Communication Complexity of Conditional Disclosure of Secrets and Attribute-Based Encryption[⋆]

Romain Gay[⋆⋆], Iordanis Kerenidis[⋆⋆⋆], and Hoeteck Wee[†]

Abstract. We initiate a systematic treatment of the communication complexity of conditional disclosure of secrets (CDS), where two parties want to disclose a secret to a third party if and only if their respective inputs satisfy some predicate. We present a general upper bound and the first non-trivial lower bounds for conditional disclosure of secrets. Moreover, we achieve tight lower bounds for many interesting setting of parameters for CDS with linear reconstruction, the latter being a requirement in the application to attribute-based encryption. In particular, our lower bounds explain the trade-off between ciphertext and secret key sizes of several existing attribute-based encryption schemes based on the dual system methodology.

1 Introduction

We revisit a fundamental question in the foundations of cryptography: what is the communication overhead of privacy in computation? This question has been considered in several different models and settings [12, 41, 2, 14]. In this work, we focus on a very simple and natural model where non-private computation requires very little communication (just a single bit), whereas the best upper bound for private computation is exponential.

Namely, we consider two-party conditional disclosure of secrets (CDS) [19] (c.f. Fig 2), a generalization of secret sharing [44, 23]: two parties want to disclose a secret to a third party if and only if their respective inputs satisfy some fixed predicate P. Concretely, Alice holds x, Bob holds y and they both share a secret $\alpha \in \{0,1\}$ (along with some additional private randomness), whereas Carol knows x, y but not α. Alice and Bob want to disclose α to Carol iff $\mathsf{P}(x,y) = 1$. How many bits do Alice and Bob need to communicate to Carol? In the non-private setting, Alice or Bob can send α to Carol, upon which Carol computes $\mathsf{P}(x,y)$ and decides whether to output α or $\perp$. This trivial protocol with one-bit communication is not private because Carol learns α even when the predicate is false; in fact, the best upper bound we have for CDS for general predicates requires that Alice and Bob each transmits $2^{\Omega(|x|+|y|)}$ bits [7]. Here, we are interested not only in the total communication from Alice and Bob to Carol, but also in trade-offs between the length of Alice's message ℓ_{A} and that of Bob's message ℓ_{B}.

Connection to Attribute-based Encryption. Attribute-based encryption (ABE) [43, 20] is a new paradigm for public-key encryption that enables fine-grained access control for encrypted data. In attribute-based encryption, ciphertexts are associated with descriptive values x in addition to a plaintext, secret keys are associated with values y, and a secret key decrypts the ciphertext if and only if $\mathsf{P}(x,y) = 1$ for some boolean predicate P. Note that x and y are public given the respective ciphertext and secret key.

[⋆] An extended abstract of this work appeared in the proceedings of CRYPTO 2015. This is the full version.

[⋆⋆] Email: `rgay@di.ens.fr`. ENS, Paris. Partially supported by ANR Project EnBiD (ANR-14-CE28-0003).

[⋆⋆⋆] Email: `jkeren@liafa.univ-paris-diderot.fr`. CNRS, LIAFA, Univ Paris Diderot. Partially supported by the ERC project QCC and the ANR project RDAM

[†] Email: `wee@di.ens.fr`. ENS, Paris, CNRS, INRIA and Columbia University. Partially supported by ANR Project EnBiD (ANR-14-CE28-0003), NSF Award CNS-1445424 and the Alexander von Humboldt Foundation.

Here, y together with P may express an arbitrarily complex access policy, which is in stark contrast to traditional public-key encryption, where access is all or nothing. The simplest example of ABE is that of identity-based encryption (IBE) [45, 8, 13] where P corresponds to equality. The security requirement for attribute-based encryption enforces resilience to collusion attacks, namely any group of users holding secret keys for different values learns nothing about the plaintext if none of them is individually authorized to decrypt the ciphertext. This should hold even if the adversary *adaptively* decides which secret keys to ask for.

In [47], Waters introduced the powerful *dual system encryption* methodology for building adaptively secure IBE in bilinear groups; this has since been extended to obtain adaptively secure ABE for a large class of predicates [31, 35, 38, 33, 30, 40]. In recent works [3, 48], Attrapadung and Wee presented a unifying framework for the design and analysis of dual system ABE schemes, which decouples the predicate P from the security proof. Specifically, the latter work puts forth the notion of *predicate encoding*, a private-key, one-time, information-theoretic primitive similar to conditional disclosure of secrets, and provides a compiler from predicate encoding for a predicate P into an ABE for the same predicate using the dual system encryption methodology. Moreover, the parameters in the predicate encoding scheme and in CDS correspond naturally to ciphertext and key sizes in the ABE. In particular, Alice's message corresponds to the ciphertext, and Bob's message to the secret key. For these applications, we require that Alice's and Bob's messages are linear functions of the shared randomness, and also that Carol computes a linear function of the messages to reconstruct the secret α. These applications consider linear functions over $\mathbb{Z}_p$ where p is the order of the underlying bilinear group; in this work, we focus on lower bounds for the case $p = 2$ although our techniques do hold for general p. Note that while the parameters for ABE schemes coming from predicate encodings are not necessarily the best known parameters, they do match the state-of-the-art in terms of ciphertext and secret key sizes for many predicates such as inner product, index, and read-once formula.

CDS Parameters. Unlike in traditional communication complexity where the primary measure is the total communication from Alice and from Bob, we make a more fine-grained distinction between the lengths of Alice's and Bob's messages ℓ_{A} and ℓ_{B}. For instance, in the application to ABE, ℓ_{A} and ℓ_{B} correspond to ciphertext and secret key sizes respectively. Note that for ABE ciphertext and key sizes, we ignore the contributions from the descriptive values x, y as well as multiplicative factors in the security parameter.[1] We are particularly interested in three regimes of parameters for $(\ell_{\mathsf{A}}, \ell_{\mathsf{B}})$:

- How small can ℓ_{B} be when ℓ_{A} is constant? This corresponding to minimizing key sizes for schemes with constant-size ciphertexts;
- How small can ℓ_{A} be when ℓ_{B} is constant? This corresponding to minimizing ciphertext sizes for schemes with constant-size keys;
- How small can $\max(\ell_{\mathsf{A}}, \ell_{\mathsf{B}})$ be? This corresponds to minimizing the overall parameter sizes of the scheme.

We also care about the complexity of the reconstruction function as computed by Carol, as a function of the messages from Alice and Bob; as noted earlier, for ABE, we will require linear reconstruction.

[1] The latter suppresses the distinction between counting bits and group elements, and also between working over $\mathbb{Z}_2$ vs $\mathbb{Z}_p$, where p is the order of the underlying bilinear group.

Predicate	ℓ_B, constant ℓ_A		ℓ_A, constant ℓ_B		$\max(\ell_\mathsf{A}, \ell_\mathsf{B})$	
	upper	lower	upper	lower	upper	lower
index, prefix	$O(n)$	$\Omega(n)^*$	$O(n)$	$\Omega(\sqrt{n})$	$O(\sqrt{n})$	$\Omega(\sqrt{n})^*$
disjointness, inner product	$O(n)$	$\Omega(n)^*$	$O(n)$	$\Omega(n)^*$	$O(n)$	$\Omega(\sqrt{n})$
read-once span programs	$O(2^n)$	$\Omega(n)$	$O(2^n)$	$\Omega(n^2)$	$O(n)$	$\Omega(n)^*$

Fig. 1. Summary of our upper and lower bounds for linear CDS, where ℓ_A and ℓ_B denote the length of the messages from Alice and Bob respectively. We marked the tight lower bounds with an asterisk *.

Prior works. There have been several works studying CDS protocols (and strengthenings thereof) for a large class of predicates [19, 3, 48, 22]: the best general upper bound achieves both linear reconstruction and communication that is linear in the size of the smallest (arithmetic) branching program computing the predicate [19, 22]. However, we basically do not have any techniques for proving lower bounds on the communication complexity of CDS protocols. Here, even the probabilistic method or a counting argument does not seem to yield meaningful lower bounds for a random function (in contrast, these techniques do yield meaningful lower bounds for circuit complexity of a random function).

1.1 Our results

We initiate a systematic treatment of the communication complexity of conditional disclosure of secrets (CDS). We present a general upper bound and the first non-trivial lower bounds for conditional disclosure of secrets, summarized in Fig 1. Moreover, we achieve tight lower bounds for many interesting setting of parameters for CDS with linear reconstruction, the latter being a requirement in the application to attribute-based encryption; this addresses an open problem posed in [48]. Very informally, for CDS with linear reconstruction, we obtain lower bounds of the form:

$$\ell_\mathsf{A} \cdot \ell_\mathsf{B} \geq \text{“communication complexity of } \mathsf{P}\text{”}$$

For example, for inner product on n-bit vectors, we have $\ell_\mathsf{A} \cdot \ell_\mathsf{B} = \Omega(n)$. Our lower bounds partially explain the trade-off between ciphertext and secret key sizes of several existing attribute-based encryption schemes based on the dual system methodology, c.f. [31, 35, 39, 48, 3, 10].

Proof techniques. Since we want to argue about the lengths of the messages of Alice and Bob to Carol, the first idea would be to look at the communication complexity of the predicate P [49, 29]. Informally, communication complexity measures how many bits of information about x and y we need to transmit in order to compute $\mathsf{P}(x, y)$ (c.f. Fig 2). Namely, Alice holds x and Bob holds y and each of them sends a message to a third party Carol who wants to compute $\mathsf{P}(x, y)$. We also allow all three parties to share public randomness w. The goal is to minimize the communication from Alice and Bob to Carol, and there is no privacy requirement. There is now a large body of works in communication complexity giving tight upper and lower bounds for a large class of predicates. For instance, a classic result from communication complexity tells us that to compute the inner product of two vectors $\mathbf{x}, \mathbf{y} \in \{0,1\}^n$, each of Alice and Bob must send $n - \Omega(1)$ bits [11]. That is, we need to know essentially all of $\mathbf{x}$ and all of $\mathbf{y}$ in order to compute their inner product.

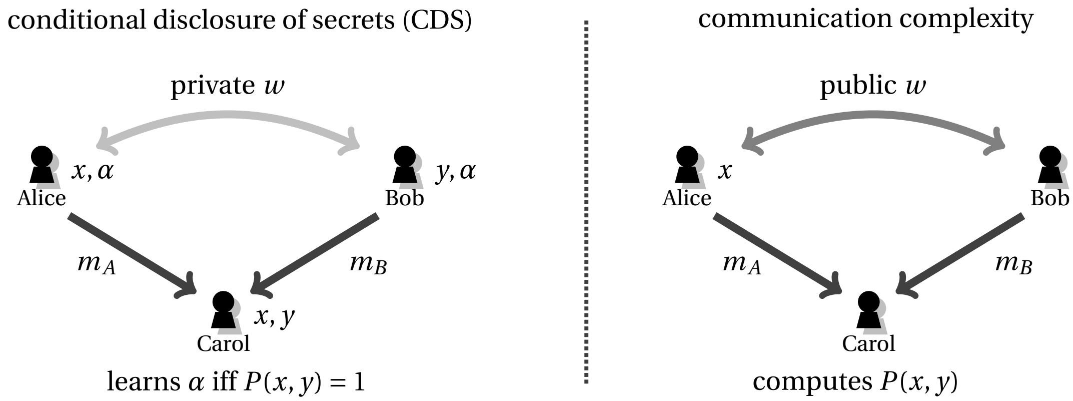

Fig. 2. Pictorial representation of CDS and communication complexity.

Our goal is to leverage the rich literature on lower bounds for communication complexity to obtain lower bounds for CDS. Namely, we want to transform any CDS Π_{cds} for a predicate P into a communication complexity protocol Π_{cc} for P with only a small blow-up in communication complexity. The crucial distinction between CDS and communication complexity is that Carol knows x, y in Π_{cds} but not in Π_{cc} (as shown in Fig 2).

The first attempt would be to show that a Π_{cds} for a predicate P is also a Π_{cc} for P. Fix x, y to denote the inputs to Π_{cc}. That is, we would like to argue that Alice's message together with Bob's message in a CDS (even without x, y) must completely determine $\mathsf{P}(x, y)$. Intuitively, this ought to be the case because if the CDS messages are consistent with both values of $\mathsf{P}(x, y)$, then they must simultaneously uniquely determine α (via correctness) and hide α (via privacy), a contradiction. Indeed, if this worked out, we would have a lower bound of the form

$$\ell_{\mathsf{A}} + \ell_{\mathsf{B}} \geq \text{“communication complexity for P”}$$

Unfortunately, the above statement is false for inner product. The above statement implies a lower bound of $2n - \Omega(1)$ bits for inner product, but we have a CDS for inner product with $n + 1$ bits! It is instructive to understand why the above attempt fails. The issue arises in using correctness of CDS to argue that Alice's and Bob's message must determine α: specifically, it is necessary for Carol to specify inputs x', y' in order to reconstruct α from Alice's and Bob's messages. In fact, different inputs (x', y') could yield different values for α. We need to fix this issue.

- The first idea is to have Alice in Π_{cc} also send the secret α; Carol then tries all possible (x', y') for which $P(x', y') = 1$ and output 1 iff for some x, y the reconstructed secret indeed equals α. By the correctness of CDS, Carol will output 1 when $P(x, y) = 1$. However, there could be false positives, since even when $P(x, y) = 0$, there could be inputs (x', y') for which $P(x', y') = 1$ and the reconstructed secret matches α, upon which Carol will incorrectly output 1. In fact, privacy tells us that Carol will recover a random value for the secret for each choice of (x', y'), and with pretty good probability, at least one of them will match α.
- The second idea is to avoid false positives by having Alice and Bob run the CDS protocol Π_{cds} N times, with fresh independent private randomness and secrets across the repetitions. As before, Carol will try all possible (x', y') for which $P(x', y') = 1$ and output 1 iff for some x', y' the reconstructed secret

equals α in all repetitions of the protocol. By the correctness of CDS, Carol will always output 1 when $P(x, y) = 1$. On the other hand, if $P(x, y) = 0$, a straight-forward union bound over $(x', y') \in P^{-1}(1)$ tells us Carol outputs 1 with probability at most $P^{-1}(1) \cdot 2^{-N}$, since Carol recovers a random value in each repetition. For inner product, we need to take a union bound over 2^{2n-1} possible pairs, which requires running $N = 2n - 1$ copies of the CDS protocol Π_{cds}; the communication complexity of Π_{cc} is then $2n - 1$ times that of Π_{cds}. This does not yield any non-trivial lower bound for Π_{cds} since we have an upper bound of $2n$ for communication complexity.

