Transformation-Based Outsourcing of Linear Equation Systems over Real Numbers

Peeter Laud[1] and Alisa Pankova[1,2,3]

[1] Cybernetica AS
[2] Software Technologies and Applications Competence Centre (STACC)
[3] University of Tartu
{peeter.laud|alisa.pankova}@cyber.ee

Abstract. This paper studies the possibility of achieving indistinguishability-based security in privately outsourcing linear equation systems over real numbers. The particular task is to solve a full-rank $n \times n$ system $A\mathbf{x} = \mathbf{b}$. Since the most complex part of this task is inverting A, the problem can be reduced to outsourcing of a square matrix inverse computation. Although outsourcing matrix inverse is trivial for matrices over finite fields, it is not so easy for matrices over real numbers. We study the class of affine transformations for matrices over real numbers, find out which forms are possible at all, and state some properties that the transformation and the initial matrices must satisfy in order to make the initial matrices perfectly (or statistically) indistinguishable after applying the transformation. This paper provides both possibility and impossibility results.

Introduction

Secure outsourcing is a simple case of privacy-preserving two-party computation. The first party (the *client*) has a task T that it wants to solve, but does not have enough computational resources for that. At the same time, the second party (the *server*) is powerful enough to solve it. The client may want to keep T private from the server. The parties execute a protocol that results in the client learning the solution to T, and does not leak any information about T to the server. Executing the protocol should be computationally much easier for the client than solving T itself. Such protocols have been proposed for various tasks such as cryptographic operations [1, 2], database operations [3], sequence matching [4], and some linear algebra tasks [5–7]. Often, secure outsourcing appears as part of secure multiparty computation (SMC) protocols, in order to speed up the solution of a particular subtask in privacy-preserving manner. In this case, the operations of the client are executed by SMC techniques, while the server's operations are run in the public view.

Solving full-rank $n \times n$ linear equation systems $A\mathbf{x} = \mathbf{b}$ over $\mathbb{R}$ is needed in various applications. For example, some linear programming algorithms need to solve such a system on each iteration. In this case, a single iteration is not allowed to be too expensive (or in turn iterative), and solving each single system has to be quite efficient. A straightforward approach is to compute the solution $\mathbf{x}_0 = A^{-1}\mathbf{b}$ by finding A^{-1} first, an then multiplying it with $\mathbf{b}$. The problem can be reduced to outsourcing square matrix inverse computation, since the most complex part is computing A^{-1}, and the consequent

multiplication by **b** can be done relatively easily even by a weak client. Hence this work treats primarily outsourcing matrix inverses over $\mathbb{R}$.

Outsourcing rank-deficient $m \times n$ linear equation systems could also be an interesting task, but the transformation can be at most as secure as the transformation for full-rank systems, since a full-rank system can be seen as an instance of a rank-deficient system whose feasible region is just a single point. Taking into account our quite pessimistic results about the full-rank systems, we do not dare to hope that we can achieve anything better for rank-deficient systems.

Outsourcing matrix inverse over a finite field $\mathbb{F}$ can be done as follows. Matrix multiplication is much easier than matrix inverse. Given an $n \times n$ invertible matrix A over $\mathbb{F}$, the client generates a random invertible matrix $R \xleftarrow{\$} \mathbb{F}^{n\times n}$. Sampling each entry uniformly from $\mathbb{F}$ gives a uniformly distributed matrix over $\mathbb{F}^{n\times n}$ which is invertible with high probability. Since invertible $n \times n$ matrices over $\mathbb{F}$ form a multiplicative group $GL(n,\mathbb{F})$, the product RA is distributed uniformly in $GL(n,\mathbb{F})$ for any $A \in GL(n,\mathbb{F})$, so RA leaks no information about A. Alternatively, a uniformly random element of $GL(n,\mathbb{F})$ can be sampled by first sampling uniformly a random lower triangular matrix L and a random upper triangular matrix U (the details of sampling L and U can be found in [8]), taking $R = LU$. If $|\mathbb{F}|$ is significantly larger than n, then only a negligible fraction of invertible matrices cannot be decomposed to LU. Any invertible matrix can be decomposed as LUP for lower triangular L, upper triangular U and a permutation matrix P, but if L, U, and P are sampled from a uniform distribution, then LUP is not distributed uniformly [8].

The client sends RA to the server. The server computes $(RA)^{-1} = A^{-1}R^{-1}$, sends it back, and the client computes $A^{-1}R^{-1} \cdot R = A^{-1}$. However, the same security argument does not pass straightforwardly for matrices over $\mathbb{R}$ since we cannot define a uniform distribution on $GL(n,\mathbb{R})$.

We show that, for a *perfectly* secure outsourcing, it is in general necessary to sample uniformly matrices from certain subgroups of $GL(n,\mathbb{R})$ or their cosets, making the situation similar to $GL(n,\mathbb{F})$. However, such groups have specific structure that provides another efficient way of inverting their elements, eliminating the need for outsourcing. It is still possible to find transformations that provide *statistical* security, but doing it straightforwardly requires the entries of matrices to grow exponentially in the security parameter η. On the other hand, for some sets of matrices to be hidden, we manage to give statistically secure hiding methods that may be less expensive than computing the matrix inverse. There are also some cases that are not fully studied yet, and *computational* security may still be possible, although cryptographic assumptions over $\mathbb{R}$ have not received much attention so far. Hence this work is not a strict impossibility result, but rather a warning that obtaining a sufficiently hiding transformation over $\mathbb{R}$ is not as easy as over $\mathbb{F}$.

Notation Throughout this paper, matrices are denoted by upper case letters (A), and vectors by lower case bold letters (**b**). We use calligraphic letters for sets ($\mathcal{S}$). A distribution over a set $\mathcal{S}$ is denoted $\mathcal{D}_{\mathcal{S}}$, and a group defined on $\mathcal{S}$ is denoted $\mathcal{G}_{\mathcal{S}}$. Operations on values are routinely lifted to operate (point-wise) on probability distributions over these values; in this case, the result is again a probability distribution.

We define *statistical distance* as

$$SD(\mathscr{D}_1, \mathscr{D}_2) = \sup_{X \subseteq \mathbb{R}} |Pr[y \in X \mid y \leftarrow \mathscr{D}_1] - Pr[y \in X \mid y \leftarrow \mathscr{D}_2]| \ .$$

We write $X \leftarrow \mathscr{D}_X$ to state that X is distributed according to $\mathscr{D}_X$, and $X \xleftarrow{\$} \mathscr{X}$ to state that X is sampled uniformly from a set $\mathscr{X}$. We write $\mathscr{D}_1 \approx \mathscr{D}_2$ to denote that $SD(\mathscr{D}_1, \mathscr{D}_2) < \varepsilon$ for a negligible ε (a function $\alpha : \mathbb{N} \to \mathbb{R}$ is *negligible* if $\forall c \exists m \forall \eta \geq m : |\alpha(\eta)| < 1/n^c$). If the distributions are exactly the same, we write $\mathscr{D}_1 = \mathscr{D}_2$. We use multiset comprehensions $\{|f(X) \mid X \leftarrow \mathscr{D}|\}$ to construct new probability distributions from the existing ones. Isomorphisms of groups and their cosets are denoted by $\cong$.

1 Related Work

Hiding affine transformations have been proposed for various linear algebra tasks [9], including solving linear equation systems. Even though the security of these transformations was proved for finite fields in [9], the follow-up papers used them freely for the field of real numbers. The ideas of [9] have been used in practical applications such as statistical analysis [10, 11], privacy-preserving polynomial interpolation [12], or just constructing particular protocols for linear equation solving [13] and more general matrix inverse outsourcing [14].

Affine transformation-based approach has also been used for outsourcing linear programming tasks [15, 6, 16–18, 7, 19–22]. Since linear programming is defined over $\mathbb{R}$, there had been no good security definitions for these transformations. Some efficient attacks based on geometrical properties of the transformed feasible region have been found for example in [20, 23, 6]. The attacks are linear programming-specific and do not extend to outsourcing a full-rank linear equation system. However, they show that transformations that are secure in $\mathbb{F}$ are in general not so easily extensible to $\mathbb{R}$.

A secure method for solving linear programming tasks based on interior point methods has been proposed in [24]. On each iteration, the algorithm uses as a black-box a method for secure inverting a matrix over $\mathbb{R}$. They suggest to use existing transformation-based matrix inverse methods, including [10, 11]. Unfortunately, the security definition used in these approaches are non-standard, and hence cannot be included into more complex composition proofs. More precisely, the security definition used in the previous works (that was first used in [25]) requires that, observing the transformation output, the adversary can reduce the set of possible inputs only to an infinite or at least computationally unfeasible set, which is not enough for example for indistinguishability-based security. Provably secure cryptographic methods for solving linear equations systems (based on Gaussian elimination and LU decomposition) have been proposed in [26].

In this work, we do not discuss the security of particular protocols that compute the transformations. We are just interested in the information that the transformed quantities may leak themselves.

2 Outsourcing by Affine Transformations

We start by defining the class of transformations for systems of linear equations that we consider in this paper. Starting from a $n \times n$ system $A\mathbf{x} = \mathbf{b}$, we transform it to an

$m \times m$ system $B\mathbf{y} = \mathbf{d}$. Let $\mathbf{x}_0$ and $\mathbf{y}_0$ be the unique solutions to $A\mathbf{x} = \mathbf{b}$ and $B\mathbf{y} = d$, respectively. The solution $\mathbf{y}_0$ has to be efficiently transformable back to $\mathbf{x}_0$.

In this paper we consider only affine transformations, for the following reasons:

- they have been the only ones considered in previous work;
- they are the most natural approach for hiding a system of linear equations;
- they are sufficiently efficient;
- in finite fields, a secure transformation is given by a highly special case of them, as described in the introduction.

2.1 Structure of Affine Transformations

Using affine transformations, the solution to the initial problem is computed from the solution to the transformed problem as $\mathbf{x}_0 = F\mathbf{y}_0 + \mathbf{r}_0$ for some $F \in \mathbb{R}^{n \times m}$ and $\mathbf{r}_0 \in \mathbb{R}^n$ that are generated by the client together with B and $\mathbf{d}$. We may assume that F is full-rank, otherwise there will have to be certain constant relationships between the variables of $\mathbf{x}_0$, and we cannot in general make any assumptions about the solution.

The entire transformed problem is $(B, \mathbf{d}, F, \mathbf{r}_0)$. In our settings, we only require that B is published, since the most complex part of solving the system of equations is finding the inverse of B. Hence the affine transformation $F\mathbf{y}_0 + \mathbf{r}_0$ does not have to be hiding, and $(\mathbf{d}, F, \mathbf{r}_0, \mathbf{y}_0)$ may leak information about $(A, \mathbf{b})$. The main questions are under which conditions $\mathbf{d}$, F, $\mathbf{r}_0$, B can be generated efficiently, and whether satisfying these conditions allows to keep B secure enough.

We consider only affine transformations for the purpose of generating B, i.e. $B = PAQ + R$, where $P \in \mathbb{R}^{m \times n}$, $Q \in \mathbb{R}^{n \times m}$, and $R \in \mathbb{R}^{m \times m}$ are random matrices over real numbers sampled independently from A according to some distribution. Here $m \geq n$, and P, Q are both full-rank, since otherwise $\mathbf{y}_0$ would not contain enough information to reconstruct $\mathbf{x}_0$.

2.2 Unavoidable Restrictions on $B = PAQ + R$

We show that the form $B = PAQ + R$ is not more general than apparently more restricted methods of computing B from A, if the transformation has to work for all $n \times n$ systems of linear equations. First we show that R cannot be full-rank if we do not make any assumptions on $\mathbf{x}_0$.

