Electronic Colloquium on Computational Complexity, Report No. 209 (2015)

On the information leakage of public-output protocols

Eli Ben-Sasson, Gal Maor

December 29, 2015

Abstract

In this paper three complexity measures are studied: (i) internal information, (ii) external information, and (iii) a measure called here “output information”. Internal information (i) measures the counter-party privacy-loss inherent in a communication protocol. Similarly, the output information (iii) measures the reduction in input-privacy that is inherent when the output of the computation is published. External information (ii) measures the privacy-loss inherent when the communication between parties is leaked.

The main result here is that for public-output protocols, i.e., protocols that require publicly posting the result of the computation, external information can be exponentially larger than the sum of internal and output information. This result can be informally summarized as saying that “viewing a process from the sidelines (external information) can be vastly more informative than either participating in it (internal information) and/or viewing only its outcome (output information)”.

The specific problem for which this exponential separation is shown, and the upper bounds on internal and output information follow closely the recent breakthrough separation of information and communication due to [Ganor, Kol, Raz; FOCS 2014]. The main technical contribution is the lower bound on external information for public-output protocols.

ISSN 1433-8092

1 Introduction

An interesting problem in complexity theory is that of understanding the interplay between communication and privacy in a distributed setting. Following [KN97, Bra12] we measure "privacy-loss" by information complexity, i.e., as the amount of information that a party learns about inputs of other parties when participating in a communication protocol aimed at solving a fixed computational problem.

The study of such privacy questions goes back to Kushilevitz [Kus92] who gave a complete characterization of the functions which can be computed with *perfect* internal privacy; namely, functions which can be calculated when the parties learn nothing except for the output itself. Later, in both [BYCKO93] and [BW12] it is shown that for the 2-party setting, most boolean functions have linear internal and external information complexities. This means that for calculating most functions, the parties must reveal roughly *all* of their inputs to each other and to an external observer while solving the task, even though the output, a single bit, reveals at most one bit of information about the inputs. Pitassi et al. [ACC+14] analyzed the second price auction function, and showed a tradeoff between the privacy and the communication complexity of calculating it. For more about information complexity, see section 1.1.

In the most well-studied model of communication, a number m of parties, each holding some private input x_i from some domain X_i, wish to jointly find a solution z that is compatible with their inputs. In the *functional* version, a function $f : X_1 \times \ldots, \times X_m \to Z$ is fixed and the parties wish to compute $z = f(x_1, \ldots, x_m)$. In the more general, *relational* version studied here, a relation $R \subseteq X_1 \times \ldots X_m \times Z$ is fixed and the parties wish to solve the *search problem* defined by R, meaning they wish to find some z such that $(x_1, \ldots, x_m, z) \in R$.

We are interested in a special case of the relational version, in which two parties, Alice and Bob, need to solve a search problem and then *publish* the solution to a third party, Charlie; we describe such protocols as *public output* protocols. When Alice and Bob participate in a public output protocol, we assume Alice wishes to reveal as little information to Bob as possible and vice versa, and both parties jointly wish to reveal as little information as possible to Charlie, beyond what can be learned from the public output (the solution).

As in [Kus92, BW12, ACC+14], we study this privacy-loss problem in the information theoretic setting, one that does not involve cryptographic assumptions, and refer the reader to the introduction of [ACC+14] for a comprehensive discussion and justification for the non-cryptographic study of privacy via information theory.

As a concrete motivating example for studying privacy loss in the public output model, consider the following natural scenario. A government office (Charlie) puts a project out to tender. Two companies (Alice and Bob) wish to cooperate in order to make their joint proposal, which will eventually be visible to the office (this is the public output). Naturally, each company wishes to disclose as little information as possible to the other company about its internal data, and both companies have a shared interest in jointly minimizing the amount of information disclosed to the government, beyond what can be deduced from their final proposal (the public output).

Our main result can be summarized as saying:

> Viewing a process from the sidelines may reveal vastly more information than either participating in it, or viewing only its outcome.

In more detail, our main result is (see Theorem 2 for the formal statement):

Theorem 1 (Main Theorem — Informal)**.** *There exists a search problem with input size parameterized by k (and a distribution on inputs) such that:*

- *There exists a 2-party communication protocol involving Alice and Bob that solves the search problem (with probability* 0.99*), while revealing to Bob at most $O(k)$ bits of information about Alice's input, and vice versa.*
- *There exists a 2-party communication protocol between Alice and Bob that solves the search problem whose output reveals to Charlie at most $O(k)$ bits of information about the inputs.*
- *The transcript of* any *public output communication protocol between Alice and Bob that solves the problem (with probability* 0.99*) must reveal to Charlie $2^{\Omega(k)}$ bits of information about the inputs.*

This is the first example of a task for which the internal and external information complexities are exponentially apart, when the mere output of the task doesn't provide significant information by itself. The internal information upper bound is due to Ganor, Kol and Raz [GKR14], as we use a simplified version of their construction (see section 3.1 for details). Our main contribution is the proof of the external information lower bound.

