

Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at SciVerse ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcsOn communication protocols that compute almost privately[☆]Marco Comi^a, Bhaskar DasGupta^{a,*}, Michael Schapira^b, Venkatakumar Srinivasan^a^a Department of Computer Science, University of Illinois at Chicago, IL 60607, United States^b School of Computer Science and Engineering, The Hebrew University of Jerusalem, 91904 Jerusalem, Israel

ARTICLE INFO

Article history:

Received 11 July 2011

Received in revised form 2 July 2012

Accepted 9 July 2012

Communicated by D.-Z. Du

Keywords:

Approximate privacy

Multi-party communication

Tiling

Binary space partition

ABSTRACT

We further investigate and generalize the approximate privacy model recently introduced by Feigenbaum et al. (2010) [7]. We explore the privacy properties of a natural class of communication protocols that we refer to as “dissection protocols”. Informally, in a dissection protocol the communicating parties are restricted to answering questions of the form “Is your input between the values α and β (under a pre-defined order over the possible inputs)?”. We prove that for a large class of functions, called *tiling functions*, there *always* exists a dissection protocol that provides a *constant average-case privacy approximation ratio* for uniform or “almost uniform” probability distributions over inputs. To establish this result we present an interesting connection between the approximate privacy framework and basic concepts in computational geometry. We show that such a good privacy approximation ratio for tiling functions does *not*, in general, exist in the *worst case*. We also discuss extensions of the basic setup to more than two parties and to non-tiling functions, and provide calculations of privacy approximation ratios for two functions of interest.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Consider the following interaction between two parties, Alice and Bob. Each of the two parties, Alice and Bob, holds a *private* input, x_{bob} and y_{alice} respectively, not known to the other party. The two parties aim to compute a function f of the two private inputs. Alice and Bob alternately query each other to make available a *small* amount of information about their private inputs, e.g., an answer to a range query on their private inputs or a few bits of their private inputs. This process ends when each of them has seen enough information to be able to compute the value of $f(x_{\text{bob}}, y_{\text{alice}})$. The central question that is the focus of this paper is:

Can we design a communication protocol whose execution reveals, to both Alice and Bob, as well as to any eavesdropper, as little information as possible about the other's private input beyond what is necessary to compute the function value?

Note that there are two conflicting constraints: Alice and Bob need to communicate sufficient information for computing the function value, but would prefer not to communicate too much information about their private inputs. This setting can be generalized in an obvious manner to $d > 1$ parties $\text{party}_1, \text{party}_2, \dots, \text{party}_d$ computing a d -ary f by querying the parties in round-robin order, allowing each party to broadcast information about its private input (via a public communication channel).

[☆] A preliminary version of this paper appeared in the 4th Symposium on Algorithmic Game Theory, G. Persiano (Ed.), LNCS 6982, Springer-Verlag, 44–56, 2011.

^{*} Corresponding author. Tel.: +1 312 355 1319; fax: +1 312 413 0024.

E-mail addresses: ingmarco85@gmail.com (M. Comi), dasgupta@cs.uic.edu, bdasgup@uic.edu (B. DasGupta), schapiram@huji.ac.il (M. Schapira), vsrini7@uic.edu (V. Srinivasan).

URL: <http://www.cs.uic.edu/~dasgupta> (B. DasGupta).

Privacy preserving computational models such as the one described above have become an important research area due to the increasingly widespread usage of sensitive data in networked environments, as evidenced by distributed computing applications, game-theoretic settings (e.g., auctions) and more. Over the years computer scientists have explored many *quantifications* of privacy in computation. Much of this research focused on designing *perfectly* privacy-preserving protocols, i.e., protocols whose execution reveals *no* information about the parties' private inputs beyond that implied by the outcome of the computation. Unfortunately, perfect privacy is often either *impossible*, or *infeasibly costly* to achieve (e.g., requiring impractically extensive communication steps). To overcome this, researchers have also investigated various notions of *approximate privacy* [6,7].

In this paper, we adopt the approximate privacy framework of [7] that quantifies approximate privacy via the *privacy approximation ratios* (PARs) of protocols for computing a deterministic function of two private inputs. *Informally*, PAR captures the objective that an observer of the transcript of the entire protocol will not be able to distinguish the real inputs of the two communicating parties from *as large a set as possible* of other inputs. To capture this intuition, [7] makes use of the machinery of communication-complexity theory to provide a geometric and combinatorial interpretation of protocols. [7] formulates both the worst-case and the average-case version of PARs and studies the tradeoff between privacy preservation and communication complexity for several functions of interest.

1.1. Motivations from mechanism designs

An original motivation of this line of research, as explained in details in [7], comes from privacy concerns in auction theory in Economics. A traditionally desired goal of designing auction mechanisms is to ensure that it is *incentive compatible*, i.e., bidders fare best by letting their truthful bids be known. However, more recently, another *complementary* goal that has gained significant attention, especially in the context of online auctions, is to preserve privacy of the bidders, i.e., bidders reveal as little information as necessary to auctioneers for optimal outcomes. To give an example, consider a 2nd-price Vickrey auction of an item via a straightforward protocol in which the price of the item is incrementally increased until the winner is determined. However, the protocol reveals more information than what is *absolutely necessary*, namely the information about the identity of the winner (with revealing his/her bid) together with the bid of the second-highest bidder, and revealing such additional information could put the winner at a disadvantage in the bidding process of a similar item in the future since the auctioneer could set a lower reserve price. In this paper, we consider a generalized version of the setting that captures applications of the above type as well as other applications in multi-party computation.

2. Summary of our contributions

Any investigation of approximate privacy for multi-party computation starts by defining how we quantify approximate privacy. In this paper, we use the combinatorial framework of [7] for quantification of approximate privacy for two parties via PARs and present its natural extension to three or more parties. Often, parties' inputs have a natural ordering, e.g., the private input of a party belongs to some range of integers $\{L, L + 1, \dots, M\}$ (as is the case when computing, say, the maximum or minimum of two inputs). When designing protocols for such environments, a natural restriction is to only allow the protocol to ask each party questions of the form "*Is your input between the values α and β (under this natural order over possible inputs)?*". We refer to this type of protocols as *dissection protocols* and study the privacy properties of this natural class of protocols. We note that the bisection and c -bisection protocols for the millionaires problem and other problems in [7], as well as the bisection auction in [8,9], all fall within this category of protocols. Our findings are summarized below.

Average- and worst-case PARs for tiling functions for two party computation. We first consider a broad class of functions, referred to as the *tiling functions* in the sequel, that encompasses several well-studied functions (e.g., Vickrey's second-price auctions). Informally, a two-variable tiling function is a function whose output space can be viewed as a collection of disjoint combinatorial rectangles in the two-dimensional plane, where the function has the same value within each rectangle. A first natural question for investigation is to classify those tiling functions for which there exists a perfectly privacy-preserving dissection protocol. We observe that for every Boolean tiling functions (i.e., tiling functions which output binary values) *this is indeed the case*. In contrast, for tiling functions with a range of just three values, perfectly privacy-preserving computation is no longer necessarily possible (even when not restricted to dissection protocols).

We next turn our attention to PARs. We prove that for every tiling function there exists a dissection protocol that achieves a constant PAR in the average case (that is, when the parties' private values are drawn from an uniform or *almost* uniform probability distribution). To establish this result, we make use of results on the binary space partitioning problems studied in the computational geometry literature. We complement this positive result for dissection protocols with the following negative result: *there exist tiling functions for which no dissection protocol can achieve a constant PAR in the worst-case*.

Extensions to non-tiling functions and three-party communication. We discuss two extensions of the above results. We explain how our constant average-case PAR result for tiling functions can be extended to a family of "almost" tiling functions. In addition, we consider the case of *more than two* parties. We show that in this setting it is *no longer true* that for every tiling function there exists a dissection protocol that achieves a constant PAR in the average case. Namely, we exhibit a three-dimensional tiling function for which every dissection protocol exhibits *exponential* average- and worst-case PARs, even when an unlimited number of communication steps is allowed.