Since $\mathbf{x}_0 = F\mathbf{y}_0 + \mathbf{r}_0$, without loss of generality we may write $\mathbf{y}_0 = \begin{pmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{pmatrix}$, and $F = (F_1 \mid F_2)$, where F_1 is invertible. If F_1 is not invertible, we may always permute the columns by defining a new matrix $F' := FT$ for a $m \times m$ permutation matrix T, and such F_1 exists since F is full-rank. Then $\mathbf{y}_1 = F_1^{-1}(\mathbf{x}_0 - \mathbf{r}_0) - F_1^{-1}F_2\mathbf{y}_2$.

We may now express the equation $(PAQ+R)\mathbf{y}_0 = \mathbf{d}$ through $\mathbf{y}_2$ only. Let $B = PAQ + R$ be divided into four blocks, indexed B_{11}, B_{12}, B_{21} and B_{22}, where the left blocks (B_{11} and B_{21}) correspond to the variables $\mathbf{y}_1$, and the right blocks (B_{12} and B_{22}) to the variables $\mathbf{y}_2$. Let $B_{ij} = (PAQ+R)_{ij} = ((PAQ)_{ij} + R_{ij})$, where $(PAQ)_{ij}$ and R_{ij} are analogous blocks. We get the following.

$$\begin{pmatrix} B_{11}F_1^{-1}F_2 + B_{12} \\ B_{21}F_1^{-1}F_2 + B_{22} \end{pmatrix} \mathbf{y}_2 = \mathbf{d} - \begin{pmatrix} B_{11} \\ B_{21} \end{pmatrix} F_1^{-1}(\mathbf{x}_0 - \mathbf{r}_0) \tag{1}$$

Since the number of rows is $m > m-n$, the upper n rows of the matrix on left hand side are a certain linear combination of the $m-n$ lower rows (again, without loss of generality we may permute the rows of B if the last $m-n$ rows are not full-rank by taking $P' := TP$ and $R' := TR$ for an $m \times m$ permutation matrix T). Let this linear combination be represented by a matrix X. Formally, X is an $(n \times (m-n))$ matrix such that

$$XB_{21}F_1^{-1}F_2 + XB_{22} = B_{11}F_1^{-1}F_2 + B_{12} \quad . \tag{2}$$

For the Equation 1 to be satisfied, the upper n entries of the right hand side must be the same linear combination of the lower $m-n$ entries. If we denote $\mathbf{d} = \begin{pmatrix} \mathbf{d}_1 \\ \mathbf{d}_2 \end{pmatrix}$, where $\mathbf{d}_1$ are the n upper entries of $\mathbf{d}$, and $\mathbf{d}_2$ are the $m-n$ lower entries, we get

$$X\mathbf{d_2} - XB_{21}F_1^{-1}(\mathbf{x}_0 - \mathbf{r}_0) = \mathbf{d_1} - B_{11}F_1^{-1}(\mathbf{x}_0 - \mathbf{r}_0) \quad . \tag{3}$$

By assumption, $\mathbf{d}$, $\mathbf{r}_0$, F do not depend on $\mathbf{x}_0$. Also, the value $\mathbf{x}_0$ is completely independent from A alone, if $\mathbf{b}$ is not taken into account, since for a fixed A, for any $\mathbf{x}_0$ there exists $\mathbf{b}$ such that $A\mathbf{x}_0 = \mathbf{b}$. Hence the Equation 3 must hold for any $\mathbf{x}_0$, and we may treat the entries of $\mathbf{x}_0$ as formal variables. We get the following polynomial vector equality.

$$X\mathbf{d}_2 - \mathbf{d}_1 + (B_{11} - XB_{21})F_1^{-1}\mathbf{x}_0 - (B_{11} - XB_{21})F_1^{-1}\mathbf{r}_0 = \mathbf{0}$$

For the left hand side to be a zero polynomial with respect to $\mathbf{x}_0$, the equality $X\mathbf{d}_2 - \mathbf{d}_1 - (B_{11} - XB_{21})F_1^{-1}\mathbf{r}_0 = \mathbf{0}$ must hold (since one possible solution is $\mathbf{x}_0 = \mathbf{0}$). Also, $(B_{11} - XB_{21})F_1^{-1} = \mathbf{0}$ must hold (to satisfy any $\mathbf{x}_0$ that does not nullify it). Since F_1^{-1} is invertible, this reduces to $B_{11} - XB_{21} = \mathbf{0}$, or equivalently $(PAQ)_{11} - X(PAQ)_{21} + R_{11} - XR_{21} = \mathbf{0}$. We have got that $R_{11} = X(PAQ)_{21} - (PAQ)_{11} + XR_{21}$.

We are also interested in relations between R_{22} and R_{12}. Starting from Equation 2, we get

$$\begin{aligned} XB_{21}F_1^{-1}F_2 + XB_{22} &= B_{11}F_1^{-1}F_2 + B_{12} \\ (XB_{21} - B_{11})F_1^{-1}F_2 &= B_{12} - XB_{22} \\ \mathbf{0}F_1^{-1}F_2 &= B_{12} - XB_{22} \\ XB_{22} &= B_{12} \quad . \end{aligned}$$

We can write this as $R_{12} = X(PAQ)_{22} - (PAQ)_{12} + XR_{22}$.
Now $(R_{11} \mid R_{12}) = X((PAQ)_{21} \mid (PAQ)_{22}) - ((PAQ)_{11} \mid (PAQ)_{12}) + X(R_{21} \mid R_{22})$.
For any $P = \begin{pmatrix} P_1 \\ P_2 \end{pmatrix}$ and $Q = (Q_1 \mid Q_2)$, the quantity PAQ looks as follows.

$$\begin{pmatrix} P_1 \\ P_2 \end{pmatrix} A \, (Q_1 \mid Q_2) = \begin{pmatrix} P_1AQ_1 & P_1AQ_2 \\ P_2AQ_1 & P_2AQ_2 \end{pmatrix} \quad .$$

Hence we have $(PAQ)_{ij} = P_iAQ_j$. We may always take P' such that $P'_2 := P_2$ and $P'_1 := XP_2$. Such a choice of P' does not depend on A: if we have achieved hiding for $B = PAQ + R$, then X is independent from A since it is computable from $B_{1j} = XB_{2j}$,

and so would leak information about A. We get $(PAQ+R) = (P'AQ+R')$ where R' is such that $(R'_{11} \mid R'_{12}) = X(R'_{21} \mid R'_{22})$ and hence R is of rank at most $(m-n)$. It is actually exactly of rank $m-n$, since otherwise $(PAQ+R)$ would not be full-rank. We get that hiding with an arbitrary R is not better than hiding with an R of rank $(m-n)$.

If R is of rank $(m-n)$, then it is of the form $\begin{pmatrix} SR'T & SR' \\ R'T & R' \end{pmatrix}$, where R' is an $(m-n)\times(m-n)$ invertible submatrix, and S and T are arbitrary (here without loss of generality we assume that the $(m-n)\times(m-n)$ invertible submatrix is located in the right lower part of R, since we again may take $R' = T_1 R T_2$ for permutation matrices T_1 and T_2). Now $B = (PAQ+R)$ can be rewritten as follows.

$$B \;=\; \begin{pmatrix} P_1AQ_1 & P_1AQ_2 \\ P_2AQ_1 & P_2AQ_2 \end{pmatrix} \;+\; \begin{pmatrix} SR'T & SR' \\ R'T & R' \end{pmatrix} \;=\; \begin{pmatrix} P_1 & S \\ P_2 & I \end{pmatrix} \cdot \begin{pmatrix} A & 0 \\ 0 & I \end{pmatrix} \cdot \begin{pmatrix} Q_1 & Q_2 \\ R'T & R' \end{pmatrix}$$

In order for B to be invertible, both $P' = \begin{pmatrix} P_1 & S \\ P_2 & I \end{pmatrix}$ and $Q' = \begin{pmatrix} Q_1 & Q_2 \\ R'T & R' \end{pmatrix}$ have to be invertible. Let $A' = \begin{pmatrix} A & 0 \\ 0 & I \end{pmatrix}$. The most general form of an affine transformation is thus $B = P'A'Q'$, where the distributions of P' and Q' do not depend on A', and A' uniquely defines A (for the fixed n,m). Hence hiding A' is equivalent to hiding A (and increasing the size of I may be used to increase the security parameter if it turns out to depend on the size of A'). In the next sections, we study the hiding properties of $B = PAQ$.

3 Security Definition

Before considering how to pick the matrices P and Q, such that B contains no information about A, let us formally define what a secure transformation is. We use the formal definition of a transformation, side information, and the security definition proposed in [27], which is based on the formalizations of transformations for garbled circuits of [28]. Let $\mathfrak{T} \subseteq \{0,1\}^*$ denote the set of all possible tasks and $\mathfrak{S} \subseteq \{0,1\}^*$ the set of all possible solutions. For $T \in \mathfrak{T}$ and $S \in \mathfrak{S}$ let $T \vDash S$ denote that S is a solution for T. A *problem transformation* is defined as a pair of functions $\mathscr{F} : \mathfrak{T} \times \{0,1\}^* \to \mathfrak{T} \times \{0,1\}^*$ and $\mathscr{G} : \mathfrak{S} \times \{0,1\}^* \to \mathfrak{S}$.

A *correct problem transformation* is a pair $(\mathscr{F},\mathscr{G})$ that satisfies

$$\forall T,r,T',S',s \;:\; \big((T',s) \;=\; \mathscr{F}(T;r) \quad\wedge\quad T' \;\vDash\; S'\big) \quad\Rightarrow\quad T \;\vDash\; \mathscr{G}(S',s)\;,$$

where r is the randomness used by the transformation. For simplicity, let $\mathscr{F}(T)$ denote a randomized function that first samples r and then runs $\mathscr{F}(T;r)$. We assume that both $\mathscr{F}$ and $\mathscr{G}$ are polynomial-time with respect to the task description size.

Some information may be intentionally leaked by the transformation, due to being impractical to hide. This is captured by the notion of *side information function* $\Phi : \mathfrak{T} \to \{0,1\}^*$. When transforming a task T, we do not try to hide the information in $\Phi(T)$. For example, if T is not transformed at all, then $\Phi(T) = T$. The function Φ is assumed to be polynomial-time.

Definition 1. *[27] A transformation $(\mathscr{F},\mathscr{G})$ for $\mathfrak{T}$ is (computationally)* Φ-private *if the advantage of any probabilistic polynomial-time adversary $\mathscr{B}=(\mathscr{B}_1,\mathscr{B}_2)$ is a negligible function of η in the following experiment* $\mathbf{Exp}^{\mathfrak{T},\Phi}_{\mathscr{B}}$*:*

$$\left[\begin{array}{l}(T_0,T_1,s)\leftarrow\mathscr{B}_1(\eta)\\ \textbf{\textit{if}}\ |T_0|\neq\eta\vee|T_1|\neq\eta\vee\Phi(T_0)\neq\Phi(T_1)\\ \quad\textbf{\textit{return}}\ \bot\\ b\xleftarrow{\$}\{0,1\}\\ (T',\_)\leftarrow\mathscr{F}(T_b)\\ b'\leftarrow\mathscr{B}_2(T',s)\\ \textbf{\textit{return}}\ (b\stackrel{?}{=}b')\end{array}\right.$$

where the advantage of the adversary is $1/2$ *less the probability of the experiment returning* true.

In the following, we work with real numbers, hence it may be unclear what it means for a probabilistic polynomial-time algorithm to process these numbers. Nevertheless, we can define *perfect* and *statistical* Φ-privacy, by letting $\mathscr{B}_1$ and $\mathscr{B}_2$ be any functions and, in case of perfect privacy, require the advantage of $\mathscr{B}$ to be 0. For a fixed η, we say that the transformation provides at least σ bits of security, if the advantage of $\mathscr{B}$ is at most $2^{-\sigma}$. This level of security is achieved iff $SD(T'_0,T'_1)\leq 2^{-\sigma}$ for all $T_0,T_1\in\mathfrak{T}$ with $\Phi(T_0)=\Phi(T_1)$ and $|T_0|=|T_1|=\eta$, where T'_0 and T'_1 are the first components of $\mathscr{F}(T_0)$ and $\mathscr{F}(T_1)$ respectively. Typically, we want the security level to be at least η^c for some constant $c>0$.