Here comes our key observation: we can substantially reduce the number of repetitions needed if the CDS protocol Π_{cds} has small communication complexity! Suppose Π_{cds} has total communication $\ell_A + \ell_B \ll n$ bits. Observe that the reconstruction function computed by Carol in Π_{cds} is a function from $\{0,1\}^{\ell_A+\ell_B}$ to $\{0,1\}$. Now, instead of having Carol in Π_{cc} enumerate over all possible (x, y), she will instead enumerate over all functions from $\{0,1\}^{\ell_A+\ell_B}$ to $\{0,1\}$, and output 1 iff for some function the reconstructed secret equals α in all N repetitions. By the correctness of CDS, Carol will always output 1 when $P(x, y) = 1$. Moreover, there are $2^{2^{\ell_A+\ell_B}}$ possible functions, which means we will need to run $2^{\ell_A+\ell_B}$ copies of Π_{cds} in Π_{cc}; this already implies a $\Omega(\log n)$ lower bound for inner product! Moreover, if the CDS Π_{cds} admits linear reconstruction, then Carol in Π_{cc} will also need to enumerate over all $2^{\ell_A+\ell_B}$ linear functions from $\{0,1\}^{\ell_A+\ell_B}$ to $\{0,1\}$, which means we only need to run $\ell_A + \ell_B$ copies of Π_{cds} in Π_{cc}; this in turn yields a $\Omega(\sqrt{n})$ lower bound for inner product.

We obtain our lower bounds on CDS for concrete predicates by instantiating the above argument with existing lower bounds in communication complexity [36, 28, 24, 42, 4, 11] (c.f. Section 5). For prefix and read-once monotone span programs, we present tight lower bounds, c.f Appendix D.

Implications for dual system ABE. As observed in [3, 48, 10], underlying most "information-theoretic" dual system ABE schemes for a predicate P is a CDS for the same predicate, and our lower bounds apply to ciphertext and secret key sizes for these dual system ABE schemes. On the other hand, we do have ABE schemes based on a "computational" dual system argument, such as those in [32, 34, 9, 3, 27], many of which are more efficient and do avoid the lower bounds in this work. Informally, underlying the "computational" dual system argument is a computational analogue of CDS, where the privacy requirement is computational rather than information-theoretic. As it turns out, formalizing the right notion of computational privacy in CDS is quite tricky.

Recall that CDS guarantees privacy of the secret α whenever $\mathsf{P}(x, y) = 0$, and in the application to ABE, we require that privacy holds even if x, y are chosen adaptively, namely Alice's input x may be chosen depending on Bob's input y and Bob's message, and vice versa. Now, if the privacy guarantee is information-theoretic and perfect, then privacy for non-adaptive choices of x, y implies privacy for adaptive choices[2]; this equivalence dissipates as soon as we relax the privacy requirement to be statistical or computational. The "right" notion of computational privacy for use in ABE schemes is that of "doubly selective" security [3, 34], where "doubly" refers to the two possibilities depending on whether x or y is chosen first. Unsurprisingly, proving[3] and using doubly selective security require substantially more

[2] The easiest way to see this is via complexity leveraging: an adaptive distinguisher with advantage ε can be converted into a non-adaptive distinguisher with an exponential loss in ε via random guessing. Since any non-adaptive distinguisher has advantage 0, we must have $\varepsilon = 0$ to begin with.

[3] Typically, this entails two separate reductions, one for x being chosen first and the other for y. In [34], these correspond to selectively secure key-policy and ciphertext-policy ABE schemes; in [3], these correspond to so-called selective and co-selective security.

delicate security reductions, and in most cases, stronger and less desirable q-type assumptions. This raises the natural question of whether the increased complexity in these proofs and assumptions are inherent, or simply a failure to find more clever and efficient CDS with information-theoretic privacy. Our work rules out the latter option.

1.2 Discussion

Perspective. Note that our set-up is quite different from previous lower bounds for private computation in the literature; to the best of our knowledge, this is the first super-constant lower bound in a setting where the price of privacy in computation is always bounded. For instance, in interactive secure two-party computation, some functions are impossible to compute securely [12], so the cost of privacy is infinite for these functions (whereas ours is bounded for all predicates). For secure computation in the FKN model [15, 14], we do not have any techniques for super-constant gaps. For locally decodable codes, there is no gap for privacy in some ranges of parameters, for instance, when we want to minimize one-way communication from the client and communication from the server is essentially "free"; here, the server needs to send the entire database, whether or not we care about client privacy.

Additional related work. There is a large body of work on lower bounds on share sizes in secret-sharing (c.f. [5, Section 5]). Most of these works rely on Shannon-type inequalities on entropy of random variables, which do not seem applicable to our setting. Roughly speaking, in secret sharing, Carol either gets a share or not, whereas Alice and Bob in CDS can do more complex computations than simply computing shares and then deciding whether to send each share to Carol. The recent work of Data, Prabhakaran and Prabhakaran [14] draws upon tools from information theory to obtain new communication complexity lower bounds for secure computation in three-party setting. In their model which allows multiple rounds of interaction, the problem we consider admits a secure protocol with a single bit of communication, and their techniques do not yield better bounds in the non-interactive setting.

Open problems. We conclude with a number of open problems:

- explore the power of non-linear reconstruction in CDS (that is, positive results, c.f. [6, 46]);
- tight lower bounds for inner product with linear reconstruction (which we conjecture to be $\Omega(n)$);
- obtain better lower bounds for multi-bit secrets (which is related to lower bounds for secret sharing for multi-bit secrets), or obtain upper bounds that are better than the naive "direct product" construction;
- improve the upper or lower bounds in CDS for read-once span programs for constant ℓ_{A} or constant ℓ_{B}. A related problem is to prove stronger communication complexity lower bounds for general span programs (which may not be read-once).

2 Preliminaries

Notations. We denote by $s \leftarrow_{\mathrm{R}} S$ the fact that s is picked uniformly at random from a finite set S or from a distribution. Throughout this paper, we denote by log the logarithm of base 2.

2.1 Conditional disclosure of secrets

We recall the notion of conditional disclosure of secrets (CDS), c.f. Fig 2. The definition we give here is for two parties Alice and Bob and a referee Carol, where Alice and Bob share randomness w and want to conditionally disclose a secret α to Carol. The general notion of conditional disclosure of secrets has first been investigated in [19]. Two-party CDS is closely related to the notions of predicate encoding [48, 10] and pairing encoding [3]; in particular, the latter two notions imply two-party CDS with linear reconstruction.

Definition 1 (conditional disclosure of secrets (CDS) [19, 48]). *Fix a predicate* $\mathsf{P} : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$. *A* $(\ell_\mathsf{A}, \ell_\mathsf{B})$-conditional disclosure of secrets (CDS) *for* P *is a triplet of deterministic functions* $(\mathsf{A}, \mathsf{B}, \mathsf{C})$

$$\mathsf{A} : \mathcal{X} \times \mathcal{W} \times \mathcal{D} \to \{0,1\}^{\ell_A}, \quad \mathsf{B} : \mathcal{Y} \times \mathcal{W} \times \mathcal{D} \to \{0,1\}^{\ell_B}, \quad \mathsf{C} : \mathcal{X} \times \mathcal{Y} \times \{0,1\}^{\ell_\mathsf{A}} \times \{0,1\}^{\ell_\mathsf{B}} \to \mathcal{D}$$

satisfying the following properties:

(reconstruction.) *For all* $(x, y) \in \mathcal{X} \times \mathcal{Y}$ *such that* $\mathsf{P}(x, y) = 1$, *for all* $w \in \mathcal{W}$, *and for all* $\alpha \in \mathcal{D}$*:*

$$\mathsf{C}(x, y, \mathsf{A}(x, w, \alpha), \mathsf{B}(y, w, \alpha)) = \alpha$$

(privacy.) *For all* $(x, y) \in \mathcal{X} \times \mathcal{Y}$ *such that* $\mathsf{P}(x, y) = 0$, *and for all* $\mathsf{C}^* : \{0,1\}^{\ell_\mathsf{A}} \times \{0,1\}^{\ell_\mathsf{B}} \to \mathcal{D}$,

$$\Pr_{w \leftarrow \mathcal{W}, \alpha \leftarrow_{\mathrm{R}} \mathcal{D}} \Big[\mathsf{C}^*\big(\mathsf{A}(x, w, \alpha), \mathsf{B}(y, w, \alpha)\big) = \alpha \Big] \le \frac{1}{|\mathcal{D}|}$$

Note that the formulation of privacy above with uniformly random secrets is equivalent to standard indistinguishability-based formulations (c.f. Section A).

A useful measure for the complexity of a CDS is the complexity of reconstruction as a function of the outputs of A, B, as captured by the function C, with (x, y) hard-wired.

Definition 2 ($\mathcal{C}$-reconstruction). *Given a set* $\mathcal{C}$ *of functions from* $\{0,1\}^{\ell_\mathsf{A}} \times \{0,1\}^{\ell_\mathsf{B}} \to \mathcal{D}$, *we say that a CDS* $(\mathsf{A}, \mathsf{B}, \mathsf{C})$ *admits* $\mathcal{C}$*-reconstruction iff for all* (x, y) *such that* $\mathsf{P}(x, y) = 1$, $\mathsf{C}(x, y, \cdot, \cdot) \in \mathcal{C}$.

Two examples of $\mathcal{C}$ of interest are:

- $\mathcal{C}_{\mathrm{all}}$ is the set of all functions from $\{0,1\}^{\ell_\mathsf{A}} \times \{0,1\}^{\ell_\mathsf{B}} \to \mathcal{D}$; that is, we do not place any restriction on the complexity of reconstruction. Note that $|\mathcal{C}_{\mathrm{all}}| = |\mathcal{D}|^{2^{\ell_\mathsf{A}+\ell_B}}$.
- $\mathcal{C}_{\mathrm{lin}}$ is the set of all *linear* functions over $\mathbb{Z}_2$ from $\{0,1\}^{\ell_\mathsf{A}} \times \{0,1\}^{\ell_\mathsf{B}} \to \mathcal{D}$; that is, we require the reconstruction to be linear as a function of the outputs of A and B as bit strings (but may depend arbitrarily on x, y). This is the analogue of linear reconstruction in linear secret sharing schemes and is a requirement for the applications to attribute-based encryption [48, 3, 10]. In Appendix B, we show that any scheme where A and B compute linear functions (again, analogue of linear secret-sharing schemes and a requirement for many cryptographic applications) also satisfies linear reconstruction. Note that $|\mathcal{C}_{\mathrm{linear}}| \le |\mathcal{D}|^{\ell_\mathsf{A}+\ell_B}$ for $|\mathcal{D}| \ge 2$.

Remark 1. Note that while looking at $\mathcal{C}$, we consider $\mathsf{C}(x, y, \cdot, \cdot)$, which has (x, y) hard-wired, and takes an input of total length $\ell_\mathsf{A} + \ell_\mathsf{B}$. In particular, it could be that C runs in time linear in $|x| = |y| = n$, and yet $\ell_\mathsf{A} = \ell_\mathsf{B} = O(\log n)$ so C has "exponential" complexity w.r.t. $\ell_\mathsf{A} + \ell_\mathsf{B}$.

Definition 3 (linear CDS). *We say that a CDS* $(\mathsf{A}, \mathsf{B}, \mathsf{C})$ *is linear if it admits* $\mathcal{C}_{\mathrm{lin}}$*-reconstruction.*

2.2 Communication complexity

The description of communication complexity in Fig 2 actually refers to the "simultaneous message" model, where A and B each sends a message to C. For our actual proof, it suffices to consider one way communication complexity, where there is no C, but either A sends a single message to B or B sends a single message to A. We now proceed to recall the basic definitions for communication complexity [49, 29], specifically one-way communication complexity with one-sided error [1, 28, 37].

Definition 4 ([28, 49]). *A* one-way $(\mathsf{A} \to \mathsf{B})$ communication protocol *for a predicate* $\mathsf{P} : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ *is a pair of deterministic functions* (A,B) *where*

$$\mathsf{A} : \mathcal{X} \times \mathcal{W} \times \{0,1\}^{\ell} \to \{0,1\}, \quad \mathsf{B} : \mathcal{Y} \times \mathcal{W} \times \{0,1\}^{\ell} \to \{0,1\},$$

and the following properties are satisfied for every $(x, y) \in \mathcal{X} \times \mathcal{Y}$*:*

- *If* $\mathsf{P}(x, y) = 1$*, then* $\Pr_{w \leftarrow_{\mathrm{R}} \mathcal{W}}[\mathsf{B}(y, w, \mathsf{A}(x, w)) = 1] = 1$
- *If* $\mathsf{P}(x, y) = 0$*, then* $\Pr_{w \leftarrow_{\mathrm{R}} \mathcal{W}}[\mathsf{B}(y, w, \mathsf{A}(x, w)) = 0] \geq 1/2$.

The one-way communication complexity *of* P*, denoted by* $\mathsf{R}^{\mathsf{A}\to\mathsf{B}}(\mathsf{P})$*, is the minimum* ℓ *over all one-way communication protocols* (A,B) *for* P.

We also denote by $\mathsf{R}^{\mathsf{B}\to\mathsf{A}}(\mathsf{P})$ *the minimum* ℓ *over all one-way* $(\mathsf{B} \to \mathsf{A})$ *communication protocols* (A,B)*, where*

$$\mathsf{A} : \mathcal{X} \times \mathcal{W} \times \{0,1\}^{\ell} \to \{0,1\}, \quad \mathsf{B} : \mathcal{Y} \times \mathcal{W} \times \{0,1\}^{\ell} \to \{0,1\},$$

and the following properties are satisfied for every $(x, y) \in \mathcal{X} \times \mathcal{Y}$*:*

- *If* $\mathsf{P}(x, y) = 1$*, then* $\Pr_{w \leftarrow_{\mathrm{R}} \mathcal{W}}[\mathsf{A}(x, w, \mathsf{B}(y, w)) = 1] = 1$
- *If* $\mathsf{P}(x, y) = 0$*, then* $\Pr_{w \leftarrow_{\mathrm{R}} \mathcal{W}}[\mathsf{A}(x, w, \mathsf{B}(y, w)) = 0] \geq 1/2$.

3 CDS for General Predicates

We present a general upper bound for linear CDS for any predicate:

Theorem 1 (generic upper bounds for linear CDS). *Given any predicate* $\mathsf{P} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$*, for any* $t \leq 2^n$*, there exists a linear* $(t, 2^n/t)$*-CDS for* P *with* $\mathcal{D} = \{0,1\}$*. In particular, there exists a* $(1, 2^n)$*-CDS, a* $(2^n, 1)$*-CDS, a* $(2^{n/2}, 2^{n/2})$*-CDS for* P*, all three of which are linear.*

The result improves upon the $(2^{n/2}, 2^{n/2})$-CDS (but not linear) given in [7]; our construction is also considerably simpler.

