LIPIcs Leibniz International Proceedings in Informatics

Non-Local Box Complexity and Secure Function Evaluation

M. Kaplan*, I. Kerenidis*, S. Laplante*, and J. Roland†

ABSTRACT. A non-local box is an abstract device into which Alice and Bob input bits x and y respectively and receive outputs a and b respectively, where a, b are uniformly distributed and $a \oplus b = x \wedge y$. Such boxes have been central to the study of quantum or generalized non-locality as well as the simulation of non-signaling distributions. In this paper, we start by studying how many non-local boxes Alice and Bob need in order to compute a Boolean function f. We provide tight upper and lower bounds in terms of the communication complexity of the function both in the deterministic and randomized case. We show that non-local box complexity has interesting applications to classical cryptography, in particular to secure function evaluation, and study the question posed by Beimel and Malkin [4] of how many Oblivious Transfer calls Alice and Bob need in order to securely compute a function f. We show that this question is related to the non-local box complexity of the function and conclude by greatly improving their bounds. Finally, another consequence of our results is that traceless two-outcome measurements on maximally entangled states can be simulated with 3 non-local boxes, while no finite bound was previously known.

1 Introduction

Communication complexity. Communication complexity is a central model of computation, which was first defined by Yao in 1979 [35] and has since found numerous applications. In this model Alice and Bob receive inputs x and y respectively and are allowed to communicate in order to compute a function $f(x, y)$. The goal is to find the minimum amount of communication needed for this task. In different variants of the model, we allow Alice and Bob to err with some probability, and to share common resources in an attempt to enable them to solve their task more efficiently.

One such resource is shared randomness. When Alice and Bob are not allowed any errors, shared randomness does not reduce the communication complexity. On the other hand, when they are allowed to err, a common random string can reduce the amount of communication needed. However, Newman's result tells us that shared randomness can be replaced by private randomness at an additional cost logarithmic in the input size.

Another very powerful shared resource is entanglement. Using teleportation, Alice and Bob can transmit quantum messages by using their entanglement and only classical communication. This model has been proven to be very powerful, in some cases exponentially more efficient than the classical one. Another way to understand the power of entanglement is by looking at the CHSH game [13], where Alice and Bob receive uniformly random bits x and y respectively and their goal is to output bits a and b resp. such that $a \oplus b = x \wedge y$ without communicating. It is not hard to conclude that even if Alice and Bob share randomness, their optimal strategy will be successful with probability 0.75 over the inputs. However, if

*LRI - Université Paris-Sud
†NEC Laboratories America

Foundations of Software Technology and Theoretical Computer Science (Kanpur) 2009.
Editors: Ravi Kannan and K. Narayan Kumar; pp 239–250
Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany.
Digital Object Identifier: 10.4230/LIPIcs.FSTTCS.2009.2322

they share entanglement, then there is a strategy that succeeds with probability 0.85. This game proves that quantum entanglement can enable two parties to create correlations that are impossible to create with classical means.

Non-local boxes. As we said, entanglement enables Alice and Bob to succeed in the CHSH game with probability 0.85. But what if they shared some resource that would enable them to win the game with probability 1? Starting from such considerations, Popescu and Rohrlich [30] defined the notion of a non-local box. A non-local box is an abstract device shared by Alice and Bob. By one use of a non-local box, we mean that Alice inputs x, Bob inputs y, Alice gets as output a and Bob gets b where a, b are uniformly distributed and more importantly $a \oplus b = x \wedge y$. The name non-local box is due to the property that one use of a non-local box creates correlations between two bits that are maximally non-local (allowing to win the CHSH game with probability one), but still does not allow to communicate, since taken separately, each bit is just an unbiased random coin. As such, a non-local box may be considered as a unit of non-locality. We note here an important property of a non-local box, namely that, similar to entanglement, one player can enter an input and receive an output even before the second player has entered an input.

The importance of the notion of a non-local box has become increasingly evident in the last years. Non-local boxes were first introduced to study (quantum or generalized) non-locality. In particular, it was shown than one of the most studied versions of the EPR experiment, where Alice and Bob perform projective measurements on a maximally entangled qubit pair, may be simulated using only one use of a non-local box [10]. More generally, it was shown that any non-signaling distribution over Boolean outputs may be exactly simulated with some finite number of non-local boxes (for finite input size) [1, 19]. This was later generalized to any non-signaling distribution, except that the simulation may not always be performed exactly for non-Boolean outputs [16]. These results rely on the fact that the set of non-signaling distributions is a polytope, so it suffices to simulate the extremal vertices to be able simulate the whole set. In the context of non-locality, another application of non-local boxes is the study of pseudo-telepathy games [7].

