An Oblivious Transfer Protocol with Log-Squared Communication

*** Third Public Draft. May 21, 2004 ***[⋆]

Helger Lipmaa

Laboratory for Theoretical CS, Department of CS&E
Helsinki University of Technology, P.O.Box 5400, FIN-02015 HUT, Espoo, Finland
{helger}@tcs.hut.fi

Abstract. We propose a $\binom{N}{1}$-computationally-private information retrieval protocol (CPIR), based on the Damgård-Jurik public-key cryptosystem from PKC 2001, with total communication $\Theta(\log^2 N) \cdot k$. We show how to extend it to a four-round $\binom{N}{1}$-OT protocol with total communication $\Theta(\log^2 N)k$. These protocols are, respectively, the first $\binom{N}{1}$-CPIR and $\binom{N}{1}$-OT protocols with logarithmic sender-side and polylogarithmic client-side communication. We show that there is no straightforward methodology to improve the asymptotic communication of our protocols.
Keywords: computationally-private information retrieval, homomorphic public-key cryptosystem, oblivious transfer.

1 Introduction

During a (single-server) $\binom{N}{1}$-computationally-private information retrieval ($\binom{N}{1}$-CPIR) protocol, Chooser retrieves an entry from Sender's database $\mu = (\mu[1], \ldots, \mu[N])$, so that a computationally bounded Sender will not obtain any information on which element was retrieved. In the case of the *(computationally chooser-private and information-theoretically sender-private) $\binom{N}{1}$-oblivious transfer protocol* ($\binom{N}{1}$-OT) it is additionally required that Chooser's retrieved information should be restricted to one entry from the database. Both CPIR and OT have important cryptographic applications and have therefore been considered by many researchers.

The first $\binom{N}{1}$-CPIR protocol with polylogarithmic (in N) communication was proposed by Cachin, Micali and Stadler in [CMS99]. The security of their protocol is based on a rather new Φ Assumption, and both its total communication is $\Omega(\log^{\geq 4} N)k$, where k is the security parameter.

Slightly earlier, in 1998, Julien P. Stern [Ste98] proposed a family—that we call $\mathsf{HomCPIR}(\alpha)$—of $\binom{N}{1}$-CPIR protocols, based on an (almost) arbitrary semantically secure additively homomorphic public-key cryptosystem. In the asymptotically near-optimal case $\alpha = \sqrt{\log_\eta N}$, $\mathsf{HomCPIR}(\alpha)$ has super-logarithmic total communication $(1 + o(1))\eta\sqrt{\log_\eta N} \cdot \eta^{\sqrt{\log_\eta N}} \cdot k$, where η is the ciphertext expansion ratio of the underlying homomorphic cryptosystem. ($\eta = 2$ in the case of Paillier's cryptosystem.)

[⋆] Corresponds to the submitted version

We propose a modification of $\mathsf{HomCPIR}(\alpha)$ that takes full advantage of the length-flexible semantically secure homomorphic cryptosystems of Damgård and Jurik [DJ01,DJ03]. In the new $\binom{N}{1}$-CPIR, that we call $\mathsf{LFCPIR}(\alpha)$, one can transfer $k_s = k^s$ bits, for some $s \geq 1$. $\mathsf{LFCPIR}(\alpha)$ is secure under the Decisional Composite Residuosity Assumption [Pai99]. The total communication of $\mathsf{LFCPIR}(\log N)$ is $\Theta(\log^2 N) \cdot k$. Moreover, $\mathsf{LFCPIR}(\log N)$ requires not more than $\Theta(N)$ k_s-bit exponentiations by Sender.

We then look at the possibilities to decrease the communication even further. We show that under two alternative additional assumptions on the used cryptosystem, there exist a CPIR with the asymptotically optimal total communication $\Theta(\log N)k$. The first assumption is that an encryption of a small constant in a smaller ciphertext-space can be transferred to an encryption of the same constant in a larger ciphertext-space even by somebody who does not know the private key. The second, alternative, assumption is that the cryptosystem is ring-homomorphic. We conjecture that such cryptosystems exist—with he first assumption being somewhat more plausible than the second assumption—, though we show that the Damgård-Jurik cryptosystem from PKC 2001 [DJ01] does not satisfy either of the two requirements.

An important albeit direct corollary from the construction of the $\mathsf{LFCPIR}(\alpha)$ is that one can construct a four-round (or two-round, assuming the existence of non-interactive zero-knowledge protocols) $\binom{N}{1}$-OT protocol with total communication $\Theta(\log^2 N) \cdot k$.

Road-map. In Section 2, we introduce the reader to notation and preliminaries. In Section 3, we propose the new CPIR protocols. In Section 4, we propose the hypothetical $\Theta(\log N) \cdot k$ protocols and study their (im)possibility. In Section 5, we compare different CPIR protocols from the efficiency viewpoint and discuss possible extensions. In Section 6, we describe the new $\binom{N}{1}$-OT protocol.

2 Preliminaries

Notation. For a natural number t, let $[t]$ denote the set $\{1, \dots, t\}$. Let $\{i \stackrel{?}{=} j\} = 1$ if $i = j$ and 0, otherwise. All logarithms in this paper will be on base 2, unless explicitly mentioned. Let e be the base of the natural logarithm. Let $x||y$ denote the concatenation of binary strings x and y. Let W be Lambert's W function, i.e., W satisfies the functional identity $W(x)e^{W(x)} = x$.

For a distribution (random variable) X, let $x \leftarrow X$ denote the assignment of x according to X. We often identify sets with the uniform distributions on them, and algorithms with their output distributions, assuming that the algorithm that outputs this distribution is clear from the context or just straightforward to construct. The statistical difference of two distributions X and Y over the discrete support U is defined as $\Delta(X||Y) := \max_{S \subseteq U} |\Pr[X \in S] - \Pr[Y \in S]|$. Let k be the security parameter. Denote a run of interactive protocol between A who has private input a and random tape r_a and between B who has private input b and a random tape r_b as $(A, B)[a, r_a; b, r_b]$. Throughout this paper, we let N to denote Sender's database size.

Semantically secure homomorphic cryptosystems. Public-key cryptosystem is a triple $\Pi = (G, E, D)$, where G is the key generation algorithm, E is the encryption algorithm and D is the decryption algorithm. For fixed Π and private key x, let $C(x)$ be the ciphertext space, let $R(x)$ be the randomness space and let $P(x)$ be the plaintext space. Define $\mathsf{Adv}^{\mathsf{pkcsem}}_{\Pi,k}(A) := |\Pr[(m_0, m_1) \leftarrow A(1^k), b \leftarrow \{0,1\} : A(1^k, m_0, m_1, E_K(m_b; R(x))) = b] - \frac{1}{2}|$. We say that Π is *semantically secure* when $\mathsf{Adv}^{\mathsf{pkcsem}}_{\Pi,k}(A)$ is negligible in k for any probabilistic polynomial-time machine A. That is, Π is *semantically secure* if it is difficult for a polynomially bounded adversary to distinguish between random encryptions of two elements, chosen by herself.

Π is an *(additively) homomorphic public-key cryptosystem* if for any key pair (K, x), any $m, m' \in P(x) = \mathbb{Z}_{|P(x)|}$ and any $r, r' \in R(x)$, $(E_K(m; r) \cdot E_K(m'; r') = E_K(m + m'; r \circ r')$, where $\circ$ is a groupoid operation in $R(x)$, and $+$ is a group operation in $P(x)$. A few homomorphic cryptosystems [OU98,NS98,Pai99,DJ01,DJ03] are believed to be semantically secure under reasonable complexity assumptions.

Throughout this paper, we assume that $k = |P(x)|$ is the security parameter. Let $\eta := \lceil |C(x)|/k \rceil$ be the *ciphertext expansion ratio* of Π; $\eta = 2$ for the Paillier cryptosystem [Pai99] and Damgård-Jurik cryptosystem from PKC 2001 [DJ01] and $\eta = 3$ for the Okamoto-Uchiyama cryptosystem [OU98] and $\eta \in \{2, 3\}$ for another cryptosystem by Damgård and Jurik [DJ03].

