A Differentially Private Mechanism of Optimal Utility for a Region of Priors[★]

Ehab ElSalamouny^{1,2}, Konstantinos Chatzikokolakis¹, and Catuscia Palamidessi¹

¹ INRIA and LIX, Ecole Polytechnique, France

² Faculty of Computer and Information Science, Suez Canal University, Egypt

Abstract. The notion of differential privacy has emerged in the area of statistical databases as a measure of protection of the participants' sensitive information, which can be compromised by selected queries. Differential privacy is usually achieved by using mechanisms that add random noise to the query answer. Thus, privacy is obtained at the cost of reducing the accuracy, and therefore the *utility*, of the answer. Since the utility depends on the user's side information, commonly modelled as a prior distribution, a natural goal is to design mechanisms that are optimal *for every prior*. However, it has been shown that such mechanisms *do not exist* for any query other than (essentially) counting queries ([1]).

Given the above negative result, in this paper we consider the problem of identifying *a restricted class of priors* for which an optimal mechanism *does exist*. Given an arbitrary query and a privacy parameter, we geometrically characterise a special region of priors as a convex polytope in the priors space. We then derive upper bounds for utility as well as for min-entropy leakage for the priors in this region. Finally we define what we call the *tight-constraints mechanism* and we discuss the conditions for its existence. This mechanism reaches the bounds for all the priors of the region, and thus it is optimal on the whole region.

1 Introduction

Statistical databases are commonly used to provide aggregate information about the individuals of a certain population, to attain a social benefit. In general, certain data of the participants in the database may be confidential, and we should not allow queries that can reveal them. On the other hand we would like to allow global queries, like, for instance, the average salary of the inhabitants of a certain region, the percentage of individuals having a certain disease, or the cities with the highest rates of crime. This kind of information can be extremely useful for e.g. financial planning, medical research, and anti-crime measures.

Unfortunately, even though these kinds of queries do not refer directly to the individual data, they still represent a major threat to the privacy of the participants in the databases. To illustrate the problem, consider a database whose records contain personal data, among which the salary, regarded as confidential. Suppose we are allowed to query the number of participants and their average salary. Then, by querying the

[★] This work is partially funded by the Inria large scale initiative CAPPRIS, the EU FP7 grant no. 295261 (MEALS), and the project ANR-09-BLAN-0169-01 (PANDA).

database before and after the insertion of a new record “Bob”, we can easily infer, by an easy calculation, the exact salary of Bob.

A successful approach to solve the above problem is to report to the user an approximate answer instead of the exact one. The approximate answer is produced by adding controlled *random noise* to the exact answer. The overall procedure, representing the sanitized query, is a (probabilistic) *mechanism* $\mathcal{K}$ which takes as input the database v and reports to the user an output o in some domain $\mathcal{O}$, according to some probabilistic distribution. Intuitively, the uncertainty introduced at the level of the global answer induces uncertainty about the value of the individual data in the database, thus making it difficult for an attacker to guess such value. However it is crucial to know *exactly* what kind of protection is achieved this way. *Differential privacy*, introduced by Dwork ([2–5]), is a formalization of the privacy property that can be guaranteed by such mechanism. It is a quantitative notion, in the sense that it depends on a parameter ϵ representing the provided level of privacy.

Following common lines (e.g. [6–8]), in this paper we assume that the mechanism $\mathcal{K}$ is *oblivious* with respect to the given query f . Namely, its output depends only on the exact query result and not on the underlying database. Furthermore, we consider only the case in which the domains of the answers (exact and reported) are finite. Under these assumptions, the mechanism $\mathcal{K}$ is determined by an underlying stochastic *noise matrix* X whose generic element x_{io} is the conditional probability of reporting the answer o when the exact query answer is i .

Besides guaranteeing differential privacy, a mechanism should of course provide an answer which is still “useful” enough to the user asking the query. This second goal is measured in terms of *utility*, which represents the average gain that a rational user obtains from the reported answer. More precisely, on the basis of the reported answer o the user can make a guess k (remapping) about the exact hidden query result i . His gain $g(i, k)$ is established by a given function g . The utility is then defined as *the expected gain under the best possible remapping*. While the gain function can take various forms, in this paper we restrict our analysis to the *binary* gain function, which evaluates to 1 when the user’s guess is the same as the query result ($k = i$) and evaluates to 0 otherwise.

The utility of a mechanism depends on the side-information which the user may have about the database. This knowledge induces a probability distribution, called ‘prior’, over the possible query results. Suppose for example that a user “Alice” knows that all people in the database have a salary of at least 20K €. Thus Alice expects the average of the salaries to be at least 20K €. This is reflected on Alice’s prior over the average-salary query results: the total probability mass is distributed on the range of values $\geq 20K$, while it is 0 on lower values. Given this prior, a mechanism X producing only outputs $\geq 20K$ is intuitively more useful to Alice than another one generating also values $< 20K$, which are less informative for Alice.

The *optimal* mechanism for a given prior and level of privacy ϵ is defined as the mechanism which maximises the utility function, while satisfying ϵ -differential privacy. Naturally, we do not want to change the mechanism depending on the user, so we would like to devise mechanisms which are *universally optimal*, i.e. optimal for *any* prior. A famous result by Gosh et al. [6] states that this is possible for the so-called *counting queries*, which are queries concerned with questions of the kind “how many

records in the database have the property $\mathcal{P}$ ” (for some $\mathcal{P}$). In [6] it was proved that the *truncated geometric mechanism* is optimal, for this type of queries, for all priors. Of course the question immediately arises whether we can obtain a similar result for other queries as well. Unfortunately Brenner and Nissim answered this question negatively, by showing that for any query other than (essentially) counting queries a universally optimal mechanism does not exist [1]. However, one can still hope that, also for other queries, by restricting the class of users (i.e. the domain of priors), one could find mechanisms that are optimal for all the users of the class. This is exactly the objective of the present paper: given a query, we aim at identifying a mechanism, and a class of users, for whom that same mechanism provides ϵ -differential privacy and maximal utility at the same time.

Given an arbitrary query and a privacy level $\epsilon > 0$, we call ϵ -*regular* the priors, for which, the probabilities of two adjacent answers (i.e. answers obtained from databases that differ for only one record) are not very different (their ratio is bounded by e^ϵ). At the same time, they may assign significantly different probabilities to “distant” answers. As an example of such prior, consider a researcher “Alice” in a medical school who is interested in the incidence of a certain disease in a statistical medical database containing 1000 records. (Each record represents a person and contains a field saying whether or not the person is infected.) Assume that Alice’s side knowledge lets her to expect that the percentage of infected people is likely to be, say, between 1% and 2%, while it is highly unlikely to be higher than 5%. Also, assume that Alice does not have “sharp” enough information to assign significantly different probabilities to adjacent answers, e.g. 1.5% (15 people affected) and 1.6% (16 people affected). It is precisely this kind of users that we target in this paper: we will see that, under certain conditions, we can design a mechanism which maximises the utility for all of them.

A related issue that we consider in this paper is the amount of information leaked by a mechanism, from the point of view of the so-called *quantitative information flow* framework. There have been various proposals for quantifying the information flow; we consider here the *information-theoretic approach*, in which the system (in this case the mechanism) is regarded as a *noisy channel*, and the leakage is defined as the difference between the a priori *entropy* of the input (the secret – in this case the database entries), and the a posteriori one, after revealing the output (in this case the reported answer). Depending on the notion of entropy adopted one can model different kinds of adversaries [9]. In particular, Shannon entropy (used, for instance, in [10–13]) is suitable for adversaries who can probe the secret repeatedly, while Rényi min-entropy (used, for instance, in [14, 15]) is suitable for one-try attacks. In both cases, the main difference with differential privacy is that the information-theoretic approaches measure the *expected* threat to confidentiality (i.e. the average amount of leakage, where each leak is weighted by its probability to occur), while differential privacy considers catastrophic any disclosure of confidential information, no matter how unlikely it is.