In our case, the set of tasks $\mathfrak{T}$ corresponds to a set $\mathscr{A}$ of $n\times n$ invertible matrices. The set $\mathscr{A}$ depends on the application in which the matrix inversion problem appears; it is quite likely that $\mathscr{A}$ is not the whole $GL(n,\mathbb{R})$. It may be possible that for some applications involving the inversion of matrices, privacy-preserving outsourcing via transformation is possible, while other applications must use other means of inverting matrices. In the rest of this paper, we describe which conditions $\mathscr{A}$ must satisfy for the existence of transformations with perfect or statistical privacy. It is also possible that some applications produce matrices A, such that there is some simple transformation f, such that $f(A)\in\mathscr{A}$ and the inverse of $f(A)$ is still sufficient for the purposes of the application. E.g. the transformation $f(A)$ might normalize the rows or the columns of A.

3.1 On the Entries of Hiding Matrices

Before starting to investigate the sets of matrices $\mathscr{A}$, we characterize the distributions of matrices P and Q on the basis of their entries. We consider that there must be an efficient procedure that samples P and Q, hence the entries of P and Q cannot be too large. Also, we show that it is definitely not sufficient to independently sample the entries from the same distribution.

3.2 Reasonable Bounds on P and Q

There may be a distribution of matrices P and Q where the entries of P and Q are much larger than the entries of the matrices in $\mathscr{A}$, such that the multiplication $B=PAQ$ gives

statistical security. In practice, we can do it only if the *sizes* of the entries of P and Q are reasonably small. By *size* we mean the number of bits required to represent an entry, which is logarithmic in the actual entry.

For $a \in \mathbb{R}$, it is actually less clear, what its *size* is; and it is something that we do not want to define in detailed manner. It is clear, however, that we do not want the transformation mechanism to introduce too large numbers. We argue that $\log a$ is still a reasonable measure for the "information content", and thus the size of a. Indeed, we are using larger entries of P and Q to statistically hide the matrices in $\mathscr{A}$, and the entropy of these entries is proportional to our definition of size.

We say that the sizes of entries P and Q are *bounded*, if they are polynomial in the security parameter η. The bound may depend on the actual set of matrices $\mathscr{A}$.

If the entries of a matrix are bounded, then some matrix properties that grow with sizes of the entries, such as *determinant* or *norms*, are also bounded. This in particular implies that the ℓ_2 operator norms of P and Q (the *spectral norms*) are bounded (an ℓ_2 operator norm of an $n \times n$ matrix A is defined as $\|A\|_2 = \sup\{\|A\mathbf{x}\|_2 \mid \mathbf{x} \in \mathbb{R}^n, \|\mathbf{x}\|_2 = 1\}$, where $\|\mathbf{x}\|_2 = \sqrt{x_1^2 + \ldots + x_n^2}$ is vector ℓ_2-norm). This fact can be used to apply some known theorems from group representation theory. An upper bound for ℓ_2-norm, defined by the sizes of entries, is $\|A\|_2 \le \sqrt{\|A\|_1 \|A\|_\infty}$, and the corresponding lower bound is $\|A\|_2 \ge 1/\|A^{-1}\|_2 \ge 1/\sqrt{\|A^{-1}\|_1 \|A^{-1}\|_\infty}$, where $\|A\|_1 = \max_{1 \le k \le n} \sum_{j=1}^{n} a_{jk}$ and $\|A\|_\infty = \max_{1 \le j \le n} \sum_{k=1}^{n} a_{jk}$ [29]. If the entries of A are at most c, then a rough bound on the determinant is $n! \cdot c^n$ according to the definition of the determinant.

3.3 Sampling P and Q similarly to finite fields

For hiding some matrix $A \in GL(n, \mathbb{F})$ for finite field $\mathbb{F}$ through its multiplication by a random matrix P, it is reasonable to sample each entry of P uniformly from $\mathbb{F}$. Different entries of P are sampled independently of each other. In the following we study, what happens if similar sampling — independently and from the same distribution — is used to generate entries of P and Q to hide elements of $\mathscr{A} \subset GL(n, \mathbb{R})$. We note that matrices P and Q sampled in such manner will be invertible with high probability, hence this approach might look reasonable.

We consider even a bit more general case. Suppose that there are any s rows indexed $c_1, \ldots, c_s$ in P such that if these rows are summed up, all the entries of the resulting vector come from the same distribution $\mathscr{D}_p$ (sampling the entries of these s rows independently is one particular case). Suppose there are t analogical columns indexed $d_1, \ldots, d_t$ in Q, and the distribution of sums of their entries is $\mathscr{D}_q$. If $B = PAQ$, then any entry of B is $b_{ij} = \sum_{k,\ell=1,1}^{n,n} p_{ik} a_{k\ell} q_{\ell j}$. Summing up the entries over s special rows and t special columns of B gives an entry that comes from the distribution

$$\begin{aligned} \mathscr{D}_{\sum_{k,\ell=1,1}^{n,n} a_{k\ell}} &= \sum_{c,d=c_1,d_1}^{c_s,d_t} \sum_{k,\ell=1,1}^{n,n} p_{ck} a_{k\ell} q_{\ell d} \\ &= \sum_{k,\ell=1,1}^{n,n} \sum_{c,d=c_1,d_1}^{c_s,d_t} p_{ck} a_{k\ell} q_{\ell d} \end{aligned}$$

$$
\begin{aligned}
&= \sum_{k,\ell=1,1}^{n,n} \left[\sum_{c=c_1}^{c_s} p_{ck} \sum_{d=d_1}^{d_t} q_{\ell d} \right] a_{k\ell} \\
&\leftarrow \{| \sum_{k,\ell=1,1}^{n,n} pq a_{k\ell} \mid p \leftarrow \mathscr{D}_p, q \leftarrow \mathscr{D}_q | \} \\
&= \{| pq \sum_{k,\ell=1,1}^{n,n} a_{k\ell} \mid p \leftarrow \mathscr{D}_p, q \leftarrow \mathscr{D}_q | \}
\end{aligned}
$$

This means that any two matrices whose sums of the corresponding $a_{k\ell}$ entries are different are distinguishable. This problem generalizes to not just the sum, but to any linear combination. In particular, it means that it would be a bad idea to sample the entries of P (or Q) independently of each other, even in just a single row (for P) or column (for Q), although it looks like the most intuitive way of sampling a random invertible matrix. In practice, leaking sums of the entries may be especially bad in the case of sparse matrices. We conclude that a more clever way of matrix sampling is needed.

4 Tolerable Side Information

In this section we are looking for the side information that multiplicative hiding $B = PAQ$ leaks. This allows us to define a suitable side information function Φ from Sec. 3.

Determinant Determinant of square matrices is multiplicative, so we have $|PAQ| = |P| \cdot |A| \cdot |Q|$. We show that, regardless of the distribution $\mathscr{D}_{(P,Q)}$ of P and Q (from which sampling is feasible), the absolute value of the determinant is leaked.

Claim. If there exist $A_1, A_2 \in \mathscr{A}$ such that $|A_1| \neq \pm|A_2|$, but the distributions of PA_1Q and PA_2Q are the same (where $(P,Q) \leftarrow \mathscr{D}_{(P,Q)}$), then the determinants of P and/or Q are unbounded.

Proof. For a fixed A we get $|PAQ| = |P| \cdot |A| \cdot |Q| \leftarrow |A| \cdot \mathscr{D}_{|P|\cdot|Q|}$ where $\mathscr{D}_{|P|\cdot|Q|} = \{||P| \cdot |Q| \mid (P,Q) \leftarrow \mathscr{D}_{(P,Q)}|\}$. Hence if $|A_1| \neq \pm|A_2|$, the distributions $|A_1| \cdot \mathscr{D}_{|P|\cdot|Q|}$ and $|A_2| \cdot \mathscr{D}_{|P|\cdot|Q|}$ are different: one is a scaled version of another. We can rewrite it as $\frac{|A_1|}{|A_2|} \cdot \mathscr{D}_{|P|\cdot|Q|} = \mathscr{D}_{|P|\cdot|Q|}$, so $\mathscr{D}_{|P|\cdot|Q|}$ has to be a scaled version of itself. For perfect security, this is possible only if the mean and the variance of $\mathscr{D}_{|P|\cdot|Q|}$ are both ∞, which implies unboundedness of determinants of P and Q.

A distribution can still be at least statistically indistinguishable from a scaled version of itself, since the difference may be caused by a negligible fraction of outliers.

Claim. Let η be a security parameter. If there exist $A_1, A_2 \in \mathscr{A}$, $|A_1| \neq \pm|A_2|$ such that $SD(\mathscr{D}_{|PA_1Q|}, \mathscr{D}_{|PA_2Q|}) < 2^{-\eta}$, then the sizes of entries of P and Q are of order $2^{\eta/n}$.

Proof. For any two distributions $\mathscr{D}$ and $c\mathscr{D}$ over $\mathbb{R}$ (where $c \geq 1$ is a constant), if $SD(\mathscr{D}, c\mathscr{D}) < \varepsilon$, then $\forall x \in \mathbb{R}^+ : \ |P[-\infty, x) - P[-\infty, cx)| < \varepsilon$, and similarly $\forall x \in \mathbb{R}^- :$

$|P[-\infty, cx) - P[-\infty, x)| < \varepsilon$ where P is the probability measure of $\mathscr{D}$. This implies $P[x, cx) < \varepsilon$ for $x \in \mathbb{R}^+$ and $P[cx, x) < \varepsilon$ for $x \in \mathbb{R}^-$. Now let us partition $\mathbb{R}$ into segments $[c^k, c^{k+1})$ and $[-c^{k+1}, -c^k)$, $k \in \mathbb{Z}$. We have $\mathbb{R} = \bigcup_{k\in\mathbb{Z}} [c^k, c^{k+1}) \cup [-c^{k+1}, -c^k)$ and since the segments are non-intersecting, $P(\mathbb{R}) = P[-\infty, \infty) = \sum_{k\in\mathbb{Z}} (P[c^k, c^{k+1}) + P[-c^{k+1}, -c^k))$. Due to the inequalities $\forall x \in \mathbb{R}^+ : P[x, cx) < \varepsilon$ and $\forall x \in \mathbb{R}^- : P[cx, x) < \varepsilon$, we have $P[c^k, c^{k+1}) < \varepsilon$ and $P[-c^{k+1}, -c^k) < \varepsilon$ ($\forall k \in \mathbb{Z}$). Since $P[-\infty, \infty) = 1$, we need $1 < \sum_{k\in\mathbb{Z}} 2\varepsilon$. Let m be the maximal number such that there exists $d \in supp(\mathscr{D})$ s.t $|c^{m-1}| \leq |d|$, but $|c^m| \geq |d|$ (such an m exists since we assume that the determinants are bounded). Then $1 < m \cdot 2\varepsilon$ (since $P(supp(\mathscr{D})) = 1$). If we want this to hold for negligible ε, then m becomes exponentially large, and at the same time $c^{m-1} \in supp(\mathscr{D})$, so we need to sample c^{m-1}.

We get that if we want to achieve $SD(\mathscr{D}_{|PA_1Q|}, \mathscr{D}_{|PA_2Q|}) < 2^{-\eta}$ for a security parameter η, then the distributions of $|P|$ and $|Q|$ take values around $(|A_1|/|A_2|)^{2^\eta}$, and hence the sizes of entries of P and Q are logarithmic to $\sqrt[n]{(|A_1|/|A_2|)^{2^\eta}} = (|A_1|/|A_2|)^{2^\eta/n}$, which is of order $2^{\eta/n}$. In this case, increasing n by extending A with an identity I to $\begin{pmatrix} A & 0 \\ 0 & I \end{pmatrix}$ does not help since it still does not change the determinant, and we would not like to increase n exponentially anyway.

