Amortizing Randomness in Private Multiparty Computations

Eyal Kushilevitz* Rafail Ostrovsky† Adi Rosén‡

August 25, 2002

Abstract

We study the relationship between the number of rounds needed to repeatedly perform a private computation (i.e. where there are many sets of inputs sequentially given to the players on which the players must compute a function privately) and the overall randomness needed for this task. For the XOR function, we show that for k sets of inputs, if instead of using totally fresh (i.e., independent) random bits for each of these k sets of inputs, we re-use the same ℓ random bits then we can significantly speedup the round-complexity of each computation compared to what is achieved by the naive strategy of partitioning the ℓ random bits between the k computations.

1 Introduction

A very basic question in the theory of computation is the *direct-sum* question: Can the complexity of solving k independent instances of a problem be smaller than the cost of independently solving the k instances? This general question was studied in various scenarios and with respect to various complexity measures, e.g., in [10, 20, 21, 22, 25, 37, 40]. To answer such a question one typically needs to consider a problem whose complexity in the single instance case is reasonably well understood.

In this work we consider a direct sum question related to the randomness complexity of private multiparty protocols. A *1-private* (or simply, *private*) protocol $\mathcal{A}$ for computing a function f is a protocol that allows n players, P_i, $1 \leq i \leq n$, each possessing an individual secret input, x_i, to compute the value of $f(\vec{x})$ in a way that no *single* player learns about the initial inputs of other players more than what is revealed by the value of $f(\vec{x})$ and its own input[1]. Private computations in this setting were the subject of a considerable amount of work, e.g., [5, 13, 2, 15, 34, 3, 21, 16, 17, 18,

*Dept. of Computer Science, Technion, Haifa, Israel. e-mail: eyalk@cs.technion.ac.il; Part of this research was done while visiting ICSI Berkeley. Supported by MANLAM Fund. URL: http://www.cs.technion.ac.il/~eyalk

†Bell Communications Research, MCC-1C365B, 445 South Street, Morristown, New Jersey 07960-6438; e-mail: rafail@bellcore.com

‡Dept. of Computer Science, University of Toronto, Toronto, Canada. Part of this work done while with the Dept. of Computer Science, Tel-Aviv University, Tel-Aviv, Israel. e-mail: adiro@cs.toronto.edu

[1]In the literature a more general definition of t-privacy is given. The above definition is the case $t = 1$.

30]. In this paper, we consider this setting for the basic XOR function, and show quite unexpected results relating the rounds complexity and the randomness complexity of such computations.

Randomness is an important resource in computation. As a result, various methods for saving in randomness were studied [6, 42, 36, 24, 19, 14, 41, 43, 35, 1, 39, 27, 28, 26]. In addition, there was a quantitative study of the role of randomness in specific contexts, e.g., [38, 33, 4, 11, 8, 9]. A particular case is the study of randomness in private multiparty computations [32, 7, 29]. In these papers the amount of randomness required for private computations of XOR (the exclusive-or function) was considered (this function was the subject of previous research in the area of privacy due to its being a basic linear operation and its relative simplicity [21, 16]).

In this paper we also concentrate on the XOR function. Previously, the following was known for privately computing XOR on a single instance (that is, where each player has a single input bit): there is no *deterministic* solution for the problem; with a single random bit the problem can be solved in $\Theta(n)$ rounds (time); with $\ell \geq 2$ random bits the problem can be solved in $O(\frac{\log n}{\log \ell})$ rounds and $\Omega(\frac{\log n}{\ell})$ rounds are required [32]. Other work on randomness in private computations appear in [31, 12], but none of the above mentioned works addressed multiple inputs.[2] We study the round-complexity problem in such multiple input settings.

Our Results: Before we make a precise statement of our general result, we start with a statement of a somewhat weaker version of out result which is simpler to state (and still is quite surprising): let us consider the case where n players are sequentially given n sets of inputs, of single bit each, and for each such set the players wish to privately compute XOR of these bits. In this case, by [32], if the players use only a single independent random bit for each set of inputs, they can compute all n xors, using n rounds for each set of inputs. It is impossible to compute the XOR's privately using the single random bits with less than $\Omega(n)$ rounds per computation. If the players wish to compute n independent XOR's with $O(1)$ rounds per computation, and using independent random bits for each computation, they will need at least $\Omega(n \log n)$ random bits overall (i.e. at least $\Omega(\log n)$ for each input set). Surprisingly, we show how to re-use the *same* $O(n)$ random bits for all computations and achieve optimal rounds-performance each time, i.e., each computation will be performed in $O(1)$ rounds. Moreover, we accompany this result by a corresponding lower bound showing that in order to privately compute the n XOR's the players need $\Omega(n)$ random bits *independently* of the number of rounds. Thus, using the minimum number of random bits possible, we achieve by a "recycling" procedure the optimal round-complexity.

