The Hardness of Being Private

ANIL ADA, McGill University
ARKADEV CHATTOPADHYAY, University of Toronto
STEPHEN A. COOK, University of Toronto
LILA FONTES, University of Toronto
MICHAL KOUCKÝ, Institute of Mathematics, Academy of Sciences, Prague
TONIANN PITASSI, University of Toronto

Kushilevitz [1989] initiated the study of information-theoretic privacy within the context of communication complexity. Unfortunately, it has been shown that most interesting functions are not privately computable [Kushilevitz 1989, Brandt and Sandholm 2008]. The unattainability of perfect privacy for many functions motivated the study of *approximate privacy*. Feigenbaum et al. [2010a, 2010b] define notions of worst-case as well as average-case approximate privacy, and present several interesting upper bounds, and some open problems for further study. In this paper, we obtain asymptotically tight bounds on the tradeoffs between both the worst-case and average-case approximate privacy of protocols and their communication cost for Vickrey-auctions.

Further, we relate the notion of average-case approximate privacy to other measures based on information cost of protocols. This enables us to prove exponential lower bounds on the subjective approximate privacy of protocols for computing the Intersection function, independent of its communication cost. This proves a conjecture of Feigenbaum et al. [2010a].

Categories and Subject Descriptors: F.2.3 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity—*Tradeoffs between Complexity Measures*

General Terms: Algorithms, Security, Theory

Additional Key Words and Phrases: approximate privacy, communication complexity, information theory, privacy, privacy tradeoff, Vickrey auction

ACM Reference Format:
Anil Ada, Arkadev Chattopadhyay, Stephen A Cook, Lila Fontes, Michal Koucký, and Toniann Pitassi, YYYY. The Hardness of Being Private. *ACM Trans. Comput. Theory* V, N, Article A (January YYYY), 23 pages.
DOI: http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

Privacy in a distributed setting is an increasingly important problem. A key application is the setting of combinatorial auctions where many agents have private information (e.g., their preferences) but would like to compute a function of their inputs without revealing any of their private information. There is a large body of research examining which functions can be computed securely, and how. Many of these results rely

This work supported by NSERC and a postdoctoral fellowship of Ontario Ministry of Research and Innovation. Partially supported by GA ČR P202/10/0854, grant IAA100190902 of GA AV ČR, by the Center of Excellence CE-ITI under the grant P202/12/G061 of GA ČR and RVO: 67985840.
Author's addresses: A. Ada, Department of Computer Science, McGill University, aada@cs.mcgill.ca; A. Chattopadhyay, School of Technology and Computer Science, Tata Institute of Fundamental Research, arkadev@cs.toronto.edu; Stephen A. Cook, Lila Fontes, and Toni Pitassi, Department of Computer Science, University of Toronto, {sacook,fontes,toni} @cs.toronto.edu; Michal Koucký, Institute of Mathematics, Academy of Sciences, Prague, koucky@math.cas.cz.

DOI: http://dx.doi.org/10.1145/0000000.0000000

on an assumption, such as a computational complexity assumption, or the assumption that more than some fixed fraction m of the players are trustworthy, or the assumption that the auctioneer (a 3^{rd} party) is trustworthy. As Brandt and Sandholm [2008] point out, privacy which is based on an assumption of hardness can become outdated as computers become faster and more powerful; security parameters (like key length) need to be continuously updated to cope with increasing computational power. Hence, ideally one would like privacy based on stronger assumptions. Auctions are a natural setting where we would doubt the trustworthiness of fellow participants or an auctioneer. We nevertheless would like to compute on the internet. In this work, we focus on situations where each player is deterministic and honest but curious. Honest, because they obey the rules of the game. Curious, as they do not miss any opportunity to gain knowledge about others' input. (Honesty, while a large assumption, will be dependable when we consider functions which are *truthful* — self-interested players will obey the protocol.)

Kushilevitz [1989] initiated the study of *information-theoretic* privacy in communication complexity, which is an appealing direction because it does not rely on computational assumptions discussed above. Informally, a multi-player communication protocol for computing a function $f(x_1, x_2, \ldots, x_k)$ is private if each player does not learn any additional information (in an information theoretic sense) beyond what follows from knowing his/her private input, and the function value $f(\vec{x})$. (A similar notion of privacy considers limiting an *eavesdropper* to learning only the function value $f(\vec{x})$ from the protocol, and nothing more.) A complete characterization of the privately computable functions was given, but unfortunately, early work ruled out private protocols for most interesting functions [Kushilevitz 1989; Brandt and Sandholm 2008]. For example, private second-price auctions are not possible with more than two participants,[1] useful and are extremely inefficient even in the setting of two bidders [Chor and Kushilevitz 1989; Brandt and Sandholm 2008].

The unattainability of perfect privacy for many functions motivated the study of *approximate* privacy. Most relevant to our work is the study of Klauck [2002] and the more recent work of Feigenbaum et al. [2010a]. The relaxation from perfect to approximate privacy is appealing because it renders more functions computable privately, and more closely mirrors real-world situations in which *some* privacy loss may be acceptable. On the other hand, it is more subtle to capture the notion of approximate privacy. While most reasonable definitions of perfect privacy turn out to be equivalent, this is not quite the case with approximate privacy. In particular, the measures of Klauck [2002] and Feigenbaum et al. [2010a] are different and each has its own advantage and characteristics. Our work here is primarily motivated by recent work of Feigenbaum et al. [2010a] and Comi et al. [2011]. A second motivation is to understand the connections between the two measures.

In the two player setting, let $f(x, y)$ be a function, and let P be a two-player deterministic communication protocol for f. The privacy loss (or privacy approximation ratio, PAR) on the input (x, y) with respect to P is defined to be the size of the monochromatic region containing (x, y) divided by the size of the protocol-induced rectangle containing (x, y): $\mathrm{PAR}(x, y) = \frac{|f^{-1}(f(x,y))|}{|P(x,y)|}$. The worst-case privacy loss of protocol P is the maximum privacy loss over all inputs (x, y), and the worst-case privacy loss of the function f is then the minimum privacy loss over all protocols for f. Perfect privacy of a protocol (as defined in 1989) requires that the privacy approximation ratio (PAR) is 1 for all inputs. (This definition is easily extended to the multi-player setting.)

[1] Even assuming that all players are honest-but-curious, secure multiparty computation is only applicable if $< 1/2$ the parties form a coalition [Ben-Or et al. 1988].

Under this relaxed notion of privacy, things are much more interesting [Feigenbaum et al. 2010a; 2010b; Comi et al. 2011]. For example, Feigenbaum et al. [2010a; 2010b] study the Vickrey auction problem, and reveal a possible *inherent* tradeoff between privacy and communication complexity: they describe a family of protocols such that the privacy loss approaches 1 (perfect privacy) as the length of the protocol approaches exponential. They also study several prominent boolean functions with respect to approximate privacy.

Feigenbaum et al. consider an *average-case* notion of approximate privacy as well. In this setting, we are interested in the average privacy loss over a distribution on inputs. Here they describe a protocol for Vickrey auction that achieves exponentially smaller average-case PAR than its worst-case PAR. A similar protocol was described by Klauck [2002].

Our Contributions

In this paper, we present several new lower bounds on the communication cost for achieving privacy and establish relationships between approximate privacy and several other known measures.

First, we prove that there is an inherent tradeoff between privacy and communication complexity, by proving a privacy/communication complexity tradeoff lower bound for the Vickrey auction problem. This shows that the upper bounds presented by Feigenbaum et al. [2010a] are essentially tight. Feigenbaum et al. [2010a] provided a lower bound only for the special case of bisection-type protocols.

THEOREM 1.1. *For all n, for all p, $2 \leq p \leq n/4$, any deterministic protocol for the two-player n-bit Vickrey auction problem with communication cost (length) less than $n2^{\frac{n}{4p}-5}$ obtains privacy loss (worst-case PAR) at least 2^{p-2}.*

This lower bound is technically interesting as it deals with super-polynomial communication protocols. The usual communication complexity techniques aim at protocols that are at most *linear* in their input size.

Our second contribution demonstrates a similar type of tradeoff for the case of average-case approximate privacy. We prove an asymptotically tight lower bound on the average-case approximate privacy of the Vickrey auction problem, showing that the upper bounds from Feigenbaum et al. [2010a] are essentially tight. This generalizes the result of Comi et al. [2011] for Vickrey auctions. Again, Feigenbaum et al. [2010a] provided lower bounds only for the special case of bisection-type protocols. As a side note, we positively resolve an open question from Feigenbaum et al. [2010a] concerning arbitrary input distributions (Proposition 4.3).

THEOREM 1.2. *For all $n, r \geq 1$, any deterministic protocol for the two-player n-bit Vickrey auction problem (over the uniform distribution of inputs) with communication cost (length) less than r obtains average-case PAR^2 at least $\Omega(\frac{n}{\log(r/n)})$.*

Our lower bounds show that the approximate privacy of any polynomial length protocol is still as large as $\Omega(n/(\log n))$. Indeed, such superlinear protocols have been devised by Klauck [2002], who proved upper bounds for his measure of approximate-privacy. To the best of our knowledge, Theorem 1.2 provides the first (tight) lower bounds on the communication cost of achieving good approximate privacy for Vickrey auctions. The proof of the theorem relates the loss of privacy to a certain Ball Partition Problem that may be of independent interest.

Furthermore, we modify the average-case privacy approximation measure of Feigenbaum et al. Our modification provides a rather natural measure that was disregarded

[2]Under the original definition [Feigenbaum et al. 2010a] or our alternate Definition 4.1.

in Feigenbaum et al. [2010a], but coincides with that of Feigenbaum et al. in the case of uniform distribution on the inputs. Our modified measure has several advantages. It allows natural alternative characterizations, and it can be directly related to the (information-theoretic) privacy measure of Klauck.[3] We can quantitatively connect Klauck's privacy measure to well studied notions of *(internal) information cost* in communication complexity. This allows us to prove a new lower bound on the average-case subjective privacy approximation measure of Feigenbaum et al. [2010a], and answers affirmatively a conjecture from their paper.

THEOREM 1.3. *For all $n \geq 1$, and any protocol P computing the Set Intersection INTERSEC$_n$ the average-case subjective PAR is exponential in n under the uniform distribution:*

$$\mathrm{avg}_U \, \mathrm{PAR}^{sub}(P) = 2^{\Omega(n)}$$

We contend that any of the mentioned measures could serve as a reasonable measure of privacy. Indeed, each of the measures seems to exhibit advantages over the other ones in some scenario, so each of the measures captures certain aspect of privacy. For example, the English auction protocol for Vickrey auction achieves perfect privacy (under any measure) but at exponential communication cost. On the other hand, the Bisection protocol achieves linear average-case PAR with merely linear communication cost. However, the difference between these two protocols is not reflected well in Klauck's privacy measure, where both protocols lose constant number of bits on the average. (And in general it is not hard to come up with examples which are "distinguishable" — very far apart — using PAR, but the same big-O using Klauck's measure.)