We conclude that that $\mathscr{D}_{(P,Q)} \approx k \cdot \mathscr{D}_{(P,Q)}$ is not acceptable for $k \neq \pm 1$. We hence add the determinant of the matrix to the side information function. Equivalently, we may demand that $\mathscr{A}$ contains only matrices with the same absolute value of the determinant.

Submultiplicative norms Another way to distinguish the distributions of PA_1Q and PA_2Q is computing any of their submultiplicative properties (a property p is called *submultiplicative* if $p(AB) \leq p(A)p(B)$). For example, any submultiplicative matrix norm satisfies the inequality $\|PAQ\| \leq \|P\|\|A\|\|Q\|$. On the other hand, since P and Q are invertible, due to the same inequality we have

$$\|P^{-1}\|^{-1}\|A\|\|Q^{-1}\|^{-1} = \|P^{-1}\|^{-1}\|P^{-1}PAQQ^{-1}\|\|Q^{-1}\|^{-1} \leq \|PAQ\| \ .$$

Hence $\|PAQ\|$ is bounded from below by $\|A\| \cdot m$, and from above by $\|A\| \cdot M$, where $m = \min_{P,Q}(\|P^{-1}\|^{-1}\|Q^{-1}\|^{-1})$ and $M = \max_{P,Q}(\|P\|\|Q\|)$ depend on P and Q only, but not on A.

Clearly, if $\|A_1\| \cdot M < \|A_2\| \cdot m$, then we get the same problems as with the determinant. Since a norm is non-negative, we would analogously require that the norms are all the same. But is it reasonable to assume that all the submultiplicative norms are the same for the entire $\mathscr{A}$? If we want A to be indistinguishable from I and A^{-1} (for example, hiding a group), then the only possibility for A and A^{-1} to have equal operator p-norms for any $p > 1$ is the case where A is a permutation matrix (the proof is straightforward, and we do not present it here). This would be a too strict requirement, and hence we do not put any constraints on the norms of $\mathscr{A}$ (except those that are already implied by equal absolute values of the determinants).

Hiding Identity Matrix Suppose that we have come up with a distribution $\mathscr{D}_{(P,Q)}$ that hides $\mathscr{A}$. Define a new distribution $\widetilde{\mathscr{D}}_{(P,Q)} = \{|(PA, Q) \mid (P, Q) \leftarrow \mathscr{D}_{(P,Q)}|\}$ for some

$A \in \mathscr{A}$. Now $\widetilde{\mathscr{D}}_{(P,Q)}$ hides the set $A^{-1}\mathscr{A}$ as well as $\mathscr{D}_{(P,Q)}$ hides $\mathscr{A}$. Note that $I \in A^{-1}\mathscr{A}$. Hence without loss of generality we are looking for (im)possibility results for $\mathscr{A}$ such that $I \in \mathscr{A}$.

Summary We have shown that the following additional assumptions on $\mathscr{A}$ are either required, or do not lessen the generality:

- $\forall A_1, A_2 \in \mathscr{A} : \ |A_1| = \pm |A_2|$;
- $I \in \mathscr{A}$ (and hence $\forall A \in \mathscr{A} : |A| = \pm 1$).

5 Perfect Secrecy

In this section we look for (im)possibilities of finding a perfectly hiding transformation for $\mathscr{A}$. Since our transformation results in sampling a random element from some set, and then applying it to $A \in \mathscr{A}$, we define more precisely what it means that a distribution (or a set) hides $\mathscr{A}$ perfectly.

Definition 2. *A distribution $\mathscr{D}_{\mathscr{S}}$ on a set $\mathscr{S}$* perfectly hides *$\mathscr{A}$ if there exists an operation $\oplus : \mathscr{S} \times \mathscr{A} \mapsto \{0,1\}^*$ such that $\mathscr{D}_{\mathscr{S} \oplus A_i} = \mathscr{D}_{\mathscr{S} \oplus A_j}$ for all $A_i, A_j \in \mathscr{A}$, where $\mathscr{D}_{\mathscr{S} \oplus A} = \{|S \oplus A \mid S \leftarrow \mathscr{D}_{\mathscr{S}}|\}$.*

Definition 3. *A set $\mathscr{S}$* perfectly hides *$\mathscr{A}$ if there exists a distribution $\mathscr{D}_{\mathscr{S}}$ on $\mathscr{S}$ that perfectly hides $\mathscr{A}$.*

5.1 Some preliminary results from group theory

Let $GL(n,\mathbb{R})$ denote the group of all invertible $(n \times n)$ matrices over $\mathbb{R}$. Let $O(n,\mathbb{R})$ denote the subgroup of $GL(n,\mathbb{R})$ that contains all the orthogonal matrices (U such that $UU^T = I$).

Theorem 1. *[30, Exercise 2.15]. Any finite subgroup of $GL(n,\mathbb{R})$ is conjugate to $O(n,\mathbb{R})$ (i.e for any finite subgroup $\mathscr{G}$ of $GL(n,\mathbb{R})$ there exists $C \in GL(n,\mathbb{R})$ and a subgroup $\mathscr{U}$ of $O(n,\mathbb{R})$ such that $\mathscr{G} = C\mathscr{U}C^{-1}$.*

This is not the most general version of this theorem. A notion of *amenability* defines the groups on which it is possible to define a probability measure that is invariant under multiplication.

Definition 4. *[31] A group G is* amenable *if it admits a right-invariant, finitely additive probability measure μ:*

$$\forall \mathscr{R} \subseteq \mathscr{G}, \forall A \in \mathscr{G}, \mu(\mathscr{R}A) = \mu(\mathscr{R}) \ .$$

Theorem 1 can be extended to amenable groups.

Definition 5. *[32] Let $\mathbb{B}(\mathscr{H})$ denote the space of bounded linear operators on some Hilbert space $\mathscr{H}$. A representation $\pi : \mathscr{G} \to \mathbb{B}(\mathscr{H})$ is called* uniformly bounded *if* $\sup_{G \in \mathscr{G}}(\|\pi(G)\|_2) < \infty$.

Theorem 2. *[32] Every uniformly bounded representation $\pi : \mathscr{G} \to \mathbb{B}(\mathscr{H})$ of an amenable group $\mathscr{G}$ is* unitarizable*, i.e there exists $S \in \mathbb{B}(\mathscr{H})$ such that $G \mapsto S \circ \pi(G) \circ S^{-1}$ is a unitary operator.*

Corollary 1. *Any amenable subgroup of $GL(n,\mathbb{R})$ with bounded ℓ_2 operator norm is conjugate to $O(n,\mathbb{R})$.*

Proof. We can take π to be an identity mapping $GL(n,\mathbb{R}) \mapsto GL(n,\mathbb{R})$, since matrices can be treated as linear operators on vectors over $\mathbb{R}$. Any subgroup $\mathscr{G}$ of $GL(n,\mathbb{R})$ is its own representation in $GL(n,\mathbb{R})$, where bounded ℓ_2 operator norm implies that it is uniformly bounded.

Note that the elements of a group $\mathscr{G}$ that is conjugate to $O(n,\mathbb{R})$ are easily invertible. Any $A \in \mathscr{G}$is of the form $C^{-1}UC$ for an orthogonal matrix U. Since C is the same for the entire group, the matrix C^{-1} has to be found only once for the entire group, and we can compute $A^{-1} = C^{-1}(CAC^{-1})^T C$.

5.2 Security of $B = PA$

First of all, we discuss the transformation $B = PA$, where the matrix $A \in \mathscr{A}$ is only multiplied from the left with a randomly generated matrix P, since in the case of finite fields it is sufficient for perfect security.

Leaking eigenvectors We present a small example that shows that the transformation $B = PA$ puts an additional restriction on the set $\mathscr{A}$. Let $\mathscr{D}_P$ be the distribution of P. Consider any eigenvector $\mathbf{x}$ of A (a vector $\mathbf{x} \in \mathbb{R}^n$ such that $A\mathbf{x} = \lambda\mathbf{x}$ for some $\lambda \in \mathbb{R}$ which is called an *eigenvalue* of A). If $A\mathbf{x} = \lambda\mathbf{x}$, then $\mathscr{D}_P \cdot A\mathbf{x} = \lambda\mathscr{D}_P \cdot \mathbf{x}$, and $\mathscr{D}_P \cdot I\mathbf{x} = \mathscr{D}_P \cdot \mathbf{x}$. Hence if we include I into the set of matrices to be hidden, the distributions $\mathscr{D}_P \cdot A\mathbf{x}$ and $\mathscr{D}_P \cdot \mathbf{x}$ must be equal, and the only possibility is $\lambda = \pm 1$.

Formalizing $B = PA$ Suppose that there exists a distribution $\mathscr{D}_P$ of P that perfectly hides $\mathscr{A}$. Let $A_i \in \mathscr{A}$. Since $I \in \mathscr{A}$, there exist $P_i, P_j \in supp(\mathscr{D}_P)$ such that $P_iA_i = P_jI = P_j$, or $P_j^{-1}P_i = A_i$. Let $\mathscr{G} = \langle\mathscr{A}\rangle$ be the subgroup of $GL(n,\mathbb{R})$ generated by $\mathscr{A}$. Since $P_j^{-1}P_i \in \mathscr{G}$, P_i and P_j belong to the same left coset $H\mathscr{G}$ of $\mathscr{G}$ for some $H \in GL(n,\mathbb{R})$. The actual value of H does not even matter, as changing H_1 to H_2 just multiplies $\mathscr{D}_P$ by the same group element $H_2H_1^{-1}$ from the left.

We now study the properties of $\mathscr{G} = \langle\mathscr{A}\rangle$ for a set $\mathscr{A}$ for which there exists a perfectly hiding distribution $\mathscr{D}_P$ over $H\mathscr{G}$. We show that there exists a distribution $\mathscr{D}_\mathscr{G}$ s.t $supp(\mathscr{D}_\mathscr{G}) = \mathscr{G}$ that hides $\mathscr{A}$ as well.

Claim. There exists a distribution $\mathscr{D}_\mathscr{G}$ on $\mathscr{G}$ such that $\mathscr{D}_\mathscr{G}A = \mathscr{D}_\mathscr{G}$ for any $A \in \mathscr{A}$.

Proof. Let $\mathscr{D}_\mathscr{F} = \{|P_j^{-1}P_i \mid P_i \leftarrow \mathscr{D}_P, P_j \leftarrow \mathscr{D}_P|\}$, and $\mathscr{F} = supp(\mathscr{D}_\mathscr{F})$. $\mathscr{D}_\mathscr{F}$ is an example of a distribution over $\mathscr{F}$ that is hiding if $\mathscr{D}_P$ is hiding, since further multiplication by an independently sampled P_j^{-1} cannot harm the hiding achieved by the multiplication with $P_i \leftarrow \mathscr{D}_P$. Multiplying $\mathscr{D}_\mathscr{F}$ by *any* distribution over $\mathscr{G}$ from the left gives a

perfectly hiding distribution $\mathscr{D}_{\mathscr{G}}$, for the same reason that this multiplication does not make hiding worse. The definition of perfect hiding and $I \in \mathscr{A}$ give $\mathscr{D}_{\mathscr{G}}A = \mathscr{D}_{\mathscr{G}}I = \mathscr{D}_{\mathscr{G}}$ for any $A \in \mathscr{A}$.

The claim proves the existence of *some* hiding $\mathscr{D}_{\mathscr{G}}$ such that $supp(\mathscr{D}_{\mathscr{G}}) = \mathscr{G}$. However, choosing such $\mathscr{D}_{\mathscr{G}}$ for hiding is sufficient, but not necessary (we may actually sample from an arbitrary coset of $\mathscr{G}$ that is not a group). We now state a necessary condition for $supp(\mathscr{D}_P)$.