More generally, we consider the setting where the players are sequentially given k input bits each, and the goal is to sequentially compute k times the XOR function immediately after the input bits for each set are provided. First, we prove that no solution exists which uses less than $(1 - \frac{2}{n})k$ random bits (hence, generalizing the claim that in the single input case there is no deterministic solution). If we have $\ell \geq k$ random bits, the naive solution would be to partition these bits to sets of $d \cong \ell/k$ random bits and to use the best single-input solution with d random bits for each sequential input. For example, if $\ell = k$ then $d = 1$ and so this solution requires $\Theta(n)$ rounds, and if $\ell = \Theta(k)$ then $d = O(1)$ and so this solution takes $\Theta(\log n)$ rounds [32]. We present much better solutions than the above. In particular, we show that if $\ell = k = c \cdot (n-1)$ the problem can be solved in $O(1)$ rounds per computation (rather than $\Theta(n)$), and for $\ell = O(k)$ we can solve the problem in $O(\frac{\log n}{\log k})$ rounds (rather than $Theta(\log n)$).

[2] Amortization of the *communication complexity* in private computation of XOR was shown in [21].

Organization: In Section 2 we provide the required definitions, including the model and the definition of privacy. In Section 3 we present our protocols that show our technique for recycling the random bits. Section 4 includes the lower bound which is technically more involved.

2 Preliminaries

In this section we give a description of the protocols we consider, and define the *privacy* property of protocols as well as the required complexity measures.

A set of n players P_i $(1 \leq i \leq n)$, each receiving sequentially k input bits $x_i^1, \ldots, x_i^k$ (known *only* to it), collaborate in a protocol to compute the k values

$$\begin{aligned} &\text{XOR}(x_1^1, x_2^1, \ldots, x_n^1) \\ &\text{XOR}(x_1^2, x_2^2, \ldots, x_n^2) \\ &\vdots \\ &\text{XOR}(x_1^k, x_2^k, \ldots, x_n^k). \end{aligned}$$

(In general, we may be interested in computing any function f.) More precisely, a protocol works in k *phases*. In phase j each player P_i gets the input bit x_i^j. Then, the players have to compute the j'th value $\text{XOR}(x_1^j, x_2^j, \ldots, x_n^j)$ and only after this computation is completed they get the $(j+1)$-st input bit. The computation in each phase operates in *rounds*. In each round each player P_i may toss some coins, and then sends messages to the other players (messages are sent over private channels so that other than the intended receiver no other player can listen to them). Player P_i then receives the messages sent to it by the other players. The content of these messages may depend on all the information available to P_i: its input (in the current and the previous phases), its random coins and the messages it received so far (in the current and the previous phases). Each player P_i receives during the execution of the protocol a sequence of messages C_i. We denote by C_i^j those messages sent to P_i during phase j. We also use x_i to denote the input seen by P_i during the whole protocol, i.e. $x_i = x_i^1 \cdots x_i^k$, and $\vec{x}$ to denote the vector of inputs seen by all players, i.e. $\vec{x} = (x_1, \ldots, x_n)$. We denote by x^j the vector of inputs received by all players in the j-th phase, i.e. $x^j = (x_1^j, \ldots, x_n^j)$. Finally, we use $f^k(\vec{x})$ to denote the vector of the k outputs of the protocol, i.e. $f^k(\vec{x}) = (f(x^1), \ldots, f(x^k))$.

Informally, *privacy* with respect to player P_i means that player P_i cannot learn anything (in particular, the inputs of the other players) from C_i, except what is implied by its input bits, and the value of the function computed. Formally,

Definition 1: (Privacy) A (k phase) protocol $\mathcal{A}$ for computing a function f is *private* with respect to player P_i if for any two input vectors $\vec{x}$ and $\vec{y}$, such that $f^k(\vec{x}) = f^k(\vec{y})$ and $x_i = y_i$, for any sequence of messages C, and for any random coins, R_i, tossed by P_i,

$$Pr[C_i = C | R_i, \vec{x}] \quad = \quad Pr[C_i = C | R_i, \vec{y}],$$

where the probability is over the random coin tosses of *all other* players.
A protocol is called *private* if it is private with respect to every P_i.