It is easy to see that one use of a non-local box can be simulated with one bit of communication and shared randomness: Alice outputs a uniform bit r and sends x to Bob, who outputs $r \oplus x \cdot y$. However, the converse cannot possibly hold, since a non-local box cannot be used for communication.

The first question is what happens if we use non-local boxes as shared resource in the communication complexity model. Van Dam showed that for any Boolean function $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, Alice and Bob can use 2^n non-local boxes and no communication at all and at the end output bits a and b such that $a \oplus b = f(x, y)$ [33]. In other words, if non-local boxes were physically implementable, then all functions would have trivial communication complexity. His results were strengthened by Brassard *et al.* who showed that even if a non-ideal non-local box existed, one that solves the CHSH game with probability 0.91, then still all functions would have trivial communication complexity [6]. Note that in these results, the number of non-local boxes needed may be exponential in the input size and do not take into account any properties of the function and more precisely its communication complexity without non-local boxes. It also follows from the work of [1, 6] that for

any Boolean function f, if f has a circuit with fan-in 2 of size s, then there is a deterministic non-local box protocol of complexity $O(s)$, where the bits of the input of f are split arbitrarily among the players. This implies that exhibiting an explicit function for which the deterministic non-local box complexity is superlinear would be a real breakthrough, since it would translate into a superlinear circuit lower bound for this function.

Secure function evaluation. Non-local boxes have also been studied in relation to cryptographic primitives such as Oblivious Transfer or Bit Commitment. Wolf and Wullschleger [34] showed that Oblivious Transfer is equivalent to a *timed* version of a non-local box (up to a factor of 2). To maintain the non-signaling property of the non-local box, one can define timed non-local box as having a predefined time limit, and if any of the players have not entered an input by this time, then some fixed input, say 0, is used instead. Subsequently, Buhrman *et al.* [8] showed how to construct Bit Commitment and Oblivious Transfer by using non-local boxes that do not need to be timed but have to be trusted.

In this paper, we are interested in secure function evaluation, which is one of the most fundamental cryptographic tasks. In this model, Alice and Bob want to evaluate some function of their inputs in a way that does not leak any more information than what follows from the output of the function. It is known that even though not all functions can be evaluated securely in the information-theoretic setting ([5, 11, 12, 26]), all functions can be computed securely in the information theoretic setting, if we have access to a black box that performs Oblivious Transfer or some other complete function, e.g. the AND function ([17, 20]).

There has been a lot of work trying to identify, in various settings, which functions can be easily evaluated in a secure way, *i.e.*, without any invocation of the black box, and which are hard to evaluate securely, *i.e.*, require at least one invocation of the black box ([12, 26, 3, 21, 23, 3, 22]). Moreover, Beaver [2] showed that there exists a hierarchy of different degrees of hardness for the information-theoretic reduction setting, in other words that for all k, there are functions that can be securely evaluated with k invocations of the AND box but cannot be computed with $k-1$ uses of the black box.

Beimel and Malkin [4] proposed a quantitative approach to secure function evaluation by studying how many calls to an Oblivious Transfer or other complete black box one needs in order to securely compute a given function f in the honest-but-curious model. For a Boolean function $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ and deterministic protocols, they provide a combinatorial characterization of the minimal number of AND calls required, which however does not lead to an efficient algorithm to determine how many ANDs are actually required. They also show that at most $2^{|\mathcal{X}|}$ ANDs are needed for any function. In the randomized case, they provide lower bounds depending on the truth-table of the function which can be at most of the order of n. They also state that "it would be very interesting to try and explore tighter connections with the communication complexity of the functions".

Finally, Naor and Nissim [29] have given some connections between the communication complexity of a function f and the communication complexity for securely computing f. These results, translated into the Beimel-Malkin and our model, only show that the number of ANDs is at most exponential in the communication complexity.

Our results. In this paper, we provide more evidence on the importance of non-local boxes by showing how they relate to different models of communication complexity as well as how they can be used as a tool to quantitatively study secure function evaluation. Our results show that non-local boxes, introduced for the study of quantum or more general non-locality, can provide a novel way of looking at questions about classical communication complexity, secure function evaluation and complexity theory.

2 Preliminaries

2.1 Communication Complexity

Let $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ be a bipartite Boolean function. Alice gets an input $x \in \mathcal{X}$ and Bob gets an input $y \in \mathcal{Y}$. We say that Alice and Bob compute $f(x,y)$ in parity if after executing a protocol, Alice outputs a bit a and Bob outputs a bit b such that $a \oplus b = f(x,y)$, where we use $\oplus$ to denote both the logical XOR and addition mod 2. This model differs from the standard model, where one of the players outputs the value of the function, by at most 1 bit.