PARs for the set covering and equality functions. [7] presents bounds on the average-case and the worst-case PARs of the bisection protocol – a special case of dissection protocols – for several functions (Yao's millionaires' problem, Vickrey's second-price auction, and others). We analyze the PARs of the bisection protocol for two well-studied Boolean functions: the set-covering and equality functions; the equality function provides a useful testbed for evaluating privacy preserving protocols [2] [10, Example 1.21] and set-covering type of functions are useful for studying the differences between deterministic and non-deterministic communication complexities [10, Section 10.4]. We show that, for both functions, the bisection protocol *fails to achieve* good PARs in both the average-case and the worst-case.

3. Summary of prior related works

3.1. Privacy-preserving computation

Privacy-preserving computation has been the subject of extensive research and has been approached from information-theoretic [2], cryptographic [4], statistical [11], communication complexity [12,16], statistical database query [6] and other perspectives [10]. Among these, most relevant to our work is the approximate privacy framework of Feigenbaum *et al.* [7] that presents a metric for quantifying privacy preservation building on the work of Chor and Kushilevitz [5] on characterizing perfectly privately computable computation and on the work of Kushilevitz [12] on the communication complexity of perfectly private computation. The bisection, c -bisection and bounded bisection protocols of [7] fall within our category of dissection protocol since we allow the input space of each party to be divided into two subsets of arbitrary size. There are also some other formulations of perfectly and approximately privacy-preserving computation in the literature, but they are inapplicable in our context. For example, the differential privacy model (see [6]) approaches privacy in a different context via adding noise to the result of a database query in such a way as to preserve the privacy of the individual database records but still have the result convey nontrivial information,

3.2. Binary space partition (BSP)

BSPs present a way to implement a *geometric divide-and-conquer* strategy and is an extremely popular approach in numerous applications such as hidden surface removal, ray-tracing, visibility problems, solid geometry, motion planning and spatial databases (*e.g.*, see [15]). However, to the best of our knowledge, a connection between BSPs bounds such as in [13,14,1,3] and approximate privacy has not been explored before.

4. The model and basic definitions

4.1. Two-party approximate privacy model of [7]

We have two parties party_1 and party_2 , each a binary string, x_1 and x_2 respectively, which represents a private value in some set \mathcal{U}^{in} . The common goal of the two parties is to compute the value $f(x_1, x_2)$ of a given public-knowledge two-variable function f . Before a communication protocol P starts, each party_i initializes its “set of maintained inputs” $\mathcal{U}_i^{\text{in}}$ to \mathcal{U}^{in} . In one step of communication, one party transmits a bit indicating in which of two parts of its input space its private input lies. The other party then updates its set of maintained inputs accordingly. The very last information transmitted in the protocol P contains the value of $f(x_1, x_2)$. The final transcript of the protocol (*i.e.*, the entire information exchanged) is denoted by $s(x_1, x_2)$.

Denoting the domain of outputs by \mathcal{U}^{out} , any function $f : \mathcal{U}^{\text{in}} \times \mathcal{U}^{\text{in}} \mapsto \mathcal{U}^{\text{out}}$ can be visualized as a $|\mathcal{U}^{\text{in}}| \times |\mathcal{U}^{\text{in}}|$ matrix with entries from \mathcal{U}^{out} in which the first dimension represents the possible values of party_1 , ordered by some permutation Π_1 , while the second dimension represents the possible values of party_2 , ordered by some permutation Π_2 , and each entry contains the value of f associated with a particular set of inputs from the two parties. This matrix will be denoted by $A_{\Pi_1, \Pi_2}(f)$, or sometimes simply by A .

We now present the following definitions from [10,7]; see Fig. 1 for a geometric illustration.

Definition 1 (Regions, Partitions). A region of A is any subset of entries in A . A partition of A is a collection of disjoint regions in A whose union equals to A .

Definition 2 (Rectangles, Tilings, Refinements). A rectangle in A is a submatrix of A . A tiling of A is a partition of A into rectangles. A tiling T_1 of A is a refinement of another tiling T_2 of A if every rectangle in T_1 is contained in some rectangle in T_2 .

Definition 3 (Monochromatic, Maximal Monochromatic and Ideal Monochromatic Partitions). A region R of A is monochromatic if all entries in R are of the same value. A monochromatic partition of A is a partition all of whose regions are monochromatic. A monochromatic region of A is a maximal monochromatic region if no monochromatic region in A properly contains it. The ideal monochromatic partition of A is made up of the maximal monochromatic regions.

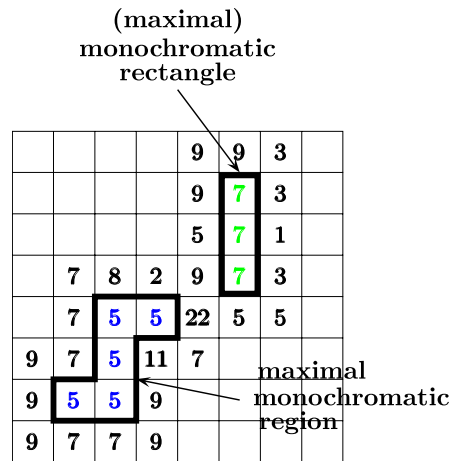


Fig. 1. An illustration of some communication-complexity definitions.

Definition 4 (Perfect Privacy). Protocol P achieves perfect privacy if, for every two sets of inputs (x_1, x_2) and (x'_1, x'_2) such that $f(x_1, x_2) = f(x'_1, x'_2)$, it holds that $s(x_1, x_2) = s(x'_1, x'_2)$. Equivalently, a protocol P for f achieves perfectly privacy if the monochromatic tiling induced by P is the ideal monochromatic partition of $A(f)$.

Definition 5 (Worst Case and Average Case PAR of a Protocol P). Let $R^P(x_1, x_2)$ be the monochromatic rectangle containing the cell $A(x_1, x_2)$ induced by P , $R^I(x_1, x_2)$ be the monochromatic region containing the cell $A(x_1, x_2)$ in the ideal monochromatic partition of A , and \mathcal{D} be a probability distribution over the space of inputs. Then P has a worst-case PAR of α_{worst} and an average case PAR of $\alpha_{\mathcal{D}}$ under distribution \mathcal{D} provided¹

$$\alpha_{\text{worst}} = \max_{(x_1, x_2) \in \mathcal{U}^{\text{in}} \times \mathcal{U}^{\text{in}}} \frac{|R^I(x_1, x_2)|}{|R^P(x_1, x_2)|} \quad \text{and} \quad \alpha_{\mathcal{D}} = \sum_{(x_1, x_2) \in \mathcal{U}^{\text{in}} \times \mathcal{U}^{\text{in}}} \Pr_{\mathcal{D}}[x_1 \& x_2] \frac{|R^I(x_1, x_2)|}{|R^P(x_1, x_2)|}.$$

Definition 6 (PAR for a Function). The worst-case (average-case) PAR for a function f is the minimum, over all protocols P for f , of the worst-case (average-case) PAR of P .

Extension to multi-party computation. In the multi-party setup, we have $d > 2$ parties $\text{party}_1, \text{party}_2, \dots, \text{party}_d$ computing a d -ary function $f : (\mathcal{U}^{\text{in}})^d \mapsto \mathcal{U}^{\text{out}}$. Now, f can be visualized as $|\mathcal{U}^{\text{in}}| \times \dots \times |\mathcal{U}^{\text{in}}|$ matrix $A_{\Pi_1, \dots, \Pi_d}(f)$ (or, sometimes simply by A) with entries from \mathcal{U}^{out} in which the i th dimension represents the possible values of party_i ordered by some permutation Π_i , and each entry of A contains the value of f associated with a particular set of inputs from the d parties. Then, all the previous definitions can be naturally adjusted in the obvious manner, i.e., the input space as a d -dimensional space, each party maintains the input partitions of all other $d - 1$ parties, the transcript of the protocol s is a d -ary function, and rectangles are replaced by d -dimensional hyper-rectangles (Cartesian product of d intervals).

