Circulant Matrices and Differential Privacy

Jalaj Upadhyay
Center for Applied Cryptographic Research, University of Waterloo.
email: jalaj.upadhyay@uwaterloo.ca

Abstract

This paper resolves an open problem raised by Blocki *et al.* (FOCS 2012), i.e., whether other variants of the Johnson-Lindenstrauss transform preserves differential privacy or not? We prove that a general class of random projection matrices that satisfies the Johnson-Lindenstrauss lemma also preserves differential privacy. This class of random projection matrices requires only n Gaussian samples and n Bernoulli trials and allows matrix-vector multiplication in $O(n \log n)$ time. In this respect, this work unconditionally improves the run time of Blocki *et al.* (FOCS 2012) without using the graph sparsification trick of Upadhyay (ASIACRYPT 2013). For the metric of measuring randomness, we stick to the norm used by earlier researchers who studied variants of the Johnson-Lindenstrauss transform and its applications, i.e., count the number of random samples made. In concise, we improve the sampling complexity by quadratic factor, and the run time of cut queries by an $O(n^{o(1)})$ factor and that of covariance queries by an $O(n^{0.38})$ factor.

Our proof for both the privacy and utility guarantee uses several new ideas. In order to improve the dimension bound, we use some known results from the domain of statistical model selection. This makes our proof short and elegant, relying just on one basic concentration inequality. For the privacy proof, even though our mechanism closely resembles that of Blocki *et al.* (FOCS 2012) and Upadhyay (ASIACRYPT 2013), we cannot use their proof idea. This is because the projection matrices we are interested in introduces non-trivial correlations between any two rows of the published matrix, and, therefore, we cannot invoke the composition theorem of Dwork, Rothblum and Vadhan (STOC 2009). We argue that the published matrix is not r-multivariate distribution; rather one matrix-variate distribution. We compute the distribution of the published matrix and then prove it preserves differential privacy.

Keywords. Circulant Matrices, Differential privacy, Sampling Complexity.

1 Introduction

In a recent work, Blocki *et al.* [9] proved that the Johnson-Lindenstrauss transform with random i.d.d. Gaussian entries preserves differential privacy, a very robust guarantee of privacy on database query. They left the question open whether other variants of the Johnson-Lindenstrauss transform, more specifically, the fast Johnson-Lindenstrauss transform, the randomness efficient Johnson-Lindenstrauss, and the sparse Johnson-Lindenstrauss transform, preserves privacy or not? In this paper, we resolve this issue. We consider a general class of random projection matrices of which the construction of Vybiral [51] is a special instance, and show that every projection matrix in this class of matrices preserves differential privacy, requires only n Gaussian samples and n Bernoulli trials, and allows fast matrix-vector multiplication.

The transform of Vybiral [51] is based on a class of matrices called *partial circulant matrices*[1], and achieves a suboptimal dimension reduction. We note that unless there is a significant improvement in the concentration properties of the partial circulant matrices, one cannot improve the dimension bound achieved

[1]Partial circulant matrices are a class of matrices indexed by a n-dimensional vector and formed as follows: the first row is the n-dimensional vector and the rest of the rows are formed iteratively by circulating the entries by shifting entries one position left with respect to the previous row.

by Vybiral [51]. We make a slight modification to their transform (more specifically, by composing a Walsh-Hadamard transform matrix) to get an almost optimal dimension reduction. We then consider the general class of matrices of which partial circulant matrices are special instance. We prove the utility and privacy bound for this general class of matrices. Using this proof, we also give a simpler proof for the recent construction of Upadhyay [50].

One of the reasons Blocki *et al.* [9] perceived the study of other variants of the Johnson-Lindenstrauss important is due to their algorithmic and practical implications [8]. As argued in a series of work by Ailon and Liberty [2, 3, 4] and Krahmer and Ward [38], the dimension of the projected space, run time of the transform, and the number of random samples used by the transform are the most important parameters for a Johnson-Lindenstrauss transform from an algorithmic as well as the applications and implementation point of view. On the other hand, major focus of research in differential privacy until now has been towards providing a tight utility and privacy tradeoff. We sought to bring the resource consideration in the domain of differential privacy as well. Similar questions were also raised by Dwork *et al.* [23] and Upadhyay [49], where the focus was on the design of efficient mechanisms with respect to the time taken to generate the sanitized data. This paper forward the study of privacy preserving mechanisms while also taking in account the amount of randomness used. An alternative way to look at this paper is the natural question whether we can use a more sampling efficient mechanisms with same utility and privacy guarantee.

One could argue that the number of random bits used is a more natural notion for considering the randomness complexity. However, in this paper, we stick with the norm used by researchers interested in the Johnson-Lindenstrauss transform and its applications, i.e., measure the randomness complexity in terms of the number of random samples. This measure has been used in the domain of dimension reduction and its application like compressed sensing [13], machine learning [5], quantum algorithms [17], and numerical analysis [16, 47].

Ailon and Liberty [2, 3, 4] and Krahmer-Ward [38] have thoroughly motivated why the number of random samples is an important parameter with respect to other applications of the Johnson-Lindenstrauss transform. Apart from all those reasons, one of the other main reasons for this choice in the domain of differential privacy is that it gives a much cleaner picture and a good quantitative estimate on the actual random bits used and the (actual) run-time of the mechanism[2]–it is the sampling process which is implemented in practice and might cause several issues. We refer the readers to Kapralov and Talwar [34] for various theoretical and Chaudhary *et al.* [15] and Mironov [41] for practical issues faced during sampling.

OUR TECHNIQUES. In this paper, we investigate whether suitable modifications and generalization to a known transform by Vybiral [51], while maintaining the number of random samples used and efficiency in terms of matrix-vector multiplication, preserves differential privacy or not. Moreover, we also give a sharper analysis to get a tighter bound than achieved by Vybiral [51] in terms of the dimension of the projected space. In concise, we achieve a quadratic improvement in the number of random samples used in comparison with all known mechanisms for answering cut queries and covariance queries, and $\mathrm{poly}\log n$ improvement in the run time over the mechanism based on graph sparsification [49]. We follow up with the generalization of the Vybiral's construction and the techniques used in this paper to prove the privacy and utility bound.

OVERVIEW OF THE CONSTRUCTION. We start by giving a brief exposition of the construction by Vybiral [51]. Vybiral [51] first pick n random Gaussian samples to form the first row and then construct the remaining $r-1$ rows by shifting the vector left-wise relative to the previous row. This matrix is also known as *partial circulant matrices* and satisfies the *Restricted Isometry Property* [13], which we define next. For for any set $T \subseteq \{1, \cdots, r\}$, we say that an $r \times n$ matrix Φ satisfies the *Restricted Isometry Property* of *order*

[2]This could be seen akin to the complexity measure used in generic as well as concrete attacks on hash functions where we just measure the number of hash computations done, and not the actual atomic operations required (see for example, the attack on SHA-1 [52] and MD5 [53], and the generic attack on collision resistance [33] and second pre-image resistance [35]).

k if there exists an $0 < \varepsilon < 1$ such that, for all set T with $|T| < k$,

$$\Pr\left[(1-\varepsilon)\|\mathbf{x}_T\|_2^2 \leq \|\Phi_T|x\rangle_T\|_2^2 \leq (1+\varepsilon)\||x\rangle_T\|_2^2\right] \geq 1-\eta \tag{1}$$

holds, where Φ_T ($\mathbf{x}_T$, respectively) is the restriction of Φ ($\mathbf{x}$, respectively) to the indices in T.

Rauhut *et al.* [44] proved that a partial circulant matrix formed as above satisfies the Restricted Isometry Property for values of $r \geq \max(\varepsilon^{-1}\sqrt{(k\log n)^3}, \varepsilon^{-2}k\log^4 n)$. Vybiral then multiply this matrix P to a diagonal matrix whose entries are ± 1 with probability $1/2$. By Theorem 1, this construction satisfies the Johnson-Lindenstrauss bound for suboptimal value of $r = O(\varepsilon^{-2}\log^2 m)$, where m is the number of vectors on which the transform is to be applied.

Theorem 1. *(Krahmer-Ward [38, Proposition 3.2])* Let ε be an arbitrary constant. Let Φ be a matrix of order k and dimension $r \times n$ that satisfies the relation $k \leq c_1\delta_k^2 r/\log(n/r)$ and equation (1). Then the matrix ΦD, where D is an $n \times n$ diagonal matrix whose entries are ± 1 with probability $1/2$ (also known as *Rademacher matrix*), is a Johnson-Lindenstrauss transform with r rows.

Therefore, unless there is a significant improvement in the understanding of concentration properties of partial circulant matrices, the dimension bound achieved by Vybiral [51] is hard to beat. The key observation here is that the diagonal Rademacher matrices does not produce a proper "mixing" of the entries of the partial circulant matrices to facilitate a strong concentration bound. For this, we need to compose it with a matrix that allows fast matrix-vector multiplication, preserves the Euclidean norm of the input vectors, and does not introduce more randomness. Our key observation is that, instead of only preconditioning by a diagonal matrix formed by a Rademacher sequence, if we also compose a Walsh-Hadamard matrix, then we achieve good enough mixing that translates to a better concentration result[3]. This in turn helps us to strengthen the Vybiral's bound [51].

To generalize this construction, a key point to note is that partial circulant matrices are nothing special. They are simply the first r rows of a fully circulant matrix, and, therefore, can be seen as a result of applying a truncated permutation matrix from the left to a fully circulant matrix. Therefore, combined with the symmetry of circulant matrices, intuitively, any r rows restricted circulant matrix should not effect the final concentration result. This intuition infact turns out to be true. This increases our sampling complexity by an additive factor of r because we need to sample r rows of a circulant matrix, but we now have a family of random projection matrices that satisfies the Johnson-Lindenstrauss lemma. Sampling r rows independently was an idea used by Rudelson and Vershynin [46], where they showed that a matrix formed by sampling r rows of a deterministic matrices with bounded orthonormal rows satisfies the Restricted Isometry Property. Here, we are sampling r rows of a random matrix. Therefore, at the cost of oversimplification, an intuitive way to see this general class of Johnson-Lindenstrauss transform is as a hybrid of the known constructions of projection matrix with Restricted Isometry Property and known constructions of the Johnson-Lindenstrauss transform.