We use the following notions of communication complexity. In probabilistic models, we assume that the players have a common source of randomness. $D(f)$ and $R_\varepsilon(f)$ denote deterministic and ε-bounded error communication complexity of $f(x,y)$ in parity. $D^{\rightarrow}(f)$ and $R_\varepsilon^{\rightarrow}(f)$ are the one-way deterministic and bounded-error communication complexity of $f(x,y)$ in parity. Finally, $D^{\parallel}(f)$ and $R_\varepsilon^{\parallel}(f)$ are the deterministic and bounded-error communication complexities in the model of simultaneous messages, where Alice and Bob each send a message to the referee and the referee outputs the value of the function $f(x,y)$.

For the model of simultaneous messages, we also consider some natural restrictions on how the referee computes the output from the messages he receives from the players. We assume the messages sent are of the same length. Suppose the referee receives bits $\mathbf{a} = (a_1, \ldots, a_t)$ from Alice, and $\mathbf{b} = (b_1, \ldots, b_t)$ from Bob. If the referee always computes a predefined function $g(\mathbf{a}, \mathbf{b})$, then we write $D^{\parallel,g}(f)$ or $R_\varepsilon^{\parallel,g}(f)$ to be the length of the message sent by the players (not the sum of these lengths, as is done in the standard model). In this paper, we will consider two functions, the inner product modulo 2, $IP_2(\mathbf{a}, \mathbf{b}) = \bigoplus_i (a_i \cdot b_i)$ (where $\cdot$ denotes the multiplication over $\mathbb{GF}_2$, which corresponds to the logical AND) and the majority function, $MAJ(\mathbf{a}, \mathbf{b}) = MAJ(a_1 \oplus b_1, \ldots, a_t \oplus b_t)$.

2.2 Non-local box Complexity

Definition 1 (Non-local box) *A non-local box is a device shared by two parties, which on one side takes Boolean input x and immediately produces Boolean output a, and on the other side takes Boolean input y and immediately produces Boolean output b, according to the following distribution:*

$$\mathbf{p}_{NL}(a,b|x,y) = \begin{cases} \frac{1}{2} & \text{if } a \oplus b = x \cdot y \\ 0 & \text{otherwise.} \end{cases}$$

We study a model akin to communication complexity, where Alice and Bob use non-local boxes instead of communication. In a non-local box protocol, Alice and Bob wish to compute some function $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ in the following way. Alice gets an input $x \in \mathcal{X}$, Bob gets an input $y \in \mathcal{Y}$, and at the end of the protocol, Alice outputs $a \in \{0,1\}$, Bob outputs $b \in \{0,1\}$, such that $a \oplus b = f(x,y)$. For a protocol P, we will write $P(x,y) = (a,b)$.

In the course of the protocol, Alice and Bob are allowed shared randomness and may use non-local boxes, but they may not communicate. Bob is not allowed to see Alice's inputs to the non-local boxes, nor does he see the outcome on Alice's side, and likewise for Alice.

Definition 2 *For any function $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$, $NL(f)$ is the smallest t such that there is a protocol that computes f exactly, using t non-local boxes.*

We will label the non-local boxes with labels from 1 to t. (Recall that in general, Alice and Bob are not required to use the t non-local boxes in the same order.) We relax the exactness condition and allow the protocol's outcome to be incorrect with constant probability ε.

Definition 3 *For any function $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$, $NL_\varepsilon(f)$ is the smallest t such that there is a protocol P using t non-local boxes, with $\Pr[P(x,y) = (a,b)$ with $a \oplus b = f(x,y)] \geq 1 - \varepsilon$.*

We will also study two variants of the general model, where the non-local boxes are used in a restricted manner. First, we assume that the non-local boxes are used in parallel, that is, the input to any non-local box does not depend on the outcome of any other. The complexity in this model is denoted $NL^{||}$ in the exact case, and $NL_\varepsilon^{||}$ in the ε error case.

Second, we define the model where both players use the non-local boxes in the same order, that is, the non-local boxes are labeled from 1 to t and Alice's input to the non-local box with label i does not depend on the outputs from the non-local boxes labeled from $i+1$ to t (similarly for Bob). The complexity in this model is denoted NL^{ord} in the exact case, and $NL_\varepsilon^{\text{ord}}$ in the ε error case. It is clear that this model is more powerful than the parallel model but less powerful than the general non-local box complexity. In fact, we will only use this last variant when we talk about secure function evaluation. Note also that in all these models, the non-local boxes are still non-signaling and Alice and Bob receive the outputs of the non-local boxes immediately after they enter their inputs.

