Placing Conditional Disclosure of Secrets in the Communication Complexity Universe

Benny Applebaum[1]
Tel Aviv University, Tel Aviv, Israel
https://www.eng.tau.ac.il/~bennyap/
benny.applebaum@gmail.com

Prashant Nalini Vasudevan[2]
UC Berkeley, Berkeley, USA
http://people.eecs.berkeley.edu/~prashvas
prashvas@berkeley.edu

Abstract

In the *conditional disclosure of secrets* (CDS) problem (Gertner et al., J. Comput. Syst. Sci., 2000) Alice and Bob, who hold n-bit inputs x and y respectively, wish to release a common secret z to Carol (who knows both x and y) if and only if the input (x, y) satisfies some predefined predicate f. Alice and Bob are allowed to send a single message to Carol which may depend on their inputs and some shared randomness, and the goal is to minimize the communication complexity while providing information-theoretic security.

Despite the growing interest in this model, very few lower-bounds are known. In this paper, we relate the CDS complexity of a predicate f to its communication complexity under various communication games. For several basic predicates our results yield tight, or almost tight, lower-bounds of $\Omega(n)$ or $\Omega(n^{1-\epsilon})$, providing an exponential improvement over previous logarithmic lower-bounds.

We also define new communication complexity classes that correspond to different variants of the CDS model and study the relations between them and their complements. Notably, we show that allowing for imperfect correctness can significantly reduce communication – a seemingly new phenomenon in the context of information-theoretic cryptography. Finally, our results show that proving explicit super-logarithmic lower-bounds for imperfect CDS protocols is a necessary step towards proving explicit lower-bounds against the class AM, or even $\mathrm{AM} \cap \mathrm{coAM}$ – a well known open problem in the theory of communication complexity. Thus imperfect CDS forms a new minimal class which is placed just beyond the boundaries of the "civilized" part of the communication complexity world for which explicit lower-bounds are known.

2012 ACM Subject Classification Theory of computation → Communication complexity, Theory of computation → Cryptographic protocols

Keywords and phrases Conditional Disclosure of Secrets, Information-Theoretic Security

Digital Object Identifier 10.4230/LIPIcs.ITCS.2019.4

Related Version The full version of this paper is available as [6].

[1] Supported by the European Union's Horizon 2020 Programme (ERC-StG-2014-2020) under grant agreement no. 639813 ERC-CLC, by an ICRC grant and by the Check Point Institute for Information Security.

[2] This work was done in part while the author was visiting Tel Aviv University. Supported in part by NSF Grant CNS-1350619, and by the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226 and W911NF-15-C-0236.

Acknowledgements We thank Pritish Kamath and Hoeteck Wee for helpful discussions, and Mika Göös for helpful pointers. We also thank the anonymous reviewers for their useful comments.

1 Introduction

Understanding the communication complexity of information-theoretically secure protocols is a fundamental research problem. Despite much effort, we have very little understanding of the communication complexity of even simple cryptographic tasks, and for most models, there are exponentially large gaps between the best known upper-bounds and the best known lower-bounds. In an attempt to simplify the problem, one may try to focus on the most basic settings with a minimal non-trivial number of players (say two or three) and the simplest possible communication pattern (e.g., single message protocols). Different cryptographic tasks have been studied in this minimal setting, including secure computation [17], and non-interactive zero-knowledge proofs [23]. In this paper we will focus on what seems to be the simplest task in this model: *Conditional Disclosure of Secrets* (CDS) [20].[3]

Conditional Disclosure of Secrets

Consider a pair of computationally unbounded parties, Alice and Bob, each holding an input, $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ respectively, to some public predicate $f : \mathcal{X} \times \mathcal{Y} \to \{0, 1\}$. Alice and Bob also hold a joint secret z (say a single bit) and have access to a joint source of randomness $r \xleftarrow{R} \mathcal{R}$. The parties wish to disclose the secret z to a third party, Carol, if and only if the predicate $f(x, y)$ evaluates to 1. To this end, Alice and Bob should each send a single message $a = a(x, z; r)$ and $b = b(y, z; r)$ to Carol. Based on the transcript (a, b) and the inputs (x, y), Carol should be able to recover the secret z if and only if $f(x, y) = 1$. (Note that Carol is assumed to know x and y.) That is, we require two properties:

- *Correctness*: There exists a deterministic decoder algorithm Dec that recovers z from (x, y, a, b) with high probability whenever x, y is a 1-input (i.e., $f(x, y) = 1$);
- *Privacy*: For every fixed 0-input (x, y) (for which the predicate evaluates to 0), regardless of the value of the secret z, the joint distribution of the transcript (a, b), induced by a choice of the shared randomness, is statistically close (up to some small deviation error) to some canonical distribution $\mathsf{Sim}(x, y)$.

The main complexity measure of CDS protocols is their communication complexity which is taken to be the total bit-length of the messages a and b. (See Figure 1 for a schematic view and Section A for formal definitions.)

Apart from being a natural basic notion, CDS has turned out to be a useful primitive with various applications in the context of private information retrieval (PIR) [20], secure multiparty computation [1, 25], secret sharing schemes [14, 15, 36, 31, 11, 2, 29], and attribute-based encryption [7, 37]. Correspondingly, the communication complexity of CDS was extensively studied in the last few years.

[3] While we do not wish to define the notions from [17] and [23], let us just mention that the complexity of a function in these two models upper-bounds the complexity in the CDS model [20, 5]. In this sense, CDS may be considered as being simpler.

