Communication-Rounds Tradeoffs for Common Randomness and Secret Key Generation

Mitali Bafna* Badih Ghazi† Noah Golowich‡ Madhu Sudan§

Abstract

We study the role of interaction in the Common Randomness Generation (CRG) and Secret Key Generation (SKG) problems. In the CRG problem, two players, Alice and Bob, respectively get samples $X_1, X_2, \ldots$ and $Y_1, Y_2, \ldots$ with the pairs $(X_1, Y_1), (X_2, Y_2), \ldots$ being drawn independently from some known probability distribution μ. They wish to communicate so as to agree on L bits of randomness. The SKG problem is the restriction of the CRG problem to the case where the key is required to be close to random even to an eavesdropper who can listen to their communication (but does not have access to the inputs of Alice and Bob). In this work, we study the relationship between the amount of communication and the number of rounds of interaction in both the CRG and the SKG problems. Specifically, we construct a family of distributions $\mu = \mu_{r,n,L}$, parametrized by integers r, n and L, such that for every r there exists a constant $b = b(r)$ for which CRG (respectively SKG) is feasible when $(X_i, Y_i) \sim \mu_{r,n,L}$ with $r+1$ rounds of communication, each consisting of $O(\log n)$ bits, but when restricted to $r/2 - 2$ rounds of interaction, the total communication must exceed $\Omega(n/\log^b(n))$ bits. Prior to our work no separations were known for $r \geq 2$.

1 Introduction

1.1 Problem Definition In this work, we study the *Common Randomness Generation (CRG)* and *Secret Key Generation (SKG)* problems — two central questions in information theory, distributed computing and cryptography — and study the need for interaction in solving these problems.

In the CRG problem, two players, Alice and Bob, have access to correlated randomness, with Alice being given $X_1, X_2, \ldots,$ and Bob being given $Y_1, Y_2, \ldots,$ where $(X_1, Y_1), (X_2, Y_2), \ldots$ are drawn i.i.d from some known probability distribution μ. Their goal is to agree on L bits of entropy with high probability while communicating as little as possible. In the SKG problem, the generated random key is in addition required to be secure against a third player, Eve, who does not have access to the inputs of Alice and Bob but who can eavesdrop on their conversation. The CRG and SKG settings are illustrated in Figures 1 and 2 respectively.

Common random keys play a fundamental role in distributed computing and cryptography. They can often be used to obtain significant performance gains that would otherwise be impossible using deterministic or private-coin protocols. Under the additional secrecy constraints, the generated keys are of crucial importance as they can be used for encryption – a central goal of cryptography.

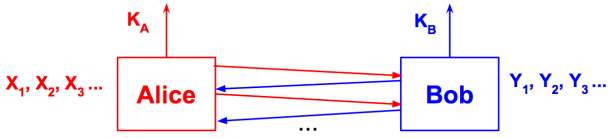

Figure 1: Common Randomness Generation (CRG)

This paper investigates the tradeoff between rounds and communication for protocols for common randomness and secret key generation: We start with some terminology needed to describe our problem. We say that a communication protocol Π is an (r, c)-*protocol* if it involves at most r rounds of interaction with Alice starting and with the total length of all the messages being at most c bits. Let $H_\infty(\cdot)$ denote the min-entropy func-

*Harvard John A. Paulson School of Engineering and Applied Sciences, 33 Oxford Street, Cambridge, MA 02138, USA. mitalibafna@g.harvard.edu. Work supported in part by a Simons Investigator Award and NSF Award CCF 1715187.

†Google Research, 1600 Amphitheatre Parkway Mountain View, CA 94043, USA. badihghazi@gmail.com This work was partly done while the author was a student at MIT. Supported in parts by NSF CCF-1650733 and CCF-1420692.

‡Harvard University. ngolowich@college.harvard.edu

§Harvard John A. Paulson School of Engineering and Applied Sciences, 33 Oxford Street, Cambridge, MA 02138, USA. madhu@cs.harvard.edu. Work supported in part by a Simons Investigator Award and NSF Award CCF 1715187.

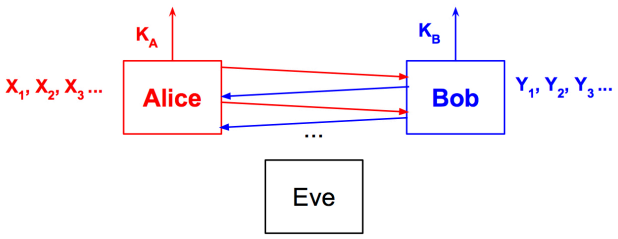

Figure 2: Secret Key Generation (SKG)

tion. A protocol is said to be an (L, ϵ)-*CRG* scheme for a correlation source μ if Alice and Bob get a finite number of i.i.d. samples of μ, and after the final round of Π, Alice outputs a key K_A and Bob outputs a key K_B, with K_A and K_B belonging to a finite set, satisfying $\min\{H_\infty(K_A), H_\infty(K_B)\} \geq L$, and with K_A and K_B being equal with probability at least $1 - \epsilon$. A protocol is said to be an (L, ϵ)-*SKG* scheme for μ if it is an (L, ϵ)-CRG scheme for μ and satisfies the additional security guarantee that $\max\{I(\Pi; K_A), I(\Pi; K_B)\} = o(1)$ where Π is also used to denote the protocol transcript and $I(\cdot;\cdot)$ is the mutual information. Then, we define the r-round communication complexity of (L, ϵ)-CRG of a correlation source μ, denoted by $CC_r(CRG_{L,\epsilon}(\mu))$, as the smallest c for which there is an (r, c)-protocol that is an (L, ϵ)-CRG scheme for μ. We similarly define the r-round communication complexity of (L, ϵ)-SKG of μ and denote it by $CC_r(SKG_{L,\epsilon}(\mu))$. In terms of the above notation we study the functions $CC_r(CRG_{L,\epsilon}(\mu))$ and $CC_r(SKG_{L,\epsilon}(\mu))$ as we vary r.

1.2 History The CRG and SKG problems have been well-studied in information theory and theoretical computer science. In information theory, they go back to the seminal work of Shannon on secrecy systems [Sha49], which was followed by the central works of Maurer [Mau93] and Ahlswede and Csiszár [AC93, AC98]. A crucial motivation for the study of SKG is the task of *secure encryption*, where a common secret key can potentially be used to encrypt/decrypt messages over an insecure channel. It turns out that without *correlated* inputs (and even allowing each party an unlimited amount of private randomness), efficiently generating common randomness is infeasible: agreeing on L bits of randomness with probability γ can be shown to require communicating at least $L - O(\log(1/\gamma))$ bits [1]. Since the original work of Shannon, the questions of how much randomness can be agreed on, with what probability, with what type of correlation and with how many rounds of interaction have attracted significant effort in both the information theory and theoretical computer science communities (e.g., [Mau93, AC93, AC98, CN00, GK73, Wyn75, CN04, ZC11, Tya13, LCV15, LCV16, BM11, CMN14, GR16, GJ18] to name a few). In particular, Ahlswede and Csiszár studied the CRG and SKG problems in the case of *one-way communication* where they gave a characterization of the ratio of the entropy of the key to the communication in terms of the *strong data processing constant* of the source (which is closely related to its hypercontractive properties [AG76, AGKN13]).

We point out that the aforementioned results obtained in the information theory community hold for the *amortized* setup where the aim is to characterize the achievable (H, C) pairs for which for every positive δ, there is a large enough N, such that there is a CRG/SKG scheme taking as input N i.i.d. copies from the source and generating $(H - \delta) \cdot N$ bits of entropy while communicating at most $(C + \delta) \cdot N$ bits. Moreover, these results mostly focus on the regime where the agreement probability gets arbitrarily close to one for sufficiently large N. The *non-amortized* setup, where the entropy of the keys and the communication are potentially independent of the number of i.i.d. samples drawn from the source, as well as the setting where the agreement probability is not necessarily close to one, have been studied in several works within theoretical computer science. In particular, for the *doubly symmetric binary source*, Bogdanov and Mossel gave a CRG protocol with a nearly tight agreement probability in the *zero-communication* case where Alice and Bob are not allowed to communicate [BM11]. This CRG setup can be viewed as an abstraction of practical scenarios where hardware-based procedures are used for extracting a unique random ID from process variations [LLG+05, SHO08, YLH+09] that can then be used for authentication [LLG+05, SD07]. Guruswami and Radhakrishnan generalized the study of Bogdanov and Mossel to the case of one-way communication (in the non-amortized setup) where they gave a protocol achieving a near-optimal tradeoff between (one-way) communication and agreement probability [GR16]. Later, [GJ18] gave explicit and sample-efficient CRG (and SKG) schemes matching the bounds of [BM11] and [GR16] for the doubly symmetric binary source and the bivariate Gaussian source.

