^{\mathrm{lg}} n$ and $\mathrm{I}(m_B : c) \leq^{\mathrm{lg}} n$. Using standard information theoretic inequalities, one can show that Alice's message m_A and Bob's message m_B determine the point c uniquely only if $\mathrm{I}(m_A : c) \overset{\mathrm{lg}}{=} n$ and $\mathrm{I}(m_B : c) \overset{\mathrm{lg}}{=} n$. Thus, Alice and Bob must send messages with large enough information on c.

The graph of possible pairs (a, c) and the graph of possible pairs (b, c) (the configurations $(\text{line}, \text{point})$) is the same as in Example 3.4. Hence, we can apply Theorem 4.1 (Alice and Bob play the roles of Speaker, and Charlie plays the role of Listener) and conclude that $\mathrm{I}(m_A : c) \leq^{\mathrm{lg}} \max\{0, \mathrm{C}(m_A) - n\}$ and $\mathrm{I}(m_B : c) \leq^{\mathrm{lg}} \max\{0, \mathrm{C}(m_B) - n\}$. In particular, $\mathrm{I}(m_A : c) \overset{\mathrm{lg}}{=} n$ and $\mathrm{I}(m_B : c) \overset{\mathrm{lg}}{=} n$ only if Kolmogorov complexities of m_A and m_B are both at least $2n$. Thus, the total communication complexity is $\geq^{\mathrm{lg}} 2n + 2n = 4n$. ◀

6 Secret key agreement: a lower bound for the most crucial profile

In this section we prove a lower bound for communication complexity of secret key agreement with three parties. Let us recall the setting. We assume that Alice, Bob, and Charlie are given inputs x, y, z respectively with the complexity profile (9). This is a pretty “generic” complexity profile; by choosing k, we control the gap between the complexities of x, y, z and the mutual informations shared by the inputs.

We consider communication protocols with public randomness. Denote by r the string of random bits accessible for all the parties (including the eavesdropper). We assume that Alice, Bob, and Charlie broadcast simultaneously messages $m_A = m_A(x,r)$, $m_B = m_B(y,r)$, $m_C = m_C(z,r)$ over a public communication channel. Then each of them computes the final result

$$\text{key}_{\text{Alice}}(x,r,m_B,m_C),\ \text{key}_{\text{Bob}}(y,r,m_A,m_C),\ \text{key}_{\text{Charlie}}(z,r,m_A,m_B).$$

We say that a protocol is successful if $\text{key}_{\text{Alice}} = \text{key}_{\text{Bob}} = \text{key}_{\text{Charlie}} = w$ (i.e., the parties agree on a common key w) and $\mathrm{C}(w \mid \langle m_A, m_B, m_C, r\rangle) \overset{\text{lg}}{=} |w|$ (i.e., the eavesdropper gets no information on this key).

Theorem 1.4 claims that for any $\epsilon > 0$ there exists a protocol that is successful with probability $(1-\epsilon)$, and the size of the key is equal to (5), which gives for the profile (9) the value $1.5n$. Moreover, this value of the key is optimal (up to an additive term $O(\log n)$).

It was shown in [24] that a secret key of this size can be obtained in an *omniscience* protocol. In this protocol, the parties broadcast messages so that each of them learns completely the entire triple of inputs (x,y,z). The total length of the broadcasted messages bits is less than $\mathrm{C}(x,y,z)$, so an eavesdropper can learn only a partial information on the inputs. More specifically, communication complexity of the omniscience protocol is (7), which is $(3k-4.5)n$ for a triple satisfying (9). The gap between $\mathrm{C}(x,y,z) \overset{\text{lg}}{=} (3k-3)n$ and the amount of the divulged information is used to produce the secret key of size $1.5n$.