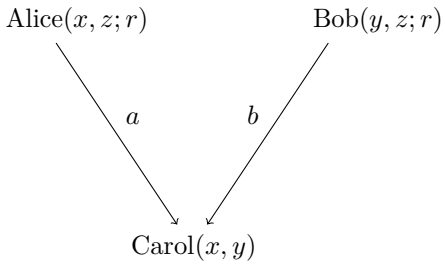

Figure 1 Schematic of a CDS protocol.

Upper bounds

On the positive side, it is known that the CDS complexity of a predicate f is at most linear in the formula complexity of f [20]. This result was extended to other (presumably stronger) computational models such as (arithmetic) branching programs [26], and (arithmetic) span programs [5]. The latter paper also shows that the CDS complexity of f is at most linear in the complexity of any zero-information Arthur Merlin (ZAM) protocol for f. (The ZAM model, introduced by [23], adds a zero-knowledge property to the standard AM communication complexity model.)[4] In a recent breakthrough, Liu, Vaikuntanathan and Wee [30] showed that the CDS complexity of any predicate $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ over n-bit inputs is at most $2^{\widetilde{O}(\sqrt{n})}$, improving over the exponential upper-bound of $O(2^{n/2})$ from [10]. Applebaum et al. [3] showed that when the secret is very long (exponential in the size of the domain of the predicate) the overhead per each bit of z can be reduced to $O(n)$; a constant-rate solution (in which the total communication is $O(|z|)$) was recently given in [2].

The quest for lower bounds

On the lower-bound front much less is known. While we have tight lower bounds for restricted forms of CDS (e.g., when the computations are restricted to linear functions [19, 9, 12]), only few, relatively weak, lower-bounds are known for general CDS. It is important to note that an insecure solution to the problem has a communication cost of 1 bit! (Let Alice send the secret in the clear regardless of her input.) Hence, any super-constant lower-bound is, in a sense, non-trivial. Indeed, unlike the case of standard communication games for which communication lower-bounds are based on the correctness properties of the protocol, the challenge here is to somehow capture the additional cost of *privacy*.

The first super-constant lower-bound was proved by Gay, Kerenidis, and Wee [19].

▶ **Theorem 1** ([19]). *For every predicate $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$,*

$$\mathrm{CDS}(f) \geq \Omega(\log(\mathrm{R}_{A\to B}(f) + \mathrm{R}_{B\to A}(f))),$$

where $\mathrm{R}_{A\to B}(f)$ denotes the one-way randomized communication complexity of f, and $\mathrm{CDS}(f)$ denotes the minimal communication complexity of a CDS protocol for f with privacy and correctness error of 0.1.[5]

[4] The theorem of [5] actually relates the communication and randomness complexity of CDS for f to the randomness and communication complexity of a ZAM protocol for the complement of f. However, using our results in this paper one can conclude that the CDS communication of f is at most linear in the ZAM communication of f.

[5] The theorem was originally proved for perfect CDS, however, the proof generalizes to the imperfect case (see [3]).

For n-bits predicates, Theorem 1 leads, at best, to a logarithmic lower-bound of $\Omega(\log n)$. Applebaum *et al.*. [3] showed that this bound is essentially tight: There are (partial) functions whose randomized communication complexity is exponentially larger than their CDS complexity. They also proved a linear n-bit lower-bound for a random (non-explicit) n-bit predicate $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$. An explicit version of this result was proved by [4].

▶ **Theorem 2** ([4]). *For every non-degenerate predicate*[6] $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ *whose largest 0-monochromatic rectangle is of size at most* L*,*

$$\mathrm{pCDS}(f) \geq \log \frac{|f^{-1}(0)|}{L} - \log \frac{|\mathcal{X} \times \mathcal{Y}|}{|f^{-1}(0)|} - 1 = 2\log|f^{-1}(0)| - \log|\mathcal{X}| - \log|\mathcal{Y}| - \log L - 1,$$

where $\mathrm{pCDS}(f)$ *denotes the minimal communication complexity of a CDS protocol for* f *with perfect privacy and perfect correctness.*

The theorem is effective for predicates whose communication matrix is rich in zeroes, and at the same time avoids large zero-monochromatic rectangles. In particular, for mod-2 inner product over n-bit inputs, we get a tight lower-bound of $n - O(1)$ and for Set-Intersection a lower-bound of $\Omega(n)$. Unfortunately, the theorem is not robust to errors, leaving the imperfect CDS complexity of these predicates wide open. Moreover, for many basic predicates the theorem does not even give logarithmic bounds either due to the lack of many zeroes (e.g., the Not-Equal predicate) or due to the existence of huge zero-rectangles (e.g., the Greater-Than predicate).

This paper

Theorems 1 and 2 provide a very partial picture, and fall short of proving meaningful and robust lower-bounds for many basic predicates, such as Not-equal, Greater-Than, Intersection, and Index.[7] We believe that a full understanding of these simple cases is necessary for the more ambitious goal of proving stronger lower bounds. Our goal in this paper is to remedy the situation by providing new lower-bound techniques. Specifically, we enrich our lower-bound toolbox by relating the CDS complexity of a function to its communication complexity under various communication games. Our results provide simple, yet effective, ways to leverage privacy to construct communication protocols. They lead to new lower-bounds for perfect and imperfect CDS protocols, and allow us to establish new results regarding the relations between different variants of the CDS model.

2 Our Contribution

2.1 Perfectly-correct CDS and coNP Games

Our first theorem relates the complexity of any perfectly-correct CDS protocol for f to the non-deterministic communication complexity of f's complement.