The study of privacy-loss in a communication setting involving three parties begs for a comparison to the seminal work of Goldwasser, Ben-Or and Wigderson [BOGW88]; they showed that, when 3 parties or more communicate via private channels, any problem can solved in a way that reveals no information at all to any of the involved parties, beyond what can be deduced from the output itself (and that party's input). At first it may seem that our result contradicts that of [BOGW88] but a moments reflection shows that the two cases are somewhat different. In our setting we assume Charlie listens to the communication between Alice and Bob, i.e., we do not assume private channels; indeed, when Bob does not listen to the communication clearly the only thing he sees is the public output. A better way of comparing the result of Ben-Or et al. to our case would be to say that if Charlie is willing to *help* Alice and Bob, they can participate in a 3-party protocol in which not only Charlie, but even Alice and Bob, learn nothing beyond the public output (and whatever can be deduced from it). We stress that this would assume that Charlie has no access to the communication channel of Alice and Bob, which is not the case in this paper.

Informally, in the search problem used to prove Theorem 1 each party receives as input a (binary) labeling of a binary tree, and the distribution we use gives Alice and Bob nearly the same labeling, but for one randomly selected "cross-section" (called also a "multi-layer") of the tree in which the labelings of Alice and Bob are independent (and uniformly distributed). The search problem of interest here is that of finding a path in the tree that nearly-agrees with Alice's labeling on the even vertices in the "important" cross-section (in which Alice and Bob's inputs diverge), and with Bob's labeling on the odd vertices in that cross-section (for complete details, see section 3.1). Our lower bound proof is rather intuitive, and relies on the fact that the parties have two options: (i) either to find the location of the "important" cross section (where the two labelings differ), and this option reveals to Alice a significant amount of information about Bob's input (and vice versa); or (ii) to use a "noisy" protocol that "hides" the important cross-section. The most technically-challenging part of our lower bound is to show that such a protocol necessarily reveals to Charlie (but not necessarily to the counter parties) a significant amount of information about the common labelling both parties received on "non-important" cross-sections.

1.1 Previous work

The notion of information complexity arose first in [BYCKO93], but gained much attention recently. Bar-Yossef et Al. [BYJKS04] used it as a method of lower bounding communication complexity. Barak et al. [BBCR13] showed several ways to compress communication protocols to a length dependent on their information complexity. Braverman and Rao [BR14] showed that the amortized communication complexity of a task is *equal* to its internal information complexity (cf. [Bra12] for a thorough survey). For product distributions, Kol [Kol15] recently showed that optimal communication complexity and information complexity are the same up to polynomial factors. Ganor, Kol and Raz [GKR14] were the first to show an exponential separation of the internal information and communication complexities, and later in [GKR15] of the external information and communication complexities. According to Braverman [Bra12] this is the largest gap possible between these measures. We use some of their constructions in our results.

2 Public-output protocols and output information complexity — definition and basic observations

In this section we define the new complexity measure of output information and state a few general comparisons between it and the more well-studied measures of information complexity.

2.1 Definition of output information complexity

We assume familiarity with general communication complexity definitions (cf. [KN97]). Basic definitions of entropy, mutual information and information complexity appear in Appendix A. Before defining the main measures studied in this paper below, we fix notation.

Notation We will use uppercase Π to denote a communication protocol, and lowercase π to denote a specific transcript of this protocol. X, Y represent the random variables of the inputs for Π. We denote by $h(p) = p \log p + (1-p)\log(1-p)$ the binary entropy function.

Definition 1. (Public output protocols solving a search problem) *Given a relation $R(X, Y, O)$, an input distribution μ on (X, Y) and a constant $\epsilon > 0$, we say that a protocol Π is a* public-output *protocol solving R with respect to μ, ϵ if there exists a function Z (called the* output *of Π) satisfying $\Pr[(x, y, Z(\Pi(x, y)) \in R)] > 1 - \epsilon$, where the probability is taken over the choice of x, y according to μ and the random coins of Π.*

Definition 2. (Output information complexity of a protocol) *Let $Z(\Pi)$ be the output of the 2-party public-output protocol Π. The output information complexity of Π, where the inputs are picked from a distribution μ is defined as $IC_\mu^{output}(\Pi) = I(X, Y; Z(\Pi))$.*

The definitions of internal and external information of relations are different here from the common ones, since we restrict our protocols to be *public-output* protocols, so that Charlie is able to learn the output at the end (for common definitions of information complexities see [Bra12]).

Definition 3. (Public-output internal and external information complexities) *The public-output internal information complexity of a relation R, under input distribution μ with error probability ϵ is defined as $IC_\mu^{\mathsf{po},int}(R, \epsilon) = \inf_\Pi IC_\mu^{int}(\Pi)$, where the infimum is taken over all public-output*

protocols solving R with error probability at most ϵ with respect to μ.
The public-output external information complexity $IC_{\mu}^{\mathsf{po},ext}(R,\epsilon)$ is defined in a similar way.

Definition 4. (Output information complexity of a relation) *The output information complexity of a relation R with error probability ϵ and input distribution μ is defined as $IC_{\mu}^{output}(R,\epsilon) = \inf_{\Pi} IC_{\mu}^{output}(\Pi)$, where the infimum is taken over all public-output protocols solving R with error probability at most ϵ with respect to μ.*

We omit the notation of μ and ϵ when they are clear from the context.