Common randomness is thus a natural model for studying how shared keys can be generated in settings where only weaker forms of correlation are available. It is one of the simplest and most natural questions within the study of *correlation distillation* and the *simulation*

[1] This fact is a special case of several known results in the literature on CRG. In particular, it follows from the lower bound on agreement distillation in [CGMS17, Theorem 2].

of joint distributions [GK73, Wyn75, Wit75, MO04, MOR+06, KA15, GKS16b, DMN18, GKR17].

Moreover, when studying the setup of *communication with imperfectly shared randomness*, Canonne et al. used lower bounds for CRG *as a black box* when proving the existence of functions having small communication complexity with public randomness but large communication complexity with imperfectly shared randomness [CGMS17]. Their setup – which interpolates between the extensively studied public-coin and private-coin models of communication complexity – was first also independently introduced by [BGI14] and further studied in [GKS16a, GJ18].

Despite substantial work having been done on CRG and SKG, some very basic questions remained open such as the the quest of this paper, namely the role of interaction in generating common randomness (or secret keys). Recently, Liu, Cuff and Verdu generalized the CRG and SKG characterizations of Ahlswede and Csiszár to the case of *multi-round communication* [LCV15, LCV16, Liu16]. Their characterization has been shown by [GJ18] to be intimately connected to the notions of *internal* and *external information costs* of protocols which were first defined by [BJKS04, BBCR13] and [CSWY01] respectively (who were motivated by the study of direct-sum questions arising in theoretical computer science). However their work does not yield sources for which randomness generation requires many rounds of interaction (to be achieved with low commununication). Their work does reveal sources where interaction does *not* help. For example, in the case where the agreement probability tends to one, Tyagi had shown that for *binary symmetric* sources, interaction does not help, and conjectured the same to be true for any (possibly asymmetric) *binary* source [Tya13]– a conjecture which was proved by Liu, Cuff and Verdu [LCV16]. Morever, Tyagi constructed a source on *ternary alphabets* for which there is a constant factor gap between the 1-round and 2-round communication complexity for Common Randomness and Secret Key Generation. This seems to be the strongest tradeoff known for communication complexity of CRG or SKG till our work.

1.3 Our Results In this work, we study the relationship between the *amount of communication* and the *number of rounds* of interaction in each of the CRG and SKG setups, namely: can Alice and Bob communicate less and still generate a random/secret key by interacting for a larger number rounds?

For every constant r and parameters n and L, we construct a family of probability distributions $\mu = \mu_{r,n,L}$ for which CRG (respectively SKG) is possible with r rounds of communication, each consisting of $O(\log n)$ bits, but when restricted to $r/2$ rounds, the total communication of any protocol should exceed $n/\log^{\omega(1)}(n)$ bits. Formally, we show that $CC_{r+1}(CRG_{L,0}(\mu)) \leq (r+1)\log n$ while for every constant $\epsilon < 1$ we have that $CC_{r/2-2}(CRG_{\ell,\epsilon}) \geq \min\{\Omega(\ell), n/\mathsf{poly}\log n\}$ (and similarly for SKG).

THEOREM 1.1. (ROUNDS TRADEOFF FOR CRG) *For all $\epsilon < 1, r \in \mathbb{Z}^+$, there exist $\eta > 0, n_0, \beta < \infty$, such that for all $n \geq n_0, L$ there exists a source $\mu_{r,n,L}$ for which the following hold:*

1. *There exists an $((r+1),(r+1)\lceil\log n\rceil)$-protocol for $(L,0)$-CRG from $\mu_{r,n,L}$.*
2. *For every $\ell \in \mathbb{Z}^+$ there is no $(r/2-2, \min\{\eta\ell - \beta, n/\log^{\beta} n\})$-protocol for (ℓ,ϵ)-CRG from $\mu_{r,n,L}$.*

We also get an analogous theorem for SKG, with the same source!

THEOREM 1.2. (ROUNDS TRADEOFF FOR SKG) *For all $\epsilon < 1, r \in \mathbb{Z}^+$, there exist $\eta > 0, n_0, \beta < \infty$, such that for all $n \geq n_0, L$ there exists a source $\mu_{r,n,L}$ for which the following hold:*

1. *There exists an $((r+1),(r+1)\lceil\log n\rceil)$-protocol for $(L,0)$-SKG from $\mu_{r,n,L}$.*
2. *For every $\ell \in \mathbb{Z}^+$ there is no $(r/2-2, \min\{\eta\ell - \beta, n/\log^{\beta} n\})$-protocol for (ℓ,ϵ)-SKG from $\mu_{r,n,L}$.*

In particular, our theorems yield a gap in the amount of communication that is almost *exponentially large* if the number of rounds of communication is squeezed by a constant factor. Note that every communication protocol can be converted to a two-round communication protocol with an exponential blowup in communication[2] – in this sense our bound is close to optimal. Prior to our work, no separations were known for any number of rounds larger than two!

1.4 Brief Overview of Construction and Proofs Our starting point for constructing the source μ is the well-known "pointer-chasing" problem [NW93] used to study tradeoffs between rounds of interaction and communication complexity. In (our variant of) this problem Alice and Bob get a series of permutations

[2] To see this claim, first notice that by incurring a constant factor blowup of the communication complexity, we may convert a c-bit communication protocol to a c round protocol in which Alice transmits one bit to Bob and Bob responds with one bit. We can then convert this to a two round protocol where Alice sends Bob at most 2^{c+1} bits specifying what bit she would send next for each possible partial transcript. From this information Bob can simulate the entire protocol and output the answer.

$\pi_1, \pi_2, \ldots, \pi_r : [n] \to [n]$ along with an initial pointer i_0 and their goal is to "chase" the pointers, i.e., compute i_r where $i_j = \pi_j(i_{j-1})$ for every $j \in \{1, \ldots, r\}$. Alice's input consists of the odd permutations $\pi_1, \pi_3, \ldots$, and Bob gets the initial pointer i_0 and the even permutations $\pi_2, \pi_4, \ldots$. The natural protocol to determine i_r takes $r+1$ rounds of communication with the jth round involving the message i_j (for $j = 0, \ldots, r$). Nisan and Wigderson show that any protocol with r rounds of interaction requires $\Omega(n)$ bits of communication [NW93].

To convert the pointer chasing instance into a correlated source, we let the source include $2n$ strings $A_1, \ldots, A_n$ and $B_1, \ldots, B_n \in \{0,1\}^L$ where $(A_1, \ldots, B_n)$ is uniform in $\{0,1\}^{2nL}$ conditioned on $A_{i_r} = B_{i_r}$. Thus the source outputs $X = (\pi_1, \pi_3, \ldots; A_1, \ldots, A_n)$ and $Y = (i_0, \pi_2, \pi_4, \ldots; B_1, \ldots, B_n)$ satisfy $A_{i_r} = B_{i_r}$ with $i_j = \pi_j(i_{j-1})$ for every $j \in \{1, \ldots, r\}$. (See Definition 2.1 and Figure 3 for more details.) The natural protocol for the pointer chasing problem also turns into a natural protocol for CRG and SKG with $r+1$ rounds of communication, and our challenge is to show that protocols with few rounds cannot extract randomness.

The lower bound does not follow immediately from the lower bound for the pointer chasing problem — and indeed we do not even give a lower bound for $r - O(1)$ rounds of communication. We explain some of the challenges here and how we overcome them.

Our first challenge is that there is a low-complexity "non-deterministic protocol" for common randomness generation in our setting. The players somehow guess i_r and then verify $A_{i_r} = B_{i_r}$ (by exchanging the first $\log 1/\epsilon$ bits of these strings) and if they do, then they output A_{i_r} and B_{i_r} respectively. While the existence of a non-deterministic protocol does not imply the existence of a deterministic one, it certainly poses hurdles to the lower bound proofs. Typical separations between non-deterministic communication complexity and deterministic ones involve lower bounds such as those for "set-disjointness" [KS92, Raz92, BJKS04] which involve different reasoning than the "round-elimination" arguments in [NW93]. Our lower bound would somehow need to combine the two approaches.