Proof (sketch). The construction follows from a standard reduction of any general predicate to the INDEX predicate on 2^n-dimensional vectors: Alice treats the truth table $\mathsf{P}(x, \cdot)$ as a vector of length 2^n and Bob treats $y \in \{0,1\}^n$ as an index, so that the INDEX predicate returns $\mathsf{P}(x, y)$. Then, we can use the $(t, 2^n/t)$-linear CDS for the INDEX predicate on 2^n-dimensional vectors in [17, 10] (c.f. Appendix C). □

More generally, for any predicate $\mathsf{P} : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$, we have a $(t, \min(|\mathcal{X}|, |\mathcal{Y}|)/t)$-linear CDS, by treating either x or y as an index depending on whether $|\mathcal{X}| \leq |\mathcal{Y}|$ or not. This is essentially optimal for linear reconstruction, since we prove a tight lower bound for INDEX: $\{0,1\}^n \times [n] \to \{0,1\}$ in Section 5.

4 Lower Bounds for CDS

In this section, we present our lower bounds on the communication complexity of CDS.

Theorem 2 (lower bounds for linear CDS). *Let* $\mathsf{P} : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ *be a predicate. For all linear* $(\ell_{\mathsf{A}}, \ell_{\mathsf{B}})$*-CDS of* P *with* $|\mathcal{D}| \geq 2$*, we have*

$$\ell_{\mathsf{A}} \cdot (\ell_{\mathsf{A}} + \ell_{\mathsf{B}} + 1) \geq \mathsf{R}^{\mathsf{A}\to\mathsf{B}}(\mathsf{P}) \quad \textit{and} \quad \ell_{\mathsf{B}} \cdot (\ell_{\mathsf{A}} + \ell_{\mathsf{B}} + 1) \geq \mathsf{R}^{\mathsf{B}\to\mathsf{A}}(\mathsf{P}).$$

We then derive our lower bounds for linear CDS by using existing lower bounds on one-way communication complexity; see Section 5. In fact, our techniques are fairly general and also yield lower bounds on non-linear CDS.

Theorem 3 (lower bounds for general CDS). *Let* $\mathsf{P} : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ *be a predicate. For all* $(\ell_{\mathsf{A}}, \ell_{\mathsf{B}})$*-predicate CDS of* P *with* $|\mathcal{D}| \geq 2$*, we have*

$$\ell_{\mathsf{A}} + \ell_{\mathsf{B}} \geq \frac{1}{2} \log\Big(\mathsf{R}^{\mathsf{A}\to\mathsf{B}}(\mathsf{P}) + \mathsf{R}^{\mathsf{B}\to\mathsf{A}}(\mathsf{P})\Big).$$

While the lower bounds for general CDS are exponentially smaller than those for linear CDS, we still do obtain non-trivial logarithmic lower bounds for many concrete predicates.

4.1 Main lemma

We obtain both lower bounds via a general reduction from CDS for a predicate P to one-way communication protocols for the same predicate; the communication cost of the reduction depends crucially on the complexity of reconstruction (c.f. Definition 2):

Lemma 1 (main technical lemma). *Let* $\mathsf{P} : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ *be a predicate. Then, any* $(\ell_{\mathsf{A}}, \ell_{\mathsf{B}})$*-CDS for* P *with* $|\mathcal{D}| \geq 2$ *and which admits* $\mathcal{C}$*-reconstruction satisfies*

$$(\log|\mathcal{C}| + 1) \cdot \ell_{\mathsf{A}} \geq \mathsf{R}^{\mathsf{A}\to\mathsf{B}}(\mathsf{P}) \cdot \log|\mathcal{D}| \quad \textit{and} \quad (\log|\mathcal{C}| + 1) \cdot \ell_{\mathsf{B}} \geq \mathsf{R}^{\mathsf{B}\to\mathsf{A}}(\mathsf{P}) \cdot \log|\mathcal{D}|$$

Theorem 2 then follows from instantiating the lemma with $\mathcal{C} := \mathcal{C}_{\text{lin}}$, where $\log|\mathcal{C}_{\text{lin}}| = (\ell_{\mathsf{A}} + \ell_{\mathsf{B}}) \cdot \log|\mathcal{D}|$. Similarly, Theorem 3 uses $\mathcal{C} := \mathcal{C}_{\text{all}}$ where $\log|\mathcal{C}_{\text{all}}| = 2^{\ell_{\mathsf{A}} + \ell_{\mathsf{B}}} \cdot \log|\mathcal{D}|$.

Proof (of Lemma 1). Let $N := \frac{\log|\mathcal{C}|+1}{\log|\mathcal{D}|}$. We build a one-way communication protocol $(\widetilde{\mathsf{A}}, \widetilde{\mathsf{B}})$ for the predicate P as follows:

- Sample $w_i \leftarrow_{\mathrm{R}} \mathcal{W}, \alpha_i \leftarrow_{\mathrm{R}} \mathcal{D}$ for $i = 1, \ldots, N$ and set

$$w := (w_1, \alpha_1, \ldots, w_N, \alpha_N)$$

- Alice computes

$$\widetilde{\mathsf{A}}(x, w) := (\mathsf{A}(x, w_1, \alpha_1), \ldots, \mathsf{A}(x, w_N, \alpha_N))$$

- Bob outputs 1 iff there exists a function $\mathsf{C}^* \in \mathcal{C}$ such that

$$\mathsf{C}^*\big(\mathsf{A}(x, w_i, \alpha_i), \mathsf{B}(y, w_i, \alpha_i)\big) = \alpha_i, \quad \forall\, i = 1, \ldots, N$$

We proceed to analyze the protocol $(\widetilde{\mathsf{A}}, \widetilde{\mathsf{B}})$.

- Completeness. Suppose $\mathsf{P}(x,y) = 1$. Then, by the reconstruction property, the function $\mathsf{C}^*(\cdot) := \mathsf{C}(x,y,\cdot) \in \mathcal{C}$ satisfies
$$\mathsf{C}^*\big(\mathsf{A}(x,w_i,\alpha_i),\mathsf{B}(y,w_i,\alpha_i)\big) = \alpha_i, \quad \forall\, i = 1,\dots,N$$
for all $(w_1,\alpha_1,\dots,w_N,\alpha_N)$. Therefore, $\widetilde{\mathsf{B}}$ outputs 1 with probability 1.
- Soundness. Suppose $\mathsf{P}(x,y) = 0$. Fix $\mathsf{C}^* \in \mathcal{C}$. For each $i = 1,\dots,N$, α-privacy implies that
$$\Pr_{w_i,\alpha_i}\Big[\mathsf{C}^*\big(\mathsf{A}(x,w_i,\alpha_i),\mathsf{B}(y,w_i,\alpha_i)\big) = \alpha_i\Big] \le \tfrac{1}{|D|}$$
Since the (w_i,α_i) are chosen independently at random, we have
$$\Pr_{w_1,\alpha_1,\dots,w_N,\alpha_N}\Big[\mathsf{C}^*\big(\mathsf{A}(x,w_i,\alpha_i),\mathsf{B}(y,w_i,\alpha_i)\big) = \alpha_i, \quad \forall\, i = 1,\dots,N\Big] \le \tfrac{1}{|\mathcal{D}|^N}$$
By a union bound over all $|\mathcal{C}|$ functions $\mathsf{C}^* \in \mathcal{C}$, we have
$$\Pr\Big[\widetilde{\mathsf{B}} \text{ outputs } 1\Big] \le |\mathcal{C}| \cdot |\mathcal{D}|^{-N} \le 1/2$$
by our choice of N.

It is straightforward to check that $\widetilde{\mathsf{A}}$ sends $\frac{\log|\mathcal{C}|+1}{\log|\mathcal{D}|} \cdot \ell_{\mathsf{A}}$ bits to $\widetilde{\mathsf{B}}$. Similarly, we can build a $(\widetilde{\mathsf{B}}, \widetilde{\mathsf{A}})$ protocol for P, where $\widetilde{\mathsf{B}}$ sends $\frac{\log|\mathcal{C}|+1}{\log|\mathcal{D}|} \cdot \ell_{\mathsf{B}}$ bits to $\widetilde{\mathsf{A}}$. This completes the proof. □

Remark 2 (extensions). It is easy to see that the reduction also works for CDS with imperfect reconstruction and weak privacy. If the gap between the probability of reconstructing α when $\mathsf{P}(x,y) = 1$ and the probability of recovering α when $\mathsf{P}(x,y) = 0$ is δ, then it suffices to take $N := O\Big(\frac{1}{\delta}\log|\mathcal{C}|\Big)$ via a straightforward application of the Chernoff bound. The ensuing randomized protocol for communication complexity will then have a two-sided error.

Remark 3 (beyond linear CDS). Note that the bounds of Theorem 2 are much more general than just for linear CDS. For instance, if we require that reconstruction be carried out by circuits of size ℓ^c for some constant c (where $\ell := \ell_{\mathsf{A}} + \ell_{\mathsf{B}}$), or by polynomials of degree c, then we get lower bounds of the form
$$\ell_{\mathsf{A}} + \ell_B = \Omega\Big((\mathsf{R}^{\mathsf{A}\to\mathsf{B}}(\mathsf{P}) + \mathsf{R}^{\mathsf{B}\to\mathsf{A}}(\mathsf{P}))^{1/(c+1)}\Big)$$

4.2 Lower bounds for multi-bit secrets

We now look at CDS where the secret α is a multi-bit string; that is, $\mathcal{D}$ is of the form $\{0,1\}^d$, for $d \ge 1$. There is a trivial upper bound for d-bit secrets obtained by running d times a CDS for single-bit secrets. Note, of course, that hiding a secret of size $d = 1$ is the easiest case, since we can simply embed this secret to a larger d-bit string by randomly adding $d-1$ bits and use the CDS for the secret of size d. Hence, the lower bounds on the message lengths of the CDS for a secret of size $d = 1$ still hold for the CDS of secret of size $d \ge 1$. We would like a lower bound that grows with d.

Here, we prove that for any non-trivial predicate P, for any $(\ell_{\mathsf{A}}, \ell_{\mathsf{B}})$-CDS of P, both ℓ_{A} and ℓ_{B} need to be at least d. A trivial predicate is one whose output is completely determined by either x or y (e.g. the output of the predicate is the first bit of x), for which there is a protocol with $\ell_A + \ell_B = d$. The intuition is that in any non-trivial predicate, Alice's message essentially serves as the secret key for a one-time pad, which is needed to "unlock" $\alpha \in \{0,1\}^d$ from Bob's message. This means that Alice's message must itself be at least d bits.

It is easy to see that the lower bound is tight for the equality predicate. For all other non-trivial predicates, it remains an open problem to close the gap between lower and upper bounds for CDS of multi-bit secrets.

Theorem 4. *Let $\mathcal{D} := \{0,1\}^d$, and let $\mathsf{P} : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ be a non-trivial predicate that depends on both inputs x and y; that is, there exists $x^* \in \mathcal{X}$, such that $\mathsf{P}(x^*, \cdot)$ is not constant on $\mathcal{Y}$, and there exists $y^* \in \mathcal{Y}$ such that $\mathsf{P}(\cdot, y^*)$ is not constant on $\mathcal{X}$. Then, for any $(\ell_\mathsf{A}, \ell_\mathsf{B})$-CDS of P, we have*

$$\ell_A \geq d \quad \textit{and} \quad \ell_\mathsf{B} \geq d.$$

Proof. We begin with the lower bound on ℓ_A. Let $x_0, x_1 \in \mathcal{X}$ be such that

$$\mathsf{P}(x_0, y^*) = 0 \quad \text{and} \quad \mathsf{P}(x_1, y^*) = 1$$

Let $\mathsf{C}^* : \{0,1\}^{\ell_\mathsf{A}+\ell_\mathsf{B}} \to \{0,1\}^d$ be a randomized function defined as follows: on input $m_\mathsf{A} \in \{0,1\}^{\ell_\mathsf{A}}$ and $m_\mathsf{B} \in \{0,1\}^{\ell_\mathsf{B}}$,

- picks a message $m \leftarrow_\mathrm{R} \{0,1\}^{\ell_\mathsf{A}}$ at random (and ignores m_A);
- outputs $\mathsf{C}(x_1, y^*, m, m_\mathsf{B})$.

By α-reconstruction for $\mathsf{P}(x_1, y^*) = 1$, for all $\alpha \in \mathcal{D}$, $w \in \mathcal{W}$, we have

$$\mathsf{C}\big(x_1, y^*, \mathsf{A}(x_1, w, \alpha), \mathsf{B}(y^*, w, \alpha)\big) = \alpha.$$

Therefore, for all $\alpha \in \mathcal{D}$, $w \in \mathcal{W}$, we have

$$\Pr_{m \leftarrow_\mathrm{R} \{0,1\}^{\ell_A}} \Big[\mathsf{C}\big(x_1, y^*, \mathsf{A}(x_1, w, \alpha), \mathsf{B}(y^*, w, \alpha)\big) = \alpha \text{ and } m = \mathsf{A}(x_1, w, \alpha)\Big] = 1/2^{\ell_\mathsf{A}}$$

Thus,

$$\begin{aligned} &\Pr_{w \leftarrow \mathcal{W}, \alpha \leftarrow_\mathrm{R} \mathcal{D}, \text{ coins of } \mathsf{C}^*} \Big[\mathsf{C}^*\big(\mathsf{A}(x_1, w, \alpha), \mathsf{B}(y^*, w, \alpha)\big) = \alpha\Big] \\ \geq &\Pr_{w \leftarrow \mathcal{W}, \alpha \leftarrow_\mathrm{R} \mathcal{D}, m \leftarrow_\mathrm{R} \{0,1\}^{\ell_A}} \Big[\mathsf{C}\big(x_1, y^*, \mathsf{A}(x_1, w, \alpha), \mathsf{B}(y^*, w, \alpha)\big) = \alpha \text{ and } m = \mathsf{A}(x_1, w, \alpha)\Big] = 1/2^{\ell_\mathsf{A}} \end{aligned}$$

Since C^* ignores m_A, this means that for all m_A, and in particular for $m_\mathsf{A} = \mathsf{A}(x_0, w, \alpha)$, we have

$$\Pr_{w \leftarrow \mathcal{W}, \alpha \leftarrow_\mathrm{R} \mathcal{D}, \text{ coins of } \mathsf{C}^*} \Big[\mathsf{C}^*\big(\mathsf{A}(x_0, w, \alpha), \mathsf{B}(y^*, w, \alpha)\big) = \alpha\Big] \geq 1/2^{\ell_\mathsf{A}}$$

On the other hand, by α-privacy for $\mathsf{P}(x_0, y^*) = 0$, we have

$$\Pr_{w \leftarrow \mathcal{W}, \alpha \leftarrow_\mathrm{R} \mathcal{D}, \text{ coins of } \mathsf{C}^*} \Big[\mathsf{C}^*\big(\mathsf{A}(x_0, w, \alpha), \mathsf{B}(y^*, w, \alpha)\big) = \alpha\Big] \leq 1/2^d$$

Combining the two preceding inequalities, we have $1/2^{\ell_{\mathsf{A}}} \le 1/2^d$ and thus,

$$\ell_{\mathsf{A}} \ge d.$$

For the same reason,

$$\ell_{\mathsf{B}} \ge d.$$

□

5 Concrete predicates

In this section, we describe how we can combine the results in the previous section with lower bounds in one-way communication complexity to obtain the results in Figure 1. Each of these predicates has been studied in prior works on attribute-based encryption. For each of these predicates, we obtain non-trivial lower bounds for general $(\ell_{\mathsf{A}}, \ell_{\mathsf{B}})$-CDS of the form:

$$\ell_{\mathsf{A}} + \ell_{\mathsf{B}} = \Omega(\log n).$$

We focus hence-forth on lower bounds for linear $(\ell_{\mathsf{A}}, \ell_{\mathsf{B}})$-CDS, where linearity is over $\mathbb{Z}_2$. In the applications to ABE, we will typically work with linear functions over $\mathcal{D} = \mathbb{Z}_p$ (where $\log p$ is linear in the security parameter), in which case we lose a multiplicative $\log p$ factor in the lower bounds.