2.2 Public and private output protocols — discussion

When performing a communication task involving 2 parties, Alice and Bob, two essential things should be required for the task to succeed:

1. Alice and Bob agree on the same output.
2. The output is *correct* with respect to the task.

As defined above, a public-output protocol is one that has a unique output when seeing the transcript. This might seem necessary even if no third party is involved, since it prevents failing in the first requirement.
In the *deterministic* model (0-error, with full support on the input space), Bauer, Moran and Yehudayoff [BMY14] proved that this is indeed the case: they proved that every deterministic protocol has a public-output[1]. However, as we will show in last paragraph in Section 2.3, Alice and Bob can still agree on the output with probability 1 even if a protocol doesn't have a public-output, possibly with no communication at all.

2.3 Basic observations regarding output information complexity

Since this work initiates the study of output information and its relation to external and internal information, a few basic observations are needed; they appear in this section.
Since we required that at the end of the protocol we have an explicit output, then from the chain rule for mutual information (fact 3 in appendix A) we get that $IC_{\mu}^{output}(R,\epsilon) \le IC_{\mu}^{\mathsf{po},ext}(R,\epsilon)$. A known result from [Bra12] is that $IC_{\mu}^{\mathsf{po},int}(R) \le IC_{\mu}^{\mathsf{po},ext}(R)$. We wish to understand the various possibilities for differences between the measures, and present them below: internal information can be arbitrarily large even for constant output information; output information can be arbitrarily large even for zero internal information; and our main result, showing that both internal and output information can be exponentially smaller than external information.

Small output information, large internal and external information For every boolean function f, $IC^{output}(f) \le 1$, since one output bit cannot reveal more than one bit of information about the inputs. A random function f has $IC^{int}(f) = \Omega(n)$ with high probability [BW12], as does the inner product function [BGPW13], and hence the gap between internal information and output complexity can be maximally large. Any protocol that computes such an extremal f, must reveal to each party a significant amount of information about the input of the counter party, even though the output reveals at most a single bit.

[1] In their notations, a protocol can either compute a function *internally* or *externally*.

Small internal information, large output information Consider the following distribution μ: x is chosen uniformly at random, and y is picked to be equal to x. The function is $f(x,y)=x$. According to μ, when knowing her input, there is no more information for Alice to learn, and likewise for Bob. Hence, $IC_{\mu}^{\mathsf{po},int}=0$, no matter what the task is. However, here the output information is linear, since the output itself reveals the entire input.
Here, the public-output requirement is essential: if the only task was to compute the function in a way that Alice and Bob *know* the output, they could have done it without communicating at all.

3 Exponential separation between external information and output and internal information for public-output protocols

In this section we prove our main result, informally presented in Theorem 1. The search problem and input distribution are defined in Section 3.1, followed by a formal statement of the main theorem. Later sections contain the proof. The search problem and distribution, as well as the upper bounds in the Main Theorem 2 are essentially due to [GKR14]. The main contribution of this paper is the lower bound on external information of public-output protocols. Its proof appears in Section 4.

3.1 Definition of the "consistent leaf" search problem and distribution

Consistent leaf search problem The *consistent leaf* search problem with parameter $k \in \mathbb{N}$, denoted R_k, is defined as follows: Let $c = 2^{2^k}$. Consider the complete binary tree of depth $c \cdot k$, and denote V the set of vertices in this tree, and $L \subset V$ the set of leaves. We treat the tree as consisting of c *multi layers* (MLs) of depth k: vertices of depth d in the range $[k(j-1)+1, kj]$ belong to the j^{th} ML of the tree and are denoted V_j.
Each party gets a binary labeling of this tree: Alice gets $x : V \to \{0,1\}$ and Bob gets $y : V \to \{0,1\}$. We can look at the labeling as separate functions according to the multi layers of the tree: Alice gets c functions $x_j : V_j \to \{0,1\}$, and Bob gets c functions $y_j : V_j \to \{0,1\}$, where $j \in [c]$.
We define the function $t : (X,Y) \to [c]$ in this way: $t(x,y)$ is the first ML for which $x_j \neq y_j$. If $x = y$ and no such ML exists, then $t(x,y) = 1$.
We say that a leaf l is *consistent* with respect to the inputs and $j \in [c]$, if at least 0.8-fraction of the odd layers in the path from the root to it on multi layer j follows the function x_j, and likewise in the even layers and y_j.
The relation $R = R_k(X,Y,L)$ is defined as follows:

$$R = \{(x,y,l) : l \text{ is consistent with respect to } x, y, t(x,y)\}$$

Distribution on inputs The distribution $\mu = \mu_k$ used here is the one induced by the following sampling algorithm:

1. pick $x \in \{0,1\}^V$ uniformly at random.
2. pick $i \in [c]$ uniformly at random.
3. for every $j \neq i$, set $y_j = x_j$. $(\forall v \in V_j, y_j(v) = x_j(v))$
4. pick $y_i \in \{0,1\}^{V_i}$ uniformly at random.

At the end, we have two labeled trees, x and y, which are the same in all MLs but one, and in this particular one, denoted i, they have independent and random labeling. We refer to ML i as the *noisy* ML. Note that for this distribution, $t(x,y)$ would be i sampled in step 2 of the algorithm, except for the rare cases when $x = y$ (in which case $t(x,y) = 1$, regardless of i).

Remark 1. In the distribution defined in [GKR14], after the noisy ML i is set, the subtrees below it are chosen to be either equal or not in x, y according to the labeling on ML i. In our case, only vertices in ML i are not equal.

Statement of main theorem Following is the formal version of Theorem 1.