We note that a natural k phase protocol for computing a function f is simply to use independently k times a private protocol for computing f on a single set of inputs. Such a solution seems to give a stronger property: no matter what happened in the first $j-1$ phases, the j-th phase gives no information. Formally,

Definition 2: (Strong Privacy) A (k phase) protocol $\mathcal{A}$ for computing a function f is *strongly private* with respect to player P_i if for any two input vectors $\vec{x}$ and $\vec{y}$, such that $f(x^j) = f(y^j)$ and $x_i^j = y_i^j$, for any sequence of messages C, for any history $C_i^1, \ldots, C_i^{j-1}$, and for any random coins, R_i, tossed by P_i,

$$Pr[C_i^j = C | R_i, \vec{x}, C_i^1, \ldots, C_i^{j-1}] \quad = \quad Pr[C_i^j = C | R_i, \vec{y}, C_i^1, \ldots, C_i^{j-1}],$$

where the probability is over the random coin tosses of *all other* players.
A protocol is called *strongly private* if it is strongly private with respect to every P_i.

Obviously, if a protocol is strongly private then it is also private.

To measure the amount of randomness used by a protocol and its round complexity we give the following definitions:

Definition 3: An ℓ-random protocol is a protocol such that for every input assignment $\vec{x}$, the total number of coins tossed by all players in *every* execution is at most ℓ.

Definition 4: An r-round protocol is a protocol such that for every input assignment $\vec{x}$, and every sequence of coin tosses the number of rounds in each phase j is at most r.

We emphasize that the definitions allow, for example, that in different executions different players will toss the coins. This may depend on both the input of the players, and previous coin tosses.

3 Upper Bound

In this section we present our positive results. First, we consider the case $k = n - 1$. By the lower bound of Section 4, at least $n - 2$ random bits are needed for such a computation, regardless of the number of rounds per phase. The protocol below uses $n - 1$ random bits. Of course, there is a naive way to perform this computation using only $n - 1$ random bits, which is to use a single random bits for each of the $n - 1$ phases. However, such computation takes $O(n)$ rounds per phase and by [32] this is the best one can do with a single random bit (per phase). Our protocol takes a different direction that allows it to use only $O(1)$ rounds per phase.

Lemma 1: There is a strongly private, k-phase protocol that computes XOR on $k = n - 1$ inputs with $\ell = n - 1$ random bits and $r = O(1)$ rounds per phase.

Proof: For the proof we present an appropriate k-phase protocol.

Initialization: Player P_n chooses $n - 1$ random bits, denoted $r_1, \ldots, r_{n-1}$. It sends bit r_i to player P_i. (This step can be performed as part of phase 1.)

Phase j:

1. Each player P_i, $1 \leq i \leq n-1$, sends a bit $b_{i,j} = x_i^j + r_i$ to player P_j.
   In addition, player P_n sends to P_j the bit $b_{n,j} = x_n^j + \sum_{m=1}^{n-1} r_m$. (The summations here and elsewhere are all modulo 2.)

2. Player P_j sums the n bits $b_{i,j}$ it received in the previous step. It announces $\sum_{i=1}^{n} b_{i,j}$ as the output for the j-th phase.

For the correctness, consider the sum computed by player P_j in Step (2). This sum equals:

$$\sum_{i=1}^{n} b_{i,j} = \sum_{i=1}^{n-1} (x_i^j + r_i) + (x_n^j + \sum_{m=1}^{n-1} r_m) = \text{XOR}(x_1^j, x_2^j, \ldots, x_n^j).$$

For the privacy, note that P_n receives no message during the protocol, hence the privacy with respect to P_n certainly holds. Also, observe that during phase j only player P_j receives any message (other than the output of the phase). Finally, C_j^j (the communication received by P_j in phase j) is a sequence of messages $b_{1,j}, \ldots, b_{n-1,j}, b_{n,j}$. Note that for every input $\vec{x}^j = (x_1^j, x_2^j, \ldots, x_n^j)$, every communication $r, c_1, \ldots, c_{n-1}, c_n$ which is consistent with the output has the same probability, $2^{-(n-1)}$. This is because each of the first c_i's determines one of the $n-1$ random bits (which are all independent) and c_n is already determined by the output and the previously determined values. The strong privacy of the protocol follows. □

The main idea in the above protocol is that we can compute the XOR of each of the $n-1$ inputs, using $n-1$ random bits, in a way that allows to use the same $n-1$ random bits for all the $n-1$ inputs. We can use the same idea, with a bit more of precaution to do it for other parameters. The following is a simple corollary of the previous construction.