We manage to do so "modularly" at the expense of a factor of 2 in the number of rounds of communication by introducing an intermediate "pointer verification (PV)" problem. In this problem Alice and Bob get permutations $\pi_1, \ldots, \pi_r$ (with Alice getting the odd ones and Bob the even ones) and additionally Bob gets pointers i and j. Their goal is to decide if the final pointer i_r equals j given that the initial pointer i_0 is equal to i. The usefulness of this problem comes from the fact that we can reduce the common randomness generation problem to the complexity of the pointer verification problem on a specific (and natural) distribution: Specifically if PV is hard on this distribution with r' rounds of communication, then we can show (using the hardness of set disjointness as a black box) that the common randomness generation problem is hard with $r' - 1$ rounds of communication.

We thus turn to showing lower bounds for PV. Before turning to our proof technique, we note that decision problem variants of pointer-chasing similar to PV have previously been studied extensively in the literature including in [PRV01, GO13]. [GO13] proves randomized communication lower bounds for the multiplayer setting in which there are $2r$ players, each holding a single function, and they must determine in $r-1$ rounds whether the end pointers of each of two sequences of r functions are equal. The work of [PRV01] proves superlinear lower bounds for the two-player setting where the players must determine only the first bit of i_r. Crucially, these works prove lower bounds on communication by only considering distributions for which the players' inputs are independent (and [PRV01] only obtains lower bounds for *deterministic* communication complexity). However, none of these works give a lower bound for PV with respect to our distributions. In particular, in the distribution for PV that we consider, the players' inputs are correlated, and this poses significant technical challenges not encountered in these prior works.

Returning to our goal of showing lower bounds for PV, we first note that we cannot expect a lower bound for r rounds of communication: PV can obviously be solved in $r/2$ rounds of communication with Alice and Bob chasing both the initial and final pointers till they meet in the middle. We also note that one can use the lower bound from [NW93] as a black box to get a lower bound of $r/2-1$ rounds of communication for PV but it is no longer on the "natural" distribution we care about and thus this is not useful for our setting.

The bulk of the proof is thus devoted to proving an $r/2 - O(1)$ round lower bound for the PV problem on our distribution. We get this lower bound by roughly following the "round elimination" strategy of [NW93]. A significant challenge in extending these lower bounds to our case is that we have to deal with distributions where Alice and Bob's inputs are dependent. This should not be surprising since the CRG problem provides Alice and Bob with correlated inputs, and so there is resulting dependency between Alice and Bob even before any messages are sent. The dependency gets more complex as Alice and Bob exchange messages, and we need to ensure that the resulting mutual information is not correlated with the

desired output, i.e., the PV value of the game. We do so by a delicate collection of conditions that allow the inputs to be correlated while guaranteeing sufficient independence to carry out a round elimination proof. We provide details in the complete version of this paper [BGGS18].

Organization of Rest of the Paper. In Section 2, we present our construction of the distribution μ alluded to in Theorem 1.1 and Theorem 1.2. In Section 3 we reduce the task of proving communication lower bounds for CRG with few rounds to the task of proving lower bounds for distinguishing some distributions. We then introduce our final problem, the Pointer Verification problem, and the distribution on which we need to analyze it in Section 4. This section includes the statement of our main technical theorem about the pointer verification problem (Theorem 4.1) and the proofs of Theorem 1.1 and Theorem 1.2 assuming this theorem.

2 Construction

We start with some basic notation used in the rest of the paper. For any positive integer n, we denote by $[n]$ the set $\{1,\ldots,n\}$. We use log to denote the logarithm to the base 2. For a distribution D on a universe Ω we use the notation $X \sim D$ to denote a random variable X sampled according to D. For any positive integer t, we denote by D^t the distribution obtained by sampling t independent identically distributed samples from D. We use the notation $X \perp\!\!\!\perp Y$ to denote that X is independent of Y and $X \perp\!\!\!\perp Y|Z$ to denote that X and Y are independent conditioned on Z. For random variables X, Y, we use $\mathbb{1}[X = Y]$ to denote the random variable that is 1 when $X = Y$, and 0 otherwise. We denote by $\mathbb{E}_{X\sim D}[X]$ the expectation of X and for an event $E \subseteq \Omega$, we denote by $\Pr_X[E]$ the probability of the event E. For $i \in \Omega$, D_i (and sometimes $D(i)$) denotes the probability of the element i, i.e., $D_i = D(i) = \Pr_{X\sim D}[X = i]$. For distributions P and Q on Ω, the total variation distance $\Delta(P,Q) \stackrel{\text{def}}{=} \frac{1}{2}\sum_{i\in\Omega}|P_i - Q_i|$. The entropy of $X \sim P$ is the quantity $H(X) = \mathbb{E}_{X\sim P}[-\log P_X]$. The min-entropy of $X \sim P$ is the quantity $H_\infty(X) = \min_{x\in\Omega}\{-\log P_x\}$. For a pair of random variables $(X,Y) \sim P$, P_X denotes the marginal distribution on X and $P_{X|y}$ denotes the distribution of X conditioned on $Y = y$. The conditional entropy $H(X|Y) \stackrel{\text{def}}{=} \mathbb{E}_{y\sim P_Y}[H(X_y)]$, where $X_y \sim P_{X|Y=y}$. The mutual information between X and Y, denoted $I(X;Y)$, is the quantity $H(X) - H(X|Y)$. The conditional mutual information between X and Y conditioned on Z, denoted $I(X;Y|Z)$, is the quantity $\mathbb{E}_{z\sim P_Z}[H(X_z) - H(X_z|Y_z)]$ where $(X_z, Y_z) \sim P_{X,Y|Z=z}$. We use standard properties of entropy and information such as the Chain rules and the fact "conditioning does not increase entropy". For further background material on information theory and communication complexity, we refer the reader to the books [CT12] and [KN97] respectively.

We start by describing the family of distributions $\mu_{r,n,L}$ that we use to prove Theorem 1.1 and Theorem 1.2. For a positive integer n, we let S_n denote the family of all permutations of $[n]$.

DEFINITION 2.1. (THE SOURCE $\mu_{r,n,L}$) *For positive integers r, n and L, the support of $\mu = \mu_{r,n,L}$ is $(S_n^{\lceil r/2\rceil} \times \{0,1\}^{nL}) \times ([n] \times S_n^{\lfloor r/2\rfloor} \times \{0,1\}^{nL})$. Denoting $X = (\pi_1, \pi_3, \ldots, \pi_{2\lceil r/2\rceil - 1}, A_1, \ldots, A_n)$ and $Y = (i, \pi_2, \pi_4, \ldots, \pi_{2\lfloor r/2\rfloor}, B_1, \ldots, B_n)$, a sample $(X,Y) \sim \mu$ is drawn as follows:*

- *$i \in [n]$ and $\pi_1, \ldots, \pi_r \in S_n$ are sampled uniformly and independently.*
- *Let $j = \pi_r(\pi_{r-1}(\cdots \pi_1(i) \cdots))$.*
- *$A_j = B_j \in \{0,1\}^L$ is sampled uniformly and independently of i and π's.*
- *For every $k \neq j$, $A_k \in \{0,1\}^L$ and $B_k \in \{0,1\}^L$ are sampled uniformly and independently.*

See Figure 3 for an illustration of the inputs to the Pointer Chasing Source.