Theorem 2 (Main Theorem)**.** *The family of consistent leaf relations* $\{R_k\}_{k\in\mathbb{N}^+}$ *and distributions* μ_k *defined above satisfy*

1. **Small internal information:** $IC^{\mathsf{po},int}_{\mu_k}(R_k) = O(k)$
2. **Small output information:** $IC^{output}_{\mu_k}(R_k) = O(k)$
3. **Large external information:** $IC^{\mathsf{po},ext}_{\mu_k}(R_k) = 2^{\Omega(k)}$

We start in the next Section with the proof of the third bullet, as it is the main contribution of this paper. The first two upper bounds in the theorem are proved in Section 5; the proofs there are essentially due to [GKR14].

4 Lower bound on external information of public-output protocols — proof of Theorem 2, part 3

Fix an arbitrary protocol Π that solves $R = R_k$ with error probability ϵ. (We use Π to also denote the random variable representing the transcript of the protocol.) Recall that the sampling algorithm that induces $\mu = \mu_k$ starts by sampling $i \in [c]$. The proof of Part 3 of Theorem 2 is split into two cases, based on the amount of information that Π reveals about i. The first, and simpler, case is when Π does reveal much information about i.

Lemma 1 (Case I — Π is informative about i)**.** *If* $I(\Pi; i) \geq \frac{\log c}{2}$ *then* $IC^{ext}(\Pi) = 2^{\Omega(k)}$.

The idea for the first case is such: Let $W = \mathit{support}(\mu)$ be the input space. Except for the rare event in which $x = y$, W can be partitioned to c subsets W_i ($i \in [c]$) in which x and y differ only on the i^{th} ML. These subspaces have equal probabilities according to μ. Thus, knowing a significant amount of information about i provides us almost the same amount of information about the inputs.

The second and harder case is when Π is not very informative about i.

Lemma 2 (Case II — Π is not informative about i)**.** *If* $I(\Pi; i) < \frac{\log c}{2}$ *then* $IC^{ext}(\Pi) = 2^{\Omega(2^k)}$.

To prove this case we use the problem construction, and claim that if not much information is learned about i and yet the protocol is correct with high probability, then the leaf output of the protocol must be *consistent* with respect to many MLs (even though correctness requires only that it will be consistent on ML i). Thus, the output $z(\pi)$ must reveal at least a constant amount of information about the labeling of each of those MLs.

The proof of part 3 of Theorem 2 follows immediately from Lemmata 1 and 2. We next provide the proofs of both lemmas, starting with the second and more interesting one.

4.1 Proof of Lemma 2: The case of non-informative Π

To prove large output information under the assumption on non-informative Π, we shall analyze the distribution of X, Y after seeing the protocol, and it is easier to do so looking at a specific transcript π. (To simplify notation we sometimes use π to describe the event $\Pi = \pi$.) Since in this section we assume $I(\Pi; i) < \frac{\log c}{2}$, define the set G of *typical* transcripts of the protocol — the ones which reveal little information about i and are also correct with high probability. Formally:

$$G = \left\{ \pi : H(i|\Pi = \pi) > \frac{\log c}{3}, \Pr_{X,Y|\Pi=\pi}[Z(\pi) \text{ is correct}] > 1 - 10\epsilon \right\} \tag{1}$$

A "correct" output for input x, y is a value $z(\pi)$ such that $(x, y, z(\pi)) \in R_k$; when $(x, y, z(\pi)) \notin R_k$ we say the output is "incorrect", or "wrong". Our first claim is that G is "large" under the distribution μ (proof deferred to next section).

Claim 1 (Most transcripts are good)**.** $\Pr[\Pi \in G] > \frac{3}{20}$.

Denote by S_π the set of indices j for which the leaf output of a transcript π is consistent with respect to multi layer j with high probability. Intuitively, these are the MLs that are *likely candidates* for being the noisy ML selected by μ, conditioned on π. Formally:

$$S_\pi = \left\{ j : \Pr_{X,Y|\pi}[Z(\pi) \text{ is consistent with respect to } j, x, y] > \frac{1}{2} \right\} \tag{2}$$

We shall prove in the next section that (i) a "typical" transcript contains $\Omega(k)$ bits of information about each $j \in S_\pi$, and (ii) S_π is large for every "good" transcript $\pi \in G$. Formally,

Lemma 3 (Typical protocols are informative for the candidate noisy multi-layers)**.** *(i) For every $j \in S_\pi$, $H(X_j) - H(X_j|\Pi = \pi) = \Omega(k)$; consequently, (ii) $H(X) - H(X|\Pi = \pi) = \Omega(k \cdot |S_\pi|)$.*

Lemma 4 (S_π is large)**.** *For every $\pi \in G$, $|S_\pi| = 2^{\Omega(\log c)}$.*

We conclude with a proof of the output information lower bound for the non-informative case.