Damgård-Jurik length-flexible cryptosystem. The Damgård-Jurik cryptosystem $\mathsf{DJ01}$ from PKC 2001 [DJ01] has flexible plaintext length parameterised by a constant $s \geq 1$. For fixed s, the plaintext space is $P_s(x) := \mathbb{Z}_{n^s}$, the randomness space is $R(x) = \mathbb{Z}_n^*$ and the ciphertext space is $C_s(x) := \mathbb{Z}_{n^{s+1}}^*$, thus $|C_s(x)|/|P_s(x)| \approx 1 + 1/s$ and $\eta = 2$. Encryption is defined as $E_K^s(m; r) := (1 + n)^m \cdot r^{n^s} \mod n^{s+1}$. To decrypt $c = E_K^s(m; r)$, one first computes $c \mod n = r^{n^s} \mod n$, and then retrieves r and thus also $c' = (1 + n)^m \mod n^{s+1}$. After that, a simple algorithm from [DJ01] is used to retrieve m from c'. The Damgård-Jurik cryptosystem is additively homomorphic since $E_K^s(m_1; r_1)E_K^s(m_2; r_2) = E_K^s(m_1 + m_2; r_1 r_2)$, and semantically secure assuming that the Decisional Composite Residuosity Problem (DCRP) is hard [DJ01]. Importantly, the same key K can be used to encrypt messages of different length, and $C_s(x) \subseteq P_{s+1}(x)$. For a fixed Π, denote encryption and decryption with public key K and a fixed s by E_K^s and D_K^s. We assume that $k = \log |P_1(x)|$, and we denote $k_s := \log |P_s(x)|$. In the case of the $\mathsf{DJ01}$ cryptosystem, $k_s \approx sk$.

Another cryptosystem by Damgård and Jurik [DJ03] has also flexible-length plaintexts, but is slightly less efficient with $|C_s(x)|/|P_s(x)| \approx 1 + 2/s$. Moreover, its security is based on both the DCRP and on the Decisional Diffie-Hellman Problem.

Computationally-private information retrieval (CPIR). During a (single-server) $\binom{N}{1}$*-computationally-private information retrieval* ($\binom{N}{1}$-CPIR) protocol, Chooser fetches $\mu[\sigma]$ from the database $\mu = (\mu[1], \ldots, \mu[N])$, $\mu[i] \in \mathcal{D}$ for some fixed $\mathcal{D}$, so that a computationally bounded Sender does not know which entry Chooser is learning.

More formally, an $\binom{N}{1}$-CPIR protocol is defined as follows. Define Chooser Cho's view $\mathsf{view}_{\mathsf{Cho}}[\sigma, r_{\mathsf{Cho}}; \mu, r_{\mathsf{Sen}}]$ in the $\binom{N}{1}$-CPIR protocol $(\mathsf{Cho}, \mathsf{Sen})[\sigma, r_{\mathsf{Cho}}; \mu, r_{\mathsf{Sen}}]$ as the concatenation of its private input σ, random tape r_{Cho}, the protocol transcript, and its private output $\mu[\sigma]$. The view of Sender Sen is defined dually. For an algorithm A executing the sender's part in the protocol $(\mathsf{Cho}, A)[\sigma, r_{\mathsf{Cho}}; \mu, r_A]$, define $\mathsf{Adv}^{\mathsf{cpir}}_{\mathsf{Cho},k}(A) := |\Pr[(\sigma_0, \sigma_1, \mu) \leftarrow A(1^k, r_A), b \leftarrow [0, 1] : A(1^k, r_A, \mathsf{view}_A[\sigma_b, r_{\mathsf{Cho}}; \mu, r_A]) = b'] - \frac{1}{2}|$ to be the advantage over random coin-tossing that A has in predicting which of the two possible choices σ_0 and σ_1 was used by the chooser, after observing an execution of the protocol $(\mathsf{Cho}, A)[\sigma_b, r_{\mathsf{Cho}}; \mu, r_A]$. We call $(\mathsf{Cho}, \mathsf{Sen})$ *a* $\binom{N}{1}$*-CPIR protocol*, if the next two conditions hold: (a) Cho's output after executing protocol $(\mathsf{Cho}, \mathsf{Sen})[\sigma, r_{\mathsf{Cho}}; \mu, r_{\mathsf{Sen}}]$ is equal to μ_σ for any $r_{\mathsf{Cho}}, r_{\mathsf{Sen}}$, and (b) $\mathsf{Adv}^{\mathsf{cpir}}_{\mathsf{Cho},k}(A)$ is negligible for any probabilistic polynomial-time algorithm A.

The first (single-server) $\binom{N}{1}$-CPIR with sublinear total communication $O(2^{\sqrt{\log N \cdot \log k}})$ was proposed by Kushilevitz and Ostrovsky [KO97]. The first $\binom{N}{1}$-CPIR with polylogarithmic total communication was proposed by Cachin, Micali and Stadler in [CMS99], but it has total communication $\Omega(\log^{\geq 4} N) \cdot k$ and its security is based on the relatively little-studied Φ Assumption. (The precise total communication of this protocol depends on a constant, existence of which is guaranteed, given the Φ Assumption.)

Stern's $\binom{N}{1}$-CPIR. The Kushilevitz-Ostrovsky $\binom{N}{1}$-CPIR was improved by Julien P. Stern [Ste98] (and later slightly refined by Chang [Cha04]). While the security of the Kushilevitz-Ostrovsky $\binom{N}{1}$-CPIR is based on the quadratic residuosity assumption and enables the chooser only to retrieve a single bit with every query, the Stern $\binom{N}{1}$-CPIR is based on an arbitrary semantically secure additively homomorphic cryptosystem Π with ciphertext expansion ratio $\eta \geq 2$, and enables the chooser to retrieve up to $k = |P(x)|$ bits with every query. We next give a self-contained description of Stern's $\binom{N}{1}$-CPIR. For this, we recast Stern's $\binom{N}{1}$-CPIR in the contemporary language of homomorphic cryptosystems, and by using the hypercube-based description from [Cha04].

Stern's $\binom{N}{1}$-CPIR is a family of CPIR-s with a parameter $\alpha \geq 1$. We call this family $\mathsf{HomCPIR}\,(\alpha)$. The main idea behind $\mathsf{HomCPIR}\,(\alpha)$ is as follows. Look at the database μ as at an α-dimensional hypercube with its elements $\mu(i_1, \dots, i_\alpha)$ addressed by coordinates $(i_1, \dots, i_\alpha)$; Chooser's query σ has thus also α coordinates, $\sigma = (\sigma_1, \dots, \sigma_\alpha)$. For every dimension i and every coordinate j along this dimension, Chooser sends an encryption $\beta_{i,j}$ of $\{j \stackrel{?}{=} \sigma_i\}$ to Sender. Fix a dimension, say 1. Sender constructs an $(\alpha - 1)$-dimensional database as follows. For all dimensions $j \neq 1$ and for all coordinates i_j along those dimensions, set the $(i_2, \dots, i_j)$th element of the new database to be equal to the product $\prod_t \beta_{1,t}^{\mu(t, i_2, \dots, i_\alpha)}$. Due to the homomorphic properties of Π, this product is an encryption of $\mu(\sigma_1, i_2, \dots, i_\alpha)$. The same process is repeated with the $(\alpha - 1)$-dimensional database of encrypted values $E_K(\mu(\sigma_1, i_2, \dots, i_\alpha))$, with the $(\alpha - 2)$-dimensional database of doubly-encrypted values $E_K(E_K(\mu(\sigma_1, \sigma_2, i_3, \dots, i_\alpha)))$, until Sender obtains an α-times encryption of $\mu(\sigma_1, \dots, \sigma_\alpha)$. Sender sends this value to the Chooser, who decrypts it α times.