Outline of Paper

In Section 2, we provide our basic notation and background on information theory. In Section 3, we review the notion of privacy approximation ratio, and in Section 3.1, we review the Vickrey auction problem. In Section 3.2, we present our lower bound tradeoff for worst-case privacy of Vickrey auctions. In Section 4, we present our lower bound on average-case PAR for Vickrey auctions, and discuss the relationship between average-case PAR and information cost, deriving several new results from this relationship. In Section 5, we summarize our comparisons and discuss different privacy measures.

2. PRELIMINARIES

In this section, we review our basic notations and concepts. For a positive integer k, we let $[k] = \{1, 2, \ldots, k\}$. We assume that the reader is familiar with communication complexity (see Kushilevitz and Nisan [1997] for more background.) We will use the following notation. Given $f : X \times Y \to Z$, each input (x, y) is associated with the **region** $R_{x,y}$ of all inputs in the preimage of $f(x, y)$, i.e.,

$$R_{x,y} = \{(x', y') \in X \times Y | f(x, y) = f(x', y')\}.$$

For any value $z \in Z$ we let $R_z = f^{-1}(z)$ be the preimage of z. The set of all regions of function f is $\mathcal{R}(f) = \{R_{x,y} : (x, y) \in X \times Y\}$. Let P be a communication protocol for the function f. For inputs $(x, y) \in X \times Y$ we let $\Pi_P(x, y)$ denote the transcript of the protocol on input x given to Alice and y given to Bob. We associate the input (x, y) with the **protocol-induced rectangle** $P_{x,y}$ of all inputs which yield the same transcript:

$$P_{x,y} = \{(x', y') \in X \times Y : \Pi_P(x, y) = \Pi_P(x', y')\}.$$

[3]Theorem 22 shows that Klauck's measure provides a lower bound for average-case PAR. This bound is not tight: upper bounds on Klauck's measure do not necessarily upper-bound PAR.

Note that $P_{x,y} \subseteq R_{x,y}$ as we assume that P correctly computes f.

2.1. Information theoretic notions

Information theory provides a highly intuitive and powerful calculus to reason about random variables. We need the following basic notions from this theory whose proofs can be found in any standard textbook on the subject (see for example Cover and Thomas [1991]).

For any random variable $\mathbf{X}$, we denote its probability distribution over its range $\mathcal{X}$ by μ_X. The entropy of $\mathbf{X}$, denoted by $H(\mathbf{X})$, is defined as follows:

$$\begin{aligned} H(\mathbf{X}) \;&=\; -\sum_{x\in\mathcal{X}} \Pr_{\mu_X}\left[\mathbf{X}=x\right]\log_2\left(\Pr_{\mu_X}\left[\mathbf{X}=x\right]\right) \\ &=\; -\mathbb{E}_{\mu_X}\big[\log(\mu_X(x))\big] \end{aligned}$$

Let $\mathbf{Y}$ be another random variable. For any y in the range of $\mathbf{Y}$, $H(\mathbf{X}|\mathbf{Y}=y)$ is defined as just the entropy of $\mathbf{X}$ under the conditional distribution, i.e.

$$H(\mathbf{X}|\mathbf{Y}=y) \equiv_{\text{def}} -\sum_{x\in\mathcal{X}} \Pr\left[\mathbf{X}=x\,|\,\mathbf{Y}=y\right]\log\left(\Pr\left[\mathbf{X}=x\,|\,\mathbf{Y}=y\right]\right).$$

Extending the above naturally, we define the notion of conditional entropy $H(\mathbf{X}|\mathbf{Y})$ as

$$H(\mathbf{X}|\mathbf{Y}) \equiv_{\text{def}} \mathbb{E}_{\mu_Y}\big[H(\mathbf{X}|\mathbf{Y}=y)\big].$$

As intuition suggests, conditioning a random variable $\mathbf{X}$ on another random variable $\mathbf{Y}$ cannot increase its uncertainty on the average. Formally,

FACT 1. *For any two random variables* $\mathbf{X}$ *and* $\mathbf{Y}$, $H\big(\mathbf{X}|\mathbf{Y}\big) \;\leq\; H\big(\mathbf{X}\big)$.

The mutual information between $\mathbf{X}$ and $\mathbf{Y}$, denoted by $I(\mathbf{X}\;:\;\mathbf{Y})$, is defined as $H(\mathbf{X})-H(\mathbf{X}|\mathbf{Y})$. It is straightforward to verify that mutual information is a symmetric quantity, i.e. $I(\mathbf{X}\;:\;\mathbf{Y}) = H(\mathbf{X}) - H(\mathbf{X}|\mathbf{Y}) = H(\mathbf{Y}) - H(\mathbf{Y}|\mathbf{X}) = I(\mathbf{Y}\;:\;\mathbf{X})$. Fact 1 implies that mutual information between two random variables is always non-negative. Just like entropy, one can define the conditional mutual information between random variables: let $\mathbf{Z}$ be another random variable with range $\mathcal{Z}$.

$$\begin{aligned} I(\mathbf{X}\;:\;\mathbf{Y}\,|\,\mathbf{Z}) \;&=\; H(\mathbf{X}\,|\,\mathbf{Z}) \;-\; H(\mathbf{X}\,|\,\mathbf{Y},\mathbf{Z}) \\ &=\; \mathbb{E}_{\mu_Z}\big[I(\mathbf{X}\;:\;\mathbf{Y}\,|\,\mathbf{Z}=z\big]. \end{aligned}$$

We will also need the following simple claim:

CLAIM 2. *Let* $\mathbf{X},\mathbf{Y},\mathbf{Z},\mathbf{W}$ *be any random variables. Then,* $\big|I(\mathbf{X}\;:\;\mathbf{Y}|\mathbf{W}) - I(\mathbf{X}\;:\;\mathbf{Y}\,|\mathbf{W},\mathbf{Z})\big| \leq H(\mathbf{Z})$.

PROOF. First, notice that

$$I(\mathbf{X}\;:\;\mathbf{Y}\,|\,\mathbf{W}) - I(\mathbf{X}:\mathbf{Y}\,|\,\mathbf{W},\mathbf{Z}) = \big(H(\mathbf{X}|\mathbf{W}) - H(\mathbf{X}|\mathbf{W},\mathbf{Z})\big) - \big(H(\mathbf{X}|\mathbf{W},\mathbf{Y}) - H(\mathbf{X}|\mathbf{Y},\mathbf{W},\mathbf{Z})\big).$$

The first quantity in brackets above is

$$\begin{aligned} \big(H(\mathbf{X}|\mathbf{W}) - H(\mathbf{X}|\mathbf{W},\mathbf{Z})\big) \;&=\; I(\mathbf{X}:\mathbf{Z}\,|\,\mathbf{W}) \\ &=\; H(\mathbf{Z}|\mathbf{W}) - H(\mathbf{Z}|\mathbf{W},\mathbf{X}) \\ &\leq\; H(\mathbf{Z}|\mathbf{W}) \\ &\leq\; H(\mathbf{Z}) \end{aligned}$$

The first equality uses the symmetry of information.

The second bracketed quantity can be likewise re-written using the symmetry of information:

$$\begin{aligned}\big(H(\mathbf{X}|\mathbf{W},\mathbf{Y}) - H(\mathbf{X}|\mathbf{Y},\mathbf{W},\mathbf{Z})\big) &= \mathbb{E}_{\mu_Y}\big[H(\mathbf{X}|\mathbf{W},\mathbf{Y}=y) - H(\mathbf{X}|\mathbf{Z},\mathbf{W},\mathbf{Y}=y)\big] \\ &= \mathbb{E}_{\mu_Y}\big[I(\mathbf{Z}:\mathbf{X}\,|\,\mathbf{W},\mathbf{Y}=y)\big] \leq H(\mathbf{Z})\end{aligned}$$

Combining the two, we are done. ■

3. WORST-CASE PRIVACY APPROXIMATION RATIO

In this paper, we are concerned with privacy-preserving communication complexity. A perfectly private communication protocol for f will reveal only the output of f and no additional information. Every two inputs (x, y) and (x', y') such that $f(x, y) = f(x', y')$ should be indistinguishable from each other [Kushilevitz 1989; Chor et al. 1994]. Approximate privacy provides a measure of how much indistinguishability has been lost. These notions are formalized as follows.

The following definition captures the privacy loss of a communication protocol with respect to a third party observer (eavesdropper) who overhears the messages sent between the players. This measure is referred to as **objective**.

Definition 3.1. [Feigenbaum et al. 2010a] A protocol P for a function f on $X \times Y$ has **worst-case objective privacy approximation ratio** (PAR) defined by

$$\mathrm{PAR}(P) = \max_{(x,y)} \frac{|R_{x,y}|}{|P_{x,y}|} = \max_{(x,y)} \mathrm{PAR}(P, x, y).$$

Each input (x, y) has its own privacy approximation ratio $\mathrm{PAR}(P, x, y) = \frac{|R_{x,y}|}{|P_{x,y}|}$. Often we do not specify the protocol P when it is clear from context.

The PAR measure of privacy can be extended to **subjective PAR**, which measures the privacy that the players lose *to each other*.

Definition 3.2. [Feigenbaum et al. 2010a] A protocol P for a function f on $X \times Y$ has **worst-case subjective privacy approximation ratio** (PAR^{sub}) defined by:

$$\mathrm{PAR}^{sub}(P) = \max\left\{ \max_{(x,y)} \frac{|R_{x,y} \cap X \times \{y\}|}{|P_{x,y} \cap X \times \{y\}|}, \max_{(x,y)} \frac{|R_{x,y} \cap \{x\} \times Y|}{|P_{x,y} \cap \{x\} \times Y|} \right\}.$$

Previous work by Kushilevitz [1989] gave a combinatorial characterization of the functions f which are computable with perfect privacy $\mathrm{PAR} = 1$. This set unfortunately excludes most auctions [Brandt and Sandholm 2008], as well as many basic functions of interest in theoretical computer science, e.g., greater than [Yao 1982] and set intersection and disjointness [Feigenbaum et al. 2010b].

As many functions are not computable with perfect privacy, it is natural to investigate the following general question for a function f: is f privately computable, and how much communication is necessary to achieve PAR less than some number c? In the next section, we focus on the case of Vickrey auctions which is one of the most studied functions in this context.