4.2. Dissection protocols and tiling functions for two-party computation

Often in a communication complexity settings the input of each party has a natural ordering, e.g., the set of input of a party from $\{0, 1\}^k$ can represent the numbers $0, 1, 2, \dots, 2^k - 1$ (as is the case when computing the maximum/minimum of two inputs, in the millionaires problem, in second-price auctions, and more). When designing protocols for such environments, a natural restriction is to only allow protocols such that each party asks questions of the form “Is your input between a and b (in this natural order over possible inputs)?”, where $a, b \in \{0, 1\}^k$. Notice that after applying an appropriate permutation to the inputs, such a protocol divides the input space into two (not necessarily equal) halves. Below, we formalize these types of protocols as “dissection protocols”.

Definition 7 (Contiguous Subset of Inputs). Given a permutation Π of $\{0, 1\}^k$, let \prec_{Π} denote the total order over $\{0, 1\}^k$ that Π induces, i.e., $\forall a, b \in \{0, 1\}^k, a \prec_{\Pi} b$ provided b comes after a in Π . Then, $I \subseteq \{0, 1\}^k$ contiguous with respect to Π if $\forall a, b \in I, \forall c \in \{0, 1\}^k: a \prec_{\Pi} c \prec_{\Pi} b \implies c \in I$.

Definition 8 (Dissection Protocol). Given a function $f : \{0, 1\}^k \times \{0, 1\}^k \mapsto \{0, 1\}^t$ and permutations Π_1, Π_2 of $\{0, 1\}^k$, a protocol for f is a dissection protocol with respect to (Π_1, Π_2) if, at each communication step, the maintained subset of inputs of each party i is contiguous with respect to Π_i .

¹ The notation $\Pr_{\mathcal{D}}[\mathcal{E}]$ denotes the probability of an event \mathcal{E} under distribution \mathcal{D} .




Fig. 2. A tiling function with respect to different permutation pairs (Π_1, Π_2) and (Π'_1, Π'_2) inducing different numbers of monochromatic rectangles.




Fig. 3. Tilability depends on Π_1 and Π_2 .

Observe that every protocol P can be regarded as a dissection protocol with respect to some permutations over inputs by simply constructing the permutation so that it is consistent with the way P updates the maintained sets of inputs. However, not every protocol is a dissection protocol with respect to specific permutations. Consider, for example, the case that both Π_1 and Π_2 are the permutation over $\{0, 1\}^k$ that orders the elements from lowest to highest binary values. Observe that a protocol that is a dissection protocol with respect to these permutations cannot ask questions of the form “Is your input odd or even?”, for these questions partition the space of inputs into non-contiguous subsets with respect to (Π_1, Π_2) .

A special case of interest of the dissection protocol is the “bisection type” protocols that have been investigated in the literature in many contexts [7,9].

Definition 9 (Bisection, c-Bisection and Bounded-bisection Protocols). For a constant $c \in [\frac{1}{2}, 1)$, a dissection protocol with respect to the permutations (Π_1, Π_2) is called a c -bisection protocol provided at each communication step each party $_i$ partitions its input space of size z into two halves of size cz and $(1 - c)z$. A bisection protocol is simply a $\frac{1}{2}$ -bisection protocol. For an integer valued function $g(k)$ such that $0 \leq g(k) \leq k$, bounded-bisection $_{g(k)}$ is the protocol that runs a bisection protocol with $g(k)$ bisection operations followed by a protocol (if necessary) in which each party $_i$ repeatedly partitions its input space into two halves one of which is of size exactly one.

We next introduce the concept of tiling functions.

Definition 10 (Tiling and Non-tiling Functions). A function $f : \{0, 1\}^k \times \{0, 1\}^k \mapsto \{0, 1\}^t$ is called a tiling function with respect to two permutations (Π_1, Π_2) of $\{0, 1\}^k$ if the monochromatic regions in $A_{\Pi_1, \Pi_2}(f)$ form a tiling, and the number of monochromatic regions in this tiling is denote by $r_f(\Pi_1, \Pi_2)$. Conversely, f is a non-tiling function if f is not a tiling function with respect to every pair of permutations (Π_1, Π_2) of $\{0, 1\}^k$.

For example, $f(x_1, \dots, x_k, y_1, \dots, y_k) \equiv \sum_{i=1}^k (x_i + y_i) \pmod{2}$ is a tiling function with respect to (Π_1, Π_2) with $r_f(\Pi_1, \Pi_2) = 4$, where each Π_i orders its inputs (z_1, \dots, z_k) in increasing order of $\sum_{i=1}^k z_i \pmod{2}$. Note that a function f that is tiling function with respect to permutations (Π_1, Π_2) may not be a tiling function with respect to a different set of permutations (Π'_1, Π'_2) ; see Fig. 3. Also, a function f can be a tiling function with respect to two distinct permutation pairs (Π_1, Π_2) and (Π'_1, Π'_2) , and the number of monochromatic regions in the two cases differ; see Fig. 2. Thus, indeed we need Π_1 and Π_2 in the definition of tiling functions and r_f .

Extensions to multi-party computation. For the multi-party computation model involving $d > 2$ parties, the d -ary tiling function f has a permutation Π_i of $\{0, 1\}^k$ for each i th argument of f (or, equivalently for each party $_i$). A dissection protocol is generalized to a “round robin” dissection protocol in the following manner. In one “mega” round of communications, parties communicate in a fixed order, say party $_1$, party $_2$, ..., party $_d$, and the mega round is repeated if necessary. Any communication by any party is made available to all the other parties. Thus, each communication of the dissection protocol partitions a d -dimensional space by an appropriate set of $(d - 1)$ -dimensional hyperplanes, where the missing dimension in the hyperplane corresponds to the index of the party communicating.

5. Two-party dissection protocol for tiling functions

5.1. Boolean tiling functions

Lemma 1. Any Boolean tiling function $f : \{0, 1\}^k \times \{0, 1\}^k \mapsto \{0, 1\}$ with respect to some two permutations (Π_1, Π_2) can be computed in a perfectly privacy-preserving manner by a dissection protocol with respect to the same permutations (Π_1, Π_2) .





Fig. 4. This configuration cannot happen in Case 2.



protocol	number of steps of communication	average case PAR for distribution D_u
c -bisection on f_1 ($c \in [1/2, 1)$)	$k / \log_2 \frac{1}{c}$	k
bounded-bisection $_{g(k)}$ on f_2 ($1 \leq g(k) \leq k$)	$g(k) + 2^{k-g(k)} - 1$	$g(k) + 2^{k-g(k)-1} - 1$

Fig. 5. Functions f_1 and f_2 with $r_{f_1}(\Pi_1, \Pi_2) = r_{f_2}(\Pi_1, \Pi_2) = 2$. The bisection-type protocols fail to achieve a good average-case PAR on them.

Proof. For any $m \times n$ Boolean matrix A with rows and columns indexed by $1, 2, \dots, m$ and $1, 2, \dots, n$, respectively, let the notation $A[i_1, i_2, j_1, j_2]$ denote the submatrix of A consisting of rows $i_1, i_1 + 1, \dots, i_2$ and columns $j_1, j_1 + 1, \dots, j_2$. Assume $m, n \geq 2$ and suppose that the zeros and ones in the matrix A form a tiling. We claim that there must exist an index $i \in \{1, 2, \dots, m - 1\}$ such that the partition of A into the two submatrices $A[1, i, 1, n]$ and $A[i + 1, m, 1, n]$ does not split any tile, or that there must exist an index $j \in \{1, 2, \dots, n - 1\}$ such that the partition of A into the two sub-matrices $A[1, m, 1, j]$ and $A[1, m, j + 1, n]$ does not split any tile. This claim, applied recursively on each submatrix of A , will prove Lemma 1.