▶ **Theorem 3.** *For every predicate* $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$*,*

$$\mathrm{pcCDS}(f) \geq \Omega(\mathrm{coNP}(f)) - O(\log(n)),$$

where n *denotes the total input length of* f*, and* $\mathrm{pcCDS}(f)$ *denotes the minimal communication complexity of a CDS protocol for* f *with perfect correctness and privacy error of* 0.1*.*

[6] A predicate is non-degenerate if for every fixing of $x \in \mathcal{X}$ the residual function $f(x,\cdot)$ is not the constant zero function.

[7] Apart of being basic examples, these predicates are motivated by some of the applications of CDS.

Proof idea

To prove the theorem, we first show that the coNP complexity is upper-bounded by the randomness complexity of the CDS, and then prove that one can always assume that the randomness complexity is comparable to the communication complexity via a new sparsification lemma (similar to that of Newman [33]). The first part relies on the following simple observation: In order to convince Alice and Bob that $f(x,y)$ evaluates to zero it suffices to prove that the joint distribution of the CDS messages for zero-secret, $(a(x,z=0;r),b(y,z=0;r))$, induced by a random choice of r, and the joint distribution of the messages for one-secret $(a(x,z=1;r),b(y,z=1;r))$, are *not disjoint*. A prover can prove this statement by sending to Alice and Bob a pair of strings r_0 and r_1 for which $(a(x,z=0;r_0),b(y,z=0;r_0))$ equals to $(a(x,z=1;r_1),b(y,z=1;r_1))$. (See full version [6] for details.)

Despite its simplicity, this theorem is quite powerful. In particular, ignoring the constants in the Omega-notation and the logarithmic loss, the bound provided by Theorem 3 subsumes the lower-bound of Theorem 2 from [4]. Indeed, the latter lower-bound is at most the logarithm of the ratio between the zero-mass of f and its largest zero-monochromatic rectangle – a quantity that cannot be larger than the non-deterministic communication complexity of the complement of f (i.e., coNP(f)). Moreover, our new theorem can be applied to predicates that have only few zero entries or to predicates with huge zero-rectangles, for which Theorem 2 becomes meaningless. For example, by plugging-in classical coNP lower-bounds, we settle the complexity of the not-equal predicate with respect to perfectly correct CDS protocols.

▶ **Corollary 4.** *Let* $\mathsf{NEQ}_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ *denote the not-equal predicate which evaluates to* 1 *if and only if* $x \neq y$*. Then,*

$$\text{pcCDS}(\mathsf{NEQ}_n) \geq \Omega(n).$$

Similar tight linear lower-bounds can be obtained for the pcCDS complexity of the Greater-Than predicate, the Set-Intersection predicate, and the Inner-Product predicate. Previously, we had no super-logarithmic lower bounds that tolerate privacy error. (As already mentioned, for Greater-Than and NEQ_n, we did not have such bounds even for perfect CDS protocols.)

pcCDS is not closed under complement

Interestingly, the *equality* function EQ_n has a very succinct perfect CDS protocol: Use the shared randomness to sample a pair-wise independent hash function $h : \{0,1\}^n \to \{0,1\}$, and let Alice output $h(x)$ and Bob output $h(y) \oplus z$. The protocol has a minimal communication complexity of 2 and randomness complexity of $O(n)$. (The latter can be reduced to $O(\log n)$ by using an almost pair-wise independent hash function and settling for a constant privacy error.) This yields a strong separation between the complexity of a predicate and its complement with respect to perfectly-correct perfectly-private CDS protocols (pCDS).

▶ **Corollary 5.** $\text{pCDS}(\mathsf{EQ}_n) = 2$ *whereas* $\text{pCDS}(\mathsf{NEQ}_n) \geq \text{pcCDS}(\mathsf{NEQ}_n) \geq \Omega(n)$*. In particular, the classes* pCDS *and* pcCDS *are not closed under complement.*[8]

Transformations from CDS protocols for f to its complement were studied in [3]. The resulting protocols either introduce a privacy error or suffer from a communication overhead that grows polynomially with the randomness complexity of the original protocol. The NEQ_n example shows that at least one of these losses is inherent.

[8] We follow the standard communication complexity terminology and write pCDS to denote the class of predicates that admit a pCDS protocol whose complexity is polylogarithmic in the input length. A similar convention will be used throughout the paper for all other variants of the CDS model.

The benefit of decoding errors

The results of [3] (together with our randomness sparsification lemma) show that imperfect CDS is closed under complement. This general result leads to a polylogarithmic CDS protocol for NEQ_n with imperfect privacy and imperfect correctness, providing a surprising separation between general imperfect CDS protocols and ones which have perfect correctness. In fact, it is not hard to directly design a CDS protocol for NEQ_n with *constant communication*, *perfect privacy*, and constant correctness error. (See the full version [6] for a more general statement.) This leads to the following stronger separation.

▶ **Corollary 6.** *There is an n-bit predicate f for which* $\mathrm{pcCDS}(f) = \Omega(n)$ *and* $\mathrm{ppCDS}(f) = O(1)$*, where* $\mathrm{ppCDS}(f)$ *denotes the minimal communication complexity of a CDS protocol for f with perfect privacy and correctness error of* 0.1*. In particular,*

$$\mathrm{ppCDS} \not\subseteq \mathrm{pcCDS}.$$

As pointed to us by Hoteck Wee, Corollary 6 provides a rare example for an information-theoretic secure protocol that can significantly benefit from a small correctness error. This phenomena seems new in the context of information-theoretic secure cryptography, and is worth further exploration.[9]

2.2 Perfectly-Private CDS and PP Games

Our next goal is to lower-bound the complexity of CDS protocols with correctness errors. We begin with the case of perfectly private protocols.