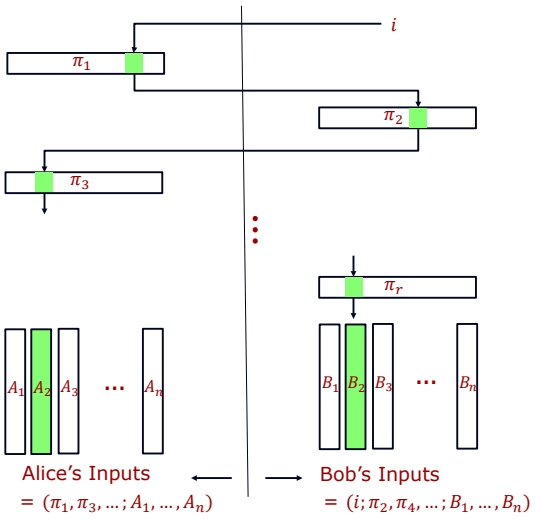

Figure 3: The Pointer Chasing Source

Informally, a sample from μ contains a common hidden block of randomness $A_j = B_j \in \{0,1\}^L$ that Alice and Bob can find by following a sequence of pointers, where Alice holds the odd pointers in the sequence and

Bob holds the even pointers. The next lemma gives (the obvious) upper bound on the r-round communication needed to generate common randomness from μ.

LEMMA 2.1. (UPPER BOUND FOR r-ROUND SKG) *For every r, n and L, there exists an $(r+1, \lceil \log n \rceil)$-protocol for $(L,0)$-SKG (and hence also for $(L,0)$-CRG) from $\mu_{r,n,L}$ with Bob speaking in the first round.*

Proof. The protocol Π is the obvious one in which Bob and Alice alternate by sending a pointer to each other starting with i and culminating in j, and the randomness they "agree on" is $A_j = B_j$.

Formally, for $t \in [r]$, let $i_t = \pi_t(i_{t-1})$ with $i_0 = i$. In odd round $t+1$, Bob sends i_t to Alice and in even round $t+1$, Alice sends i_t to Bob. At the end of $r+1$ rounds of communication Alice outputs A_{i_r} and Bob outputs B_{i_r}.

Note that by the construction of μ, we have that $i_r = j$ and $A_j = B_j$. Note further that at the beginning of the $(t+1)$st round of communication both Alice and Bob know i_{t-1}. Furthermore if $t+1$ is odd, then Bob also knows π_t and hence can compute $i_t = \pi_t(i_{t-1})$ (and similarly Alice knows her message in even rounds).

Thus we conclude that the above is a valid $(r+1, \lceil \log n \rceil)$-protocol for $(L,0)$-CRG. Furthermore since $A_{i_r} = B_{i_r}$ is independent of $i_0, \ldots, i_r$ it follows that $I((i_0, \ldots, i_r); A_{i_r}) = I(\Pi; K_A) = 0$ (and similarly for $I(\Pi; K_B)$) and so this is also a valid protocol $(L,0)$-SKG.

In the rest of the paper we show that no $r/2 - O(1)$ round protocol can solve CRG from $\mu_{r,n,L}$ with non-trivial communication.

3 Related Indistinguishability Problems

Our lower bound on the number of rounds needed to generate common randomness comes from an "indistinguishability argument". We show that to protocols with a small number of rounds and small amount of communication, the distribution μ is indistinguishable from the distribution $\mu_X \times \mu_Y$, where Alice and Bob's inputs are independent. Using the well-known fact that generating L bits of common randomness essentially requires L bits of communication in the absence of correlated inputs, this leads us to conclude that CRG is hard with limited number of rounds of communication.

In this section we simply set up the stage by defining the notion of indistinguishability and connecting it to the task of common randomness generation, leaving the task of proving the indistinguishability to later sections.

3.1 The Main Distributions and Indistinguishability Claims

We start by defining the indistinguishability of inputs to protocols.

DEFINITION 3.1. *We say that two distributions D_1 and D_2 on (X,Y) are* ϵ-indistinguishable *to a protocol Π if the distributions of transcripts (the sequence of messages exchanged by Alice and Bob) generated when $(X,Y) \sim D_1$ has total variation distance at most ϵ from the distribution of transcripts when $(X,Y) \sim D_2$.*

We say that distributions D_1 and D_2 are (ϵ, c, r)-indistinguishable *if they are ϵ-indistinguishable to every (r,c)-protocol Π using public randomness. Conversely, we say that the distributions D_1 and D_2 are* (ϵ, c, r)-distinguishable *if they are not (ϵ, c, r)-indistinguishable.*

Fix r, n, L and let $\mu = \mu_{r,n,L}$. Now let μ_X denote the marginal distribution of X under μ, i.e., $X = (\pi_1, \pi_3, \ldots, \pi_{2\lceil r/2 \rceil - 1}, A_1, \ldots, A_n)$ have all coordinates chosen independently and uniformly from their domains. Similarly let μ_Y denote the marginal on Y, and let $\mu_X \times \mu_Y$ denote the distribution where $X \sim \mu_X$ and $Y \sim \mu_Y$ are chosen independently.

Our main technical result (Theorem 4.1 and in particular its implication Lemma 4.1) shows that μ and $\mu_X \times \mu_Y$ are $(\epsilon, r/2 - O(1), n/\mathsf{poly}\log n)$-indistinguishable, even to protocols with common randomness. In the rest of this section, we explain why this rules out common randomness generation.

3.2 Reduction to Common Randomness Generation

PROPOSITION 3.1. *There exists a constant $\eta > 0$ such that for every r, r', n, L, ℓ, t and $\epsilon < 1$, there is no $(r', \eta\ell - 3/2 \cdot \log(1/1-\epsilon) - O(1))$-protocol for (ℓ, ϵ)-CRG from $\mu_X^t \times \mu_Y^t$, where $\mu = \mu_{r,n,L}$ with μ_X and μ_Y being its marginals.*

Proof. This is essentially folklore. For instance it follows immediately from [CGMS17, Theorem 2.6] using $\rho = 0$ (which corresponds to private-coin protocols). In particular, any $(r', \eta\ell - 3/2 \cdot \log(1/1-\epsilon) - O(1))$-protocol Π for (ℓ, ϵ)-CRG from $\mu_X^t \times \mu_Y^t$ gives a protocol with communication $\eta\ell - 3/2 \cdot \log(1/1-\epsilon) - O(1)$ for (ℓ, ϵ)-CRG with no inputs since Alice and Bob can draw μ_X^t and μ_Y^t independently using private randomness and then simulate Π.

PROPOSITION 3.2. *There is a sufficiently large absolute constant ξ such that the following holds. Let η be the constant from Proposition 3.1. If there exists an (r', c)-protocol that solves the $(\ell, 1-\gamma)$-CRG problem from $\mu = \mu_{r,n,L}$ with $c < \eta(\ell - 3) - 3/2 \cdot \log 1/\gamma - \xi$, then there exists some positive integer t for which μ^t and $\mu_X^t \times \mu_Y^t$ are $(\gamma/10, r'+1, c + \xi \log 1/\gamma)$-distinguishable.*

Proof. Let Π be an (r', c) protocol with private randomness for $(\ell, 1-\gamma)$-CRG from μ and let D_1 denote the distribution of K_A conditioned on $K_A = K_B$. Let t be the

number of samples of μ used by Π. Let $I = \mathbb{1}[K_A = K_B]$ be the indicator variable determining if $K_A = K_B$. Let D_1^A be the distribution of (K_A, I) when Π is run on samples from μ^t. Let D_2^A be the distribution of the (K_A, I) when Π is run on samples from $\mu_X^t \times \mu_Y^t$. Define D_1^B and D_2^B analogously. We distinguish between the cases where $\Delta(D_1^A, D_2^A)$ and $\Delta(D_1^B, D_2^B)$ are both small from the cases where one of them is large.

Case 1: D_1^A is $\gamma/4$-far from D_2^A (in total variation distance). We argue that in this case, μ^t and $\mu_X^t \times \mu_Y^t$ are distinguishable. Let T be the optimal distinguisher of D_1^A from D_2^A (i.e., T is a 0/1 valued function with $\mathbb{E}_{(K_A,I)\sim D_1^A}[T(K_A, I)] - \mathbb{E}_{(K_A,I)\sim D_2^A}[T(K_A, I)] \geq \gamma/4$). Let α denote $\mathbb{E}_{(K_A,I)\sim D_2^A}[T(K_A, I)]$. We now describe a protocol Π' which uses public randomness and augments Π by including a bit I' (which is usually equal to I) and $T(K_A, I')$ as part of the transcript. We consider two subcases: (1) If Bob is the last speaker in Π, then Π' executes Π and then at the conclusion of Π, Bob sends a random hash $h_B = h(K_B)$ which is $O(\log 1/\gamma)$ bits long (so that for $K_A \neq K_B$ we have $\Pr_h[h(K_A) = h(K_B)] \leq \gamma/20$). Alice then sends $I' = \mathbb{1}[h(K_A) = h_B]$ and the bit $b_{I'} = T(K_A, I')$. (2) If Alice is the last speaker in Π, then Π' executes Π and then Alice sends $h_A = h(K_A)$ to Bob, as well as $b_0 = T(K_A, 0)$ and $b_1 = T(K_A, 1)$. Bob then sends $I' = \mathbb{1}[h_A = h(K_B)]$ and $b_{I'}$.