We prove our claim by induction on n . The basis case of $n = 2$ follows trivially. Suppose that our claim is true for all $n \in \{2, \dots, q\}$ and consider the case of $n = q + 1$.

Case 1: there exists an index $j \in \{1, 2, \dots, q - 1\}$ such that the partition of $A[1, m, 1, q]$ into the two sub-matrices $A[1, m, 1, j]$ and $A[1, m, j + 1, q]$ does not split any tile. Then, the same index j works for $A[1, m, 1, q + 1]$ also.

Case 2: there is no such index j as in Case 1 above, but there exists an index $i \in \{1, 2, \dots, m - 1\}$ such that the partition of $A[1, m, 1, q]$ into the two submatrices $A[1, i, 1, q]$ and $A[i + 1, m, 1, q]$ does not split any tile. Suppose that the index i does split a tile in the partition $A[1, i, 1, q + 1]$ and $A[i + 1, m, 1, q + 1]$ of $A[1, m, 1, q + 1]$. Then, we must have the situation as shown in Fig. 4, which shows that the zeros and ones of $A[1, m, 1, q + 1]$ do not form a tiling. \square

Remark 1. As Fig. 4 shows, the claim of Lemma 1 is false if f outputs three values.

5.2. Average and worst case PAR for non-Boolean tiling functions

Let $f : \{0, 1\}^k \times \{0, 1\}^k \mapsto \{0, 1\}^t$ be a given tiling function with respect to permutations (Π_1, Π_2) . Neither the c -bisection nor the bounded-bisection protocol performs well in terms of average PAR on arbitrary tiling functions; see Fig. 5 for an illustration. In this section, we show that any tiling function f admits a dissection protocol that has a small constant average case PAR. Moreover, we show that this result cannot be extended to the case of worst-case PARs.

5.2.1. Constant average-case PAR for non-Boolean functions

Let D_u denote the uniform distribution over all input pairs. We define the notion of a c -approximate uniform distribution D_u^{c} ; note that $D_u^0 \equiv D_u$.

Definition 11 (*c-Approximate Uniform Distribution*). A c -approximate uniform distribution D_u^{c} is a distribution in which the probabilities of the input pairs are close to that for the uniform distribution as a linear function of c , namely

$$\max_{(\mathbf{x}, \mathbf{y}), (\mathbf{x}', \mathbf{y}') \in \{0, 1\}^k \times \{0, 1\}^k} \left| \Pr_{D_u^c} [\mathbf{x} \& \mathbf{y}] - \Pr_{D_u^c} [\mathbf{x}' \& \mathbf{y}'] \right| \leq c 2^{-2k}.$$

Theorem 1.

(a) A tiling function f with respect to permutations (Π_1, Π_2) admits a dissection protocol P with respect to the same permutations (Π_1, Π_2) using at most $4r_f(\Pi_1, \Pi_2)$ communication steps such that $\alpha_{D_u^c} \leq 4 + 4c$.

(b) For all $0 \leq c < 1$, there exists a tiling function $f : \{0, 1\}^2 \times \{0, 1\}^2 \mapsto \{0, 1\}^4$ such that, for any two permutations (Π_1, Π_2) of $\{0, 1\}^2$, every dissection protocol with respect to (Π_1, Π_2) using any number of communication steps has $\alpha_{D_u^c} \geq \frac{9+c}{8}$.

Proof. Let $\mathcal{S} = \{S_1, S_2, \dots, S_{r_f}\}$ be the set of $r_f = r_f(\Pi_1, \Pi_2)$ ideal monochromatic rectangles in the tiling of f induced by the permutations (Π_1, Π_2) and consider a protocol P that is a dissection protocol with respect to (Π_1, Π_2) . Suppose that the ideal monochromatic rectangle $S_i \in \mathcal{S}$ has y_i elements, and P partitions this rectangle into t_i rectangles $S_{i,1}, \dots, S_{i,t_i}$ having $Z_{i,1}, \dots, Z_{i,t_i}$ elements, respectively. Then, using the definition of α_{D_u} it follows that

$$\begin{aligned} \alpha_{D_u} &= \sum_{(x_1, x_2) \in \mathcal{U} \times \mathcal{U}} \Pr_{D_u} [x_1 \& x_2] \frac{|R^f(x_1, x_2)|}{|R^P(x_1, x_2)|} \\ &= \sum_{i=1}^{r_f} \sum_{j=1}^{t_i} \sum_{(x_1, x_2) \in S_{i,j}} \Pr_{D_u} [x_1 \& x_2] \frac{y_i}{Z_{i,j}} = \sum_{i=1}^{r_f} \sum_{j=1}^{t_i} \frac{y_i}{2^{2k}} = \sum_{i=1}^{r_f} \frac{t_i y_i}{2^{2k}}. \end{aligned}$$

Similarly, it follows that

$$\alpha_{D_u^c} \leq \sum_{i=1}^{r_f} \sum_{j=1}^{t_i} \sum_{(x_1, x_2) \in S_{i,j}} \frac{1+c}{2^{2k}} \times \frac{y_i}{Z_{i,j}} = \sum_{i=1}^{r_f} \sum_{j=1}^{t_i} \frac{(1+c)y_i}{2^{2k}} = \sum_{i=1}^{r_f} \frac{(1+c)t_i y_i}{2^{2k}}.$$

A binary space partition (BSP) for a collection of *disjoint* rectangles in the two-dimensional plane is defined as follows. The plane is divided into two parts by cutting rectangles with a line if necessary. The two resulting parts of the plane are divided recursively in a similar manner; the process continues until at most one fragment of the original rectangles remains in any part of the plane. This division process can be naturally represented as a binary tree (BSP-tree) where a node represents a part of the plane and stores the cut that splits the plane into two parts that its two children represent and each leaf of the BSP-tree represents the final partitioning of the plane by storing at most one fragment of an input rectangle; see Fig. 7 for an illustration. The size of a BSP is the number of leaves in the BSP-tree. The following result is known.

Fact 1 ([3]). Assume that we have a set \mathcal{S} of disjoint axis-parallel rectangles in the plane. Then, there is a BSP of \mathcal{S} such that every rectangle in \mathcal{S} is partitioned into at most 4 rectangles.²

(a) Consider the dissection protocol corresponding to the BSP in Fact 1. Then, using $\max_i \{t_i\} \leq 4$ we get $\alpha_{D_u^c} \leq \sum_{i=1}^{r_f} \frac{4(1+c)y_i}{2^{2k}} = 4(1+c)$. Also, the number of communication steps in this protocol is the height of the BSP-tree, which is at most $4r_f$.

(b) Consider the function f whose ideal monochromatic rectangles are shown in Fig. 6. Any correct protocol for computing f must partition at least one rectangle of two elements, giving

$$\alpha_{D_u^c} \geq 4 \times \left(\frac{1+c}{16} \right) + \frac{7-c}{8} = \frac{9+c}{8}. \quad \square$$

5.2.2. Large worst-case PAR for non-Boolean functions

Can one extend the results of the last section to show that every tiling function admits a dissection protocol that achieves a good PAR even in the worst case? We answer this question in the negative by presenting a tiling function for which every dissection protocol has large worst-case PAR.

Theorem 2. Let $k > 0$ be an even integer. Then, there exists a tiling function $f : \{0, 1\}^k \times \{0, 1\}^k \mapsto \{0, 1\}^3$ with respect to some two permutations (Π_1, Π_2) such that, for any two permutations Π'_1 and Π'_2 of $\{0, 1\}^k$, every dissection protocol for f with respect to (Π'_1, Π'_2) has $\alpha_{\text{worst}} > 2^{k/2} - 1$.

² The stronger bounds by Berman et al. [1] apply to average number of fragments only.