3.1. Vickrey auctions

Vickrey auctions (also known known as 2^{nd}-price auctions) arise in mechanism design, and are a canonical example of a *truthful* mechanism: neither player has incentive to cheat, as long as the auction is computed correctly. For a positive integer N, the two-player $(\log N)$-bit Vickrey auction is defined as $f : X \times Y \to Z \times \{A, B\}$ where

	1	2	3	4	...	2^n-1	2^n
1	$(1,B)$	$(1,B)$	$(1,B)$	$(1,B)$	...	$(1,B)$	$(1,B)$
2	$(1,A)$	$(2,B)$	$(2,B)$	$(2,B)$	...	$(2,B)$	$(2,B)$
3	$(1,A)$	$(2,A)$	$(3,B)$	$(3,B)$	...	$(3,B)$	$(3,B)$
4	$(1,A)$	$(2,A)$	$(3,A)$	$(4,B)$	...	$(4,B)$	$(4,B)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	...	$\vdots$	$\vdots$
2^n-1	$(1,A)$	$(2,A)$	$(3,A)$	$(4,A)$	...	$(2^n-1,B)$	$(2^n-1,B)$
2^n	$(1,A)$	$(2,A)$	$(3,A)$	$(4,A)$	...	$(2^n-1,A)$	$(2^n,B)$

Fig. 1. The matrix M_f for two-player n-bit Vickrey auction.

$X = Y = Z = \{1, 2, \ldots, N\}$ and

$$f(x,y) = \begin{cases} (x,B), & \text{if } x \le y \\ (y,A) & \text{if } y < x \end{cases}$$

Two players, Alice and Bob, have private values x and y, respectively. These private values indicate the amount of money that the item is worth to each of them. If $x \le y$, then Bob wins, and the price that he pays is x. (Thus, $f(x,y) = (x,B)$ means that Bob wins and pays x for the item.) Similarly, if $x > y$, then Alice wins, and the price that she pays is y. This mechanism is also called "2^{nd}-price auction" because the winner's price is the 2^{nd}-highest bid. Vickrey auctions remain truthful for more than two players, but are not computable with perfect privacy ($\mathrm{PAR} = 1$) for more than two players [Brandt and Sandholm 2008].

The matrix M_f of the n-bit Vickrey auction is shown in Figure 1.

Perfect privacy for two-player Vickrey auctions is achieved by the successive English bidding protocol, in which bids start at 1 and increase by 1 in each round, and the first player to drop out of bidding reveals his entire private value. (Note that this incurs no loss of privacy, since that value is part of the function output.) The protocol tree for this protocol is given in Figure 2. This protocol takes 2^{n+1} rounds for the n-bit Vickrey auction, and is known to be the only protocol which obtains perfect privacy $\mathrm{PAR} = 1$ for Vickrey auctions [Kushilevitz 1989].

THEOREM 3.3. *[Kushilevitz 1989] Perfect privacy for two-player n-bit Vickrey is only achievable by the 2^{n+1}-length English auction.*

Notice that the range of f is of size 2^{n+1} and that f is surjective, so that there must be at least 2^{n+1} distinct leaves in any protocol tree for f. Thus any protocol for f requires at least $n+1$ rounds. An example of such a protocol is the Bisection Protocol that proceeds by binary search on an interval containing the smaller input [Feigenbaum et al. 2010a]. Bisection Protocol obtains $\mathrm{PAR} = 2^n$, the worst possible loss of privacy for this function.

These two extremes – on the one hand $\mathrm{PAR} = 1$ at exponential communication cost, and on the other, exponential PAR at linear communication cost – suggest that there is a tradeoff between privacy and communication for Vickrey auctions. The structure of the function itself suggests this tradeoff as well. Any move which differs from the English protocol must divide some monochromatic region into two pieces. Thus inputs in the same monochromatic region are distinguishable by the protocol, and some privacy is lost.

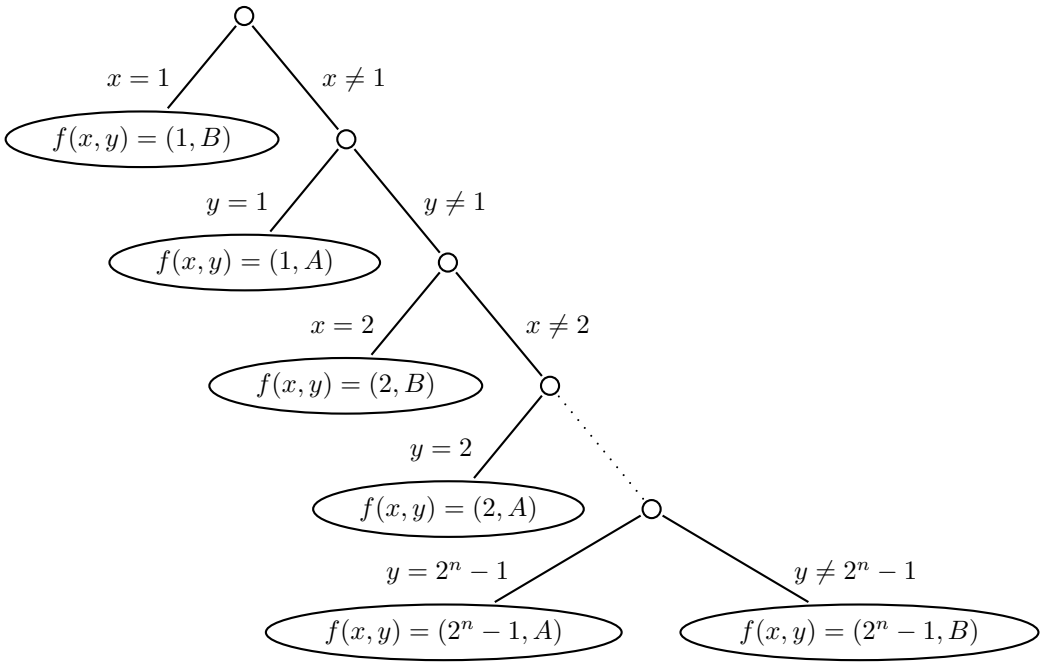

Fig. 2. The protocol tree for an English auction computing f.

Different PAR is achievable depending on the nature of the protocol. Feigenbaum et al. [2010a] examine a family of Bisection-type protocols and use average-case PAR to differentiate amongst them. Such protocols obtain worst-case PAR varying from 1 to 2^n, inversely related to their length. This observation inspired the results below.

3.2. Worst-case lower bound for Vickrey auction

The two algorithms discussed in the previous section suggest that any protocol computing Vickrey auctions should have a tradeoff between length and privacy. Protocol steps which resemble those of the ascending English bidding protocol partition the inputs in an unbalanced way, so that most inputs follow one branch of the protocol tree, and few inputs follow the other branch. Such steps preserve privacy but do not make much progress. (In an imbalanced partition, on the larger side the protocol still has a lot of work to do in order to compute the function.) On the other hand, protocol steps that resemble binary search partition the inputs in a nearly balanced way. Such steps make good progress, but are bad for privacy. (Dividing the remaining inputs in half increases the PAR by a factor of 2.) This is the intuition behind Theorem 1.1, stated again below.

THEOREM 1.1. *For all n, for all p, $2 \leq p \leq n/4$, any deterministic protocol for the two-player n-bit Vickrey auction problem with communication cost (length) less than $n2^{\frac{n}{4p}-5}$ obtains privacy loss (worst-case PAR) at least 2^{p-2}.*

Here the variable p serves as a parameter, explicitly linking the protocol length to the achievable PAR. For instance, if we put $p = \sqrt{n}$, then we conclude by Theorem 1.1 that either the protocol communicates $2^{\Omega(\sqrt{n})}$ bits in the worst case, or the worst-case privacy loss is $2^{\Omega(\sqrt{n})}$. This theorem shows that for Vickrey auctions, there is an inherent tradeoff between communication complexity and privacy.

PROOF. We will assume without loss of generality that in the protocol, the players take turns and send one bit per message. (Any protocol can be put into this form by at most doubling the length of the protocol.) Moreover, our protocol is assumed to be deterministic and to have zero error.

The Vickrey auction function has a corresponding matrix M such that entry (x, y) of the matrix is the value of the Vickrey auction on inputs (x, y) (Figure 1). A submatrix of M is called a *rectangle*; a rectangle is "monochromatic" if the matrix is constant on inputs in that submatrix. Every communication protocol can be visualized as a binary decision tree [Yao 1979]. Each node v of the tree is associated with a rectangle (submatrix) $T(v) = T_A(v) \times T_B(v) \subseteq X \times Y$. The root node r is associated with the entire matrix $T_A(r) \times T_B(r) = X \times Y = M$. Each leaf node l is associated with a monochromatic submatrix $T_A(l) \times T_B(l)$. Each internal node v has two children, v_0 and v_1. If the protocol calls for Alice to speak at node v, then the bit sent by Alice at v induces a partition of $T_A(v)$ into two pieces, $T_A(v_0)$ and $T_A(v_1)$. The submatrix associated with v_0 is $T_A(v_0) \times T_B(v)$, and the submatrix associated with v_1 is $T_A(v_1) \times T_B(v)$. Similarly if Bob speaks at node v, then the submatrix associated with v_0 is $T_A(v) \times T_B(v_0)$ and the submatrix associated with v_1 is $T_A(v) \times T_B(v_1)$.

Traversing the tree from the root to a leaf l generates a transcript of the bits of communication sent for some input $(x, y) \in T_A(l) \times T_B(l)$. The depth of the tree is the worst-case communication cost of the protocol.

Any deterministic protocol consists of a series of partitions of the matrix M into rectangles. The resulting protocol-induced tiling of the matrix M is a partition into monochromatic rectangles, which are precisely the rectangles associated with the leaves of the protocol's decision tree.

For every correct communication protocol, we describe an adversary strategy that follows a path through the protocol tree and finds some input (x, y) such that either: (i) the privacy loss of (x, y) is large i.e., $\mathrm{PAR}(x, y) \geq 2^{p-2}$, or (ii) the communication protocol on (x, y) requires at least $\frac{\ln 2}{4}(n - 2p)2^{n/4p}$ bits to compute.

Let M denote the matrix corresponding to the Vickrey auction problem, as drawn in Figure 1. Bob wins for inputs in the horizontal regions; Alice wins for inputs in the vertical regions. Fix a communication protocol, and corresponding protocol tree, P. For every node v in P that our adversary strategy selects, we will maintain three sets: $S(v), A^L(v), B^L(v) \subseteq [2^n]$. At node v, the adversary will be interested in tracking the privacy loss on the set of inputs $S(v) \times S(v)$. The privacy loss for these inputs will be measured with the help of the two auxiliary sets $A^L(v)$ and $B^L(v)$, respectively.

Initially, at the root r of the protocol tree, $S(r) = [2^{n-p}]$. This initial set of inputs $S(r) \times S(r)$ are the "small" inputs that sit in the upper left submatrix of M. As we move down the protocol tree, we will update $S(v)$ so that it is always a subset of $[2^{n-p}] \cap T_A(v) \cap T_B(v)$. We are interested in these small inputs since the regions that they are contained in are very large, and thus have the potential to incur a large (exponential) privacy loss.