Finally, we consider a restriction where the players always output the same predefined function g of the outputs of the non-local boxes. Let $(a_1, b_1), \ldots, (a_t, b_t)$ be the outcomes of the t non-local boxes in some particular run of a protocol. Of particular interest are protocols where Alice outputs $a = a_1 \oplus \cdots \oplus a_t$ and Bob outputs $b = b_1 \oplus \cdots \oplus b_t$. The function g is used in a superscript to denote the complexity of a function f in this model, NL^g in the determinstic case, and NL_ε^g in the ε error case, and in particular, $NL^{||,\oplus}$ and $NL_\varepsilon^{||,\oplus}$ when the non-local boxes are in parallel and $g = \oplus$.

2.3 Secure Function Evaluation

We will consider the following cryptographic primitives.

Definition 4 (Oblivious transfer) *A 2-1 Oblivious Transfer (OT) is a device which on input bits p_0, p_1 for Alice and q for Bob, outputs bit b to Bob, such that $b = p_q$.*

Definition 5 (Secure AND) *A secure AND is a device which on input bits p for Alice and q for Bob, outputs bit a to Alice, such that $a = p \cdot q$.*

While at first view, these definitions seem similar to the definition of the non-local box, note that the timing properties are different: for the cryptographic primitives, the outputs are produced only after all the inputs have been entered into the device. It is precisely this subtlety that has led to confusion when trying to use non-local boxes to implement cryptographic primitives, in particular for bit commitment, when timing is particularly important, since a cheating Alice could wait until the reveal phase before committing her bit into the

non-local box, without Bob ever realizing it [8]. However, we will see that this is not an issue for our results on secure computation.

Let $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ be a bipartite Boolean function. We study the number of cryptographic primitives necessary to compute f. In all the models we consider, we require perfect privacy. In the honest-but-curious model, perfect privacy means that when a player follows the protocol, he should not learn *more than required* about the other player's input. In the malicious model, this condition must still hold even if the player does not follow the protocol. *Not more than required* means, for models where the function must be computed in parity, that the players should learn nothing about the other's input, while for models where one of the player should output the function, it means that this player should learn nothing more than what he can infer from his input and the value of the function, while the other player should learn nothing.

Let us note that AND may not be used as a primitive in the malicious model, so we will consider the OT primitive instead. Moreover, in this model, it is known that randomness is necessary to achieve perfect privacy [15], so in this setting we do not consider the deterministic model. Our bounds in the randomized malicious model also hold for the weaker honest-but-curious model. We define $AND(f)$ to be the number of secure AND gates required to securely compute $f(x,y)$ (*not* in parity) in the deterministic, honest-but-curious model. We note that we can allow free two-way communication without in fact changing the complexity [4]. Similarly, $OT_\varepsilon(f)$ is the number of 2-1 Oblivious Transfer calls required to compute $f(x,y)$ in parity with perfect privacy and ε error over the players' private coins, assisted with (free) two-way communication, in the malicious model.

2.4 Complexity Measures

We will compare non-local box complexity to traditional models of communication complexity and prove upper and lower bounds for this new model. Some of these bounds are in terms of the factorization norms of the communication matrix [28] and related measures.

DEFINITION 1. *Let M be a real matrix. The γ_2 norm of M is $\gamma_2(M) = \min_{X^TY=M} col(X)col(Y)$, where $col(N)$ is the largest Euclidian norm of a column of N.*

It is known that $2\log(\gamma_2(M))$ gives a lower bound on deterministic communication complexity of M, where M is a sign matrix of the Boolean function to be computed [28]. In order to lower bound the randomized and quantum communication complexity, we have to consider a "smoothed" version of this measure.

DEFINITION 2. *Let M be a sign matrix and $\alpha \geq 1$. $\gamma_2^\alpha(M) = \min\{\gamma_2(N) : \forall i,j\, 1 \leq M_{i,j}N_{i,j} \leq \alpha\}$. In particular, $\gamma_2^\infty(M)$ is the minimum γ_2 norm over all matrices N such that $1 \leq M_{i,j}N_{i,j}$.*

The measures γ_2^α and γ_2^∞ give upper and lower bounds for bounded-error communication complexity [28]: $2\log(\gamma_2^\alpha(f)/\alpha) \leq R_\varepsilon(f)$ and $R_\varepsilon^{||,MAJ}(f) \leq O((\gamma_2^\infty(f))^2)$ (implicit in [28]), where $\alpha = \frac{1}{1-2\varepsilon}$. The discrepancy of a sign matrix M over inputs $X \times Y$ with respect to distribution μ over the inputs is $Disc_\mu(M) = \max_R \sum_{(x,y)\in R} \mu(x,y)M(x,y)$, where R is taken from all possible rectangles. It is known that $\gamma_2^\infty(f) = \Theta(\frac{1}{Disc(f)})$, and for any α, $\gamma_2^\infty(f) \leq \gamma_2^\alpha(f)$ [28]. Finally, for a Boolean function, the L_1 norm is defined as the sum

of the absolute values of its Fourier coefficients. We can think of the $2n$ bits of input of the function as equally split between Alice and Bob. Grolmusz uses this notion to upper bound the randomized communication complexity by proving that $R_\epsilon(f) \leq O(L_1^2(f))$ [18].