Proof. Recall the example in Fig. 6 that essentially showed that there exists functions that cannot be computed in a perfectly private manner. Our construction of the function f is based on the tiling shown in Fig. 6. We consider the specific permutations Π_1, Π_2 over $\{0, 1\}^k$ that order the elements in $\{0, 1\}^k$ by binary value (from 0 to $2^k - 1$). We now use the construction in Fig. 6 “recursively” to create a tiling of the input space. We first embed $\frac{2^k-2}{2} = 2^{k-1} - 1$ instances of the construction in Fig. 6 recursively within one another, as shown in Fig. 8(a), leaving a 1×1 square at the center. The vertical level i and the horizontal level i rectangles have dimension $1 \times (2^k - (2i - 1))$ and $(2^k - (2i - 1)) \times 1$, respectively, for $i = 1, 2, \dots, 2^{k-1} - 1$. We then partition each of the level 1 rectangle in Fig. 8(a) into two “nearly” equal-sized rectangles as shown in Fig. 8(b). Consider the function f such that the monochromatic rectangles of $A_f(\Pi_1, \Pi_2)$ are the tilings in Fig. 8(b) (f outputs a different outcome for each (minimal) rectangle in the figure). Clearly, f is a tiling function with respect to (Π_1, Π_2) and, moreover, since every rectangle shares a side with no more than 8 rectangles, at most 8 output values of f suffice.

Let Π'_1, Π'_2 be any two arbitrary permutations of $\{0, 1\}^k$ and consider any dissection protocol P with respect to (Π'_1, Π'_2) . Consider the first meaningful step in the execution of P and suppose that this step was executed by party₁ (the case that the step was executed by party₂ is analogous). This step partitions the total input space $\mathcal{S} = \{0, 1, 2, \dots, 2^k - 1\}$ into two nonempty subsets, say $I \subset \mathcal{S}$ and $I' = \mathcal{S} \setminus I$ such that $0 \in I$. Let $0 < i < 2^k - 1$ be the least integer such that $i \in I$ but $i + 1 \notin I$; such an i must exist since both the sets are non-empty. Consider the rectangles A, B, C and D in Fig. 8(b). We have the following cases.

Case 1: $i \leq 2^{k-1} - 2^{k/2}$. Observe that, for every such i , there exists a level $i + 1$ vertical rectangle of size $2^k - 2i - 1$ that is partitioned into two rectangles, one of which is of size exactly 1. Thus, $\alpha_{\text{worst}} \geq 2^k - 2i + 1 > 2^{(k/2)+1} - 1 > 2^{k/2} - 1$.

Case 2: $2^{k-1} - 2^{k/2} < i < 2^{k-1} - 1$. Observe that, for every such value of i , rectangle A , which is of size 2^{k-1} , is partitioned into two rectangles, of which one is of size at most $2^{k/2}$. Thus, in this case $\alpha_{\text{worst}} \geq \frac{2^k-1}{2^{k/2}} > 2^{k/2} - 1$.

Case 3: $2^{k-1} - 1 \leq i \leq 2^{k-1} + 1$. In this case, at least one of the rectangles B, C or D is partitioned into two parts, one of which is of size at most 2 and thus $\alpha_{\text{worst}} \geq \frac{2^{k-1}-1}{2} > 2^{k-2} - \frac{1}{2} > 2^{k/2} - 1$.

Case 4: $2^{k-1} + 1 < i < 2^k - 2^{k/2}$. Similar to Case 2.

Case 5: $i \geq 2^k - 2^{k/2}$. Similar to Case 1. \square

6. Extensions of the basic two-party setup

6.1. Non-tiling functions

A natural extension of the class of tiling functions involves relaxing the constraint that each monochromatic region *must* be a rectangle.

Definition 12 (δ -Tiling Function). A function $f : \{0, 1\}^k \times \{0, 1\}^k \mapsto \{0, 1\}^t$ is called a δ -tiling function with respect to permutations (Π_1, Π_2) of $\{0, 1\}^k$ if each maximal monochromatic region of $A_{\Pi_1, \Pi_2}(f)$ is an union of at most δ disjoint rectangles.

For example, the function whose tiling is as shown in Fig. 4 is a 2-tiling Boolean function.

Proposition 1. For any δ -tiling function f with respect to (Π_1, Π_2) with r maximal monochromatic regions, there is a dissection protocol P with respect to (Π_1, Π_2) using at most $4r\delta$ communication steps such that $\alpha_{\text{Duc}} \leq (4 + 4c)\delta$.

Proof. We use the algorithm of Theorem 1 on the set of at most $r\delta$ rectangles obtained by partitioning each monochromatic region into rectangles. Since each rectangle is partitioned at most 4 times, each maximal monochromatic region of $A_f(\Pi_1, \Pi_2)$ will be partitioned at most 4δ times. \square

6.2. Multi-party computation

How good is the average PAR for a dissection protocol on a d -dimensional tiling function? For a general d , it is non-trivial to compute precise bounds because each party _{i} has her/his own permutation Π_i of the input, the tiles are boxes of full dimension and hyperplanes corresponding to each step of the dissection protocol is of dimension exactly $d - 1$. Nonetheless, we show that the average PAR is very high for dissection protocols even for 3 parties and uniform distribution, thereby suggesting that this quantification of privacy may not provide good bounds for three or more parties.

Theorem 3 (Large Average PAR for Dissection Protocols with 3 Parties). There exists a tiling function $f : \{0, 1\}^k \times \{0, 1\}^k \times \{0, 1\}^k \mapsto \{0, 1\}^{3k}$ such that, for any three permutations Π_1, Π_2, Π_3 of $\{0, 1\}^k$, every dissection protocol with respect to (Π_1, Π_2, Π_3) must have $\alpha_{\text{Duc}} = \Omega(2^k)$.

$\frac{1-c}{16}$	$\frac{1-c}{16}$	$\frac{1-c}{16}$	$\frac{1-c}{16}$
$\frac{1-c}{16}$	$\frac{1+c}{16}$	$\frac{1+c}{16}$	$\frac{1+c}{16}$
$\frac{1-c}{16}$	$\frac{1+c}{16}$	$\frac{1-c}{16}$	$\frac{1+c}{16}$
$\frac{1-c}{16}$	$\frac{1+c}{16}$	$\frac{1+c}{16}$	$\frac{1+c}{16}$

Fig. 6. Example for $\alpha_{D_0^c} \geq \frac{9+c}{8}$. The tiles are shown by thick black lines. The numbers shown at each cell is the associated probability of that input.




Fig. 7. BSP and Bsp-tree.

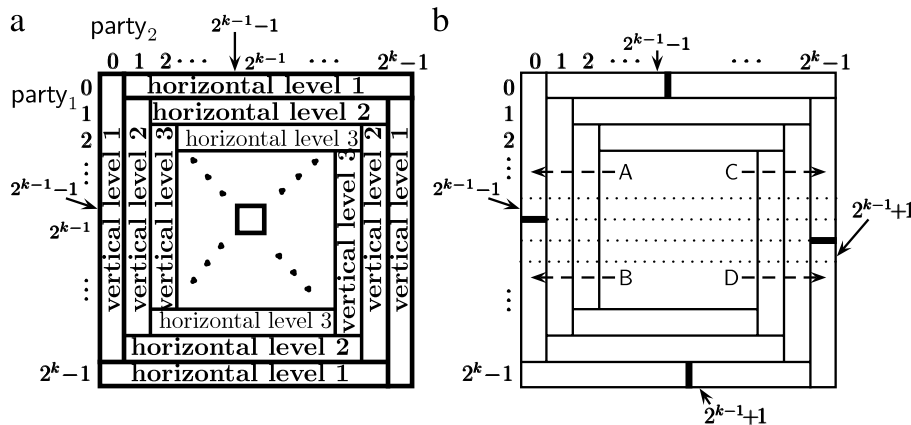


Fig. 8. Illustrations of the arguments in the proof of Theorem 2. The dotted lines in (b) are shown for visual clarity only.