TECHNIQUES USED FOR THE UTILITY PROOF. The general idea to prove the Johnson-Lindenstrauss lemma is to first bound the expectation of the random variables corresponding to the output of an application of the projection matrix, and then use the standard concentration bound to prove the result. For example, in the simplified proof of the Johnson-Lindenstrauss transform [19], the above method gives a failure bound of at most $1 - 1/n$. They then repeat the experiment a required number of times to get the failure bound closer to the desired constant. We cannot rely on repetition because it would increase the random samples required. Therefore, we have to give a tighter bound. For this, we rely on a result from statistical model selection.

We break our analysis in two parts. We first use the isometry of Ailon and Chazelle [1], which precondition the input vector $\mathbf{x}$ to get a vector $\tilde{\mathbf{x}}$ with bounded co-ordinates. Then, we use this promise to prove that when we multiply a restricted circulant matrix, formed by n-dimensional Gaussian vector, from the right, then the Euclidean norm of $\mathbf{x}$ is preserved with high probability. For this, we use known concentration inequalities from the area of model selection. Unlike the earlier results on randomness-efficient fast Johnson-Lindenstrauss

[3]This composed matrix is the isometry matrix of Ailon and Chazelle [1].

Method	Cut-queries	Covariance-Queries	Run-time	# Random Samples
Randomized Response [11]	$O(\sqrt{sn^2/\varepsilon})$	$\tilde{O}(\sqrt{nd}/\varepsilon)$	$\Theta(n^2)$	$O(n^2)$
Exponential [12, 40]	$O(n \log n/\varepsilon)$	$O(n \log n/\varepsilon)$	Intractable	$O(n^2)$
Multiplicative Weight [31]	$\tilde{O}(n^2/\varepsilon)$	$\tilde{O}(nd\sqrt{1/\varepsilon})$	$O(n^2)$	$O(n^2)$
Johnson-Lindenstrauss [9]	$O(s\sqrt{n}/\varepsilon)$	$O(\varepsilon^{-2}n)$	$O(rn^2)$	$O(rn)$
Graph Sparsification [49]	$O(s\sqrt{n}/\varepsilon)$	–	$\tilde{O}(n^2)$	$O(n^2)$
This paper	$O(s\sqrt{n}/\varepsilon)$	$O(\varepsilon^{-1}n)$	$O(n^2 \log n)$	$2n + r$

Table 1: Comparison between our mechanism and other mechanism when answering all possible queries.

transform that uses matrices satisfying the Restricted Isometry Property, the proof in this paper is elementary and relies on basic concentration inequalities.

TECHNIQUES USED FOR THE PRIVACY PROOF. The proof of differential privacy is far more involved. We use the same notion of neighbouring data as in [9, 49]. So, it is tempting to assume that multiplying our projection matrix (because of the form it has) results in r multivariate Gaussian and proof of [9] can be applied. However, there are subtle correlations between two rows (or two columns) of our projection matrices, and, applying these projection matrices to a private matrix does not yield r independent multivariate distribution, but one matrix-variate distribution. We first compute the distribution of the published matrix and then prove it is differentially private. Our proof uses various characterization of positive-definite matrices and Hermittian matrices along with the properties of the trace of a matrix. Along the line, we need to prove concentration result for a distribution which is the sum of the squares of n independent Gaussian variables.

Using the above technique for privacy proof, we also give a simpler proof for the construction given by Upadhyay [50]. We recall that the author used a different isometry matrix than that of Ailon and Chazelle [1] for projection that preserves differential privacy. We give our proof for the original construction. Our proof involves reducing the proof of differential privacy for their original construction to that used in this paper.

One can also implement our mechanisms *as distributed algorithms*, a desirable feature as argued by [6]. This is because our mechanism uses operations that have efficient distributed algorithms. For example, one could use Jacobi method for SVD [37] and Cannon's algorithm for multiplication [14].

We summarize our results and its comparison with previous works in Table 1. The second and third column is the noise added by the respective mechanisms when answering all possible queries. In the table, s denote the size of a single query and r is the dimension of the projected space in our transform. Note that, except for the random projection based mechanisms, all the other mechanisms are interactive. Since comparing the noise bound for interactive and non-interactive mechanisms is not that straightforward, in our comparison, we follow Blocki *et al.* [9] method: compare answering all set of adaptive queries for interactive mechanims and all predetermined queries for non-interactive mechanisms (see [9, Section 3.2 and 4.2]).

RELATED WORK. The first formal definition of Differential Privacy was given by Dwork et al. [22] to address the privacy concern of any participants. The key idea used in Dwork et al. [22] is to add noise according to a Laplace distribution to the output of a query; the Gaussian variant was proven to preserve differential privacy by Dwork et al. [21] in a follow-up work. Since then, many sanitizer for preserving differential privacy has been proposed in the literature, including the Exponential mechanism [12, 40], the Multiplicative Update mechanism [26, 27, 28, 31], the Median mechanism [45], the Boosting mechanism [24], and the Random Projection mechanism [36]. All these mechanisms have a common theme: they perturb the output before responding to queries. Blocki et al. [9, 10] and Upadhyay [48, 49] took a complementary approach. They perturb the input by performing a random projection of the input and show that existing algorithms preserves differential privacy if the input is perturbed in a reversible manner.

2 Preliminaries, Notations, and Basic Definitions

NOTATIONS. We fix the letter n to denote the space of the input vectors, m to denote the number of vectors, and r to denote the subspace to which the vectors are projected. We use the symbol η to denote the approximation parameter in the statement of the Johnson-Lindenstrauss transform. We use the notation $\langle a_1, \cdots, a_n \rangle$ to denote the individual entries of an n-dimensional vector $|a\rangle$. We use the symbol W to denote an $n \times n$ Walsh-Hadamard matrix. We use $A_{1..r}$ to denote the matrix formed by taking the first r rows of A and $A_{i:}$ ($A_{:j}$) to denote i-th row (column, respectively) of matrix A. We use Dirac notation to represent vectors, i.e., $\langle \cdot |$ to represent row vector and $| \cdot \rangle$ to represent a column vector. We use bold faced capital letters, like $\mathbb{A}$, to denote n copies of matrix A stacked together row-wise. For a vector $|x\rangle$, we use the notation $\mathsf{Diag}(|x\rangle)$ to represent a diagonal matrix with non-zero entries $\langle x_1, \cdots, x_n \rangle$.

PRIVACY MODEL USED IN THIS PAPER. In this work, we deal with privacy-preserving mechanisms for answering cut-queries on a graph and directional covariance queries on a matrix. We work with the natural relaxed notion of differential privacy, known as *approximate differential privacy*.

Definition 1. A randomized mechanism, $\mathcal{K}$, gives (ε, δ)-*differential privacy*, if for all neighboring data-sets D_1 and D_2, and all range $S \subset Range(\mathcal{K})$, $\mathsf{Pr}[\mathcal{K}(D_1) \in S] \leq \exp(\varepsilon)\mathsf{Pr}[\mathcal{K}(D_2) \in S] + \delta$, where the probability is over the coin tosses of $\mathcal{K}$. When $\delta = 0$, we get the traditional definition of *differential privacy*.

We call two data-sets D_1 and D_2 are *neighboring* if $\|D_1 - D_2\| \leq 1$. The following lemma is key to our analysis in Section 3.

Lemma 2. Let $M(D)$ be a (ε, δ)-differential private mechanism for a database D , and let h be any function, then any mechanism $M' := h(M(D))$ is also (ε, δ)-differentially private for the same set of queries.

STATISTICAL MODEL SELECTION AND PROBABILITY THEORY. The main ingredients in our utility proof are inequalities from model selection. We review some of its basics and probability theory that are required to understand our proof. One of the main methods to prove concentration inequalities is the following two step process: control the moment generating function of a random variable and then minimize the upper bound resulting from the Markov's inequality. Though simple, it is extremely powerful.

Let ζ be a real valued centered random variable, then the log-moment generating function is defined as $\psi_\zeta(\lambda) := \ln(\mathbb{E}[\exp(\lambda\zeta)]), \forall \lambda \in \mathbb{R}_+$, and the *Cramer's transform* is defined as $\psi^*_\zeta(x) := \sup_{\lambda \in \mathbb{R}_+}(\lambda x - \psi_\zeta(\lambda)$. The *generalized inverse* of ψ^* at a point t is defined by $\psi^{*-1}(f) := \inf\{x \geq 0 : \psi^*(x) > f\}$.

The log generating function for centered random variable has some nice properties. It is continuously differentiable in a half-open interval $I = [0, b)$, where $0 < b \leq \infty$, and both ψ_ζ and its differentiation at 0 equals 0. There is a nice characterization of the generalized inverse in the form of following lemma.

Lemma 3. Let ψ be a convex continuously differentialable function on I. Assume that $\psi(0) = \psi'(0) = 0$. Then ψ^* is non-negative non-decreasing convex function on $\mathbb{R}_+$. Moreover, its generalized inverse can be written as $\psi^{*-1} = \inf_{\lambda \in I} \left[(f + \psi(\lambda))/\lambda\right]$.

This lemma follows from the definition and basic calculus. In the area of model selection, Lemma 3 is often used to control the expectation of the supremum of a finite family of exponentially integrable variables. Pisier [43] proved the following fundamental lemma.

Lemma 4. (Pisier [43]) Let $\{\zeta_f\}_{f \in F}$ be a finite family of random variables and ψ be as in Lemma 3. Let $\mathbb{E}^A[\zeta] = \mathbb{E}[\zeta \chi_A]/\mathsf{Pr}[A]$ for a non-zero measurable set A. Then, for any non-zero measurable set A, we have $\mathbb{E}^A\left[\sup_{f \in F} \zeta_f\right] \leq \psi^{*-1}\left(\ln\left(|F|/\mathsf{Pr}[A]\right)\right)$.