▶ **Theorem 7.** *For every predicate $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$,*

$$\mathrm{ppCDS}(f) \geq \Omega(\mathrm{PP}(f)) - O(\log(n)),$$

where n denotes the total input length of f, and $\mathrm{ppCDS}(f)$ *denotes the minimal communication complexity of a CDS protocol for f with perfect privacy and correctness error of* 0.1*.*

The complexity measure $\mathrm{PP}(f)$ essentially corresponds to the sum of the communication complexity and number of private random bits used by a communication protocol that computes f correctly with probability more than $1/2$, where shared randomness is not allowed. (See the full version [6] for a formal definition.) The discrepancy method implies that the PP complexity of the mod-2 inner-product predicate IP_n is $\Omega(n)$ (cf. [28]) and so we get the following.

▶ **Corollary 8.** *Let $\mathsf{IP}_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ denote the inner-product predicate on n-bit inputs. Then,*

$$\mathrm{ppCDS}(\mathsf{IP}_n) \geq \Omega(n).$$

This is the first linear lower-bound on CDS with imperfect correctness. (Previous arguments fail to achieve such a result even for a non-explicit predicate.)

[9] Compare this, for example, to Shannon's classical lower-bound for perfectly-secure one-time symmetric encryption [35] in which a constant decryption error has a minor effect on the key/ciphertext length [16].

Proof idea

In order to prove Theorem 7, we turn a ppCDS protocol into a PP protocol. Loosely speaking, the idea is to construct a randomized protocol that accepts the input (x, y) based on collisions between random CDS transcripts that correspond to a zero-secret and random CDS transcripts that correspond to a one-secret. This idea, which was employed in the query setting by [13], leads to the desired result. (Details appear in the full version [6].)

2.3 Imperfect CDS, Interactive Proofs, and Zero Knowledge

We move on to the most general case of imperfect CDS protocols with both constant privacy error and correctness error. We show that the complexity of such protocols is at least polynomial in the AM communication complexity of f. (The latter class is the communication complexity analogue of Arthur-Merlin proofs.)

▶ **Theorem 9.** *There exists some universal constant $c > 0$, such that for any Boolean function f it holds that*

$$\mathrm{CDS}(f) \geq \mathrm{AM}(f)^c - \mathsf{polylog}(n),$$

where n denotes the total input length of f, and $\mathrm{CDS}(f)$ denotes the minimal communication complexity of a CDS protocol for f with correctness and privacy errors of 0.1.

Since (imperfect) CDS is closed under complement (by [3, Theorem 2] and [6, Lemma 1]), it holds that $\mathrm{CDS}(\bar{f}) \leq \mathsf{poly}(\mathrm{CDS}(f))$, and so we conclude the following.

▶ **Corollary 10.** *There exists some universal constant $c > 0$, such that for any Boolean function f it holds that*

$$\mathrm{CDS}(f) \geq \max(\mathrm{AM}(f), \mathrm{coAM}(f))^c - \mathsf{polylog}(n),$$

where n denotes the total input length of f.

Explicit CDS lower-bounds?

Corollary 10 can be used to show that the CDS complexity of most n-bit predicates must be at least polynomial in n, even when the protocol is imperfect. Unfortunately, it falls short of providing explicit lower-bounds; Finding an explicit function outside $\mathrm{AM} \cap \mathrm{coAM}$ is a central open problem in the theory of communication complexity. In fact, $\mathrm{AM} \cap \mathrm{coAM}$ forms a minimal class for which no explicit lower-bounds are known [24]. Corollary 10 places CDS as a weaker (and perhaps more accessible) target for explicit lower-bounds.

Proof idea

To prove Theorem 9 we show that a CDS protocol can be transformed into a constant-round private-coins interactive-proof. Then, we note that, just like in the computational setting, such interactive proofs can be converted to an AM protocol with polynomial overhead [8, 22].[10] The first step is obtained by imitating the standard interactive proof of Graph Nonisomorphism [21]. Indeed, the AM protocol constructed in Theorem 9 turns out to satisfy a *statistical zero-knowledge* property; That is, the view of Alice and Bob can be simulated via a low complexity 2-party randomized protocol. (See the full version [6] for details.)

[10] This reduction has a polynomial dependency in the randomness. In order to avoid such an overhead in the final statement, we prove a randomness sparsification lemma for constant-round interactive protocols. This requires some care due to the use of private coins.

CDS vs. SZK

Recall that, by definition, a CDS protocol yields a (distributed mapping) from the input (x, y) and the secret z to a distribution D_z over the transcript (a, b) such that the distributions, D_0 and D_1, are either statistically-close or statistically-far depending on the value of $f(x, y)$. This resembles the Statistical Difference problem [34], which is known to be complete for the computational complexity class SZK (consisting of problems that have interactive proofs that are statistically zero-knowledge). One may therefore hope to prove that in the communication complexity setting CDS complexity is characterized by SZK complexity. As already mentioned, Theorem 9 actually shows that CDS $\subseteq$ SZK, however, we do not know whether the reverse direction holds. Roughly speaking, such a result faces two obstacles. Firstly, the completeness result from [34] has an overhead that depends on the randomness complexity of the protocol, and we do not know how to get rid of this dependency. (In particular, it is not clear how to prove a proper sparsification lemma for SZK without sacrificing the zero-knowledge property.) Secondly, even if the randomness complexity is small, we do not know how to obtain a CDS protocol without allowing some interaction between Alice and Bob. Indeed, in the full version [6] we show that SZK$' \subseteq$ CDS$'$ where the "prime" version of SZK charges randomness towards the total complexity and the "prime" version of CDS allows short interaction between Alice and Bob. The problem of proving that SZK $\subseteq$ CDS (and therefore SZK = CDS) remains as an interesting open problem.