Proof. In the sequel, for convenience we refer to 3-dimensional hyper-rectangles simply by rectangles and refer to the arguments of function f via decimal equivalent of the corresponding binary numbers. The tiling function for this theorem is adopted from an example of the paper by Paterson and Yao [13,14] with appropriate modifications. The three arguments of f are referred to as dimensions 1, 2 and 3, respectively. Define the volume of a rectangle $R = [x_1, x'_1] \times [x_2, x'_2] \times [x_3, x'_3] \subseteq \{0, 1, \dots, 2^k - 1\}^3$ as $\text{Volume}(R) = \max\{0, \prod_{i=1}^3 (x'_i - x_i + 1)\}$. For convenience, let $[*]$ denote the interval $[0, 2^k - 1]$. We provide the tiling for the function f ; see Fig. 9 for a graphical illustration (note that Fig. 9 is not drawn to scale):

- For each dimension, we have a set of $\Theta(2^{2k})$ rectangles; we refer to these rectangles as non-trivial rectangles for this dimension.
 - For dimension 1, these rectangles are of the form $[*] \times [2y, 2y] \times [2z, 2z]$ for every integral value of $0 \leq 2y, 2z < 2^k$.
 - For dimension 2, these rectangles are of the form $[2x, 2x] \times [*] \times [2z + 1, 2z + 1]$ for every integral value of $0 \leq 2x, 2z + 1 < 2^k$.
 - For dimension 3, these rectangles are of the form $[2x + 1, 2x + 1] \times [2y + 1, 2y + 1] \times [*]$ for every integral value of $0 \leq 2x + 1, 2y + 1 < 2^k$.
- The remaining “trivial” rectangles are each of unit volume such that they together cover the remaining input space.

Let $\mathcal{S}_{\text{non-trivial}}$ be the set of all non-trivial rectangles. Observe that:

- Rectangles in $\mathcal{S}_{\text{non-trivial}}$ are mutually disjoint since any two of them do not intersect in at least one dimension.
- Each rectangle in $\mathcal{S}_{\text{non-trivial}}$ has a volume of 2^k and thus the sum of their volumes is $\Theta(2^{3k})$.




Fig. 9. (not drawn to scale) (a) The tiling function in the proof of Theorem 3. The non-trivial rectangles for dimensions 1, 2 and 3 are colored by black, dark gray and light gray, respectively; the trivial rectangles, each having a distinct value, cover the region colored magenta. (b) Rectangles (in light gray) corresponding to a hypothetically first meaningful step of the protocol.

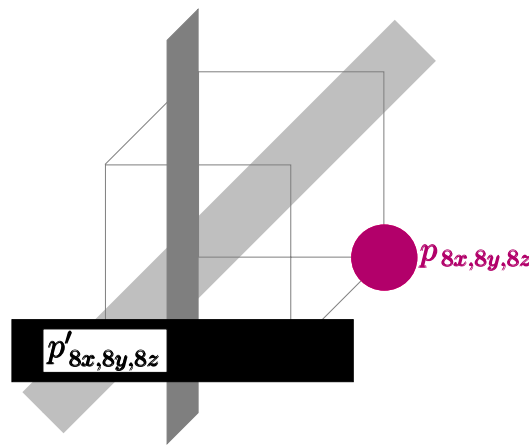


Fig. 10. Separating $p_{8x,8y,8z}$ from $p'_{8x,8y,8z}$.

It now also follows that the number of monochromatic regions is $O(2^{3k})$. Suppose that a dissection protocol partitions, for $i = 1, 2, \dots, |\mathcal{S}_{\text{non-trivial}}|$, the i th non-trivial rectangle $R_i \in \mathcal{S}_{\text{non-trivial}}$ into t_i rectangles, say $R_{i,1}, R_{i,2}, \dots, R_{i,t_i}$. Then,

$$\begin{aligned} \alpha_{D_u} &\stackrel{\text{def}}{=} \sum_{\substack{(x,y,z) \in \\ \{0,1\}^k \times \{0,1\}^k \times \{0,1\}^k}} \Pr_{D_u}[x \& y \& z] \frac{|R^l(x, y, z)|}{|R^p(x, y, z)|} \geq \sum_{i=1}^{|\mathcal{S}_{\text{non-trivial}}|} \sum_{j=1}^{t_i} \sum_{(x,y,z) \in R_{i,j}} \Pr_{D_u}[x \& y \& z] \frac{\text{Volume}(R_i)}{\text{Volume}(R_{i,j})} \\ &= \sum_{i=1}^{|\mathcal{S}_{\text{non-trivial}}|} \sum_{j=1}^{t_i} \frac{2^k}{2^{3k}} = \sum_{i=1}^{|\mathcal{S}_{\text{non-trivial}}|} (t_i / 2^{2k}). \end{aligned}$$

Thus, it suffices to show that $\sum_{i=1}^{|\mathcal{S}_{\text{non-trivial}}|} t_i = \Omega(2^{3k})$. Let \mathcal{Q} be the set of maximal *monochromatic* rectangles that produced the partitioning of the entire protocol. Consider the two entries $p_{x,y,z} = (2x + 1, 2y, 2z + 1)$ and $p'_{x,y,z} = (2x, 2y, 2z)$ (see Fig. 10). Note that $p_{x,y,z}$ belongs to a trivial rectangle since their third, first and second coordinates do not lie within *any* non-trivial rectangle of dimension 1, 2 and 3, respectively, whereas $p'_{x,y,z}$ belongs to the non-trivial rectangle $[*] \times [2 \times (8y), 2 \times (8y)] \times [2 \times (8z), 2 \times (8z)]$ of dimension 1. Thus, $p_{x,y,z}$ and $p'_{x,y,z}$ cannot belong to the same rectangle in \mathcal{Q} . Let $T = \bigcup \{ \{ p_{8x,8y,8z}, p'_{8x,8y,8z} \} \mid 64 < 16x, 16y, 16z < 2^k - 64 \}$. Clearly, $|T| = \Theta(2^{3k})$. For an entry (x_1, x_2, x_3) , let its neighborhood be defined by the ball $\text{Nbr}(x_1, x_2, x_3) = \{ (x'_1, x'_2, x'_3) \mid \forall i : |x_i - x'_i| \leq 4 \}$, i.e., the neighborhood of an entry is the set of all entries (x'_1, x'_2, x'_3) such that each x'_i lies in the range $[x_i - 4, x_i + 4]$ for $i = 1, 2, 3$. Note that $\text{Nbr}(p_{8x,8y,8z}) \cap \text{Nbr}(p_{8x',8y',8z'}) = \emptyset$ provided $(x, y, z) \neq (x', y', z')$. Next, we show that, to ensure that the two entries $p_{8x,8y,8z}$ and $p'_{8x,8y,8z}$ are in two different rectangles in \mathcal{Q} , the protocol must produce an *additional* fragment of one of the non-trivial rectangles in the neighborhood $\text{Nbr}(p_{8x,8y,8z})$; this would directly imply $\sum_i t_i = \Omega(2^{3k})$.

Consider the step of the protocol *before* which $p_{8x,8y,8z}$ and $p'_{8x,8y,8z}$ were contained inside the same rectangle, namely a rectangle Q that includes the rectangle $[16x, 16x + 1] \times [16y, 16y] \times [16z, 16z + 1]$, but after which they are in two different rectangles $Q_1 = [a'_1, b'_1] \times [a'_2, b'_2] \times [a'_3, b'_3]$ and $Q_2 = [a''_1, b''_1] \times [a''_2, b''_2] \times [a''_3, b''_3]$. Remember that both Q_1 and Q_2 must have the same two dimensions and these two dimensions must be the same as the corresponding dimensions of Q . The following cases arise.

Case 1 (split via the first coordinate): $[a'_2, b'_2] = [a''_2, b''_2] \supseteq [16y, 16y]$, $[a'_3, b'_3] = [a''_3, b''_3] \supseteq [16z, 16z + 1]$, $b'_1 = 16x$ and $a''_1 = 16x + 1$. Then, a new fragment of a non-trivial rectangle of dimension 2 is produced at $[16x, 16y, 16z] \in \text{Nbr}(p_{8x,8y,8z})$.