Note that in both cases Π' has $r' + 1$ rounds of communication and the total number of bits of communucation is $c + O(\log 1/\gamma)$. We now show that Π' distinguishes μ^t from $\mu_X^t \times \mu_Y^t$ with probability $\Omega(\gamma)$. To see this note that $\Pr_{(K_A,I)\sim D_1^A}[b_{I'} = 1] \geq \Pr_{(K_A,I)\sim D_1^A}[T(K_A, I) = 1] - \Pr_h[I' \neq \mathbb{1}[K_A = K_B]] \geq (\alpha + \gamma/4) - \gamma/20 = \alpha + \gamma/5$. On the other hand we also have $\Pr_{(K_A,I)\sim D_2^A}[b_{I'} = 1] \leq \Pr_{(K_A,I)\sim D_2^A}[T(K_A, I) = 1] + \Pr_h[I' \neq \mathbb{1}[K_A = K_B]] \leq \alpha + \gamma/20$. We conclude that $\Pr_{(K_A,I)\sim D_1^A}[b_{I'} = 1] - \Pr_{(K_A,I)\sim D_2^A}[b_{I'} = 1] \geq \gamma/5 - \gamma/20 \geq \gamma/10$. And since $b_{I'}$ is a part of the transcript of Π' we conclude that the two distributions are $\gamma/10$-distinguished by Π'.

Case 2: D_1^B is $\gamma/4$-far from D_2^B. This is similar to the above and yields that μ^t and $\mu_X^t \times \mu_Y^t$ are $(\gamma/10, r' + 1, c + O(\log 1/\gamma))$-distinguishable.

Case 3: $\Delta(D_1^A, D_2^A) \leq \gamma/4$ and $\Delta(D_1^B, D_2^B) \leq \gamma/4$. We argue that this case can not happen since this allows a low-communication protocol to solve CRG with private randomness, thereby contradicting Proposition 3.1. The details are the following.

Our main idea here is to run Π on $\mu_X^t \times \mu_Y^t$ (which, being a product distribution involves only private randomness). The proximity of D_1^A to D_2^A implies that the probability that $K_A = K_B$ when Π is run on $\mu_X^t \times \mu_Y^t$ is at least $3\gamma/4$ (since the probability that $K_A = K_B$ on μ^t is at least γ and the probability that $I = \mathbb{1}[K_A = K_B]$ is different under μ^t than under $\mu_X^t \times \mu_Y^t$ is at most $\gamma/4$). But we are not done since the min-entropy of K_A or K_B when Π is run on $\mu_X^t \times \mu_Y^t$ might not be lower-bounded by ℓ. So we modify Π to get a protocol Π' as follows: Run Π and let (K_A, K_B) be the output of Π. (The output of Π' will be different as we see next.) If the probability of outputting K_A is more than $4 \cdot 2^{-\ell}$ then let K'_A be a uniformly random string in $\{0,1\}^\ell$, else let $K'_A = K_A$. Similarly if the probability of outputting K_B is more than $4 \cdot 2^{-\ell}$ then let K'_B be a uniformly random string in $\{0,1\}^\ell$, else let $K'_B = K_B$. (Note that when $K'_A \neq K_A$ then K'_A and K'_B are independent.) Let (K'_A, K'_B) be the outputs of Π'. We claim below that Π' solves the $(\ell-3, 1-\gamma/12)$-CRG from $\mu_X^t \times \mu_Y^t$ which contradicts Proposition 3.1 if $c < \eta(\ell-3) - 3/2 \cdot \log(12/\gamma) - \xi$ for a sufficiently large constant ξ. First note that by design the probability of outputting any fixed output k'_A is at most $4 \cdot 2^{-\ell} + 2^{-\ell} < 2^{-(\ell-3)}$. (If $\Pr[K_A = k'_A] \geq 4 \cdot 2^{-\ell}$ then $\Pr[K'_A = k'_A] \leq 2^{-\ell}$, else $\Pr[K'_A = k'_A] \leq \Pr[K_A = k'_A] + 2^{-\ell}$.) It remains to see that $\Pr[K'_A = K'_B] \geq \gamma/12$. First note that $\Pr[K_A \neq K'_A] \leq \gamma/3$. This is so since every k'_A such that $\Pr[K_A = k'_A] \geq 4 \cdot 2^{-\ell}$ contributes at least $\Pr[K_A = k'_A] - 2^{-\ell} \geq (3/4) \cdot \Pr[K_A = k'_A]$ to $\Delta(D_1^A, D_2^A)$ (the probability of k'_A on μ^t is at most $2^{-\ell}$). Thus using $\Delta(D_1^A, D_2^A) \leq \gamma/4$, we conclude $\Pr[K_A \neq K'_A] \leq (4/3)\Delta(D_1^A, D_2^A) \leq \gamma/3$. But now we have $\Pr[K'_A = K'_B] \geq \Pr[K_A = K_B] - (\Pr[K_A \neq K'_A] + \Pr[K_B \neq K'_B]) \geq 3\gamma/4 - 2\gamma/3 = \gamma/12$.

3.3 Reduction to the Case $t = 1$

Next we show that we can work with the case $t = 1$ without loss of generality. Roughly the intuition is that all permutations look the same, and so chasing one series of pointers $\pi_1, \ldots, \pi_r$ is not harder than chasing a sequence of t pointers of the form $(\pi'_{1,\tau} \ldots, \pi'_{r,\tau})_{\tau \in [t]}$. Informally, even if the players in latter problem are given the extra information $(\pi'_{\ell,\tau})^{-1}\pi_\ell$, for every $\ell \in [r]$ and $\tau \in [t]$, they still have to effectively chase the pointers $\pi_1, \ldots, \pi_r$. This intuition is formalized in the reduction below.

PROPOSITION 3.3. *Fix r, n, L and let $\mu = \mu_{r,n,L}$ and μ_X and μ_Y be its marginals. If there exists ϵ, r', c, t such that μ^t and $\mu_X^t \times \mu_Y^t$ are (ϵ, r', c)-distinguishable, then $\mu' = \mu_{r,n,Lt}$ and $(\mu')_X \times (\mu')_Y$ are (ϵ, r', c)-distinguishable.*

Proof. Suppose Π is a (r', c)-protocol that ϵ-distinguishes μ^t from $\mu_X^t \times \mu_Y^t$. We show how to distinguish μ' from $(\mu')_X \times (\mu')_Y$ using Π. Let (X, Y) be an instance of the μ' vs. $(\mu')_X \times (\mu')_Y$ distinguishability problem. We now show how Alice and Bob can use common randomness to generate $(X'_1, Y'_1), \ldots, (X'_t, Y'_t)$ such

that $((X'_1, Y'_1), \ldots, (X'_t, Y'_t)) \sim \mu^t$ if $(X, Y) \sim \mu'$ and $((X'_1, Y'_1), \ldots, (X'_t, Y'_t)) \sim \mu_X^t \times \mu_Y^t$ if $(X, Y) \sim \mu'_X \times \mu'_Y$. It follows that by applying Π to $((X'_1, Y'_1), \ldots, (X'_t, Y'_t))$, Alice and Bob can distinguish μ' from $\mu'_X \times \mu'_Y$.