Proof of Lemma 2.

$$\begin{aligned}
IC^{ext}(\Pi) &= I(X, Y; \Pi) \geq I(X; \Pi) = && \text{(Chain rule)}\\
&= \mathbb{E}_\pi[H(X) - H(X|\pi)] = \\
&= \sum_\pi \Pr[\Pi = \pi] \cdot (H(X) - H(X|\pi)) \geq \\
&\geq \sum_{\pi \in G} \Pr[\Pi = \pi] \cdot (H(X) - H(X|\pi)) \\
&\geq \sum_{\pi \in G} \Pr[\Pi = \pi] \Omega(k \cdot |S_\pi|) \geq && \text{(Lemma 3.(ii))}\\
&\geq \sum_{\pi \in G} \Pr[\Pi = \pi] 2^{\Omega(\log c)} = && \text{(Lemma 4)}\\
&= 2^{\Omega(2^k)} \Pr[\Pi \in G] = \\
&= 2^{\Omega(2^k)} && \text{(Claim 1)}
\end{aligned}$$

□

4.1.1 Proofs of sub-lemmas

Here we give proofs of the three claims stated in the previous section and needed in the proof of Lemma 2, in the same order as stated there.

Proof of Claim 1. Recalling the protocol succeeds with probability $\geq 1-\epsilon$, apply Markov's inequality to deduce that the second condition of (1) holds with probability at least $\frac{9}{10}$. We now show that the first condition of (1) holds with probability p which is at least $\frac{1}{4}$. We first bound $\mathbb{E}_\pi H(i|\pi)$ from below,

$$\mathbb{E}_\pi H(i|\pi) = H(i|\Pi) = H(i) - I(\Pi; i) > \frac{\log c}{2}, \tag{3}$$

and then from above,

$$\begin{aligned}\mathbb{E}_\pi H(i|\pi) = &\sum_{\pi: H(i|\pi) > \frac{\log c}{3}} \Pr[\pi] \cdot H(i|\pi) + \sum_{\pi: H(i|\pi) \leq \frac{\log c}{3}} \Pr[\pi] \cdot H(i|\pi) \leq \\ &\leq p \log c + (1-p)\frac{\log c}{3}\end{aligned} \tag{4}$$

Combining equations (1) and (2) we get that $p \geq \frac{1}{4}$. We have seen that the second condition of (1) holds with probability at least 1/4 and the second does not hold with probability at most 1/10. Subtracting the two probabilities show that both conditions hold with probability at least 3/20, as claimed. □

Proof of Lemma 3. We start with (i). $H(X_j) = |V_j|$, since every ML of X is picked uniformly at random. With probability at least $\frac{1}{2}$, a 0.8-fraction of the $\frac{k}{2}$ odd nodes in the unique path from the root to the leaf $Z(\pi)$ are correct with respect to x. Hence we get that with probability at least $\frac{1}{2}$, $H(X_j|\Pi=\pi) \leq |V_j| - k + h(0.8)k$, and in any case $H(X_j|\Pi=\pi) \leq |V_j|$. Hence

$$H(X_j) - H(X_j|\Pi=\pi) \geq \frac{1}{4}(1-h(0.8)) \cdot k = \Omega(k)$$

This proves (i).

To prove (ii), denote $X = X_1X_2...X_c$, where X_j denote the random variable of the labeling of the j^{th} ML of X. Similarly let $X_{<j} = X_1X_2...X_{j-1}$.

$$\begin{aligned}
H(X) - H(X|\Pi=\pi) &= \sum_{j=1}^{c} [H(X_j|X_{<j}) - H(X_j|\Pi=\pi, X_{<j})] = && \text{(Chain rule)}\\
&= \sum_{j=1}^{c} [H(X_j) - H(X_j|\Pi=\pi, X_{<j})] \geq && \text{(independence of multi-layers)}\\
&\geq \sum_{j=1}^{c} [H(X_j) - H(X_j|\Pi=\pi)] \geq \\
&\geq \sum_{j\in S_\pi} [H(X_j) - H(X_j|\Pi=\pi)] = \\
&= \Omega(k \cdot |S_\pi|) && \text{(Lemma 3.(i))}
\end{aligned}$$

□

To prove Lemma 4 we require a two preliminary claims. The first shows that the probability that i (the multi-layer selected by μ) belongs to S_π with high probability.

Claim 2 (S_π has large measure)**.** *Let i be the multi layer picked in step 2 of the sampling algorithm of μ (section 3.1). For every $\pi \in G$, $\Pr[i \notin S_\pi|\pi] \leq 20\epsilon$.*

Proof. From the definition of G:

$$10\epsilon > \Pr_{x,y}[Z(\pi) \text{ is wrong}|\pi]$$

From the definition of S_π:

$$\Pr_{x,y}[Z(\pi) \text{ is wrong}|\pi] \geq \Pr_{x,y}[Z(\pi) \text{ is wrong}|\pi, i \notin S_\pi] \cdot \Pr[i \notin S_\pi|\pi] \geq \frac{1}{2}\Pr[i \notin S_\pi|\pi]$$

□

The second claim is of a rather general nature.

Claim 3. *Let J be a random variable taking values in $[c]$, and let $S \subseteq [c]$ be a set such that $\Pr[J \in S] > 1 - \delta$. Then $|S| \geq 2^{H(J)-\delta \log c - h(\delta)}$.*

Proof. Let T be an indicator for the event $j \in S$. By the chain rule we get that $H(J,T) = H(J) + H(T|J) = H(J)$ and also $H(J,T) = H(T) + H(J|T)$. Thus we have:

$$\begin{aligned} H(J) &= H(T) + H(J|T) \leq \\ &\leq h(\delta) + \delta H(J|T=0) + H(J|T=1) \leq \\ &\leq h(\delta) + \delta \log c + \log|S| \end{aligned}$$

□

Proof of Lemma 4. From the definition of G we get that $H(i|\pi) > \frac{\log c}{3}$, and from Claim 2 we know that $\Pr[i \notin S_\pi|\pi] \leq 20\epsilon$.
Using Claim 3 (with $J = i|\pi$, $S = S_\pi$, $\delta = 20\epsilon$ and ϵ small enough) we get the desired property. □

4.2 Proof of Lemma 1 — $I(\Pi; i) \geq \frac{\log c}{2}$

First, we wish to show that i from step 2 of the sampling algorithm *almost* induces a partition of the input space; namely, when seeing x, y we know i except for a very small probability.