Case 2 (split via the second coordinate): $[a'_1, b'_1] = [a''_1, b''_1] \supseteq [16x, 16x + 1]$ and $[a'_3, b'_3] = [a''_3, b''_3] \supseteq [16z, 16z + 1]$. This case is not possible.

Case 3 (split via the third coordinate): $[a'_1, b'_1] = [a''_1, b''_1] \supseteq [16x, 16x + 1]$, $[a'_2, b'_2] = [a''_2, b''_2] \supseteq [16y, 16y]$, $b'_3 = 16z$ and $a''_3 = 16z + 1$. Then, a new fragment of a non-trivial rectangle of dimension 1 is produced at $[16x, 16y, 16z] \in \text{Nbr}(p_{8x,8y,8z})$. \square

Remark 2. A generalized version of the example in d dimensions can be used to provide a slightly improved lower bound on α_{D_u} for dissection protocols with more than three parties; the bound asymptotically approaches $\Omega(2^{2k})$ for large d .

7. Analysis of the bisection protocol for two functions

In Section 5.1 we showed that any Boolean tiling function can be computed with perfect privacy by a dissection protocol. In [7] the authors provided calculated bounds on α_{worst} and α_{D_u} for the *bisection* protocol, a special case of the general dissection protocol (see Definition 9), on a few functions. In this section, we analyze the bisection protocol [8,9], for two Boolean functions that appear in the literature. As before, D_u denotes the uniform distribution. Letting $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0, 1\}^k$ and $\mathbf{y} = (y_1, y_2, \dots, y_n) \in \{0, 1\}^k$, the functions that we consider are the following:

set-covering: $f_{\wedge, \vee}(\mathbf{x}, \mathbf{y}) = \bigwedge_{i=1}^n (x_i \vee y_i)$. To interpret this as a set-covering function, suppose that the universe \mathcal{U} consists of n elements e_1, e_2, \dots, e_n and the vectors \mathbf{x} and \mathbf{y} encode membership of the elements in two sets S_x and S_y , i.e., x_i (respectively, y_i) is 1 if and only if $e_i \in S_x$ (respectively, $e_i \in S_y$). Then, $f_{\wedge, \vee}(\mathbf{x}, \mathbf{y}) = 1$ if and only if $S_x \cup S_y = \mathcal{U}$.

equality: $f_{=}(\mathbf{x}, \mathbf{y}) = \begin{cases} 1 & \text{if } \forall i : x_i = y_i \\ 0 & \text{otherwise} \end{cases}$. The equality function provides a useful testbed for evaluating privacy preserving protocols, e.g., see [2].

As we already noted in Section 2, both of these functions are studied in the context of evaluating privacy preserving protocols and communication complexity settings [2,10]. A summary of our bounds are as follows.

$f_{\wedge, \vee}$	$\alpha_{\text{worst}} \geq \alpha_{D_u} \geq \left(\frac{3}{2}\right)^{2k}$
$f_{=}$	$\alpha_{D_u} = 2^k - 2 + 2^{1-k}$ $\alpha_{\text{worst}} = 2^{2k-1} - 2^{k-1}$

We will use the formula for α_{D_u} that we derived in the proof of Theorem 1: letting r denote the number of monochromatic regions in an ideal partition of the function if, for $i = 1, 2, \dots, r$, the i th monochromatic region contain $y_i \times 2^{2k}$ elements and the bisection protocol partitions this region into $t_i \geq 1$ rectangles containing z_1, \dots, z_{t_i} elements, respectively, then $\alpha_{D_u} = \sum_{i=1}^r t_i y_i$. In the sequel, by “contribution of a rectangle (of the bisection protocol) to the (average PAR)” we mean the size of the ideal monochromatic region that the rectangle is a part.

7.1. Set covering function

Theorem 4. $\alpha_{D_u} \geq (3/2)^{2k}$.

Proof. We begin by showing the geometry of the tilings for small values of k , which easily generalizes to larger k . The ideal tiling for $f_{\wedge, \vee}$ is shown in Fig. 11(a) for $k = 3$ with the value of the function for each input pair. The sizes of the ideal monochromatic partition are shown in Fig. 11(b) for $k = 1, 2, 3, 4$. The contributions to the average PAR of various inputs after applying the bisection protocol are illustrated in Fig. 12 for $k = 1, 2, 3, 4$. We observe the following:

- The tiles colored *light gray* for the case when $k = 4$ are referred to as the “background tiles”. For $k = 1, 2, 3, 4$ each such tile contributes 3, 9, 27 and 81, respectively, to the average PAR. In general, this contribution is given by 3^k and all these tiles have size 1.
- The contributions of the tiles in the upper-left region of the matrix are given by the sum of the first $2^k - 1$ natural numbers; thus each of these tiles contribute $2^{2k-1} - 2^{k-1}$.
- For any k , observe that the matrix can be decomposed into 4 quadrants; the following observations can be repeated recursively on each resulting quadrant, except for the first quadrant:
 - The first quadrant is a monochromatic region that contributes $2^{2k-1} - 2^{k-1}$ to the average PAR.
 - The fourth quadrant has the same structure as the original matrix, but the contributions for the *non-background* tiles will be related to the case of a matrix with j bits instead of k , where the size of the quadrant is 2^j . For example, notice that the fourth quadrant of a matrix with $k = 4$ is the same as a whole matrix with $k = 3$, except for the “background tiles”, that always contribute for 3^k , with the original value of k .
 - The second and third quadrants are similar to the fourth quadrant case, but in this case the values in the upper-left portion of the quadrants will remain the same as the original matrix, instead of going down as with the fourth quadrant case.

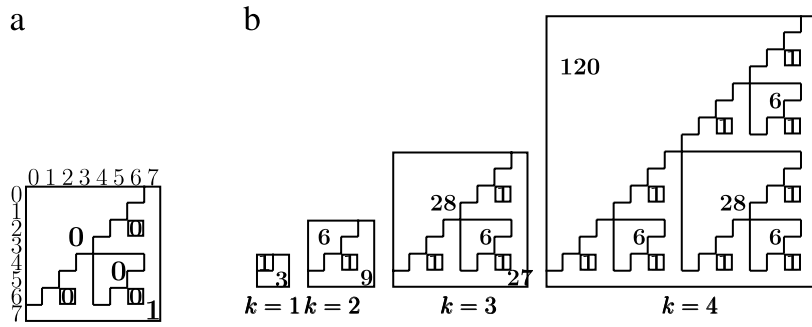


Fig. 11. (a) Ideal monochromatic partition for $f_{\wedge, \vee}$ when $k = 3$. (b) Sizes of ideal monochromatic partition for $f_{\wedge, \vee}$.

1 3 3 3		6 9 1 9 9 9 9 9		28 27 28 27 27 27 27 27 27 27 27 27 27 27									
(a) $k = 1$		(b) $k = 2$		(c) $k = 3$									
												120 81 81 81	
												120 81 81 81	
												120 81 81 81	
												28 81 81 81	
												6 81 81 81	
												28 81 81 81	
												6 81 81 81	
												120 81 81 81	
												6 81 81 81	
												28 81 81 81	
												6 81 81 81	
												120 81 81 81	
												6 81 81 81	
												28 81 81 81	
												6 81 81 81	
												120 81 81 81	
												6 81 81 81	
												28 81 81 81	
												6 81 81 81	
												120 81 81 81	
												6 81 81 81	
												28 81 81 81	
												6 81 81 81	
												120 81 81 81	
												6 81 81 81	
												28 81 81 81	
												6 81 81 81	

Fig. 12. Contribution to PAR for $k = 0, 1, 2, 3, 4$.

Based on these observations, we can obtain a recurrence for the total contribution to the average PAR of all the tiles in a generic matrix. We need the following parameters:

- The number of bits in the original matrix, which we denote by k ;
- The number of bits corresponding to the size of the matrix, or submatrix being considered, which we denote by i ;
- The number of bits to be used in the calculation of the contribution of the upper-left portion of the matrix, or submatrix, being considered; we denote this by j .