The omniscience protocol used in [24] provides an *upper bound* on the communication complexity of secret key agreement. In what follows we prove the matching *lower bound* (for protocols with simultaneous messages) and show that $(3k-4.5)n$ is the optimal communication complexity for a protocol of secret key agreement protocols with simultaneous messages for inputs satisfying (9). The proof follows the scheme sketched in Section 3. The first ingredient of this proof is Lemma 3.2 (see p. 8).

Sketch of proof of Lemma 3.2. This lemma is a relativized version of [24, Theorem 4.2]. One can follow the argument from [24] step by step, substituting s as a supplementary condition in each term of Kolmogorov complexity appearing in the proof. ◀

▶ **Corollary 6.1.** *Consider a communication protocol with three parties where Alice is given x, Bob is given y, and Charlie is given z. Denote by m_C the concatenation of all messages broadcasted by Charlie during the communication. If the parties agree on a secret key w on which the eavesdropper gets no information (even given access to the messages sent by all parties), then $\mathrm{C}(w) \leq^{\text{lg}} \mathrm{I}(x:y \mid r, m_C)$.*

Proof. We apply Lemma 3.2 substituting m_C instead of the public information s. ◀

▶ **Theorem 3.1 rephrased.** *Let Alice, Bob, and Charlie be given x, y, and z respectively such that (x,y,z) is a hyperedge of the hypergraph $G = (V_1, V_2, V_3, H)$ from Proposition 3.6 (the pairwise disjoint self-orthogonal directions in a $(k+2)$-dimensional vector space over $\mathbb{F}_{2^n}$). We consider non-interactive communication protocols where Alice, Bob, and Charlie send messages m_A, m_B, and m_C respectively and produce a secret key w with the optimal complexity $\mathrm{C}(w) \overset{\text{lg}}{=} 1.5n$. Then $\mathrm{C}(m_A) \geq^{\text{lg}} (k-1.5)n$, $\mathrm{C}(m_B) \geq^{\text{lg}} (k-1.5)n$, $\mathrm{C}(m_C) \geq^{\text{lg}} (k-1.5)n$, and the communication complexity of the protocol is at least $(3k-4.5)n - O(\log n)$.*

Proof. To simplify the notation, we ignore the bits r provided by the public source of randomness and explain the proof for deterministic protocols. Our argument trivially relativizes given any instance of random bits r independent of (x,y) (which is true with a probability close to 1), cf. the proof of Theorem 5.1 in the full version of the paper.

From Corollary 6.1 we know that the size of the key (in our case $1.5n$) cannot be greater than $\mathrm{I}(x : y \mid m_C)$. By the construction of the tri-expander, $\mathrm{I}(x : y) \overset{\lg}{=} n$. Therefore, the difference between $\mathrm{I}(x : y)$ and $\mathrm{I}(x : y \mid m_C)$ is at least $0.5n$.

▶ **Lemma 6.2.** *For all binary strings x, y, s it holds $\mathrm{I}(x : y \mid s) - \mathrm{I}(x : y) \leq^{\lg} \mathrm{I}(s : xy)$.*

(See the proof of the lemma in the full version of the paper.) We combine Corollary 6.1 with Lemma 6.2 and obtain $\mathrm{I}(m_C : xy) \geq^{\lg} 0.5n$.

Then, we apply Theorem 4.1 to the bipartite graph G_3 associated with the tri-expander G (see p. 11); here Charlie plays the role of Speaker, and Alice and Bob together play the role of Listener. Since $\mathrm{I}(m_C : xy) \geq^{\lg} 0.5n$, we obtain $\mathrm{C}(m_C) \geq^{\lg} \mathrm{C}(z \mid x, y) + 0.5n \overset{\lg}{=} (k - 1.5)kn$. A similar argument applies to $\mathrm{C}(m_A)$ and $\mathrm{C}(m_B)$, and we are done. ◀