Claim. If $supp(\mathscr{D}_P)$ contains any element of a coset $H\mathscr{G}$ of $\mathscr{G}$, then it contains that coset entirely.

Proof. We prove the claim by induction on the length of the generating sequence of an element of $\mathscr{G}$.

- **Base:** We have shown that all $P_i \in supp(\mathscr{D}_P)$ belongs to the same coset $H\mathscr{G}$. Taking $H' := P_i$ for any $P_i \in supp(\mathscr{D}_P)$, we get $H\mathscr{G} = H'\mathscr{G}$, so for the base we may take *any* $P_i \in supp(\mathscr{D}_P)$.
- **Step:** Let $G \in \mathscr{G}$, then $G = A_1 \cdots A_n$ for $A_1, \ldots, A_n \in \mathscr{A}$. Then $P_i A_1 \ldots A_{n-1} A_n = P_i A_1 \cdots A_n$. By induction hypothesis, $P_i A_1 \cdots A_{n-1} \in supp(\mathscr{D}_P)$. Since $\mathscr{D}_P$ is perfectly hiding, there exists a matrix from $supp(\mathscr{D}_P)$ that makes I indistinguishable from A_n, and it is uniquely defined due to invertibility. At the same time, $P_i A_1 \cdots A_n$ is such a matrix. Hence $P_i A_1 \cdots A_n = P_i A_1 \cdots A_n I \in supp(\mathscr{D}_P)$.

Since $\mathscr{D}_{\mathscr{G}}$ is perfectly hiding, and $I \in \mathscr{A}$, we have $\mathscr{D}_{\mathscr{G}}A_i = \mathscr{D}_{\mathscr{G}}$ for all $A_i \in \mathscr{A}$. This can be extended to any $A \in \mathscr{G}$.

Claim. $\mathscr{D}_{\mathscr{G}}A = \mathscr{D}_{\mathscr{G}}$ for all $A \in \mathscr{G}$.

Proof. For all $A \in \mathscr{G}$, we have $A = A_1, \ldots, A_k$, where $A_i \in \mathscr{A}$ are generators of $\mathscr{G}$. Applying $\mathscr{D}_{\mathscr{G}}A_i = \mathscr{D}_{\mathscr{G}}$ k times, we get $\mathscr{D}_{\mathscr{G}}A = \mathscr{D}_{\mathscr{G}}A_1 \ldots A_k = \mathscr{D}_{\mathscr{G}}A_2 \ldots A_k = \ldots = \mathscr{D}_{\mathscr{G}}$.

We are interested in the distribution $\mathscr{D}_{\mathscr{G}}A = \mathscr{D}_{\mathscr{G}}$ for $A \in \mathscr{G}$. If $\mathscr{G}$ is a finite group, then $\mathscr{D}_{\mathscr{G}}$ must be a uniform distribution on $\mathscr{G}$, since in a group each product is a distinct element. By Theorem 1, a finite group is conjugate to $O(n, \mathbb{R})$. We are more interested in the cases where $\mathscr{G}$ is infinite.

Claim. Let $\mathscr{D}_{\mathscr{G}}A = \mathscr{D}_{\mathscr{G}}$ for all $A \in \mathscr{G}$. Then $\mathscr{G}$ is amenable.

Proof. We take the probability measure $\mu_{\mathscr{D}_{\mathscr{G}}}$ of $\mathscr{D}_{\mathscr{G}}$. Take any $\mathscr{R} \subseteq \mathscr{G}$. Then $\forall A \in \mathscr{G} :\ \mu_{\mathscr{D}_{\mathscr{G}}}(\mathscr{R}) = \mu_{\mathscr{D}_{\mathscr{G}}A}(\mathscr{R}) = \mu_{\mathscr{D}_{\mathscr{G}}}(\mathscr{R}A^{-1})$, which is equivalent to $\forall A \in \mathscr{G} :\ \mu_{\mathscr{D}_{\mathscr{G}}}(\mathscr{R}) = \mu_{\mathscr{D}_{\mathscr{G}}}(\mathscr{R}A)$. The measure $\mu_{\mathscr{D}_{\mathscr{G}}}$ is suitable for the amenability definition.

The definition of amenability is not too interesting on its own, but it tells something more due to the ℓ_2-norm boundedness requirement.

Claim. The elements of $\mathscr{A}$ are conjugate to $O(n, \mathbb{R})$.

Proof. By Claim 5.2, we need to sample from a distribution $\mathscr{D}_P$ such that $supp(\mathscr{D}_P) = H\mathscr{G}$. We have agreed in Sec. 3.2 that we deal with P with bounded ℓ_2-norms, so the ℓ_2 norms of $H\mathscr{G}$ should be bounded. Hence the norms of $\mathscr{G}$ are also bounded: $\forall G \in \mathscr{G}$: $\|G\|_2 = \|H^{-1}HG\|_2 \leq \|H^{-1}\|_2\|HG\|_2$.

By Claim 5.2, $\mathscr{G}$ is amenable. By Corollary 1, the elements of $\mathscr{G}$ are conjugate to $O(n,\mathbb{R})$. We have $\mathscr{A} \subseteq \mathscr{G}$.

We conclude that it is indeed possible to find a perfect hiding of the form $B = PA$ for certain sets $\mathscr{A}$ for which the group $\langle\mathscr{A}\rangle$ is amenable, but knowing that $A \in \mathscr{A}$ would already give enough information about the inverse of A^{-1} so that it could be computed easily without the transformation.

5.3 Security of $B = PAQ$

From Claim 5.2, we conclude that the transformation $B = PA$ is not powerful enough for perfect secrecy. We proceed with $B = PAQ$.

Positive Examples First, we give small examples that demonstrates additional hiding that the transformation $B = PAQ$ can give compared to just $B = PA$. Consider 2×2 matrices. We show perfectly hiding examples for both $B = PA$ and $B = PAQ$.

Example 1. Let $A_1 = \begin{pmatrix} \lambda & 0 \\ 0 & 1/\lambda \end{pmatrix}$ for some $\lambda \in \mathbb{R}$ and $A_2 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. Let $\mathscr{D}_{\mathscr{P}}$ be the distribution that uniformly chooses between $P_1 = \begin{pmatrix} 1/\lambda & 0 \\ 0 & 1 \end{pmatrix}$ and $P_2 = \begin{pmatrix} 0 & 1 \\ 1/\lambda & 0 \end{pmatrix}$. In this case, the distributions $\mathscr{D}_{\mathscr{P}}A_1$ and $\mathscr{D}_{\mathscr{P}}A_2$ are equal. However, if we took $A_2 = I$ instead, perfect hiding would be possible only if $\lambda = \pm 1$, due to the $\lambda\mathscr{D}_{\mathscr{P}}\mathbf{x} = \mathscr{D}_{\mathscr{P}}\mathbf{x}$ requirement for an eigenvector $\mathbf{x}$ of A.

Example 2. Let $A_1 = \begin{pmatrix} \lambda & 0 \\ 0 & 1/\lambda \end{pmatrix}$ and $A_2 = I$. Let $(P_1, Q_1) = \left(\begin{pmatrix} 1/\lambda & 0 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}\right)$ and $(P_2, Q_2) = \left(\begin{pmatrix} 0 & 1 \\ 1/\lambda & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}\right)$. Let $\mathscr{D}_{(P,Q)}$ be a distribution over pairs of matrices that uniformly chooses between (P_1, Q_1) and (P_2, Q_2). Then $\mathscr{D}_{(P,Q)}$ perfectly hides $\{A_1, A_2\}$. This would be impossible to achieve with the $B = PA$ transformation.

Example 3. Let $\mathscr{G}_P = B^{-1}\mathscr{U}B$ and $\mathscr{G}_Q = C^{-1}\mathscr{V}C$ for some subgroups $\mathscr{U}, \mathscr{V}$ of $O(n,\mathbb{R})$, and $B, C \in GL(n,\mathbb{R})$. These groups are amenable, and we can sample from them uniformly (for example, taking $\mathscr{U} = O(n,\mathbb{R})$, first sample uniformly $U \xleftarrow{\$} O(n,\mathbb{R})$ [33], and then multiply it by B^{-1} and B). Take $\mathscr{D}_{(P,Q)} = \{|(P,Q) \mid P \xleftarrow{\$} \mathscr{G}_P, Q \xleftarrow{\$} \mathscr{G}_Q|\}$. Such a distribution hides perfectly the set $\mathscr{A} = \{|PQ \mid P \xleftarrow{\$} \mathscr{G}_P, Q \xleftarrow{\$} \mathscr{G}_Q|\}$. Note that the elements of $\mathscr{A}$ are no longer easily invertible in the same way as matrices in an amenable group (unless $B = C$). Hence such a construction could be useful theoretically, although the question is how exactly $A \in \mathscr{A}$ could be generated in practice (if the client already knows (P,Q) such that $A = PQ$, then he already knows the inverse). We also cannot claim that there exists no other efficient algorithm of inverting matrices of this form.

Formalizing $B = PAQ$ Suppose that there exists a perfectly hiding distribution $\mathscr{D}_{(P,Q)}$ of P and Q for $\mathscr{A}$. Similarly to the $B = PA$ case, we may rewrite $P_iA_kQ_i = P_jA_\ell Q_j$ as $P_j^{-1}P_iA_kQ_iQ_j^{-1} = A_\ell$. If multiplying by (P_i, Q_i) once provides sufficient hiding, then multiplying the result again with an independently sampled (P_j^{-1}, Q_j^{-1}) is hiding as well. Let $\mathscr{D}_{\mathscr{F}} = \{|(P_j^{-1}P_i, Q_iQ_j^{-1}) \mid (P_i, Q_i), (P_j, Q_j) \leftarrow \mathscr{D}_{(P,Q)}|\}$ and $\mathscr{F} = supp(\mathscr{D}_{\mathscr{F}})$. Define a group $\mathscr{G} := \langle \mathscr{F} \rangle$, where the inverse is defined as $(P,Q)^{-1} = (P^{-1}, Q^{-1})$, and the multiplication as $(P_1, Q_1) * (P_2, Q_2) = (P_1P_2, Q_2Q_1)$. We may now consider the hiding process as an action of the group $\mathscr{G}$ onto the set $\mathscr{A}$, defined as $(P,Q).A = PAQ$ for $A \in \mathscr{A}$, $(P,Q) \in \mathscr{G}$. Define the following sets:

- $\mathscr{X}_0 := \mathscr{A}$,
- $\mathscr{X}_i := \mathscr{G}.\mathscr{X}_{i-1}$,
- $\mathscr{X} := \bigcup_{i=0}^{\infty} \mathscr{X}_i$.

By construction, $\mathscr{A} \subseteq \mathscr{X}$.

(Im)possibility results for $B = PAQ$ In all the claims proven it this section, we assume that there exists a distribution $\mathscr{D}_{(P,Q)}$ that hides $\mathscr{A}$ perfectly. We study the action of $\mathscr{G}$ on $\mathscr{X}$ and find some its interesting properties.

Claim. Let $\mathscr{G}$ and $\mathscr{X}$ be defined as in Sec. 5.3. The set $\mathscr{X}$ is isomorphic (as an action of the group $\mathscr{G}$) to the set of cosets of $\mathscr{H}_{X_0} := \{G | G \in \mathscr{G}, G.X_0 = X_0\}$ (for any $X_0 \in \mathscr{X}$). The isomorphism is $G\mathscr{H}_{X_0} \leftrightarrow G.X_0$.