The results described so far are summarised in Figure 2, which shows the relationship between perfect and imperfect CDS and various measures from communication complexity. In Table 1, we list the current state of knowledge of the various CDS complexities of a number of commonly studied predicates. (See Section 3.)

2.4 Asymmetry in CDS and One-Way Communication

We shift gears, and turn to study the communication tradeoffs between Alice's and Bob's messages. Suppose that Alice's message is restricted to a short string of length t_A. Can we prove that Bob's message must be very long? We prove such tradeoffs based on the one-way randomized communication complexity of f.

▶ **Theorem 11.** *In any perfectly correct* 0.1*-private* CDS *protocol for f in which Alice and Bob communicate t_A and t_B bits respectively and the total input length of the function is n, it holds that:*

$$2^{t_A}(t_A + t_B + \log n) \geq \Omega(\mathrm{R}_{B\to A}(f)).$$

(In fact, the result holds even if one considers one-way randomized protocols that err only over zero inputs.) Recall that Theorem 1 (which is from [19]) shows that the total communication complexity $t_A + t_B$ is at least logarithmic in $(\mathrm{R}_{A\to B}(f) + \mathrm{R}_{B\to A}(f))$, which is tight for some predicates [3]. Theorem 11 provides a more accurate picture. If the total communication complexity is dominated by t_A, then one gets a logarithmic bound, similar to Theorem 1; however, when t_A is small (e.g., constant), we get a strong linear lower-bound of

$$t_B = \Omega(\mathrm{R}_{B\to A}(f)) - O(\log n).$$

In fact, when $\mathrm{R}_{B\to A}(f) = \Omega(n)$, for any constant $\alpha < 1$ if $t_A \leq \alpha \log n$ then

$$t_B = \Omega(n^{1-\alpha}).$$

Concretely, consider the Index_n predicate in which Bob holds an n-bit database $x \in \{0,1\}^n$ and Alice holds an index $i \in [n]$ (encoded as a string of length $\log n$) and the output is the i-th bit of x. Since $\mathrm{R}_{B\to A}(\mathsf{Index}_n) = \Omega(n)$ [27] we get:

▶ **Corollary 12.** *In any perfectly correct* 0.1*-private* CDS *protocol for* Index_n *in which Alice communicates at most* $\alpha \log n + O(1)$ *bits for some constant* $0 \leq \alpha < 1$*, the database owner, Bob, must communicate at least* $\Omega(n^{1-\alpha})$ *bits.*

Similar results can be obtained for predicates like Greater-Than, Set-Disjointness and Set-Intersection, based on classical lower-bounds for randomized one-way communication complexity (cf. [32, 27]).

The Index_n predicate plays an important role in CDS constructions and applications. First, it is complete for CDS in the sense that any n-bit predicate can be reduced to Index_N for $N = 2^n$. Indeed, the best known general CDS protocols were obtained by improving the pCDS complexity of Index [30]. In addition, a CDS for the index function can be viewed as a one-time version of the well-studied notion of *Broadcast Encryption*, and the lower-bound of Corollary 12 becomes especially appealing under this framework. Details follow.

Broadcast Encryption [18]

Suppose that we have a single sender and n receivers. The sender has a private encryption key r and each receiver $i \in [n]$ has its own private decryption key k_i. All the keys were collectively generated and distributed in an offline phase. In an online phase, the sender gets a message z together with a public list of authorized users $y \subseteq [n]$, represented by an n-bit characteristic vector. The sender should broadcast a ciphertext $b = b(y, z; r)$ to all the receivers (who also know y) so that an *authorized* receiver will be able to decrypt the ciphertext, and an unauthorized (computationally unbounded) receiver will learn nothing about the message z. The goal is to minimize the length of the ciphertext b, and the length of the keys k_i.

Information-theoretic one-time secure Broadcast Encryption turns to be equivalent to the CDS problem with respect to the Index_n predicate: Identify the ciphertext with Bob's message $b = b(y, z; r)$ and the i-th key with Alice's message $a(i; r)$.[11] The problem can be solved with n-bit ciphertext and 1-bit keys, and with 1-bit ciphertext and n-bit keys. In fact, [19] showed that one can smoothly get any tradeoff as long as the product of the ciphertext length and the key length is n. Corollary 12 shows that when the key-length is sub-logarithmic the ciphertext must be almost linear, confirming a conjecture of Wee [38].

Proof idea (of Theorem 11)

The idea is to let Bob send to Alice a pair of random strings r_0 and r_1 that are mapped to the same Bob's message b under the zero-secret and under the one-secret respectively. Alice then uses the string r_z and the secret z to compute a corresponding message a_z, and accepts if the zero message a_0 equals to the one message a_1. Perfect correctness guarantees that Alice will never err on 0-inputs. We further show that, when $f(x, y) = 1$, Alice accepts with probability which is at least inverse-exponential in her message length (up to a loss that is proportional to the privacy error of the protocol). See the full version [6] for details.

[11] Here we assume that we have a CDS in which only Bob holds the secret. However, any CDS can be transformed into this form with an additional communication cost of $O(|z|) = O(1)$.

Table 1 The CDS complexity of some simple functions. By definition, an upper-bound in the leftmost column (pCDS) implies an upper-bound in all other columns, and a lower-bound in the rightmost column (CDS) implies a lower-bound in all other columns. All the linear upper-bounds for pCDS follow from the fact that all of these predicates can be computed by a linear-size formula. The logarithmic lower-bounds for CDS follow from Theorem 1 (and the fact that the corresponding predicates have linear randomized one-way communication complexity.) The linear lower-bounds for pcCDS and ppCDS follow from Theorems 3 and 7 respectively.