Now, $\mathsf{HomCPIR}\,(\alpha)$ deviates from this basic idea since Π must have ciphertext expansion ratio $\eta \geq 2$. To account with this, $\mathsf{HomCPIR}\,(\alpha)$ creates η^j $(\alpha - j)$-dimensional intermediate databases of j-times encrypted ciphertexts. Sender transfers η^α α-times encrypted ciphertexts to Chooser, who recursively combines their decryptions to obtain $\mu(\sigma_1, \dots, \sigma_\alpha)$. Thus, $\mathsf{HomCPIR}\,(\alpha)$ has chooser-side communication $\eta\alpha N^{1/\alpha} k$ bits and sender-side communication $\eta^\alpha k$ bits. One can minimise sender-side communication by minimising α. However, even in the close-to-optimal case when $\alpha = \sqrt{\log N}$, the total communication $\Theta\left(\sqrt{\log N} \cdot 2^{\sqrt{\log N}}\right) \cdot k$ is still super-logarithmic.

([Ste98] erroneously claims that the total communication of $\mathsf{HomCPIR}\left(\sqrt{\log N}\right)$ is $\Theta\left(\eta^{\sqrt{\log N}}\right) \cdot k$.) Thus,

Proposition 1 ([Ste98,Cha04]). *Let Π be a semantically secure homomorphic cryptosystem. (a) Let $\alpha \geq 1$. W.l.o.g., assume that Sender's database $\mu = (\mu[1], \dots, \mu[N])$ contains $N = \ell^{\alpha}$ entries of length $\leq k$. There exists an $\binom{N}{1}$-CPIR protocol $\mathsf{HomCPIR}(\alpha)$ with chooser-side communication $\alpha N^{1/\alpha}\eta k$ and sender-side communication $\eta^{\alpha} k$ bits. (b) Fix $\alpha = \sqrt{\log_\eta N}$. There exists an $\binom{N}{1}$-CPIR protocol $\mathsf{HomCPIR}\left(\sqrt{\log_\eta N}\right)$ with total communication $(\eta\sqrt{\log_\eta N} + 1)\eta^{\sqrt{\log_\eta N}} k$ bits.*

Here, $\alpha = \sqrt{\log_\eta N}$ is the "on limit" solution of the equation $\alpha N^{1/\alpha} = \eta^{\alpha}$. Note that for any $N \in \mathbb{Z}$ and $\alpha = \sqrt{\log_\eta N}$, the sender-side communication is smaller than the chooser-side communication.

Oblivious Transfer. An $\binom{N}{1}$-CPIR is an *(computationally chooser-private and information-theoretically sender-private) $\binom{N}{1}$-OT protocol* (with computational chooser-privacy) when also Sender's privacy is guaranteed. For the formal definition, we make the comparison to the ideal implementation, using a trusted third party that receives μ from Sender, receives σ from Chooser, and tells Chooser $\mu[\sigma]$. Let $\mathcal{I}$ be the set of legal indices, $\mathcal{I} = [1, N]$. We assume that $\mu[\sigma]$ is garbage (i.e., a random value from some μ-independent set T) if $\sigma \notin \mathcal{I}$. (Alternatively, one could require that Chooser halts when $\sigma \notin \mathcal{I}$.)

We define the security by showing that for every algorithm A, one can define a simulator S that, given only private input σ, random tape r_A, and private output $\mu[\sigma]$ of A, generates output that is statistically indistinguishable from the view of A that reacts with the honest sender Sen. More precisely, for a sender Sen and an algorithm S, define $\mathsf{Adv}^{\mathsf{otsen}}_{\mathsf{Sen},k}(A, S) := \Delta\left(S(1^k, \sigma, r_A, \mu[\sigma]) || \mathsf{view}_A[\sigma, r_A; \mu, r_{\mathsf{Sen}}]\right)$. We say that the $\binom{N}{1}$-CPIR protocol $(\mathsf{Cho}, \mathsf{Sen})$ is a *statistically sender-private $\binom{N}{1}$-oblivious transfer protocol* if for every (not necessarily probabilistic polynomial-time) A there exists a (not necessarily probabilistic polynomial-time) S, such that $\mathsf{Adv}^{\mathsf{otsen}}_{\mathsf{Sen},k}(A, S)$ is negligible in k. As usually, sender-privacy is perfect when $\mathsf{Adv}^{\mathsf{otsen}}_{\mathsf{Sen},k}(A, S) = 0$.

As argued, e.g., in [NP01, Section 2.1.2], an oblivious transfer protocol does not have to guarantee the correctness (even if Cho is honest but Sen is not, Cho will still receive Sen's input $\mu[\sigma]$). Following this convention, also we will leave it up to the application protocols to provide security in this sense.

Stern [OP98] also proposed an $\binom{N}{1}$-OT protocol [OP98] that requires four rounds and $\Theta\left(\eta\sqrt{\log_\eta N} \cdot \eta^{\sqrt{\log_\eta N}}\right) \cdot k$ communication. Some other $\binom{N}{1}$-OT protocols [NP01,AIR01,Lip03] require two rounds but $\Theta(N) \cdot k$ communication. Asymptotically, the most communication-efficient—although with $\Omega(\log^4 N) \cdot k$ communication—previous $\binom{N}{1}$-OT protocol results when applying the Naor-Pinkas methodology from [NP99] to the $\binom{N}{1}$-CPIR from [CMS99]. However, Stern's $\binom{N}{1}$-OT protocol is more communication-efficient for all relevant database sizes.

3 New $\binom{N}{1}$-CPIR with Log-Squared Communication

Next, we use the length-flexibility property of the DJ01 cryptosystem to improve both the concrete and the asymptotic total communication of HomCPIR (α), by presenting a family LFCPIR(α) (*length-flexible private information retrieval*) of $\binom{N}{1}$-CPIR protocols.

Let $\Pi = (G, E, D)$ be the DJ01 cryptosystem. Assume that $\mathcal{D} \subseteq P_s(x)$ for some fixed s. That is, $k_s = s \log n$ for some public key n, and $|\mathcal{D}| \leq n^s$. Fix $\alpha \geq 1$. Consider the database $\mu = (\mu[1], \dots, \mu[N])$ as an α-dimensional $\ell_1 \times \cdots \times \ell_\alpha$ hyper-rectangle. Thus, every element of μ is indexed by a tuple $(i_1, \dots, i_\alpha)$. If $\alpha > 1$, Sender computes an $(\alpha - 1)$-dimensional database w_1 as follows: for every $(i_2, \dots, i_\alpha)$, let $w_1(i_2, \dots, i_\alpha) \leftarrow \prod_{t \in [\ell_1]} \beta_{1,t}^{\mu(t, i_2, \dots, i_\alpha)}$, where $D_K^s(\beta_{1,t}) = \{t \stackrel{?}{=} \sigma_1\}$ and $\beta_{1,t}$ is obtained from Chooser. Thus, $D_K^s(w_1(i_2, \dots, i_\alpha)) = \mu(\sigma_1, i_2, \dots, i_\alpha)$. If $\alpha > 2$, Sender repeats this process recursively on the database w_1 to obtain an $(\alpha - 2)$-dimensional database w_2, such that $D_K^{s+1}(w_2(i_3, \dots, i_\alpha)) = w_1(\sigma_2, i_3, \dots, i_\alpha)$, then—if $\alpha > 3$—an $(\alpha - 3)$-dimensional database w_3, such that $D_K^{s+2}(w_3(i_4, \dots, i_\alpha)) = w_2(\sigma_3, i_4, \dots, i_\alpha)$, etc.

Thus, the basic idea of the new protocol is similar to that of HomCPIR (α), as described in Section 2. After α steps, Sender has a single element w_α, an α-times encrypted $\mu(\sigma_1, \dots, \sigma_\alpha)$. He then returns w_α to Chooser, who decrypts it α times to obtain $\mu(\sigma_1, \dots, \sigma_\alpha)$. Since the Damgård-Jurik cryptosystem has flexible-length plaintexts, instead of dividing every intermediate ciphertext into η chunks (as in the case of Stern's $\binom{N}{1}$-CPIR), we just increase the length of the plaintexts. Therefore, it suffices for Sender to just send one value w_α, with length $|w_\alpha| = \frac{s+\alpha}{s} \cdot k_s = (s + \alpha)k$, to Chooser.