Index, Prefix. We consider the following predicates:

- Index: $\mathcal{X} := \{0,1\}^n, \mathcal{Y} := [n]$ and
$$\mathsf{P}_{\text{index}}(\mathbf{x}, i) = 1 \text{ iff } x_i = 1$$
That is, $\mathbf{x}$ is the characteristic vector of a subset of $[n]$. In the context of ABE, this corresponds to broadcast encryption [16].
- Prefix: $\mathcal{X} := \{0,1\}^n, \mathcal{Y} := \{0,1\}^{\le n}$ and
$$\mathsf{P}_{\text{prefix}}(\mathbf{x}, \mathbf{y}) = 1 \text{ iff } \mathbf{y} \text{ is a prefix of } \mathbf{x}$$
In the context of ABE, this corresponds to hierarchical identity-based encryption [18, 21].

For both predicates, we have tight bounds for one-way communication complexity:

$$\mathsf{R}^{\mathsf{A}\to\mathsf{B}}(\mathsf{P}) = \Theta(n) \quad \text{and} \quad \mathsf{R}^{\mathsf{B}\to\mathsf{A}}(\mathsf{P}) = \Theta(\log n)$$

given in [36, 28] for index, and Appendix D for prefix.

By Theorem 2, this means that any linear $(\ell_{\mathsf{A}}, \ell_{\mathsf{B}})$-CDS for any of the two predicates must satisfy

$$\ell_{\mathsf{A}}(\ell_{\mathsf{A}} + \ell_{\mathsf{B}} + 1) = \Omega(n).$$

This immediately yields

- $\ell_{\mathsf{B}} = \Omega(n)$ if $\ell_{\mathsf{A}} = O(1)$ and more generally, $\ell_{\mathsf{B}} = \Omega(n/\ell_A)$ for any $\ell_A = o(\sqrt{n})$;
- $\ell_{\mathsf{A}} = \Omega(\sqrt{n})$ if $\ell_{\mathsf{B}} = O(1)$;

- $\max(\ell_{\mathsf{A}}, \ell_{\mathsf{B}}) = \Omega(\sqrt{n})$.

The first and third lower bounds are tight, as we have a linear $(t, n/t)$-CDS for any $t \in [n]$, in [10, 48, 3] for index, and in Appendix C for prefix.

Disjointness, Inner product. We consider the following predicates:

- Disjointness: $\mathcal{X} = \mathcal{Y} := \{S \subseteq [n]\}$ and

$$\mathsf{P}_{\text{disj}}(X, Y) = 1 \text{ iff } X \cap Y = \emptyset$$

  In the context of ABE, this is related to a special case of fuzzy IBE [43].
- Inner Product [26]: $\mathcal{X} = \mathcal{Y} := \mathbb{Z}_p^n$ and

$$\mathsf{P}_{\text{IP}}(\mathbf{x}, \mathbf{y}) = 1 \text{ iff } \mathbf{x}^{\top}\mathbf{y} = 0$$

For both predicates, we have tight bounds for one-way communication complexity:

$$\mathsf{R}^{\mathsf{A}\to\mathsf{B}}(\mathsf{P}) = \Theta(n) \quad \text{and} \quad \mathsf{R}^{\mathsf{B}\to\mathsf{A}}(\mathsf{P}) = \Theta(n)$$

given in [24, 42, 4] for disjointness, in [11] for inner product. By Theorem 2, this means that any linear $(\ell_{\mathsf{A}}, \ell_{\mathsf{B}})$-CDS for any of the two predicates must satisfy

$$\ell_{\mathsf{A}}(\ell_{\mathsf{A}} + \ell_{\mathsf{B}} + 1) = \Omega(n) \quad \text{and} \quad \ell_{\mathsf{B}}(\ell_{\mathsf{A}} + \ell_{\mathsf{B}} + 1) = \Omega(n).$$

This immediately yields

- $\ell_{\mathsf{B}} = \Omega(n)$ if $\ell_{\mathsf{A}} = O(1)$;
- $\ell_{\mathsf{A}} = \Omega(n)$ if $\ell_{\mathsf{B}} = O(1)$;
- $\max(\ell_{\mathsf{A}}, \ell_{\mathsf{B}}) = \Omega(\sqrt{n})$.

The first and second lower bounds are tight, as we have matching upper bounds in [10, 48, 3]. In Appendix C, we exhibit a linear $(t, n-t+O(1))$-CDS for these predicates, for any $t \in [n]$. It is open whether a CDS with overall parameter size of $O(\sqrt{n})$ is possible.

Read-once monotone span programs. We consider the following predicate:

- Read-once monotone span program: $\mathcal{X} := \{0,1\}^n$, $\mathcal{Y} := \mathbb{Z}_p^{n\times n}$ is a collection of read-once monotone span programs [25] specified by a matrix $\mathbf{M}$ of height n and

$$\mathsf{P}_{\text{MSP}}(\mathbf{x}, \mathbf{M}) = 1 \text{ iff } \mathbf{x} \text{ satisfies } \mathbf{M}$$

  Here, $\mathbf{x}$ satisfies $\mathbf{M}$ iff $(1, 0, \dots, 0)$ lies in the row span of $\{\mathbf{M}_j : x_j = 1\}$ where $\mathbf{M}_j$ is the j'th row of $\mathbf{M}$. In the context of ABE, this corresponds to key-policy ABE for access structures [20].

In Appendix D, we prove tight lower bounds for one-way communication complexity:

$$\mathsf{R}^{\mathsf{A}\to\mathsf{B}}(\mathsf{P}) = \Theta(n) \quad \text{and} \quad \mathsf{R}^{\mathsf{B}\to\mathsf{A}}(\mathsf{P}) = \Theta(n^2).$$

By Theorem 2, this means that any linear $(\ell_\mathsf{A}, \ell_\mathsf{B})$-CDS for both predicates must satisfy

$$\ell_\mathsf{A}(\ell_\mathsf{A} + \ell_\mathsf{B} + 1) = \Omega(n) \quad \text{and} \quad \ell_\mathsf{B}(\ell_\mathsf{A} + \ell_\mathsf{B} + 1) = \Omega(n^2).$$

This immediately yields

- $\ell_\mathsf{B} = \Omega(n)$ if $\ell_\mathsf{A} = O(1)$;
- $\ell_\mathsf{A} = \Omega(n^2)$ if $\ell_\mathsf{B} = O(1)$;
- $\max(\ell_\mathsf{A}, \ell_\mathsf{B}) = \Omega(n)$.

The third lower bound is tight, as we have matching upper bounds in [10, 48, 3] exhibiting a linear (n, n)-CDS for the predicate. It is open what the optimal parameters are when we keep either the key or the ciphertext size constant.

Acknowledgments. We would like to thank Amos Beimel, Yuval Ishai and Sophie Laplante for insightful and inspiring discussions.

References

[1] F. M. Ablayev. Lower bounds for one-way probabilistic communication complexity and their application to space complexity. *Theor. Comput. Sci.*, 157(2):139–159, 1996.

[2] A. Ada, A. Chattopadhyay, S. A. Cook, L. Fontes, M. Koucký, and T. Pitassi. The hardness of being private. *TOCT*, 6(1):1, 2014.

[3] N. Attrapadung. Dual system encryption via doubly selective security: Framework, fully secure functional encryption for regular languages, and more. In *EUROCRYPT*, pages 557–577, 2014.

[4] Z. Bar-Yossef, T. S. Jayram, R. Kumar, and D. Sivakumar. Information theory methods in communication complexity. In *IEEE Conference on Computational Complexity*, pages 93–102, 2002.

[5] A. Beimel. Secret-sharing schemes: A survey. In *Coding and Cryptology - Third International Workshop, IWCC 2011*, pages 11–46, 2011.

[6] A. Beimel and Y. Ishai. On the power of nonlinear secrect-sharing. In *Conference on Computational Complexity*, pages 188–202, 2001.

[7] A. Beimel, Y. Ishai, R. Kumaresan, and E. Kushilevitz. On the cryptographic complexity of the worst functions. In *TCC*, pages 317–342, 2014.

[8] D. Boneh and M. K. Franklin. Identity-based encryption from the Weil pairing. *SIAM J. Comput.*, 32(3):586–615, 2003.

[9] J. Chen and H. Wee. Fully, (almost) tightly secure IBE and dual system groups. In *CRYPTO (2)*, pages 435–460, 2013.

[10] J. Chen, R. Gay, and H. Wee. Improved dual system ABE in prime-order groups via predicate encodings. In *Eurocrypt*, pages 595–624, 2015.

[11] B. Chor and O. Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. Comput.*, 17(2):230–261, 1988.

[12] B. Chor and E. Kushilevitz. A zero-one law for boolean privacy. *SIAM J. Discrete Math.*, 4(1):36–47, 1991.

[13] C. Cocks. An identity based encryption scheme based on quadratic residues. In *IMA Int. Conf.*, pages 360–363, 2001.

[14] D. Data, M. Prabhakaran, and V. M. Prabhakaran. On the communication complexity of secure computation. In *CRYPTO (2)*, pages 199–216, 2014.

[15] U. Feige, J. Kilian, and M. Naor. A minimal model for secure computation. In *STOC*, pages 554–563, 1994.

[16] A. Fiat and M. Naor. Broadcast encryption. In *CRYPTO*, pages 480–491, 1993.

[17] S. Garg, A. Kumarasubramanian, A. Sahai, and B. Waters. Building efficient fully collusion-resilient traitor tracing and revocation schemes. In *ACM Conference on Computer and Communications Security*, pages 121–130, 2010.

[18] C. Gentry and A. Silverberg. Hierarchical ID-based cryptography. In *ASIACRYPT*, pages 548–566, 2002.

[19] Y. Gertner, Y. Ishai, E. Kushilevitz, and T. Malkin. Protecting data privacy in private information retrieval schemes. *J. Comput. Syst. Sci.*, 60(3):592–629, 2000.

[20] V. Goyal, O. Pandey, A. Sahai, and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *ACM Conference on Computer and Communications Security*, pages 89–98, 2006.

[21] J. Horwitz and B. Lynn. Toward hierarchical identity-based encryption. In *EUROCRYPT*, pages 466–481, 2002.

[22] Y. Ishai and H. Wee. Partial garbling schemes and their applications. In *ICALP (1)*, pages 650–662, 2014.

[23] M. Ito, A. Saito, and T. Nishizeki. Secret sharing schemes realizing general access structure. In *GLOBECOM*, pages 99–102, 1987.

[24] B. Kalyanasundaram and G. Schnitger. The probabilistic communication complexity of set intersection. *SIAM J. Discrete Math.*, 5(4):545–557, 1992.

[25] M. Karchmer and A. Wigderson. On span programs. In *Structure in Complexity Theory Conference*, pages 102–111, 1993.

[26] J. Katz, A. Sahai, and B. Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *EUROCRYPT*, pages 146–162, 2008.

[27] L. Kowalczyk and A. B. Lewko. Bilinear entropy expansion from the decisional linear assumption. In *CRYPTO*, 2015. To appear.

[28] I. Kremer, N. Nisan, and D. Ron. On randomized one-round communication complexity. *Computational Complexity*, 8 (1):21–49, 1999.

[29] E. Kushilevitz and N. Nisan. *Communication complexity*. Cambridge University Press, 1997. ISBN 978-0-521-56067-2.

[30] A. B. Lewko. Tools for simulating features of composite order bilinear groups in the prime order setting. In *EUROCRYPT*, pages 318–335, 2012. Also Cryptology ePrint Archive, Report 2011/490.

[31] A. B. Lewko and B. Waters. New techniques for dual system encryption and fully secure HIBE with short ciphertexts. In *TCC*, pages 455–479, 2010.

[32] A. B. Lewko and B. Waters. Unbounded HIBE and attribute-based encryption. In *EUROCRYPT*, pages 547–567, 2011.

[33] A. B. Lewko and B. Waters. Decentralizing attribute-based encryption. In *EUROCRYPT*, pages 568–588, 2011.

[34] A. B. Lewko and B. Waters. New proof methods for attribute-based encryption: Achieving full security through selective techniques. In *CRYPTO*, pages 180–198, 2012.

[35] A. B. Lewko, T. Okamoto, A. Sahai, K. Takashima, and B. Waters. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In *EUROCRYPT*, pages 62–91, 2010.

[36] A. Nayak. Optimal lower bounds for quantum automata and random access codes. In *FOCS*, pages 369–377, 1999. Also, CoRR quant-ph/9904093.

[37] I. Newman and M. Szegedy. Public vs. private coin flips in one round communication games. In *STOC*, pages 561–570, 1996.

[38] T. Okamoto and K. Takashima. Fully secure functional encryption with general relations from the decisional linear assumption. In *CRYPTO*, pages 191–208, 2010.

[39] T. Okamoto and K. Takashima. Achieving short ciphertexts or short secret-keys for adaptively secure general inner-product encryption. In *CANS*, pages 138–159, 2011. Also, Cryptology ePrint Archive, Report 2011/648.