Predicate	pCDS	pcCDS	ppCDS	CDS
Equality	$\Theta(1)$	$\Theta(1)$	$\Theta(1)$	$\Theta(1)$
Non-Equality	$\Theta(n)$	$\Theta(n)$	$\Theta(1)$	$\Theta(1)$
Inner-Product	$\Theta(n)$	$\Theta(n)$	$\Theta(n)$	$O(n)$ & $\Omega(\log n)$
Greater-Than	$\Theta(n)$	$\Theta(n)$	$O(n)$ & $\Omega(\log n)$	$O(n)$ & $\Omega(\log n)$
Set-Intersection	$\Theta(n)$	$\Theta(n)$	$O(n)$ & $\Omega(\log n)$	$O(n)$ & $\Omega(\log n)$
Set-Disjointness	$O(n)$ & $\Omega(\log n)$	$O(n)$ & $\Omega(\log n)$	$O(n)$ & $\Omega(\log n)$	$O(n)$ & $\Omega(\log n)$

3 Conclusion and Open Questions

In this paper we studied the relations between CDS protocols and standard communication complexity games. We established new connections between CDS communication complexity (with perfect and imperfect privacy and correctness) to well-known communication complexity measures for non-deterministic protocols, randomized unbounded-error protocols, and one-way protocols. This leads to new CDS bounds for various simple functions. These results are summarized in Figure 2 and Table 1.

We end by listing the immediate interesting questions left open following our work.

1. Prove an explicit polynomial lower-bound on (imperfect) CDS complexity. (A natural candidate would be Inner-Product.)
2. Our current ppCDS lower-bounds are based on PP complexity, which corresponds to discrepancy. Can we derive such bounds on weaker, easier-to-establish, properties? In particular, can we prove non-trivial ppCDS lower-bounds for predicates that have low randomized bounded-error communication complexity like Greater-Than?
3. Unlike all the other communication complexity measures considered here, CDS complexity is not necessarily upper-bounded by the length of the inputs. But we have no super-linear (or even linear with a large constant factor) lower-bounds for even perfect CDS protocols. Can any of the existing lower-bound techniques from communication complexity be used to obtain such bounds?
4. If not, can this difficulty be explained, perhaps by relating the problem of proving such lower bounds for CDS to more well-studied problems that are still unsolved?
5. Following the paradigm of lifting query complexity lower bounds to the communication setting, is there a natural query complexity measure that can be lifted to CDS complexity?
6. One simple predicate that has eluded all our bounds is Set-Disjointness, for which the best (imperfect) CDS protocol we know has $O(n)$ complexity, and the best lower bound we can prove, even for perfect CDS, is $\Omega(\log(n))$. Are either of these tight?

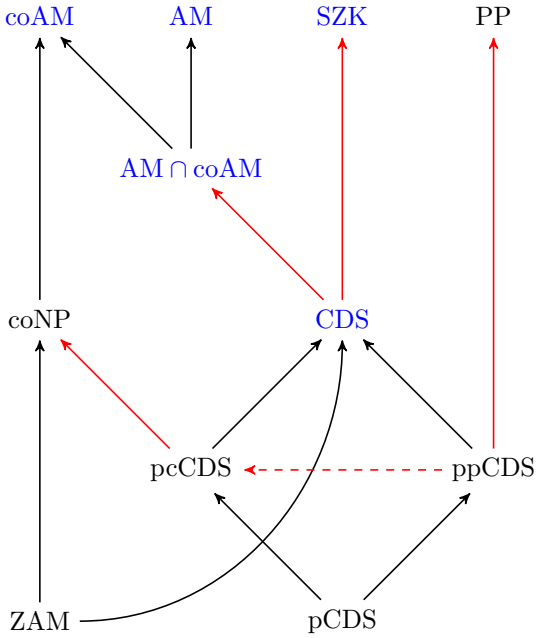

Figure 2 As is standard, we use the name of a complexity measure to also denote the class of functions with $\mathsf{polylog}(n)$ complexity under the measure. For classes C_1 and C_2, a solid arrow $C_1 \to C_2$ indicates that $C_1 \subseteq C_2$, and a dashed arrow $C_1 \dashrightarrow C_2$ indicates that $C_1 \not\subseteq C_2$. Red arrows indicate new results from this paper. Blue text indicates classes for which explicit bounds are not known.

References

1. William Aiello, Yuval Ishai, and Omer Reingold. Priced Oblivious Transfer: How to Sell Digital Goods. In Birgit Pfitzmann, editor, *Advances in Cryptology - EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, Innsbruck, Austria, May 6-10, 2001, Proceeding*, volume 2045 of *Lecture Notes in Computer Science*, pages 119–135. Springer, 2001. doi:10.1007/3-540-44987-6_8.
2. Benny Applebaum and Barak Arkis. On the Power of Amortization in Secret Sharing: d-Uniform Secret Sharing and CDS with Constant Information Rate. In Amos Beimel and Stefan Dziembowski, editors, *Theory of Cryptography - 16th International Conference, TCC 2018, Panaji, India, November 11-14, 2018, Proceedings, Part I*, volume 11239 of *Lecture Notes in Computer Science*, pages 317–344. Springer, 2018. doi:10.1007/978-3-030-03807-6_12.
3. Benny Applebaum, Barak Arkis, Pavel Raykov, and Prashant Nalini Vasudevan. Conditional Disclosure of Secrets: Amplification, Closure, Amortization, Lower-Bounds, and Separations. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 727–757. Springer, 2017. doi:10.1007/978-3-319-63688-7_24.
4. Benny Applebaum, Thomas Holenstein, Manoj Mishra, and Ofer Shayevitz. The Communication Complexity of Private Simultaneous Messages, Revisited. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 261–286. Springer, 2018. doi:10.1007/978-3-319-78375-8_9.