Proof. Since any A_i must be indistinguishable from any A_j, $\mathscr{G}$ is able to map any $A_i \in \mathscr{A}$ to any $A_j \in \mathscr{A}$. Thus, for all $X \in \mathscr{X}$, there exists $G \in \mathscr{G}$ such that $G.A_i = X$, and hence for all $X_i, X_j \in \mathscr{X}$ there exists $G \in \mathscr{G}$ such that $G.X_i = X_j$. That is, the given group action has only one orbit that contains the entire $\mathscr{X}$. Such group actions are called transitive. By the orbit-stabilizer theorem, for any subgroup $\mathscr{H}_{X_0} := \{G | G \in \mathscr{G}, G.X_0 = X_0\}$ ($X_0 \in \mathscr{X}$), $\mathscr{G}$ acts on $\mathscr{X}$ similarly as it would act on the set of its left cosets $\mathscr{G}/\mathscr{H}_{X_0}$ by multiplication from the left, and the isomorphism between $\mathscr{X}$ and $\mathscr{G}/\mathscr{H}_{X_0}$ is $G\mathscr{H}_{X_0} \leftrightarrow G.X_0$ for $G \in \mathscr{G}$. These group-theoretical results can be found for example in [34].

The next claim is similar to Claim 5.2 which states that $\mathscr{G}$ itself perfectly hides $\mathscr{A}$.

Claim. Let $\mathscr{G}$ be defined as in Sec. 5.3. There exists a distribution $\mathscr{D}_{\mathscr{G}}$ on $\mathscr{G}$, such that $\mathscr{D}_{\mathscr{G}.A_i} = \mathscr{D}_{\mathscr{G}.A_j}$ for all $A_i, A_j \in \mathscr{A}$, where $\mathscr{D}_{\mathscr{G}.A} := \{|G.A \mid G \leftarrow \mathscr{D}_{\mathscr{G}}|\}$.

Proof. For any $A_i, A_j \in \mathscr{A}$, the distributions $\mathscr{D}_{\mathscr{F}.A_i} = \{|F.A_i \mid F \leftarrow \mathscr{D}_{\mathscr{F}}|\}$ and $\mathscr{D}_{\mathscr{F}.A_j} = \{|F.A_j \mid F \leftarrow \mathscr{D}_{\mathscr{F}}|\}$ are indistinguishable. Since $\mathscr{G} = \langle \mathscr{F} \rangle$, similarly to the proof of Claim 5.2 this implies that there exists a distribution $\mathscr{D}_{\mathscr{G}}$ on $\mathscr{G}$, such that $\mathscr{D}_{\mathscr{G}.I} = \mathscr{D}_{\mathscr{G}.A_i}$.

In the case of $B = PA$, we observed properties of the group $\langle \mathscr{A} \rangle$. It would be interesting to find something similar in $B = PAQ$. For each $A_i \in \mathscr{A}$ there exists $(P_i, Q_i) \in supp(\mathscr{D}_{(P,Q)})$ such that $A_i = P_iIQ_i$.

Claim. It is possible to sample uniformly from $\mathscr{A}$.

Proof. Let $\mathscr{G}$ be defined as in Sec. 5.3. According to Claim 5.3, there exists a perfectly hiding distribution $\mathscr{D}_{\mathscr{G}}$ on $\mathscr{G}$. In terms of cosets, according to Claim 5.3 we have a mapping $\mathscr{G} \times \mathscr{G}/\mathscr{H}_{X_0} \mapsto \mathscr{G}/\mathscr{H}_{X_0}$, and for $G_i, G_j \in \mathscr{G}$ such that $A_i = G_i.X_0$ and $A_j = G_j.X_0$, the distributions over cosets $\mathscr{D}_{\mathscr{G}G_i\mathscr{H}_{X_0}} = \{|GG_i\mathscr{H}_{X_0} \mid G \leftarrow \mathscr{D}_{\mathscr{G}}|\}$ and $\mathscr{D}_{\mathscr{G}G_j\mathscr{H}_{X_0}} = \{|GG_j\mathscr{H}_{X_0} \mid G \leftarrow \mathscr{D}_{\mathscr{G}}|\}$ are the same.

Let $G_i = (I, I)$. Take any $G_k \in \mathscr{G}$. We have $Pr[G_k\mathscr{H}_{X_0} \in \mathscr{D}_{\mathscr{G},\mathscr{H}_{X_0}}] = Pr[G_k\mathscr{H}_{X_0} \in \mathscr{D}_{\mathscr{G}G_j\mathscr{H}_{X_0}}]$. Let $G_k = (I, I)$. The elements of $\mathscr{G}$ that map $\mathscr{H}_{X_0}$ to $\mathscr{H}_{X_0}$ are exactly those from $\mathscr{H}_{X_0}$. The elements of $\mathscr{G}$ that map $G_j\mathscr{H}_{X_0}$ to $\mathscr{H}_{X_0}$ are exactly those from $\mathscr{H}_{X_0}G_j^{-1}$.

More formally, for all $G_j \in \mathscr{G}$ s.t $A_j = G_j.X_0$ we have $Pr[G \in \mathscr{H}_{X_0} \mid G \leftarrow \mathscr{D}_{\mathscr{G}}] = Pr[G \in \mathscr{H}_{X_0}G_j^{-1} \mid G \leftarrow \mathscr{D}_{\mathscr{G}}]$. Since there is an isomorphism $G\mathscr{H}_{X_0} \leftrightarrow \mathscr{H}_{X_0}G^{-1}$ between the left and the right cosets, define a distribution $\widetilde{\mathscr{D}}_{\mathscr{G}} = \{|G^{-1} \mid G \leftarrow \mathscr{D}_{\mathscr{G}}|\}$, getting $Pr[G \in \mathscr{H}_{X_0} \mid G \leftarrow \widetilde{\mathscr{D}}_{\mathscr{G}}] = Pr[G \in G_j\mathscr{H}_{X_0} \mid G \leftarrow \widetilde{\mathscr{D}}_{\mathscr{G}}]$. Hence there exists a distribution of $\mathscr{G}$ where the elements are distributed uniformly among the left cosets that are isomorphic to $\mathscr{A}$.

Similarly to the $B = PA$ case, being able to sample uniformly from $\mathscr{G}$ would be sufficient for hiding. Due to the isomorphism with cosets of some $\mathscr{H}_{A_0}$, each $A \in \mathscr{A}$ can be written out as $A = G_i.A_0$. Given a uniform distribution $\mathscr{D}_{\mathscr{G}}$ on $\mathscr{G}$, we have $\mathscr{D}_{\mathscr{G}}(G_i.A_0) = (\mathscr{D}_{\mathscr{G}}G_i).A_0 = \mathscr{D}_{\mathscr{G}}.A_0$. This hides the entire $\mathscr{X}$. However, it may happen that hiding the entire $\mathscr{X}$ is a too strong requirement. If we could sample uniformly from $\mathscr{G}$, then we could also sample from $\mathscr{X}$, as group elements are distributed uniformly among its cosets. The question is whether we can always sample uniformly from $\mathscr{X}$.

Claim. Let $\mathscr{G}$ and $\mathscr{X}$ be defined as in Sec. 5.3. Being able to sample uniformly from $\mathscr{A}$ does not imply sampling uniformly from $\mathscr{X}$.

Proof. A simple counterexample is the construction of Example 2. It is easy to sample uniformly from $\mathscr{A}$ which is a finite set. We have $A_1 = P_2^{-1}P_1A_1Q_1Q_2$. Defining the group $\mathscr{G} = \langle (P_2^{-1}P_1, Q_1Q_2^{-1}) \rangle$ as before, we get $\mathscr{G} = \left\langle \left(\begin{pmatrix} 0 & \lambda \\ 1/\lambda & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \right) \right\rangle$, where the first component gets infinite powers of λ in $\mathscr{G}$, and multiplying it by the second component, which is either I or $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, does not help to decrease these powers.

A more interesting question is if there *exists* some $\mathscr{G}'$ such that $\mathscr{X}'$ is uniformly distributed. At least for $n = 1$, the only A s.t $|A| = 1$ is $A = (1)$. We take $P = (1)$ and $Q = (1)$. However, for $n = 2$, this is already not possible in general.

Claim. Being able to sample uniformly from $\mathscr{A}$ does not imply existence of $\mathscr{D}_{(P,Q)}$ such that $\mathscr{X}$ (as constructed in Sec. 5.3) admits a uniform distribution.

Proof. Take the same Example 2. In the previous choice, the entries of $\langle Q \rangle$ are bounded by 1 (as $Q^2 = I$), but the entries of $\langle P \rangle$ are unbounded. Let us now take the decomposition of A_1 where the sizes of entries of P and Q are as close as possible, which is the square root of A_1. A 2×2 matrix with two distinct nonzero eigenvalues has exactly four

square roots, which are of the form $\begin{pmatrix} \pm\sqrt{\lambda} & 0 \\ 0 & \pm 1/\sqrt{\lambda} \end{pmatrix}$. In all these cases, we get that some elements of $\mathscr{X}$ are of the form $(P,Q)^k.I = \begin{pmatrix} \pm\sqrt{\lambda} & 0 \\ 0 & \pm 1/\sqrt{\lambda} \end{pmatrix}^{2k} = \begin{pmatrix} \lambda^k & 0 \\ 0 & 1/\lambda^k \end{pmatrix}$, which still have infinitely growing entries. We conclude that the construction of Example 3 is not the most general possible, as Example 2 is not its instance. We have shown that Claim 5.2 of Sec. 5.2 does not extend to the $B = PAQ$ case.

Due to Claim 5.3, we cannot conclude that $\mathscr{G}$ is necessarily amenable in general, as it was in the case of $B = PA$.

5.4 Conclusion for Perfect Hiding

Secure transformation of the form $B = PA$, that works in finite fields, implies being able to sample uniformly from the group $\langle \mathscr{A} \rangle$. If the ℓ_2-norms of P and Q are bounded, then the ℓ_2-norms of $\langle \mathscr{A} \rangle$ are also bounded. According to known facts about groups, this group is conjugate to $O(n,\mathbb{R})$, and hence there is an easier way of inverting its elements.

If the transformation $B = PAQ$ achieves perfect hiding, then we must be able to sample uniformly from $\mathscr{A}$. Since the elements of $\mathscr{A}$ come from real tasks, we may assume that their ℓ_2-norms are bounded. If $\langle \mathscr{A} \rangle = \mathscr{A}$, then $\mathscr{A}$ is conjugate to $O(n,\mathbb{R})$, and hence there is an easier way of inverting its elements.

6 Statistical hiding

The previous results are partially extensible to statistical hiding. Instead of taking perfectly hiding $\mathscr{D}$, we might take just some part of it, leaving behind a certain quantity that comes with negligible probability. However, if we had negligible statistical distance instead of equality of distributions in Claim 5.2, then the statistical distance between $\mathscr{D}_\mathscr{G}$ and $\mathscr{D}_\mathscr{G}A$ might grow with k, and hence this argument does not pass straightforwardly for statistical security. Claim 5.2 also does not pass: it no longer makes sense to require that the ℓ_2-norms of matrices of $\mathscr{D}$ are bounded, since we do not sample all its elements anyway. In this section, we present some possibility results for statistical hiding.

6.1 Hiding two arbitrary matrices by $B = PA$

We start from the simplest case. It is easy to achieve statistical security for a set $\mathscr{A}$ consisting of just two matrices A_1 and A_2, without even requiring that the absolute values of their determinants must be the same. Starting from $P = (A_1 A_2^{-1})$, sequentially generate P_{i+1} such that $P_i A_1 = P_{i+1} A_2$, getting $P^k = (A_1 A_2^{-1})^k$. Let $\mathscr{D}_\mathscr{G}$ uniformly pick from $P_1, \ldots, P_l$ for some limit l. Even if there is no t such that $P^t = P$, the statistical difference will be at most $\frac{2}{|supp(\mathscr{D}_\mathscr{G})|}$, where $\mathscr{D}_\mathscr{G}$ is uniform distribution on the finite number of matrices that we take into the support.