Computing and bounding the information leakage has been pursued in several papers, we mention for instance [16, 17]. Recently, researchers have investigated the relation between differential privacy and information leakage [18–20, 8], and in particular it has been proved in [20] that differential privacy induces a bound on the min-entropy

leakage, which is met by a certain mechanism for the uniform prior (for which min-entropy leakage is always maximum). In this paper, we extend the above result so to provide a more accurate bound for any fixed ϵ -regular prior distribution. More precisely, we provide a bound to the leakage specific to the prior and that can be met, under a certain condition, by a suitable mechanism. It is worth noting that this mechanism is defined similarly to the one that is optimal for the ϵ -regular priors. In fact, min-entropy leakage and utility are strongly related: the main difference is what we regard as the input of the channel. For the former is the database, for the latter the exact answer to the query. Correspondingly, min-entropy leakage measures the correlation between the reported answer and the database entries, while utility measures the correlation between the reported answer and the exact answer.

Contribution

- We identify, for an arbitrary query and a privacy parameter ϵ , the class of the ϵ -regular prior distributions on the exact answers. The interest of this class is that for each prior in it we are able to provide a specific upper bound to the utility of any ϵ -differentially-private mechanism. We characterise this class as a geometric region, and we study its properties.
- We describe an ϵ -differentially-private mechanism, called “tight-constraints mechanism”, which meets those upper bounds for every ϵ -regular prior, and is therefore universally optimal in this region. We provide necessary and sufficient conditions for the existence of such mechanism, and an effective method to test the conditions and to construct the mechanism.
- Switching view, and considering the correlation between the databases and the reported answers (instead than between the exact and reported answers) we recast the above definitions and results in terms of quantitative information flow. The outcome is that we are able to improve the upper bounds for the min-entropy leakage of an ϵ -differentially-private mechanism, for all the ϵ -regular prior distributions on the databases. A construction similar to the one in previous point yields the tight-constraints mechanism which reaches those upper bounds.

Plan of the paper. In the next section we recall the basic definitions of differential privacy and utility. Section 3 introduces the notion of ϵ -regular prior, investigates the properties of these priors, and gives a geometric characterisation of their region. Section 4 shows that for all ϵ -regular priors on the exact answers (resp. databases), ϵ -differential privacy induces an upper bound on the utility (resp. on the min-entropy leakage). Section 5 identifies a mechanism which reaches the above bounds for every ϵ -regular prior, and that is therefore the universally optimal mechanism (resp. the maximally leaking mechanism) in the region. Section 6 illustrates our methodology and results using the example of the sum queries. Section 7 concludes and proposes some directions for future research.

For reasons of space we have omitted several proofs. The interested reader can find them in the report version of this paper [21].

2 Preliminaries

2.1 Differential Privacy

The notion of ϵ -differential privacy, introduced by Dwork in [2], imposes constraints on data reporting mechanisms so that the user is unable to distinguish, from an output, between two databases differing only for one record. This indistinguishability property represents a protection for the individual corresponding to that record. In the following, the mechanism is represented as a probabilistic function $\mathcal{K}$ from the set of possible databases $\mathcal{V}$ to the set of possible reported outputs $\mathcal{O}$. The relation of ‘differing only for one record’ for two databases v and v' is represented by the *adjacency* relation and written as $v \sim v'$.

Definition 1 (Differential privacy [2]). *A probabilistic mechanism $\mathcal{K}$ from $\mathcal{V}$ to $\mathcal{O}$ satisfies ϵ -differential privacy if for all pairs $v, v' \in \mathcal{V}$, with $v \sim v'$, and all $S \subseteq \mathcal{O}$, it holds that*

$$P(\mathcal{K}(v) \in S) \leq e^\epsilon P(\mathcal{K}(v') \in S).$$

Note that the indistinguishability property is independent from the a priori knowledge the user may have about the database.

Consider a query $f : \mathcal{V} \rightarrow \mathcal{R}_f$, where $\mathcal{R}_f$ is the set of the query results. Then a mechanism $\mathcal{K}$ is said to be *oblivious* if for every database $v \in \mathcal{V}$, the output of the mechanism, $\mathcal{K}(v)$, depends only on $f(v)$, the result of applying the query to the database v , regardless of v itself. More formally,

Definition 2 ([1]). *Let $f : \mathcal{V} \rightarrow \mathcal{R}_f$ be a query. A mechanism $\mathcal{K} : \mathcal{V} \rightarrow \mathcal{O}$ is oblivious if there exists a randomised function $\mathcal{M} : \mathcal{R}_f \rightarrow \mathcal{O}$ such that, for all $v \in \mathcal{V}$, and all $S \subseteq \mathcal{O}$, it holds that*

$$P(\mathcal{K}(v) \in S) = P(\mathcal{M}(f(v)) \in S).$$

According to the above definition, any oblivious mechanism $\mathcal{K}$ can be seen as a cascade of two functions: the deterministic query f and a randomised function $\mathcal{M}$. The role of $\mathcal{M}$ is to add random noise to the exact query result $f(v)$ and produce a ‘noisy’ output $o \in \mathcal{O}$ to the user. The privacy guarantees are therefore provided by the function $\mathcal{M}$ which we implement by a stochastic matrix $X = (x_{io})$, called the *noise matrix*. The rows of X are indexed by the elements of $\mathcal{R}_f$ and the columns are indexed by the elements of $\mathcal{O}$. With this representation x_{io} is the probability of giving the output o when the exact query result is i . In this paper, we consider only oblivious mechanisms and therefore our results concern the design of the noise matrix X . Similarly, the query function f and the mechanism can be represented as matrices and hence it holds by Def. 2 that $\mathcal{K} = fX$.

Given a query f , The adjacency relation on databases $\mathcal{V}$ induces another adjacency relation on the set of query results $\mathcal{R}_f$ as follows.

Definition 3 (Adjacent query results). *Given a query function f with a range $\mathcal{R}_f$, two different results $i, h \in \mathcal{R}_f$ are said to be ‘adjacent’, and written as $i \sim_f h$, if and only if there exists two databases v, v' such that $f(v) = i$ and $f(v') = h$, and $v \sim v'$.*

Fig. 1. Examples for the graph structures of different queries

Informally, $i, h \in \mathcal{R}_f$ are adjacent if they discriminate between two adjacent databases. Using the introduced notion of adjacency between query results, a graph structure can be used to model these results along with their adjacency relationship. More precisely, the set of nodes in this graph represents the set of query results $\mathcal{R}_f$, while edges represent the adjacency relationship among them. It is worth noting that this graph structure of queries have been used also in [8, 1] to analyse the differentially private mechanisms. Figure 1 shows examples of the graph structures of different queries. In these examples $\text{count}(v, p)$ refers to a counting query which returns the number of records in the database v which satisfy a certain property p . Other queries in the figure are expressed using the count function.

Let $\mathcal{K}$ be an oblivious mechanism for which X is the underlying noise matrix. It is intuitive to see that satisfying the indistinguishability between adjacent databases (i.e. satisfying differential privacy) corresponds to the satisfying indistinguishability (by means of X) between adjacent query results. Formally,

Lemma 1. *Given a noise matrix X , An oblivious mechanism $\mathcal{K}$ satisfies ϵ -differential privacy if and only if for all query results i, h where $i \sim_f h$ and all outputs $o \in \mathcal{O}$, it holds that $x_{io} \leq e^\epsilon x_{ho}$.*

Note that Lemma 1 provides an equivalent characterisation for differential privacy in terms of adjacent query results rather than adjacent databases.

With the graph structure of a query, the ‘distance’ between two query results i, h , denoted by $d(i, h)$ is defined as the shortest graph distance between i and h . Using this distance measure, differential privacy constraints can be further lifted from conditions on pairs of adjacent query results (Lemma 1) to a general condition on any pair of query results according to the following proposition.