5 Benny Applebaum and Pavel Raykov. From Private Simultaneous Messages to Zero-Information Arthur-Merlin Protocols and Back. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part II*, pages 65–82, 2016. doi:10.1007/978-3-662-49099-0_3.

6 Benny Applebaum and Prashant Nalini Vasudevan. Placing conditional disclosure of secrets in the communication complexity universe. *Technical Report*, 2018. Full version of this paper. URL: https://www.eng.tau.ac.il/~bennyap/publications.html.

7 Nuttapong Attrapadung. Dual System Encryption via Doubly Selective Security: Framework, Fully Secure Functional Encryption for Regular Languages, and More. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 557–577. Springer, 2014. doi:10.1007/978-3-642-55220-5_31.

8 László Babai. Trading Group Theory for Randomness. In Robert Sedgewick, editor, *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, May 6-8, 1985, Providence, Rhode Island, USA*, pages 421–429. ACM, 1985. doi:10.1145/22145.22192.

9 Amos Beimel, Oriol Farràs, Yuval Mintz, and Naty Peter. Linear Secret-Sharing Schemes for Forbidden Graph Access Structures. To appear in TCC 2017, 2017.

10 Amos Beimel, Yuval Ishai, Ranjit Kumaresan, and Eyal Kushilevitz. On the Cryptographic Complexity of the Worst Functions. In Yehuda Lindell, editor, *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings*, volume 8349 of *Lecture Notes in Computer Science*, pages 317–342. Springer, 2014. doi:10.1007/978-3-642-54242-8_14.

11 Amos Beimel, Eyal Kushilevitz, and Pnina Nissim. The Complexity of Multiparty PSM Protocols and Related Models. To appear in Eurocrypt 2018, 2018. URL: https://eprint.iacr.org/2018/148.

12 Amos Beimel and Naty Peter. Optimal Linear Multiparty Conditional Disclosure of Secrets Protocols. Cryptology ePrint Archive, Report 2018/441, 2018. URL: https://eprint.iacr.org/2018/441.

13 Adam Bouland, Lijie Chen, Dhiraj Holden, Justin Thaler, and Prashant Nalini Vasudevan. On the Power of Statistical Zero Knowledge. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 708–719. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.71.

14 Ernest F. Brickell and Daniel M. Davenport. On the Classification of Ideal Secret Sharing Schemes. *J. Cryptology*, 4(2):123–134, 1991. doi:10.1007/BF00196772.

15 Renato M. Capocelli, Alfredo De Santis, Luisa Gargano, and Ugo Vaccaro. On the Size of Shares for Secret Sharing Schemes. *J. Cryptology*, 6(3):157–167, 1993. doi:10.1007/BF00198463.

16 Yevgeniy Dodis. Shannon Impossibility, Revisited. In Adam D. Smith, editor, *Information Theoretic Security - 6th International Conference, ICITS 2012, Montreal, QC, Canada, August 15-17, 2012. Proceedings*, volume 7412 of *Lecture Notes in Computer Science*, pages 100–110. Springer, 2012. doi:10.1007/978-3-642-32284-6_6.

17 Uriel Feige, Joe Kilian, and Moni Naor. A minimal model for secure computation (extended abstract). In Frank Thomson Leighton and Michael T. Goodrich, editors, *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, pages 554–563. ACM, 1994. doi:10.1145/195058.195408.

18 Amos Fiat and Moni Naor. Broadcast Encryption. In Douglas R. Stinson, editor, *Advances in Cryptology - CRYPTO '93, 13th Annual International Cryptology Conference, Santa Barbara, California, USA, August 22-26, 1993, Proceedings*, volume 773 of *Lecture Notes in Computer Science*, pages 480–491. Springer, 1993. doi:10.1007/3-540-48329-2_40.

19 Romain Gay, Iordanis Kerenidis, and Hoeteck Wee. Communication Complexity of Conditional Disclosure of Secrets and Attribute-Based Encryption. In Rosario Gennaro and Matthew Robshaw, editors, *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part II*, volume 9216 of *Lecture Notes in Computer Science*, pages 485–502. Springer, 2015. doi:10.1007/978-3-662-48000-7_24.

20 Yael Gertner, Yuval Ishai, Eyal Kushilevitz, and Tal Malkin. Protecting Data Privacy in Private Information Retrieval Schemes. *J. Comput. Syst. Sci.*, 60(3):592–629, 2000. doi:10.1006/jcss.1999.1689.

21 Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that Yield Nothing But Their Validity for All Languages in NP Have Zero-Knowledge Proof Systems. *J. ACM*, 38(3):691–729, 1991. doi:10.1145/116825.116852.

22 Shafi Goldwasser and Michael Sipser. Private Coins versus Public Coins in Interactive Proof Systems. *Advances in Computing Research*, 5:73–90, 1989.

23 Mika Göös, Toniann Pitassi, and Thomas Watson. Zero-Information Protocols and Unambiguity in Arthur-Merlin Communication. In Tim Roughgarden, editor, *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science, ITCS 2015, Rehovot, Israel, January 11-13, 2015*, pages 113–122. ACM, 2015. doi:10.1145/2688073.2688074.