Let $X = (\pi_1, \pi_3, \ldots, \pi_{2\lceil r/2 \rceil - 1}, A_1, \ldots, A_n)$ and $Y = (i, \pi_2, \pi_4, \ldots, \pi_{2\lfloor r/2 \rfloor}, B_1, \ldots, B_n)$, where $\pi_\ell \in S_n$ and $A_k, B_k \in \{0,1\}^{Lt}$. Further, let $A_k = A_{k,1} \circ \cdots \circ A_{k,t}$ and $B_k = B_{k,1} \circ \cdots \circ B_{k,t}$ where $A_{k,\tau}, B_{k,\tau} \in \{0,1\}^L$ and $\circ$ denotes concatenation. Alice and Bob use their common randomness to generate permutations $\sigma_{\ell,\tau}$, for $\ell \in \{0, \ldots, r\}$ and $\tau \in [t]$, uniformly and independently from S_n. Now let $\pi'_{\ell,\tau} = \sigma_{\ell\tau} \cdot \pi_\ell \cdot \sigma^{-1}_{\ell-1,\tau}$. Let $i'_\tau = \sigma_{0,\tau}(i)$. And let $A'_{k,\tau} = A_{\sigma_{r,\tau}(k),\tau}$ and $B'_{k,\tau} = B_{\sigma_{r,\tau}(k),\tau}$. Finally, let $X'_\tau = (\pi'_{1,\tau}, \pi'_{3,\tau}, \ldots, \pi'_{2\lceil r/2 \rceil - 1,\tau}, A'_{1,\tau}, \ldots, A'_{n,\tau})$ and $Y'_\tau = (i'_\tau, \pi'_{2,\tau}, \pi'_{4,\tau}, \ldots, \pi'_{2\lfloor r/2 \rfloor,\tau}, B'_{1,\tau}, \ldots, B'_{n,\tau})$. We claim that this sequence (X'_τ, Y'_τ) has the claimed properties.

First note that the permutations $\pi'_{\ell,\tau}$ are uniform and independent from S_n due to the fact that the $\sigma_{\ell,\tau}$'s are uniform and independent. Similarly i'_τ's are uniform and independent of the $\pi'_{\ell,\tau}$s. If $(X, Y) \sim \mu'_X \times \mu'_Y$ then the $A'_{k,\tau}$'s and $B'_{k,\tau}$'s are also uniform and independent of i's and π''s, estabilishing that $((X'_1, Y'_1), \ldots, (X'_t, Y'_t)) \sim \mu_X^t \times \mu_Y^t$ if $(X, Y) \sim \mu'_X \times \mu'_Y$. If $(X, Y) \sim \mu'$ then note that $j'_\tau = \pi'_{r,\tau}(\cdots(\pi'_{1,\tau}(i'_\tau))) = \sigma_{r,\tau}(\pi_r(\cdots(\pi_1(i)))) = \sigma_{r,\tau}(j)$. We thus have that $A'_{j'_\tau,\tau} = A_{j,\tau} = B_{j,\tau} = B'_{j'_\tau}$ and otherwise the $A'_{k,\tau}$'s and $B'_{k,\tau}$'s are uniform and independent. This establishes that $((X'_1, Y'_1), \ldots, (X'_t, Y'_t)) \sim \mu^t$ if $(X, Y) \sim \mu'$, and thus the proposition is proved.

4 The Pointer Verification Problem

When L is very large compared to n, there are two possible natural options for trying to distinguish μ from $\mu_X \times \mu_Y$. One option is for Alice and Bob to ignore the pointers $(\pi_1, \ldots, \pi_r)$ and simply try to see if there exists $j \in [n]$ such that $A_j = B_j$. The second option is for Alice and Bob to ignore the $A's$ and the $B's$ while communicating and simply try to find the end of the chain of pointers $i_0 = i, \ldots, i_\ell = \pi_\ell(i_{\ell-1}), \ldots, i_r$ and then check to see if $A_{i_r} = B_{i_r}$.

The former turns out to be a problem that is at least as hard as Set Disjointness on n bit inputs (and so requires $\Omega(n)$ bits of communication). The latter requires $\tilde{\Omega}(n)$ bits of communication with fewer than r rounds. But combining the two lower bounds seems like a non-trivial challenge. In this section we introduce an intermediate problem, that we call the pointer verification (PV) problem, that allows us to modularly use lower bounds on the set disjointness problem and on the (small-round) communication complexity of PV, to prove that μ is indistinguishable from $\mu_X \times \mu_Y$.

The main difference between PV and pointer chasing is that here Alice and Bob are given both a source pointer i_0 and a target pointer j_0 and simply need to decide if chasing pointers from i_0 leads to j_0. We note that the problem is definitely easier than pointer chasing in that for a sequence of r pointers, Alice and Bob can decide PV in $r/2$ rounds (by "chasing i_0 forward and j_0 backwards simultaneously"). This leads us to a bound that is weaker in the round complexity by a factor of 2, but allows us the modularity alluded to above. Finally the bulk of the paper is devoted to proving a communication lower bound for $r/2 - O(1)$ round protocols for solving PV (or rather again, an indistinguishability result for two distributions related to PV). This lower bound is similar to the lower bound of Nisan and Wigderson [NW93] though the proofs are more complex due to the fact that we need to reason about settings where Alice's input and Bob's input are correlated.

We start with the definition of a distributional version of the Pointer Verification Problem and then relate it to the complexity of distinguishing μ from $\mu_X \times \mu_Y$.

DEFINITION 4.1. *For integers r and n with r being odd, the distributions $D^{\mathrm{Y}}_{\mathrm{PV}} = D^{\mathrm{Y}}_{\mathrm{PV}}(r,n)$ and $D^{\mathrm{N}}_{\mathrm{PV}} = D^{\mathrm{N}}_{\mathrm{PV}}(r,n)$ are supported on $((S_n^{\lceil r/2 \rceil}) \times ([n]^2 \times S_n^{\lfloor r/2 \rfloor})$. $D^{\mathrm{N}}_{\mathrm{PV}}$ is just the uniform distribution over this domain. On the other hand, $(X, Y) \sim D^{\mathrm{Y}}_{\mathrm{PV}}$ is sampled as follows: Sample $\pi_1, \ldots, \pi_r$ uniformly and independently from S_n and further sample $i_0 \in [n]$ uniformly and independently. Finally let $j_0 = \pi_r(\cdots(\pi_1(i_0)))$, and let $X = (\pi_1, \pi_3, \ldots, \pi_r)$ and $Y = (i_0, j_0, \pi_2, \pi_4, \ldots, \pi_{r-1})$.*

Our main theorem about Pointer Verification is the following:

THEOREM 4.1. *For every $\epsilon > 0$ and odd r there exists β, n_0 such for every $n \geq n_0$, $D^{\mathrm{Y}}_{\mathrm{PV}}(r,n)$ and $D^{\mathrm{N}}_{\mathrm{PV}}(r,n)$ are $(\epsilon, (r-1)/2, n/\log^\beta n)$-indistinguishable.*

The proof of Theorem 4.1 is provided in the full version of our paper [BGGS18]. We now show that this suffices to prove our main theorem. First we prove in Lemma 4.1 below that μ is indistinguishable from $\mu_X \times \mu_Y$. This proof uses the theorem above, and the fact that set disjointness cannot be solved with $o(n)$ bits of communication, that we recall next.

THEOREM 4.2. ([RAZ92]) *For every $\epsilon > 0$ there exists $\delta > 0$ such that for all n the following holds: Let $\mathrm{Disj}_{\mathrm{Y}}$, respectively $\mathrm{Disj}_{\mathrm{N}}$, be the uniform distribution on pairs (U, V) with $U, V \subseteq [n]$ and $|U| = |V| = n/4$ such that $|U \cap V| = 1$ (respectively $|U \cap V| = 0$). Then $\mathrm{Disj}_{\mathrm{Y}}$ and $\mathrm{Disj}_{\mathrm{N}}$ are $(\epsilon, \delta n, \delta n)$-indistinguishable to Alice and Bob, if Alice gets U and Bob gets V as inputs.*

REMARK 4.1. *We note that the theorem in [Raz92] explicitly only rules out* $(1-\epsilon_0, \Omega(n), \Omega(n))$*-distinguishability of* Disj_Y *and* Disj_N *for some* $\epsilon_0 > 0$*. But we note that the distinguishability gap of any protocol can be amplified in this case (even though we are in the setting of distributional complexity) since by applying a random permutation to* $[n]$*, Alice and Bob can simulate independent inputs from* Disj_Y *(or* Disj_N*) given any one input from its support. Thus an* (r, c) *protocol that* ϵ*-distinguishes* Disj_Y *from* Disj_N *can be converted to an* $(r, (c/\epsilon^2)\log(1/\epsilon_0))$*-protocol that* $(1-\epsilon_0)$*-distinguishes* Disj_Y *from* Disj_N*, implying the version of the theorem above.*

