εKTELO: A Framework for Defining Differentially Private Computations

DAN ZHANG and RYAN MCKENNA, University of Massachusetts, Amherst
IOS KOTSOGIANNIS, Duke University
GEORGE BISSIAS, University of Massachusetts, Amherst
MICHAEL HAY, Colgate University
ASHWIN MACHANAVAJJHALA, Duke University
GEROME MIKLAU, University of Massachusetts, Amherst

The adoption of differential privacy is growing, but the complexity of designing private, efficient, and accurate algorithms is still high. We propose a novel programming framework and system, εKTELO, for implementing both existing and new privacy algorithms. For the task of answering linear counting queries, we show that nearly all existing algorithms can be composed from operators, each conforming to one of a small number of operator classes. While past programming frameworks have helped to ensure the privacy of programs, the novelty of our framework is its significant support for authoring accurate and efficient (as well as private) programs.

After describing the design and architecture of the εKTELO system, we show that εKTELO is expressive, allows for safer implementations through code reuse, and allows both privacy novices and experts to easily design algorithms. We provide a number of novel implementation techniques to support the generality and scalability of εKTELO operators. These include methods to automatically compute lossless reductions of the data representation, implicit matrices that avoid materialized state but still support computations, and iterative inference implementations that generalize techniques from the privacy literature.

We demonstrate the utility of εKTELO by designing several new state-of-the-art algorithms, most of which result from simple re-combinations of operators defined in the framework. We study the accuracy and scalability of εKTELO plans in a thorough empirical evaluation.

CCS Concepts: • **Security and privacy** → **Privacy-preserving protocols**; **Database and storage security**; • **Theory of computation** → **Theory of database privacy and security**;

Additional Key Words and Phrases: Differential privacy, programming framework, privacy-preserving query answering

This work was supported by the National Science Foundation under grants 1253327, 1408982, 1409125, 1443014, 1421325, and 1409143; and by DARPA and SPAWAR under contract N66001-15-C-4067. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes not withstanding any copyright notation thereon. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

Authors' addresses: D. Zhang, R. McKenna, G. Bissias, and G. Miklau, College of Information and Computer Sciences, University of Massachusetts, Amherst, MA 01003; emails: {dzhang, rmckenna, gbiss, miklau}@cs.umass.edu; I. Kotsogiannis and A. Machanavajjhala, Department of Computer Science, Duke University, Durham, NC 27708; emails: {iosk, ashwin}@cs.duke.edu; M. Hay, Computer Science Department, Colgate University, Hamilton, NY 13346; email: mhay@colgate.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

0362-5915/2020/02-ART2 \$15.00

<https://doi.org/10.1145/3362032>

ACM Reference format:

Dan Zhang, Ryan Mckenna, Ios Kotsogiannis, George Bissias, Michael Hay, Ashwin Machanavajjhala, and Jerome Miklau. 2020. ϵ KTELO: A Framework for Defining Differentially Private Computations. *ACM Trans. Database Syst.* 45, 1, Article 2 (February 2020), 44 pages.
<https://doi.org/10.1145/3362032>

1 INTRODUCTION

As the collection of personal data has increased, many institutions face an urgent need for reliable privacy protection mechanisms. They must balance the need to protect individuals with demands to use collected data for new applications, to model their users' behavior, or share data with external partners. Differential privacy [13, 14] is a rigorous privacy definition that offers a persuasive assurance to individuals, provable guarantees, and the ability to analyze the impact of combined releases of data. Informally, an algorithm satisfies differential privacy if its output does not change too much when any one record in the input database is added or removed.

The research community has actively investigated differential privacy and algorithms are known for a variety of tasks ranging from data exploration to query answering to machine learning. However, the adoption of differentially private techniques in real-world applications remains rare. This is because implementing programs that provably satisfy privacy and ensure sufficient utility for a given task is still extremely challenging for non-experts in differential privacy. In fact, the few real world deployments of differential privacy—like OnTheMap [2, 20] (a U.S. Census Bureau data product), RAPPOR [17] (a Google Chrome extension), and Apple's private collection of emoji's and HealthKit data—have required teams of privacy experts to ensure that implementations meet the privacy standard and that they deliver acceptable utility. There are at least three important challenges in implementing and deploying differentially private algorithms.