If we take $A = (\zeta \geq \phi(x))$ and applying Markov's inequality, then using the property that ϕ is an increasing function, this immediately gives us that $x \leq \ln(1/\mathsf{Pr}[A])$. This gives the following key lemma.

Lemma 5. Let A be a set with non-zero measure and ζ be a centered random variable. Let ϕ be an increasing function on positive reals such that $\mathbb{E}^A[\zeta] \leq \phi(\ln(1/\mathsf{Pr}[A]))$. Then $\mathsf{Pr}[\zeta \geq \phi(x)] \leq \exp(-x)$.

We refer the interested readers to the book by Massart and Picard [39]. In this paper, we use the following result by Birge and Massart [7] for bounding the utility, which follows from application of Lemma 5.

Theorem 6. (Birge-Massart [7]) Let $(\zeta_f)_{\mathcal{F}}$ be a finite family of random variable and ψ be a convex and continuously differentiable function on $[0, b)$ with $0 \leq b \leq \infty$ such that $\psi(0) = \psi'(0) = 0$ and for every $u \in [0, b)$ and $f \in \mathcal{F}$, we have $\log(\mathbb{E}[\exp(u\zeta_f)]) \leq \psi(u)$. If N denotes the cardinality of $\mathcal{F}$. Then $\mathbb{E}\left[\sup_{f\in\mathcal{F}} \zeta_f\right] \leq \psi^{*-1}(\ln N)$, where ψ^* is the Cramer's transformation.

Using Lemma 5 and Talagrand inequality, the authors also proved the following corollary.

Corollary 7. (Birge-Massart [7]) Let $0 < \lambda < 1/b$ for some b. If ζ be a real valued integrable variable, and a and b be constants such that $\log(\mathbb{E}[\exp(\lambda\zeta)]) \leq \frac{a\lambda^2}{2(1-b\lambda)}$. Then $\mathsf{Pr}\left[\zeta \geq \sqrt{2a\tau} + b\tau\right] \leq \exp(-\tau)$.

Linear Algebra and Probability Distributions. Our analysis of privacy makes extensive use of linear algebra and statistical properties of Gaussian distribution. We give an exposition to the level required to understand this paper. Let A be an $n \times d$ matrix. The singular value decomposition (SVD) of A is $A = V\Lambda U^{\mathsf{T}}$, where U, V are unitary matrices and Λ is a diagonal matrix consisting of the *singular values* of A. Since U and V are unitary matrices, one can write $A^i = V\Lambda^i U^{\mathsf{T}}$ for any real value i. We use standard Walsh-Hadamard matrix and discrete Fourier transform matrix. A Walsh-Hadamard matrix W_m is a $2^m \times 2^m$ matrix formed recursively as follows: $W_0 = 1$ and $W_m = \frac{1}{\sqrt{2}}\begin{pmatrix} W_{m-1} & W_{m-1} \\ W_{m-1} & -W_{m-1} \end{pmatrix}$. Where it is clear from the context, we drop the subscript. We use the symbol $\mathcal{F}$ to denote discrete Fourier transform.

A Rademacher sequence is a sequence of random variables having value ± 1 with probability $1/2$. Given a random variable, X, we denote by $X \sim \mathcal{N}(\mu, \sigma^2)$ the fact that X is distributed according to a Gaussian distribution with the probability density function, $\mathsf{PDF}_X(x) = \frac{1}{\sqrt{2\pi\sigma}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$. The Gaussian distribution is invariant under affine transformation. This is called *spherical symmetry* of Gaussian variable. The multivariate Gaussian distribution is a generalization of univariate Gaussian distribution. Given a m dimensional multivariate random variable, $X \sim \mathcal{N}(\mu, \Sigma)$ with mean $\mu \in \mathbb{R}^m$ and covariance matrix $\Sigma = \mathbb{E}[(X-\mu)(X-\mu)^{\mathsf{T}}]$, the PDF of a multivariate Gaussian is given by $\mathsf{PDF}_{\mathbf{X}}(|x\rangle) := \frac{1}{\sqrt{2\pi \det(\Sigma)}} \exp\left(-\frac{1}{2}\mathrm{Tr}\left(\langle x|\Sigma|x\rangle\right)\right)$. It is easy to see from the description of the PDF that, in order to define the PDF corresponding to a multivariate Gaussian distribution, Σ has to have full rank and is positive definite matrix (see Appendix B.2).

3 Circulant Matrices and Differential Privacy

In this section, we show that a general class of projection matrices of which Vybiral [51] is a special instance also preserves differential privacy. We also note a slight modification that allows us to get a tighter dimension bound, which is optimal up to a logarithmic factor. Our random projections have the form $\Phi = PWD$, where W is a Walsh-Hadamard transform, D is a diagonal matrix formed by Rademacher sequence, and P is formed by independently sampling r rows of a circulant matrix. The mechanisms for answering cut-queries and covariance queries follows by substituting our projection matrix instead of the random Gaussian matrix based Johnson-Lindenstrauss transform used in Blocki *et al.* [9]. We assume that the private matrix has a dimension $n \times d$ where n is a power of 2. This is without any loss of generality because we can simply append block of 0 matrix to make the number of rows a power of 2 while incurring at most constant overhead.

Description of the Matrix P and the Family of Projection Matrices. Let $\alpha := \langle \alpha_1, \cdots, \alpha_n \rangle$ be n i.d.d. Gaussian samples and C be a circulant matrix formed using the vector α, i.e, for $1 \leq i \leq n$, $C_{i:} = \langle \alpha_i, \cdots, \alpha_n, \alpha_1, \cdots, \alpha_{i-1} \rangle$. Then the matrix P corresponding to a permutation matrix Π truncated to

Construction of Φ: Set $r = O(\eta^{-2} \log m \log n)$. Construct the matrices D and P as below.

1. D is an $n \times n$ diagonal matrix such that $\mathsf{Pr}[D_{ii} = +1] = \mathsf{Pr}[D_{ii} = -1] = 1/2$.
2. P is a matrix formed as follows: Construct a circulant matrix C with entries picked from a Gaussian distribution and then independently sample r rows from C.

Compute $\Phi = PWD$, where W is the normalized $n \times n$ Walsh-Hadamard transform matrix.

Figure 1: A Family of Random Projection Matrices

r rows and circulant matrix C is formed by choosing r rows of the circulant matrix (i.e., $P = \Pi_{1..r}C$). This approach has been used by Rudelson and Vershynin [46] to prove that certain matrix satisfies the Restricted Isometry Property. This choice of sampling r rows combined by sampling a Gaussian vector gives us a family of matrices $\mathcal{P} = \{P\}_{\Pi}$, the size of this family being $\binom{n}{r}$. Of this, a special case is the partial circulant matrix, which was used by Vybiral [51], in which the first r rows of the matrix C is chosen deterministically to form the matrix P. Note that this is not the only option. One can also pick r rows deterministically by specifying some known fixed permutations, like Affine transformation or combinatorial designs, like r random rows of Latin squares[4], but they are also a special case of the family $\mathcal{P}$.

We first note few salient features of this class of projection matrices. The matrix P alone cannot be used for random projection because for some bad input vector $|x\rangle$, the estimate of $\|P|x\rangle\|_2$ can be really bad. For example, when $|x\rangle$ is along a single coordinate, then only the non-zero values of P along this coordinate would contribute to $P|x\rangle$, giving a very bad variance bound. This is why we need some preconditioning on the inputs–Vybiral [51] does this by using diagonal Rademacher matrix. However, as we argued earlier using the result by Krahmer and Ward [38] and Rauhut *et al.* [44], unless there is a significant improvement in the concentration bounds on partial circulant matrices, the Vybiral [51] result seems hard to improve. The intuitive reasoning is that the diagonal Rademacher matrix does not preconditions the input to a proper degree of isometry. For this, we need to precondition the input with WD instead of just D. The extra W helps in spreading out the vector in all direction. At a very high level, this isometry allows us to mimic a projection matrix with every entries picked i.d.d.

We conclude this section by giving the formal description of the projection matrix (Figure 1), privacy guarantee (Theorem 8), and a proof of the utility (Theorem 11).

Theorem 8. *Privacy Guarantee.* Let Φ be a $n \times r$ projection matrix constructed by transposing the construction in Figure 1. If the singular values of a private $n \times d$ matrix A is at least $\left(16 \ln(1/\delta)\sqrt{n/\varepsilon}\right)$. Then $A^{\mathsf{T}}\Phi$ preserves (ε, δ)-differential privacy. Moreover, the computation requires $O(nd \log n)$ basic operations.

A remark about the above theorem is due here. Note that we have a factor of n instead of a factor of r as in Blocki *et al.* [9]. However, as mentioned by the authors [9], in order to answer all cut (or covariance) queries, r has to be set at least equal to n. In other words, the mechanism does not perform dimension reduction, rather it increases the dimension. In that respect, differential privacy is distinct from all the other known applications of the Johnson-Lindenstrauss transform. However, we loss an extra $\sqrt{1/\varepsilon}$ factor in the singular value term than in Blocki *et al.* [9]. This is not surprising as we expect to pay the price for faster and randomness efficient computation in some or the other way.

Proof. Before we give our proof, we argue why the proof of Blocki *et al.* [9] does not extend to our case. One of the reasons why the proof of Blocki *et al.* [9] does not generalize to any Johnson-Lindenstrauss transform

[4]Latin squares are combinatorial design in which a square $n \times n$ matrices have entries from 1 to n such that every row and column have all the entries.

in general is due to its strong dependency on the fact that each samples in a dense Gaussian matrices is picked i.d.d. More concretely, each row of their published matrix is a multivariate Gaussian and preserves differential privacy. This allows them to use the composition theorem of Dwork, Rothblum and Vadhan [24] to prove differential privacy of the entire published matrix. Unfortunately, we cannot invoke the composition theorem because, as we reuse the random samples, we introduce correlations between the entries of our projection matrices. Therefore, applying our projection matrix to a private matrix does not yield n independent multivariate distribution, but one matrix-variate distribution. We compute the resulting distribution and then prove that it preserves differential privacy. In this sense, our proof uses the same idea as used by Upadhyay [48]; however, the analogy ends here as the probability distribution are very different and requires a fresh analysis.