Fix α. To get optimal communication, we define ℓ_i as $\ell_i := \left(\frac{(s+\alpha)!}{s!}\right)^{1/\alpha} \cdot \frac{1}{s+i} \cdot N^{1/\alpha}$; this results in the minimal $\sum_i i\ell_i$ under the constraint that $\prod \ell_i = N$. (In practice, we must round ℓ_i-s to the nearest integers. For the simplicity of exposition, we will not explicitly mention such issues anymore.) Thus, we assume that $N = \prod_{i=1}^{\alpha} \ell_i$ and define $\mu(i_1, \dots, i_\alpha) := \mu[(i_1 - 1) \prod_{i=2}^{\alpha} \ell_i + (i_2 - 1) \prod_{i=3}^{\alpha} \ell_i + \cdots + (i_{\alpha-1} - 1)\ell_\alpha + (i_\alpha - 1) + 1]$. Protocol 1 depicts the new LFCPIR(α) protocol.

Theorem 1. *Let $\Pi = (G, E, D)$ be the* DJ01 *cryptosystem. Assume that the Decisional Composite Residuosity Problem (DCRP) is hard. Let $n \in G(1^k)$ be a valid Chooser's public key, that has been securely distributed to* Sen. *Assume that $\mathcal{D} \subseteq n^s$ for some fixed $s \geq 1$. For every $\alpha \geq 1$, there exists a CPIR protocol* LFCPIR(α) *with chooser-side and sender-side communication being respectively $((s+\alpha)!/s!)^{1/\alpha} \cdot \frac{\alpha}{s} \cdot N^{1/\alpha} k_s$ and $(1 + \alpha/s) \cdot k_s$ bits.*

Proof. *Communication*: Chooser-side communication (the total length of ciphertexts $\beta_{j,t}$) is $\sum_{i=1}^{\alpha} \ell_i \frac{s+i}{s} k_s = \sum_{i=1}^{\alpha} \sqrt[\alpha]{\frac{(s+\alpha)!}{s!}} \cdot \frac{1}{s} N^{1/\alpha} k_s = \sqrt[\alpha]{\frac{(s+\alpha)!}{s!}} \cdot \frac{\alpha}{s} \cdot N^{1/\alpha} k_s$. *Correctness*: clear, since $D_K^{s+\alpha-1}(w_\alpha) = w_{\alpha-1}(\sigma_\alpha)$, $D_K^{s+\alpha-2}(w_{\alpha-1}(\sigma_\alpha)) = w_{\alpha-2}(\sigma_{\alpha-1}, \sigma_\alpha)$, $\dots$, $D_K^{s+i-1}(w_i(\sigma_{i+1}, \dots, \sigma_\alpha)) = w_{i-1}(\sigma_i, \dots, \sigma_\alpha)$, $\dots$, and $D_K^s(w_1(\sigma_2, \dots, \sigma_\alpha)) = \mu(\sigma_1, \dots, \sigma_\alpha)$.

PRIVATE INPUT: Chooser has $\sigma = (\sigma_1, \dots, \sigma_\alpha)$, Sender has μ.
PRIVATE OUTPUT: Chooser has $\mu(\sigma_1, \dots, \sigma_\alpha)$.

Round 1, Chooser: For all $j \in [\alpha]$ do, for all $t \in [\ell_j]$ do:
Generate $r_{j,t} \leftarrow_r R(x)$.
Send $\beta_{j,t} \leftarrow E_K^{s+j-1}(\{t \stackrel{?}{=} \sigma_j\}; r_{j,t}) \in \mathbb{Z}_{n^{s+j}}$ to Sender.
Round 2, Sender:

For all $i_2 \in [\ell_2], \dots, i_\alpha \in [\ell_\alpha]$ do:
Set $w_1(i_2, \dots, i_\alpha) \leftarrow \prod_{t \in [\ell_1]} \beta_{1,t}^{\mu(t, i_2, \dots, i_\alpha)} \mod n^{s+1}$.
For all $i_3 \in [\ell_3], \dots, i_\alpha \in [\ell_\alpha]$ do:
Set $w_2(i_3, \dots, i_\alpha) \leftarrow \prod_{t \in [\ell_2]} \beta_{2,t}^{w_1(t, i_3, \dots, i_\alpha)} \mod n^{s+2}$.
$\vdots$
For all $i_{\alpha-1} \in [\ell_{\alpha-1}], i_\alpha \in [\ell_\alpha]$ do:
Set $w_{\alpha-2}(i_{\alpha-1}, i_\alpha) \leftarrow \prod_{t \in [\ell_{\alpha-2}]} \beta_{\alpha-2,t}^{w_{\alpha-3}(t, i_{\alpha-1}, i_\alpha)} \mod n^{s+\alpha-2}$.
For all $i_\alpha \in [\ell_\alpha]$ do:
Set $w_{\alpha-1}(i_\alpha) \leftarrow \prod_{t \in [\ell_{\alpha-1}]} \beta_{\alpha-1,t}^{w_{\alpha-2}(t, i_\alpha)} \mod n^{s+\alpha-1}$.
Set $w_\alpha \leftarrow \prod_{t \in [\ell_\alpha]} \beta_{\alpha,t}^{w_{\alpha-1}(t)} \mod n^{s+\alpha}$.
Send $w_\alpha \in \mathbb{Z}_{n^{s+\alpha}}$ to Chooser.
Reconstruction, Chooser: Output $D_K^s(D_K^{s+1}(\dots(D_K^{s+\alpha-1}(w_\alpha))\dots))$.

Protocol 1: Protocol $\mathsf{LFCPIR}(\alpha)$, for fixed public key $K = n$ and fixed s.

Privacy: Since the DCRP is hard, Π is semantically secure. Sender sees only random encryptions of 0-s and 1-s, and thus privacy of Chooser is guaranteed by the semantical security of Π. The formal proof of this result is slightly more involved. Assume that for some machine A that works in time t, $\varepsilon = \mathsf{Adv}^{\mathsf{cpir}}_{\mathsf{Cho},k}(A)$. Next, we construct a machine M that uses A as an oracle to break the semantical security of Π, with $\mathsf{Adv}^{\mathsf{pkcsem}}_{\Pi,k}(M^A) = \varepsilon$. M does the next: Set $m_0 \leftarrow 0$ and $m_1 \leftarrow 1$, ask for a challenge ciphertext $c = E_K(m_b; R(x))$ for $b \leftarrow \{0,1\}$. (M must now guess the value b.) Set $c' \leftarrow E_K(0; R(x))/c$. Invoke A to obtain the inputs $(\sigma_0, \sigma_1, \mu)$. Write σ_0 and σ_1 as tuples of coordinates on the $\ell_1 \times \cdots \times \ell_\alpha$ hypercube, $\sigma_j = (\sigma_{j1}, \dots, \sigma_{jk})$. Toss a random coin $\overline{b} \leftarrow \{0,1\}$.For $j \in [\alpha]$, $t \in [\ell_j]]$, define

$$\beta_{jt} \leftarrow \begin{cases} E_K(0; r_{j,t}) \text{ for } r_{j,t} \leftarrow R(x) \ , & t \notin \{\sigma_{0,j}, \sigma_{1,j}\} \ , \\ c \cdot E_K(0; r_{j,t}) \ , & t = \sigma_{\overline{b},j} \ , \\ c' \cdot E_K(0; r_{j,t}) \ , & t = \sigma_{1-\overline{b},j} \ . \end{cases}$$

Send $\{\beta_{jt}\}$ to A and obtain her guess $\overline{b}'$. Answer $b = 1$ iff $\overline{b} = \overline{b}'$.

Now, $\{\beta_{jt}\}$ has the same distribution as Cho's part of the protocol corresponding to the query $\overline{b}$ (if $b = 1$) or $1 - \overline{b}$ (if $b = 0$.) Thus, $\mathsf{Adv}^{\mathsf{cpir}}_{\mathsf{Cho},k}(M^A) = \mathsf{Adv}^{\mathsf{pkcsem}}_{\Pi,k}(A)$. Moreover, M works in time that is slightly longer than the working time of A and of Cho summed together. □

Asymptotically, the total communication of $\mathsf{LFCPIR}(\alpha)$ is $(\alpha+O(\log\alpha))\cdot\frac{\alpha}{s^2e}N^{1/\alpha}\cdot k_s$. In the asymptotically optimal case $\alpha=\frac{\log N}{2}$ this is equal to $(1+o(1))\frac{\log^2 N}{s^2 e}\cdot k_s$.