LEMMA 4.1. *For every* $\epsilon > 0$ *and odd* r *there exists* β, n_0 *such for every* $n \geq n_0$ *and* L*, the distributions* $\mu = \mu_{r,n,L}$ *and* $\mu_X \times \mu_Y$ *are* $(2\epsilon, (r-1)/2, n/\log^\beta n)$ *indistinguishable.*

Proof. We use a new distribution μ_{mid} which is a hybrid of μ and $\mu_X \times \mu_Y$ where $(X, Y) \sim \mu_{\mathrm{mid}}$ is sampled as follows: Sample $\pi_1, \ldots, \pi_r \in S_n$ independently and uniformly. Further sample $i, j \in [n]$ uniformly and independently (of each other and the π's). Finally sample $A_j = B_j \in \{0,1\}^L$ uniformly and A_{-j} and B_{-j} uniformly and independently from $\{0,1\}^{(n-1)L}$. Let $X = (\pi_1, \pi_3, \ldots, \pi_r, A_1, \ldots, A_n)$ and $Y = (i, \pi_2, \pi_4, \ldots, \pi_{r-1}, B_1, \ldots, B_n)$. (So μ_{mid} does force a correlation between A and B, but the permutations do not lead to this correlated point.)

We show below that μ_{mid} and $\mu_X \times \mu_Y$ are indistinguishable to low-communication protocols (due to the hardness of Set Disjointness), while μ and μ_{mid} are indistinguishable to low-round low-communication protocols, due to Theorem 4.1. The lemma follows by the triangle inequality for indistinguishability (which follows from the triangle inequality for total variation distance).

We now use the fact (Theorem 4.2) that disjointness is hard, and in particular $o(n)$-bit protocols cannot distinguish between $(U, V) \sim \mathrm{Disj}_Y$ and $(U, V) \sim \mathrm{Disj}_N$. Note in particular that Disj_Y is supported on pairs (U, V) such that $U \cap V = \{j\}$ where $j \in [n]$ is distributed uniformly. Specifically, we have that for every $\epsilon > 0$ there exists $\delta > 0$ such that Disj_Y and Disj_N are $(\epsilon, \delta n, \delta n)$-indistinguishable.

We now show how to reduce the above to the task of distinguishing μ_{mid} and $\mu_X \times \mu_Y$ (using shared randomness and no communication). Alice and Bob share $W_1, \ldots, W_n \in \{0,1\}^L$ distributed uniformly and independently. Given $U \subseteq [n]$, Alice picks $\pi_1, \pi_3, \ldots$ uniformly and independently, lets $A_\ell = W_\ell$ if $\ell \in U$ and samples $A_\ell \in \{0,1\}^L$ uniformly otherwise, and lets $X = (\pi_1, \pi_3, \ldots, \pi_r, A_1, \ldots, A_n)$. Similarly Bob samples $i \in [n]$ uniformly, and $\pi_2, \pi_4, \ldots, \pi_{r-1} \in S_n$ uniformly and independently. Let $B_\ell = W_\ell$ if $\ell \in V$ and let B_ℓ be drawn uniformly from $\{0,1\}^L$ otherwise. Let $Y = (i, \pi_2, \pi_4, \ldots, \pi_{r-1}, B_1, \ldots, B_n)$. It can be verified that $(X, Y) \sim \mu_{\mathrm{mid}}$ if $(U, V) \sim \mathrm{Disj}_Y$ and $(X, Y) \sim \mu_X \times \mu_Y$ if $(U, V) \sim \mathrm{Disj}_N$. Thus we conclude that μ_{mid} and $\mu_X \times \mu_Y$ are $(\epsilon, \delta n, \delta n)$-indistinguishable.

Next we turn to the (in)distinguishability of μ vs. μ_{mid}. We reduce the task of distinguishing D^Y_{PV} and D^N_{PV} to distinguishing μ and μ_{mid}. Given an instance (X, Y) of pointer verification with $X = (\pi_1, \pi_3, \ldots, \pi_r)$ and $Y = (i, j, \pi_2, \pi_4, \ldots, \pi_{r-1})$, we generate an instance (X', Y') as follows: Let $W_1, \ldots, W_n$ be uniformly and independently chosen elements of $\{0,1\}^L$ shared by Alice and Bob. Alice lets $A_\ell = W_\ell$ for every ℓ and lets $X' = (\pi_1, \ldots, \pi_r, A_1, \ldots, A_n)$. Bob lets $B_j = W_j$ and samples B_ℓ uniformly and independently for $\ell \in [n] - \{j\}$, and lets $Y' = (i, \pi_2, \ldots, \pi_{r-1}, B_1, \ldots, B_n)$. It can be verified that $(X', Y') \sim \mu$ if $(X, Y) \sim D^Y_{\mathrm{PV}}$ and $(X', Y') \sim \mu_{\mathrm{mid}}$ if $(X, Y) \sim D^N_{\mathrm{PV}}$. It follows from Theorem 4.1 that μ and μ_{mid} are $(\epsilon, (r-1)/2, n/\log^\beta n)$-indistinguishable.

Combining the two we get that μ and $\mu_X \times \mu_Y$ are $(2\epsilon, (r-1)/2, n/\log^\beta n)$-indistinguishable (assuming $(r-1)/2 < \delta n$ and $n/\log^\beta n < \delta n$, which are both true for sufficiently large n).

We are ready to prove Theorem 1.1, which says that we cannot generate ℓ bits of common randomness from $\mu_{r,n,L}$ in $r/2-2$ rounds using only $\min(O(\ell), n/\log^\beta n)$ communication.

Proof. (of Theorem 1.1) We start with the case of odd r. We use the distribution $\mu = \mu_{r,n,L}$ in this case. Part (1) of the theorem which says that one can generate common randomness using an $(r+1, r+1\lceil \log n \rceil)$ protocol, follows from Lemma 2.1. Part (2) of Theorem 1.1 claims that using $r/2$ rounds and insufficient communication one cannot generate common randomness. This follows by combining Lemma 4.1 with Proposition 3.3 and Proposition 3.2. In particular, let η be the constant from Proposition 3.2 (and also Proposition 3.1), ξ be the constant from Proposition 3.2, and β_0 be the constant β from Lemma 4.1 given the number of rounds r and $(1-\epsilon)/40$ for the variational distance parameter. Finally let β be a constant such that $\beta > \max\{\beta_0, 3\eta + 3/2 \cdot \log 1/(1-\epsilon) + \xi\}$ and $n/\log^\beta n + \xi \log 1/(1-\epsilon) \leq n/\log^{\beta_0} n$, which is possible for sufficiently large n. Suppose for the purpose of contradiction that for some $\ell \in \mathbb{Z}^+$, there were a $((r-3)/2, \min\{\eta\ell - \beta, n/\log^\beta n\})$-protocol for (ℓ, ϵ)-CRG from $\mu_{r,n,L}$. By Proposition 3.2, there is some positive integer t for which μ^t and $\mu^t_X \times \mu^t_Y$ are $((1-\epsilon)/10, (r-1)/2, \min\{\eta\ell, n/\log^\beta n\} + \xi \log 1/(1-\epsilon))$-distinguishable. But now let $\mu' = \mu_{r,n,Lt}$. Then

by Proposition 3.3 and our assumption on β, μ' and $(\mu')_X \times (\mu')_Y$ are $((1-\epsilon)/10, (r-1)/2, n/\log^{\beta_0} n)$-distinguishable. But this contradicts Lemma 4.1, which states that μ' and $(\mu')_X \times (\mu')_Y$ are $((1-\epsilon)/20, (r-1)/2, n/\log^{\beta_0} n)$-indistinguishable.

For even r, we just use the distribution $\mu_{r-1,n,L}$. Part (1) continues to follow from Lemma 2.1. And for Part (2) we can reason as above, with the caveat that the bound on round complexity from Lemma 4.1 now is "only" $((r-1)-1)/2$. The additional loss from Proposition 3.2 is one more round, leading to a final lower bound of $r/2-2$.

Proof. (of Theorem 1.2) Part (1) of the theorem follows from Lemma 2.1. Part (2) follows from Part (2) of Theorem 1.1 since SKG is a strictly harder task.