7 Secret key agreement: a lower bound for all symmetric profiles

Proof of Theorem 1.6. If the complexity profile of (x, y, z) is symmetric then it can be specified by three real numbers:

$$\begin{cases} \mathrm{C}(x \mid y, z) \overset{\lg}{=} \mathrm{C}(y \mid x, z) \overset{\lg}{=} \mathrm{C}(z \mid x, y) \overset{\lg}{=} \alpha, \\ \mathrm{I}(x : y \mid z) \overset{\lg}{=} \mathrm{I}(x : z \mid y) \overset{\lg}{=} \mathrm{I}(y : z \mid x) \overset{\lg}{=} \beta, \ \mathrm{I}(x : y : z) \overset{\lg}{=} \gamma. \end{cases} \tag{13}$$

In what follows we say that the triple of numbers (α, β, γ) represent *symmetric complexity profile* of the triple (x, y, z).

In Theorem 3.1 we proved that communication complexity (7) of the omniscience protocol is optimal in case $\alpha = (k-2)n$, $\beta = n$, and $\gamma = 0$. We reduce the problem with arbitrary α, β, γ to the special case settled in Theorem 3.1.

▶ **Lemma 7.1.** *If communication complexity* (7) *is optimal (in the worst case) for some triples of inputs (x, y, z) with complexity profile* (13)*, then*

(i) *for every positive $\delta \leq n$, the omniscience protocol is also optimal for some triples of inputs (x', y', z') with symmetric complexity profile $(\alpha', \beta', \gamma') = (\alpha - \delta, \beta, \gamma)$;*

(ii) *for every positive δ, the omniscience protocol is also optimal for some triples of inputs (x', y', z') with symmetric complexity profile $(\alpha', \beta', \gamma') = (\alpha, \beta, \gamma + \delta)$;*

(iii) *if $\alpha \overset{\lg}{=} (k-2)n$, $\beta \overset{\lg}{=} n$, $\gamma \overset{\lg}{=} 0$ (as in Theorem 3.1), then for every positive $\delta \leq \beta/2$ the omniscience protocol is also optimal for some triples of inputs (x', y', z') with symmetric complexity profile $(\alpha', \beta', \gamma') = (\alpha, \beta + \delta, \gamma - 3\delta)$.*

In Lemma 7.1 we show that the existence of an "excessively efficient" protocol for triples of inputs with modified symmetric profiles $(\alpha', \beta', \gamma')$ would imply an "excessively efficient" protocol for the original symmetric profile (9), which is impossible due to Theorem 3.1. The proof of this lemma uses mostly techniques of Kolmogorov complexity that are not specific for communication problems. The argument is based on repeated application of Muchnik's theorem on conditional descriptions ([23]). Due to the lack of space, the proof of the lemma is deferred to full version of the paper.

It is not hard to verify that starting with a triple (x, y, z) from Theorem 3.1 and then applying the reductions from Lemma 7.1, we can obtain any realizable profiles (2). Indeed, we begin with a triple of pairwise orthogonal directions (x, y, z) with $\alpha = (k-2)n, \beta = n, \gamma = 0$ for a suitable n and k, then apply Lemma 7.1 (ii) or Lemma 7.1 (iii) to get a triple (x', y', z') with a suitable $I(x' : y' : z')$ (case (ii) serves to make the triple mutual information positive, and case (iii) is needed if we want to make it negative), and further apply Lemma 7.1 (i) to trim the value of α. Thus, Theorem 3.1 implies optimality of (7) for triples of inputs (x, y, z) with arbitrary symmetric complexity profile (2). ◀

8 Conclusion and open problems

We proved that the standard omniscience protocol provides the optimal worst-case communication complexity of the problem of secret key agreement (with three parties) in the class of protocols with simultaneous messages. A natural direction for further research is the study of the limits of our approach. A more specific open problem is to settle the communication complexity of multi-party secret key agreement for multi-round protocols. Indeed, in the multi-party settings, the existing communication protocols can be actually improved at the cost of increasing the number of rounds. In particular, the communication complexity (7) is no longer the optimal for multi-round protocols:

▶ **Proposition 8.1.** *In the setting of Theorem 3.1 there is a* multi-round *communication protocol* (*not a* simultaneous messages *protocol*) *with communication complexity* $(2k-2.5)n+O(\log n)$*, where the parties agree on a secret key of the optimal size* $1.5n - O(\log n)$*.*

(See the proof in the full version of the paper.) Our technique implies *some* lower bounds for communication complexity of interactive protocols, but this bound does not match the known upper bounds.