The set $A^L(v)$ is a subset of $T_A(v)$, and similarly $B^L(v)$ is a subset of $T_B(v)$. The sets $A^L(r)$ and $B^L(r)$ are initially $[2^n] \backslash [2^{n-p}]$, the "large" inputs. At vertex v, the set $A^L(v)$ describes the set of large inputs of Alice that have survived so far; thus $A^L(v) = T_A(v) \cap [2^n] \backslash [2^{n-p}]$. Similarly, $B^L(v)$ describes the set of large inputs of Bob that have survived so far; thus $B^L(v) = T_B(v) \cap [2^n] \backslash [2^{n-p}]$. As we traverse the protocol tree, these sets track the loss of privacy for Alice and Bob (respectively) on inputs in $S(v) \times S(v)$.

We can measure the loss of privacy *so far* in the protocol. For any $(x, y) \in T(v)$,

$$\mathrm{PAR}_v(x, y) = \frac{|R_{x,y}|}{|R_{x,y} \cap T(v)|}.$$

If v is a leaf, then for any $(x, y) \in T(v)$, $\mathrm{PAR}_v(x, y) = \mathrm{PAR}(x, y)$. The following simple claim will be useful:

CLAIM 3. $\forall (x, y) \in T(v)$, $\mathrm{PAR}(x, y) \geq \mathrm{PAR}_v(x, y)$.

In particular the following fact is crucial to our argument. For any (x, y) in $S(r) \times S(r) \cap T(v)$, if (x, y) is in a vertical region ($y < x$, a win for Alice), then

$$\mathrm{PAR}(x, y) = \frac{|R_{x,y}|}{|P_{x,y}|} \geq \mathrm{PAR}_v(x, y) \geq \frac{2^n - 2^{n-p}}{|A^L(v)| + 2^{n-p}}.$$

This holds because $|R_{x,y}| \geq 2^n - 2^{n-p}$ and $|R_{x,y} \cap T(v)| \leq |A^L(v)| + 2^{n-p}$. Similarly, if $(x, y) \in S(r) \times S(r)$ is in a horizontal region ($x \leq y$, a win for Bob), then

$$\mathrm{PAR}(x, y) \geq \mathrm{PAR}_v(x, y) \geq \frac{2^n - 2^{n-p}}{|B^L(v)| + 2^{n-p}}.$$

The above inequality shows how $A^L(v)$ and $B^L(v)$ track the privacy loss of inputs $S(v) \times S(v)$: for those inputs $(x, y) \in S(v) \times S(v)$ where Alice wins, the privacy loss for (x, y) increases as $A^L(v)$ decreases, and similarly for those inputs where Bob wins, the privacy loss increases as $B^L(v)$ decreases.

Adversary Strategy: We are now ready to describe the adversary strategy. There are two cases, depending on whether it is Alice's or Bob's turn to send a message. We will first describe the case where at node v, it is Alice's turn to speak. Alice sends Bob some bit b which partitions her inputs $T_A(v)$ into two pieces. Since $S(v)$ and $A^L(v)$ are always subsets of $T_A(v)$, this induces a partition of $S(v)$ into $S_0(v)$ and $S_1(v)$ and $A^L(v)$ into $A^L_0(v)$ and $A^L_1(v)$.

Let $\alpha = 2^{\frac{-n}{4p}}$. We determine if a step made *progress* or was *useless* in the following way:

— If $\alpha|S(v)| \leq |S_0(v)| \leq (1 - \alpha)|S(v)|$ (hence $\alpha|S(v)| \leq |S_1(v)| \leq (1 - \alpha)|S(v)|$), then we say this step made **progress** on $S(v)$. In this case, the set $S(v)$ is partitioned into roughly balanced pieces. Select i such that $|A^L_i(v)| \leq \frac{1}{2}|A^L(v)|$.

— Otherwise, pick i such that $|S_i(v)| \geq (1-\alpha)|S(v)|$. In this case, we call it a **useless** step.

We update sets in the obvious way: if w is the new node in the protocol tree that we traverse to, then $S(w) = S_i(v)$ and $A^L(w) = A^L_i(v)$.

The second case is when it is Bob's turn to speak. Our adversary strategy is entirely symmetric. Now $T_B(v)$ is partitioned into two pieces, inducing a partition of $S(v)$ into $S_0(v)$ and $S_1(v)$, and a partition of $B^L(v)$ into $B^L_0(v)$ and $B^L_1(v)$. We pick i as above, but with A^L_i replaced with B^L_i.

The strategy continues as described above, traversing the protocol tree until one of the two events happens for the first time:

— Alice (or Bob) has made p progress steps, so $A^L(v)$ (or $B^L(v)$) has been halved at least p times.

— The strategy reaches a leaf node, and can go no further.

This completes the description of the strategy.

The following are the two main ideas in analyzing our strategy.

LEMMA 3.4. *Let our strategy reach node v and find Alice (or Bob) took p progress steps on the way. Then, for each $(x, y) \in S(v) \times S(v)$ such that $x > y$ (or $x \leq y$) $\mathrm{PAR}_v(x, y) \geq 2^{p-2}$.*

We can exit our strategy at this point and invoke Claim 3 to finish the argument. In the other case, we make the following claim:

LEMMA 3.5. *If our strategy reaches a leaf node v without Alice or Bob taking p progress steps, then for every $(x,y) \in T(v)$, the protocol communicates at least $\frac{\ln 2}{4}(n-2p)^{n/4p}$ bits.*

Thus, we would conclude that in this case the cost of the protocol is larger than $n2^{\frac{n}{4p}-4}$. Hence, all that remains to finish our argument is to prove Lemma 3.4 and Lemma 3.5.

Proof of Lemma 3.4: Let r be the root node of our protocol tree. For each input $(x,y) \in S(r) \times S(r)$, note that $R_{x,y} \geq 2^n - 2^{n-p}$ and $|A^L(r)| = 2^n - 2^{n-p}$. Let φ be the path in the protocol tree from r to v that our strategy chooses such that Alice takes p progress steps along φ. Consider any pair of adjacent nodes u, w in path φ such that Alice makes progress in going from u to w. Then, by definition of our strategy, $|A^L(w)| \leq \frac{1}{2}|A^L(u)|$. Hence, $|A^L(v)| \leq \frac{1}{2^p}|A^L(r)|$. Thus, any input (x,y) in $S(v) \times S(v)$ on which Alice would win is contained in an induced rectangle of size at most $2^{n-p} + |A^L(v)|$. Claim 3 yields:

$$\mathrm{PAR}_v(x,y) \geq \frac{2^n - 2^{n-p}}{\frac{2^n - 2^{n-p}}{2^p} + 2^{n-p}} \geq 2^{p-2}$$

The analysis when Bob makes p progress steps proceeds very similarly. ■

Proof of Lemma 3.5: The strategy reaches a leaf node v traversing a path φ, and $|S(v)| = 1$. (If $|S(v)| > 1$, then there is more than one possible answer, and so the computation is not yet finished.) In this case, Alice and Bob each took fewer than p progress steps. Let q be the total number of useless steps followed to get to v. (The protocol is at most $2p+q$ long.) On each progress step (u,w) in path φ, by definition, $|S(w)| \geq \alpha|S(u)|$. On each useless step (u,w), the updated size of $|S(w)| \geq (1-\alpha)|S(u)|$. This gives a lower bound on the size of set $S(v)$. Hence $|S(v)| \geq 2^{n-p}\alpha^{2p}(1-\alpha)^q$.

Assume that $q < \frac{\ln 2}{4}(n-2p)2^{\frac{n}{4p}}$. The calculation below shows that $|S(v)| > 1$, thus deriving a contradiction to the fact that v is a leaf node where the protocol ends.

$$\begin{aligned}
|S(v)| &\geq 2^{n-p}\alpha^{2p}(1-\alpha)^q \\
&> 2^{n-p}(2^{-n/4p})^{2p}(1-2^{-n/4p})^{\frac{\ln 2}{4}(n-2p)2^{\frac{n}{4p}}} && \text{by assumption about } q \\
&= 2^{n/2-p}(1-2^{-\frac{n}{4p}})^{\frac{\ln 2}{4}(n-2p)2^{n/4p}} && \text{simplifying algebra} \\
&> 2^{\frac{n}{2}-p}e^{-2^{-n/4p}\cdot\frac{\ln 2}{2}(n-2p)2^{n/4p}} && \text{as } (1-x) > e^{-2x} \text{ for } x \in (0,1/2] \\
&= 2^{\frac{n}{2}-p}e^{-(\ln 2)\cdot(\frac{n}{2}-p)} && \text{by simplification} \\
&= 1
\end{aligned}$$

■

Thus we know that either there is an input with privacy loss at least 2^{p-2} or an input with communication at least $\frac{\ln 2}{4}(n-2p)2^{n/4p} \geq \frac{n}{12}2^{n/4p}$ bits. Since we might have doubled the communication cost by alternating Alice and Bob bits, we obtain the resulting lower bound.

■

Note. The tradeoff of Theorem 1.1 holds for both the objective PAR and subjective PAR. For Vickrey auctions they coincide (Lemma 3.6), because all regions are rectangles with width or depth one.

LEMMA 3.6. *For $N \geq 1$, let P be any protocol for two-player $(\log N)$-bit Vickrey auction. Then* $\mathrm{PAR}(P) = \mathrm{PAR}^{sub}(P)$.

PROOF. $\mathrm{PAR}(P) \leq \mathrm{PAR}^{sub}(P)$. Let (x, y) be the input which maximizes $\frac{|R_{x,y}|}{|P_{x,y}|}$. If $x \leq y$, then $\mathrm{PAR}(P) = \frac{|R_{x,y}|}{|P_{x,y}|} \leq \max_{y'}(\frac{|R_{x,y'} \cap \{x\} \times Y|}{|P_{x,y'} \cap \{x\} \times Y|}) \leq \mathrm{PAR}^{sub}(P)$. The case for $y < x$ is similar.

$\mathrm{PAR}^{sub}(P) \leq \mathrm{PAR}(P)$. Let (x, y) be the input which maximizes $\mathrm{PAR}^{sub}(P)$. If $x \leq y$, then $\mathrm{PAR}^{sub}(P) = (\frac{|R_{x,y'} \cap \{x\} \times Y|}{|P_{x,y'} \cap \{x\} \times Y|}) \leq \max_{x,y} \frac{|R_{x,y}|}{|P_{x,y}|} = \mathrm{PAR}(P)$. The case for $y < x$ is similar. ■

4. AVERAGE-CASE PAR

In this section we consider the average-case privacy approximation ratio. For a probability distribution D on $X \times Y$ and a protocol P for a function $f : X \times Y \to Z$, Feigenbaum et al. [2010a] define the average-case PAR as follows:

$$\mathrm{avg}\,\mathrm{PAR}(P) = \mathbb{E}_D\left[\frac{|R_{x,y}|}{|P_{x,y}|}\right].$$

In this paper we will consider the following alternative definition.