Proposition 1. *Given a noise matrix X , the oblivious mechanism $\mathcal{K}$ satisfies ϵ -differential privacy if and only if for all query results i, h and all outputs $o \in \mathcal{O}$, it holds that $x_{io} \leq e^{\epsilon d(i, h)} x_{ho}$.*

That is, the ratio between the probability of reporting an answer o given that the query result is r and the probability of reporting the same output o given that the query result

is h does not exceed $e^{\epsilon d(i,h)}$. We call the noise matrix that satisfies this condition ϵ -differentially private. Note that, while Lemma 1 describes differential privacy in terms of only adjacent query results, the equivalent characterisation given by Proposition 1 specifies the privacy constraints imposed on any pair of results (whether or not they are adjacent to each other). This feature abstracts our analysis to arbitrary pairs of graph nodes rather than reasoning about only adjacent ones.

2.2 Utility Model

For an oblivious mechanism $\mathcal{K}$, the objective of the underlying noise matrix X is to guarantee the differential privacy of the database, while providing the user with ‘useful’ information about the true query result. That is to satisfy a trade-off between the privacy and utility. For quantifying the utility of $\mathcal{K}$ we follow the model adopted in [6]. Given a query f , let $i \in \mathcal{R}_f$ be the result of executing f on some database. After processing i by the noise matrix X , let o be the reported output to the user. In practice, the user may use the output o , to ‘guess’ the value of the real query result. Therefore she may apply a *remap* (or guess) function which maps the mechanism output o to a guess $k \in \mathcal{R}_f$ for the exact query answer. The remap function (or simply ‘remap’) can be described as a stochastic matrix R , where its entry r_{ok} is the probability of guessing k when the observed mechanism output is o . With this representation, it can be easily seen that the probabilities of the user’s guesses given individual query results are described by the matrix product XR . We say here that X is remapped to XR by the remap R . Note that this remapping procedure models the post-processing done by the user for the mechanism output o . Now, with the user’s guessed value k , a real-valued *gain function* $g : (\mathcal{R}_f \times \mathcal{R}_f) \rightarrow \mathbb{R}$ quantifies how informative k is compared to the real result i .

The *utility* of a given mechanism to the user is described as the expected value of the gain function g . The evaluation of this expected value depends on the a priori probability distribution π over the real query results, which models the side knowledge of the user about the database. The utility of the mechanism depends therefore on the definition of the gain function g , the mechanism’s underlying noise matrix X , the user’s remap R , and also the probability distribution π over the real query results.

One choice for the gain function is the binary gain defined as $g_b(i, j) = 1$ iff $i = j$ and 0 otherwise. The binary gain function formalises the requirement of a user to guess the exact query result using the mechanism output. In the current work we restrict our analysis to this gain function. An important feature of this function, is that it is applicable to the ranges of various queries including numerical and non-numerical one. Moreover, it will be shown that this gain function is strongly connected to the information theoretic notions of conditional entropy and information leakage. Hence, our results about the utility of private mechanism imply corresponding results regarding quantifying information leaked by these mechanisms. These results go inline with a recent trend of research aiming at quantifying information leaked by security protocols, and privacy mechanisms specifically (see e.g. [16, 17, 8, 18]). We leave considering other gain functions to future work.

Now, for formulating the utility we represent the a priori probability distribution (called the ‘prior’) over the real query results by a row vector π , indexed by $\mathcal{R}_f$, where

π_i is the probability that the query in hand yields the result i . The prior is therefore relative to the user and depends on her knowledge. With a generic gain function g , the utility of a mechanism for a prior π using the remap R is denoted by $\mathcal{U}(X, \pi, R)$, and defined as follows.

$$\mathcal{U}(X, \pi, R) = \mathbf{E}[g(i, k)] = \sum_{i,k} \pi_i (X R)_{ik} g(i, k), \quad (1)$$

where X is the noise matrix of the given mechanism. In our case, where the binary gain function g_b is used, the utility reduces to a convex combination of the diagonal elements of $X R$ as follows.

$$\mathcal{U}(X, \pi, R) = \sum_i \pi_i (X R)_{ii}. \quad (2)$$

Accordingly, for a given prior π , an oblivious ϵ -differentially private mechanism, with a noise matrix X , is said to be *optimal* if and only if there is a remap R such that the above function is maximised over all ϵ -differentially private mechanisms and all remaps¹. As exemplified in the introduction, the optimality of a mechanism depends, in general, on the prior (user); that is a mechanism can be optimal for a prior while it is not for another one. It has been proved by [1] that for arbitrary queries (except the counting ones), there is no such a mechanism that is optimal for all priors simultaneously. Nevertheless, we identify in the following section a region of priors, where it is possible to find a single mechanism which is optimal to all of them.

3 ϵ -Regular Priors

In this section we describe a region of priors, called ‘ ϵ -regular’. These priors are determined by the given query f and privacy parameter ϵ . In our way to specify these priors, we first represent the ϵ -differential privacy constraints in a matrix form. By Proposition 1, observe that each ϵ -differential privacy constraint imposed on a noise matrix X can be written as $x_{io}/x_{ho} \geq e^{-\epsilon d(i,h)}$. Since the lower bound $e^{-\epsilon d(i,h)}$ depends only on i, h , all constraints can be described altogether by a square matrix Φ formed by such lower bounds. We refer to this matrix as the *privacy-constraints* matrix. Note that the rows, and also columns of Φ are indexed by the elements of $\mathcal{R}_f$, the set of query results.

Definition 4 (privacy-constraints matrix). *The privacy-constraints matrix Φ of a query f with a range $\mathcal{R}_f$, and a privacy parameter $\epsilon > 0$ is a square matrix, indexed by $\mathcal{R}_f \times \mathcal{R}_f$, where $\phi_{ih} = e^{-\epsilon d(i,h)}$ for all $i, h \in \mathcal{R}_f$.*

Note that Φ is symmetric ($\phi_{ih} = \phi_{hi}$) due to the symmetry of the distance function $d(i, h)$. Observe that when $\epsilon \rightarrow \infty$, i.e. exclude privacy at all, Φ converges to the identity matrix where each diagonal entry is 1 and other entries are zeros. In terms of the privacy-constraints matrix of a query and ϵ , we define now the ϵ -regular priors as follows (note that we use $y \geq 0$ to denote $\forall i : y_i \geq 0$).

Definition 5 (ϵ -regular prior). *For a given query f and a privacy parameter $\epsilon > 0$, a prior π is called ϵ -regular iff there exists a row vector $y \geq 0$ such that $\pi = y \Phi$.*

¹ Note that there may exist many optimal mechanism for a given prior.

In the following we describe the common properties of these priors and also give a geometric characterisation for their region comparing it to the whole prior space. As the first observation, note that, as privacy is excluded ($\epsilon \rightarrow \infty$), this region converges to the entire prior space. This is because Φ approaches the identity matrix where the vector $\mathbf{y}$ exists for each prior.

An important property of any ϵ -regular prior is that the ratio between any two of its entries π_i, π_j is always bound as follows, depending on ϵ and the distance $d(i, j)$. Because of this property, such a prior is called ϵ -regular.

Proposition 2. *Consider a query f and $\epsilon > 0$. Then for any ϵ -regular prior π , it holds for all $i, j \in \mathcal{R}_f$: $\pi_i/\pi_j \leq e^{\epsilon d(i,j)}$.*

While the above property restricts the ratio between probabilities of adjacent query results, this restriction, in practice, holds for a large class of users who have no sharp information suggesting discrimination between adjacent results. This class is exemplified in the introduction. Note that the above property is not equivalent to Definition 5. Namely, it is not true that all priors having such a property are ϵ -regular.

A consequence of the above proposition is that for any ϵ -regular prior π , the probability π_i associated with any query result i is restricted by upper and lower bounds as follows.