Another open problem is to establish the trade-off between the size of the secret key and the optimal communication complexity. For two-parties protocols, communication complexity of secret key agreement cannot be reduced even if Alice and Bob agree on a very small secret key, see the "threshold phenomenon" in [15]; for protocols with $k \geq 3$ parties the situation is different, and communication complexity may be improved if the size of the secret key is suboptimal. It would be also interesting to extend our results to the communication model with private sources of randomness.

References

1 Rudolf Ahlswede and Imre Csiszár. Common randomness in information theory and cryptography - I: secret sharing. *IEEE Trans. Inf. Theory*, 39(4):1121–1132, 1993. `doi:10.1109/18.243431`.

2 Luis Antunes, Sophie Laplante, Alexandre Pinto, and Liliana C. M. Salvador. Cryptographic security of individual instances. In Yvo Desmedt, editor, *Information Theoretic Security - Second International Conference, ICITS 2007, Madrid, Spain, May 25-29, 2007, Revised Selected Papers*, volume 4883 of *Lecture Notes in Computer Science*, pages 195–210. Springer, 2007. `doi:10.1007/978-3-642-10230-1_17`.

3 László Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory (preliminary version). In *27th Annual Symposium on Foundations of Computer Science, Toronto, Canada, 27-29 October 1986*, pages 337–347. IEEE Computer Society, 1986. `doi:10.1109/SFCS.1986.15`.

4 Charles H. Bennett, François Bessette, Gilles Brassard, Louis Salvail, and John A. Smolin. Experimental quantum cryptography. *J. Cryptol.*, 5(1):3–28, 1992. `doi:10.1007/BF00191318`.

5 Matthieu Bloch, Onur Günlü, Aylin Yener, Frédérique Oggier, H. Vincent Poor, Lalitha Sankar, and Rafael F. Schaefer. An overview of information-theoretic security and privacy: Metrics, limits and applications. *IEEE Journal on Selected Areas in Information Theory*, 2(1):5–22, 2021. `doi:10.1109/JSAIT.2021.3062755`.

6 Imre Csiszár and Prakash Narayan. Secrecy capacities for multiple terminals. *IEEE Trans. Inf. Theory*, 50(12):3047–3061, 2004. `doi:10.1109/TIT.2004.838380`.

7 Igor Devetak and Andreas Winter. Distillation of secret key and entanglement from quantum states. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 461(2053):207–235, January 2005. `doi:10.1098/rspa.2004.1372`.

8 Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Trans. Inf. Theory*, 22(6):644–654, 1976. doi:10.1109/TIT.1976.1055638.

9 Yan Zong Ding. Error correction in the bounded storage model. In Joe Kilian, editor, *Theory of Cryptography*, pages 578–599, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.

10 Yevgeniy Dodis, Bhavana Kanukurthi, Jonathan Katz, Leonid Reyzin, and Adam Smith. Robust fuzzy extractors and authenticated key agreement from close secrets. *IEEE Transactions on Information Theory*, 58(9):6207–6222, 2012. doi:10.1109/TIT.2012.2200290.

11 Yevgeniy Dodis and Adam D. Smith. Correcting errors without leaking partial information. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 654–663. ACM, 2005. doi:10.1145/1060590.1060688.

12 Shai Evra, Konstantin Golubev, and Alexander Lubotzky. Mixing properties and the chromatic number of ramanujan complexes, 2014. arXiv:1407.7700.