In the case of the hypercube (i.e., when $\ell_i=N^{1/\alpha}$), the chooser-side communication is $\sum_{i=1}^{\alpha}(1+i/s)\cdot N^{1/\alpha}k_s=(\alpha s+\frac{\alpha(\alpha+1)}{2s})N^{1/\alpha}k_s=\alpha(1+\frac{\alpha+1}{2s})N^{1/\alpha}k_s=(1+o(1))\frac{\alpha^2}{2s}\cdot N^{1/\alpha}k_s$. In the asymptotically optimal case $\alpha=\frac{\log N}{2}$, this will be equal to $(1+o(1))\frac{\log^2 N}{2s}\cdot k_s$. Thus, using the hyper-rectangle instead of the hypercube results in the asymptotic win of $\frac{se}{2}$ times.

Computation. In Protocol 1, Chooser must do ℓ_j encryptions E_K^s, and Sender—$\prod_{i=j+1}^{\alpha}\ell_i$ ℓ_j-term multi-exponentiations modulo n^{s+j} for every $j\in[2,\alpha]$. First, for simplicity, assume that $\alpha=\log N/2$ and $\ell_i=N^{1/\alpha}=4$ for every i. Second, assume that x-bit 4-term multi-exponentiation can be done in time x^a for some a. Then, Sender's workload is dominated by $\sum_{i=1}^{\alpha-1}4^{\alpha-i}\cdot(s+i)^a k^a$ bit-operations. Asymptotically in N, this is equal to $\Theta(s^a N)k^a$. Conservative estimation $a=3$ yields time $\Theta(s^3N)k^3$; fast multi-exponentiation algorithms result in Sender's time $\Theta(N)\cdot k_s^{2+o(1)}$. Therefore, the computation cost of $\mathsf{LFCPIR}(\log N/2)$ is comparable to that of the $\binom{N}{1}$-oblivious transfer protocols from [NP01,AIR01,Lip03] that have linear total communication. This is an important property of the $\mathsf{LFCPIR}(\alpha)$ protocol since otherwise, as argued e.g. in [AIR01], in some applications one would prefer a protocol with linear communication but with a smaller computational complexity.

4 Two Hypothetical Logarithmic $\binom{N}{1}$-CPIR-s

The constructed $\binom{N}{1}$-CPIR $\mathsf{LFCPIR}(\alpha)$ has total communication $\Theta(\log^2 N)\cdot k$. However, the communication complexity of the non-private information retrieval is $\log N+\log|\mathcal{D}|$ (Chooser sends a $\log N$-bit index σ and Sender responds with a $\log|\mathcal{D}|$-bit value μ_σ).

Next, we modify $\mathsf{LFCPIR}(\alpha)$ so as to have near-optimal communication $\Theta(\log N)\cdot k$ under some feasible additional conjectures on the used homomorphic cryptosystem. We do not know whether such cryptosystems exist but we would like to think that they do. The common motivation behind both modifications is based on the next observation: $\mathsf{LFCPIR}(\alpha)$ has log-squared—and not logarithmic—communication mostly because the length of the values $\beta_{j,t}$ depends (linearly) on j. Thus, one way to improve the communication-efficiency $\mathsf{LFCPIR}(\alpha)$ to logarithmic is to make $|\beta_{j,t}|$ independent of j. We propose two different modifications to the $\mathsf{LFCPIR}(\alpha)$ that overcome this obstacle, although under two different conjectures.

In the first case, we assume that $\beta_{j,t}=E_K^{s+j-1}(\{t\stackrel{?}{=}\sigma_j\};r_{j,t})$ can be securely communicated to Sender by using ηk_s bits, i.e., that $|\beta_{j,t}|$ does not depend on j at all. This involves an hypothesis of the existence of a certain succinct cryptosystem. In the second solution, we assume that $\beta_{j,t}=E_K^s(\{t\stackrel{?}{=}\sigma_j\};r_{j,t})$ is used as a multiplicand; this involves an alternative hypothesis of the existence of a ring-homomorphic cryptosystem. We conjecture that such cryptosystems exist—although we also show that the $\mathsf{DJ01}$ cryptosystem does not have the conjectured properties. We hope that spelling out these two conjectures gives an extra motivation for studying cryptosystems with such properties.

First case. We now make the next conjecture.

Conjecture 1. There exists a length-flexible semantically secure homomorphic cryptosystem $\Pi = (G, E, D)$, such that (a) $C_s(x) \subseteq P_{s+1}(x)$, (b) $|C_s(x)| \approx (s+1)k$ and (c) for every s, z there exist efficient functions $g_s : P_s(x) \times R(x) \to Z_s(x)$ and $f_{s,z} : Z_s(x) \to C_{s+z}(x)$, such that $f_{s,z}(g_s(m; r)) = E_K^{s+z}(m; r')$, for $m \in \{0, 1\}$ and for any $r \in R(x)$. Moreover, $|Z_s(x)| = |C_s(x)| = (1 + o(1))s \log n$ for some n. We call such a cryptosystem *succinct*.

(Here, (b) can be softened to the requirement that $|C_s(x)| \approx (s+\gamma)\cdot k$ for some constant γ.) Based on a succinct cryptosystem Π, we can construct a new $\binom{N}{1}$-CPIR protocol, that looks almost exactly the same as Protocol 1, except that (a) $\beta_{j,t}$ is defined as $\beta_{j,t} \leftarrow g_s(\{t \stackrel{?}{=} \sigma_j\}; r_{j,t})$, and (b) when using $\beta_{i,t}$ in the i-th sub-round of round 2, Sender first applies the function f_i on it. As an example, he sets $w_\alpha \leftarrow \prod_{t \in [\ell_\alpha]} f_\alpha(\beta_{\alpha,t})^{w_{\alpha-1}(t)} \mod n^{s+\alpha}$. Thus, Chooser communicates every $\beta_{j,t}$ to Sender by using ηk_s bits. The sender just sends back the value w_α.

Theorem 2. *Let Π be a (hypothetical) succinct cryptosystem. (a) For every $\alpha \geq 1$, there exists a CPIR protocol $\mathsf{HypoCPIR}^1(\alpha)$ with chooser-side and sender-side communication being respectively $\alpha N^{1/\alpha}(1 + 1/s)k_s$ and $(1 + \alpha/s)k_s$ bits. (a) Let $a = (W(\frac{1}{e(s+1)}) + 1) \cdot \log e$. There exists a CPIR protocol $\mathsf{HypoCPIR}^1(\frac{\log 2}{a} \log N)$ with chooser-side and sender-side communication being respectively $\frac{2^{a+1}}{a} \log N \cdot k_s$ and $(1 + \frac{1}{a} \cdot \log N) \cdot k_s$ bits.*

Proof. *Correctness*: By Conjecture 1, $f_j(\beta_{j,t}) \mod n^{s+j} = E_K^{s+j-1}(\{t \stackrel{?}{=} \sigma_j\}; r_{j,t})$. *Communication*: Chooser-side communication (the total length of ciphertexts $\beta_{j,t}$) is $\alpha \cdot N^{1/\alpha}(1 + 1/s) \cdot k_s$. *Privacy*: Sender sees only random encryptions of 0-s and 1-s, and thus Chooser's privacy is guaranteed by the semantical security of Π. □

The constant a has been chosen to optimise the total communication. For $s = 1$, the chooser-side communication is approximately $3.81076 \cdot \log N \cdot k$ and the sender-side communication is approximately $(1 + 0.59899 \log N)k$.

On the (im)possibility of succinct cryptosystems. Since Π is a semantically secure homomorphic cryptosystem, $E_K^s(m_1; r_1) \cdot E_K^s(m_2; r_2) = E_K^s(m_1 + m_2; r_1 \circ r_2)$. For all well-known semantically secure homomorphic cryptosystems, $\cdot$ corresponds to the modular multiplication in $C_s(x) = \mathbb{Z}_{n_s}^*$ for some n_s. Thus, in such a case $E_K^s(m; r) = g_s^m f(r, s) \mod n_s$ for some element $g_s \in \mathbb{Z}_{n_s}^*$ and for some function f. Moreover, when Π is length-flexible then it holds—at least for the currently known length-flexible cryptosystems Π [DJ01,DJ03]—that $\overline{n} := |C_{s+1}(x)|/|C_s(x)|$ is an integer, and that $E_K^{s+1}(m; r) = g_{s+1}^m f(r, s+1) \mod n_{s+1}$, where $n_{s+1} = \overline{n} n_s$ for some integer $\overline{n}$.