3 Deterministic non-local box complexity

We start by studying a restricted model of non-local box complexity, where the non-local boxes are used in parallel and at the end of the protocol, Alice and Bob output the parity of the outputs of their non-local boxes respectively. We will show that the complexity of f in this model is equal to the rank of the communication matrix of f over $\mathbb{GF}_2$. Note that this rank is equal to the minimum m, such that $f(x,y)$ can be written as $f(x,y) = \bigoplus_{i=1}^{m} a_i(x) \cdot b_i(y)$ (see also [9]). This restricted variant of non-local box complexity is exactly the one that appears in van Dam's work [33], where he shows that any Boolean function f can be computed by such a protocol of complexity 2^n. Moreover, we prove that the restriction that the players output the XOR of the outcomes of the non-local boxes is without loss of generality.

THEOREM 3. $NL^{||,\oplus}(f) = \text{rank}_{\mathbb{GF}_2}(M_f) = D^{\|,IP_2}(f)$.

PROOF. We start by showing that $NL^{||,\oplus}(f) \leq \text{rank}_{\mathbb{GF}_2}(M_f)$. Let $\text{rank}_{\mathbb{GF}_2}(M_f) = t$, *i.e.*, $f(x,y) = \bigoplus_{i\in[t]} p_i(x) \cdot q_i(y)$. Then we construct a protocol that uses t non-local boxes in parallel, where Alice and Bob output the parity of the outcomes of the non-local boxes and for every input (x,y) the output of the protocol is equal to $f(x,y)$. The inputs of Alice and Bob to the i-th non-local box are the bits $p_i(x)$ and $q_i(y), i \in [t]$ respectively and let a_i, b_i the outputs of the non-local box such that $a_i \oplus b_i = p_i(x) \cdot q_i(y)$. Alice and Bob output at the end of the protocol the value $(\bigoplus_{i\in[t]} a_i) \oplus (\bigoplus_{i\in[t]} b_i) = \bigoplus_{i\in[t]} p_i(x) \cdot q_i(y) = f(x,y)$.

Conversely, if there exists a protocol where Alice and Bob use t non-local boxes in parallel with inputs $p_i(x), q_i(y)$ and outputs a_i, b_i, their final output is $(\bigoplus_{i\in[t]} a_i) \oplus (\bigoplus_{i\in[t]} b_i)$ and it always equals $f(x,y)$, then we have $f(x,y) = (\bigoplus_{i\in[t]} a_i) \oplus (\bigoplus_{i\in[t]} b_i) = \bigoplus_{i\in[t]} p_i(x) \cdot q_i(y)$ and hence $\text{rank}_{\mathbb{GF}_2}(M_f) \leq t$. From this last argument, we get $D^{\|,IP_2}(f) \leq NL^{||,\oplus}(f)$ since the players can send p_i and q_i to the referee who computes the inner product. For the converse, if the referee receives m_A, m_B from each player and computes their inner product mod 2, the players can instead input each bit of the message into a non-local box and output the parity of the outputs to obtain the same result.

For the next corollary, we use the fact that $\log(2\text{rank}_{\mathbb{F}}(M_f) - 1) \leq D(f) + 1$ for any field $\mathbb{F}$ (see [27]).

COROLLARY 4. $NL^{||,\oplus}(f) \leq 2^{D(f)}$.

On the other hand, it is easy to see that the one-way communication complexity is a lower bound on the non-local box complexity. Alice can send all her inputs to Bob, and since the non-local box protocol is always correct, they can simulate it, assuming that Alice received only zeros from all non-local boxes.

LEMMA 5. $D^{\rightarrow}(f) \leq NL(f)$.

Moreover, we show that both in the general and in the parallel model of deterministic non-local box complexity, we can assume without loss of generality that the players output the XOR of the outcomes of the non-local boxes. Unlike the general case, showing that

in the parallel case we can assume that the players output the XOR of the outputs of the non-local boxes is not a trivial statement.

THEOREM 6. *$NL(f) \leq NL^{\oplus}(f) \leq NL(f)+2$, and $NL^{||}(f) \leq NL^{||,\oplus}(f) \leq NL^{||}(f)+2$.*

Last, our bounds for deterministic non-local box complexity are tight as can be shown by looking at the Inner Product and Disjointness functions. Indeed, for Inner Product we have $D(IP) = NL^{||}(IP) = n$, while for Disjointness, $NL^{||,\oplus}(DISJ) = 2^{D(DISJ)} = 2^n$. For Disjointness, the circuit size upper bound also implies that $NL(DISJ) = O(n)$, so there is an exponential separation between NL and $NL^{||,\oplus}$.