Our starting point is an alternate way to look at any matrix $P \in \mathcal{P}$, i.e., a matrix formed by sampling r rows of a fully circulant matrix independently. In other words, one can see P as a product of a truncated permutation matrix and a circulant matrix formed by α. Let $\mathbb{I}_k$ denote the $k \times k$ identity matrix. Since differential privacy is preserved under any arbitrary post-processing, we can just concentrate on the distribution when a fully circulant matrix is used instead of the matrix P (the truncated permutation can be seen as a post-processing step). For example, for a partial circulant matrix, we have,

$$P = \begin{pmatrix} \alpha_1 & \alpha_2 & \cdots & \alpha_n \\ \alpha_2 & \cdots & \alpha_n & \alpha_1 \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_r & \alpha_1 & \cdots & \alpha_{r-1} \end{pmatrix} = \underbrace{\begin{pmatrix} \mathbb{I}_r & 0 \end{pmatrix}}_{n \text{ columns}} \begin{pmatrix} \alpha_1 & \alpha_2 & \cdots & \alpha_n \\ \alpha_2 & \cdots & \alpha_n & \alpha_1 \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_n & \alpha_1 & \cdots & \alpha_{n-1} \end{pmatrix} = \begin{pmatrix} \mathbb{I}_r & 0 \end{pmatrix} C. \tag{2}$$

Therefore, for the rest of this proof, we just concentrate on fully-circulant matrix. Let denote by $\mathsf{vec}(C)$ the vector formed by the entries of C. Then, the covariance matrix of $\mathsf{vec}(C)$ is,

$$\Lambda := \mathsf{COV}(\mathsf{vec}(C)) = \underbrace{\begin{pmatrix} \mathbb{I}_{n/2} & 0 & 0 & \mathbb{I}_{n-1} & 0 & \mathbb{I}_{n-2} & \cdots & 0 & \mathbb{I}_1 \\ 0 & \mathbb{I}_{n/2} & \mathbb{I}_1 & 0 & \mathbb{I}_2 & 0 & \cdots & \mathbb{I}_{n-1} & 0 \\ 0 & \mathbb{I}_1 & \mathbb{I}_{n/2} & 0 & \cdots & \cdots & \cdots & \vdots & \vdots \\ \mathbb{I}_{n-1} & 0 & 0 & \mathbb{I}_{n/2} & & \cdots & \cdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & \mathbb{I}_{n-2} & \vdots & \vdots & \vdots & \ddots & \cdots & \mathbb{I}_{n/2+1} & 0 \\ \mathbb{I}_2 & 0 & \cdots & \cdots & \cdots & \cdots & \ddots & 0 & \mathbb{I}_{n/2-1} \\ 0 & \mathbb{I}_{n-1} & 0 & \mathbb{I}_{n-2} & \cdots & \cdots & \cdots & \mathbb{I}_{n/2} & 0 \\ \mathbb{I}_1 & 0 & \mathbb{I}_2 & 0 & \cdots & & \cdots & 0 & \mathbb{I}_{n/2} \end{pmatrix}}_{n^2 \text{ columns}}, n^2 \text{ rows} \tag{3}$$

where 0 are block zero matrices of appropriate dimensions. We first note that WD does not effect the privacy. This is because of the spherical symmetry of a vector of Gaussian distribution. Also note that WD is norm-preserving; therefore, WDA and WDA' are also neighbouring matrices if A and A' are. Therefore, without any loss of generality, we can analyze the distribution $A^{\mathsf{T}}C$ instead of $A^{\mathsf{T}}\Phi$. Another way to look at it is that differential privacy is preserved under arbitrary post-processing.

In order to compute the PDF of the matrix-variate distribution corresponding to the published matrix, we follow the standard technique. We look at the published matrix as a vector and analyze the corresponding multivariate distribution. Recall that the published matrix is not an n independent multi-variate distribution; rather one matrix-variate distribution and there are non-trivial correlations between the entries of two rows of the published matrix as clear by equation (3). In Lemma 18, we prove that a covariance matrix is a positive semi-definite matrix; therefore, we can write equation (3) succinctly in form of its Cholesky decomposition, say $\Lambda = LL^{\mathsf{T}}$.

Note that $\Lambda\Lambda^T = n\mathbb{I}$. Using the left spherical symmetry of Gaussian distribution, and since the Jacobian of the transformation $A^{\mathsf{T}}C$ is $\sqrt{\det(A^{\mathsf{T}}A)}$, the resulting matrix variate distribution for $X \sim A^{\mathsf{T}}Y$ for Y picked from a distribution with mean vector 0 and covariance matrix Λ has the covariance matrix $\mathbb{A}^{\mathsf{T}}\Lambda\mathbb{A}$, where $\mathbb{A}$ is matrix formed by stacking n copies of A row-wise. Therefore,

$$\mathsf{PDF}_{A^{\mathsf{T}}C}(X) = \frac{1}{\sqrt{\det(\mathbb{A}^{\mathsf{T}}\Lambda\mathbb{A})}} \exp\left(-\frac{1}{2}\mathrm{Tr}\left(X^{\mathsf{T}}(\mathbb{A}^{\mathsf{T}}\Lambda\mathbb{A})^{-1}X\right)\right). \tag{4}$$

For the sake of simplicity, let us denote by $\mathbb{B} = L^{\mathsf{T}}\mathbb{A}$. Let the singular value decomposition of $\mathbb{B} = \mathbb{U}\Sigma\mathbb{V}^{\mathsf{T}}$. Similarly, let define $\tilde{\mathbb{B}} = L^{\mathsf{T}}\tilde{\mathbb{A}} = \tilde{\mathbb{U}}\tilde{\Sigma}\tilde{\mathbb{V}}^{\mathsf{T}}$. Then from equation (4), we can write the distribution of the published matrices corresponding to the neighbouring matrices A and $\tilde{A}$ as follows.

$$\mathsf{PDF}_{A^{\mathsf{T}}C}(X) = \frac{1}{\sqrt{\det(\mathbb{B}^{\mathsf{T}}\mathbb{B})}} \exp\left(-\frac{1}{2}\mathrm{Tr}(X^{\mathsf{T}}(\mathbb{B}^{\mathsf{T}}\mathbb{B})^{-1}X)\right),$$
$$\mathsf{PDF}_{\tilde{A}^{\mathsf{T}}C}(X) = \frac{1}{\sqrt{\det(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})}} \exp\left(-\frac{1}{2}\mathrm{Tr}(X^{\mathsf{T}}(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})^{-1}X)\right).$$

In order to prove the differential privacy, we prove the following lemma.

Lemma 9. For a matrix A with all singular values greater than $\Omega\left(\sqrt{\frac{n}{\varepsilon}}\log(4/\delta)\right)$, the following holds

$$\sqrt{\frac{\det(\mathbb{B}^{\mathsf{T}}\mathbb{B})}{\det(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})}} \in \exp(\pm\varepsilon). \tag{5}$$

If $X = \mathbb{A}^{\mathsf{T}}C$, then

$$\mathsf{Pr}\left[\left|\mathrm{Tr}\left(X^{\mathsf{T}}\left((\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})^{-1} - (\mathbb{B}^{\mathsf{T}}\mathbb{B})^{-1}\right)X\right)\right| \leq \varepsilon\right] \geq 1-\delta. \tag{6}$$

Proof. The first part of the proof follows simply as in Blocki *et al.* [9]. More concretely, we have $\det(\mathbb{B}^{\mathsf{T}}\mathbb{B}) = \left(\prod_i \sigma_i^2\right)n$, where $\sigma_1 \geq \cdots \geq \sigma_d \geq \sigma_{\mathsf{min}}$ are the singular values of B. Let $\tilde{\sigma}_1 \geq \cdots \geq \tilde{\sigma}_d \geq \sigma_{\mathsf{min}}$ be its singular value for $\tilde{B}$. Since the singular values of $B - \tilde{B}$ and $\tilde{B} - B$ are the same, $\sum_i(\sigma_i - \tilde{\sigma}_i) \leq 1$ using Linskii's theorem. Therefore,

$$\frac{\det(\mathbb{B}^{\mathsf{T}}\mathbb{B})}{\det(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})} = \left(\prod_i \frac{\tilde{\sigma}_i^2}{\sigma_i^2}\right) \leq \exp\left(\frac{\varepsilon}{8\log(2/\delta)}\right)\sum_i(\tilde{\sigma}_i - \sigma_i) \leq \exp(\varepsilon).$$

Similarly, we can bound $\frac{\det(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})}{\det(\mathbb{B}^{\mathsf{T}}\mathbb{B})} \leq \exp(\varepsilon)$.