[40] T. Okamoto and K. Takashima. Adaptively attribute-hiding (hierarchical) inner product encryption. In *EUROCRYPT*, pages 591–608, 2012. Also, Cryptology ePrint Archive, Report 2011/543.

[41] R. Ostrovsky and W. E. Skeith III. Communication complexity in algebraic two-party protocols. In *CRYPTO*, pages 379–396, 2008.

[42] A. A. Razborov. On the distributional complexity of disjointness. *Theor. Comput. Sci.*, 106(2):385–390, 1992.

[43] A. Sahai and B. Waters. Fuzzy identity-based encryption. In *EUROCRYPT*, pages 457–473, 2005.

[44] A. Shamir. Factoring numbers in O(log n) arithmetic steps. *Inf. Process. Lett.*, 8(1):28–31, 1979.

[45] A. Shamir. Identity-based cryptosystems and signature schemes. In *CRYPTO*, pages 47–53, 1984.

[46] V. Vaikuntanathan and P. N. Vasudevan. From statistical zero knowledge to secret sharing. Cryptology ePrint Archive, Report 2015/281, 2015.

[47] B. Waters. Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In *CRYPTO*, pages 619–636, 2009.

[48] H. Wee. Dual system encryption via predicate encodings. In *TCC*, pages 616–637, 2014.

[49] A. C. Yao. Some complexity questions related to distributive computing. In *STOC*, pages 209–213, 1979.

[50] A. C.-C. Yao. Lower bounds by probabilistic arguments. In *FOCS*, pages 420–428, 1983.

A Equivalence of two definitions of α-privacy

In [48] the notion of α-privacy is stated differently than in Section 2.1. We show that the two definition of α-privacy are in fact equivalent. That is, we show that the two following statements are equivalent:

1. For all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ such that $\mathsf{P}(x, y) = 0$, for all $\mathsf{C}^* : \{0,1\}^{\ell_{\mathsf{A}}} \times \{0,1\}^{\ell_{\mathsf{B}}} \to \mathcal{D}$,
$$\Pr_{w \leftarrow \mathcal{W}, \alpha \leftarrow_{\mathrm{R}} \mathcal{D}}\Big[\mathsf{C}^*\big(\mathsf{A}(x, w, \alpha), \mathsf{B}(y, w, \alpha)\big) = \alpha\Big] \leq \frac{1}{|\mathcal{D}|}.$$

2. For all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ such that $\mathsf{P}(x, y) = 0$, and for all $\alpha \in \mathcal{D}$, the joint distribution $\mathsf{A}(x, w, \alpha), \mathsf{B}(y, w, \alpha)$ *perfectly* hides α. That is, for all $\alpha, \alpha' \in \mathcal{D}$, the following joint distributions are *identically* distributed:
$$\big\{x, y, \alpha, \mathsf{A}(x, w, \alpha), \mathsf{B}(y, w, \alpha)\big\} \quad \text{and} \quad \big\{x, y, \alpha, \mathsf{A}(x, w, \alpha'), \mathsf{B}(y, w, \alpha')\big\}$$
where the randomness is taken over $w \leftarrow_{\mathrm{R}} \mathcal{W}$.

Proof. Throughout the proof, fix $(x, y) \in \mathcal{X} \times \mathcal{Y}$ such that $\mathsf{P}(x, y) = 0$.

1. implies 2. : Suppose that there exist $\alpha, \alpha' \in \mathcal{D}$ such that the following joint distributions are not identically distributed:

$$\big\{x, y, \alpha, \mathsf{A}(x, w, \alpha), \mathsf{B}(y, w, \alpha)\big\} \quad \text{and} \quad \big\{x, y, \alpha, \mathsf{A}(x, w, \alpha'), \mathsf{B}(y, w, \alpha')\big\}$$

where the randomness is taken over $w \leftarrow_{\mathrm{R}} \mathcal{W}$. Then, there exists an adversary $\mathcal{A}$ that can distinguish between these two distributions with advantage $\varepsilon > 0$, that is, such that

$$\Pr_{w \leftarrow_{\mathrm{R}} \mathcal{W}}[\mathcal{A}\big(x, y, \alpha, \mathsf{A}(x, w, \alpha), \mathsf{B}(y, w, \alpha)\big) = 1] - \Pr_{w \leftarrow_{\mathrm{R}} \mathcal{W}}[\mathcal{A}\big(x, y, \alpha, \mathsf{A}(x, w, \alpha'), \mathsf{B}(y, w, \alpha')\big) = 1] = \varepsilon$$

From $\mathcal{A}$, we build a function $\mathsf{C}^* : \{0,1\}^{\ell_{\mathsf{A}}} \times \{0,1\}^{\ell_{\mathsf{B}}} \to \mathcal{D}$ such that

$$\Pr_{w \leftarrow \mathcal{W}, b \leftarrow_{\mathrm{R}} \mathcal{D}}\Big[\mathsf{C}^*\big(\mathsf{A}(x, w, b), \mathsf{B}(y, w, b)\big) = b\Big] > \frac{1}{|\mathcal{D}|}.$$

The function C^* has (α, α', x, y) hard-wired into it and is defined as follows:

$$\mathsf{C}^*(m_{\mathsf{A}}, m_{\mathsf{B}}) := \begin{cases} \alpha \text{ if } \mathcal{A}(x, y, \alpha, m_{\mathsf{A}}, m_{\mathsf{B}}) = 1 \\ \alpha' \text{ if } \mathcal{A}(x, y, \alpha, m_{\mathsf{A}}, m_{\mathsf{B}}) = 0 \end{cases}$$

We have:

$$\begin{aligned}
&\Pr_{w\leftarrow\mathcal{W},b\leftarrow_{\mathrm{R}}\mathcal{D}}\Big[\mathsf{C}^*\big(\mathsf{A}(x,w,b),\mathsf{B}(y,w,b)\big)=b\Big]\\
&=\Pr_{w\leftarrow\mathcal{W},b\leftarrow_{\mathrm{R}}\mathcal{D}}\Big[\mathcal{A}\big(x,y,\alpha,\mathsf{A}(x,w,b),\mathsf{B}(y,w,b)\big)=1 \text{ and } b=\alpha\Big]\\
&+\Pr_{w\leftarrow\mathcal{W},b\leftarrow_{\mathrm{R}}\mathcal{D}}\Big[\mathcal{A}\big(x,y,\alpha,\mathsf{A}(x,w,b),\mathsf{B}(y,w,b)\big)=0 \text{ and } b=\alpha'\Big]\\
&=\Pr_{w\leftarrow\mathcal{W}}\Big[\mathcal{A}\big(x,y,\alpha,\mathsf{A}(x,w,\alpha),\mathsf{B}(y,w,\alpha)\big)=1\Big]\cdot 1/|\mathcal{D}|\\
&+\Pr_{w\leftarrow\mathcal{W}}\Big[\mathcal{A}\big(x,y,\alpha,\mathsf{A}(x,w,\alpha'),\mathsf{B}(y,w,\alpha')\big)=0\Big]\cdot 1/|\mathcal{D}|\\
&=\Pr_{w\leftarrow\mathcal{W}}\Big[\mathcal{A}\big(x,y,\alpha,\mathsf{A}(x,w,\alpha),\mathsf{B}(y,w,\alpha)\big)=1\Big]\cdot 1/|\mathcal{D}|\\
&+\Big(1-\Pr_{w\leftarrow\mathcal{W}}\Big[\mathcal{A}\big(x,y,\alpha,\mathsf{A}(x,w,\alpha'),\mathsf{B}(y,w,\alpha')\big)=1\Big]\Big)\cdot 1/|\mathcal{D}|\\
&=1/|\mathcal{D}|+1/|\mathcal{D}|\cdot\varepsilon\\
&>1/|\mathcal{D}|
\end{aligned}$$

2. implies 1. : Suppose that there exists a function $\mathsf{C}^*:\{0,1\}^{\ell_\mathsf{A}}\times\{0,1\}^{\ell_\mathsf{B}}\to\mathcal{D}$ such that

$$\Pr_{w\leftarrow\mathcal{W},b\leftarrow_{\mathrm{R}}\mathcal{D}}\Big[\mathsf{C}^*\big(\mathsf{A}(x,w,b),\mathsf{B}(y,w,b)\big)=b\Big]>\frac{1}{|\mathcal{D}|}.$$

We need to exhibit a pair $(\alpha,\alpha')\in\mathcal{D}^2$ together with a distinguisher $\mathcal{A}$ that is able to distinguish the two following distributions:

$$\big\{x,y,\alpha,\mathsf{A}(x,w,\alpha),\mathsf{B}(y,w,\alpha)\big\}\quad\text{and}\quad\big\{x,y,\alpha,\mathsf{A}(x,w,\alpha'),\mathsf{B}(y,w,\alpha')\big\}$$

where the randomness is taken over $w\leftarrow_{\mathrm{R}}\mathcal{W}$.
We choose an arbitrary $\alpha'\in\mathcal{D}$ (that is, the following is true for all $\alpha'\in\mathcal{D}$). Observe that

$$\begin{aligned}
&\Pr_{\alpha\leftarrow_{\mathrm{R}}\mathcal{D},w\leftarrow\mathcal{W}}\Big[\mathsf{C}^*\big(\mathsf{A}(x,w,\alpha),\mathsf{B}(y,w,\alpha)\big)=\alpha\Big]-\Pr_{\alpha\leftarrow_{\mathrm{R}}\mathcal{D},w\leftarrow\mathcal{W}}\Big[\mathsf{C}^*\big(\mathsf{A}(x,w,\alpha'),\mathsf{B}(y,w,\alpha')\big)=\alpha\Big]\\
&=\Pr_{\alpha\leftarrow_{\mathrm{R}}\mathcal{D},w\leftarrow\mathcal{W}}\Big[\mathsf{C}^*\big(\mathsf{A}(x,w,\alpha),\mathsf{B}(y,w,\alpha)\big)=\alpha\Big]-\frac{1}{|\mathcal{D}|}>0
\end{aligned}$$

By an averaging argument, this means that there exists $\alpha\in\mathcal{D}$ for which

$$\Pr_{w\leftarrow\mathcal{W}}\Big[\mathsf{C}^*\big(\mathsf{A}(x,w,\alpha),\mathsf{B}(y,w,\alpha)\big)=\alpha\Big]-\Pr_{w\leftarrow\mathcal{W}}\Big[\mathsf{C}^*\big(\mathsf{A}(x,w,\alpha'),\mathsf{B}(y,w,\alpha')\big)=\alpha\Big]>0$$

We can then construct a distinguisher $\mathcal{A}$ as follows:

$$\mathcal{A}(x,y,\alpha,m_\mathsf{A},m_\mathsf{B}):=\begin{cases}1 \text{ if } \mathsf{C}^*(m_\mathsf{A},m_\mathsf{B})=\alpha\\ 0 \text{ otherwise}\end{cases}$$

Clearly, the distinguisher $\mathcal{A}$ has a positive advantage in distinguishing the two distributions

$$\big\{x,y,\alpha,\mathsf{A}(x,w,\alpha),\mathsf{B}(y,w,\alpha)\big\}\quad\text{and}\quad\big\{x,y,\alpha,\mathsf{A}(x,w,\alpha'),\mathsf{B}(y,w,\alpha')\big\}$$

□

B Strongly linear CDS satisfy linear reconstruction

It is natural to restrict the functions A and B to be linear functions in w when $\mathcal{W}$ is a vector space. We give the notion of *strongly linear* CDS, similar to the predicate encoding definition given in [48].

Fix a prime p. Let $(\mathsf{A},\mathsf{B},\mathsf{C})$ be an $(\ell_{\mathsf{A}},\ell_{\mathsf{B}})$-CDS for $\mathsf{P}:\mathcal{X}\times\mathcal{Y}\to\{0,1\}$, where $\mathcal{X}$ and $\mathcal{Y}$ may depend on p. We say that $(\mathsf{A},\mathsf{B},\mathsf{C})$ is *strongly p-linear* if it satisfies the following properties:

(input domains.) $\mathcal{D} := \mathbb{Z}_p$, and $\mathcal{W} := \mathbb{Z}_p^{\ell_{\mathcal{W}}}$ for some integer $\ell_{\mathcal{W}}$.

(output domains.) The output of A and B are vectors over $\mathbb{Z}_p$: they belong to $\mathbb{Z}_p^{\ell_{\mathsf{A}}}$ and $\mathbb{Z}_p^{\ell_{\mathsf{B}}}$, respectively.

(Alice linearity.) For all $x\in\mathcal{X}$, $\mathsf{A}(x,\cdot,\cdot)$ is a linear function of (w,α).

(Bob linearity.) For all $y\in\mathcal{Y}$, $\mathsf{B}(y,\cdot,\cdot)$ is a linear function of (w,α).

Lemma 2. *For any predicate* P*, any prime* p*, and any strongly* p*-linear CDS* $(\mathsf{A},\mathsf{B},\mathsf{C})$ *for* P*, we have the following property:*

(linear $\boldsymbol{\alpha}$-reconstruction.) *For all* (x,y) *such that* $\mathsf{P}(x,y)=1$*, there is a* $\mathbb{Z}_p$*-linear function* $\mathsf{L}_{x,y}:\mathbb{Z}_p^{\ell_A}\times\mathbb{Z}_p^{\ell_B}\to\mathbb{Z}_p$ *such that for all* $w\in\mathcal{W}$ *and for all* $\alpha\in\mathcal{D}$*:*

$$\mathsf{L}_{x,y}(\mathsf{A}(x,w,\alpha),\mathsf{B}(y,w,\alpha))=\alpha$$

Proof. By p-linearity, we know that for every $x\in\mathcal{X},y\in\mathcal{Y}$, there exists matrices $\mathbf{A}_x,\mathbf{B}_y$ such that for every $\alpha\in\mathcal{D},\mathbf{w}\in\mathcal{W}$:

$$\mathsf{A}(x,w,\alpha)=\mathbf{A}_x\begin{pmatrix}\alpha\\ \mathbf{w}\end{pmatrix}\quad\text{and}\quad\mathsf{B}(y,w,\alpha)=\mathbf{B}_y\begin{pmatrix}\alpha\\ \mathbf{w}\end{pmatrix}$$

This means that

$$\begin{pmatrix}\alpha\\ m_{\mathsf{A}}\\ m_{\mathsf{B}}\end{pmatrix}=\begin{pmatrix}\mathbf{e}_1\\ \mathbf{A}_x\\ \mathbf{B}_y\end{pmatrix}\begin{pmatrix}\alpha\\ \mathbf{w}\end{pmatrix}$$

By a standard argument from linear algebra, this means that for all x,y:

- if $\mathbf{e}_1\in\mathrm{span}\begin{pmatrix}\mathbf{A}_x\\ \mathbf{B}_y\end{pmatrix}$, α is uniquely determined given $m_{\mathsf{A}},m_{\mathsf{B}},x,y$, for all $\alpha,\mathbf{w}$. This would violate α-privacy, which means $\mathsf{P}(x,y)=1$.
- if $\mathbf{e}_1\notin\mathrm{span}\begin{pmatrix}\mathbf{A}_x\\ \mathbf{B}_y\end{pmatrix}$, α is uniformly random given $m_{\mathsf{A}},m_{\mathsf{B}},x,y$, for a uniformly random $(\alpha,\mathbf{w})$. This would violate α-reconstruction, which means $\mathsf{P}(x,y)=0$.