Lemma 2: There is a strongly private, k-phase protocol that computes XOR on $k = d(n-1)$ inputs with $\ell = d(n-1)$ random bits and $r = O(1)$ rounds.

Proof: Simply partition the $d(n-1)$ inputs to d sets of size $n-1$. For each set of $n-1$ inputs use the $(n-1)$-phase protocol of Lemma 1 that requires $n-1$ random bits and $O(1)$ rounds. If we do this each time with new and independent random bits, we get the desired result. □

Again note that, by the results of Section 4, the above lemma is (almost) optimal in terms of the number of random bits required for this computation. Moreover, if k (the number of inputs) is not divisible by $n-1$ we can always add some dummy inputs to each player so as to make the number of inputs be some k' which is divisible by $n-1$. For example, if $\frac{n-1}{2} \leq k < n-1$, then $\ell = n-1$ random bits and $r = O(1)$ rounds suffice. The only case where this is inefficient is when $k \ll n-1$ as increasing the number of inputs to $k' = n-1$ would be inefficient. For such cases, we use the following construction:

Lemma 3: There is a strongly private, k-phase protocol that computes XOR on k inputs, $k < (n-1)/2$ with $\ell = 2k + s$ random bits and $r \leq \log n / \log s$ rounds per computation.

Proof: As in the previous protocols, P_n will be the player that makes the random choices. We partition the other $n-1$ players into g groups of size k. Assume for simplicity that $n-1$ is divisible by k. Moreover that $g = (n-1)/k$ is even (later we describe the modifications required when this is not the case).

Initialization: Player P_n chooses k random bits, denoted $r_1, \ldots, r_k$. It sends the bit r_i to the i-th player of each of the g groups. P_n chooses additional s bits $\alpha_1, \ldots, \alpha_s$ to be used later.

Phase j:

1. The i-th player in each group t ($1 \leq i \leq k$, $1 \leq t \leq g$), denoted $P_{i,t}$, sends a bit $b^j_{i,t} = x^j_{i,t} + r_i$ to player $P_{j,t}$ (where $x^j_{i,t}$ denotes the input of $P_{i,t}$ in the current phase, j).

2. Each player $P_{j,t}$ computes $Y^j_t = \sum_{i=1}^{k} b^j_{i,t}$.

3. The j-th players of all groups together with P_n involve in a private protocol to compute the sum of $g+1$ bits: $Y^1_j, \ldots, Y^g_j$ and x^j_n. They announce the output as the XOR of the j-th input. The players do this computation using the protocol of [32]. This protocol, when using s random bits, terminates within $\log(g+1)/\log s$ rounds. In addition, all the random bits in this protocol are chosen by one player which other than that receives no message during the protocol; we choose this player to be P_n and $\alpha_1, \ldots, \alpha_s$ to be these random bits (note that P_n uses the *same* s random bits in all k phases).

For the correctness note that

$$\sum_{t=1}^{g} Y^t_j \;+\; x^j_n = \sum_{t=1}^{g}\sum_{i=1}^{k} b^j_{i,t} + \; x^j_n = \sum_{t=1}^{g}\sum_{i=1}^{k}(x^j_{i,t} + r_i) + \; x^j_n = \sum_{i=1}^{n} x^j_i + g \cdot (\sum_{i=1}^{k} r_i).$$

Since we assumed that g is even, the last term contributes 0 to the sum (modulo 2) and so the j-th output is $\sum_{i=1}^{n} x^j_i$, as needed. The number of random bits used is $k+s$.

For the privacy, note again that P_n only sends messages during the whole protocol and that in phase j only the j-th player of each group receives messages. Each of the players $P_{j,t}$ receives in Step 1 from the members of his group, k bits $b^j_{1,t}, \ldots, b^j_{k,t}$ which are distributed uniformly and independently. Then, it involves in a private protocol which guarantees that no matter what is the input to the protocol each player sees the same distribution of communications, for all possible inputs $P_{j,t}$ may have and all possible outputs. Altogether, the strong privacy follows.