PROOF OF EQUATION (6). In this part, we bound the following expression.

$$\left|\mathrm{Tr}\left(X^{\mathsf{T}}\left(\left(\mathbb{B}^{\mathsf{T}}\mathbb{B}\right)^{-1} - \left(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}}\right)^{-1}\right)X\right)\right|. \tag{7}$$

We can write $\tilde{A} = A + |v\rangle\langle e_i|$ for some i and a unit vector v. Let $\mathcal{E}$ be a matrix formed by n-copies of $|e_i\rangle\langle v|$ stacked together. The following is immediate.

$$\begin{aligned}
\mathrm{Tr}\left(X^{\mathsf{T}}\left(\left(\mathbb{B}^{\mathsf{T}}\mathbb{B}\right)^{-1}-\left(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}}\right)^{-1}\right)X\right) &= \mathrm{Tr}\left(X^{\mathsf{T}}\left((\mathbb{B}^{\mathsf{T}}\mathbb{B})^{-1}(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})^{-1}-(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})^{-1}\right)X\right)\\
&= \mathrm{Tr}\left(X^{\mathsf{T}}\left((\mathbb{B}^{\mathsf{T}}\mathbb{B})^{-1}(\mathbb{B}+\mathcal{E})^{\mathsf{T}}(\mathbb{B}+\mathcal{E})(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})^{-1}-(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})^{-1}\right)X\right)\\
&= \mathrm{Tr}\left(X^{\mathsf{T}}\left((\mathbb{B}^{\mathsf{T}}\mathbb{B})^{-1}(\mathbb{B}^{\mathsf{T}}\mathcal{E}+\mathcal{E}^{\mathsf{T}}\tilde{\mathbb{B}})(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})^{-1}\right)X\right)\\
&= \mathrm{Tr}\left(C^{\mathsf{T}}A\left((\mathbb{B}^{\mathsf{T}}\mathbb{B})^{-1}(\mathbb{B}^{\mathsf{T}}\mathcal{E}+\mathcal{E}^{\mathsf{T}}\tilde{\mathbb{B}})(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})^{-1}\right)A^{\mathsf{T}}C\right)\\
&\leq \mathrm{Tr}\left(CC^{\mathsf{T}}\right)\mathrm{Tr}\left(A\left((\mathbb{B}^{\mathsf{T}}\mathbb{B})^{-1}(\mathbb{B}^{\mathsf{T}}\mathcal{E}+\mathcal{E}^{\mathsf{T}}\tilde{\mathbb{B}})(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})^{-1}\right)A^{\mathsf{T}}\right)\\
&= \underbrace{\mathrm{Tr}\left(CC^{\mathsf{T}}\right)}_{S}\underbrace{\mathrm{Tr}\left(A^{\mathsf{T}}A\left((\mathbb{B}^{\mathsf{T}}\mathbb{B})^{-1}(\mathbb{B}^{\mathsf{T}}\mathcal{E}+\mathcal{E}^{\mathsf{T}}\tilde{\mathbb{B}})(\tilde{\mathbb{B}}^{\mathsf{T}}\tilde{\mathbb{B}})^{-1}\right)\right)}_{Q},
\end{aligned}$$

where the inequality follows from the fact that $\mathrm{Tr}(XY) \leq \mathrm{Tr}(X)\mathrm{Tr}(Y)$ for Hermittian matrices X and Y. In fact, the two matrices in question are positive semi-definite. We bound each of the above trace terms. To bound S, we recall the fundamental relation between discrete Fourier transform and circulant matrices. Recall that C is made by circulating the Gaussian vectors $\langle \alpha_1, \cdots, \alpha_n \rangle$ to form an $n \times n$ matrix. Then

$$P = \mathcal{F}_n \mathsf{Diag}(\sqrt{n}\mathcal{F}_n\alpha)\mathcal{F}_n^{-1}. \tag{8}$$

Therefore, to bound the trace of CC^{T}, we have to bound the following.

$$PP^{\mathsf{T}} = \mathcal{F}_n \mathsf{Diag}(\sqrt{n}\mathcal{F}_n\alpha)\mathcal{F}_n^{-1}\left[\mathcal{F}_n \mathsf{Diag}(\sqrt{n}\mathcal{F}_n\alpha)\mathcal{F}_n^{-1}\right]^{\mathsf{T}} = \mathcal{F}_n \mathsf{Diag}(n|\mathcal{F}_n\alpha|^2)\mathcal{F}_n^{-1}. \tag{9}$$

Since α and $Z\alpha$ are equidistributed when the rows of Z are orthonormal, we need the following lemma to bound S.

Lemma 10. Let $\beta_1, \cdots, \beta_n$ be n i.d.d. $\mathcal{N}(0,1)$ random variables. Then,

$$\mathsf{Pr}\left[\sum_{i=1}^{n}\beta_i^2 > 2(1+\theta)n\right] \leq 2^{-\theta n/2}.$$

Proof. First, from the definition of normal distribution, we know that $\mathsf{Pr}[\beta_i = t] = \frac{1}{\sqrt{2\pi}}\exp\left(-t^2/2\right)$. Then consider the random variable $Z_i = \exp(\beta_i^2/4)$. Then

$$\mathbb{E}[Z_i] = \int_{-\infty}^{\infty}\frac{1}{\sqrt{2\pi}}\exp\left(-t^2/2\right)\exp(-t^2/4)\,\mathrm{d}t = \sqrt{2}.$$

Now, observe that,

$$\begin{aligned}
\mathsf{Pr}_{\beta_1,\cdots\beta_n}[\beta_1^2+\cdots\beta_n^2 > \lambda] &= \mathsf{Pr}_{\beta_1,\cdots,\beta_n}\left[\frac{\beta_1^2+\cdots\beta_n^2}{4} > \frac{\lambda}{4}\right]\\
&= \mathsf{Pr}_{\beta_1,\cdots,\beta_n}\left[\exp\left(\frac{\beta_1^2+\cdots\beta_n^2}{4}\right) > \exp\left(\frac{\lambda}{4}\right)\right]\\
&\leq \exp(-\lambda/4)\mathbb{E}_{\beta_1,\cdots,\beta_n}\left[\exp\left(\frac{\beta_1^2+\cdots\beta_n^2}{4}\right)\right].
\end{aligned}$$

Since all β_i are i.d.d., the above expression is bounded as

$$\prod_{i=1}^{n} \mathbb{E}\left[\exp\left(\frac{\beta_i^2}{4}\right)\right] = \prod_{i=1}^{n} \mathbb{E}\left[Z_i\right] = 2^{n/2}.$$

Putting $\lambda = 2(1+\theta)n$, the lemma follows. □

The term Q is easy to compute once we note the following.

$$n\mathsf{Tr}\left(A^{\mathsf{T}}A\right) \leq \mathsf{Tr}\left(\mathbb{V}\Sigma^{-2}\mathbb{V}^{\mathsf{T}}\right) = \mathsf{Tr}\left(\mathbb{B}^{\mathsf{T}}\mathbb{B}\right) \leq n^{3/2}\mathsf{Tr}\left(A^{\mathsf{T}}A\right)$$

since $\Lambda\Lambda^{\mathsf{T}} = n\mathbb{I}_{n^2\times n^2}$ and $\mathrm{Tr}(\Lambda) = n^{3/2}$. Now e_i and v are unit vectors which forms the marix $\mathcal{E}$, and the singular values of $A, \tilde{A}$ are at least σ_{min} with $\tilde{A} - A = |v\rangle\langle e_i|$. Combining with the invariance of trace of matrices under cyclic permutations, it is easy to see that Q by bounded by at most $1/\sigma_{\mathsf{min}}^2\left(1/\sigma_{\mathsf{min}} + 1/\sigma_{\mathsf{min}}^2\right)$. Using Lemma 10, equation (9), and the bound on σ_{min}, we have

$$\mathsf{Pr}\left[(7) \leq \frac{1}{\sigma_{\mathsf{min}}^2}\left(\frac{1}{\sigma_{\mathsf{min}}} + \frac{1}{\sigma_{\mathsf{min}}^2}\right)\frac{4n^2\ln(4n/\delta)}{n} \leq 5\varepsilon\right] \geq 1-\delta.$$

Rescaling the value of ε, the lemma follows. □

It is straightforward to see that Lemma 9 implies the privacy guarantee in Theorem 8. For the run-time guarantee, note that a circulant matrix has a singular value decomposition in form of discrete Fourier transform. Therefore, in disguise, all the matrices used in the projection matrix allow fast matrix-vector multiplication. The final truncated permutation matrix takes $O(r)$ time to sample the corresponding entry after the application of CWD; therefore, the run-time of a single matrix-vector multiplication takes $O(n\log n)$ time. Since, there are d columns in the matrix A, it takes $O(nd\log n)$ time to publish the sanitized matrix. □

Theorem 11. *Utility Guarantee.* Let Φ be as in Figure 1. Then for any set of m vectors S in $\mathbb{R}^n$, there is an $r = O(\eta^{-2}\log n\log m)$ such that the following holds with the probability at least $2/3$,

$$(1-\eta)\sqrt{r}\||x\rangle\|_2^2 \leq \|\Phi|x\rangle\|_2^2 \leq (1+\eta)\sqrt{r}\||x\rangle\|_2^2, \qquad \forall|x\rangle \in S. \tag{10}$$

Proof. The usual idea in proving the Johnson Lindenstrauss lemma is to first bound the expectation of the random variables corresponding to the output of an application of the projection matrix, and then use the standard concentration bound to prove the result. We use the same idea. We first proof the result when we use a Walsh-Hadamard matrix instead of discrete-Fourier transform.

We break the analysis in two parts. We first use the fact that WD is an isometry [1], i.e., for any vector $|x\rangle$ of unit length, $WD|x\rangle$ has bounded co-ordinates. Then, we use this promise to prove the following: when we multiply a circulant matrix formed by Gaussian vector from the right to this smoothen vector and sample r rows, then this preserves the Euclidean norm with high probability.

Since the transformation is linear, without loss of generality, we can assume $|x\rangle$ is a unit vector. Fix a $|x\rangle \in \mathcal{S}$. The first step follows simply from the following result by Ailon and Chazelle [1].

Theorem 12. *(Ailon-Chazelle [1], Wolff [54, Proposition 4.2])* Let $|x\rangle \in \mathbb{R}^n$ and $t > 0$. Let W and D be as above. Then, for any $\kappa > 0$, we have $\mathsf{Pr}\left[\|WD|x\rangle\|_\infty \geq \sqrt{2e/n}\log\left(2n/\kappa\right)\langle x, x\rangle\right] \leq \kappa$.

Bounding the expectation. The second step is to use the guarantee that $\||\tilde{x}\rangle\|_\infty = O(n^{-1/2}\sqrt{\log m})$ to get the desired expectation bound. The naive method to work with the permutation in the matrix Φ to get the concentration result makes the proof very lengthy. The crucial observation here is that a circulant matrix formed by a vector of i.d.d. Gaussian is very symmetric; therefore, picking any set of r would have the same concentration properties as picking the first r rows. We follow the approach taken by Upadhyay [50]. This

is possible because of the nice representation of a partial circulant matrices by a discrete Fourier transform matrix. We first get around the problem of dealing with the permutation matrices by making a substitution, i.e., for the matrix Φ and any vector $|x\rangle \in \mathbb{R}^n$, $\|\Phi|x\rangle\|_2 = \|Z\alpha\|_2$, where the entries of Z are $z_{i,j} = (\Pi_{1..r})_{i:}(\mathsf{Diag}(|\tilde{x}\rangle))_{:j}$. The rest of the proof is very similar to [51, 50] owing to our observation in equation (9); we include it for the sake of completion.