An analogical construction for k matrices would produce statistical difference $\approx 1 - \frac{1}{k-1}$, so it cannot be generalized straightforwardly to k matrices without setting any

constraints on $\mathscr{A}$. Additionally, if the absolute values of determinants of A_1 and A_2 are different (let $|A_1A_2^{-1}| = d$ for $|d| > 1$), then $|P_i| = d^i$ and the sizes of entries of P_l are of the order $2^{\eta/n}$ for η bits of security, as shown in Sec. 4.

6.2 Hiding triangular matrices by $B = PA$

In Sec. 5.2 we have shown that we cannot perfectly hide matrices with eigenvalues $\lambda \neq \pm 1$. We can show that it extends to statistical hiding. Let $\mathscr{D}_P$ be the distribution of the matrices P.

Claim. Let η be a security parameter. If there exists $A \in \mathscr{A}$ with an eigenvalue $\lambda \neq \pm 1$ such that $SD(\mathscr{D}_P A, \mathscr{D}_P) < 2^{-\eta}$, then the sizes of entries of P are of order $2^{\eta-1}$.

Proof. In the proof of Claim 4, we have proven a more general result: if $SD(\mathscr{D}, c\mathscr{D}) < 2^{-\eta}$ for some $c > 1$ such that $\mathscr{D}$ is a bounded distribution over $\mathbb{R}$, we have that the sizes of the values in the support of $\mathscr{D}$ should be of order 2^{η}. Now consider the distributions $\mathscr{D}_P A\mathbf{x} = \lambda \cdot \mathscr{D}_P\mathbf{x}$ and $\mathscr{D}_P\mathbf{x}$ for any eigenvector $\mathbf{x}$ of A. Let $\mathscr{D}'$ be the distribution of a fixed entry of $\mathscr{D}_P\mathbf{x}$. In order to get statistical hiding, we must achieve $SD(\mathscr{D}', \lambda\mathscr{D}') < 2^{-\eta}$. Similarly to the proof of Claim 4, since the sizes of entries of $\mathscr{D}_P$ (and hence also $\mathscr{D}_P\mathbf{x}$ for a fixed $\mathbf{x}$, and $\mathscr{D}'$), are bounded, we get that the sizes of values from support of $\mathscr{D}'$ must be of the order 2^{η}. Each value of $\mathscr{D}'$ is an entry of $\mathscr{D}_P\mathbf{x}$.

The scaling of an eigenvector results in an eigenvector. Hence, without lessening the generality, we may assume $\|\mathbf{x}\|_2 = 1$. For all $P \in supp(\mathscr{D}_P)$, we have $\|P\mathbf{x}\|_2 \leq \|P\|_2$. All the entries of $P\mathbf{x}$ are of order 2^{η}, and so also $\|P\mathbf{x}\|_2$ is of order 2^{η}, and $\|P\|_2 \leq \sqrt{\|P\|_1\|P\|_\infty} = \sqrt{\sum_{ij} p_{ij}^2}$, where p_{ij} are entries of P. This means that the sizes of entries of $\mathscr{D}_P$ should be of order $2^{\eta}/2$, which is not acceptable.

Let A be a lower triangular matrix with $a_{ii} = \pm 1$ (this is exactly the case when we have $\lambda = \pm 1$ for each eigenvalue λ of A). Let $P = (\mathbf{p}_1|\ldots|\mathbf{p}_n)$. Due to triangularity of A, We can write $PA = (a_{11}\mathbf{p}_1 + \cdots + a_{n1}\mathbf{p}_n|\ldots|a_{n-1,n-1}\mathbf{p}_{n-1} + a_{n,n-1}\mathbf{p}_n, a_{nn}\mathbf{p}_n)$. Since $a_{nn} = 1$, the last vector of PA does not depend on A.

We see that, starting from an arbitrary $\mathbf{p}_n$, we can generate $\mathbf{p}_i$ in such a way that it hides $\mathbf{p}_{i+1}$. Namely, if the entries of $\mathbf{p}_n$ are bounded by 1 (which is the least possible bound), let the entries of $\mathbf{p}_{n-1}$ be uniformly distributed between 0 and $c \cdot 2^{\eta}$, where $c \geq |a_{n,n-1}|$. In this way, $a_{n-1,n-1}\mathbf{p}_{n-1} + a_{n,n-1}\mathbf{p}_n$ statistically hides $\mathbf{p}_n$ with at least η bits of security. On the next step, to hide the $(n-2)$-th column with at least η bits of security, we have to sample the entries of $\mathbf{p}_{n-2}$ uniformly from between 0 and $c(2^{\eta} + 2^{2\eta})$, where c is an upper bound for both $|a_{n,n-2}|$ and $|a_{n-1,n-2}|$. Proceeding to $\mathbf{p}_1$, we get that its entries should be sampled uniformly from between 0 and $c(2^{\eta} + 2^{2\eta} + \cdots + 2^{(n-1)\eta})$, where c is an upper bound for the absolute values of all entries in the first column of A.

The preceding discussion shows that in order to statistically hide matrices from a set $\mathscr{A}$ satisfying the following:

- a matrix $A \in \mathscr{A}$ is lower triangular with the main diagonal containing only the values ± 1;
- the absolute values of the entries of the matrices in $\mathscr{A}$ are bounded by c;

it is sufficient to sample the matrix P according to a distribution described in the previous paragraph. The sizes of the entries of P are bounded by $(\log c)n\eta$. These entries, although not small, are still of the acceptable size according to our discussion in Sec. 3.2. On the other hand, lower triangular matrices can be inverted much more easily than general matrices, and it does not make any practical sense to outsource it like that.

6.3 Hiding arbitrary matrices by $B = PAQ$

Hiding based on QR-decomposition Let $\mathscr{A}$ be an arbitrary set of $n \times n$ matrices with bounded entries, and determinant ± 1. Any matrix can be represented as $A = Q_A R_A$, where $Q_A \in O(n,\mathbb{R})$ and R_A is an upper triangular matrix with positive entries on the diagonal (such a decomposition is unique). Note that $|Q_A| = 1$ in any case. We put the following further restrictions on $\mathscr{A}$:

- there exists $c \in \mathbb{R}$ that upper bounds the absolute values of the entries of R_A for any $A \in \mathscr{A}$;
- the entries on the main diagonal of R_A are ± 1.

Taking $\mathscr{D}_P$ from Sec. 6.2, we may now take $\mathscr{D}_{(P,Q)} = \{|(U, P^T) \mid U \xleftarrow{\$} O(n,\mathbb{R}), P \leftarrow \mathscr{D}_P|\}$. The distribution $\mathscr{D}_{(P,Q)}$ statistically hides $\mathscr{A}$, because for any $A \in \mathscr{A}$, the matrix UQ_A is a uniformly distributed orthogonal matrix, and the product $R_A P^T$ statistically hides R_A.

There exist algorithms that allow to sample uniformly from $O(n,\mathbb{R})$. The question is whether we can sample efficiently uniformly from $O(n,\mathbb{R})$ in practice. Some algorithms can be found in [33], where the best complexity of generating pseudorandom orthogonal matrices is $O(n^2 \log n)$. Some of the more straightforward methods involve the generation of a random matrix (with certain distributions for its entries) and finding its QR-decomposition. This is not easier than inverting the matrix A.

If our setting is secure multiparty computation (SMC), then the proposed hiding method can be useful. Indeed, in this setting, a random orthogonal matrix can be generated simply by each computing party locally generating a random orthogonal matrix, entering it into the computation, and multiplying these matrices using SMC protocols.

Hiding based on LU-decomposition Again, let $\mathscr{A}$ be a set of $n \times n$ matrices with bounded entries, and determinant ± 1. Almost any matrix A can be decomposed as $A = L_A U_A$, where L_A is lower and R_A upper diagonal [35, Chapter 3]. If the set $\mathscr{A}$ satisfies the condition that there exists $c \in \mathbb{R}$, such that each $A \in \mathscr{A}$ has a LU decomposition, where

- the absolute values of the entries in L_A and U_A are upper-bounded by c,
- the entries on the main diagonal of L_A and U_A are ± 1,

then the following distribution $\mathscr{D}_{(P,Q)}$ provides statistical hiding for $\mathscr{A}$. Let $\mathscr{D}_P$ be defined as in Sec. 6.2. Let $\mathscr{D}_{(P,Q)} = \{|(P_1, P_2^T) \mid P_1, P_2 \leftarrow \mathscr{D}_P|\}$. Indeed, $P_1 L_A$ statistically hides L_A and $U_A P_2^T$ statistically hides U_A for any $A \in \mathscr{A}$.

Hiding based on SVD-decomposition Let now $\mathscr{A}$ be a set of matrices such that $\forall A \in \mathscr{A} : \ \sigma = 1$, where σ is a *singular value* of A (an eigenvalue of the matrix $A^T A$). Any invertible matrix can be decomposed as $A = U_A S V_A$, where $U_A, V_A \in O(n, \mathbb{R})$, and S is a diagonal matrix of singular values that are the same for all A. The distribution $\mathscr{D}_{(P,Q)} = \{|(P,Q) \mid P \xleftarrow{\$} O(n,\mathbb{R}), Q \xleftarrow{\$} O(n,\mathbb{R})|\}$ provides perfect hiding for $\mathscr{A}$: PU_A perfectly hides U_A, $V_A Q$ perfectly hides V_A, and S is the same for all $A \in \mathscr{A}$ up to permutation which is taken into account by both P and Q. This hiding is somewhat similar to the more general hiding of Example 3 of Sec. 5.3, which can be applied assuming that the singular values of $\mathscr{A}$ are the same.

More generally, this hiding method is applicable for any set of matrices $\mathscr{A}$, where all elements of $\mathscr{A}$ have the same singular values. In effect, this amounts to adding the singular values of the matrix to the output of the side information function Φ.

Hiding based on eigendecomposition Let $\mathscr{A}$ be a set of symmetric matrices such that $\forall A \in \mathscr{A} : \ \lambda = \pm 1$, where λ is an eigenvalue of A. Then each $A \in \mathscr{A}$ can be decomposed as $U_A D U_A^T$ where $U_A \in O(n,\mathbb{R})$, and D is a diagonal matrix of eigenvalues. The distribution $\mathscr{D}_{(P,Q)} = \{|(P,P^T) \mid P \xleftarrow{\$} O(n,\mathbb{R})|\}$ provides perfect hiding for $\mathscr{A}$: PU_A perfectly hides U_A, $U_A^T P^T$ perfectly hides U_A^T, and D is the same for all $A \in \mathscr{A}$.

6.4 Conclusion for Statistical Hiding

As a conclusion, we can say that it is indeed possible to achieve statistical multiplicative hiding for an arbitrary set $\mathscr{A}$. In practice, such solutions would be not too efficient in general, and the matrix entries should be extremely large if we want to hide matrices with different determinant absolute values. Still, for some interesting classes of low-dimensional matrices, statistical hiding is possible with an effort that may be smaller than directly inverting the matrix.

7 Conclusion

We have studied the methods for outsourcing the inversion of a non-singular matrix over real numbers, based on affine transformations. We have shown that most general type of affine transformation is of the form $A \mapsto P \begin{pmatrix} A & 0 \\ 0 & I \end{pmatrix} Q$. Only matrices with the same determinant absolute values can be perfectly indistinguishable. If we want to achieve statistical indistinguishability for arbitrary A_1 and A_2, the entries of P and Q grow in $\Omega(||A_1|/|A_2||^{2^\eta})$ for a security parameter η.

We have found that over reals, it is much more difficult to hide matrices by multiplying them with random matrices. If we try to limit hiding to just multiplication by P with bounded operator ℓ_2 norms, which works well in matrices over finite fields, being able to achieve perfect secrecy implies that we should be able to sample uniformly from a certain group $\mathscr{G}$ that depends on the set of matrices $\mathscr{A}$ that we want to hide. According to known facts about the properties of groups, if we can sample from a group uniformly, this group is conjugate to an orthogonal group, and hence there is an easier

way of inverting its elements. The set of matrices $\mathscr{A}$ is related to $\mathscr{G}$ in such a way that inverting the elements of $\mathscr{G}$ make it easy to invert the elements of $\mathscr{A}$.