Now, there are some technical issues that we still need to take care of. First, if the number of groups g is odd then P_n can always make sure that there will be no contribution of random bits to the result by using as its input in Step 3 of phase j the bit x^j_n xored with these random bits. Another technical issue that has to be dealt with is the case that $n-1$ is not divisible by k. In this case the g-th group is of size $k' < k$ and hence cannot use the above protocol. We solve this by letting P_n choose for them k' special random bits. The messages $b^j_{i,g}$ will be sent to the j-th player of group 1 instead of the j-th player of group g (who may not exist). Since this is done with new random bits the privacy still holds. □

Combining Lemma 2 and Lemma 3 we get the following theorem:

Theorem 4: There is a strongly private k-phase protocol that computes XOR on $k = d(n-1)+q$ inputs, with $\ell = d(n-1) + 2q + s$ random bits and $r \leq \log n / \log s$ rounds.

An interesting case is the following:

Corollary 5: There is a strongly private, k-phase protocol that computes XOR on k inputs, with $\ell = O(k)$ random bits and $r \leq \log n / \log k$ rounds per phase.

This means that we can use $O(k)$ random bits, use for each phase the time which is the best known upper bound for computing XOR on a single input using k bits (i.e., $O(\log n/\log k)$ rounds), and while the lower bound to compute XOR on a single input with k random bits is $\Omega(\log n/k)$ rounds [32].

4 Lower Bound

In this section we prove a lower bound on the number of random bits required for the private computation of k instances of XOR, in a sequential manner. The result that we prove is for any number of rounds used by the protocol. We prove the following theorem.

Theorem 6: Let A be a d-random n-player private protocol to compute sequentially k instance of XOR. Then $d \geq (1 - \frac{2}{n}) \cdot k$.

The proof of this theorem is based on the following two lemmas. The proof of the first lemma is immediate from the proofs in [32], and we omit it from this abstract. Before we proceed we give a technical definition and a technical proposition.

Definition 5: Denote by $View_i^t(\vec{x}, \vec{b})$ be the view of player P_i at round t on input $\vec{x}$ and vector of random tapes $\vec{b}$. This view consists of the inputs to player P_i, the random tape of plater P_i, R_i, and the communication received by player P_i until, and including, round $t-1$.

Claim 1: For any non-negative values $a_{i,j}$,

$$\Pi_{i=1}^{q}(\sum_{j=1}^{p} a_{i,j}) \geq p^q \min_{1 \leq j \leq p} \{\Pi_{i=1}^{q} a_{i,j}\}$$

Proof:

$$\begin{aligned}
\prod_{j=1}^{q}\sum_{i=1}^{p} a_{i,j} &= p^q \prod_{j=1}^{q}\left(\frac{\sum_{i=1}^{p} a_{i,j}}{p}\right) \\
&\geq p^q \prod_{j=1}^{q}\left(\prod_{i=1}^{p} a_{i,j}\right)^{\frac{1}{p}} \\
&= p^q \left(\prod_{i=1}^{p}\prod_{j=1}^{q} a_{i,j}\right)^{\frac{1}{p}} \\
&\geq p^q \min_{1 \leq i \leq p}\left(\prod_{j=1}^{q} a_{i,j}\right)
\end{aligned}$$

□

Lemma 7: Consider a private d-random protocol to compute k instances of a boolean function f. Fix the random tapes of the players to be $\vec{R}$. Then, for any P_i, $View_i^k(\vec{y}, \vec{R})$ can assume at most 2^{2k+d} different values (over the 2^{kn} values of $\vec{y}$).

Lemma 8: Let A be a deterministic (non-private) n-player protocol to sequentially compute k instances of XOR. Then there is at least one player that has at least $2^{(3-\frac{2}{n})k}$ views over the 2^{kn} input assignments.

Proof of Theorem 6. By Lemma 7, if we fix the random tapes of the players then each player can see (over the different inputs) at most 2^{2k+d} different views. But by Lemma 8 for the protocol to be correct there must be at least one player that sees at least $2^{(3-\frac{2}{n})k}$ views. Thus $2^{(3-\frac{2}{n})k} \leq 2^{2k+d}$, and the theorem follows. □

To prove Lemma 8, we now turn to the main technical claim of this section. For the purpose of the proof, we consider protocols that compute k instances of XOR, but such that for the first instance only m of the n players get inputs (alternatively, we can assume that the input of $n-m$ of the players for the first instance is 0). For $k \geq 1$ and $1 \leq m \leq n$ let $\mathcal{A}(k,m)$ be the set of protocols that correctly compute this function (sequentially). We prove the following lemma.

Lemma 9: Let $A \in \mathcal{A}(k,m)$. Let V_i^A be the number of different views player P_i can see over the $2^{(k-1)n+m}$ inputs. Then $\Pi_{s=1}^n V_i^A \geq 2^{(3n-2)(k-1)+n+2(m-1)}$.