The first and foremost challenge is the difficulty of designing utility-optimal algorithms: i.e., algorithms that can extract the maximal accuracy given a fixed "privacy budget." While there are a number of general-purpose differentially private algorithms, such as the Laplace Mechanism [13], they typically offer suboptimal accuracy if applied directly. A carefully designed algorithm can improve on general-purpose methods by an order of magnitude or more—without weakening privacy: that is, accuracy is improved by careful engineering and sophisticated algorithm design.

One might hope for a single dominant algorithm for each task, but a recent empirical study [22] showed that the accuracy of existing algorithms is complex: No single algorithm delivers the best accuracy across the range of settings in which it may be deployed. The choice of the best algorithm may depend on the particular task, the available privacy budget, and properties of the input data. Therefore, to achieve state-of-the-art accuracy, a practitioner currently has to make a host of complex algorithm choices, which may include choosing a low-level representation for the input data, translating their queries into that representation, choosing among available algorithms, and setting parameters. The best choices will vary for different input data and different analysis tasks.

The second challenge is that the tasks in which practitioners are interested are diverse and may differ from those considered in the literature. Hence, existing algorithms need to be adapted to new application settings, a non-trivial task. For instance, techniques used by modern privacy algorithms include optimizing error over multiple queries by identifying common sub-expressions, obtaining noisy counts from the data at different resolutions and using complex inference techniques to reconstruct answers to target queries from noisy, inconsistent, and incomplete measurement queries. But different algorithms use different specialized operators for these sub-tasks, and it can be challenging to adapt them to new situations. Thus, designing utility-optimal algorithms requires significant expertise in a complex and rapidly evolving research literature.

A third equally important challenge is that correctly implementing differentially private algorithms can be difficult. There are known examples of algorithm pseudocode in research papers not satisfying differential privacy as claimed. For instance, Zhang et al. [56] showed that many variants of a primitive called the sparse vector technique do not in fact meet their claims of differential privacy. Differential privacy can also be broken through incorrect implementations of valid algorithms. For example, Mironov [36] showed that standard implementations of basic algorithms like the Laplace Mechanism [13] can violate differential privacy because of their use of floating point arithmetic. Privacy-oriented programming frameworks such as PINQ [15, 35, 42], Fuzz [19], PrivInfer [9], and LightDP [53] help implement programs whose privacy can be verified with relatively little human intervention. While they help to ensure the privacy criterion is met, they may impose their own restrictions and offer little or no support for designing utility-optimal programs. In fact, in PINQ [35], some state-of-the-art algorithms involving inference and domain reduction cannot be implemented.

To address the aforementioned challenges, we have developed ϵ KTELO, a programming framework and system that aids programmers in developing differentially private programs with high utility. ϵ KTELO programs can be used to solve a core class of statistical tasks that involve answering counting queries over a table of arbitrary dimension (described in Section 3). Tasks supported by ϵ KTELO include releasing contingency tables, multi-dimensional histograms, answering OLAP and range queries, and implementing private machine learning algorithms. ϵ KTELO is an open source system under active development; the source code is publicly available at github.com/ektelo/ektelo [6].

This article makes the following contributions.

First, we recognize that, for the tasks we consider, virtually all algorithms in the research literature can be described as combinations of a small number of operators that perform basic functions. Our first contribution is to abstract and unify key subroutines into a small set of operator classes in ϵ KTELO—transformations, query selection, partition selection, measurement, and inference. Different algorithms differ in (i) the sequence in which these operations are performed on the data and (ii) the specific implementation of operations from these classes. In our system, differentially private programs are described as *plans* over a high-level library of *operator* implementations supported by ϵ KTELO. Plans described in ϵ KTELO are expressive enough to reimplement all state-of-the-art algorithms from DPBench [22].