Proposition 3. *Consider a query f and $\epsilon > 0$. Then for any ϵ -regular prior π , it holds for all $i \in \mathcal{R}_f$ that*

$$1 / \sum_{j \in \mathcal{R}_f} e^{\epsilon d(i,j)} \leq \pi_i \leq 1 / \sum_{j \in \mathcal{R}_f} e^{-\epsilon d(i,j)}.$$

One implication is that any ϵ -regular prior must have full support, that is $\pi_i > 0$ for all $i \in \mathcal{R}_f$.

In the following we go further and describe the region of ϵ -regular priors as a region of points in the prior space, where each point represents a member in this region. For doing so, we identify by the following definition a set of priors which describe the ‘corner points’ or vertices of the region.

Definition 6 (corner priors). *Given a query f and a privacy parameter $\epsilon > 0$, then for each query result $i \in \mathcal{R}_f$, a corresponding corner prior, denoted by $\mathbf{c}^i$, is defined as*

$$c_j^i = \frac{\phi_{ij}}{\sum_{k \in \mathcal{R}_f} \phi_{ik}} \quad \forall j \in \mathcal{R}_f.$$

Note that the above definition is sound, i.e. $\mathbf{c}^i$ is a probability distribution. By the above definition, for a given query with the domain $\mathcal{R}_f$ of results, the region of ϵ -regular priors has $|\mathcal{R}_f|$ corner priors. Each one corresponds to a query result $i \in \mathcal{R}_f$. Note that each corner prior $\mathbf{c}^i$ is maximally biased (relative to the region) to the query result i ; that is the entry c_j^i meets its maximum value given in Proposition 3. It can be seen that each corner prior is ϵ -regular. Namely for any corner $\mathbf{c}^i$, define the vector $\mathbf{y}$ as $y_i = 1 / \sum_{k \in \mathcal{R}_f} \phi_{ik}$ and $y_j = 0$ for all $j \neq i$; thus it holds that $\mathbf{c}^i = \mathbf{y} \Phi$.

The region of the ϵ -regular priors can be characterised in terms of the corner priors. More precisely, this region consists of all priors that can be composed as a convex combination of the corner priors.

Fig. 2. Regions of ϵ -regular priors for queries described in Example 1

Proposition 4 (convexity). *For a given a query f and privacy parameter $\epsilon > 0$, a prior π is ϵ -regular iff there exist real numbers $\gamma_i \geq 0$, $i \in \mathcal{R}_f$ such that*

$$\pi = \sum_{i \in \mathcal{R}_f} \gamma_i \mathbf{c}^i.$$

It is easy to see that it must hold that $\sum_{i \in \mathcal{R}_f} \gamma_i = 1$ for any ϵ -regular prior. This is obtained by summing the components of the π as follows.

$$\sum_{j \in \mathcal{R}_f} \pi_j = \sum_i \gamma_i \sum_j \mathbf{c}_j^i \quad \text{and} \quad \sum_j \pi_j = 1, \forall \pi.$$

From Proposition 4 and the above observation, the region of ϵ -regular priors is a convex set, where each point (prior) in this region is a convex combination of the corner priors. This region is therefore geometrically regarded as a convex polytope in the prior space. Since the corner points always exists, this region is never empty.

For a prior π in this region, the coefficients γ_i model the ‘proximity’ of π to each corner prior $\mathbf{c}^i$. Observe that $0 \leq \gamma_i \leq 1$, and $\gamma_i = 1$ iff $\pi = \mathbf{c}^i$. We demonstrate this geometric interpretation using the following examples.

Example 1. Priors having 3 entries can be represented as points in the 3-dimensional euclidean space. These priors correspond to queries whose graph structures contain 3 nodes. These nodes can be arranged in either a sequence or a cycle, corresponding to queries f_1 and f_2 respectively shown in Figure 1, with $n = 2$ in both cases. Figure 2 shows - for each of these queries - the region of ϵ -regular priors. The corner priors of each region are represented by points $\mathbf{c}^1, \mathbf{c}^2, \mathbf{c}^3$. For each query in Fig. 2, we depict the regions for $e^{-\epsilon} = 0.5$ and $e^{-\epsilon} = 0.2$. Note that the level of privacy set by ϵ imposes a restriction on the region of ϵ -regular priors. With $e^{-\epsilon} = 0.2$ (less privacy), this region is larger than the one with $e^{-\epsilon} = 0.5$. In fact, as $e^{-\epsilon} \rightarrow 0$ (i.e. no privacy), the region of ϵ -regular priors converges to the entire region of priors defined by the corner points $\{(0, 0, 1), (0, 1, 0), (0, 0, 1)\}$.

Fig. 3. Regions of ϵ -regular priors for the query described in Example 2

Example 2. Let v be database containing at most one record. Consider a bundle of two counting queries $f_3 = (\text{count}(v, p_1), \text{count}(v, p_2))$ which counts the records satisfying properties p_1 and p_2 respectively in the database v . The graph structure of this query is depicted in Figure 1 (with $n = 1$). Note that in this case the adjacency graph (and also the set $\mathcal{R}_f$ of query results) consists of 4 nodes: $\{(0, 0), (1, 0), (0, 1), (1, 1)\}$. Any prior π corresponds therefore to a point in a 4-dimensional space. However, since the 4th component of the prior is redundant ($\sum_i \pi_i = 1$), each prior is defined by its ‘projection’ onto the 3-dimensional subspace. Given this observation, Figure 3 shows the projection of the ϵ -regular prior region for different values of $e^{-\epsilon}$. It is again seen that the region is getting larger as the level of privacy $e^{-\epsilon}$ decreases, and coincides with the full space of priors when $e^{-\epsilon} \rightarrow 0$ (i.e. when no privacy is provided).

4 Upper Bounds for Utility and Min-mutual Information

In this section, we further describe the ϵ -regular priors in terms of the utility that can be achieved for these priors by ϵ -differentially private mechanisms. We also describe the amount of information that can be conveyed by these mechanisms to users with such priors. More precisely, we identify for any ϵ -regular prior π upper bounds for the utility and min-mutual information, considering all ϵ -differentially private mechanisms and all possible remaps. These bounds are indeed induced by the privacy constraints parameterised by ϵ and the query f as stated by Proposition 1. They also depend on the given prior π .

4.1 Utility

Given a query f and a privacy parameter $\epsilon > 0$, let π be a prior on the set $\mathcal{R}_f$ of the query results. For any noise matrix X satisfying ϵ -differential privacy (as in Proposition 1), and a remap R , we derive in the following a linear algebraic expression for $\mathcal{U}(X, \pi, R)$, the

utility of X for π using the remap R . Such an expression will play the main role in the subsequent results. We start by observing that the matrix product of the noise matrix X and the remap R describes an ϵ -differentially private noise matrix $XR : \mathcal{R}_f \rightarrow \mathcal{R}_f$. Hence the entries of XR satisfy (by Proposition 1) the following subset of constraints.

$$e^{-\epsilon d(i,k)} (XR)_{kk} \leq (XR)_{ik}$$

for all $i, k \in \mathcal{R}_f$. Using Definition 4 of the privacy-constraints matrix Φ , and taking into account that $\sum_{k \in \mathcal{R}_f} (XR)_{ik} = 1$ for all i (as both X and R are stochastic), we imply the following inequalities.

$$\sum_{k \in \mathcal{R}_f} \phi_{ik} (XR)_{kk} \leq 1, \quad \forall i \in \mathcal{R}_f.$$

The inequality operators can be replaced by equalities while introducing *slack* variables $0 \leq s_i \leq 1$ for all $i \in \mathcal{R}_f$. The above inequalities can therefore be written as follows.

$$\sum_{k \in \mathcal{R}_f} \phi_{ik} (XR)_{kk} + s_i = 1, \quad \forall i \in \mathcal{R}_f.$$

Let the slack variables s_i form a column vector s indexed by $\mathcal{R}_f$. Let also $\mathbf{1}$ denote another column vector of the same size having all entries equal to 1. Using these vectors and the privacy-constraints matrix Φ (for the given query and ϵ), the above equations can be rewritten in the following matrix form.