Definition 4.1. For a probability distribution D on $X \times Y$ and a protocol P for a function $f : X \times Y \to Z$, let the **average-case objective privacy approximation ratio** of protocol P for function f be:

$$\mathrm{avg}_D\,\mathrm{PAR}(P) = \mathbb{E}_{(x,y)\in D}\left[\frac{|R_{x,y}|_D}{|P_{x,y}|_D}\right],$$

where for $S \subseteq X \times Y$, $|S|_D = \sum_{(x,y)\in S} D(x, y)$. Furthermore, we let the **average-case subjective privacy approximation ratio** of protocol P for function f be:

$$\mathrm{avg}_D\,\mathrm{PAR}^{sub}(P) = \max\left\{\mathbb{E}_{(x,y)\in D}\left[\frac{|R_{x,y} \cap X \times \{y\}|_D}{|P_{x,y} \cap X \times \{y\}|_D}\right], \mathbb{E}_{(x,y)\in D}\left[\frac{|R_{x,y} \cap \{x\} \times Y|_D}{|P_{x,y} \cap \{x\} \times Y|_D}\right]\right\}.$$

As opposed to Feigenbaum et al. we measure the size of subsets of $X \times Y$ relative to the measure D. This "corrected" definition coincides with the definition of Feigenbaum et al. [2010a] for the uniform distribution. Their paper does not give any results for distributions other than uniform, so our definition is consistent with their results. Similarly, most of our results for concrete functions are for the uniform distribution, so they hold under both definitions.

Despite the fact that Feigenbaum et al. argue the opposite, we believe this definition by probability mass is a natural choice. When the players (or an eavesdropper) have previous knowledge of D, the loss of privacy of a protocol should be related to the perceived size of regions and not their actual size. This is true for example when all the measure in each protocol rectangle is concentrated on a single input. Then knowing D and the protocol transcript reveals everything about the input. Another example is when all the measure of each region of a function is concentrated in a single protocol rectangle (where we assume that all these rectangles are of the same size). Then the protocol reveals very little about its actual input drawn from D except for the function value, and so ought to be considered "private". The original measure of Feigenbaum et al. does not make any distinction between these examples.

Definition 4.1 is motivated by an attempt to prove Theorem 1.2, and will be convenient and useful in that proof (see Proposition 4.2). Our definition has interesting mathematical properties and (as we will see in a moment) it is related to other known

measures. For further discussion of alternative definitions of average-case PAR, see section 8.1 of Feigenbaum et al. [2010a].

One benefit of Definition 4.1 is that one can relate average-case PAR to another natural measure on protocols. Consider a protocol P for a function f. For a region $R \in \mathcal{R}(f)$ let $\mathrm{cut}_P(R) = |\{P_{x,y} \,|\, (x,y) \in R\}|$ be the number of protocol-induced rectangles contained within R. The following statement is implicit in Feigenbaum et al. [2010a] for the case of uniform distribution and objective PAR.

PROPOSITION 4.2. *For any function $f : X \times Y \to Z$, protocol P for f and any probability distribution D on $X \times Y$,*

$$\mathrm{avg}_D \,\mathrm{PAR}(P) = \sum_{R \in \mathcal{R}(f)} |R|_D \cdot \mathrm{cut}_P(R)$$

$$\mathrm{avg}_D \,\mathrm{PAR}^{sub}(P) \;=\; \max\Bigg\{ \sum_{y \in Y, R \in \mathcal{R}(f)} |R \cap X \times \{y\}|_D \cdot \mathrm{cut}_P(R \cap X \times \{y\}),$$
$$\sum_{x \in X, R \in \mathcal{R}(f)} |R \cap \{x\} \times Y|_D \cdot \mathrm{cut}_P(R \cap \{x\} \times Y) \Bigg\}.$$

PROOF. For any protocol-induced rectangle A, $\sum_{(x,y) \in A} D(x,y) \cdot \frac{1}{|A|_D} = 1$. Hence,

$$\begin{aligned}
\mathrm{avg}_D \,\mathrm{PAR}(P) \;&=\; \mathbb{E}_D \left[\frac{|R_{x,y}|_D}{|P_{x,y}|_D} \right] \\
&= \sum_{(x,y) \in X \times Y} D(x,y) \cdot \frac{|R_{x,y}|_D}{|P_{x,y}|_D} \\
&= \sum_{R \in \mathcal{R}(f)} \sum_{(x,y) \in R} \frac{D(x,y) \cdot |R|_D}{|P_{x,y}|_D} \\
&= \sum_{R \in \mathcal{R}(f)} |R|_D \Bigg(\sum_{(x,y) \in R} \frac{D(x,y) \cdot 1}{|P_{x,y}|_D} \Bigg) \\
&= \sum_{R \in \mathcal{R}(f)} |R|_D \cdot \mathrm{cut}_P(R).
\end{aligned}$$

The case of subjective PAR is analogous. ■

In the setting of our definition, this characterization of average-case PAR provides a simple answer to the conjecture [Feigenbaum et al. 2010a] that for any probability distribution D on inputs, there is a protocol that has average-case PAR at most n for the n-bit Vickrey auction. Recall that the Bisection Protocol for the Vickrey auction proceeds by binary search on the input domain [Feigenbaum et al. 2010a].

PROPOSITION 4.3. *For any probability distribution D on $[2^n] \times [2^n]$, the Bisection Protocol for the two-player n-bit Vickrey auction satisfies:*

$$\mathrm{avg}_D \,\mathrm{PAR}(\textit{Bisection Protocol}) \leq n + 1.$$

PROOF. Each region R of the n-bit Vickrey auction is covered by at most $n+1$ rectangles induced by the Bisection Protocol, i.e., $\mathrm{cut}_{\text{Bisection Protocol}}(R) \leq n + 1$. The claim follows by the previous proposition. ■

The relation between objective and subjective privacy approximation ratios (Lemma 3.6) for Vickrey auctions extends to the average-case setting.

LEMMA 4.4. *For $N \geq 1$, let P be any protocol for two-player $(\log N)$-bit Vickrey auction. If U is the uniform probability distribution on $[N] \times [N]$ then*

$$\operatorname{avg}_U \operatorname{PAR}^{sub}(P) \leq \operatorname{avg}_U \operatorname{PAR}(P) \leq 2 \operatorname{avg}_U \operatorname{PAR}^{sub}(P).$$

PROOF. To prove the relationship for the average-case PAR, consider input $(x, y) \in [N] \times [N]$. If $x \leq y$ then $R_{x,y} \cap \{x\} \times Y = R_{x,y}$ and $R_{x,y} \cap X \times \{y\} = \{(x, y)\}$. If $x > y$ then $R_{x,y} \cap \{x\} \times Y = \{(x, y)\}$ and $R_{x,y} \cap X \times \{y\} = R_{x,y}$. Identically for $P_{x,y}$ instead of $R_{x,y}$. Hence,

$$\frac{|R_{x,y} \cap X \times \{y\}|}{|P_{x,y} \cap X \times \{y\}|} = \frac{|R_{x,y}|}{|P_{x,y}|}$$

if $x \leq y$, and

$$\frac{|R_{x,y} \cap X \times \{y\}|}{|P_{x,y} \cap X \times \{y\}|} = 1 \leq \frac{|R_{x,y}|}{|P_{x,y}|}$$

otherwise. On the other hand

$$\frac{|R_{x,y} \cap \{x\} \times Y|}{|P_{x,y} \cap \{x\} \times Y|} = 1 \leq \frac{|R_{x,y}|}{|P_{x,y}|}$$

if $x \leq y$ and if $x > y$ then

$$\frac{|R_{x,y} \cap \{x\} \times Y|}{|P_{x,y} \cap \{x\} \times Y|} = \frac{|R_{x,y}|}{|P_{x,y}|}.$$

Thus, $\operatorname{avg}_U \operatorname{PAR}^{sub}(P) \leq \operatorname{avg}_U \operatorname{PAR}(P)$. For the upper bound

$$\begin{aligned}\sum_{x,y} \frac{1}{N^2} \cdot \frac{|R_{x,y}|}{|P_{x,y}|} = &\sum_{x \leq y} \frac{1}{N^2} \cdot \frac{|R_{x,y} \cap X \times \{y\}|}{|P_{x,y} \cap X \times \{y\}|} \\ &+ \sum_{x > y} \frac{1}{N^2} \cdot \frac{|R_{x,y} \cap \{x\} \times Y|}{|P_{x,y} \cap \{x\} \times Y|}.\end{aligned}$$

Hence, $\operatorname{avg}_U \operatorname{PAR}(P) \leq 2 \operatorname{avg}_U \operatorname{PAR}^{sub}(P)$. ■

For the uniform distribution we can prove the following tradeoff between the length and average-case PAR of any protocol. This is one of our main results.

THEOREM 1.2. *For all $n, r \geq 1$, any deterministic protocol for the two-player n-bit Vickrey auction problem (over the uniform distribution of inputs) with communication cost (length) less than r obtains average-case PAR at least $\Omega(\frac{n}{\log(r/n)})$.*

This bound is asymptotically tight for the uniform distribution (the n/r-Bisection Protocol achieves asymptotically the same upper-bound). Our lower bound holds only for the uniform distribution on inputs. This is not surprising; if the distribution is concentrated say on a single input one should not expect large loss of privacy. Using Lemma 4.4 one obtains a similar tradeoff also for the average-case subjective PAR.

The rest of this section (up to subsection 4.1) is devoted to the proof of Theorem 1.2. Proposition 4.2 characterizes the average-case PAR as the weighted sum of $\operatorname{cut}_P(R)$ over all regions R of the function. We will use this characterization but simplify the calculation a little bit.

— We will sum only over regions $R_{x,y}$ for $x, y \leq 2^{n-1}$. Call this collection of regions L. These are the largest regions in $X \times Y$, and together cover $\frac{3}{4}$ the area of $X \times Y$. Hence the loss of privacy on these regions will be significant. Each of the regions is of size between 2^{n-1} and 2^n, so they all have the same weight up to a factor of at most 2.

— To estimate $\mathrm{cut}_P(R)$ for various regions R we will track only the set of "diagonal" inputs $Diag = \{(x,x) \mid x \in [2^{n-1}]\}$ as they progress in the protocol tree, and count protocol-induced rectangles that intersect regions $R_{x,x}$ and $R_{x+1,x}$.

Combining these two simplifications gives a lower bound on the average-case PAR for the uniform distribution:

$$\frac{2^{n-1}}{4^n} \sum_{R \in L} \mathrm{cut}_P(R). \tag{1}$$

Note that each input pair $(x,x) \in Diag$ must finish the protocol in a separate induced rectangle.