Second, if operator implementations are vetted and shown to satisfy differential privacy, then plans implemented in ϵ KTELO come with a proof of privacy. This proof requires a non-trivial extension of a formal analysis of a past framework [15]. This relieves the algorithm designer of the burden of proving their programs are private. By isolating privacy critical functions in operators, ϵ KTELO reduces the amount of code that needs to be verified for privacy. In future work, we hope to implement operators in ϵ KTELO using programming frameworks like LightDP to eliminate this burden too.

Third, we describe a number of novel implementation techniques that support the generality and efficiency of ϵ KTELO. We design sophisticated matrix support into ϵ KTELO, which allows plan authors to represent and operate on matrix objects that would be infeasible to represent otherwise. We describe a general-purpose, efficient, and scalable inference engine that subsumes customized inference subroutines from the literature, and we describe a new dimensionality reduction operator that is applicable to plans that answer a workload of linear counting queries and can reduce error by at most 3 $\times$ and runtime at most 5 $\times$. These systems innovations are central to the goals of ϵ KTELO, because they give plan authors more freedom to describe plans without being limited by efficiency concerns or bound by the need to design custom inference techniques.

The operator-based approach to implementing differentially private programs has the following benefits:

- *Modularity*: ϵ KTELO enables code reuse, since the same operator can be used in multiple algorithms. This helps safety, as there is less code to verify the correctness of an implementation, and it amplifies innovation, as any improvement to an operator is inherited by all plans containing it.
- *Transparency*: By expressing algorithms as plans with operators from operator classes, differences/similarities of competing algorithms can be discerned. Moreover, algorithm modifications easier to explore. Further, it is possible to identify general rules for restructuring plans (like heuristics in query optimizers).
- *Flexibility*: Practitioners can now use existing operators from different algorithms and recombine them in arbitrary ways—allowing them to invent new algorithms that borrow ideas from the state of the art—without the need for a custom privacy analysis.

We demonstrate the benefits of ϵ KTELO by re-implementing a wide range of algorithms from the literature, using ϵ KTELO to design original algorithms with improved error and efficiency, and by using ϵ KTELO to address two case studies. For a use-case of releasing Census data tabulations, we define a new algorithm that offers a $10\times$ improvement over the best competitor from the literature. For building a private classifier, we used ϵ KTELO to design algorithms that beat all available baselines.

Organization. We provide an overview of ϵ KTELO and highlight its design in the next section. After providing background in Section 3, we describe the execution framework and methods for privacy enforcement Section 4. The operator classes are described in Section 5, and we show the expressiveness of ϵ KTELO plans in Section 6 by re-implementing existing algorithms. Efficient matrix support is described in Section 7, including improvements to inference. Section 8 describes our method for selecting and optimal partition given a workload. We then put ϵ KTELO into action by designing new algorithms for cases studies Section 9, followed by a thorough experimental evaluation in Section 10. We discuss related work and limitations in Sections 11 and 12 and conclude in Section 13.

2 OVERVIEW AND DESIGN PRINCIPLES

In this section, we provide an overview of ϵ KTELO by presenting an example algorithm written in the framework. Then we discuss the principles guiding the design of ϵ KTELO.

2.1 An Example Plan: CDF Estimation

In ϵ KTELO, differentially private algorithms are described using *plans* composed over a rich library of *operators*. Most of the plans described in this article are linear sequences of operators, but ϵ KTELO also supports plans with iteration, recursion, and branching. Operators supported by ϵ KTELO perform a well-defined task and typically capture a key algorithm design idea from the state of the art. Each operator belongs to one of five *operator classes* based on its input–output specification. These are as follows: (a) transformation, (b) query, (c) inference, (d) query selection, and (e) partition selection. Operators are fully described in Section 5 and listed in Figure 1.