24 Mika Göös, Toniann Pitassi, and Thomas Watson. The Landscape of Communication Complexity Classes. *Computational Complexity*, 27(2):245–304, 2018. doi:10.1007/s00037-018-0166-6.

25 Yuval Ishai, Eyal Kushilevitz, and Anat Paskin. Secure Multiparty Computation with Minimal Interaction. In Tal Rabin, editor, *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, volume 6223 of *Lecture Notes in Computer Science*, pages 577–594. Springer, 2010. doi:10.1007/978-3-642-14623-7_31.

26 Yuval Ishai and Hoeteck Wee. Partial Garbling Schemes and Their Applications. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*, pages 650–662. Springer, 2014. doi:10.1007/978-3-662-43948-7_54.

27 Ilan Kremer, Noam Nisan, and Dana Ron. On Randomized One-Round Communication Complexity. *Computational Complexity*, 8(1):21–49, 1999. doi:10.1007/s000370050018.

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

29 Tianren Liu and Vinod Vaikuntanathan. Breaking the Circuit-Size Barrier in Secret Sharing. To appear in STOC2018, 2018. URL: https://eprint.iacr.org/2018/333.

30 Tianren Liu, Vinod Vaikuntanathan, and Hoeteck Wee. Conditional Disclosure of Secrets via Non-linear Reconstruction. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 758–790. Springer, 2017. doi:10.1007/978-3-319-63688-7_25.

31 Tianren Liu, Vinod Vaikuntanathan, and Hoeteck Wee. Towards Breaking the Exponential Barrier for General Secret Sharing. To appear in Eurocrypt 2018, 2017. URL: https://eprint.iacr.org/2017/1062.

32 Peter Bro Miltersen, Noam Nisan, Shmuel Safra, and Avi Wigderson. On Data Structures and Asymmetric Communication Complexity. *J. Comput. Syst. Sci.*, 57(1):37–49, 1998. doi:10.1006/jcss.1998.1577.

33 Ilan Newman. Private vs. Common Random Bits in Communication Complexity. *Inf. Process. Lett.*, 39(2):67–71, 1991. doi:10.1016/0020-0190(91)90157-D.

34 Amit Sahai and Salil P. Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 50(2):196–249, 2003. doi:10.1145/636865.636868.

35 Claude E. Shannon. Communication Theory of Secrecy Systems. *Bell Systems Technical Journal*, 28:656–715, 1949.

36 Hung-Min Sun and Shiuh-Pyng Shieh. Secret Sharing in Graph-Based Prohibited Structures. In *Proceedings IEEE INFOCOM '97, The Conference on Computer Communications, Sixteenth Annual Joint Conference of the IEEE Computer and Communications Societies, Driving the Information Revolution, Kobe, Japan, April 7-12, 1997*, pages 718–724. IEEE, 1997. URL: http://ieeexplore.ieee.org/xpl/mostRecentIssue.jsp?punumber=4979, doi:10.1109/INFCOM.1997.644525.

37 Hoeteck Wee. Dual System Encryption via Predicate Encodings. In Yehuda Lindell, editor, *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings*, volume 8349 of *Lecture Notes in Computer Science*, pages 616–637. Springer, 2014. doi:10.1007/978-3-642-54242-8_26.

38 Hoteck Wee. Personal Communication, 2018.

A Formal Setup

For a finite set A we write $a \xleftarrow{R} A$ to denote a random variable which is sampled uniformly from A. The *statistical distance* between two discrete random variables, X and Y, denoted by $\Delta(X;Y)$ is defined by $\Delta(X;Y) := \frac{1}{2}\sum_z |\Pr[X=z] - \Pr[Y=z]|$. We will also use statistical distance for probability distributions, where for a probability distribution D the value $\Pr[D=z]$ is defined to be $D(z)$.

▶ **Definition 13** (CDS)**.** Let $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ be a predicate. Let $\mathsf{F}_A : \mathcal{X} \times \mathcal{Z} \times \mathcal{R} \to \mathcal{T}_A$ and $\mathsf{F}_B : \mathcal{Y} \times \mathcal{Z} \times \mathcal{R} \to \mathcal{T}_B$ be deterministic encoding algorithms, where $\mathcal{Z}$ is the *secret domain*. Then, the pair $(\mathsf{F}_A, \mathsf{F}_B)$ is a CDS scheme for f with *correctness error* c and *privacy error* s if the function $\mathsf{F}(x,y,z,r) = (\mathsf{F}_A(x,z,r), \mathsf{F}_B(y,z,r))$ that corresponds to the joint computation of F_A and F_B on a common z and r, satisfies the following properties:

1. (c-Correctness) There exists a deterministic algorithm Dec, called a *decoder*, such that for every 1-input (x,y) of f and any secret $z \in \mathcal{Z}$ we have that:

$$\Pr_{r \xleftarrow{R} \mathcal{R}} [\mathsf{Dec}(x, y, \mathsf{F}(x,y,z,r)) \neq z] \leq c$$

2. (s-Privacy) There exists a randomized simulator Sim such that for every 0-input (x,y) of f, every secret $z \in \mathcal{Z}$, and uniformly chosen randomness $r \xleftarrow{R} \mathcal{R}$ the following holds:

$$\Delta\left(\mathsf{Sim}(x,y)\ ;\ \mathsf{F}(x,y,z,r)\right) \leq s.$$

The *communication complexity* of the CDS protocol is $(\log|\mathcal{T}_A| + \log|\mathcal{T}_B|)$ and its *randomness complexity* is $\log|\mathcal{R}|$. If c and s are zeros, such a CDS scheme is called *perfect*.