The problem of counting the cuts of interest (in order to get a lower bound) can be abstracted away into the Ball Partition Problem. By Lemma 4.8, a lower bound on the Ball Partition Problem will yield a lower bound on the average-case PAR for the uniform distribution on Vickrey auctions.

Definition 4.5 (*Ball Partition Problem*). For integers N and $r \geq 1$, there are N balls and r rounds. All of the balls begin in one big set. In each round, the balls in each current set are partitioned into (at most) two new sets. The cost of partitioning the balls in any set S into sets S_1 and S_2 is $\min(|S_1|, |S_2|)$. After r rounds, each of the N balls shall be in a singleton set. The total cost of the game is the sum of the cost, over all r rounds, of every partition made during each round. We denote the minimal possible cost by $B(N,r)$.

The interesting values of r lie in a particular range. For $r < \log_2 N$, the game cannot be finished at any cost. For $r > N$, the game can easily be finished with minimal cost $B(N,r) = N - 1$: cut away 1 ball from the largest set at every round. However, for intermediate values $\log N \leq r \leq N$, one might ask: what is the smallest possible cost c achievable in r rounds?

THEOREM 4.6. *For the Ball Partition Problem,* $B(N,r) \geq \frac{N \log N}{4 \log(\frac{4r}{\log N})}$.

The above lower bound is asymptotically optimal as shown in the following proposition.

PROPOSITION 4.7. *Let* N *and* r *be integers such that* $2 \log N \leq r$. *For the Ball Partition Problem,* $B(N,r) \leq O\left(\frac{N \log N}{\log(\frac{r}{\log N})}\right)$.

PROOF. Ignoring the rounding issues, at each round we can split each non-singleton set S into two sets of sizes $\alpha|S|$ and $(1-\alpha)|S|$, for $\alpha = (\log N)/r \leq 1/2$. It follows that within r rounds, each set contains at most one element as $(1-\alpha)^r N \leq N e^{-\alpha r} < 1$. The total cost of the ball partitioning is the sum of sizes of all the smaller sets obtained in each partition. This corresponds to the number of elements in these sets (counting multiplicity). Each element can appear in at most $\log_{1/\alpha} N = (\log N)/\log(r/\log N)$ of the smaller sets as the size of the set containing the element shrinks by factor of α on each such occasion. Hence, the total cost is at most $N \cdot (\log N)/\log(r/\log N)$. Always rounding the size of the smaller set up will introduce a constant factor in the final bound. ■

Lemma 4.8 relates a lower bound for the Ball Partition Problem (N balls in r rounds) to a lower bound for the average-case Vickrey auction on the uniform distribution (N possible inputs for each player and r bits of communication).

LEMMA 4.8. *Let $N, r \geq 1$ be integers where N is a power of two. Let $B(N,r)$ be the minimal cost of the Ball Partition Problem on N balls in r rounds. Then for any deterministic r-bit protocol P for two-player $(\log N)$-bit Vickrey auction, the average-case PAR is $\operatorname{avg}\mathrm{PAR}(P) \geq \frac{B(N/2,r)}{2N}$ under the uniform distribution.*

Proof of Lemma 4.8: Our goal is to establish that $\sum_{R\in L} \mathrm{cut}_P(R) \geq B(N/2,r)$. The lemma easily follows from this since each region R in L contains probability mass at least $1/2N$ under the uniform distribution.

The Ball Partition Problem is an abstraction of the calculation of average-case PAR for Vickrey auctions. Recall the following notation used in the proof of Theorem 1.1. Protocol P is associated with a protocol tree where each node v corresponds to a combinatorial rectangle $T(v) = T_A(v) \times T_B(v) \subseteq X \times Y$. For $t = 0, \ldots, r$, let $\mathcal{R}(P,t)$ be the set of rectangles associated with nodes at level t of the tree, level 0 consisting of the root. For $R \subseteq X \times Y$, let $\mathrm{cut}_P(R,t) = |\{S \in \mathcal{R}(P,t); S \cap R \neq \emptyset\}|$ be the number of rectangles intersecting R after round t of the protocol. Clearly, $\mathrm{cut}_P(R,r) = \mathrm{cut}_P(R)$. We want to estimate from below $\sum \mathrm{cut}_P(R)$ over $R \in L$.

We associate every node v of the protocol tree with sets $D_v = [N/2] \cap T_A(v) \cap T_B(v)$ and $L_v = \{R_{x,y}; x, y \in D_v\}$. For each leaf node v, $|D_v| \leq 1$ as no two distinct inputs (x,x) and (x',x') can finish in the same protocol-induced rectangle of the leaf. Notice, $L_v \subseteq L$. It is easy to see by induction on the level of the tree that sets D_v associated with nodes at the same level partition $[N/2]$ and hence, sets L_v associated with nodes at the same level are disjoint. Let v be a node at level t, $0 \leq t < r$, with $D_v \neq \emptyset$. Let v_1 and v_2 be its two children. If $D_{v_1} \neq \emptyset \neq D_{v_2}$ then we claim that

$$\sum_{R\in L_v} \mathrm{cut}_P(R,t+1) \geq \sum_{R\in L_v} \mathrm{cut}_P(R,t) + \min(|D_{v_1}|, |D_{v_2}|) - 1.$$

We prove the claim. Assume that v is a node where Alice speaks. Hence, $T_A(v) = T_A(v_1) \dot{\bigcup} T_A(v_2)$ and $T_B(v) = T_B(v_1) = T_B(v_2)$. Clearly, $D_v = D_{v_1} \dot{\bigcup} D_{v_2}$. Let $x_1 = \max(D_{v_1})$ and $x_2 = \max(D_{v_2})$. WLOG, $x_1 < x_2$. For every $y \in D_{v_1}, y \neq x_1$, $(x_1, y) \in R_{y+1,y} \cap T(v_1)$ and also $(x_2, y) \in R_{y+1,y} \cap T(v_2)$, so both are non-empty. Hence, $\mathrm{cut}_P(R_{y+1,y}, t+1) \geq \mathrm{cut}_P(R_{y+1,y}, t) + 1$. As there are $|D_{v_1}| - 1$ such $y's$, the claim follows in this case.

If v is a node where Bob speaks, the argument is similar. Let $y_1 = \max(D_{v_1})$ and $y_2 = \max(D_{v_2})$, and assume WLOG $y_1 < y_2$. Then for every $x \in D_{v_1}$, $(x, y_1) \in R_{x,x} \cap T(v_1)$ and also $(x, y_2) \in R_{x,x} \cap T(v_2)$. Thus in this case one does not even lose the -1 additive term.

Hence, each node v, for which D_v is split into two non-empty sets D_{v_1} and D_{v_2}, contributes by at least $\min(|D_{v_1}|, |D_{v_2}|) - 1$ to the increase of $\sum_{R\in L} \mathrm{cut}_P(R)$ overall. There are exactly $N/2-1$ nodes like that as $|D_{\mathrm{root}}| = N/2$. These sets D_v constitute a solution to the Ball Partition Problem in r rounds, and given the cost function for the Ball Partition Problem it is immediate that the overall increase of $\sum_{R\in L} \mathrm{cut}_P(R)$ is thus at least $B(N/2,r) - (N/2-1)$ as the -1 terms add up to $N/2-1$. Since $\sum_{R\in L} \mathrm{cut}_P(R,0) = N-1$ we get $\sum_{R\in L} \mathrm{cut}_P(R) \geq B(N/2,r)$. ■

All that remains to prove the lower bound on average-case PAR for Vickrey auctions (Theorem 1.2) is to prove the lower bound on the Ball Partition Problem (Theorem 4.6).

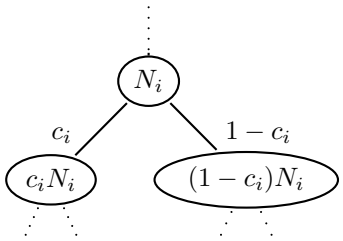

Fig. 3. An arbitrary node in the ball-partitioning tree.

Proof of Theorem 4.6: We will examine the entropy of the partitions at each round. This permits an abstraction away from a particular ball-partitioning instance, in order to obtain general properties. This will lead to a lower bound on the objective function $B(N, r)$, the cost of the Ball Partition Problem.

It will be useful to associate with the Ball Partition Problem in r rounds a full binary tree of depth r where each set obtained at round t is associated to a distinct node at level t, and remaining nodes are associated with the empty set. The association should be so that a node associated with a set S has its children associated with sets S_1 and S_2 obtained from S during the partitioning. We label each node i, by the size of the associated set N_i, and we label edges by the fraction of balls that travel "over" that edge from the parent to the child node. (See Figure 3: a node labelled N_i with children labelled $c_i N_i$ and $(1 - c_i)N_i$ will have edges to those children labelled c_i and $1 - c_i$, respectively.)

The tree's root node is labelled N; each leaf is labelled 1 or 0. (The 0 leaves are a result of assuming the binary tree is full; if some ball is partitioned into a singleton set in round $i < r$, then in each subsequent round it is "partitioned" into two sets: the singleton set and the empty set.)

Remark 4.9. At each level of the tree, the sum of the node labels $= N$. Thus the sum of labels of all the non-leaf nodes in the tree is rN.

Consider the path followed by any ball b from the root to a leaf. It traverses edges labelled $d_1^b, d_2^b, \ldots, d_r^b$, where $\prod_{i=1}^{r} d_i^b = \frac{1}{N}$.

Multiplying this number for all balls gives a nice symmetrization which is true for all trees representing solutions to the Ball Partition Problem.

$$\left(\frac{1}{N}\right)^N = \prod_{b \text{ a ball}} \prod_{i=1}^{r} d_i^b \tag{2}$$

Consider some non-leaf node i of the tree, with edges to its children labelled c_i and $1 - c_i$ (Figure 3). Together, these edges contribute $(c_i)^{c_i N_i}(1 - c_i)^{(1-c_i)N_i}$ to the right-hand side of equation (2). (If $c_i = 0$ this term equals 1 by definition.) WLOG assume each $c_i \leq 1/2$. Equation (2) can be rewritten as:

$$\begin{aligned} \left(\frac{1}{N}\right)^N &= \prod_{\text{non-leaf node } i} (c_i)^{c_i N_i}(1 - c_i)^{(1-c_i)N_i} \\ -N \log N &= \sum_i N_i(-H(c_i)) \end{aligned} \tag{3}$$

Here $H(x) = x \log \frac{1}{x} + (1 - x) \log \frac{1}{1-x}$ is the binary entropy of x.