First, we describe an example ϵ KTELO plan and use it to introduce the different operator classes. Algorithm 1 shows the pseudocode for a plan authored in ϵ KTELO, which takes as input a table D with schema [Age, Gender, Salary] and returns the differentially private estimate of the empirical cumulative distribution function (CDF) of the Salary attribute for males in their 30s. The plan is fairly sophisticated and works in multiple steps. First the plan uses *transformation* operators on the input table D to filter out records that do not correspond to males in their 30's (Line 2), selecting

ALGORITHM 1: εKTELO CDF Estimator

1: $D \leftarrow \text{PROTECTED}(\text{source_uri})$	▷ Init
2: $D \leftarrow \text{WHERE}(D, \text{sex} == \text{'M'} \text{ AND } \text{age} \in [30, 39])$	▷ Transform
3: $D \leftarrow \text{SELECT}(\text{salary})$	▷ Transform
4: $\mathbf{x} \leftarrow \text{T-VECTORIZE}(D)$	▷ Transform
5: $\mathbf{P} \leftarrow \text{AHPARTITION}(\mathbf{x}, \epsilon/2)$	▷ Partition Select
6: $\bar{\mathbf{x}} \leftarrow \text{V-REDUCEBYPARTITION}(\mathbf{x}, \mathbf{P})$	▷ Transform
7: $\mathbf{M} \leftarrow \text{IDENTITY}(\bar{\mathbf{x}})$	▷ Query Select
8: $\mathbf{y} \leftarrow \text{VECLAPLACE}(\bar{\mathbf{x}}, \mathbf{M}, \epsilon/2)$	▷ Query
9: $\hat{\mathbf{x}} \leftarrow \text{NNLS}(\mathbf{P}, \mathbf{y})$	▷ Inference
10: $\mathbf{W}_{\text{pre}} \leftarrow \text{PREFIX}(\mathbf{x})$	▷ Query Select
11: return $\mathbf{W}_{\text{pre}} \cdot \hat{\mathbf{x}}$	▷ Output

only the salary attribute (Line 3). Then it uses another transformation operator to construct a vector of counts $\mathbf{x}$ that contains one entry for each value of salary. x_i represents the number of rows in the input (in this case males in their 30s) with salary equal to i .

Before adding noise to this histogram, the plan uses a *partition selection* operator, AHPARTITION (Line 5). Operators in this class choose a partition of the data vector that is later used in a transformation. AHPARTITION uses the sensitive data to identify a partition $\mathbf{P}$ of the counts in $\mathbf{x}$ such that counts within a partition group are close. Since AHPARTITION uses the input data, it expends part of the privacy budget (in this case $\epsilon/2$). AHPARTITION is a key subroutine in AHP [57], which was shown to have state-of-the-art performance for histogram estimation [22].

Next the plan uses V-REDUCEBYPARTITION (Line 6), another transformation operator on $\mathbf{x}$, to apply the partition $\mathbf{P}$ computed by AHPARTITION. This results in a new reduced vector $\bar{\mathbf{x}}$ that contains one entry for each partition group in $\mathbf{P}$ and the entry is computed by adding up counts within each group.

The plan now specifies a set of measurement queries $\mathbf{M}$ on $\bar{\mathbf{x}}$ using the IDENTITY *query selection* operator (Line 7). The identity matrix corresponds to querying all the entries in $\bar{\mathbf{x}}$ (since $\mathbf{M}\bar{\mathbf{x}} = \bar{\mathbf{x}}$). Query selection operators do not answer any query, but rather specify which queries should be estimated. (This is analogous to how partition selection operators only select a partition but do not apply it.) Next, VECTOR LAPLACE returns differentially private answers to all the queries in $\mathbf{M}$. It does so by automatically calculating the sensitivity of the vectorized queries—which depends on all upstream data transformations—and then using the standard Laplace mechanism (Line 8) to add noise. This operator consumes the remainder of the privacy budget (again $\epsilon/2$).