If we use the $A \mapsto PAQ$ transformation with bounded ℓ_2 operator norms for hiding, then we should be able to sample uniformly from $\mathscr{A}$. If we are using matrices P and Q with unbounded ℓ_2 operator norms, then we can still achieve statistical security. We have given some possibility results that nevertheless leak the determinant, unless the entries grow in $\Omega(c^{2^\eta})$ where c is the maximal ratio of different determinants.

A summary of our findings is depicted in Fig. 1.

References

1. Susan Hohenberger and Anna Lysyanskaya. How to securely outsource cryptographic computations. In Joe Kilian, editor, *TCC*, volume 3378 of *Lecture Notes in Computer Science*, pages 264–282. Springer, 2005.
2. Xiaofeng Chen, Jin Li, Jianfeng Ma, Qiang Tang, and Wenjing Lou. New algorithms for secure outsourcing of modular exponentiations. In Sara Foresti, Moti Yung, and Fabio Martinelli, editors, *ESORICS*, volume 7459 of *Lecture Notes in Computer Science*, pages 541–556. Springer, 2012.
3. Sergei Evdokimov and Oliver Günther. Encryption techniques for secure database outsourcing. In Joachim Biskup and Javier Lopez, editors, *ESORICS*, volume 4734 of *Lecture Notes in Computer Science*, pages 327–342. Springer, 2007.
4. Mikhail J. Atallah and Jiangtao Li. Secure outsourcing of sequence comparisons. *Int. J. Inf. Sec.*, 4(4):277–287, 2005.
5. Mikhail J. Atallah and Keith B. Frikken. Securely outsourcing linear algebra computations. In Dengguo Feng, David A. Basin, and Peng Liu, editors, *ASIACCS*, pages 48–59. ACM, 2010.
6. Jannik Dreier and Florian Kerschbaum. Practical privacy-preserving multiparty linear programming based on problem transformation. In *SocialCom/PASSAT*, pages 916–924. IEEE, 2011.
7. Cong Wang, Kui Ren, and Jia Wang. Secure and practical outsourcing of linear programming in cloud computing. In *INFOCOM, 2011 Proceedings IEEE*, pages 820–828, 2011.
8. Ronald Cramer and Ivan Damgård. Secure distributed linear algebra in a constant number of rounds. In *Proceedings of the 21st Annual International Cryptology Conference on Advances in Cryptology*, CRYPTO '01, pages 119–136, London, UK, UK, 2001. Springer-Verlag.
9. Wenliang Du and Mikhail J. Atallah. Privacy-preserving cooperative scientific computations. In *Proceedings of the 14th IEEE Workshop on Computer Security Foundations*, CSFW '01, pages 273–, Washington, DC, USA, 2001. IEEE Computer Society.
10. Wenliang Du, Shigang Chen, and Yunghsiang S. Han. Privacy-preserving multivariate statistical analysis: Linear regression and classification. In *In Proceedings of the 4th SIAM International Conference on Data Mining*, pages 222–233, 2004.
11. Shuguo Han and Wee Keong Ng. Privacy-preserving linear fisher discriminant analysis. In *Proceedings of the 12th Pacific-Asia conference on Advances in knowledge discovery and data mining*, PAKDD'08, pages 136–147, Berlin, Heidelberg, 2008. Springer-Verlag.
12. Xuan Yang, Zhaoping Yu, and Bin Kang. Privacy-preserving cooperative linear system of equations protocol and its application. In *Wireless Communications, Networking and Mobile Computing, 2008. WiCOM '08. 4th International Conference on*, pages 1–4, Oct 2008.
13. Ju sung Kang and Dowon Hong. A practical privacy-preserving cooperative computation protocol without oblivious transfer for linear systems of equations.

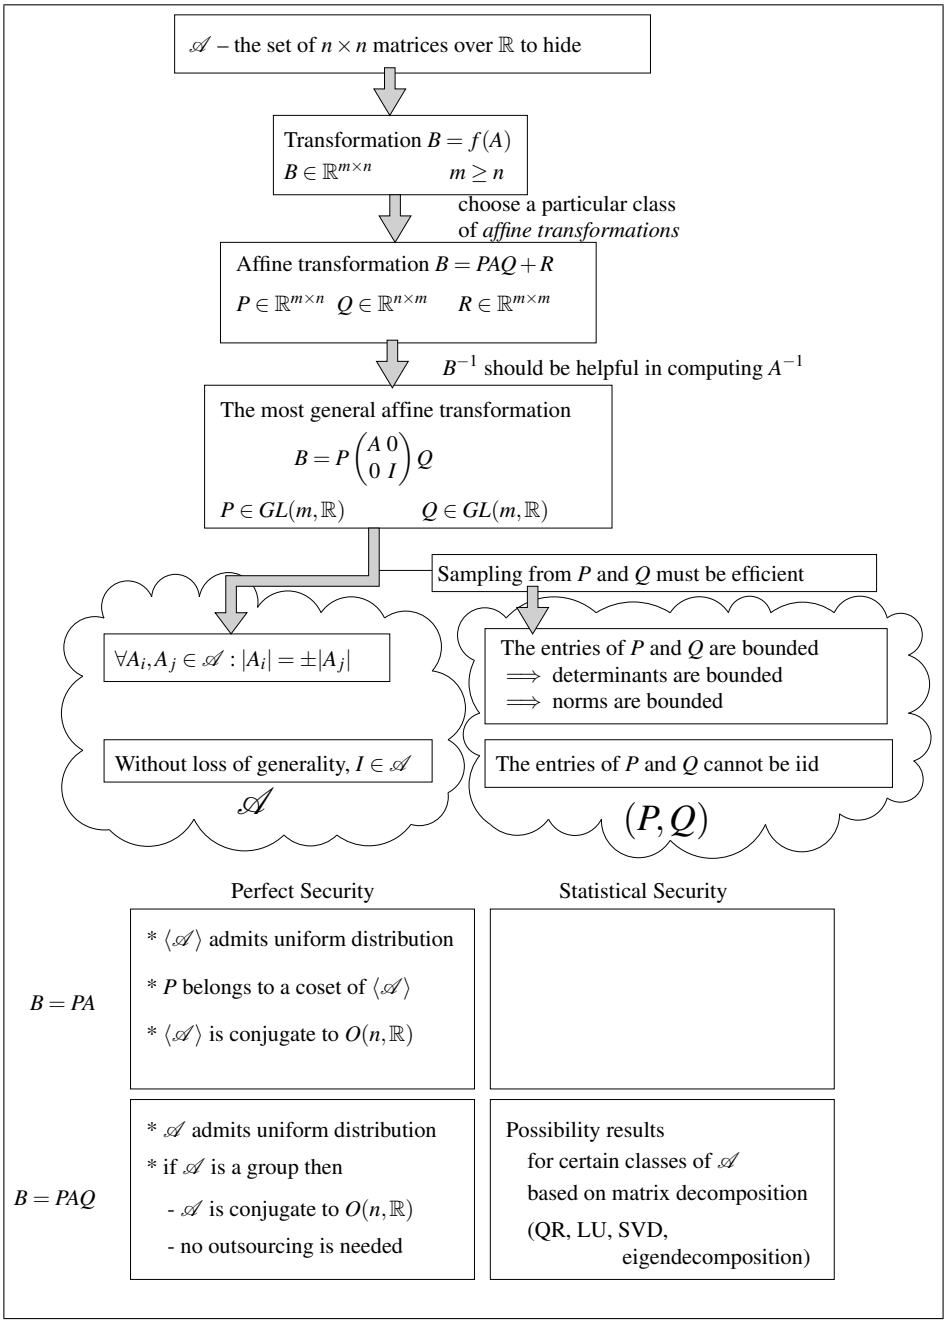

Fig. 1: Possibility and impossibility results for outsourcing the inversion of real matrices in a set $\mathscr{A}$ with help of affine transformations

14. Xinyu Lei, Xiaofeng Liao, Tingwen Huang, Huaqing Li, and Chunqiang Hu. Outsourcing large matrix inversion computation to a public cloud. *Cloud Computing, IEEE Transactions on*, 1(1):1–1, Jan 2013.
15. Wenliang Du. *A Study Of Several Specific Secure Two-Party Computation Problems*. PhD thesis, Purdue University, 2001.
16. Jaideep Vaidya. Privacy-preserving linear programming. In Sung Y. Shin and Sascha Ossowski, editors, *SAC*, pages 2002–2007. ACM, 2009.
17. Olvi L. Mangasarian. Privacy-preserving linear programming. *Optimization Letters*, 5(1):165–172, 2011.
18. Olvi L. Mangasarian. Privacy-preserving horizontally partitioned linear programs. *Optimization Letters*, 6(3):431–436, 2012.
19. Yuan Hong, Jaideep Vaidya, and Haibing Lu. Secure and efficient distributed linear programming. *Journal of Computer Security*, 20(5):583–634, 2012.
20. Alice Bednarz, Nigel Bean, and Matthew Roughan. Hiccups on the road to privacy-preserving linear programming. In *Proceedings of the 8th ACM workshop on Privacy in the electronic society*, WPES '09, pages 117–120, New York, NY, USA, 2009. ACM.
21. Wei Li, Haohao Li, and Chongyang Deng. Privacy-preserving horizontally partitioned linear programs with inequality constraints. *Optimization Letters*, 7(1):137–144, 2013.
22. Yuan Hong and Jaideep Vaidya. An inference-proof approach to privacy-preserving horizontally partitioned linear programs. *Optimization Letters*, 2013. To appear. Published online 05 October 2012.
23. Peeter Laud and Alisa Pankova. New Attacks against Transformation-Based Privacy-Preserving Linear Programming. In Rafael Accorsi and Silvio Ranise, editors, *Security and Trust Management (STM) 2013, 9th International Workshop*, volume 8203 of *Lecture Notes in Computer Science*, pages 17–32. Springer, 2013.
24. Alice Bednarz. *Methods for two-party privacy-preserving linear programming*. PhD thesis, University of Adelaide, 2012.
25. Wenliang Du and Zhijun Zhan. A practical approach to solve secure multi-party computation problems. In *New Security Paradigms Workshop*, pages 127–135. ACM Press, 2002.
26. Dan Bogdanov, Liina Kamm, Sven Laur, and Ville Sokk. Rmind: a tool for cryptographically secure statistical analysis. *IACR Cryptology ePrint Archive*, 2014:512, 2014.
27. Peeter Laud and Alisa Pankova. On the (Im)possibility of Privately Outsourcing Linear Programming. In Ari Juels and Bryan Parno, editors, *Proceedings of the 2013 ACM Workshop on Cloud computing security, CCSW 2013*, pages 55–64. ACM, 2013.
28. Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM Conference on Computer and Communications Security*, pages 784–796. ACM, 2012.
29. W. Hundsdorfer and J.G. Verwer. *Numerical Solution of Time-Dependent Advection-Diffusion-Reaction Equations*. Springer Series in Computational Mathematics. Springer, 2003.
30. A. Szczepański. *Geometry of Crystallographic Groups*. Algebra and discrete mathematics. World Scientific, 2012.
31. Tullio Ceccherini-Silberstein and Michel Coornaert. Amenable groups. In *Cellular Automata and Groups*, Springer Monographs in Mathematics, pages 77–110. Springer Berlin Heidelberg, 2010.
32. N. Monod and N. Ozawa. The dixmier problem, lamplighters and burnside groups. *ArXiv e-prints*, February 2009.
33. Alan Genz. Methods for generating random orthogonal matrices.
34. J.S. Rose. *A Course on Group Theory*. Dover Publications, 1978.
35. R.A. Horn and C.R. Johnson. *Matrix Analysis*. Cambridge University Press, 2012.