Let $U\Sigma V^{\mathsf{T}}$ be the SVD of Z, $\gamma = V^{\mathsf{T}}\alpha$. Let $\sigma := \langle \sigma_1, \cdots, \sigma_r \rangle$ be the singular values of Z and $\langle \gamma_1, \cdots, \gamma_r \rangle$ be the co-ordinates of $V^{\mathsf{T}}\alpha$. Making this substitution, we have the following equalities.

$$\begin{aligned}
\mathsf{Pr}_\alpha \left[\|\Phi|x\rangle\|_2^2 \geq (1+\eta)\right] &= \mathsf{Pr}_\alpha \left[\|WZ\alpha\|_2^2 \geq (1+\eta)r\right] = \mathsf{Pr}_\alpha \left[\|Z\alpha\|_2^2 \geq (1+\eta)r\right] \\
&= \mathsf{Pr}_\alpha \left[\|U\Sigma V^{\mathsf{T}}\alpha\|_2^2 \geq (1+\eta)r\right] = \mathsf{Pr}_\gamma \left[\|U\Sigma\gamma\|_2^2 \geq (1+\eta)r\right] \\
&= \mathsf{Pr}_\gamma \left[\|\Sigma\gamma\|_2^2 \geq (1+\eta)r\right]. \qquad (11)
\end{aligned}$$

In other words, if we can prove the concentration bound on $\sum \sigma_i^2 |\gamma_i|^2$, we are done. We use the Corollary 7 to Theorem 6, for which we need to find the function ψ corresponding to our case. Let $0 < \lambda < 1/2a$. The following proposition, which we prove in Appendix A, follows by simple arithmetic and linearity of expectation.

Proposition 13. Let $Y_1, \cdots Y_r$ be picked from $\mathcal{N}(0,1)$ and $\sigma = \langle \sigma_1, \cdots, \sigma_r \rangle$ be an r dimensional vector. Let λ be an arbitrary constant such that $0 < \lambda < 1/2\|\sigma\|_\infty$. Then

$$\sum_{i=1}^{r} \log\left(\mathbb{E}_{Y_i}\left[\exp\left(\lambda\sigma_i^2(Y_i^2-1)\right)\right]\right) \leq \frac{\lambda^2 \sum_{i=1}^r \sigma_i^4}{1 - 2\lambda \max_i |\sigma_i|^2}.$$

Since $\Sigma = \mathsf{Diag}(\sigma_1, \cdots, \sigma_r)$, $Z = U\Sigma V^T$, and Gaussian distribution is invariant if we multiply with a matrix with orthonormal rows on the right, we can restate Proposition 13 as

$$\sum_{i=1}^{r} \log\left(\mathbb{E}_{Y_i}\left[\exp(\lambda\sigma_i^2(\gamma_i^2-1))\right]\right) \leq \frac{\lambda^2\|Z\|_2^4}{1 - 2\|Z\|_\infty^2\lambda} = \frac{2\lambda^2\|Z\|_2^4}{2(1 - 2\|Z\|_\infty^2\lambda)}.$$

The right hand side has the form $\psi(u) = \frac{a\lambda^2}{2(1-b\lambda)}$ for $a = 2\|Z\|_2^4$ and $b = 2\|Z\|_\infty^2$. Using Corollary 7, we have

$$\mathsf{Pr}_\gamma \left[\sum_{i=1}^{r} \sigma_i^2\left(\gamma_i^2 - 1\right) \geq 2\|Z\|_\infty^2\tau + 2\|Z\|_2^2\sqrt{\tau}\right] \leq \exp(-\tau). \qquad (12)$$

We need to estimate $\|Z\|_\infty$ and $\|Z\|_2$. This is where the guarantee on $\||\tilde{x}\rangle\|_\infty$ is useful. Using Theorem 12, with probability $19/20$ and the symmetry of the matrices, we have

$$\|Z\|_\infty^2 = \max_{|x\rangle \in \mathbb{R}^n, \||x\rangle\|_2 = 1} \|Z|x\rangle\|_2^2 \leq n\|WD|x\rangle\|_\infty^2 = n\||\tilde{x}\rangle\|_\infty^2 = 2\log(40n). \qquad (13)$$

Since $\|Z\|_F = \sum_{i=1}^r \sigma_i^2 = r$. Thus,

$$\|Z\|_2^2 \leq \|Z\|_F \cdot \|Z\|_\infty = 2r\log(40n). \qquad (14)$$

Since $\sum_j \sigma_j^2 = r$, by setting $\tau = cr\eta^2/\log(40n)$ for a small constant c, and using equations (11), (12), (13), and (14), we have

$$\mathsf{Pr}_\alpha[\|\Phi|x\rangle\|_2^2 \geq (1+\eta)] = \mathsf{Pr}_\gamma \left[\sum_{j=1}^{r} \sigma_j^2(|\gamma_j| - 1) \geq \eta r\right] < \exp\left(-\frac{r\eta^2}{\log(40n)}\right). \qquad (15)$$

For $(15) < 1/10m$, we need $r = O\left(\eta^{-2} \log n \log m\right)$. The result follows using the union bound and similar analysis for the negative side of the tail, i.e., for the value of r, we have

$$\mathsf{Pr}_\alpha[\|\Phi|x\rangle\|_2^2 \leq (1-\eta)] = \mathsf{Pr}_\gamma\left[\sum_{j=1}^{r} \sigma_j^2(|\gamma_j| - 1) \leq -\eta r\right] < \frac{1}{10m}. \tag{16}$$

Combining equation (15) and equation (16) and using union bound over all $x \in \mathcal{S}$, the result follows. □

As we mentioned in Section 1, our proof extends to give a simpler privacy proof for the projection matrix given by Upadhyay [50]. We first recall their construction.

CONSTRUCTION 2. Let D be a diagonal Bernoulli matrix and M be a diagonal Gaussian matrix. Let Π and Π' be permutation matrices. Then Upadhyay [50] showed that $\Pi_{1..r} M \Pi' W D$ satisfies the Johnson-Lindenstrauss bound, where $\Pi_{1..r}$ is a matrix restricted to the first r rows of the permutation matrix Π. We follow up with the details as to how we can mould the proof of Theorem 8 to the case of Construction 2.

Theorem 14. Let Φ be a $n \times r$ projection matrix constructed by transposing the matrix of Construction 2. If the singular values of an $n \times d$ matrix A is at least $\left(16\sqrt{n/\varepsilon}\ln(1/\delta)\right)$. Then for any private input matrix A, $A^\mathsf{T}\Phi$ preserves (ε, δ)-differential privacy. Moreover, the computation requires $O(nd \log r)$ basic operations.

Proof. Our proof reduces the problem of proving privacy for Construction 2 to that for Theorem 8. First note that for two neighbouring matrices A and A', $\|A - \tilde{A}\| = \|WD(A - \tilde{A})\| \leq 1$. Also, note that $A^\mathsf{T}\Pi_{1..r}$ and $\tilde{A}\Pi_{1..r}$ differs by at most one row by a unit entry depending on whether $\Pi_{1..r}$ picks that row or not. Therefore, $\|(A^\mathsf{T} - \tilde{A}^\mathsf{T})\Pi_{1..r}\| \leq 1$ given that $\|A - \tilde{A}\| \leq 1$. Moreover, for discrete Fourier transform $\mathcal{F}_n$, we also have $\|(A^\mathsf{T} - \tilde{A}^\mathsf{T})\Pi_{1..r}\mathcal{F}_n\| \leq 1$ given that $\|(A^\mathsf{T} - \tilde{A}^\mathsf{T})\Pi_{1..r}\| \leq 1$ because $\|\mathcal{F}_n\| = 1$. Let $B = A^\mathsf{T}\Pi_{1..r}\mathcal{F}_n$ and $\tilde{B} = \tilde{A}\Pi_{1..r}\mathcal{F}_n$. Therefore, proving Theorem 14 reduces to that of proving that for $B^\mathsf{T}\mathcal{F}_n^\mathsf{T} M \Pi' W D$. Also, from Lemma 2, we have that proving privacy of $B^\mathsf{T}\mathcal{F}_n^\mathsf{T} M \Pi' W D$ is equivalent to proving privacy for $B^\mathsf{T}\mathcal{F}_n^\mathsf{T} M \mathcal{F}_n$. Now recall that M is $\mathsf{Diag}(\alpha)$ for Gaussian vector α. Therefore, M is distributed equivalent to $\mathsf{Diag}(\mathcal{F}_n\alpha)$. In other words, proving privacy for $B^\mathsf{T}\mathcal{F}_n^\mathsf{T} M \mathcal{F}_n$ is equivalent to proving privacy for $B^\mathsf{T}\mathcal{F}_n^\mathsf{T}\mathsf{Diag}(\mathcal{F}_n\alpha)\mathcal{F}_n$. Using equation (9), Theorem 14 follows. The run-time efficiency is straightforward to compute. □

4 Implications of our Projection Matrices

In this section, we give two applications of our projection matrix where we improve the run time and the random samples required in the case of [9, 49]. Our mechanism gives an improvement wherever random projections have been used to sanitize data-base, but we restrict our attention to just these two cases.

Cut Queries on a Graph. Blocki *et al.* [9] gave the first mechanism that uses random projection to answer cut-queries on a graph. They achieve the best additive error bound; however, their mechanism takes $O(n^{2.38})$ basic operations to publish a sanitized graph assuming we use Coppersmith-Winograd's matrix multiplication. Upadhyay [49] showed that we can use graph sparsification in composition with random projection to improve the run time to $O(n^{2+o(1)})$, but this comes at the price of increased sampling cost. Using our projection matrix shown in Figure 1, we achieve the following guarantee (see Figure 2 for the formal description of the mechanism).