4 Randomized non-local box complexity

In this section, we consider protocols that use shared randomness and have success probability at least $2/3$. We start by comparing the parallel non-local box complexity to communication complexity. In the full paper, we also exactly characterize $NL_{\varepsilon}^{||,\oplus}$ in terms of the approximate rank (over $\mathbb{GF}_2$) of the communication matrix.

THEOREM 7. *$R_{\epsilon}^{\rightarrow}(f) \leq NL_{\varepsilon}^{||}(f) \leq NL_{\varepsilon}^{||,\oplus}(f) \leq 2^{R_{\epsilon}(f)}$.*

Next, we relate the general non-local box complexity to the following model of communication: Alice and Bob send to a referee one message each and the referee outputs 1 if for the majority of indices, the two messages are equal. We denote the communication complexity in this model by $R_{\epsilon}^{||,MAJ}(f)$. This is a natural model of communication complexity that has appeared repeatedly in the simulation of quantum protocols by classical ones, as well as various upper bounds on simultaneous messages [25, 18, 32, 28].

THEOREM 8. *$R_{\varepsilon}^{\rightarrow}(f) \leq NL_{\varepsilon}(f) \leq O(R_{\epsilon}^{||,MAJ}(f))$.*

PROOF. The lower bound proof follows directly from the deterministic case. For the upper bound, fix a t-bit simultaneous protocol for f, where the referee receives two messages $\mathbf{a}$ and $\mathbf{b}$ of size t from Alice and Bob and outputs $MAJ(a_1 \oplus b_1, \ldots, a_t \oplus b_t)$. It is well-known, by using an addition circuit, that the majority of t bits can be computed by a circuit of size $O(t)$ with AND, NOT gates. Moreover, the distributed AND of two bits can be computed using two non-local boxes [6]. We conclude that the non-local box complexity of the distributed Majority is $O(t)$ and hence the theorem follows.

COROLLARY 9. *$2\log(\gamma_2^{\alpha}(f)/\alpha) \leq NL_{\varepsilon}(f) \leq O((\gamma_2^{\infty}(f))^2)$, where $\alpha = \frac{1}{1-2\varepsilon}$.*

It is known that $\gamma_2^{\infty}(f) = \Theta(\frac{1}{Disc(f)})$, and also that for any α, $\gamma_2^{\infty}(f) \leq \gamma_2^{\alpha}(f)$ [28]. Hence, since discrepancy gives a lower bound on the quantum communication complexity with entanglement $Q_{\varepsilon}^{*}(f)$ [24], we get the following corollary.

COROLLARY 10. *$NL_{\epsilon}(f) \leq O(2^{2Q_{\varepsilon}^{*}(f)})$.*

Finally, we can relate the non-local box complexity of a function f, to the L_1 norm of the Fourier coefficients of f by using a result by Grolmusz. Grolmusz showed that for any Boolean function f, there exists a randomized public coin protocol that solves f with complexity $O(L_1^2(f))$. This protocol can be easily transformed into a simultaneous messages protocol where the referee outputs the distributed majority of the message bits. Hence,

COROLLARY 11. $NL_\epsilon(f) \leq O(L_1^2(f))$.

Let us make here a last remark about the proof of Theorem 8. We started from a Simultaneous Messages protocol where the referee outputs a Majority function and we constructed a non-local box protocol with complexity equal to the communication complexity. If we look at this protocol, we can see that Alice and Bob can use their non-local boxes in the same order. This will be useful when we relate non-local boxes to secure function evaluation.

COROLLARY 12. $R_\varepsilon^{\rightarrow}(f) \leq NL_\varepsilon(f) \leq NL_\varepsilon^{\text{ord}}(f) \leq O(R_\varepsilon^{||,MAJ}(f))$.

Our bounds are almost tight for the general case, but the case of parallel non-local box complexity is more interesting. We can give a simple $O(n)$ parallel protocol for Disjointness, showing that the exponential separation does not hold anymore. It is open whether parallel and general randomized non-local box complexity are polynomially related.

5 Non-local boxes and measurement simulation

In this section we present another application of our results on non-local boxes. Using the recent breakthrough of Regev and Toner [31], who give a two-bit one-way protocol for simulating two-outcome measurements on entangled states for arbitrary dimensions, we show that this can be done with 3 non-local boxes. Previously, no finite upper bound was known for this problem. In the full paper, we prove the following, more general, theorem.

THEOREM 13. *For any non-signaling distribution over binary outputs with uniform marginals, any t-bit communication protocol can be simulated with $2^t - 1$ non-local boxes in parallel.*

The proof builds on an idea presented in [14] to replace communication by non-local boxes, which is here used recursively, and is given in the full paper.