Assume now that $E_K^s(m; r) = a \mod n_s$ (e.g., $a = g_s^m f(r, s)$ without modular reduction), that is, $a = E_K^s(m; r) + b n_s$ for some b. We want to establish for which functions h, $h(a) \mod n_{s+1}$ can be computed from $a \mod n_s$—that is, for which functions h, $h(a) \mod n_{s+1}$ does not depend on b.

When we assume that h does not involve modular reductions, we can easily find that h must be one of the next basic operations or their composite. First, $h_1(a) = a\overline{n}$. Then $h_1(a) \equiv a\overline{n} \equiv E^s_K(m;r)\overline{n} \mod n_{s+1}$, thus $h_1(a) \mod n_{s+1}$ can be computed from $a \mod n_s$. Second, $h_2(a) = a^{\overline{n}}$, since $h_2(a) \equiv E^s_K(m;r)^{\overline{n}} \mod n_{s+1}$. Assuming $E^{s+1}_K(m;r) = g^m_s f(r, s+1) \mod n_{s+1}$, we get $a^{\overline{n}} \equiv g^{\overline{n}m}_s f(m, s+1)^{\overline{n}} \mod n_{s+1}$.

No other function on $a \mod n_s$, a is an arbitrary integer, except the two cases above (h_1 and h_2) and their compositions with each other and with functions in $\mathbb{Z}_{n_{s+1}}$, can return a meaningful function $h(a) \mod n_{s+1}$: e.g., for $\overline{n} \nmid j$, $a^j \equiv (E^s_K(m;r) + bn_s)^j \equiv (E^s_K(m;r)^j + \binom{j}{2} E^s_K(m;r)^{j-1} b + \dots) \mod n_{s+1}$ depends also on b—a value that is unknown to the potential Sender.

Thus, $\mathsf{HypoCPIR}^1(\alpha)$ exists iff some composition of h_1 and h_2 can transform a valid ciphertext in $C_s(x)$ to a valid ciphertext of the same element in $C_{s+1}(x)$. Let us look at the case of the $\mathsf{DJ01}$ cryptosystem.

Lemma 1. *Let $\Pi = (G, E, D)$ be the* $\mathsf{DJ01}$ *cryptosystem. Let $\alpha = E^s_K(m;r)$ for some $m \in \mathbb{Z}_{n^s}$ and $r \in \mathbb{Z}^*_n$. Then $E^{s+1}_K(mn;r) = \alpha^n \mod n^{s+2}$ and thus $E^{s+\alpha}_K(mn^\alpha;r) = \alpha^{n^\alpha} \mod n^{s+\alpha+1}$ for any $\alpha > 0$.*

In this case, $\overline{n} = n$. Then, $E^s_K(m;r) \cdot n \mod n^{s+2}$ is not a valid ciphertext and $E^s_K(m;r)^n \equiv E^{s+1}_K(mn;r) \mod n^{s+2}$. The latter value cannot be used to compute $E^{s+1}_K(m;r)$ by somebody who does not know the factorisation of n, since n is not an invertible element of $P_{s+1}(x)$. Thus, the $\mathsf{DJ01}$ cryptosystem is not succinct.

An hypothetical CPIR based on a ring-homomorphic cryptosystem. Assume that Π is a ring-homomorphic cryptosystem, that is, $D^s_K(E^s_K(m_1;r_1) \otimes E^s_K(m_2;r_2)) = m_1 \cdot m_2$ and $D^s_K(E^s_K(m_1;r_1) \oplus E^s_K(m_2;r_2)) = m_1 + m_2$ where $C_s(x)$ is a finite ring with operations $(\otimes, \oplus)$ and $P_s(x)$ is a finite ring with operations $(\cdot, +)$. In this case, we again modify Protocol 1 just by changing the definitions of $\beta_{j,t}$ and w_i. First, set $\beta_{j,t} \leftarrow E^s_K(\{t \stackrel{?}{=} \sigma_j\}; r_{j,t})$. Then, set, e.g., $w_1(i_2, \dots, i_\alpha) \leftarrow \bigoplus_t \left(E^s_K(\mu(t, i_2, \dots, i_\alpha); r'_{i_2,\dots,i_\alpha}) \otimes \beta_{1,t}\right)$, thus $D^s_K(w_1(i_2, \dots, i_\alpha)) = \mu(\sigma_1, i_2, \dots, i_\alpha)$. Computation of w_i for $i > 1$ is analogous.

The resulting hypothetical $\binom{N}{1}$-CPIR $\mathsf{HypoCPIR}^2(\alpha)$ is slightly more efficient than $\mathsf{HypoCPIR}^1(\alpha)$ since the sender only returns one ciphertext from $C_s(x)$ and not from $C_{s+\alpha}(x)$, as in the case of $\mathsf{HypoCPIR}^1(\alpha)$.

Theorem 3. *Let Π be a (hypothetical) ring-homomorphic cryptosystem. (a) For every $\alpha \geq 1$, there exists a CPIR protocol* $\mathsf{HypoCPIR}^2(\alpha)$ *with chooser-side and sender-side communication being respectively $\alpha N^{1/\alpha}(1 + 1/s)k_s$ and $(1 + 1/s)k_s$ bits. (a) Let $\alpha = \log N$. There exists a CPIR protocol* $\mathsf{HypoCPIR}^2(\log N)$ *with chooser-side and sender-side communication being respectively $2(1 + 1/s) \log N \cdot k_s$ and $(1 + 1/s) \cdot k_s$ bits.*

For $s = 1$, the chooser-side communication of $\mathsf{HypoCPIR}^2(\log N)$ is $4 \cdot \log N \cdot k$ and the sender-side communication is $2 \cdot k$. However, the Damgård-Jurik cryptosystem is not ring-homomorphic and moreover, no such cryptosystems are currently known. Our intuition is that it is "easier" to construct a succinct cryptosystem than to construct a ring-homomorphic cryptosystem. One of the reasons for this intuition is that it is well-known that a ring-homomorphic cryptosystems would enable to implement

Protocol	Chooser-side	Sender-side
Without privacy	$\log N$	$\log \lvert\mathcal{D}\rvert \leq s \log n = k_s$
$\mathsf{HomCPIR}\,(\alpha)$	$\eta\alpha N^{1/\alpha} \cdot k_s$	$\eta^\alpha \cdot k$
$\mathsf{HomCPIR}\,(\log_\eta N)$	$\eta^2 \log_\eta N \cdot k$	$N \cdot k$
$\mathsf{HomCPIR}\,(\sqrt{\log_\eta N})$	$\eta\sqrt{\log_\eta N}\eta^{\sqrt{\log_\eta N}} \cdot k$	$\eta^{\sqrt{\log_\eta N}} \cdot k$
$\mathsf{HomCPIR}\,(1)$	$\eta N \cdot k$	ηk
$\mathsf{LFCPIR}(\alpha)$	$\left(\frac{(s+\alpha)!}{s!}\right)^{1/\alpha} \cdot \frac{\alpha}{s} N^{1/\alpha} k_s$	$(\alpha+1) \cdot k_s$
$\mathsf{LFCPIR}(\frac{\log N}{2})$	$\frac{1+o(1)}{se} \log^2 N \cdot k_s$	$(\frac{\log N}{2}+1) \cdot k_s$
$\mathsf{LFCPIR}(1)$	$(1+1/s)Nk_s$	$2k_s$
$\mathsf{HypoCPIR}^1(\alpha)$	$\frac{s+1}{s} \cdot \alpha N^{1/\alpha} \cdot k_s$	$(1+\alpha/s)k_s$
$\mathsf{HypoCPIR}^1(\frac{\ln N}{a})$	$\frac{2^{a+1}}{a} \cdot \log N \cdot k_s$	$(1+\log N/a)k_s$
$\mathsf{HypoCPIR}^2(\alpha)$	$\frac{s+1}{s} \cdot \alpha N^{1/\alpha} \cdot k_s$	$(1+1/s)k_s$
$\mathsf{HypoCPIR}^2(\log N)$	$2\frac{s+1}{s} \cdot \log N \cdot k_s$	$(1+1/s)k_s$

Table 1. Total communication of $\mathsf{HomCPIR}\,(\alpha)$, $\mathsf{LFCPIR}(\alpha)$, $\mathsf{HypoCPIR}^1(\alpha)$ and $\mathsf{HypoCPIR}^2(\alpha)$ protocols when retrieving $< k_s$ bits with a single query. (For simplicity, we omit the ceiling signs, but it must be remembered that α and α are integers.) Here, a is as defined in Thm. 2

very efficiently many other tasks (including general multi-party computation), while the existence of a succinct cryptosystem has no so many direct corollaries.