$$\Phi \text{diag}(XR) + s = \mathbf{1}, \tag{3}$$

where $\text{diag}(XR)$ is the column vector consisting of the diagonal entries of XR . Now, for any noise matrix $X : \mathcal{R}_f \rightarrow \mathcal{O}$ and a remap $R : \mathcal{O} \rightarrow \mathcal{R}_f$ satisfying Eq. (3), and for a prior π , we want to refine the generic expression (2) of the utility by taking Eq. (3) into account. We start by rewriting Eq. (2) in the following matrix form.

$$\mathcal{U}(X, \pi, R) = \pi \text{diag}(XR). \tag{4}$$

Now, let y be a row vector such that

$$\pi = y \Phi. \tag{5}$$

Note that, the above matrix equation is in fact a system of $|\mathcal{R}_f|$ linear equations. The k th equation in this system is formed by the k th column of Φ , and the k th entry of π as follows.

$$y \Phi_k = \pi_k \quad \forall k \in \mathcal{R}_f.$$

Solving this system of equations for the row vector y has the following possible outcomes: If the matrix Φ is invertible, then, for any prior π , Eq. (5) has exactly one solution. If Φ is not invertible (i.e. it contains linearly dependent columns), then there are either 0 or an infinite number of solutions, depending on the prior π : If the entries of π respect the linear dependence relation then there are infinitely many solutions. Otherwise, the equations are ‘*inconsistent*’, in which case there are no solutions.

Since the matrices Φ have a precise format, one may wonder whether it could be that they are all invertible or all non invertible. In fact, this is not the case: In the report

version of this paper [21] we show an example of a matrix Φ that, for certain values of ϵ is invertible, while for others is non invertible.

Whether Φ is invertible or not, we consider here only the priors where the matrix equation (5) has at least one solution $\mathbf{y}$. Note that, by definition, all the ϵ -regular priors have this property, but there can be others for which the solution $\mathbf{y}$ has some negative components. In some of the results below (in particular in Lemma 2) we consider this larger class of priors, for the sake of generality.

Multiplying Equation (3) by $\mathbf{y}$ yields

$$\mathbf{y} \Phi \text{diag}(X R) + \mathbf{y} \mathbf{s} = \mathbf{y} \mathbf{1}. \quad (6)$$

Substituting Equations (5) and (4) in the above equation consecutively provides the required expression for the utility and therefore proves the following lemma.

Lemma 2. *For a given query f and a privacy parameter $\epsilon > 0$, let π be any prior. Then for every row vector $\mathbf{y}$ satisfying $\pi = \mathbf{y} \Phi$, the utility of any ϵ -differentially private mechanism with a noise matrix X for the prior π using a remap R is given by*

$$\mathcal{U}(X, \pi, R) = \mathbf{y} \mathbf{1} - \mathbf{y} \mathbf{s}, \quad (7)$$

for a vector $\mathbf{s}$ satisfying $0 \leq s_i \leq 1$ for all $i \in \mathcal{R}_f$.

Lemma 2 expresses the utility function for any ϵ -private noise matrix X for a prior π with a remap R as a function of the vector $\mathbf{y}$ and the slack vector $\mathbf{s}$. Although the matrix X and the remap R do not explicitly appear on the right hand side of Equation (7), the utility still depends on them indirectly through the vector $\mathbf{s}$. Namely, according to Equation (3), the choice of X and R determines the slack vector $\mathbf{s}$. The utility function depends also on the prior π , because the choice of π determines the set of vectors satisfying Eq. (5). Substituting any of these vectors $\mathbf{y}$ in Eq. (7) yields the same value for $\mathcal{U}(X, \pi, R)$.

By Definition 5, of ϵ -regular priors, the above lemma specifies the utility for any of them. Therefore, we use Lemma 2, and obtain an upper bound for the utility of ϵ -differentially private mechanisms for ϵ -regular priors.

Theorem 1 (utility upper bound). *For a given query f and a privacy parameter $\epsilon > 0$, let π be an ϵ -regular prior and X be an ϵ -differentially private noise matrix. Then for all row vectors $\mathbf{y} \geq 0$ satisfying $\mathbf{y} \Phi = \pi$, it holds for any remap R that*

$$\mathcal{U}(X, \pi, R) \leq \sum_{i \in \mathcal{R}_f} y_i, \quad (8)$$

where the equality holds iff $\Phi \text{diag}(X R) = \mathbf{1}$.

The above result can be also seen from the geometric perspective. As shown by Proposition 4, each member in the region of ϵ -regular priors is described as a convex combination of the corner priors. That is there are coefficients $\gamma_i \geq 0$ for $i \in \mathcal{R}$ which form this combination. It can be shown (as in the proof of Proposition 4) that $\gamma_i = y_i \left(\sum_{k \in \mathcal{R}_f} \phi_{ik} \right)$. Hence, the upper bound given by Theorem 1 can be written as follows using the coefficients γ_i .

$$\mathcal{U}(X, \pi, R) \leq \sum_{i \in \mathcal{R}_f} \frac{\gamma_i}{\sum_{k \in \mathcal{R}_f} \phi_{ik}}.$$

Inspecting the above result for corner priors, recall that for a corner $\mathbf{c}^i$, γ_j is 1 for $j = i$ and is 0 otherwise; thus, the utility upper bound for $\mathbf{c}^i$ is therefore $1 / \sum_k \phi_{ik}$. Moreover, the upper bound for each ϵ -regular prior π can be regarded (according to the above equation) as a convex combination of the upper bounds for the corner priors. That is, from the geometric perspective, the utility upper bound for π linearly depends on its proximity to the corner priors.

4.2 Min-mutual Information

In this section, we employ an information-theoretic notion, namely mutual information, to quantify the amount of information conveyed by a noise matrix X as an information theoretic channel. We use this notion in two distinct ways: first, mutual information is used to measure the information conveyed about the result of a specific query, similarly to the use of “utility” in the previous section. Mutual information and utility (under the binary gain function) are closely related, which allows us to transfer the bound obtained in the previous section to the information-theoretic setting.

Second, we use mutual information to quantify the information *about the database* that is revealed by a mechanism, a concept known in the area of quantitative information flow as “information leakage”. This allows us to obtain bounds on the information leaked by any mechanism, even non-oblivious ones, independently from the actual query. For arbitrary priors, we obtain in a more natural way the bound conjectured in [18] and proven in [8]. Moreover, if we restrict to specific (ϵ -regular) priors, then we are able to provide more accurate bounds.

Following recent works in the area of quantitative information flow ([14–17, 8, 18]), we adopt Rényi’s *min-entropy* ([22]) as our measure of uncertainty. The min-entropy $\mathcal{H}_\infty(\pi)$ of a prior π , defined as $\mathcal{H}_\infty(\pi) = -\log_2 \max_i \pi_i$, measures the user’s uncertainty about the query result. Then, the corresponding notion of *conditional* min-entropy, defined as $\mathcal{H}_\infty(X, \pi) = -\log_2 \sum_o \max_i \pi_i x_{io}$, measures the uncertainty about the query result after observing the output of the noise matrix X . Finally, subtracting the latter from the former brings us to the notion of min-mutual information:

$$\mathcal{L}(X, \pi) = \mathcal{H}_\infty(\pi) - \mathcal{H}_\infty(X, \pi)$$

which measures the amount of information about the query result conveyed by the noise matrix. In the area of quantitative information flow this quantity is known as *min-entropy leakage*; the reader is referred to [14] for more details about this notion.

Min-mutual information is closely related to the notion of utility under the binary gain function and using an *optimal* remap. A remap $\hat{R}$ is optimal for X, π if it gives the best utility among all possible remaps for this noise matrix and prior. The following result from [8] connects min-mutual information and utility:

Proposition 5. *Given a noise matrix X and a prior π , let $\hat{R}$ be an optimal remap for π, X . Then, it holds*

$$\mathcal{L}(X, \pi) = \log_2 \frac{\mathcal{U}(X, \pi, \hat{R})}{\max_i \pi_i}$$

This connection allows us to transfer the upper-bound given by Theorem 1 to min-mutual information.