So far the plan has computed an estimated histogram of partition group counts $\mathbf{y}$, while our goal is to return the empirical CDF on the original salary domain. Hence, the plan uses the noisy counts on the reduced domain $\mathbf{y}$ to infer non-negative counts in the original vector space of $\mathbf{x}$ by invoking an *inference* operator NNLS (short for non-negative least squares) (Line 9). NNLS($\mathbf{P}, \mathbf{y}$) finds a solution, $\hat{\mathbf{x}}$, to the problem $\mathbf{P}\hat{\mathbf{x}} = \mathbf{y}$, such that all entries of $\hat{\mathbf{x}}$ are non-negative. Last, the plan constructs the set of queries, $\mathbf{W}_{\text{pre}}$, needed to compute the empirical CDF (a lower triangular $n \times n$ matrix representing prefix sums) by calling the query selection operator PREFIX(n) (Line 10) and returns the output (Line 11).

2.2 εKTELO Design Principles

The design of εKTELO is guided by the following principles. With each principle, we include references to future sections of the article where the consequent benefits are demonstrated.

Expressiveness. ϵ KTELO is designed to be expressive, meaning that a wide variety of state-of-the-art algorithms can be written succinctly as ϵ KTELO plans. To ensure expressiveness, we carefully designed a foundational set of operator classes that cover features commonly used by leading differentially private algorithms. We demonstrate the expressiveness of our operators by showing in Section 6 that the algorithms from the recent DPBench benchmark [22] can be readily re-implemented in ϵ KTELO.

Privacy “for free”. ϵ KTELO is designed so that any plan written in ϵ KTELO automatically satisfies differential privacy. The formal statement of this privacy property is in Section 4.3. This means that plan authors are not burdened with writing privacy proofs for each algorithm they write. Furthermore, when invoking privacy-critical operators that take noisy measurements of the data, the magnitude of the noise is automatically calibrated. As described in Section 4, this requires tracking all data transformations and measurements and using this information to handle each new measurement request.

Reduced privacy verification effort. Ensuring that an algorithm implementation satisfies differential privacy requires verifying that it matches the algorithm specification. The design of ϵ KTELO reduces the amount of code that must be vetted each time an algorithm is crafted. First, since an algorithm is expressed as a plan and all plans automatically satisfy differential privacy, the code to be vetted is solely the individual operators. Second, operators need to be vetted only once but may be reused across multiple algorithms. Finally, it is not necessary to vet every operator, but only the privacy-critical ones (as discussed in Section 4, ϵ KTELO mandates a clear distinction between privacy-critical and non-private operators). This means that verifying the privacy of an algorithm requires checking fewer lines of code. In Section 6, we compare the verification effort to vet the DPBench codebase¹ against the effort required to vet these algorithms when expressed as plans in ϵ KTELO.

Transparency. In ϵ KTELO, all algorithms are expressed in the same form: Each is a plan, consisting a sequence of operators where each operator is selected from a class of operators based on common functionality. This facilitates algorithm comparison and makes differences between algorithms more apparent. In Section 6, we summarize the plan signatures of a number of state-of-the-art algorithms (pictured in Figure 2). These plan signatures reveal similarities and common idioms in existing algorithms. These are difficult to discover from the research literature or through code inspection.

Efficiency and Scalability. Many ϵ KTELO plans compute on data vectors formed from projections of an input table. The current implementation of ϵ KTELO relies on storing these vectors in memory on a single machine. Even under this restriction, it is challenging to get all ϵ KTELO operators to run efficiently. Our specialized matrix representation techniques, presented in Section 7, allow many of the key operators to scale to large data vectors without imposing undue restrictions on plan authors.

We believe that ϵ KTELO, by supporting the design principles described above, provides an improved platform for designing and deploying differentially private algorithms.

3 PRELIMINARIES

The input to ϵ KTELO is a database instance of a single-relation schema $T(A_1, A_2, \dots, A_\ell)$. Each attribute A_i is assumed to be discrete (or suitably discretized). A *condition formula*, ϕ , is a Boolean condition that can be evaluated on any tuple of T . We use $\phi(T)$ to denote the number of tuples in

¹ Available at: https://github.com/dpcomp-org/dpcomp_core.