5 Comparisons and Discussion

Communication. Fix $s = 1$. Table 1 summaries the next discussion. The difference between the communications of $\mathsf{HomCPIR}\,(\sqrt{\log N})$, $\mathsf{LFCPIR}(\frac{\log N}{2})$ and the hypothetical $\mathsf{HypoCPIR}^1(\ln N/a)$ (where a is as defined in Thm. 2) is depicted by the Figure 1. Note that $\mathsf{LFCPIR}(\frac{\log N}{2})$ is more communication-efficient than $\mathsf{HomCPIR}\,(\sqrt{\log N})$ for all N-s.

Chooser-side versus sender-side communication. One of the nice properties of the $\mathsf{HomCPIR}\,(\alpha)$ is that by varying α, one can achieve either chooser-side communication $\Theta(\log N)k$ (although then the sender-side communication is linear in N) or sender-side communication ηk (although then the chooser-side communication is linear). In the $\mathsf{LFCPIR}(\alpha)$, the minimal chooser-side communication is $\Theta(\log^2 N)k$ (with logarithmic sender-side communication) and the minimal sender-side communication is $2k_s$ (with linear chooser-side communication). On the other hand, the $\binom{N}{1}$-OT protocols of [NP01,AIR01,Lip03] have constant chooser-side communication and linear sender-side communication. It is an interesting open question to construct a (non-hypothetical) $\binom{N}{1}$-CPIR with logarithmic chooser-side communication and polylogarithmic sender-side communication.

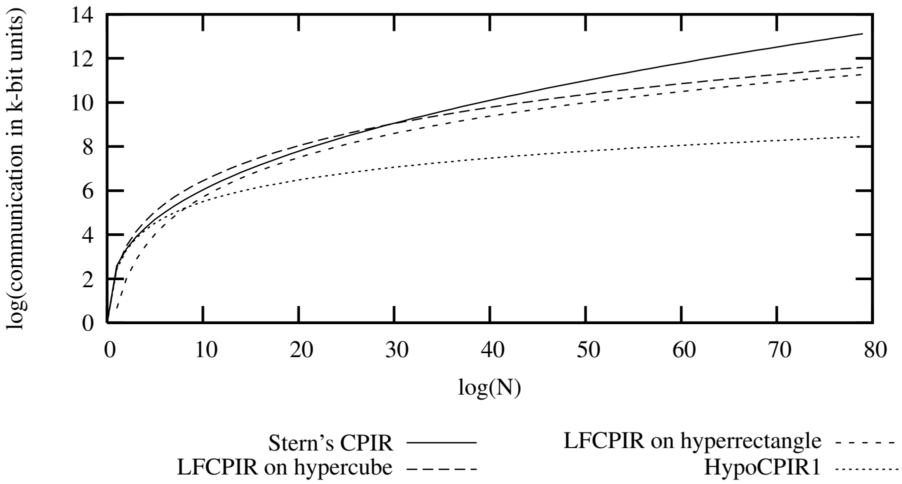

Fig. 1. Communication of $\mathsf{HomCPIR}\left(\sqrt{\log N}\right)$, $\mathsf{LFCPIR}(\frac{1}{2}\log N)$ (on hypercube) and $\mathsf{LFCPIR}(\frac{\ln N}{2})$ (on hyper-rectangle) in logarithm of k-bit units on logarithmic scale in N, assuming that $\eta = 2$ and $s = 1$. Comparison is also given with the hypothetical $\mathsf{HypoCPIR}^1(\frac{\ln N}{a})$

6 Oblivious Transfer with Log-Squared Communication

The $\mathsf{LFCPIR}(\alpha)$ protocol can be modified to an oblivious transfer protocol as follows. First, Sender should mask all values $w_j(i_{j+1}, \dots, i_\alpha)$ by multiplying them with $E_K^{s+j-1}(0; R(x))$. Second, Chooser must accompany the first round of the $\mathsf{LFCPIR}(\alpha)$ protocol with a standard honest-verifier zero-knowledge proof of knowledge (HVSZK POK) that he chose the inputs correctly. This means (in the case of hypercube) $\alpha N^{1/\alpha}$ HVSZK POK-s that an encrypted value is either 0 or 1, and $N^{1/\alpha}$ HVSZK POK-s that the product of α ciphertexts decrypts to 1. Again, some of these HVSZK POK-s are done with a larger s, so that the total overhead of the HVSZK POK-s is $\Theta(\log^2 N)k$. This does not increase the asymptotic communication complexity of the protocol. Therefore, we can construct a four-round $\binom{N}{1}$-OT protocol with total communication $\Theta(\log^2 N)k$ that can be made two-round in the random-oracle model by using the Fiat-Shamir heuristic. Somewhat more detailed exposition of the corresponding HVSZK POK-s is given in the Appendix.

Acknowledgements

We would like to thank Yan-Cheng Chang, Ivan Damgård and Rafail Ostrovsky for useful comments. This work was partially supported by the Finnish Defence Forces Institute for Technological Research and by the Finnish Academy of Sciences.

References

[AIR01] William Aiello, Yuval Ishai, and Omer Reingold. Priced Oblivious Transfer: How to Sell Digital Goods. In Birgit Pfitzmann, editor, *Advances in Cryptology — EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 119–135, Innsbruck, Austria, 6–10 May 2001. Springer-Verlag.

[Cha04] Yan-Cheng Chang. Single Database Private Information Retrieval with Logarithmic Communication. In Josef Pieprzyk and Huaxiong Wang, editors, *The 9th Australasian Conference on Information Security and Privacy (ACISP 2004)*, volume ? of *Lecture Notes in Computer Science*, pages ?–?, Sydney, Australia, 13–15 July 2004. Springer-Verlag. Accepted.

[CMS99] Christian Cachin, Silvio Micali, and Markus Stadler. Computational Private Information Retrieval with Polylogarithmic Communication. In Jacques Stern, editor, *Advances in Cryptology — EUROCRYPT '99*, volume 1592 of *Lecture Notes in Computer Science*, pages 402–414, Prague, Czech Republic, 2–6 May 1999. Springer-Verlag.

[DJ01] Ivan Damgård and Mads Jurik. A Generalisation, a Simplification and Some Applications of Paillier's Probabilistic Public-Key System. In Kwangjo Kim, editor, *Public Key Cryptography 2001*, volume 1992 of *Lecture Notes in Computer Science*, pages 119–136, Cheju Island, Korea, 13–15 February 2001. Springer-Verlag.

[DJ03] Ivan Damgård and Mads Jurik. A Length-Flexible Threshold Cryptosystem with Applications. In Rei Safavi-Naini, editor, *The 8th Australasian Conference on Information Security and Privacy*, volume 2727 of *Lecture Notes in Computer Science*, pages 350–364, Wollongong, Australia, July 9-11 2003. Springer-Verlag.

[KO97] Eyal Kushilevitz and Rafail Ostrovsky. Replication is Not Needed: Single Database, Computationally-Private Information Retrieval. In *38th Annual Symposium on Foundations of Computer Science*, pages 364–373, Miami Beach, Florida, 20–22 October 1997. IEEE.