Claim 4. *$H(i|X,Y) = o(1)$.*

Proof. For x and y to be equal, they need to be equal on the entire i^{th} multi layer. The first ML has 2^k vertices, so that the probability to be equal on all of them is 2^{-2^k}. For every other layer the probability is smaller, and thus $\Pr[x = y] = 2^{-\Omega(2^k)}$. When $x \neq y$, the noisy ML is decided uniquely when knowing the inputs, and thus its entropy is zero. Together we get:

$$\begin{aligned} H(i|X,Y) &= \mathbb{E}_{x,y} H(i|X=x, Y=y) = \\ &= \sum_{x,y:x\neq y} \Pr(x,y) H(i|X=x,Y=y) + \sum_{x,y:x=y} \Pr(x,y) H(i|X=x,Y=y) \leq \\ &\leq 0 + 2^{-\Omega(2^k)} \cdot \log c = o(1) \end{aligned}$$

□

Next, we wish to show that the information revealed about X, Y is at least as large as the information revealed about the identity of the noisy ML.

Claim 5. $I(\Pi; X, Y) \geq I(\Pi; i) - o(1)$

Proof. By definition $I(\Pi; X, Y) = H(X, Y) - H(X, Y|\Pi)$. Bounding the first term we get:

$$\begin{aligned} H(X,Y) &= H(i,X,Y) - H(i|X,Y) = && \text{(Chain rule)} \\ &= H(i) + H(X,Y|i) - H(i|X,Y) = && \text{(Chain rule)} \\ &= H(i) + H(X,Y|i) - o(1) && \text{(Claim 4)} \end{aligned}$$

Bounding the second:

$$\begin{aligned} H(X,Y|\Pi) &\leq H(i,X,Y|\Pi) = && \text{(Chain rule)} \\ &= H(i|\Pi) + H(X,Y|i,\Pi) \leq && \text{(Chain rule)} \\ &\leq H(i|\Pi) + H(X,Y|i) \end{aligned}$$

Subtracting the two we get: $I(\Pi; X, Y) \geq H(i) - o(1) - H(i|\Pi) = I(\Pi; i) - o(1)$, and this completes the proof. □

proof of Lemma 1. By claim 5 we get: $IC^{ext}(\Pi) = I(\Pi; X, Y) \geq I(\Pi; i) - o(1)$, and by the assumption $I(\Pi; i) \geq \frac{\log c}{2}$ we get $IC^{ext}(\Pi) = \frac{\log c}{2} - o(1) = 2^{\Omega(k)}$. □

5 Upper bounds on internal information and output information

In this section we prove the first two parts of theorem 2. The first part follows essentially from the work of Ganor, Kol and Raz [GKR14]; The second is rather simple, and follows from the multiple possibilities for a correct answer for every input pair.

5.1 Upper bound on internal information of public-output protocols — proof of Theorem 2, part 1

To prove that $IC_\mu^{\mathsf{po},int}(R) = O(k)$, we will show how Alice and Bob can communicate while revealing little information to each other. This method was first shown in [GKR14]. The idea of the protocol is to travel from the root down the tree, following Alice's labels on the odd depth vertices, and Bob's on the even depth vertices. However, Doing this deterministically would reveal the noisy multi layer i, which reveals 2^k bits of information about the other party's input ($i \in [c]$, and $\log c = 2^k$). Thus, we want to somehow *disguise* the noisy multi layer, by adding noise to the communication. Formally, Π_1 is defined as follows: set the current vertex to be the root. At each odd step, Alice transmits one bit to determine which child vertex of the current one we go to. This bit is chosen to be the one according to her labeling with probability 0.9, and the *wrong* one with probability 0.1. Bob does the same for the even layers. This is done until we reach a leaf. This protocol has a public-output, since its transcript defines a unique leaf.

Claim 6. $\Pr[\Pi_1 \textit{ gives a wrong output}] \leq 2^{-\Omega(k)}$.

Proof. Note that the definition of a consistent leaf depends only on the layer i, so correctness is independent of what the parties have communicated apart from it. Since we require 0.8-consistency with the inputs along i, and the parties follow their labeling with probability 0.9 at each step, the probability that they fail in achieving 0.8-consistency is exponentially small in k (the depth of the multi layer), according to the Chernoff bound. □

For the proof of the low information of this protocol we will follow several steps from [GKR14]. The intuition behind it is the following: Although the communication is double exponential in k, in almost all multi layers (all but the noisy one), Bob can *predict* the distribution of Alice's messages by guessing that she has the same label as he has in the corresponding vertices. This prediction would fail in only $O(k)$ vertices along the path: the ones that are in the path along the i^{th} ML.
Formally, for every vertex v in the protocol tree, we define $P_v = (p_v, 1 - p_v)$ as the probability distribution of the next bit to be sent in the protocol, conditioned on reaching v and the input of the player who owns v. This distribution is well defined from the protocol's definition: if Alice owns v and $x_v = 1$, then $P_v = (0.1, 0.9)$. If $x_v = 0$ then $P_v = (0.9, 0.1)$. We think of $Q_v = (q_v, 1 - q_v)$ as an *estimation* of this distribution, which in general is every distribution on the space $\{0, 1\}$ which doesn't depend on the input of v's owner when given the other party's input. Let T be the protocol tree, rooted by the vertex r, and denote T_v as the subtree of T rooted by v.