Hence, $\mathsf{P}(x,y)=1$ iff $\mathbf{e}_1\in\mathrm{span}\begin{pmatrix}\mathbf{A}_x\\ \mathbf{B}_y\end{pmatrix}$. Moreover, if $\mathbf{e}_1\in\mathrm{span}\begin{pmatrix}\mathbf{A}_x\\ \mathbf{B}_y\end{pmatrix}$, then there exists a row vector $\mathbf{v}\in\mathbb{Z}_p^{\ell_{\mathsf{A}}+\ell_{\mathsf{B}}}$ such that

$$\mathbf{e}_1=\mathbf{v}\begin{pmatrix}\mathbf{A}_x\\ \mathbf{B}_y\end{pmatrix}\Longrightarrow\mathbf{v}\begin{pmatrix}m_{\mathsf{A}}\\ m_{\mathsf{B}}\end{pmatrix}=\mathbf{e}_1\begin{pmatrix}\alpha\\ \mathbf{w}\end{pmatrix}=\alpha$$

The linear function $\mathsf{L}_{x,y}$ is that given by multiplying by the vector $\mathbf{v}$. □

C Concrete Strongly Linear CDS

For completeness, we extend here several strongly linear CDS (as defined in Appendix B) from [3, 48] in order to provide a trade-off between the lengths of the two messages. In these CDS, it suffices for Bob to know α, and we remove α from Alice's inputs (this only makes the upper bounds stronger).

Throughout this appendix, p is a given prime. For all vectors $\mathbf{u}$, we denote by $|\mathbf{u}|$ the size of $\mathbf{u}$, by u_i the i'th coordinate of vector $\mathbf{u}$, and for any matrix $\mathbf{M}$, we denote by $M_{i,j}$ the (i,j)'th entry of $\mathbf{M}$. Finally, we denote by $\mathbb{Z}_p^{\leq n}$ (resp. $\mathbb{Z}_p^{<n}$) the set of vectors of length at most n (resp. less than n).

Equality. Here, $\mathcal{X} = \mathcal{Y} := \mathbb{Z}_p$ and

$$\mathsf{P}_{\mathrm{Eq}}(x,y) = 1 \text{ iff } x = y$$

strongly linear $(1,1)$*-CDS:*

- $\mathcal{W} := \mathbb{Z}_p \times \mathbb{Z}_p$.
- $\mathsf{A}(x,(u,v)) := ux + v \in \mathbb{Z}_p$
- $\mathsf{B}(y,(u,v),\alpha) := uy + v + \alpha \in \mathbb{Z}_p$
- $\mathsf{C}(x,y,c,d) := d - c$

For α-privacy, we exploit the fact that $ux + v, uy + v$ are pairwise independent when $x \neq y$.

Remark 4 (multi-bit secret CDS). Note that we can handle larger space $\mathcal{D}$ such as $\mathcal{D} := \mathbb{Z}_p^d$, for $d > 1$, by simply using d CDS for equality, each with independent randomness. That is:

- $\mathcal{W} := (\mathbb{Z}_p \times \mathbb{Z}_p)^d$.
- $\mathsf{A}\big(x,((u_1,v_1),\dots,(u_d,v_d))\big) := (u_1x + v_1,\dots,u_dx + v_d) \in \mathbb{Z}_p^d$
- $\mathsf{B}\big(y,((u_1,v_1),\dots,(u_d,v_d)),(\alpha_1,\dots,\alpha_d)\big) := (\alpha_1 + u_1y + v_1,\dots,\alpha_d + u_dx + v_d) \in \mathbb{Z}_p^d$
- $\mathsf{C}(x,y,\mathbf{c},\mathbf{d}) := \mathbf{d} - \mathbf{c}$

Note that this multi-bit secret CDS for equality is optimal, according to Theorem 4.

Inner Product (IP)

Predicate [26]: Here, $\mathcal{X} = \mathcal{Y} := \mathbb{Z}_p^n$ and

$$\mathsf{P}(\mathbf{x},\mathbf{y}) = 1 \text{ iff } \mathbf{x}^\top\mathbf{y} = 0$$

strongly linear $(n-t+1,t+1)$*-CDS:*

- $\mathcal{W} := \mathbb{Z}_p^n \times \mathbb{Z}_p \times \mathbb{Z}_p$;
- $\mathsf{A}(\mathbf{x},(\mathbf{w},u,u')) := \left(\sum_{i=1}^t x_i w_i + u', u\begin{pmatrix} x_{t+1} \\ \vdots \\ x_n \end{pmatrix} + \begin{pmatrix} w_{t+1} \\ \vdots \\ w_n \end{pmatrix}\right) \in \mathbb{Z}_p^{n-t+1}$
- $\mathsf{B}(\mathbf{y},(\mathbf{w},u,u')) := \left(u\begin{pmatrix} y_1 \\ \vdots \\ y_t \end{pmatrix} + \begin{pmatrix} w_1 \\ \vdots \\ w_t \end{pmatrix}, \sum_{i=t+1}^n y_i w_i - u' + \alpha\right) \in \mathbb{Z}_p^{t+1}$
- $\mathsf{C}(\mathbf{x},\mathbf{y},(c',\mathbf{c}),(\mathbf{d},d')) := d' + c' - \begin{pmatrix} x_1 \\ \vdots \\ x_t \end{pmatrix}^\top \mathbf{d} - \mathbf{c}^\top \begin{pmatrix} y_{t+1} \\ \vdots \\ y_n \end{pmatrix}$

Correctness: Suppose $\mathsf{P}(\mathbf{x},\mathbf{y}) = 1$. Observe that

$$\begin{aligned}\mathsf{C}\big(\mathbf{x},\mathbf{y},\mathsf{A}(\mathbf{x},(\mathbf{w},u,u')),\mathsf{B}(\mathbf{y},(\mathbf{w},u,u'))\big) &= (\sum_{i=t+1}^{n} y_i w_i + \alpha - u') + (\sum_{i=1}^{t} x_i w_i + u') - \sum_{i=1}^{t} x_i w_i - \sum_{i=t+1}^{n} y_i w_i - u\mathbf{x}^\top\mathbf{y} \\ &= \alpha - u\mathbf{x}^\top\mathbf{y} \\ &= \alpha\end{aligned}$$

We use the fact that $\mathbf{x}^\top\mathbf{y} = 0$ in the last equality.

Privacy: The constructions exploit the following simple algebraic fact: given $\mathbf{x},\mathbf{y}, u\mathbf{x}+\mathbf{w},\mathbf{y}^\top\mathbf{w}+\alpha$,

- if $\mathbf{x}^\top\mathbf{y} = 0$, then we can recover α.
- if $\mathbf{x}^\top\mathbf{y} \neq 0$, then α is masked by $u\mathbf{x}^\top\mathbf{y}$.

Index predicate (Broadcast encryption)

Predicate [16]: Here, $\mathcal{X} := \{0,1\}^n, \mathcal{Y} := [n]$ and

$$\mathsf{P}(\mathbf{x}, i) = 1 \text{ iff } x_i = 1$$

That is, $\mathbf{x}$ is the characteristic vector of a subset of $[n]$. For notational convenience, we rewrite the predicate as follows: $\mathcal{X} := (\{0,1\}^{n/t})^t, \mathcal{Y} := [t]\times[n/t]$ and

$$\mathsf{P}((\mathbf{x}_1,\ldots,\mathbf{x}_t),(i_1,i_2)) = 1 \text{ iff } \mathbf{x}_{i_1}^\top\mathbf{e}_{i_2} = 1$$

where (i_1, i_2) is the unique pair of integers satisfying $i = (i_1 - 1)\cdot n/t + i_2$ and $0 < i_2 \leq n/t$, and $(\mathbf{e}_1,\ldots,\mathbf{e}_{n/t})$ is the standard basis of $\mathbb{Z}_q^{n/t}$.

strongly linear $(t, n/t)$-CDS [17]:

- $\mathcal{W} := \mathbb{Z}_p^t \times \mathbb{Z}_p^{n/t}$.
- $\mathsf{A}(\mathbf{x},(\mathbf{w},\mathbf{u})) := (w_1 + \mathbf{x}_1^\top\mathbf{u},\ldots,w_t + \mathbf{x}_t^\top\mathbf{u}) \in \mathbb{Z}_p^t$
- $\mathsf{B}((i_1,i_2),(\mathbf{w},\mathbf{u})) := (w_{i_1} + \alpha)\cdot\mathbf{e}_{i_2} + \mathbf{u} \in \mathbb{Z}_p^{n/t}$
- $\mathsf{C}(\mathbf{x},(i_1,i_2),\mathbf{c},\mathbf{d}) := \mathbf{x}_{i_1}^\top\mathbf{d} - c_{i_1}$

Correctness. Suppose $\mathsf{P}(\mathbf{x},(i_1,i_2)) = 1$. Observe that

$$\mathsf{C}\big(\mathbf{x},(i_1,i_2),\mathsf{A}(\mathbf{x},(\mathbf{w},\mathbf{u})),\mathsf{B}((i_1,i_2),(\mathbf{w},\mathbf{u}))\big) = \mathbf{x}_{i_1}^\top\big((w_{i_1}+\alpha)\cdot\mathbf{e}_{i_2} + \mathbf{u}\big) - w_{i_1} - \mathbf{x}_{i_1}^\top\mathbf{u} = \alpha$$

Privacy. Privacy follows readily from the fact that:

- For all $j \neq i_1$, $w_j + \mathbf{x}_j^\top\mathbf{u}$ reveals no information about $\mathbf{u}$;
- If $\mathbf{x}_{i_1}^\top\mathbf{e}_{i_2} = 0$, then α is perfectly hidden given $\mathbf{x}_{i_1}$, $(\alpha + w_{i_1})\cdot\mathbf{e}_{i_2} + \mathbf{u}$, and $w_{i_1} + \mathbf{x}_{i_1}^\top\mathbf{u}$.

Prefix. Here, $\mathcal{X} := \{0,1\}^n, \mathcal{Y} := \{0,1\}^{\leq n}$ and

$$\mathsf{P}_{\text{prefix}}(\mathbf{x},\mathbf{y}) = 1 \text{ iff } \mathbf{y} \text{ is a prefix of } \mathbf{x}$$

For notational convenience, we rewrite the predicate as follows: $\mathcal{X} := (\{0,1\}^{n/t})^t$, $\mathcal{Y} = (\{0,1\}^{n/t})^{<t} \times \{0,1\}^{\leq n/t}$ and

$$\mathsf{P}((\mathbf{x}_1,\ldots,\mathbf{x}_t),(\mathbf{y}_1,\ldots,\mathbf{y}_{t_2})) = 1 \text{ iff } t \geq t_2 \wedge \mathbf{x}_i = \mathbf{y}_i \text{ for } i = 1,\ldots,t_2-1 \ \wedge \mathbf{y}_{t_2} \text{ is a prefix of } \mathbf{x}_{t_2}.$$

We also write $\mathcal{W} := \mathbb{Z}_p^n \times \mathbb{Z}_p^t$ as $(\mathbb{Z}_p^{n/t} \times \mathbb{Z}_p)^t$.

strongly linear $(t, n/t)$-CDS:

- $\mathcal{W} := (\mathbb{Z}_p^{n/t} \times \mathbb{Z}_p)^t$.
- $\mathsf{A}(\mathbf{x},(\mathbf{w},\mathbf{u})) := (\mathbf{x}_1^\top\mathbf{w}_1 + u_1,\ldots,\mathbf{x}_t^\top\mathbf{w}_t + u_t) \in \mathbb{Z}_p^t$.
- $\mathsf{B}(\mathbf{y},(\mathbf{w},\mathbf{u}),\alpha) := \left(\sum_{i=1}^{t_2-1}(\mathbf{y}_i^\top\mathbf{w}_i + u_i) + \binom{\mathbf{y}_{t_2}}{\mathbf{0}}^\top \mathbf{w}_{t_2} + u_{t_2} + \alpha, (\mathbf{w}_{t_2})_{|\mathbf{y}_{t_2}|+1},\ldots,(\mathbf{w}_{t_2})_{n/t}\right) \in \mathbb{Z}_p^{n/t-|\mathbf{y}_{t_2}|+1}$, where $\binom{\mathbf{y}_{t_2}}{\mathbf{0}} \in \mathbb{Z}_p^{n/t}$, and $(\mathbf{w}_{t_2})_{|\mathbf{y}_{t_2}|+1},\ldots,(\mathbf{w}_{t_2})_{n/t}$ are the last $n/t - |\mathbf{y}_{t_2}|$ coordinates of $\mathbf{w}_{t_2}$.
- $\mathsf{C}(\mathbf{x},\mathbf{y},\mathbf{c},(d',\mathbf{d})) := d' + \sum_{j=1}^{n/t-|\mathbf{y}_{t_2}|} d_j(\mathbf{x}_{t_2})_{|\mathbf{y}_{t_2}|+j} - \sum_{i=1}^{t_2} c_i$, where $(\mathbf{x}_{t_2})_{|\mathbf{y}_{t_2}|+1},\ldots,(\mathbf{x}_{t_2})_{n/t}$ are the last $n/t - |\mathbf{y}_{t_2}|$ coordinates of $\mathbf{x}_{t_2}$.