T for which ϕ is true. A number of operators in εKTELO answer linear queries over the table. A linear query is the linear combination of any finite set of condition counts.

Definition 3.1 (Linear Counting Query (Declarative)). A linear query q on T is defined by conditions $\phi_1 \dots \phi_k$ and coefficients $c_1 \dots c_k \in \mathbb{R}$ and returns $q(T) = c_1\phi_1(T) + \dots + c_k\phi_k(T)$.

It is common to consider a vector representation of the database, denoted $\mathbf{x} = [x_1 \dots x_n]$, where x_i is equal to the number of tuples of type i for each possible tuple type in the relational domain of T . The size of this vector, n , is the product of the attribute domains. Then it follows that any linear counting query has an equivalent representation as a vector of n coefficients and can be evaluated by taking a dot product with $\mathbf{x}$. Abusing notation slightly, let $\phi(i) = 1$ if ϕ evaluates to true for the tuple type i and 0 otherwise.

Definition 3.2 (Linear Counting Query (Vector)). For a linear query q defined by $\phi_1 \dots \phi_k$ and $c_1 \dots c_k$, its equivalent vector form is $\mathbf{q} = [q_1 \dots q_n]$, where $q_i = c_1\phi_1(i) + \dots + c_k\phi_k(i)$. The evaluation of the linear query is $\mathbf{q} \cdot \mathbf{x}$, where $\mathbf{x}$ is vector representation of T .

In the sequel, we will use vectorized representations of the data frequently. We refer to the *domain* as the size of $\mathbf{x}$, the vectorized table. This vector is sometimes large and a number of methods for avoiding its materialization are discussed later.

Let T and T' denote two tables of the same schema, and let $T \oplus T' = (T - T') \cup (T' - T)$ denote the symmetric difference between them. We say that T and T' are neighbors if $|T \oplus T'| = 1$.

Definition 3.3 (Differential Privacy [13]). A randomized algorithm $\mathcal{A}$ is ϵ -differentially private if for any two instances T, T' such that $|T \oplus T'| = 1$, and any subset of outputs $S \subseteq \text{Range}(\mathcal{A})$,

$$\Pr[\mathcal{A}(T) \in S] \leq \exp(\epsilon) \times \Pr[\mathcal{A}(T') \in S].$$

Differentially private algorithms can be composed with each other and other algorithms using composition rules, such as sequential and parallel composition [35] and post-processing [14]. Let f be a function on tables that outputs real numbers. The *sensitivity* of the function is defined as $\max_{|T \oplus T'|=1} |f(T) - f(T')|$.

Definition 3.4 (Stability). Let g be a transformation function that takes a data source (table or vector) as input and returns a new data source (of the same type) as output. For any pair of sources S and S' let $|S \oplus S'|$ denote the distance between sources. If the sources are both tables, then this distance is the size of the symmetric difference; if the sources are both vectors, then this distance is the L_1 norm; if the sources are of mixed type, it is undefined. Then the stability of g is $\max_{S, S': |S \oplus S'|=1} |g(S) \oplus g(S')|$. When the stability of g is at most c for some constant c , we say that g is c -stable.

4 EXECUTION FRAMEWORK AND PRIVACY ENFORCEMENT

This section describes the execution environment and then formalizes the claim that any program executed in εKTELO satisfies differential privacy.

4.1 Protected Kernel and Client Space

The execution framework consists of an untrusted client space and a protected kernel that encloses the private data. An εKTELO program, which we call a plan, runs in the unprotected client space. When the plan needs to interact with the private data, it does so through privileged operators that can issue requests to the protected kernel. Such operators may, for example, request that protected kernel apply a data transformation or perhaps return a noisy measurement. The protected kernel services requests from privileged operators, only executing them if their cost is within the available privacy budget. The distinction between the client space and the protected kernel is a fundamental

one in ϵKTELO . It allows authors to write plans that consist of operator calls embedded in otherwise arbitrary code (which may freely include conditionals, loops, recursion, etc.).