Theorem 15. Let G be a graph on n-vertices. There exists a mechanism that published a sanitized graph $\tilde{G}$ in $O(n^2 \log n)$ time using $2n + r$ random samples such that, for every $\alpha, \beta > 0$, one can compute all possible cut-queries on G with an additive error at most $\tilde{O}(s\sqrt{n/\varepsilon})$ while preserving (ε, δ)-differential privacy.

Proof. The proof of privacy follows from Theorem 8 and the construction of Blocki *et al.* [9], while the proof of utility follows from plugging in the guarantee of Theorem 11 in the computation done by Blocki *et al.* [9, Theorem 3.2]. The details are in Appendix B.1. □

Note that almost the same additive bound is achieved by Blocki *et al.* [9] and Upadhyay [49] with a factor of $\tilde{O}(n)$ more random samples (we have a loss of factor $\sqrt{1/\varepsilon}$). Moreover, Blocki *et al.* [9] requires time $O(n^{2.38})$ in comparison to $O(n^2 \log n)$ time taken by the mechanism in Figure 2 to compute the sanitized graph.

Directional Covariance Queries on a Matrix. Blocki *et al.* [9] also gave a mechanism that uses random projection to answer covariance-queries with least additive error among all existing mechanisms. However, their mechanism requires $O(n^{2.38})$ and uses nr Gaussian samples, which amounts to n^2 if one wishes to answer all covariance queries. Plugging in our projection matrix instead of random Gaussian matrix in their mechanism (see Figure 3) amounts to the following.

Theorem 16. Let A be an $n \times d$ matrix. There exists a mechanism that published a sanitized matrix $\tilde{A}$ in $O(nd \log n)$ time and using $2n + r$ random samples such that, for any unit vector $|x\rangle$, one can compute its covariance with A with an additive error at most $\tilde{O}(n/\varepsilon)$ while preserving (ε, δ)-differential privacy.

Proof. The proof of privacy follows from Theorem 8 and the construction of Blocki *et al.* [9], while the proof of utility follows from plugging in the guarantee of Theorem 11 in the computation done by Blocki *et al.* [9, Theorem 4.2]. The details are in Appendix B.2. □

5 Conclusion and Future Works

In this paper, we modified and generalized a known construction of the Johnson-Lindenstrauss transform by Vybiral [51] and proved that it preserves differential privacy with the same additive error as in comparison to Blocki *et al.* [9] and Upadhyay [49], which by far achieve the best bound compared to other mechanisms (see [9, 49] or Table 1 for more details). We exhibited a counter-intuitive result that less randomness than the number of entries of the data-base also preserve differential privacy. We believe the reason why this is true is because multiplying two matrices distribute the noise throughout over the private matrix.

This work leaves several open questions. Of particular interest is whether sparse variants of the Johnson-Lindenstrauss transform preserves differential privacy or not. Few constructions of such transforms, like [18], use linear sampling to achieve sparsity. We do not hope to see any improvement on the utility guarantee, but they would improve the running time by $\log n$ factor if the sparse transform under study achieves optimal dimension bound. An interesting problem relates to the problem of error amplification. The question is whether we can introduce some error-correction techniques to the problem? Any positive result in this direction would help reduce the additive error.

In the context of this work, one major open problem is to find a non-trivial lower bounds on the sampling complexity and multiplicative noise. There are tight lower bounds known for the sampling complexity of the samplers [25] and for additive noise in a differentially-private mechanisms [20, 32, 34]. We believe we could use some ideas from these lower bound results to give lower bounds on the sampling complexity and multiplicative noise of differentially private mechanisms. Any such lower bounds, even in the non-interactive setting, would help in our understanding of the gap, if any, between traditional privacy and differential privacy.

As Blocki *et al.* [9] mentioned, one of the open problem is whether we can use our mechanism to compute differentially private low-rank approximation of a matrix. There have been some recent activity using Gaussian matrices, starting with the work of Hardt and Roth [29, 30]. In low-rank approximation, one first computes the range of the projection and then perform the actual projection. Our result already gives a differentially private mechanism for the range finding step, but the private second step is still elusive.

Acknowledgements. I would like to thank the anonymous reviewers of TCC-2014 for their useful feedbacks, especially citation to Mironov [41], and finding a flaw in the earlier draft.

References

[1] Nir Ailon and Bernard Chazelle. The Fast Johnson–Lindenstrauss Transform and Approximate Nearest Neighbors. *SIAM J. Comput.*, 39(1):302–322, 2009. 3, 4, 11

[2] Nir Ailon and Edo Liberty. Fast dimension reduction using Rademacher series on dual BCH codes. In *SODA*, pages 1–9, 2008. 2

[3] Nir Ailon and Edo Liberty. Fast Dimension Reduction Using Rademacher Series on Dual BCH Codes. *Discrete & Computational Geometry*, 42(4):615–630, 2009. 2

[4] Nir Ailon and Edo Liberty. An Almost Optimal Unrestricted Fast Johnson-Lindenstrauss Transform. *ACM Transactions on Algorithms*, 9(3):21, 2013. 2

[5] Richard G Baraniuk and Michael B Wakin. Random projections of smooth manifolds. *Foundations of computational mathematics*, 9(1):51–77, 2009. 2

[6] Amos Beimel, Kobbi Nissim, and Eran Omri. Distributed private data analysis: Simultaneously solving how and what. In *Advances in Cryptology–CRYPTO 2008*, pages 451–468. Springer, 2008. 4

[7] Lucien Birgé and Pascal Massart. *From Model Selection to Adaptive Estimation*, pages 55–87. Springer New York, 1997. 6

[8] Jeremiah Blocki. Personal communication. 2013. 2

[9] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. The Johnson-Lindenstrauss Transform Itself Preserves Differential Privacy. In *FOCS*, pages 410–419. IEEE Computer Society, 2012. 1, 2, 4, 6, 7, 9, 13, 14, 18

[10] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. Differentially private data analysis of social networks via restricted sensitivity. In Robert D. Kleinberg, editor, *ITCS*, pages 87–96. ACM, 2013. 4

[11] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: the sulq framework. In Chen Li, editor, *PODS*, pages 128–138. ACM, 2005. 4

[12] Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to noninteractive database privacy. *J. ACM*, 60(2):12, 2013. 4

[13] Emmanuel J. Candès and Terence Tao. Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies? *IEEE Transactions on Information Theory*, 52(12):5406–5425, 2006. 2

[14] Lynn E Cannon. A cellular computer to implement the kalman filter algorithm. Technical report, DTIC Document, 1969. 4

[15] Kamalika Chaudhuri, Anand D Sarwate, and Kaushik Sinha. Near-optimal differentially private principal components. In *NIPS*, pages 998–1006, 2012. 2

[16] Kenneth L. Clarkson and David P. Woodruff. Numerical linear algebra in the streaming model. In Mitzenmacher [42], pages 205–214. 2

[17] Richard Cleve, Peter Høyer, Benjamin Toner, and John Watrous. Consequences and Limits of Nonlocal Strategies. In *IEEE Conference on Computational Complexity*, pages 236–249. IEEE Computer Society, 2004. 2

[18] Anirban Dasgupta, Ravi Kumar, and Tamás Sarlós. A sparse Johnson: Lindenstrauss transform. In *STOC*, pages 341–350, 2010. 14

[19] Sanjoy Dasgupta and Anupam Gupta. An elementary proof of a theorem of johnson and lindenstrauss. *Random Structures & Algorithms*, 22(1):60–65, 2003. 3

[20] Anindya De. Lower bounds in differential privacy. In *Theory of Cryptography*, pages 321–338. Springer, 2012. 14

[21] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *EUROCRYPT*, pages 486–503, 2006. 4

[22] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating Noise to Sensitivity in Private Data Analysis. In *TCC*, pages 265–284, 2006. 4

[23] Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil P. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In Mitzenmacher [42], pages 381–390. 2

[24] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and Differential Privacy. In *FOCS*, pages 51–60, 2010. 4, 8

[25] Oded Goldreich. A sample of samplers: A computational perspective on sampling. *def*, 1:2n, 1997. 14

[26] Anupam Gupta, Moritz Hardt, Aaron Roth, and Jonathan Ullman. Privately Releasing Conjunctions and the Statistical Query Barrier. *SIAM J. Comput.*, 42(4):1494–1520, 2013. 4

[27] Anupam Gupta, Aaron Roth, and Jonathan Ullman. Iterative constructions and private data release. In *TCC*, pages 339–356, 2012. 4

[28] Moritz Hardt, Katrina Ligett, and Frank McSherry. A Simple and Practical Algorithm for Differentially Private Data Release. In *NIPS*, pages 2348–2356, 2012. 4

[29] Moritz Hardt and Aaron Roth. Beating randomized response on incoherent matrices. In *STOC*, pages 1255–1268, 2012. 14

[30] Moritz Hardt and Aaron Roth. Beyond worst-case analysis in private singular vector computation. In *STOC*, pages 331–340, 2013. 14

[31] Moritz Hardt and Guy N. Rothblum. A Multiplicative Weights Mechanism for Privacy-Preserving Data Analysis. In *FOCS*, pages 61–70, 2010. 4

[32] Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *STOC*, pages 705–714, 2010. 14

[33] Antoine Joux. Multicollisions in iterated hash functions. application to cascaded constructions. In *Advances in Cryptology–CRYPTO 2004*, pages 306–316. Springer, 2004. 2

[34] Michael Kapralov and Kunal Talwar. On differentially private low rank approximation. In *SODA*, volume 5, page 1. SIAM, 2013. 2, 14

[35] John Kelsey and Bruce Schneier. Second preimages on n-bit hash functions for much less than 2 n work. In *Advances in Cryptology–EUROCRYPT 2005*, pages 474–490. Springer, 2005. 2

[36] Krishnaram Kenthapadi, Aleksandra Korolova, Ilya Mironov, and Nina Mishra. Privacy via the johnson-lindenstrauss transform. *arXiv preprint arXiv:1204.2606*, 2012. 4

[37] Erricos John Kontoghiorghes. *Handbook of parallel computing and statistics*. CRC Press, 2010. 4