The recurrence that computes the total contribution to the PAR of all the tiles in the matrix is:

$$g(i, j, k) = \begin{cases} 3^k, & \text{if } i = 0 \\ 2^{2^{j-1}} - 2^{j-1} + 2g(i-1, j, k) + g(i-1, i-1, k), & \text{otherwise.} \end{cases}$$

The values of i and j are initially set to the value of k . The interpretation of each term in the above recurrence is as follows:

- 3^k is the contribution of each ‘‘background tile’’;
- $2^{2^{j-1}} - 2^{j-1}$ is the contribution of the first quadrant;
- $g(i-1, j, k)$ is the contribution of each one of the second and third quadrants and
- $g(i-1, i-1, k)$ is the contribution of the fourth quadrant.




Fig. 13. (a) Ideal tiling for equality function. (b) The induced tiling by the bisection protocol (shown for $k = 3$). (c) Contribution of each rectangle in protocol-induced tiling where $*$ $\equiv 2^{2k-1} - 2^{k-1}$. The numbers in the figure denote the size of each tile.

Remember that, for a given k , the recurrence equation is initialized with $i = j = k$. Thus, we have:

Case: $k = 0$: $g(k, k, k) = 3^k = 3^{2k}$.

Case: $k > 0$: $g(k, k, k) = g(k - 1, k - 1, k) + 2g(k - 1, k, k) + t(k)$. The second parameter to the function indicates how to generate the $t(k)$ terms; the value of such terms is proportional to that parameter. Thus, for $a \geq b$, $g(k, a, k) \geq g(k, b, k)$. For our lower bound, we can neglect the terms $t(k)$. Thus, we obtain:

$$g(k, k, k) \geq 3g(k - 1, k - 1, k) \geq 3g(k - 2, k - 2, k) \geq \dots \geq 3g(1, 1, k) \geq 3g(0, 0, k).$$

For each step, the value of the first parameter decreased exactly by one unit, so after k iterations the value of the first parameter will be zero. Hence we have $g(k, k, k) \geq 3^k g(0, 0, k)$. Since $g(0, 0, k) = 3^k$ we finally obtain $g(k, k, k) \geq 3^k \times 3^k = 3^{2k}$.

Thus, $\alpha_{D_u} = g(k, k, k)/2^{2k} \geq (3/2)^{2k}$. \square

7.2. Equality function

Theorem 5. $\alpha_{D_u} = 2^k - 2 + 2^{1-k}$ and $\alpha_{\text{worst}} = 2^{2k-1} - 2^{k-1}$.

Proof. An illustration of the ideal partition into monochromatic regions for the equality function is shown in Fig. 13(a). After running the bisection protocol, the induced tiling is (for $k = 3$) is shown in Fig. 13(b). Excluding the diagonal, we have 2 tiles of size 16, 4 tiles of size 4, and 8 tiles of size 1. In general, it is easy to observe that, for each $0 \leq i < k$, we have exactly 2^{k-i} tiles of size 2^{2i} .

The following accounting scheme can be used to simplify calculation. For uniform distribution D_u , α_{D_u} is the sum of the ratio $\frac{|R^I(i,j)|}{|R^P(i,j)|}$ over each element (i, j) in the matrix divided by the number of total elements 2^{2k} in the matrix, where $R^I(i, j)$ and $R^P(i, j)$ is the size of the ideal and protocol-induced tiling that contains the cell (i, j) . Consider a rectangle A of size m in the protocol-induced tiling and suppose that A is contained in a monochromatic region of the ideal partition of size m' . Then, the sum of contributions of the elements of A is $\sum_{i=1}^m \frac{m'}{m} = m'$. Thus, the total contribution of the rectangle A is simply the size of region of the ideal partition containing it.

Fig. 13(c) illustrates the contribution of each rectangle in the protocol-induced tiling to average PAR. We can calculate the total contribution to the average PAR of all the tiles in the matrix, except the diagonal, by multiplying $2^{2k-1} - 2^{k-1}$ by the number of tiles. The number of tiles is given by: $\sum_{i=0}^{k-1} 2^{k-i} = 2^{k+1} - 2$. The total contribution of those tiles is $(2^{k+1} - 2) \times (2^{2k-1} - 2^{k-1}) = 2^{3k} - 2^{2k+1} + 2^k$. The contribution of the diagonal is $\underbrace{1 + 1 + \dots + 1}_{2^k \text{ times}} = 2^k$. Since the

average objective PAR α_{D_u} is the sum of the total contributions divided by the number of cells in the matrix, we have

$$\alpha_{D_u} = \frac{2^{3k} - 2^{2k+1} + 2^k + 2^k}{2^{2k}} = \frac{2^{3k} - 2^{2k+1} + 2^{k+1}}{2^{2k}} = 2^k - 2 + 2^{1-k}.$$

It can be seen from the ideal and protocol tilings that the worst case for PAR is the one in which the ideal tile size is $2^{2k-1} - 2^{k-1}$, and the protocol tile size is 1. Thus $\alpha_{\text{worst}} = 2^{2k-1} - 2^{k-1}$. \square

References

- [1] P. Berman, B. DasGupta, S. Muthukrishnan, On the exact size of the binary space partitioning of sets of isothetic rectangles with applications, SIAM Journal of Discrete Mathematics 15 (2) (2002) 252–267.
- [2] R. Bar-Yehuda, B. Chor, E. Kushilevitz, A. Orlitsky, Privacy, additional information, and communication, IEEE Transactions on Information Theory 39 (1993) 55–65.
- [3] F. d'Amore, P.G. Franciosa, On the optimal binary plane partition for sets of isothetic rectangles, Information Processing Letters 44 (1992) 255–259.
- [4] D. Chaum, C. Crépeau, I. Damgaard, Multiparty, unconditionally secure protocols, in: 22th ACM Symposium on Theory of Computing, 1988, pp. 11–19.
- [5] B. Chor, E. Kushilevitz, A zero-one law for boolean privacy, SIAM Journal of Discrete Mathematics 4 (1991) 36–47.

- [6] C. Dwork, Differential privacy, in: 33rd International Colloquium on Automata, Languages and Programming, 2006, pp. 1–12.
- [7] J. Feigenbaum, A. Jaggard, M. Schapira, Approximate privacy: foundations and quantification, *ACM Conference on Electronic Commerce* (2010) 167–178.
- [8] E. Grigorievaa, P.J.-J. Heringsb, R. Müllera, D. Vermeulena, The communication complexity of private value single-item auctions, *Operations Research Letters* 34 (2006) 491–498.
- [9] E. Grigorievaa, P.J.-J. Heringsb, R. Müllera, D. Vermeulena, The private value single item bisection auction, *Economic Theory* 30 (2007) 107–118.
- [10] E. Kushilevitz, N. Nisan, *Communication Complexity*, Cambridge University Press, 1997.
- [11] D. Kifer, B.-R. Lin, An axiomatic view of statistical privacy and utility, *Journal of Privacy and Confidentiality* (conference version appeared in 2010 ACM SIGMOD/PODS Conference) (in press).
- [12] E. Kushilevitz, Privacy and communication complexity, *SIAM Journal of Discrete Mathematics* 5 (2) (1992) 273–284.
- [13] M. Paterson, F.F. Yao, Efficient binary space partitions for hidden-surface removal and solid modeling, *Discrete and Computational Geometry* 5 (1) (1990) 485–503.
- [14] M. Paterson, F.F. Yao, Optimal binary space partitions for orthogonal objects, *Journal of Algorithms* 13 (1992) 99–113.
- [15] C.D. Tóth, Binary space partitions: recent developments, in: J.E. Goodman, J. Pach, E. Welzl (Eds.), *Combinatorial and Computational Geometry*, in: MSRI Publications, vol. 52, Cambridge University Press, 2005, pp. 529–556.
- [16] A.C. Yao, Some complexity questions related to distributive computing, in: 11th ACM Symposium on Theory of Computing, 1979, pp. 209–213.