6 Secure Function Evaluation

6.1 Honest-but-curious model

As a starting point, we consider the most basic model, namely deterministic secure computation with ANDs in the honest-but-curious model. Beimel and Malkin [4] have shown that $AND(f) \leq 2^{|\mathcal{X}|}$. We show that it is characterized by the one-way communication complexity of f. (The proof is given in the full paper.)

THEOREM 14. $AND(f) = 2^{D^{\rightarrow}(f)}$.

One can say that this shows that for most functions, randomization is necessary in order to construct efficient protocols even in the honest-but-curious model.

6.2 Malicious model

Due to their non-signaling property, protocols using non-local boxes only and no communication, such as those presented in the previous sections, are trivially secure even against malicious players. Indeed, the non-signaling property implies that the view of the protocol by a possibly dishonest player is always independent from the actions of the other player.

We show that certain such protocols may be transformed into protocols using OTs, namely the protocols where Alice and Bob use their non-local boxes in the same order. At this point, we don't know if this type of protocols are strictly weaker than general non-local box protocols. Nevertheless, our upper bounds in terms of communication complexity hold for such protocols as well (Corollary 12) and hence they translate into upper bounds on $OT_\varepsilon(f)$.

THEOREM 15. *For any* $\varepsilon \geq 0$, $OT_\varepsilon(f) \leq NL_\varepsilon^{\text{ord}}(f)$.

The proof, which will be given in the full version of the paper, consists in first showing how to simulate the non-local box protocol using OTs, following a construction due to Wolf and Wullschleger [34]. The security of the OT protocol then follows from the non-signaling property of the non-local boxes. From the above theorem we can conclude that all the upper bounds that we had for the $NL_\varepsilon^{\text{ord}}$ complexity (see Corollaries 9-12) translate into upper bounds for $OT_\varepsilon(f)$.

We now turn our attention to lower bounds. For this we need to restrict ourselves to what we call 'optimal' secure protocols. An 'optimal' secure protocol is one where the function is computed securely in the usual sense, but we also require that for all the OT calls, there is always an input that remains perfectly secure throughout the protocol. Intuitively, since we try to minimize the number of OTs that we use, it should be the case that these OT calls are really necessary, in the sense that one of the two inputs should always remain secure. If for example both inputs are revealed at some point during the protocol, then one may not use this OT at all, resulting into a more efficient protocol. Even though intuitively our definition seems natural, at this point, we do not know whether this assumption can be done without loss of generality. We denote by $\widehat{OT}_\varepsilon(f)$ the minimum number of OT calls of an 'optimal' secure protocol. In the full paper we provide the formal definition and prove the following

THEOREM 16. $\widehat{OT}_\varepsilon(f) = \Omega(R_\varepsilon(f))$.

7 Conclusion and open questions

We have shown various upper and lower bounds on non-local box complexity, and shown how the upper bounds could be translated into bounds for secure function evaluation. We have also shown how to simulate quantum correlations arising from binary measurements on bipartite entangled states using 3 non-local boxes.

During our investigations, we have come across a series of interesting open questions. 1) While the disjointness function provides an example of exponential gap between parallel and general deterministic non-local box complexity, the gap disappears in the randomized model. Are parallel and general randomized non-local box complexities polynomially related? 2) Are there functions for which $NL_\varepsilon^{\text{ord}}(f) > NL_\varepsilon(f)$? 3) We proved that the communication complexity is a lower bound on OT complexity only under some optimality assumption. Can this assumption be made without loss of generality? 4) Can we prove an analogue of Theorem 16 for non-local boxes? Ideally, we would like to prove that for secure computation with non-local boxes, communication does not help. Indeed, due to the

reduction from non-local boxes to OT boxes and vice versa, this would imply that $NL_\varepsilon^{\text{ord}}(f)$ is exactly $OT_\varepsilon(f)$, and not just an upper bound.

Acknowledgements

We would like to thank Troy Lee and Falk Unger for pointing out the deterministic protocol for Disjointness. We thank Ronald de Wolf for suggesting a randomized protocol for disjointness with bias $1/\log(n)$, using Valiant-Vazirani.

References

[1] Jonathan Barrett and Stefano Pironio. Popescu-Rohrlich correlations as a unit of nonlocality. *Phys. Rev. Lett.*, 95:140401, 2005.

[2] D. Beaver. Correlated pseudorandomness and the complexity of private computations. In *Proc. 28th STOC*, pages 479–488, 1996.

[3] A. Beimel, T. Malkin, and S. Micali. The all-or-nothing nature of two-party secure computation. In *Proc. CRYPTO '99*, pages 80–97, 1999.

[4] Amos Beimel and Tal Malkin. A quantitative approach to reductions in secure computation. In *Proc. TCC'04*, volume 2951, pages 238–257, 2004.