13 Noah Golowich and Madhu Sudan. Round complexity of common randomness generation: The amortized setting. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 1076–1095. SIAM, 2020. doi:10.1137/1.9781611975994.66.

14 Peter Grünwald and Paul M. B. Vitányi. Shannon information and kolmogorov complexity. *CoRR*, cs.IT/0410002, 2004. URL: http://arxiv.org/abs/cs.IT/0410002.

15 Emirhan Gürpinar and Andrei E. Romashchenko. Communication complexity of the secret key agreement in algorithmic information theory. In Javier Esparza and Daniel Král', editors, *45th International Symposium on Mathematical Foundations of Computer Science, MFCS 2020, August 24-28, 2020, Prague, Czech Republic*, volume 170 of *LIPIcs*, pages 44:1–44:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICS.MFCS.2020.44.

16 Ryszard Horodecki, Paweł Horodecki, Michał Horodecki, and Karol Horodecki. Quantum entanglement. *Reviews of Modern Physics*, 81(2):865–942, June 2009. doi:10.1103/revmodphys.81.865.

17 Tanya Ignatenko and Frans M. J. Willems. Biometric security from an information-theoretical perspective. *Found. Trends Commun. Inf. Theory*, 7:135–316, 2012. URL: https://api.semanticscholar.org/CorpusID:51848802.

18 A. Wigderson J. Friedman. On the second eigenvalue of hypergraphs. *Combinatorica 15*, pages 43–65, 1995. doi:10.1007/BF01294459.

19 Eyal Kushilevitz and Noam Nisan. *Communication complexity.* Cambridge University Press, 1997.

20 Ming Li and Paul M. B. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications, 4th Edition.* Texts in Computer Science. Springer, 2019. doi:10.1007/978-3-030-11298-1.

21 Ueli M. Maurer. Secret key agreement by public discussion from common information. *IEEE Trans. Inf. Theory*, 39(3):733–742, 1993. doi:10.1109/18.256484.

22 Ralph C. Merkle. Secure communications over insecure channels. *Commun. ACM*, 21(4):294–299, 1978. doi:10.1145/359460.359473.

23 Andrei A. Muchnik. Conditional complexity and codes. *Theor. Comput. Sci.*, 271(1-2):97–109, 2002. doi:10.1016/S0304-3975(01)00033-0.

24 Andrei E. Romashchenko and Marius Zimand. An operational characterization of mutual information in algorithmic information theory. *J. ACM*, 66(5):38:1–38:42, 2019. doi:10.1145/3356867.

25 A. Wigderson S. Hoory, N. Linial. Expander graphs and their applications. *Bulletin of the American Mathematical Society 43*, pages 439–561, August 2006.

26 Alexander Shen, Vladimir Andreevich Uspensky, and Nikolay Vereshchagin. *Kolmogorov Complexity and Algorithmic Randomness.* American Mathematical Society, 2017. URL: https://hal-lirmm.ccsd.cnrs.fr/lirmm-01803620.

27 Benjamin Smith. Pre- and post-quantum diffie-hellman from groups, actions, and isogenies. In Lilya Budaghyan and Francisco Rodríguez-Henríquez, editors, *Arithmetic of Finite Fields - 7th International Workshop, WAIFI 2018, Bergen, Norway, June 14-16, 2018, Revised Selected Papers*, volume 11321 of *Lecture Notes in Computer Science*, pages 3–40. Springer, 2018. doi:10.1007/978-3-030-05153-2_1.

28 Madhu Sudan, Badih Ghazi, Noah Golowich, and Mitali Bafna. Communication-rounds tradeoffs for common randomness and secret key generation. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1861–1871. SIAM, 2019. doi:10.1137/1.9781611975482.112.

29 Madhu Sudan, Himanshu Tyagi, and Shun Watanabe. Communication for generating correlation: A unifying survey. *IEEE Trans. Inf. Theory*, 66(1):5–37, 2020. doi:10.1109/TIT.2019.2946364.