The protected kernel is initialized by specifying a single protected data object—an input table T —and a global privacy budget, which we denote as ϵ_{tot} . Note that requests for data transformations may cause the protected kernel to derive additional data sources. Thus, the protected kernel maintains a *data source environment*, which consists of a mapping between data source variables, which are exposed to the client, and the protected data objects, which are kept private. In addition, the data source environment tracks the transformation lineage of each data source. It also maintains the stability of each transformation (defined in Section 3). Note that in describing operators (Section 5), we speak informally of operators having data sources as inputs and outputs rather than data source *variables*. A layer of indirection is always maintained in the implementation but sometimes elided in our descriptions to simplify the presentation.

4.2 Operator Types

Operators have one of three types, based on their interaction with the protected kernel. The first type is a **Private** operator, which requests that the protected kernel perform some action on the private data (e.g., a transformation) but receives only an acknowledgement that the operation has been performed. The second type is a **Private→Public** operator, which receives information about the private data (e.g., a measurement) and thus consumes privacy budget. The last type is a **Public** operator, which does not interact with the protected kernel at all and can be executed entirely in client space. An example of a public operator would operators that perform inference on the noisy measurements received from the protected kernel. When describing operators in Section 5, we color code them based on their type.

4.3 Privacy Guarantee

In this section, we state the privacy guarantee offered by ϵKTELO . Informally, ϵKTELO ensures that if the protected kernel is initialized with a source database T and a privacy budget ϵ_{tot} , then any plan (chosen by the client) will satisfy ϵ_{tot} -differential privacy with respect to T . Note that if the client exhausts the privacy budget, subsequent calls to **Private→Public** operators will return an exception, indicating that they are not permitted. Importantly, an exception itself does not leak sensitive information—i.e., the decision to return an exception does not depend on the private state.

A *transcript* is a sequence of operator calls and their responses. Formally, let $r_k = \langle op_1, a_1, \dots, op_k, a_k \rangle$ denote a length k sequence where op_i is an operator call and a_i the response. We assume that the value of op_i is a deterministic function of $a_1, \dots, a_{i-1}$. We use $R_k = r_k$ to denote the event that the first k operations result in transcript r_k . Let $\mathcal{R}_k$ be the set of all possible transcripts of length k . We assume that all **Private→Public** operators output values from an arbitrary but finite set. Thus, the set of possible transcripts is finite. Let $P(R_k = r_k \mid \mathbf{Init}(T, \epsilon_{tot}))$ be the conditional probability of event $R_k = r_k$ given that the system was initialized with input T and a privacy budget of ϵ_{tot} .

THEOREM 4.1 (PRIVACY OF ϵKTELO PLANS). *Let T, T' be any two instances such that $|T \oplus T'| = 1$. For all $k \in \mathbb{N}^+$ and $r_k \in \mathcal{R}_k$,*

$$P(R_k = r_k \mid \mathbf{Init}(T, \epsilon_{tot})) \leq \exp(\epsilon_{tot}) \times P(R_k = r_k \mid \mathbf{Init}(T', \epsilon_{tot})).$$

The proof of Theorem 4.1, which appears in the sequel, extends the proof in Reference [15] to support the V-SPLITBYPARTITION operator.

While ϵKTELO ensures differential privacy, private information could be leaked via side-channel attacks (e.g., timing attacks). Privacy engineers who design operators are responsible for protecting against such attacks; an analysis of this issue is beyond the scope of this article.

4.4 Privacy Proof

This section presents a proof of Theorem 4.1. We start by introducing some supporting concepts and notation. (Some notation is adapted from Reference [15].)