Proposition 6 (min-mutual information upper bound). *Let f be a query, let $\epsilon > 0$, let π be an ϵ -regular prior and let X be the noise matrix of any ϵ -differentially private mechanism. Then for all row vectors $\mathbf{y} \geq 0$ satisfying $\mathbf{y} \Phi = \pi$, it holds that:*

$$\mathcal{L}(X, \pi) \leq \log_2 \frac{\sum_{i \in \mathcal{R}_f} y_i}{\max_i \pi_i} \quad (9)$$

The above bound holds only for ϵ -regular priors. However, it is well-known ([15]) that min-mutual information is maximised by the uniform prior $\mathbf{u}$, i.e. $\mathcal{L}(X, \pi) \leq \mathcal{L}(X, \mathbf{u})$ for all X, π . Thus, in cases when $\mathbf{u}$ is ϵ -regular, we can extend the above bound to any prior.

Corollary 1. *Let f be a query, let $\epsilon > 0$ such that the uniform prior $\mathbf{u}$ is ϵ -regular, and let X be the noise matrix of any ϵ -differentially private mechanism. Then for all row vectors $\mathbf{y} \geq 0$ satisfying $\mathbf{y} \Phi = \mathbf{u}$, and for all priors π , it holds that:*

$$\mathcal{L}(X, \pi) \leq \log_2(|\mathcal{R}_f| \sum_{i \in \mathcal{R}_f} y_i)$$

4.3 Quantifying the Leakage about the Database

In the previous section we considered the information about the query result conveyed by an oblivious mechanism. This information was measured by the min-mutual information $\mathcal{L}(X, \pi)$, where X is noise matrix, mapping query results $\mathcal{R}_f$ to outputs.

We now turn our attention to quantifying the information about the *database* that is conveyed by the complete mechanism $\mathcal{K}$ (even in the case of non-oblivious mechanisms). Intuitively, we wish to minimise this information to protect the privacy of the users, contrary to the utility which we aim at maximising. Quantifying this information can be done in a way very similar to the previous section. The only difference is that we use a stochastic matrix Y that models the mechanism $\mathcal{K}$, mapping databases $\mathcal{V} = V^u$ to outputs (recall that u is the number of individuals in the database and V the set of possible values for each individual). Moreover, the underlying graph $\sim$ is the *Hamming graph*, induced by the adjacency relation on databases, and ϵ -regularity concerns priors π on databases.

In this case, $\mathcal{L}(Y, \pi)$ measures the information about the database conveyed by the mechanism, which we refer to as “min-entropy leakage”, and the bounds from the previous section can be directly applied. However, since we now work on a specific graph $(\mathcal{V}, \sim)$, we can obtain a closed expression for the bound of Corollary 1. We start by observing that due to the symmetry of the graph, the uniform prior $\mathbf{u}$ is ϵ -regular for all $\epsilon > 0$. More precisely, we can show that the vector $\mathbf{y}$, defined as

$$y_i = \left(\frac{e^\epsilon}{|V|(|V| - 1 + e^\epsilon)} \right)^u \quad i \in \mathcal{V}$$

satisfies $\mathbf{y} \Phi = \mathbf{u}$ and $\mathbf{y} \geq 0$. Thus, applying Corollary 1 we get the following result.

Theorem 2 (min-entropy leakage upper bound). *Let $\mathcal{V} = V^u$ be a set of databases, let $\epsilon > 0$, and let Y be an ϵ -differentially private mechanism. Then for all priors π , it holds that:*

$$\mathcal{L}(Y, \pi) \leq u \log_2 \frac{|V| e^\epsilon}{|V| - 1 + e^\epsilon}$$

Fig. 4. Leakage bounds for various values of ϵ

This bound determines the maximum amount of information that *any* differentially privacy mechanism can leak about the database (independently from the underlying query). The bound was first conjectured in [18] and independently proven in [8]; our technique gives an alternative and arguably more intuitive proof of this result.

Note that the above bound holds for *all* priors. If we restrict to a *specific* ϵ -regular prior π , then we can get better results by using the bound of Proposition 6 which depends on the actual prior. This is demonstrated in the following example.

Example 3. Consider a database of 5 individuals, each having one of 4 possible values, i.e. $\mathcal{V} = V^u$ with $V = \{1, 2, 3, 4\}$ and $u = 5$. Assume that each individual selects a value independently from the others, but not all values are equally probable; in particular the probabilities of values 1, 2, 3, 4 are 0.3, 0.27, 0.23, 0.2 respectively. Let π be the corresponding prior on $\mathcal{V}$ that models this information. We have numerically verified that for all $0.48 \leq \epsilon \leq 1$ (with step 0.01) π is ϵ -regular. Thus we can apply Proposition 6 to get an upper bound of $\mathcal{L}(Y, \pi)$ for this prior.

The resulting bound, together with the general bound for all priors from Theorem 2, are shown in Figure 4. We see that restricting to a specific prior provides a significantly better bound for all values of ϵ . For instance, for $\epsilon = 0.5$ we get that $\mathcal{L}(Y, \pi) \leq 1.2$ for this π , while $\mathcal{L}(Y, \pi) \leq 2.5$ for all priors π .

Note that, in general, the above bounds for the utility and the min-mutual information are not tight. For a given query and a privacy parameter ϵ , there may be no noise matrix X that meets these bounds. Nevertheless, they provide ultimate limits, induced by the privacy constraints, for all ϵ -differential private mechanisms and ϵ -regular priors. Note also that these bounds are simultaneously tight if the *common* condition $\Phi \text{diag}(XR) = \mathbf{1}$ is satisfied (note that this condition is independent of the underlying prior). From this point we investigate the mechanisms that, whenever exist, they satisfy such a condition and are therefore optimal for the entire class of ϵ -regular priors.

5 Tight-Constraints Mechanisms

In this section we introduce the notion of *tight-constraints* mechanism and we investigate their properties in terms of privacy guarantees and optimality for ϵ -regular priors.

Definition 7 (A tight-constraints mechanism). *For a given query f with range $\mathcal{R}_f$, and a given privacy parameter $\epsilon > 0$, an oblivious mechanism with a noise matrix $X : \mathcal{R}_f \rightarrow \mathcal{R}_f$ is called a tight-constraints mechanism iff it satisfies the following conditions for all $i, k \in \mathcal{R}_f$.*

$$e^{-\epsilon d(i,k)} x_{kk} = x_{ik}. \quad (10)$$

It is important to note that, in general, there may exist zero, one or more tight-constraints mechanisms for a given query f and a privacy parameter $\epsilon > 0$. The above definition enforces $|\mathcal{R}_f| (|\mathcal{R}_f| - 1)$ linearly independent equations, referred to as the ‘*tight constraints*’. Additionally it must also hold that $\sum_{k \in \mathcal{R}_f} x_{ik} = 1$ for all $i \in \mathcal{R}_f$. Thus we have, in total, $|\mathcal{R}_f| |\mathcal{R}_f|$ equations. If these equations are linearly independent, then they solve to unique values. If these values are non-negative, then they determine a *unique* tight-constraints mechanism. On the other hand, if these equations are not linearly independent, then there may be multiple solutions with non-negative entries, in which case we have multiple tight-constraints mechanisms for the given query and privacy parameter ϵ .

5.1 Properties

It has been seen from Definition 7, that the choice of a query f and a value $\epsilon > 0$ correspond to a set of tight-constraints mechanisms. The first important feature of these mechanisms is that they satisfy ϵ -differential privacy as confirmed by the following proposition.