Correctness. Suppose $\mathbf{y}$ is a prefix of $\mathbf{x}$. Observe that

$$\begin{aligned}
\mathsf{C}(\mathbf{x},\mathbf{y},\mathsf{A}(\mathbf{x},(\mathbf{w},\mathbf{u})),\mathsf{B}(\mathbf{y},(\mathbf{w},\mathbf{u}),\alpha)) &= \alpha + \sum_{i=1}^{t_2-1}(\mathbf{y}_i^\top\mathbf{w}_i + u_i) + \binom{\mathbf{y}_{t_2}}{\mathbf{0}}^\top \mathbf{w}_{t_2} + u_{t_2} + \sum_{i=|\mathbf{y}_{t_2}|+1}^{n/t}(\mathbf{x}_{t_2})_i(\mathbf{w}_{t_2})_i - \sum_{i=1}^{t_2}(\mathbf{x}_i^\top\mathbf{w}_i + u_i) \\
&= \alpha + \binom{\mathbf{y}_{t_2}}{\mathbf{0}}^\top \mathbf{w}_{t_2} + u_{t_2} + \sum_{i=|\mathbf{y}_{t_2}|+1}^{n/t}(\mathbf{x}_{t_2})_i(\mathbf{w}_{t_2})_i - (\mathbf{x}_{t_2}^\top\mathbf{w}_{t_2} + u_{t_2}) \\
&= \alpha
\end{aligned}$$

We use the fact that $\mathbf{y}_i = \mathbf{x}_i$ for all $i = 1,\ldots,t_2-1$ in the second equality, and the fact that $\mathbf{y}_{t_2}$ is a prefix of $\mathbf{x}_{t_2}$ in the last equality.

Privacy. Suppose $\mathbf{y}$ is not a prefix of $\mathbf{x}$. One of the following is true:

1. $\mathbf{y}_j \neq \mathbf{x}_j$ for some $j \leq t_2 - 1$, which implies that $\mathbf{x}_j^\top\mathbf{w}_j + u_j$ and $\mathbf{y}_j^\top\mathbf{w}_j + u_j$ are pairwise independent.
2. $\mathbf{y}_{t_2}$ is not a prefix of $\mathbf{x}_{t_2}$, which implies that $\mathbf{x}_{t_2}^\top\mathbf{w}_{t_2} + u_{t_2}$ and $\binom{\mathbf{y}_{t_2}}{\mathbf{0}}^\top \mathbf{w}_{t_2} + u_{t_2}$ are pairwise independent.

Privacy follows readily.

Disjointness. Here, $\mathcal{X} = \mathcal{Y} := \{S \subseteq [n]\}$ and

$$\mathsf{P}_{\text{disj}}(X,Y) = 1 \text{ iff } X \cap Y = \emptyset.$$

As a warm-up, here is a linear $(n, n+1)$-CDS: Alice sends $\{w_i : i \notin X\}$, Bob sends $\{w_j : j \notin Y\}$ along with $\alpha + w_1 + \cdots + w_n$. (The messages are padded to vectors of length n and $n+1$ respectively.) Correctness and privacy follow readily from the fact that

$$X \cap Y = \emptyset \Longleftrightarrow \bar{X} \cup \bar{Y} = [n]$$

where $\bar{X}$ is the set complement of X.

strongly linear $(t+1, n-t+1)$*-CDS:*

- $\mathcal{W} := \mathbb{Z}_p^{n+1}$
- $\mathsf{A}(X, (\mathbf{w}, u)) := \big(\{w_i : i \in \bar{X} \cap \{1, \ldots, t\}\}, u + \sum_{i \in X \cap \{t+1,\ldots,n\}} w_i\big)$ (padded to a vector in $\mathbb{Z}_p^{t+1}$)
- $\mathsf{B}(Y, (\mathbf{w}, u), \alpha) := \big(\alpha + u + \sum_{i \in Y \cap \{1,\ldots,t\}} w_i, \{w_i : i \in \bar{Y} \cap \{t+1, \ldots, n\}\}\big)$ (padded to a vector in $\mathbb{Z}_p^{n-t+1}$)
- $\mathsf{C}(X, Y, (\mathbf{c}, c'), (d', \mathbf{d})) := d' - \sum_{i \in Y \cap \{1,\ldots,t\}} w_i - c' + \sum_{j \in X \cap \{t+1,\ldots,n\}} w_j$.

Correctness is straightforward.

Privacy. Suppose $\mathsf{P}(X, Y) = 0$ so that $X \cap Y \neq \emptyset$. Fix $j \in X \cap Y$. One of the following is true:

1. $j \in \{1, \ldots, t\}$, therefore, α is masked by w_j.
2. $j \in \{t+1, \ldots, n\}$, therefore, u is masked by w_j, and α is masked by u.

D Lower bound on the communication complexity of predicates.

For completeness, we describe here the lower bounds on the one-way communication complexity of $\mathsf{P}_{\text{prefix}}$ and P_{MSP}. Note that the upper bounds are given by the trivial protocols.

First, we present a straightforward lemma that allows us to lower bound the randomized one-way communication complexity of a predicate, where the correctness is over the public coins used in the protocol, by its *distributional one-way communication complexity*, where we consider deterministic protocols whose correctness is over a probability distribution over the *inputs*. This lemma is implied by a theorem due to Yao [50].

Lemma 3 ([50]). *Let* $\mathsf{P} : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ *be a predicate, and* (A, B) *be a one-way* $(\mathsf{A} \to \mathsf{B})$ *communication protocol for* P*, with*

$$\mathsf{A} : \mathcal{X} \times \mathcal{W} \to \{0,1\}^\ell, \quad \mathsf{B} : \mathcal{Y} \times \mathcal{W} \times \{0,1\}^\ell \to \{0,1\}.$$

For all distributions μ *over* $\mathcal{X} \times \mathcal{Y}$*, there exists a deterministic protocol* $(\widetilde{\mathsf{A}}, \widetilde{\mathsf{B}})$ *of same communication complexity as* (A, B)*, such that*

$$\Pr_{(x,y) \leftarrow_{\mathrm{R}} \mu}[\widetilde{\mathsf{B}}(y, \widetilde{\mathsf{A}}(x)) = \mathsf{P}(x, y)] \geq 1 - \delta/2$$

where $\delta := \Pr_{(x,y) \leftarrow_{\mathrm{R}} \mu}[\mathsf{P}(x, y) = 0]$*. A similar statement holds for one-way* $(\mathsf{B} \to \mathsf{A})$ *protocols.*

Proof. We know that

- for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ such that $\mathsf{P}(x, y) = 1$,

$$\Pr_{w \leftarrow_{\mathrm{R}} \mathcal{W}}[\mathsf{B}(y, w, \mathsf{A}(x, w)) = 1] = 1$$

- for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$ such that $\mathsf{P}(x, y) = 0$,

$$\Pr_{w \leftarrow_{\mathrm{R}} \mathcal{W}}[\mathsf{B}(y, w, \mathsf{A}(x, w)) = 0] \geq 1/2.$$

Therefore,

$$\begin{aligned}
&\Pr_{w \leftarrow_{\mathrm{R}} \mathcal{W}, (x,y) \leftarrow_{\mathrm{R}} \mu}[\mathsf{B}(y, w, \mathsf{A}(x, w)) = \mathsf{P}(x, y)] \\
&= \Pr_{w \leftarrow_{\mathrm{R}} \mathcal{W}, (x,y) \leftarrow_{\mathrm{R}} \mu}[\mathsf{B}(y, w, \mathsf{A}(x, w)) = 1 | \mathsf{P}(x, y) = 1] \cdot \Pr_{(x,y) \leftarrow_{\mathrm{R}} \mu}[\mathsf{P}(x, y) = 1] \\
&+ \Pr_{w \leftarrow_{\mathrm{R}} \mathcal{W}, (x,y) \leftarrow_{\mathrm{R}} \mu}[\mathsf{B}(y, w, \mathsf{A}(x, w)) = 0 | \mathsf{P}(x, y) = 0] \cdot \Pr_{(x,y) \leftarrow_{\mathrm{R}} \mu}[\mathsf{P}(x, y) = 0] \\
&\geq 1 \cdot (1 - \delta) + 1/2 \cdot \delta \\
&= 1 - \delta/2
\end{aligned}$$

By an averaging argument, there exists a $w^* \in \mathcal{W}$ such that

$$\Pr_{(x,y) \leftarrow_{\mathrm{R}} \mu}[\mathsf{B}(y, w^*, \mathsf{A}(x, w^*)) = \mathsf{P}(x, y)] \geq 1 - \delta/2.$$

The deterministic protocol $(\widetilde{\mathsf{A}}, \widetilde{\mathsf{B}})$ is the protocol (A, B) where the random string $w \in \mathcal{W}$ is fixed to w^*. □

The second lemma, due to Nayak [36], gives a lower bound on the distributional complexity of the *augmented* index predicate, which is an easier version of the index predicate, and which will be used to prove lower bounds on the one-way communication complexity of $\mathsf{P}_{\text{prefix}}$.

Augmented index. Here, $\mathcal{X} := \{0,1\}^n, \mathcal{Y} := [n] \times \{0,1\}^{\leq n}$ and

$$\mathsf{P}_{\text{Aug.index}}(\mathbf{x}, (i, \mathbf{y})) = 1 \text{ iff } (\mathbf{y} = (x_1, \ldots, x_{i-1})) \wedge (x_i = 1).$$

Lemma 4 ([36]). *Let $\mathcal{X} := \{0,1\}^n, \mathcal{Y} := [n] \times \{0,1\}^{\leq n}$ and $\varepsilon \in [0,1]$. Let μ^* the following input distribution*

$$\mu^* : \mathbf{x} \leftarrow_{\mathrm{R}} \{0,1\}^n, i \leftarrow_{\mathrm{R}} [n], \mathbf{y} := (x_1, \ldots, x_{i-1}).$$

Then, for any deterministic one-way $(\mathsf{A} \to \mathsf{B})$ protocol (A, B) of communication complexity ℓ with

$$\Pr_{(\mathbf{x},(i,\mathbf{y})) \leftarrow_{\mathrm{R}} \mu^*}\Big[\mathsf{B}\big((i, \mathbf{y}), \mathsf{A}(\mathbf{x})\big) = \mathsf{P}_{\text{Aug.index}}\big(\mathbf{x}, (i, \mathbf{y})\big)\Big] \geq 1 - \varepsilon,$$

we have

$$\ell \geq (1 - H(\varepsilon))n$$

where H is the binary entropy function.

Also, for any deterministic one-way $(\mathsf{B} \to \mathsf{A})$ *protocol* (A, B) *of communication complexity* ℓ *with*

$$\Pr_{(\mathbf{x},(i,\mathbf{y}))\leftarrow_{\mathrm{R}}\mu^*}\left[\mathsf{A}\big(\mathbf{x}, \mathsf{B}(i,\mathbf{y})\big) = \mathsf{P}_{\text{Aug.index}}\big(\mathbf{x},(i,\mathbf{y})\big)\right] \geq 1-\varepsilon,$$

we have

$$\ell = \Omega(\log n)$$

The same bounds hold for $\mathsf{P}_{\text{index}}$ *on the uniform distribution* $\mathbf{x} \leftarrow_{\mathrm{R}} \{0,1\}^n, i \leftarrow_{\mathrm{R}} [n]$.

We can now prove the lower bounds on the one-way communication complexity of $\mathsf{P}_{\text{prefix}}$ and P_{MSP}.

Prefix. Here, $\mathcal{X} := \{0,1\}^n$, $\mathcal{Y} := \{0,1\}^{\leq n}$ and

$$\mathsf{P}_{\text{prefix}}(\mathbf{x},\mathbf{y}) = 1 \text{ iff } \mathbf{y} \text{ is a prefix of } \mathbf{x}$$

Theorem 5. *For all n, we have*

$$\mathsf{R}^{\mathsf{A}\to\mathsf{B}}(\mathsf{P}_{\text{prefix}}) \geq (1-H(1/4))n \;\; \textit{and} \;\; \mathsf{R}^{\mathsf{B}\to\mathsf{A}}(\mathsf{P}_{\text{prefix}}) = \Omega(\log n)$$

Proof. We reduce $\mathsf{P}_{\text{Aug.index}}$ to $\mathsf{P}_{\text{prefix}}$. In high level, we show that a deterministic one-way communication protocol for $\mathsf{P}_{\text{prefix}}$ for a specific input distribution μ, implies a deterministic one-way communication protocol with the same communication for $\mathsf{P}_{\text{Aug.index}}$ and the distribution μ^* in Lemma 4. Hence, the lower bound on the communication complexity of $\mathsf{P}_{\text{Aug.index}}$, implies a lower bound for the distributional one-way communication complexity of $\mathsf{P}_{\text{prefix}}$ for the distribution μ. By Lemma 3, this implies the same lower bound for the randomized one-way communication complexity of $\mathsf{P}_{\text{prefix}}$.

We start by defining the following input distribution μ for $\mathsf{P}_{\text{prefix}}$,

$$\mu : \mathbf{x} \leftarrow_{\mathrm{R}} \{0,1\}^n, i \leftarrow_{\mathrm{R}} [n], \mathbf{y} := (x_1,\ldots,x_{i-1},1).$$

Note that for all $(\mathbf{x},\mathbf{y})$ in the support of μ, we have $(\mathbf{x},\mathbf{y}) = (\mathbf{x},(x_1,\ldots,x_{i-1},1))$ for some $i \in [n]$ and

$$\mathsf{P}_{\text{prefix}}(\mathbf{x},(x_1,\ldots,x_{i-1},1)) = 1 \text{ iff } \mathsf{P}_{\text{Aug.index}}(\mathbf{x},(i,(x_1,\ldots,x_{i-1})) = 1$$

First, we prove the statement about $\mathsf{R}^{\mathsf{A}\to\mathsf{B}}(\mathsf{P}_{\text{prefix}})$. Suppose there exists a deterministic one-way protocol (A, B) for $\mathsf{P}_{\text{prefix}}$ with communication ℓ, such that

$$\Pr_{(\mathbf{x},\mathbf{y})\leftarrow_{\mathrm{R}}\mu}[\mathsf{B}((x_1,\ldots,x_{i-1},1),\mathsf{A}(\mathbf{x})) = \mathsf{P}_{\text{prefix}}(\mathbf{x},(x_1,\ldots,x_{i-1},1))] \geq 3/4.$$

From (A, B), we build a deterministic one-way protocol $(\widetilde{\mathsf{A}}, \widetilde{\mathsf{B}})$ for $\mathsf{P}_{\text{Aug.index}}$ on the input distribution

$$\mu^* : \mathbf{x} \leftarrow_{\mathrm{R}} \{0,1\}^n, i \leftarrow_{\mathrm{R}} [n], \mathbf{y} := (x_1,\ldots,x_{i-1}).$$

For all $(\mathbf{x},(i,\mathbf{y})) = (\mathbf{x},(i,(x_1,\ldots,x_{i-1})))$ in the support of μ^*, $(\widetilde{\mathsf{A}}, \widetilde{\mathsf{B}})$ is defined by:

- $\widetilde{\mathsf{A}}(\mathbf{x}) := \mathsf{A}(\mathbf{x})$;
- for all messages $m \in \{0,1\}^\ell$, $\widetilde{\mathsf{B}}((i,(x_1,\ldots,x_{i-1})),m) := \mathsf{B}((x_1,\ldots,x_{i-1},1),m)$.