Proof: We prove the claim by induction on both k and m, where the base case being $k=1, m=1$. Let $A \in \mathcal{A}(1,1)$. That is, one player has an input bit and A has to ensure that all players "compute" the value of this bit. Obviously for all P_i we have $V_i^A \geq 2$ (as there are two output values), which gives $\Pi_{s=1}^n V_i^A \geq 2^n$, as required. For the induction step, let $A \in \mathcal{A}(k,m)$, for $k > 1$ *or* $m > 1$. We consider two cases, $m > 1$ and $m = 1$.

$m > 1$**:** Before the first XOR value is computed by any player there must be one non-constant message sent in the protocol. That is, there must be some player P_i that sends a message to player P_j, and this message is not constant over all input assignments. Consider the first such non-constant message, and w.l.o.g., let it be sent from P_i to P_j. Denote this message by M. Since M is the first non-constant message it depends on the current input of P_i only. W.l.o.g., assume that P_i sends the value of its input bit. Let ℓ_s^0 (resp. ℓ_s^1) be the number of possible views of player P_s given that the value of M is 0 (resp. 1). We get that

- $V_i^A = \ell_i^0 + \ell_i^1$.
- $V_j^A = \ell_j^0 + \ell_j^1$.
- $\forall k \neq i,j, V_k^A \geq \max(\ell_k^0, \ell_k^1)$.

Therefore,

$$\begin{aligned}
\Pi_{s=1}^n V_s^A &\geq (\ell_i^0 + \ell_i^1)(\ell_j^0 + \ell_j^1)\Pi_{s\neq i,j} \max(\ell_s^0, \ell_s^1) \\
&\geq 4\min(\ell_i^0\ell_j^0, \ell_i^1\ell_j^1)\Pi_{s\neq i,j} \max(\ell_s^0, \ell_s^1) \\
&= 4\ell_i^0\ell_j^0\Pi_{s\neq i,j} \max(\ell_s^0, \ell_s^1) \\
&\geq 4\Pi_{s=1}^n \ell_s^0
\end{aligned}$$

Where the equality is by assuming w.l.o.g. that $\ell_i^0\ell_j^0 \leq \ell_i^1\ell_j^1$.

Now consider a protocol A_0 defined as follows. It is the protocol A with the modification that P_i has no input, and behaves as if its input is 0. Since we assume that A is a correct protocol, then

A_0 is a correct protocol as well in the class $\mathcal{A}(k, m-1)$. [3] Also, we know that A_0 sends 0 as the value of M. Therefore for any s, $1 \leq s \leq n$, we have $V_s^{A_0} = \ell_s^0$. We get

$$\Pi_{s=1}^n V_s^A \geq 4\Pi_{s=1}^n \ell_s^0 = 4\Pi_{s=1}^n V_s^{A_0} \geq 4 \cdot 2^{(3n-2)(k-1)+n+2(m-2)} ,$$

where the last inequality follows from the induction hypothesis. We get that

$$\Pi_{s=1}^n V_s^A \geq 2^{(3n-2)(k-1)+n+2(m-1)} ,$$

which concludes the proof of the first case.

$m = 1$: This is the case where in the first iteration there is a single player who has an input bit. The value of the function on this input has to be computed by all players before they get the next input to be computed. Therefore, the first step of the protocol must be that all players receive messages from which each player can conclude if this first input is 0 or 1. It follows for each P_i, that $V_i^A = \ell_i^0 + \ell_i^1$, where ℓ_i^0 (resp. ℓ_i^1) is the number of different views of player P_i given that the first input bit is 0 (resp. 1). Also note that all players agree on the output. We get,

$$\Pi_{s=1}^n V_i^A = \Pi_{s=1}^n (\ell_i^0 + \ell_i^1) .$$

Using Proposition 1 we have

$$\Pi_{s=1}^n (\ell_i^0 + \ell_i^1) \geq 2^n \min(\Pi_{s=1}^n \ell_i^0, \Pi_{s=1}^n \ell_i^1) ,$$

and assuming w.l.o.g. that $\Pi_{s=1}^n \ell_i^0 \geq \Pi_{s=1}^n \ell_i^1$ we get

$$\Pi_{s=1}^n V_i^A \geq 2^n \Pi_{s=1}^n \ell_i^0 .$$

By the same arguments as those for the first case we now consider a protocol $A_0 \in \mathcal{A}(k-1, n)$ defined using protocol A, and have that $V_i^{A_0} = \ell_i^0$, for any P_i. Using the induction hypothesis we have

$$\Pi_{s=1}^n V_i^A \geq 2^n \Pi_{s=1}^n \ell_i^0 = 2^n \Pi_{s=1}^n V_i^{A_0} \geq 2^n s^{(3n-1)(k-2)+n+2(n-1)} = 2^{(3n-2)(k-1)+n} ,$$

which concludes the proof for the second case. □

We can now complete the proof of Lemma 8:

Proof of Lemma 8. Let $A \in \mathcal{A}(k, n)$. Then, by Lemma 9, $\Pi_{s=1}^n V_i^A \geq 2^{(3n-2)(k-1)+n+2(n-1)} = 2^{(3n-2)k}$. Therefore there is at least one player P_i such that $V_i^A \geq 2^{(3-\frac{2}{n})k}$. □

Acknowledgments

We wish to thank Pino Persiano for several interactions regarding this paper. We also wish to thank the referees for useful comments.

[3] In case $\ell_i^0 \ell_j^0 > \ell_i^1 \ell_j^1$ we consider a protocol A_1, that behaves as if the input to P_i is 1, but also negates the outputs of the first set of inputs.

References

[1] N. Alon, O. Goldreich, J. Hastad, and R. Peralta, "Simple constructions of almost k-wise independent random variables", Proc. of 31st FOCS, 1990, pp. 544–553.

[2] J. Bar-Ilan, and D. Beaver, "Non-Cryptographic Fault-Tolerant Computing in a Constant Number of Rounds", Proc. of 8th PODC, 1989, pp. 201–209.

[3] D. Beaver, "Perfect Privacy for Two-Party Protocols", TR-11-89, Harvard University, 1989.

[4] M. Bellare, O. Goldreich, and S. Goldwasser, "Randomness in Interactive Proofs", Proc. of 31st FOCS, 1990, pp. 563–571.

[5] M. Ben-or, S. Goldwasser, and A. Wigderson, "Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation", Proc. of 20th STOC, 1988, pp. 1–10.

[6] M. Blum, and S. Micali "How to Generate Cryptographically Strong Sequences Of Pseudo- Random Bits", FOCS 82 and *SIAM J. on Computing,* Vol 13, 1984, pp. 850–864.

[7] C. Blundo, A. De-Santis, G. Persiano, and U. Vaccaro, "On the Number of Random Bits in Totally Private Computations", *Proc. of 22nd ICALP*, 1995.

[8] C. Blundo, A. Giorgio Gaggia, and D. R. Stinson, "On the Dealer's Randomness Required in Secret Sharing Schemes", *Proc. of EuroCrypt94.*

[9] C. Blundo, A. De-Santis, and U. Vaccaro, "Randomness in Distribution Protocols", *Proc. of 21st ICALP*, 1994.

[10] Bshouty, N. H., "On The Extended Direct Sum Conjecture", *Proc. of 21th ACM Symposium on Theory of Computing*, 1989, pp. 177-185.

[11] R. Canetti, and O. Goldreich, "Bounds on Tradeoffs between Randomness and Communication Complexity", FOCS 90 and *Computational Complexity* Vol. 3, (1993), 141-167.

[12] R. Canetti, E. Kushilevitz, R. Ostrovsky, and A. Rosén, "Randomness vs. Fault-Tolerance", *Proc. of 16th PODC*, pp. 35-44, August 1997.

[13] D. Chaum, C. Crepeau, and I. Damgard, "Multiparty Unconditionally Secure Protocols", Proc. of 20th STOC, 1988, pp. 11–19.

[14] B. Chor, and O. Goldreich, "Unbiased Bits from Sources of Weak Randomness and Probabilistic Communication Complexity", FOCS 85 and *SIAM J. Computing* Vol. 17, (1988), 230–261.

[15] B. Chor, and E. Kushilevitz, "A Zero-One Law for Boolean Privacy", STOC 89 and *SIAM J. Disc. Math.* Vol. 4, (1991), 36–47.

[16] B. Chor, and E. Kushilevitz, "A Communication-Privacy Tradeoff for Modular Addition", *Information Processing Letters*, Vol. 45, 1993, pp. 205–210.

[17] B. Chor, M. Geréb-Graus, and E. Kushilevitz, "Private Computations Over the Integers", Proc. of 31st FOCS, 1990, pp. 335–344.

[18] B. Chor, M. Geréb-Graus, and E. Kushilevitz, "On the Structure of the Privacy Hierarchy", *Journal of Cryptology*, Vol. 7, No. 1, 1994, pp. 53-60.