Since the leaf nodes are not included in the sum, $\sum_{\text{non-leaf node } i} N_i = rN$ (by Remark 4.9). Let $c = \sum_i \frac{c_i N_i}{rN}$ be the average cost of a cut in the Ball Partition Problem. Then the cost of the entire tree is $B(N,r) = crN$. Since H is concave, $\sum_i \frac{N_i}{rN} H(c_i) \leq H(\sum_i \frac{c_i N_i}{rN}) = H(c)$.

$$N \log N = rN \sum_i \frac{N_i}{rN} H(c_i) \leq rNH(c) \tag{4}$$

For the sake of contradiction, suppose that the cost of the tree $B(N,r) = crN < \frac{N \log N}{4 \log(\frac{4r}{\log N})}$. Then the average cost of a cut is $c < \frac{\log N}{4r \log(\frac{4r}{\log N})}$. This c can be rewritten as $c = \frac{x}{-\log x}$ for $x = \frac{\log N}{4r}$. Combining equation (4) and Lemma 4.10 (below),

$$\frac{\log N}{r} \leq H(c) = H\Big(\frac{x}{-\log x}\Big) < 4x = 4\frac{\log N}{4r} = \frac{\log N}{r}$$

The inequality makes this a contradiction. Therefore every tree of depth $\leq r$ must incur cost $\geq \frac{N \log N}{4 \log(\frac{4r}{\log N})}$. ■

LEMMA 4.10. *For $0 < x \leq \frac{1}{2}$, the binary entropy $H\big(\frac{x}{-\log x}\big) < 4x$.*

PROOF. For $0 < x \leq \frac{1}{2}$, $\log \frac{1}{x} \geq 1$ so clearly $0 < \big(\frac{x}{-\log x}\big) \leq \frac{1}{2}$. Let $y = \frac{x}{-\log x}$. Expanding,

$$H(y) = y \log \frac{1}{y} + (1-y) \log \frac{1}{1-y}$$

For $0 < y \leq \frac{1}{2}$, it is not difficult to see that $-\log(1-y) \leq 2y$ and $1 - y < 1$.

$$H(y) \leq y \log \frac{1}{y} + (1-y)2y < y \log \frac{1}{y} + 2y$$

Substituting for y and expanding,

$$H\left(\frac{x}{\log \frac{1}{x}}\right) < x\left(\frac{\log \log \frac{1}{x}}{\log \frac{1}{x}}\right) + x\left(\frac{\log \frac{1}{x}}{\log \frac{1}{x}}\right) + 2x\left(\frac{1}{\log \frac{1}{x}}\right)$$

Examination reveals that for $0 < x \leq \frac{1}{2}$, the parenthesized coefficients are each ≤ 1. Hence $H(\frac{x}{\log \frac{1}{x}}) < 4x$. ■

4.1. Mutual information

The definition of average-case PAR is closely related to previously studied concepts in communication complexity such as information content [Barak et al. 2010] and (information-theoretic) privacy [Klauck 2002]. The main distinction is that these concepts measure in terms of bits, and PAR does not. Next we recapitulate some of these measures, show their relationship to the average-case PAR, and use this connection to prove new lower bounds for the average-case PAR.

Among these notions, Klauck's privacy measure [2002] is most closely related to average-case PAR. Let D be a probability distribution on $X \times Y$. Let $(\mathbf{X}, \mathbf{Y}) \sim D$ be the random variable obtained by sampling according to D. For a function f on $X \times Y$, its protocol P, and inputs $(x,y) \in X \times Y$, recall that we let $\Pi_P(x,y)$ be the transcript of the protocol on input (x,y). Then $\Pi_P(\mathbf{X}, \mathbf{Y})$ is the random variable obtained by sampling a random input according to D. Klauck [2002] gives the following definition of privacy of a protocol.

$$\text{PRIV}_D(P) = \max\{I(\mathbf{X} : \Pi_P(\mathbf{X}, \mathbf{Y}) | \mathbf{Y}, f(\mathbf{X}, \mathbf{Y})), I(\mathbf{Y} : \Pi_P(\mathbf{X}, \mathbf{Y}) | \mathbf{X}, f(\mathbf{X}, \mathbf{Y}))\}.$$

The relationship between this measure and our average-case PAR is given by the following theorem.

THEOREM 4.11. *For a probability distribution D on $X \times Y$ and a protocol P for a function $f : X \times Y \to Z$, the following holds:*

$$\mathrm{PRIV}_D(P) \leq \log(\mathrm{avg}_D\, \mathrm{PAR}^{sub}(P)).$$

PROOF. By symmetry, it suffices to show that $I(\mathbf{X} : \Pi_P(\mathbf{X},\mathbf{Y})|\mathbf{Y}, f(\mathbf{X},\mathbf{Y})) \leq \log(\mathrm{avg}_D\, \mathrm{PAR}^{sub}(P))$.

$$\begin{aligned} I(\mathbf{X} : \Pi_P(\mathbf{X},\mathbf{Y})|\mathbf{Y}, f(\mathbf{X},\mathbf{Y})) &\leq H(\Pi_P(\mathbf{X},\mathbf{Y})|\mathbf{Y}, f(\mathbf{X},\mathbf{Y})) \\ &\leq \sum_{y\in Y, z\in Z} |R_z \cap X \times \{y\}|_D \cdot \log(\mathrm{cut}_P(R_z \cap X \times \{y\})) \\ &\leq \log(\mathrm{avg}_D\, \mathrm{PAR}^{sub}(P)), \end{aligned}$$

The first inequality holds by simple algebra. The second inequality holds because, for any $y \in Y$ and $z \in Z$, $\Pr[\mathbf{Y} = y, f(\mathbf{X},\mathbf{Y}) = z] = |R_z \cap X \times \{y\}|_D$ and $H(\Pi_P(\mathbf{X},\mathbf{Y})|\mathbf{Y} = y, f(\mathbf{X},\mathbf{Y}) = z) \leq \log(\mathrm{cut}_P(R_z \cap X \times \{y\}))$. The final inequality follows from concavity of logarithm. ■

Hence, one can use lower bounds on PRIV to derive lower bounds for average-case PAR. For example, consider the function $\mathrm{DISJ}_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ on inputs $x, y \in \{0,1\}^n$, which is defined to be one if $\{i \in [n]; x_i = y_i = 1\}$ is empty and zero otherwise. Klauck [2002] shows that for any protocol P for the disjointness problem, $\mathrm{PRIV}_D(P) \in \Omega(\sqrt{n}/\log n)$, where D is uniform on strings of hamming weight $\sqrt{n}$. Using the above lower bound, we immediately obtain $\mathrm{avg}_D\, \mathrm{PAR}^{sub}(P) \in 2^{\Omega(\sqrt{n}/\log n)}$ for any protocol P for DISJ_n.

There are two other well studied measures that are closely related to our average-case PAR: the *external* and *internal information cost* (IC^{ext} and IC, resp.). The external information cost was defined in Chakrabarti et al. [2001] where the internal cost was also used implicitly. Later, using this measure, Bar-Yossef et al. [2004] obtained $\Omega(n)$ lower bounds on the randomized communication complexity of DISJ_n. The internal information cost was formalized in Barak et al. [2010]. For a protocol P for function $f : X \times Y \to Z$ and a distribution D on $X \times Y$, they are defined respectively as follows:

$$\mathrm{IC}^{ext}_D(P) = I(\mathbf{X},\mathbf{Y} : \Pi_P(\mathbf{X},\mathbf{Y}))$$

$$\mathrm{IC}_D(P) = I(\mathbf{X} : \Pi_P(\mathbf{X},\mathbf{Y})|\mathbf{Y}) + I(\mathbf{Y} : \Pi_P(\mathbf{X},\mathbf{Y})|\mathbf{X}).$$

As one can see the internal information cost is closely related to the privacy measure PRIV of Klauck. The only substantial difference is that PRIV is conditioned on the value of the function whereas IC is not. When f is a Boolean function, they are asymptotically identical.

PROPOSITION 4.12. *For any probability distribution D on $X \times Y$ and any protocol P for a function $f : X \times Y \to Z$:*

$$\mathrm{PRIV}_D(P) - \log|Z| \leq \mathrm{IC}_D(P) \leq 2 \cdot (\mathrm{PRIV}_D(P) + \log|Z|).$$

The proposition follows from Claim 2.

4.2. Set Intersection

The relationship described in Proposition 4.12, together with the known lower bounds on internal information cost of DISJ_n, allows us to prove one of the conjectures of Feigenbaum et al. [2010a] for the intersection function $\mathrm{INTERSEC}_n$. Function

$\text{INTERSEC}_n : \{0,1\}^n \times \{0,1\}^n \to \mathcal{P}([n])$ on inputs $x, y \in \{0,1\}^n$ gives the set $\{i \in [n]; x_i = y_i = 1\}$.

Feigenbaum et al. conjecture that the average-case subjective PAR for the intersection function under the uniform distribution is exponential in n. This can be proven using the above tools and the following result, which strengthens an earlier work by Bar-Yossef et al. [2004]. Let ν be the uniform distribution supported on $\{(0,1),(1,0),(0,0)\}$. Let τ be the distribution generated by taking the n-fold product of ν. In other words, τ is the uniform distribution supported on pairs of strings that are disjoint.

THEOREM 4.13. *[Braverman 2011] Let P be any randomized protocol that computes disjointness DISJ_n with error probability $< 1/3$. Then,* $\text{IC}_\tau\left(P\right) = \Omega(n)$.

Using the above theorem, we show the following bound for Intersection.

THEOREM 4.14. *Let P be any deterministic protocol that computes set intersection INTERSEC_n. Then, for U the uniform distribution,* $\text{PRIV}_U\left(P\right) = \Omega(n)$.

PROOF. We prove this by a contradiction. Assume that we have a protocol P to solve INTERSEC_m on m-bit inputs with little privacy loss under the uniform distribution. The main idea of the argument is to come up with an appropriate reduction from set disjointness DISJ_n on n bits to set intersection INTERSEC_m. This reduction will need to satisfy the following features: solving intersection on the reduced instance should solve set-disjointness on the original input instance. The reduced instance should not blow up too much in size, i.e. $m = \Theta(n)$. Finally, and most importantly, distribution τ on input instances to set-disjointness should generate (by our reduction) the uniform distribution on Intersection. This last step seems difficult to do via a deterministic reduction. So we aim to get a workaround as follows.

Let $\mathbf{\Pi}$ be the random variable denoting the transcript generated by P. Then, our assumption on P gives the following for some constant β which we fix at the end: $\beta\, m > I_U\big(\mathbf{X} : \mathbf{\Pi} \,|\, \mathbf{Y}, \text{INTERSEC}(\mathbf{X},\mathbf{Y})\big) + I_U\big(\mathbf{Y} : \mathbf{\Pi} \,|\, \mathbf{X}, \text{INTERSEC}(\mathbf{X},\mathbf{Y})\big)$.