References

[AC93] Rudolf Ahlswede and Imre Csiszár. Common randomness in information theory and cryptography. part i: secret sharing. *IEEE Transactions on Information Theory*, 39(4), 1993. 2

[AC98] Rudolf Ahlswede and Imre Csiszár. Common randomness in information theory and cryptography. ii. cr capacity. *Information Theory, IEEE Transactions on*, 44(1):225–240, 1998. 2

[AG76] Rudolf Ahlswede and Peter Gács. Spreading of sets in product spaces and hypercontraction of the markov operator. *The annals of probability*, pages 925–939, 1976. 2

[AGKN13] Venkat Anantharam, Amin Gohari, Sudeep Kamath, and Chandra Nair. On maximal correlation, hypercontractivity, and the data processing inequality studied by Erkip and Cover. *arXiv preprint arXiv:1304.6133*, 2013. 2

[BBCR13] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. How to compress interactive communication. *SIAM Journal on Computing*, 42(3):1327–1363, 2013. 3

[BGGS18] Mitali Bafna, Badih Ghazi, Noah Golowich, and Madhu Sudan. Communication-rounds tradeoffs for common randomness and secret key generation. *arXiv:1808.08907*, 2018. 5, 8

[BGI14] Mohammad Bavarian, Dmitry Gavinsky, and Tsuyoshi Ito. On the role of shared randomness in simultaneous communication. In *Automata, Languages, and Programming*, pages 150–162. Springer, 2014. 3

[BJKS04] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *Journal of Computer and System Sciences*, 68(4):702–732, 2004. 3, 4

[BM11] Andrej Bogdanov and Elchanan Mossel. On extracting common random bits from correlated sources. *Information Theory, IEEE Transactions on*, 57(10):6351–6355, 2011. 2

[CGMS17] Clément L Canonne, Venkatesan Guruswami, Raghu Meka, and Madhu Sudan. Communication with imperfectly shared randomness. *IEEE Transactions on Information Theory*, 63(10):6799–6818, 2017. 2, 3, 6

[CMN14] Siu On Chan, Elchanan Mossel, and Joe Neeman. On extracting common random bits from correlated sources on large alphabets. *Information Theory, IEEE Transactions on*, 60(3):1630–1637, 2014. 2

[CN00] Imre Csiszár and Prakash Narayan. Common randomness and secret key generation with a helper. *Information Theory, IEEE Transactions on*, 46(2):344–366, 2000. 2

[CN04] Imre Csiszár and Prakash Narayan. Secrecy capacities for multiple terminals. *IEEE Transactions on Information Theory*, 50(12):3047–3061, 2004. 2

[CSWY01] Amit Chakrabarti, Yaoyun Shi, Anthony Wirth, and Andrew Yao. Informational complexity and the direct sum problem for simultaneous message complexity. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pages 270–278. IEEE, 2001. 3

[CT12] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012. 5

[DMN18] Anindya De, Elchanan Mossel, and Joe Neeman. Non interactive simulation of correlated distributions is decidable. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2728–2746. SIAM, 2018. 3

[GJ18] Badih Ghazi and TS Jayram. Resource-efficient common randomness and secret-key schemes. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1834–1853. Society for Industrial and Applied Mathematics, 2018. 2, 3

[GK73] Peter Gács and János Körner. Common information is far less than mutual information. *Problems of Control and Information Theory*, 2(2):149–162, 1973. 2, 3

[GKR17] Badih Ghazi, Pritish Kamath, and Prasad Raghavendra. Dimension reduction for polynomials over gaussian space and applications. *arXiv preprint arXiv:1708.03808*, 2017. 3

[GKS16a] Badih Ghazi, Pritish Kamath, and Madhu Sudan. Communication complexity of permutation-invariant functions. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1902–1921, 2016. 3

[GKS16b] Badih Ghazi, Pritish Kamath, and Madhu Sudan. Decidability of non-interactive simulation of joint distributions. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 545–554. IEEE, 2016. 3

[GO13] Venkatesan Guruswami and Krzysztof Onak. Superlinear lower bounds for multipass graph processing. In *2013 IEEE Conference on Computational Complex-*

ity, pages 287–298. IEEE, 2013. 4

[GR16] Venkatesan Guruswami and Jaikumar Radhakrishnan. Tight bounds for communication-assisted agreement distillation. In *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan*, pages 6:1–6:17, 2016. 2

[KA15] Sudeep Kamath and Venkat Anantharam. On non-interactive simulation of joint distributions. *arXiv preprint arXiv:1505.00769*, 2015. 3

[KN97] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997. 5

[KS92] Bala Kalyanasundaram and Georg Schintger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5(4):545–557, 1992. 4

[LCV15] Jingbo Liu, Paul Cuff, and Sergio Verdú. Secret key generation with one communicator and a one-shot converse via hypercontractivity. In *2015 IEEE International Symposium on Information Theory (ISIT)*, pages 710–714. IEEE, 2015. 2, 3

[LCV16] Jingbo Liu, Paul W. Cuff, and Sergio Verdú. Common randomness and key generation with limited interaction. *CoRR*, abs/1601.00899, 2016. 2, 3

[Liu16] Jingbo Liu. Rate region for interactive key generation and common randomness generation. *Manuscript available at* http://www.princeton.edu/~jingbo/preprints/RateRegionInteractiveKeyGen120415.pdf (visited on 02/13/2017), 2016. 3

[LLG+05] Daihyun Lim, Jae W Lee, Blaise Gassend, G Edward Suh, Marten Van Dijk, and Srinivas Devadas. Extracting secret keys from integrated circuits. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 13(10):1200–1205, 2005. 2

[Mau93] Ueli M Maurer. Secret key agreement by public discussion from common information. *Information Theory, IEEE Transactions on*, 39(3):733–742, 1993. 2

[MO04] Elchanan Mossel and Ryan O'Donnell. Coin flipping from a cosmic source: On error correction of truly random bits. *arXiv preprint math/0406504*, 2004. 3

[MOR+06] Elchanan Mossel, Ryan O'Donnell, Oded Regev, Jeffrey E Steif, and Benny Sudakov. Non-interactive correlation distillation, inhomogeneous markov chains, and the reverse bonami-beckner inequality. *Israel Journal of Mathematics*, 154(1):299–336, 2006. 3

[NW93] Noam Nisan and Avi Wigderson. Rounds in communication complexity revisited. *SIAM Journal on Computing*, 22(1):211–219, 1993. 3, 4, 8

[PRV01] Stephen J. Ponzio, Jaikumar Radhakrishnan, and S. Venkatesh. The communication complexity of pointer chasing. *Journal of Computer and System Sciences*, 62:323–355, 2001. 4

[Raz92] Alexander A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, 1992. 4, 8, 9

[SD07] G Edward Suh and Srinivas Devadas. Physical unclonable functions for device authentication and secret key generation. In *Proceedings of the 44th annual Design Automation Conference*, pages 9–14. ACM, 2007. 2

[Sha49] Claude E Shannon. Communication theory of secrecy systems. *Bell system technical journal*, 28(4):656–715, 1949. 2

[SHO08] Ying Su, Jeremy Holleman, and Brian P Otis. A digital 1.6 pj/bit chip identification circuit using process variations. *IEEE Journal of Solid-State Circuits*, 43(1):69–77, 2008. 2

[Tya13] Himanshu Tyagi. Common information and secret key capacity. *IEEE Transactions on Information Theory*, 59(9):5627–5640, 2013. 2, 3

[Wit75] Hans S Witsenhausen. On sequences of pairs of dependent random variables. *SIAM Journal on Applied Mathematics*, 28(1):100–113, 1975. 3

[Wyn75] Aaron D. Wyner. The common information of two dependent random variables. *IEEE Transactions on Information Theory*, 21(2):163–179, 1975. 2, 3

[YLH+09] Haile Yu, Philip Heng Wai Leong, Heiko Hinkelmann, L Moller, Manfred Glesner, and Peter Zipf. Towards a unique FPGA-based identification circuit using process variations. In *2009 International Conference on Field Programmable Logic and Applications*, pages 397–402. IEEE, 2009. 2

[ZC11] Lei Zhao and Yeow-Kiang Chia. The efficiency of common randomness generation. In *2011 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2011. 2