[Lip03] Helger Lipmaa. Verifiable Homomorphic Oblivious Transfer and Private Equality Test. In Chi Sung Laih, editor, *Advances on Cryptology — ASIACRYPT 2003*, volume 2894 of *Lecture Notes in Computer Science*, pages 416–433, Taipei, Taiwan, 30 November–4 December 2003. Springer-Verlag.

[NP99] Moni Naor and Benny Pinkas. Oblivious Transfer and Polynomial Evaluation. In *Proceedings of the Thirty-First Annual ACM Symposium on the Theory of Computing*, pages 245–254, Atlanta, Georgia, USA, 1–4 May 1999. ACM Press.

[NP01] Moni Naor and Benny Pinkas. Efficient Oblivious Transfer Protocols. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 448–457, Washington, DC, USA, January 7–9 2001. ACM Press.

[NS98] David Naccache and Jacques Stern. A New Public Key Cryptosystem Based on Higher Residues. In *5th ACM Conference on Computer and Communications Security*, pages 59–66, San Francisco, CA, USA, 3–5 November 1998. ACM Press.

[OP98] Kazuo Ohta and Dingyi Pei, editors. *Advances on Cryptology — ASIACRYPT '98*, volume 1514 of *Lecture Notes in Computer Science*, Beijing, China, 18–22 October 1998. Springer-Verlag.

[OU98] Tatsuaki Okamoto and Shigenori Uchiyama. A New Public-Key Cryptosystem as Secure as Factoring. In Kaisa Nyberg, editor, *Advances in Cryptology — EUROCRYPT '98*, volume 1403 of *Lecture Notes in Computer Science*, pages 308–318, Helsinki, Finland, May 31 – June 4 1998. Springer-Verlag.

[Pai99] Pascal Paillier. Public-Key Cryptosystems Based on Composite Degree Residuosity Classes. In Jacques Stern, editor, *Advances in Cryptology — EUROCRYPT '99*, volume 1592 of *Lecture Notes in Computer Science*, pages 223–238, Prague, Czech Republic, 2–6 May 1999. Springer-Verlag.

[Ste98] Julien P. Stern. A New and Efficient All or Nothing Disclosure of Secrets Protocol. In Ohta and Pei [OP98], pages 357–371.

A Communication of $\binom{N}{1}$-OT Protocol

Next, we give a detailed description of the HVSZK POK that $D_K^s(c) \in \{0,1\}$. For this, one first computes $c_1 = c/E_K(0;0)$ and $c_2/E_K(1;0)$ and then proves that either $D_K(c_1) = 0$ or $D_K(c_2) = 0$, as follows. Let Π be a semantically secure additively homomorphic cryptosystem and let $\tau \approx 80$. Assume that $c_1 = E_K(0;r)$. Let c_2 be another ciphertext. The next protocol for proving that either $D_K(c_1) = 0$ or $D_K(c_2) = 0$ was proposed in [DJ01]:

1. Chooser chooses a random $r_1 \leftarrow R(x)$ and sets $a_1 \leftarrow E_K^s(0;r_1)$. He generates a view (a_2, e_2, z_2) by setting $e_2 \leftarrow \{0,1\}^\tau$, $z_2 \leftarrow \mathbb{Z}_n$ and $a_2 \leftarrow E_K^s(0;z_2) \cdot c_2^{-e_2} \mod n^{s+1}$. He sends (a_1, a_2) to Sender.
2. Sender chooses a random τ-bit number d, and sends d to Chooser.
3. Chooser computes $e_1 \leftarrow d - e_2 \mod 2^\tau$ and $z_1 \leftarrow r_1 r^{e_1} \mod n$. He sends (e_1, z_1, e_2, z_2) to Sender.
4. Sender checks that $d = e_1 + e_2 \mod 2^\tau$, $E_K^s(0;z_1) = a_1 c_1^{e_1} \mod n^{s+1}$, $E_K^s(0;z_2) = a_2 c_2^{e_2} \mod n^{s+1}$, and that $u_1, u_2, a_1, a_2, z_1, z_2$ are relatively prime to n. He accepts iff all verifications succeed.

Total communication of this HVSZK POK is $2|C_s(x)| + 3t + 2\lceil \log(n+1) \rceil = 2(1 + 1/s)k_s + 3t + \frac{2}{s}k_s = 3t + 2(1+2/s) \cdot k_s$.

In total, the HVSZK POK-s of this type add $\sum_{i=1}^{\alpha} \ell_i(3\tau + 2(1+(i+1)/s) \cdot k_s)$ bits of communication to the $\sum_{i=1}^{\alpha} \ell_i((1+i/s) \cdot k_s)$ bits of communication of the CPIR. Let us assume for simplicity that $\ell_i = N^{1/\alpha}$ for all i. Then the communication overhead of the $\binom{N}{1}$-OT protocol is

$$N^{1/\alpha} \sum_{i=1}^{\alpha} (3\tau + 2(1 + (i+2)/s) \cdot k_s) = N^{1/\alpha}(3\tau\alpha + \frac{1}{s} \cdot (2\alpha(s+2) + \alpha^2 + \alpha)k_s)$$
$$= (1 + o(1)) N^{1/\alpha} \cdot \frac{\alpha^2}{s} \cdot k_s \ .$$

In the near-optimal case $\alpha = \frac{\log N}{2}$, this is $(1/s + o(1)) \log^2 N \cdot k_s$ bits. Recall that in the case of the hypercube-$\mathsf{LFCPIR}(\frac{\log N}{2})$, the total communication is $(1 + o(1)) \frac{\log^2 N}{2s} k_s$, thus the communication only asymptotically triples after adding the HVSZK POK-s. (The $N^{1/\alpha}$ proofs that some product of ciphertexts decrypts to 1 does not add a significant amount of communication.)

Theorem 4. *Let $\Pi = (G, E, D)$ be the* $\mathsf{DJ01}$ *cryptosystem. Assume that the DCRP is hard. For every $\alpha \geq 1$, there exists an $\binom{N}{1}$-oblivious transfer protocol* $\mathsf{DJOT}(\alpha)$ *with total communication $\frac{3}{2s}(1 + o(1))\alpha^2 N^{1/\alpha} k_s$ bits. In particular, there exists an $\binom{N}{1}$-OT protocol* $\mathsf{DJOT}(\frac{\log N}{2})$ *with total communication $\frac{3}{2s}(1 + o(1)) \log^2 N \cdot k_s$ bits.*

Another alternative to the proof that the decryption of $\beta_{j,t}$ is either 0 or 1 is to prove that its decryption is equal to its own square. In this case, Chooser proves that $D_K^{s+j-1}(\beta_{jt}) = D_K^{s+j-1}(\beta_{jt})^2 = \{t \stackrel{?}{=} \sigma_j\}$ by using the next honest-verifier zero-knowledge proof for multiplicative relationship.

- Chooser generates random $m \leftarrow M(x)$, $r_1, r_2 \leftarrow R(x)$ and sends $a_1 \leftarrow E_K^{s+j-1}(m; r_1)$ and $a_2 \leftarrow E_K^{s+j-1}(m \cdot \{t \stackrel{?}{=} \sigma_j\}; r_2)$ to Sender.
- Sender generates a random τ-bit number e and sends it to Chooser.
- Chooser sends $z_m = m_1 + e \cdot \{t \stackrel{?}{=} \sigma_j\}$, $z_r \leftarrow r_{jt}^e r_1$ and $z_r' \leftarrow r_{j,t}^{z_m - e} r_2^{-1}$ to Sender.
- Sender verifies that $\beta_{jt}^e a_1 = E_K^{s+j-1}(z_m; z_r)$, $\beta_{jt}^{z_m - e} a_2^{-1} = E_K^{s+j-1}(0; z_r')$, and that all elements sent by Chooser are relatively prime to n.

However, its communication, $2|C_{s+j-1}(x)| + \tau + |P_{s+j-1}(x)| + 2|R(x)| = 2(s+j)k + \tau + (s+j-1)k + k = 3(1+j/s)k_s + \tau$ is actually worse.

Alternative constructions. Alternative $\binom{N}{1}$-OT protocol constructions are possible (e.g., based on the methodology of [NP99]) and will be analysed in the final version of this paper.