[5] M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for noncryptographic fault-tolerant distributed computations. In *Proc. 20th STOC*, pages 1–10, 1988.

[6] Gilles Brassard, Harry Buhrman, Noah Linden, Andre Allan Méthot, Alain Tapp, and Falk Unger. Limit on nonlocality in any world in which communication complexity is not trivial. *Phys. Rev. Lett.*, 96(25):250401, 2006.

[7] Anne Broadbent and André Allan Méthot. On the power of non-local boxes. *Theoretical Computer Science*, 358(1):3–14, 2005.

[8] Harry Buhrman, Matthias Christandl, Falk Unger, Stephanie Wehner, and Andreas Winter. Implications of superstrong nonlocality for cryptography. *Proc. Roy. Soc. A*, 462:1919–1932, 2007.

[9] Harry Buhrman and Ronald de Wolf. Communication complexity lower bounds by polynomials. In *Proc. 16th CCC*, pages 120–130, 2001.

[10] Nicolas J. Cerf, Nicolas Gisin, Serge Massar, and Sandu Popescu. Simulating Maximal Quantum Entanglement without Communication. *Phys. Rev. Lett.*, 94(22):220403, 2005.

[11] D. Chaum, C. Crepeau, and I. Damgard. Multiparty unconditionally secure protocols. In *Proc. 20th STOC*, pages 11–19, 1988.

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

[13] John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. Proposed Experiment to Test Local Hidden-Variable Theories. *Phys. Rev. Lett.*, 23:880–884, 1969.

[14] Julien Degorre, Sophie Laplante, and Jérémie Roland. Classical simulation of traceless binary observables on any bipartite quantum state. *Phys. Rev. A*, 75(012309), 2007.

[15] Yevgeniy Dodis and Silvio Micali. Lower bounds for oblivious transfer reductions. In *Proc. EUROCRYPT*, pages 42–55, 1999.

[16] Manuel Forster and Stefan Wolf. The universality of non-local boxes. In *Proc. 9th QCMC*, 2008. To appear.

[17] O. Goldreich and R. Vainish. How to solve any protocol problem - an efficiency improvement. In *Proc. CRYPTO '87*, pages 73–86, 1988.

[18] Vince Grolmusz. On the power of circuits with gates of low l_1 norms. *Theoretical Computer Science A*, 188:117–127, 1997.

[19] Nick S. Jones and Lluis Masanes. Interconversion of nonlocal correlations. *Phys. Rev. A*, 72(5):052312, 2005.

[20] J. Kilian. Basing cryptography on oblivious transfer. In *Proc. 20th STOC*, pages 20–31, 1988.

[21] J. Kilian. A general completeness theorem for two-party games. In *Proc. 23th STOC*, pages 553–560, 1991.

[22] J. Kilian. More general completeness theorems for two-party games. In *Proc. STOC*, pages 316–324, 2000.

[23] J. Kilian, E. Kushilevitz, S. Micali, and R. Ostrovsky. Reducibility and completeness in private computations. *SIAM J. Comput.*, 28(4):1189–1208, 2000.

[24] Ilan Kremer. Quantum communication. Master's thesis, The Hebrew University of Jerusalem, 1995.

[25] Ilan Kremer, Noam Nisan, and Dana Ron. On randomized one-way communication complexity. *Computational Complexity*, 8(1):21–49, 1999.

[26] E. Kushilevitz. Privacy and communication complexity. *SIAM J. Discrete Mathematics*, 5(2):273–284, 1992.

[27] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, New York, 1997.

[28] Nati Linial and Adi Shraibman. Lower bounds in communication complexity based on factorization norms. *Random Structures and Algorithms*, 2008.

[29] Moni Naor and Kobbi Nissim. Communication preserving protocols for secure function evaluation. In *Proc. 33rd STOC*, 2001.

[30] Sandu Popescu and Daniel Rohrlich. Causality and nonlocality as axioms for quantum mechanics. *Foundations of Physics*, pages 379–385, 1994.

[31] Oded Regev and Benjamin Toner. Simulating quantum correlations with finite communication. In *Proc. 48th FOCS*, pages 384–394, 2007.

[32] Yaoyun Shi and Yufan Zhu. Tensor norms and the classical communication complexity of bipartite quantum measurements. *SIAM J. Comput.*, 38(3):753–766, 2008.

[33] Wim van Dam. Implausible Consequences of Superstrong Nonlocality. Technical Report quant-ph/0501159, arXiv e-Print archive, 2005.

[34] Stefan Wolf and Jürg Wullschleger. Oblivious transfer and quantum non-locality. In *Proc. ISIT*, pages 1745–1748, 2005.

[35] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing. In *Proc. 11th STOC*, pages 209–213, 1979.