Definition 5. *(divergence cost of the protocol tree,* [BR14]*). Let Q be some estimation distribution for T. The divergence of T according to Q is defined as:*

$$D_Q(T) = D(P_r||Q_r) + \mathbb{E}_{v\sim P_r} D_Q(T_v)$$

Where $D(T) = 0$ if the tree is empty.

Since in Π_1 every bit sent in the protocol depends on a specific labeled vertex in the input tree, we can identify each vertex in the protocol tree with the corresponding vertex in the input tree.
We define Q to be the following: if the labeling of v in the tree of the party who *doesn't* own v is 1, then $Q_v = (0.1, 0.9)$. Otherwise $Q_v = (0.9, 0.1)$. Note that this estimation is *independent* of the owner's input for all vertices, given the other party's input.

Lemma 5. *$D_Q(T) = O(k)$.*

Proof. The recursive definition of $D_Q(T)$ can be written as $D_Q(T) = \sum_v [\tilde{p}_v \cdot D(P_v||Q_v)]$, where $\tilde{p}_v$ is the overall probability to reach the vertex v in the protocol.
Since for every vertex in any ML apart from i, $P_v = Q_v$, from Fact 2 we know they contribute nothing to this sum.
For every vertex v *in* ML i, $D(P_v||Q_v) \leq 0.9 \log 9 - 0.1 \log 9 < 3$. For each layer in this ML, the total of $\tilde{p}_v$ in this layer is 1, since this sum is the probability to reach *any* vertex in this depth. The number of layers in every ML is k, and hence $D_Q(T) \leq 3k$. □

Let Q^* be the estimation of the party *receiving* the messages, defined as the expectation over distributions according to the inputs of the speaking party, these being sampled conditioned on the input of the other party. Intuitively, this should be the best estimation possible, as was proved by Kol at al. in [GKR14]:

Lemma 6. *For every estimation Q, $D_{Q^*}(T) \leq D_Q(T)$.*

In [BR14], Braverman and Rao prove the relation between the divergence cost and the internal information[2]:

Lemma 7. $IC_{\mu}^{int}(\Pi) = \mathbb{E}_{x,y\sim\mu} D_{Q^*}(T)$.

proof of part 1 of Theorem 2. Using Lemmas 7 and 6, 5 respectively:

$$IC_{\mu}^{int}(\Pi) = \mathbb{E}_{x,y\sim\mu} D_{Q^*}(T) \leq \mathbb{E}_{x,y\sim\mu} D_{Q}(T) = \mathbb{E}_{x,y\sim\mu} O(k) = O(k).$$

□

5.2 Upper bound on output information of public-output protocols — proof of Theorem 2, part 2

To prove that $IC_{\mu}^{output}(R) = O(k)$, Consider the following protocol Π_2: Alice sends Bob her entire input x. Bob now recognizes which multi layer is the noisy one, calculates the set of consistent leaves, and outputs a uniformly chosen one of them. By definition, this protocol always succeeds. For calculating the output information complexity, consider this: The probability of a leaf to be consistent with the inputs is $2^{-\Theta(k)}$, since the consistency is determined only by the path in the noisy multi layer. Before seeing the output, the random variable l representing the output leaf is uniform over L, and after seeing the input it is uniform on a subset of fraction $2^{-\Theta(k)}$ of it. We get:

$$\begin{aligned} IC_{\mu}^{output}(\Pi_2) = I(X,Y;l) = H(l) - H(l|X,Y) = \\ = \log|L| - \log(2^{-\Theta(k)} \cdot |L|) = \\ = \log|L| - \log|L| + \Theta(k) = \Theta(k) \end{aligned}$$

This proves part 2, and together with the proofs of the other two parts in previous sections, the proof of Theorem 2 is complete.

6 Open problem — separating internal and external information for general protocols

Open Question 1. *Is there a family of relations* $\{R_k\}_{k\in\mathbb{N}}$ *and distributions* $\{\mu_k\}_{k\in\mathbb{N}}$ *having* $IC_{\mu_k}^{int}(R_k) = I$ *and* $IC_{\mu_k}^{ext}(R_k) = 2^{\Omega(I)}$*?*

We conjecture that the answer to this question is *yes*, using the exact relation and distribution defined in section 3.1. However, the techniques reported here do not prove this; in particular the crucial variable $Z(\pi)$ defined in section 4.1 would not be well defined if one drops the requirement of reporting an explicit output. Notice that an affirmative answer to the question above generalized and implies the main Theorem 2.