[38] Felix Krahmer and Rachel Ward. New and Improved Johnson-Lindenstrauss Embeddings via the Restricted Isometry Property. *SIAM J. Math. Analysis*, 43(3):1269–1281, 2011. 2, 3, 7

[39] Pascal Massart and Jean Picard. *Concentration inequalities and model selection*, volume 1896. Springer, 2007. 6

[40] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 94–103. IEEE, 2007. 4

[41] Ilya Mironov. On significance of the least significant bits for differential privacy. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 650–661. ACM, 2012. 2

[42] Michael Mitzenmacher, editor. *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*. ACM, 2009. 15

[43] Gilles Pisier. Some applications of the metric entropy condition to harmonic analysis. In *Banach Spaces, Harmonic Analysis, and Probability Theory*, pages 123–154. Springer, 1983. 5

[44] Holger Rauhut, Justin K. Romberg, and Joel A. Tropp. Restricted Isometries for Partial Random Circulant Matrices. *CoRR*, abs/1010.1847, 2010. 3, 7

[45] Aaron Roth and Tim Roughgarden. Interactive privacy via the median mechanism. In *STOC*, pages 765–774, 2010. 4

[46] Mark Rudelson and Roman Vershynin. On sparse reconstruction from Fourier and Gaussian measurements. *Communications on Pure and Applied Mathematics*, 61(8):1025–1045, 2008. 3, 7

[47] Tamas Sarlos. Improved approximation algorithms for large matrices via random projections. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 143–152. IEEE, 2006. 2

[48] J. Upadhyay. Differentially Private Linear Algebra in the Streaming Model. *ArXiv e-prints*, September 2014. 4, 8

[49] Jalaj Upadhyay. Random Projections, Graph Sparsification, and Differential Privacy. In *ASIACRYPT (1)*, pages 276–295, 2013. 2, 4, 13, 14

[50] Jalaj Upadhyay. Random Projection, Restricted Isometry Property, and More. *Submitted to ITCS*, 2015. 2, 4, 11, 12, 13

[51] Jan Vybíral. A variant of the Johnson–Lindenstrauss lemma for circulant matrices. *Journal of Functional Analysis*, 260(4):1096–1105, 2 2011. 1, 2, 3, 6, 7, 12, 14

[52] Xiaoyun Wang, Yiqun Lisa Yin, and Hongbo Yu. Finding collisions in the full sha-1. In *Advances in Cryptology–CRYPTO 2005*, pages 17–36. Springer, 2005. 2

[53] Xiaoyun Wang and Hongbo Yu. How to break md5 and other hash functions. In *Advances in Cryptology–EUROCRYPT 2005*, pages 19–35. Springer, 2005. 2

[54] P. Wolff. On randomness reduction in the Johnson-Lindenstrauss lemma. *ArXiv e-prints*, February 2012. 11

A Deferred Proofs

A.1 Proof of Proposition 13

We start by proving a one-dimensional analogue of the proposition.

Proposition 17. Let $Y \sim \mathcal{N}(0,1)$ and $S := \log(\mathbb{E}[\exp(a\lambda(Y^2-1))])$. Then $S \leq \frac{a^2\lambda^2}{1-2a\lambda}$.

Proof. $S := \log(\mathbb{E}_Y[\exp(\lambda a(Y^2-1))])$, where $Y \sim \mathcal{N}(0,1)$. A simple calculation shows that when $Y \sim \mathcal{N}(0,1)$, then

$$S = a^2\lambda^2 \sum_{i\geq 0} \frac{(2\lambda a)^i}{i+1} \leq a^2\lambda^2 \sum_{i\geq 0} (2a\lambda)^i = \frac{a^2\lambda^2}{1-2a\lambda}.$$

□

The proof of Proposition 13 now follows from linearity of expectation. Let $Y_1, \cdots, Y_r$ be random variables picked using the distribution $\mathcal{N}(0,1)$. From the linearity of expectation, a simple extension of Proposition 17 to a vector of Gaussian variables results in Proposition 13. More concretely, from the linearity of expectation, we have

$$\begin{aligned}\sum_{j=1}^{r} \log\left(\mathbb{E}_{Y_i}\left[\exp(\lambda\sigma_j^2(Y_j^2-1))\right]\right) &= \sum_{j=1}^{r} \lambda^2\sigma_j^4 \sum_{i\geq 0} \frac{(2\lambda\sigma_j^2)^i}{i+1} \\ &\leq \lambda^2 \sum_{j=1}^{r} \sigma_j^4 \sum_{i\geq 0} (2\lambda\sigma_j^2)^i \leq \frac{\lambda^2 \sum_{i=1}^{r} \sigma_i^4}{1-2\lambda \max_i |\sigma_i|^2}.\end{aligned}$$

A.2 Covariance Matrices are Positive Semi-Definite

Lemma 18. Suppose that Σ is the covariance matrix corresponding to some random vector $|x\rangle$. Then Σ is symmetric positive semi-definite.

Proof. For any vector $|x\rangle \in \mathbb{R}^n$, we have

$$\langle x|\Sigma|x\rangle = \sum\sum (\Sigma_{ij}x_ix_j) = \sum\sum (\mathsf{COV}[x_i,x_j])\, x_ix_j = \mathbb{E}\left[\sum\sum (x_i - \mathbb{E}[x_i])(x_j - \mathbb{E}[x_j])x_ix_j\right].$$

Now the quantity under the summation is of form $\sum\sum x_ix_jz_iz_j = (|x\rangle\langle z|) \geq 0$. Therefore, the quantity inside the expectation is always non-negative; therefore, the expectation is non-negative. This proves the proposition. Now, for the definition of PDF for the above multivariate distribution, Σ^{-1} should exists; therefore, $\Sigma \in S^n_{++}$. □

B Details of the Application of Our Projection Matrices

B.1 Cut Queries on a Graph

We first give the utility proof. For any any set of vertex of size n, assuming equation (10) holds, we have $\chi_S^\mathsf{T}\tilde{L}_G\chi_S \leq (1+\eta)\chi_S^\mathsf{T}L_\mathcal{G}\chi_S$. In other word, the approximation can be bounded from the above by

$$\begin{aligned}&\frac{1}{1-\frac{w}{n}}\left((1+\eta)\chi_S^\mathsf{T}L_\mathcal{G}\chi_S - \frac{ws(n-s)}{n}\right) \\ &\quad= \frac{1}{1-\frac{w}{n}}\left((1+\eta)\frac{w}{n}s(n-s) + (1+\eta)\left(1-\frac{w}{n}\right)\chi_S^\mathsf{T}L_\mathcal{G}\chi_S - \frac{ws(n-s)}{n}\right) \\ &\quad\leq (1+\eta)\mathsf{CUT}(S) + 2\eta ws,\end{aligned}$$

INPUT. A n-node graph G, parameters $\varepsilon, \delta, \eta, \nu$.
OUTPUT. A Laplacian of a graph $\tilde{L}_G$.

1. Set $w = 16\sqrt{\frac{n}{\varepsilon}}\ln(4n/\delta)$ and $L'_G = \frac{w}{n}K_n + \left(1 - \frac{w}{n}\right)L_G$, where K_n is the complete graph on n-vertices.
2. Construct a matrix Φ as in Figure 1 and publish $\frac{1}{r}\Phi L'_G\Phi^{\mathsf{T}}$.

ANSWERING A CUT QUERY.

INPUT. A non-empty $S \subseteq V(G)$, parameters n, w, and $\tilde{L}_G$.
OUTPUT. Return $\frac{1}{1-\frac{w}{n}}\left(\chi_S^{\mathsf{T}}\tilde{L}_G\chi_S - \frac{ws(n-s)}{n}\right)$.

Figure 2: Answering Cut Queries of a Graph while Preserving Differential Privacy

where $\mathsf{CUT}(S)$ is the correct answer. Now, just as in Blocki *et al.* [9], if we wish to answer correctly a set of all possible queries, we need to set $\nu' = \nu/2^n$, and deduce that the amount of noise added to each query is $\tilde{O}(s\sqrt{n}/\varepsilon)$. For the privacy proof, note that Step 1 assures that the singular values of L_G which is now the private graph are at least $w = 16\sqrt{\frac{n}{\varepsilon}}\ln(4n/\delta) = \sigma_{\mathsf{min}}$. This is because it has a complete graph as a subgraph, and complete graph has second largest eigenvalue n. Note that in this case, we are only consider the space orthogonal to the kernel space of the Laplacian of a graph, i.e., as stated in Lemma 19 below. In other words, we have X with support over $\mathbf{1}^{\perp}$.

Lemma 19. The kernel space of a connected graph is $\mathsf{Span}\{\mathbf{1}\}$, the span of all one vector.

B.2 Covariance Queries on a Matrix

INPUT. A $n \times d$-matrix A, parameters $\varepsilon, \delta, \eta, \nu$.
OUTPUT. A sanitized matrix $\tilde{A}$.

1. Set $w = 16\sqrt{\frac{n}{\varepsilon}}\ln(4n/\delta)$.
2. Subtract the mean from A by setting $A_1 = A - \frac{1}{n}|\mathbf{1}\rangle\langle\mathbf{1}|A$.
3. For $A_1 = U\Sigma V^{\mathsf{T}}$ computed above, set $A' = U(\sqrt{\Sigma^2 + w^2\mathbb{I}})V^{\mathsf{T}}$.
4. Construct a matrix Φ as in Figure 1 and return $\tilde{A} = \frac{1}{r}A'^{\mathsf{T}}\Phi^{\mathsf{T}}\Phi A$.

ANSWERING A COVARIANCE QUERY.

INPUT. A unit vector $|x\rangle$, parameters n, w, and $\tilde{A}$.
OUTPUT. Return $\langle x|A|x\rangle - w^2$.

Figure 3: Answering Covariance Queries on a Matrix while Preserving Differential Privacy

The utility proof is straightforward just as in Appendix B.1. For the privacy proof, note that Step 3 assures that the singular values of A' which is now the private matrix are at least $w = 16\sqrt{\frac{n}{\varepsilon}}\ln(4n/\delta) = \sigma_{\mathsf{min}}$.