Proposition 7 (differential privacy). *For a given query f and a privacy parameter $\epsilon > 0$, every tight-constraints mechanism is ϵ -differentially private.*

Thanks to the above fact, we can give a further useful characterisation of the tight-constraints mechanisms in comparison to other ϵ -differentially private mechanisms. More precisely, the following proposition identifies a linear algebraic condition that is satisfied *only by* the tight-constraints mechanisms for given f, ϵ :

Lemma 3 (diagonal characterisation). *Let f be a query and let $\epsilon > 0$. Then for any oblivious ϵ -differentially private mechanism $\mathcal{K}$ with a noise matrix $X : \mathcal{R}_f \rightarrow \mathcal{R}_f$, the following equation holds iff $\mathcal{K}$ is a tight-constraints mechanism.*

$$\Phi \text{ diag}(X) = \mathbf{1}. \quad (11)$$

The above proposition provides a way to check the existence of, and also compute, the tight-constraints mechanisms for given f, ϵ . Since Condition (11) is satisfied only by these mechanisms, then there is at least one tight-constraints mechanism if there is a vector $\mathbf{z}$, with non-negative entries, that satisfies the equation $\Phi \mathbf{z} = \mathbf{1}$. In this case the

noise matrix $\hat{X}$ of a tight-constraints mechanism is obtained by setting its diagonal to $\mathbf{z}$, and evaluating the non-diagonal entries from the diagonal using Eqs. (10).

Now we turn our attention to the region of ϵ -regular priors and we identify the oblivious mechanisms which are optimal wrt both utility and min-mutual information in this region. It turns out that the set of these optimal mechanisms consists exactly of all the mechanisms that can be *mapped* to a tight-constraints one using a remap R .

Theorem 3 (Optimality). *Let f be a query and let $\epsilon > 0$ such that at least one tight-constraints mechanism exists. Then any oblivious mechanism $\mathcal{K} : \mathcal{V} \rightarrow \mathcal{O}$ is optimal (wrt both utility and min-mutual information) for every ϵ -regular prior π iff there is a remap $R : \mathcal{O} \rightarrow \mathcal{R}_f$ such that $\mathcal{K}R$ is a tight-constraints mechanism for f, ϵ .*

Proof. If there exists a tight-constraints mechanism for given f, ϵ , then its noise matrix $\hat{X}$ must satisfy Eq (11). This implies that the upper-bound in Theorem 1 is reachable by $\hat{X}$ and the identity remap. Thus that upper-bound, in this case, is tight. By Theorem 1, a mechanism $\mathcal{K}$ with a noise matrix X meets such an upper bound for the utility (and therefore is optimal) iff it satisfies the condition $\Phi \text{diag}(XR) = \mathbf{1}$, with some remap R . Since any mechanism with noise matrix XR is ϵ -differentially private, then by Lemma 3, this condition is satisfied iff XR is the noise matrix of a tight-constraints mechanism (for f, ϵ). That is iff $fXR = \mathcal{K}R$ is a tight-constraints mechanism. Using the relation, given by Proposition 5, between utility and min-mutual information, the same argument holds for the latter. $\square$

Note that tight-constraints mechanisms are themselves optimal as they are mapped to themselves by the identity remap.

As a consequence of the above general result, we consider the special case of the uniform prior, denoted by $\mathbf{u}$, where all exact query results in $\mathcal{R}_f$ are equally likely. Note that this prior corresponds to users having unbiased knowledge about the query results, i.e. they assume that all the exact results $\mathcal{R}_f$ are yielded, by executing the query, with the same probability. Firstly, the following lemma proves an equivalence between the existence of at least one tight-constraints mechanism on one hand and the uniform prior $\mathbf{u}$ being ϵ -regular on the other hand.

Lemma 4. *For a given query f and privacy parameter $\epsilon > 0$, there exists at least one tight-constraints mechanism iff the uniform prior $\mathbf{u}$ is ϵ -regular.*

It is worth noticing that in general the region of ϵ -regular priors may or may not include the uniform prior. However, as shown earlier in Section 3, this region is enlarged and converges to the entire prior space as less privacy is imposed (that is as ϵ increases). This means that for the values of ϵ above certain threshold ϵ^* , depending on the query, the region of ϵ -regular priors accommodates the uniform prior $\mathbf{u}$, and therefore (by Lemma 4), there is at least one tight-constraints mechanism. This provides a design criteria to *select* a setting for ϵ such that we have an optimal mechanism for the whole region.

Using Lemma 4, we can describe the optimal mechanisms for the uniform prior as a corollary of Theorem 3.

Corollary 2. *Let f be a query and let $\epsilon > 0$ such that there exists at least one tight-constraints mechanism. Then a mechanism $\mathcal{K} : \mathcal{V} \rightarrow \mathcal{O}$ is optimal for the uniform prior iff $\mathcal{K}R$ is a tight-constraints mechanism for some remap $R : \mathcal{O} \rightarrow \mathcal{R}_f$.*

Fig. 5. Adjacency graphs

In fact when we consider arbitrary queries, we find that our specification for the tight-constraints mechanisms covers other well known differentially-private mechanisms. In particular, when we consider a counting query, we find that a tight-constraints mechanism for this query is exactly the *truncated-geometric mechanism*, which is shown by [6] to be optimal for every prior. Furthermore, we are able to show that this mechanism, as a tight-constraints one, exists for the selected query with any $\epsilon > 0$.

Another class of queries, studied in [8] are the ones whose graph structures are either *vertex-transitive* or *distance-regular*. The authors in [8] were able to construct a mechanism which is optimal for the uniform prior for any $\epsilon > 0$. In the context of our results, when we consider a query f in this class, it is easy to show that such an optimal mechanism is in fact a tight-constraints mechanism for f . We can also show that this tight-constraints mechanism exists for all $\epsilon > 0$.

6 Case-Study: Sum and 2-Count Queries

In this section we show the usefulness of the tight-constraints mechanism by applying it to two particular families of queries, namely sum and 2-count queries. For each family, we evaluate the tight-constraint mechanism on databases consisting of u individuals each having an integer value between 0 and v , and we compare its utility to the geometric mechanism.

It is well-known that no universally optimal mechanism exists for these families; in particular, the geometric mechanism, known to be optimal for a single counting query, is not guaranteed to be optimal for sum queries or multiple counting queries. On the other hand, as discussed in the previous section, tight-constraints mechanisms, whenever they exist, are guaranteed to be optimal within the region of ϵ -regular priors.

The comparison is made as follows: for each query, we numerically compute the smallest ϵ (using a step of 0.01) for which a tight-constraints mechanism exists (i.e. for which the uniform prior $\mathbf{u}$ is ϵ -regular, see Lemma 4). Then we compute the utility (using an optimal remap) of both the tight constraints and the geometric mechanisms, for a range of ϵ starting from the minimum one. Note that the tight constraint mechanism exists for any ϵ greater than the minimum one.

Sum query. Let f be the query returning the sum of the value for all individuals, thus it has range $\mathcal{R}_f = \{0, \dots, vu\}$. By modifying the value of a single individual, the outcome

Fig. 6. Utility for various values of ϵ

of the query can be altered by at most v (when changing the value from 0 to v), thus two elements $i, j \in \mathcal{R}_f$ are adjacent iff $|i - j| \leq v$. The induced graph structure on $\mathcal{R}_f$ is shown in Figure 5(a) (for the case $v = 3$).

For our case-study we numerically evaluate this query for $u = 150, v = 5$ and for the uniform prior. We found that the minimum ϵ for which a tight-constraints mechanism exists (and is in fact unique since Φ is invertible) is 0.8. Figure 6(a) shows the utility of the tight-constraint mechanism, as well as that of the geometric mechanism, for values of ϵ between 0.8 and 1.3, the uniform prior and using and optimal remap. We see that the tight-constraint mechanism provides significantly higher utility than the geometric mechanism in this case.

2-count query. Consider now the query f consisting of 2 counting queries (i.e. reporting the number of users satisfying properties p_1 and p_2), thus it has range $\mathcal{R}_f = \{0, \dots, u\} \times \{0, \dots, u\}$. By modifying the value of a single individual, the outcome of each counting query can be altered by at most 1, thus two answers $(i_1, i_2), (j_1, j_2) \in \mathcal{R}_f$ are adjacent iff $|i_1 - j_1| \leq 1$ and $|i_2 - j_2| \leq 1$. The induced graph structure on $\mathcal{R}_f$ is shown in Figure 5(b).