[19] A. Cohen, and A. Wigderson, "Dispersers, Deterministic Amplification, and Weak Random Sources", Proc. of 30th FOCS, 1989, pp. 14–19.

[20] T. Feder, E. Kushilevitz, M. Naor, and N. Nisan, "Amortized Communication Complexity", *SIAM J. Computing*, Vol. 24, No. 4, pp. 736-750, 1995. (Early version in *Proceedings of the 32nd FOCS*, pp. 239-248, 1991.)

[21] M. Franklin, and M. Yung, "Communication Complexity of Secure Computation", Proc. of 24th STOC, 1992, pp. 699–710.

[22] Galibati G., and M. J. Fischer, "On The Complexity of 2-Output Boolean Networks", *Theoretical Computer Science*, Vol 16, 1981, pp. 177–185.

[23] Hardy, Littlewood, Polya, **inequalities**, Cambridge University Press.

[24] R. Impagliazzo, and D. Zuckerman, "How to Recycle Random Bits", Proc. of 30th FOCS, 1989, pp. 248–253.

[25] M. Karchmer, E. Kushilevitz, and N. Nisan, "Fractional Covers and Communication Complexity", *SIAM J. Discrete Math.*, Vol. 8, No. 1, pp. 76-92, 1995. (Early version in *7th IEEE Structure in Complexity Theory*, pp. 262-274, 1992.)

[26] D. Karger, and D. Koller, "(De)randomized Construction of Small Sample Spaces in NC", Proc. of 35th FOCS, 1994.

[27] D. Koller, and N. Megiddo, "Constructing Small Sample Spaces Satisfying Given Constraints", Proc. of 25th STOC, 1993, pp. 268–277.

[28] H. Karloff, and Y. Mansour, "On Construction of k-wise Independent Random Variables", Proc. of 26th STOC, 1994, pp. 564–573.

[29] E. Kushilevitz, and Y. Mansour, "Randomness in Private Computations", manuscript 1995.

[30] E. Kushilevitz, S. Micali, and R. Ostrovsky, "Reducibility and Completeness in Multi-Party Private Computations", Proc. of 35th FOCS, 1994, pp. 478-489.

[31] E. Kushilevitz, R. Ostrovsky, and A. Rosén, "Characterizing Linear Size Circuits in Terms of Privacy", Proc. of 28th STOC, pp. 541–550, 1996.

[32] E. Kushilevitz, and A. Rosén, "A Randomness-Rounds Tradeoff in Private Computation", *Advances in Cryptology*, Proceedings of Crypto '94, Y. Desmedt, ed., Springer-Verlag, Lecture Notes in Computer Science, Vol. 839, pp. 397-410, 1994.

[33] D. Krizanc, D. Peleg, and E. Upfal, "A Time-Randomness Tradeoff for Oblivious Routing", Proc. of 20th STOC, 1988, pp. 93–102.

[34] E. Kushilevitz, "Privacy and Communication Complexity", FOCS 89, and SIAM Jour. on Disc. Math., Vol. 5(2), pp. 273–284, 1992.

[35] J. Naor, and M. Naor, "Small-Bias Probability Spaces: Efficient Constructions and Applications", STOC 90, and *SIAM J. on Computing,* Vol 22, No. 4, 1993, pp. 838–856.

[36] N. Nisan, "Pseudorandom Generator for Space Bounded Computation", Proc. of 22nd STOC, 1990, pp. 204–212.

[37] Paul W., "Realizing Boolean Function on Disjoint Sets of Variables", *Theoretical Computer Science* 2, 1976, pp. 383-396.

[38] P. Raghavan, and M. Snir, "Memory vs. Randomization in On-Line Algorithms", Proc. of 16th ICALP, 1989, pp. 687–703.

[39] L. J. Schulman, "Sample Spaces Uniform on Neighborhoods", Proc. of 24th STOC, 1992, pp. 17–25.

[40] Q. F. Stout, "Meshes with multiple buses", *Proc. of 27th IEEE Symposium on Foundations of Computer Science,* 1986, pp. 264-273.

[41] U. Vazirani, and V. Vazirani, "Random Polynomial Time is Equal to Slightly-Random Polynomial Time", Proc. of 26th FOCS, 1985, pp. 417–428.

[42] A. C. Yao, "Theory and Applications of Trapdoor Functions", Proc. of 23rd FOCS, 1982, pp. 80–91.

[43] D. Zuckerman, "Simulating BPP Using a General Weak Random Source", Proc. of 32nd FOCS, 1991, pp. 79–89.