It is easy to check that

$$\begin{aligned}
&\Pr_{(\mathbf{x},(i,\mathbf{y}))\leftarrow_{\mathrm{R}}\mu^*}\left[\widetilde{\mathsf{B}}((i,(x_1,\dots,x_{i-1})),\widetilde{\mathsf{A}}(\mathbf{x}))=\mathsf{P}_{\text{Aug.index}}(\mathbf{x},(i,(x_1,\dots,x_{i-1})))\right]\\
&=\Pr_{(\mathbf{x},\mathbf{y})\leftarrow_{\mathrm{R}}\mu}\left[\mathsf{B}((x_1,\dots,x_{i-1},1),\mathsf{A}(\mathbf{x}))=\mathsf{P}_{\text{prefix}}(\mathbf{x},(x_1,\dots,x_{i-1},1))\right]\\
&\geq 3/4
\end{aligned}$$

Therefore, by Lemma 4,

$$\ell \geq (1-H(1/4))n.$$

Second, we prove the statement about $\mathsf{R}^{\mathsf{B}\to\mathsf{A}}(\mathsf{P}_{\text{prefix}})$. Suppose there exists a deterministic one-way protocol (A,B) for $\mathsf{P}_{\text{prefix}}$ with communication ℓ, such that

$$\Pr_{(\mathbf{x},\mathbf{y})\leftarrow_{\mathrm{R}}\mu}[\mathsf{A}(\mathbf{x},\mathsf{B}(x_1,\dots,x_{i-1},1))=\mathsf{P}_{\text{prefix}}(\mathbf{x},(x_1,\dots,x_{i-1},1))]\geq 3/4.$$

From (A,B), we build a deterministic one-way protocol $(\widetilde{\mathsf{A}},\widetilde{\mathsf{B}})$ for $\mathsf{P}_{\text{Aug.index}}$ on the input distribution

$$\mu^*:\mathbf{x}\leftarrow_{\mathrm{R}}\{0,1\}^n, i\leftarrow_{\mathrm{R}}[n],\mathbf{y}:=(x_1,\dots,x_{i-1}).$$

For all $(\mathbf{x},(i,\mathbf{y}))=(\mathbf{x},(i,(x_1,\dots,x_{i-1})))$ in the support of μ^*, $(\widetilde{\mathsf{A}},\widetilde{\mathsf{B}})$ is defined by:

- $\widetilde{\mathsf{B}}(i,(x_1,\dots,x_{i-1})):=B((x_1,\dots,x_{i-1},1))$;
- for all messages $m\in\{0,1\}^\ell$, $\widetilde{\mathsf{A}}(\mathbf{x},m):=\mathsf{A}(\mathbf{x},m)$.

It is easy to check that

$$\begin{aligned}
&\Pr_{(\mathbf{x},(i,\mathbf{y}))\leftarrow_{\mathrm{R}}\mu^*}\left[\widetilde{\mathsf{A}}(\mathbf{x},\widetilde{\mathsf{B}}(i,(x_1,\dots,x_{i-1})))=\mathsf{P}_{\text{Aug.index}}(\mathbf{x},(i,(x_1,\dots,x_{i-1})))\right]\\
&=\Pr_{(\mathbf{x},\mathbf{y})\leftarrow_{\mathrm{R}}\mu}\left[\mathsf{A}(\mathbf{x},\mathsf{B}(x_1,\dots,x_{i-1},1)=\mathsf{P}_{\text{prefix}}(\mathbf{x},(x_1,\dots,x_{i-1},1))\right]\\
&\geq 3/4
\end{aligned}$$

Therefore, by Lemma 4,

$$\ell \geq \Omega(\log n).$$

□

Read-once monotone span programs [25]. Here, $\mathcal{X}:=\{0,1\}^n,\mathcal{Y}:=\mathbb{Z}_p^{n\times n}$, and

$$\mathsf{P}_{\text{MSP}}(\mathbf{x},\mathbf{M})=1 \text{ iff } \mathbf{x} \text{ satisfies } \mathbf{M}$$

where $\mathbf{x}$ satisfies $\mathbf{M}$ iff $(1,0,\dots,0)$ lies in the row span of $\{\mathbf{M}_j : x_j=1\}$ where $\mathbf{M}_j$ is the j'th row of $\mathbf{M}$.

Theorem 6. *For all n, we have*

$$\mathsf{R}^{\mathsf{B}\to\mathsf{A}}(\mathsf{P}_{\text{MSP}})\geq\frac{1-H(1/4)}{4}n^2.$$

Proof. We reduce $\mathsf{P}_{\text{index}}$ to P_{MSP}. In high level, we show that a deterministic one-way ($\mathsf{B}\to\mathsf{A}$) communication protocol for P_{MSP} for a specific input distribution μ, implies a deterministic one-way communication protocol with the same communication for $\mathsf{P}_{\text{index}}$ with input size $\Omega(n^2)$ and the uniform input distribution. By Lemma 4, the lower bound on the communication complexity of $\mathsf{P}_{\text{index}}$ implies a lower bound for the distributional one-way communication complexity of P_{MSP} for the distribution μ. By Lemma 3, this implies the same lower bound for the randomized one-way communication complexity of P_{MSP}.

For notational convenience, let n be even. We define the following input distribution μ for P_{MSP}.

The matrix $\mathbf{M} \in \mathbb{Z}_p^{n\times n}$, is distributed as follows:

- For all $i \leq n/2$, $M_{i,1} = 1$.
- For all $i \leq n/2$, and all $2 \leq j \leq n/2+1$, $M_{i,j} \leftarrow_{\text{R}} \{0,1\}$.
- For all $i > n/2$, $M_{i,i-n/2+1} = 1$.
- For all $i > n/2$, and all $j \neq i - n/2 + 1$, $M_{i,j} = 0$.
- for all $i \in [n]$, and all $j > n/2+1$, $M_{i,j} = 0$.

The vector $\mathbf{x} \in \{0,1\}^n$ is distributed as follows:

- For a random $i' \leftarrow_{\text{R}} [n/2]$, we have $x_{i'} := 1$. For all $i \leq n/2$, $i \neq i'$, we have $x_i := 0$.
- For a random $i'' \leftarrow_{\text{R}} [n/2]$, we have $x_{n/2+i''} := 0$. For all $i > n/2$, $i \neq i''$, $x_i := 1$.

The matrix $\mathbf{M}$ and the vector $\mathbf{x}$ are shown in figure 3.

$$\mathbf{M} := \left(\begin{array}{c|ccc|cc} 1 & * & * & * & 0 & 0 \\ 1 & * & * & * & 0 & 0 \\ 1 & * & * & * & 0 & 0 \\ \hline 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{array}\right) \qquad \mathbf{x} := \left(\begin{array}{c} 0 \\ 1 \\ 0 \\ \hline 1 \\ 1 \\ 0 \end{array}\right)$$

Fig. 3. Examples of $\mathbf{M}$ with $n = 6$ and $\mathbf{x}$ with $i' = 2$ and $i'' = 3$.

Note that for all $(\mathbf{x}, \mathbf{M})$ in the support of μ, for which $x_{i'} = 1$ and $x_{n/2+i''} = 0$ for some $i', i'' \in [n/2]$, we have

$$\mathsf{P}_{\text{MSP}}(\mathbf{x}, \mathbf{M}) = 1 \text{ iff } M_{i',i''+1} = 0.$$

Now, suppose there exists a deterministic one-way ($\mathsf{B} \to \mathsf{A}$) communication protocol (A, B) for P_{MSP} with communication ℓ, where $\mathsf{A} : \mathcal{X} \times \mathcal{W} \times \{0,1\}^\ell \to \{0,1\}$, $\mathsf{B} : \mathcal{Y} \times \mathcal{W} \to \{0,1\}^\ell$, and

$$\Pr_{(\mathbf{x},\mathbf{M})\sim\mu}[\mathsf{A}(\mathbf{x}, \mathsf{B}(\mathbf{M})) = \mathsf{P}_{\text{MSP}}(\mathbf{x}, \mathbf{M})] \geq 3/4.$$

From (A, B), we build a deterministic one-way ($\mathsf{A} \to \mathsf{B}$) protocol $(\widetilde{\mathsf{A}}, \widetilde{\mathsf{B}})$ for $\mathsf{P}_{\text{index}}$ on the uniform input distribution μ^* defined as follows:

- $\widetilde{\mathbf{M}} \leftarrow_{\text{R}} \{0,1\}^{n/2\times n/2}$

- $(i', i'') \leftarrow_{\mathrm{R}} [n/2] \times [n/2]$

Note that this is a distribution which is uniform on strings of length $n^2/4$ and uniform on an index in $[n^2/4]$. Moreover, for all inputs in the support of μ^* we have

$$\mathsf{P}_{\text{index}}(\widetilde{\mathbf{M}}, (i', i'')) = 1 \text{ iff } \mathsf{P}_{\text{MSP}}(\mathbf{x}, \mathbf{M}) = 1.$$

We can now define the protocol $(\widetilde{\mathsf{A}}, \widetilde{\mathsf{B}})$, for all $(\widetilde{\mathbf{M}}, (i', i''))$ in the support of μ^*, as:

- $\widetilde{\mathsf{A}}(\widetilde{\mathbf{M}}) := \mathsf{B}(\mathbf{M})$, where $\mathbf{M}$ is in the support of μ and it is such that for all $i \le n/2$, and all $2 \le j \le n/2 + 1$, $M_{i,j} = \widetilde{\mathbf{M}}_{i,j-1}$.
- for all messages $m \in \{0,1\}^{\ell}$, $\widetilde{\mathsf{B}}((i', i''), m) := \mathsf{A}(\mathbf{x}, m)$, where $\mathbf{x}$ is in the support of μ and it is such that $x_{i'} := 1$ and $x_{n/2+i''} := 0$.

It is easy to check that

$$\Pr_{(\widetilde{\mathbf{M}}, (i', i'')) \leftarrow_{\mathrm{R}} \mu^*} \left[\widetilde{\mathsf{B}}((i', i''), \widetilde{\mathsf{A}}(\widetilde{\mathbf{M}})) = \mathsf{P}_{\text{index}}(\widetilde{\mathbf{M}}, (i', i'')) \right] = \Pr_{(\mathbf{x}, \mathbf{M}) \leftarrow_{\mathrm{R}} \mu} \left[\mathsf{A}(\mathbf{x}, \mathsf{B}(\mathbf{M})) = \mathsf{P}_{\text{MSP}}(\mathbf{x}, \mathbf{M}) \right] \ge 3/4.$$

Therefore, by Lemma 4,

$$\ell \ge \frac{1 - H(1/4)}{4} n^2.$$

□

Theorem 7. *For all n, we have*

$$\mathsf{R}^{\mathsf{A} \to \mathsf{B}}(\mathsf{P}_{\text{MSP}}) \ge (1 - H(1/4))n.$$

Proof. We again reduce $\mathsf{P}_{\text{index}}$ to P_{MSP}. In high level, we show that a deterministic one-way ($\mathsf{A} \to \mathsf{B}$) communication protocol for P_{MSP} for a specific input distribution μ, implies a deterministic one-way communication protocol with the same communication for $\mathsf{P}_{\text{index}}$ with input size n and the uniform input distribution. By Lemma 4, the lower bound on the communication complexity of $\mathsf{P}_{\text{index}}$ implies a lower bound for the distributional one-way communication complexity of P_{MSP} for the distribution μ. By Lemma 3, this implies the same lower bound for the randomized one-way communication complexity of P_{MSP}.

We define the following input distribution μ for P_{MSP}.

The matrix $\mathbf{M} \in \mathbb{Z}_p^{n \times n}$, is distributed as follows:

- For a uniformly random $i \in [n]$, $M_{i,1} = 1$. All other elements of the matrix are 0.

The vector $\mathbf{x} \in \{0,1\}^n$ is distributed uniformly.

Note that for all $(\mathbf{x}, \mathbf{M})$ in the support of μ with $i \in [n]$ such that $M_{i,1} = 1$

$$\mathsf{P}_{\text{MSP}}(\mathbf{x}, \mathbf{M}) = 1 \text{ iff } x_i = 1.$$

Now, suppose there exists a deterministic one-way ($\mathsf{A} \to \mathsf{B}$) communication protocol (A, B) for P_{MSP} with communication ℓ, where $\mathsf{A} : \mathcal{X} \times \mathcal{W} \to \{0,1\}^{\ell}$ and $\mathsf{B} : \mathcal{Y} \times \mathcal{W} \times \{0,1\}^{\ell} \to \{0,1\}$, and

$$\Pr_{(\mathbf{x},\mathbf{M})\sim\mu}[\mathsf{B}(\mathbf{M},\mathsf{A}(\mathbf{x})) = \mathsf{P}_{\mathrm{MSP}}(\mathbf{x},\mathbf{M})] \geq 3/4.$$

From (A,B), we build a deterministic one-way $(\mathsf{A}\to\mathsf{B})$ protocol $(\widetilde{\mathsf{A}},\widetilde{\mathsf{B}})$ for $\mathsf{P}_{\mathrm{index}}$ on the uniform input distribution μ^* defined as follows:

$$\mu^* : \mathbf{x} \leftarrow_{\mathrm{R}} \{0,1\}^n, i \leftarrow_{\mathrm{R}} [n]$$

We can easily define the protocol $(\widetilde{\mathsf{A}},\widetilde{\mathsf{B}})$, for all $(\mathbf{x}, i)$ in the support of μ^*, as:

- $\widetilde{\mathsf{A}}(\mathbf{x}) := \mathsf{A}(\mathbf{x})$.
- for all messages $m \in \{0,1\}^\ell$, $\widetilde{\mathsf{B}}(i,m) := \mathsf{B}(\mathbf{M},m)$, where $\mathbf{M}$ is in the support of μ, and it is such that $M_{i,1} = 1$.

It is easy to check that

$$\Pr_{(\mathbf{x},i)\leftarrow_{\mathrm{R}}\mu^*}\left[\widetilde{\mathsf{B}}(i,\widetilde{\mathsf{A}}(\mathbf{x})) = \mathsf{P}_{\mathrm{index}}(\mathbf{x},i)\right] = \Pr_{(\mathbf{x},\mathbf{M})\leftarrow_{\mathrm{R}}\mu}\left[\mathsf{B}(\mathbf{M},\mathsf{A}(\mathbf{x})) = \mathsf{P}_{\mathrm{MSP}}(\mathbf{x},\mathbf{M})\right] \geq 3/4.$$

Therefore, by Lemma 4,

$$\ell \geq (1 - H(1/4))n.$$

□