The uniformly distributed pairs of m-bit random strings $(\mathbf{X},\mathbf{Y})$ can be alternatively generated by first selecting a random subset $\mathbf{A}$ of $[m]$ where each element is in the set independently with probability $1/4$. For each $i \in \mathbf{A}$, we set $(\mathbf{X}_i, \mathbf{Y}_i) = (1,1)$. Then, for each coordinate $i \in \overline{\mathbf{A}} = [m] - \mathbf{A}$, $(\mathbf{X}_i, \mathbf{Y}_i)$ is picked independently according to ν. Let τ denote the joint distribution $(\mathbf{X},\mathbf{Y},\mathbf{A})$ sampled as described. Let $(\mathbf{X},\mathbf{Y}|\mathbf{A})$ denote pair of random variables that are distributed according to $\mathbf{X},\mathbf{Y}$ conditioned on $\mathbf{A}$ as above and the underlying distribution on this pair be denoted by $\tau_{\mathbf{A}}$. Thus, our assumption becomes equivalently:

$$\mathbb{E}_{\mu_A}\left[I_{\tau_{\mathbf{A}}}\big(\mathbf{X} : \mathbf{\Pi} \,|\, \mathbf{Y}, \mathbf{A}\big) + I_{\tau_{\mathbf{A}}}\big(\mathbf{Y} : \mathbf{\Pi} \,|\, \mathbf{X}, \mathbf{A}\big)\right] < \beta m,$$

where μ_A is the distribution on $\mathbf{A}$. Applying the Chernoff bound on the deviation of $|\mathbf{A}|$ from its expectation, one concludes:

$$\mathbb{E}_{\mu_A}\left[I_{\tau_{\mathbf{A}}}\big(\mathbf{X} : \mathbf{\Pi} \,|\, \mathbf{Y}, \mathbf{A}\big) + I_{\tau_{\mathbf{A}}}\big(\mathbf{Y} : \mathbf{\Pi} \,|\, \mathbf{X}, \mathbf{A}\big) \,\middle|\, |\mathbf{A}| \le m/2\right] < \frac{\beta m}{1 - \exp(-\Omega(m))}$$

Thus, there exists some fixed set a of size at most $m/2$ such that

$$I_{\tau_a}\big(\mathbf{X} : \mathbf{\Pi} \,|\, \mathbf{Y}, \mathbf{A} = a\big) + I_{\tau_a}\big(\mathbf{Y} : \mathbf{\Pi} \,|\, \mathbf{X}, \mathbf{A} = a\big) \;<\; \beta' m. \tag{5}$$

This set a is going to provide us with the workaround needed for the deterministic reduction. We define our reduction now w.r.t a. Set $n = m - |a| \ge m/2$. Let P' be a protocol that solves set-disjointness as follows: Given two n-bit strings (u,v), protocol P' first embeds u and v naturally into $a^c = [m] - a$. Let the embedded strings be called

$X(u)$ and $Y(v)$ which each player can generate privately on its own. Then, the players run the protocol P on $\big(X(u), Y(v)\big)$. Let J be the intersection set that P returns. Clearly, $\mathrm{DISJ}_n(u,v) = 1$ iff $|J| = |a|$. Finally, note if $(\mathbf{U}, \mathbf{V})$ are generated according to τ, then the mapped strings $\big(\mathbf{X}(\mathbf{U}), \mathbf{Y}(\mathbf{V})\big) \sim (\mathbf{X}, \mathbf{Y} | \mathbf{A} = a)$. Hence, (5) implies that $\mathrm{IC}_\tau(P) \leq \beta' m \leq 2\beta' n$. By setting β' to be a small enough constant, we derive a contradiction to Theorem 4.13. This completes the argument. ■

By using Theorem 4.11, this immediately yields the following theorem, conjectured by Feigenbaum et al. [2010b].

THEOREM 1.3. *For all $n \geq 1$, and any protocol P computing the Set Intersection INTERSEC$_n$ on n bits, the average-case subjective PAR is exponential in n under the uniform distribution:* $\mathrm{avg}_U \, \mathrm{PAR}^{sub}(P) = 2^{\Omega(n)}$.

5. SUMMARY OF PRIVACY MEASURES

Given the plethora of approaches to privacy, it seems useful to compare the differing approximate privacy measures and their motivations.

The first difference in privacy models stems from the question: privacy from whom? Participants can be concerned with privacy from an eavesdropper (as in external PAR and $\mathrm{IC}^{\mathrm{ext}}$) or privacy from other participants (as in internal PAR, IC, and PRIV). In general, the eavesdropper measure provides an upper bound.

$$\mathrm{PAR}^{\mathrm{sub}}(P) \leq \mathrm{PAR}(P)$$
$$\mathrm{avg}_D \, \mathrm{PAR}^{sub}(P) \leq \mathrm{avg}_D \, \mathrm{PAR}^{sub}(P)$$
$$\mathrm{IC} \leq \mathrm{IC}^{\mathrm{ext}}$$

Two-player Vickrey auctions are a special case where objective and subjective PAR are identical; see Lemma 3.6. For most functions this bound is not tight.

Additionally, differences arise from whether the function output is considered a loss of privacy. This is the main difference betweet IC and PRIV; recall Proposition 4.12:

$$\mathrm{PRIV}_D(P) - \log|Z| \leq \mathrm{IC}_D(P) \leq 2 \cdot (\mathrm{PRIV}_D(P) + \log|Z|)$$

The choice of measure here depends largely on the setting, and seems to be one of personal preference. The proposition says that the IC and PRIV measures are essentially the same.

Our proposed modification to the definition of average-case PAR (Definition 4.1) is more natural than the original in Feigenbaum et al. [2010a]. While the two definitions coincide on the uniform distribution of inputs, the original measure exhibits strange behaviors on non-uniform distributions, often not in sync with the semantic notion of measuring privacy loss. Further supporting our "revised" definition is the fact that it allows PAR to be directly related to information-theoretic measures like PRIV and IC. This permits the powerful and well-developed tools of information theory to be brought to bear on questions of privacy.

Recall Theorem 4.11:

$$\mathrm{PRIV}_D(P) \leq \log(\mathrm{avg}_D \, \mathrm{PAR}^{sub}(P)).$$

Thus, information-theoretic lower bounds on PRIV or IC yield lower bounds for average-case PAR. This bound is not tight, however.

Another difference arises from the choice of a worst-case measure (like PAR) or an average-case measure (average-case PAR, IC, and PRIV). Here the semantics of privacy again determine which model is preferable. Although average-case measures

permit achievable privacy guarantees (e.g., our matching bounds for two-player Vickrey auction), an individual participant might care more about the worst-case measure ("what's the most privacy I can lose?").

Further considering the viewpoint of players adds additional wrinkles to the measure of privacy. The bits of input may not be equally private (e.g., the most significant bit of a salary is much more privacy-revealing than the least significant), and which bits or parts of the input are private can depend on the context of the problem as well as the other participants. (Such an approach was proposed in **?**].) PAR seems much more complicated to adjust for such scenarios than the information-theoretic measures.

6. CONCLUSION

These techniques hold the promise of similar length-privacy tradeoffs for other functions. Further, it seems that one can readily extend this work to include randomized and ϵ-error settings. With the restriction of perfect privacy for two-player functions, Kushilevitz [1989] shows that the set of functions with deterministic protocols and the set of functions with randomized protocols are the same. Perhaps there is a similar result for any fixed constant PAR, or perhaps as the PAR requirement is relaxed, the two sets gradually differ.

REFERENCES

Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. 2004. An information statistics approach to data stream and communication complexity. *J. Comput. System Sci.* 68, 4 (June 2004), 702–732. DOI:http://dx.doi.org/10.1016/j.jcss.2003.11.006

Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. 2010. How to compress interactive communication. *Proceedings of the 42nd ACM symposium on Theory of Computing* (2010). DOI:http://dx.doi.org/10.1145/1806689.1806701

Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. 1988. Completeness Theorems for Non-Cryptographic Fault-Tolerant Distributed Computation. *Proceedings of the 20th Annual ACM Symposium on Theory of Computing* (1988), 1–10.

Felix Brandt and Tuomas Sandholm. 2008. On the Existence of Unconditionally Privacy-Preserving Auction Protocols. *ACM Transactions on Information and System Security* 11, 2 (May 2008), 1–21. DOI:http://dx.doi.org/10.1145/1330332.1330338

Mark Braverman. 2011. Interactive information complexity. *Electronic Colloquium on Computational Complexity* 123 (2011).

Amit Chakrabarti, Anthony Wirth, Andrew Yao, and Yaoyun Shi. 2001. Informational complexity and the direct sum problem for simultaneous message complexity. *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science* (2001), 270–278. DOI:http://dx.doi.org/10.1109/SFCS.2001.959901

Benny Chor, Mihály Geréb-Graus, and Eyal Kushilevitz. 1994. On the structure of the privacy hierarchy. *Journal of Cryptology* 7, 1 (1994), 53–60. DOI:http://dx.doi.org/10.1007/BF00195209

Benny Chor and Eyal Kushilevitz. 1989. A Zero-One Law for Boolean Privacy (extended abstract). In *Proceedings of the 21st Annual ACM Symposium on the Theory of Computing*. 62–72.

Marco Comi, Bhaskar Dasgupta, Michael Schapira, and Venkatakumar Srinivasan. 2011. On Communication Protocols That Compute Almost Privately. *Symposium on Algorithmic Game Theory* (2011), 44–56.

Thomas M. Cover and Joy A. Thomas. 1991. *Elements of Information Theory*. Wiley-Interscience, New York. 542 pages.

Joan Feigenbaum, Aaron D Jaggard, and Michael Schapira. 2010a. Approximate Privacy: Foundations and Quantification. *Proceedings of the 11th Conference on Electronic Commerce* (2010), 167–178. DOI:http://dx.doi.org/10.1145/1807342.1807369

Joan Feigenbaum, Aaron D Jaggard, and Michael Schapira. 2010b. Approximate Privacy: PARs for Set Problems. *DIMACS Technical Report 2010-01* (2010), 1–34. http://dimacs.rutgers.edu/~adj/Research/type.html

Hartmut Klauck. 2002. On quantum and approximate privacy. *Proc. STACS* (2002).

Eyal Kushilevitz. 1989. Privacy and communication complexity. *30th Annual Symposium on Foundations of Computer Science* (1989), 416–421. DOI:http://dx.doi.org/10.1109/SFCS.1989.63512

Eyal Kushilevitz and Noam Nisan. 1997. *Communication Complexity*. Cambridge University Press.

Andrew Chi-chih Yao. 1979. Some Complexity Questions Related to Distributive Computing. *Proceedings of the 11th Annual ACM Symposium on Theory of Computing* (1979), 209–213.

Andrew Chi-chih Yao. 1982. Protocols for Secure Computations. *Proceedings of the 23rd Annual Foundations of Computer Science* (1982), 160–164.

Received ; revised ; accepted