One example of a non public-output protocol for the consistent leaf relation is the following: Assume that the noisy ML is not in the first $d = 2^k$ MLs (this happens with probability $1 - \frac{2^k}{2^{2^k}} = 1 - o(1)$, and in that case the protocol below will likely err.). In this case, the parties have the same labels on the first $2^{\Theta(2^k)}$ vertices of the tree. These labels can be used as a *one time pad* for the rest of the communication: Alice and Bob start at the left most vertex in layer $d \cdot k + 1$, and proceed like

[2] For more details see [BR14, section 5.2].

in protocol Π_1 described in section 5.1, but XOR every bit with the next bit in this one time pad. Charlie, when listening to the communication, gains no information about what the output is, but Alice and Bob agree with high probability.
Even in this case, Charlie learns a double exponential amount of information: in every bit communicated, he learns (with high probability) the XOR of 2 input bits, which is roughly one bit of information. However, it is still open to find a lower bound proof that will hold for *any* non public-output protocol.

Acknowledgements

Our preliminary writeup contained a weaker version of Claim 3 (that was sufficient to prove our main result). We thank Eldar Fischer for generalizing that version to the one appearing here. We also thank Yadu Vasudev and Gregory Schwartzman for useful discussions.
The research leading to these results has received funding from: the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement number 240258 and the Israeli Science Foundation (grant 1501/14).

References

[ACC+14] Anil Ada, Arkadev Chattopadhyay, Stephen A. Cook, Lila Fontes, Michal Koucký, and Toniann Pitassi. The hardness of being private. *ACM Trans. Comput. Theory*, 6(1):1:1–1:24, March 2014.

[BBCR13] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. How to compress interactive communication. *SIAM Journal on Computing*, 42(3):1327–1363, 2013.

[BGPW13] Mark Braverman, Ankit Garg, Denis Pankratov, and Omri Weinstein. Information lower bounds via self-reducibility. In *Computer Science–Theory and Applications*, pages 183–194. Springer, 2013.

[BMY14] Balthazar Bauer, Shay Moran, and Amir Yehudayoff. Internal compression of protocols to entropy. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 21, page 101, 2014.

[BOGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, pages 1–10, New York, NY, USA, 1988. ACM.

[BR14] Mark Braverman and Akhila Rao. Information equals amortized communication. *Information Theory, IEEE Transactions on*, 60(10):6058–6069, 2014.

[Bra12] Mark Braverman. Interactive information complexity. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 505–524. ACM, 2012.

[BW12] Mark Braverman and Omri Weinstein. A discrepancy lower bound for information complexity. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 459–470. Springer, 2012.

[BYCKO93] Reuven Bar-Yehuda, Benny Chor, Eyal Kushilevitz, and Alon Orlitsky. Privacy, additional information and communication. *Information Theory, IEEE Transactions on*, 39(6):1930–1943, 1993.

[BYJKS04] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *J. Comput. Syst. Sci.*, 68(4):702–732, 2004.

[GKR14] Anat Ganor, Gillat Kol, and Ran Raz. Exponential separation of information and communication. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 176–185. IEEE, 2014.

[GKR15] Anat Ganor, Gillat Kol, and Ran Raz. Exponential separation of communication and external information. *Electronic Colloquium on Computational Complexity (ECCC)*, 22:88, 2015.

[KN97] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, New York, NY, USA, 1997.

[Kol15] Gillat Kol. Interactive compression for product distributions. *Electronic Colloquium on Computational Complexity (ECCC)*, 2015.

[Kus92] Eyal Kushelvitz. Privacy and communication complexity. *SIAM Journal on Discrete Mathematics*, 5(2):273–284, 1992.

A Information theory preliminaries

These definitions are included here for the paper to be self contained.

A.1 Information theory

Definition 6. *(Entropy). Let X be a discrete random variable, and denote $p_x = \Pr[X = x]$. The entropy of X is $H(X) = \sum_x p_x \log \frac{1}{p_x}$.*
The conditional entropy of X given Y, denoted $H(X|Y)$ is defined to be $\mathbb{E}_y H(X|Y = y)$.

Fact 1. *$H(A|B) \geq H(A|BC)$.*

Definition 7. *(Kullback-Leibler divergence). The KL divergence between two distributions P, Q is $D(P||Q) = \sum_x P(x) \log \frac{P(x)}{Q(x)}$.*

Fact 2. *$D(P||P) = 0$*

Definition 8. *(Mutual information). The mutual information of two random variables A, B is $I(A;B) = H(A) - H(A|B) = H(B) - H(B|A)$.*

Fact 3. *(Chain rules for entropy and mutual information).*
$H(A, B) = H(A) + H(B|A)$.
$I(A_1, A_2; B) = I(A_1; B) + I(A_2; B|A_1)$.

A.2 Information and communication complexuty

Definition 9. *The internal information complexity of a 2-party protocol Π, where the inputs are picked from a distribution μ is defined as $IC_\mu^{int}(\Pi) = I(X; \Pi|Y) + I(Y; \Pi|X)$.*

Definition 10. *The external information complexity of a 2-party protocol Π, where the inputs are picked from a distribution μ is defined as $IC_\mu^{ext}(\Pi) = I(X, Y; \Pi)$.*

Definition 11. *The internal information complexity of a relation R, under input distribution μ with error probability ϵ is defined as $IC_\mu^{int}(R, \epsilon) = \inf_\Pi IC_\mu^{int}(\Pi)$, where the infimum is taken over all protocols solving R with error probability at most ϵ with respect to μ.*
The external information complexity $IC_\mu^{ext}(R, \epsilon)$ is defined in a similar way.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de