We evaluate this query for $u = 30$ and for the uniform prior. We found that the minimum ϵ for which a tight-constraints mechanism exists is 0.9. Figure 6(b) shows the utility of the tight-constraint mechanism, as well as that of the geometric mechanism (applied independently to each counting query), for values of ϵ between 0.9 and 1.3 and the uniform prior. Similarly to the sum query, we see that the tight-constraint mechanism provides significantly higher utility than the geometric mechanism in this case.

7 Conclusion and Future Work

In this paper we have continued the line of research initiated by [6, 1] about the existence of universally-optimal differentially-private mechanisms. While the positive result of [6] (for counting queries) and the negative one of [1] (for essentially all other queries) answer the question completely, the latter sets a rather dissatisfactory scenario for differential privacy and the typical mechanisms used in the community, since counting queries are just one of the (infinitely many) kinds of queries one can be interested

in. In practice the result of [1] says that for essentially any query other than counting queries one cannot devise a mechanism that gives the best trade-off between privacy and utility for all users. Hence one has to choose: either design a different mechanism for every user, or be content with a mechanism that, depending on the user, can be far from providing the best utility. We have then considered the question whether, for a generic query, the optimality is punctual or can actually be achieved with the same mechanism for a class of users. Fortunately the answer is positive: we have identified a class of priors, called ϵ -regular, and a mechanism which is optimal for all the priors in this class. We have also provided a complete and effectively checkable characterisation of the conditions under which such mechanism exists, and an effective method to construct it. As a side result, we have improved on the existing bounds for the min-entropy leakage induced by differential privacy. More precisely, we have been able to give specific and tight bounds for each ϵ -regular prior, in general smaller than the bound existing in literature for the worst-case leakage (achieved by the uniform prior [20]).

So far we have been studying only the case of utility for binary gain functions. In the future we aim at lifting this limitation, i.e. we would like to consider also other kinds of gain. Furthermore, we intend to study how the utility decreases when we use a tight-constraints mechanism outside the class of ϵ -regular priors. In particular, we aim at identifying a class of priors, larger than the ϵ -regular ones, for which the tight-constraints mechanism is close to be optimal.

The definition of tight-constraints mechanism is related to the connectivity condition of the column graphs introduced by Kifer and Lin [23, 24]. They show that this property implies maximality w.r.t. the postprocessing preorder. Following the suggestion of an anonymous reviewer, we will try to exploit the above result to strengthen our results in Section 5, in particular Theorem 3.

The negative result of [1] is stated in terms of the graph induced by $\sim_f$: a universally optimal mechanism can exist only if such graph is a line. This is the case of counting queries, but not only: any composition of a counting query with a function that preserves the graph structure would induce the same kind of graph, and it is for this reason that the authors of [1] write “*essentially* counting queries”. As pointed out by an anonymous reviewer, we can use the techniques of our paper to prove that a universally optimal mechanism exists if and only if the query is derivable from a counting query by a bijection, thus making the result of [1] more precise, and extending the result of [6].

Acknowledgement. We would like to thank the anonymous reviewers for their valuable recommendations for improving the paper, and their suggestion for future work.

References

1. Brenner, H., Nissim, K.: Impossibility of differentially private universally optimal mechanisms. In: Proc. of FOCS, pp. 71–80. IEEE (2010)
2. Dwork, C.: Differential Privacy. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006, Part II. LNCS, vol. 4052, pp. 1–12. Springer, Heidelberg (2006)
3. Dwork, C., Lei, J.: Differential privacy and robust statistics. In: Proc. of STOC, pp. 371–380. ACM (2009)
4. Dwork, C.: Differential privacy in new settings. In: Proc. of SODA, pp. 174–183. SIAM (2010)

5. Dwork, C.: A firm foundation for private data analysis. *Communications of the ACM* 54(1), 86–96 (2011)
6. Ghosh, A., Roughgarden, T., Sundararajan, M.: Universally utility-maximizing privacy mechanisms. In: *Proc. of STOC*, pp. 351–360. ACM (2009)
7. Gupte, M., Sundararajan, M.: Universally optimal privacy mechanisms for minimax agents. In: *Proc. of PODS*, pp. 135–146. ACM (2010)
8. Alvim, M.S., Andrés, M.E., Chatzikokolakis, K., Palamidessi, C.: On the Relation between Differential Privacy and Quantitative Information Flow. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) *ICALP 2011, Part II. LNCS*, vol. 6756, pp. 60–76. Springer, Heidelberg (2011)
9. Köpf, B., Basin, D.A.: An information-theoretic model for adaptive side-channel attacks. In: *Proc. of CCS*, pp. 286–296. ACM (2007)
10. Clark, D., Hunt, S., Malacaria, P.: Quantitative information flow, relations and polymorphic types. *J. of Logic and Computation* 18(2), 181–199 (2005)
11. Boreale, M.: Quantifying Information Leakage in Process Calculi. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) *ICALP 2006, Part II. LNCS*, vol. 4052, pp. 119–131. Springer, Heidelberg (2006)
12. Malacaria, P.: Assessing security threats of looping constructs. In: *Proc. of POPL*, pp. 225–235. ACM (2007)
13. Chatzikokolakis, K., Palamidessi, C., Panangaden, P.: Anonymity protocols as noisy channels. *Inf. and Comp.* 206(2–4), 378–401 (2008)
14. Smith, G.: On the Foundations of Quantitative Information Flow. In: de Alfaro, L. (ed.) *FOSSACS 2009. LNCS*, vol. 5504, pp. 288–302. Springer, Heidelberg (2009)
15. Braun, C., Chatzikokolakis, K., Palamidessi, C.: Quantitative notions of leakage for one-try attacks. In: *Proc. of MFPS. ENTCS*, vol. 249, pp. 75–91. Elsevier (2009)
16. Köpf, B., Smith, G.: Vulnerability bounds and leakage resilience of blinded cryptography under timing attacks. In: *Proc. of CSF*, pp. 44–56. IEEE (2010)
17. Andrés, M.E., Palamidessi, C., van Rossum, P., Smith, G.: Computing the Leakage of Information-Hiding Systems. In: Esparza, J., Majumdar, R. (eds.) *TACAS 2010. LNCS*, vol. 6015, pp. 373–389. Springer, Heidelberg (2010)
18. Barthe, G., Köpf, B.: Information-theoretic bounds for differentially private mechanisms. In: *Proc. of CSF*, pp. 191–204. IEEE (2011)
19. Clarkson, M.R., Schneider, F.B.: Quantification of integrity. *Tech. Rep.* (2011), <http://hdl.handle.net/1813/22012>
20. Alvim, M.S., Andrés, M.E., Chatzikokolakis, K., Degano, P., Palamidessi, C.: Differential Privacy: On the Trade-Off between Utility and Information Leakage. In: Barthe, G., Datta, A., Etalle, S. (eds.) *FAST 2011. LNCS*, vol. 7140, pp. 39–54. Springer, Heidelberg (2012)
21. ElSalamouny, E., Chatzikokolakis, K., Palamidessi, C.: A differentially private mechanism of optimal utility for a region of priors. *Technical report*, INRIA (2013), <http://hal.inria.fr/hal-00760735/>
22. Rényi, A.: On Measures of Entropy and Information. In: *Proc. of the 4th Berkeley Symposium on Mathematics, Statistics, and Probability*, pp. 547–561 (1961)
23. Kifer, D., Lin, B.R.: Towards an axiomatization of statistical privacy and utility. In: *Proc. of PODS*, pp. 147–158. ACM (2010)
24. Kifer, D., Lin, B.R.: An axiomatic view of statistical privacy and utility. *Journal of Privacy and Confidentiality* 4(1), 5–49 (2012)