Information tracked by the protected kernel. The protected kernel maintains the following state, which we denote as S_{kernel} :

- A set of source variables SV .
- A data source environment E maps each source variable $sv \in SV$ to an actual data source S , as in $E(sv) = S$. (Recall that sources can be tables or vectors.)
- A transformation graph: the nodes are SV and there is an edge from sv to sv' if sv' was derived via transformation from sv . (Note: a partition transformation introduces a special dummy data source variable whose parent is the source variable being partitioned and whose children are the variables associated with each partition.)
- A stability tracker St maps each source variable $sv \in SV$ to a non-negative number: $St(sv)$ represents the stability (Definition 3.4) of the transformation that derived data source sv from the initial source, or 1 if sv is the initial source.
- A budget consumption tracker B that maps each source variable $sv \in SV$ to a non-negative number: $B(sv)$ represents the total budget consumption made by queries to sv or to any source derived from sv .
- A query history Q that maps each source variable to information about the state of queries asked about sv or any of its descendants. Specifically, for $sv \in SV$, $Q(sv)$ returns of a set of tuples (q, s, σ, v) where the meaning of the tuple is that query q was executed on data source s (which is sv or one of its descendants) with σ noise, the result was v . In the context of the proof a query is any **Private**→**Public** operator. Such an operator is assumed to satisfy ϵ -differential privacy with respect to the data source on which it is applied.
- The global privacy budget, denoted ϵ_{tot} .

When the protected kernel is initialized, as in **Init**(T, ϵ_{tot}), it sets global budget to ϵ_{tot} , creates new source variable sv_{root} , sets $E(sv_{root}) = T$, sets $St(sv_{root}) = 1$, and $B(sv_{root}) = 0$, and adds sv_{root} to the transformation graph.

Budget Management. When a query request is issued to the protected kernel, the protected kernel uses Algorithm 2 to check whether the query can be answered given the available privacy budget.

Configurations. A configuration, denoted $\mathbb{C} = \langle S_{client}, S_{kernel} \rangle$, captures the state of the client, denoted S_{client} , and the state of the protected kernel, denoted S_{kernel} . The client state can be arbitrary, but state updates are assumed to be deterministic.

We can define the similarity of two configurations $\mathbb{C}$ and $\mathbb{C}'$ as follows. (Notation: We use X' to refer to component X of configuration $\mathbb{C}'$.) We say that $\mathbb{C} \sim \mathbb{C}'$ iff $S_{client} = S'_{client}$ and $S_{kernel} \sim S'_{kernel}$, where $S_{kernel} \sim S'_{kernel}$ iff $SV = SV'$ and the transformation graphs are identical and for each $sv \in SV$ the following conditions hold:

- $St(sv) = St'(sv)$, $B(sv) = B'(sv)$, $Q(sv) = Q'(sv)$, and $\epsilon_{tot} = \epsilon'_{tot}$.
- $|E(sv) \oplus E'(sv)| \leq St(sv) = St'(sv)$ where $|x \oplus y|$ is measured as symmetric difference when the sources x and y are tables and L_1 distance for vectors; see Definition 3.4.)

We introduce a lemma that bounds the difference probability between query answers. Let $P(q(E(s), \sigma) = v)$ denote the probability that query operator q when applied to data source $E(s)$ with noise σ returns answer v .

ALGORITHM 2: An algorithm for budget requests

```

1: procedure REQUEST( $sv, \sigma$ )
2:   if  $sv$  is the root then
3:     If  $B(sv) + \sigma > \epsilon_{tot}$ , return FALSE. Otherwise  $B(sv) += \sigma$  and return TRUE.
4:   else if  $sv$  is a partition variable then
5:     Let  $sv_{child}$  be the child from which the request came.
6:     Let  $r = \max\{B(sv_{child}) + \sigma - B(sv), 0\}$ 
7:     Let  $ans = \text{REQUEST}(\text{parent}(sv), r)$ .
8:     If  $ans = \text{FALSE}$ , return FALSE. Otherwise,  $B(sv) += r$  and return TRUE.
9:   else
10:     $ans = \text{REQUEST}(\text{parent}(sv), s \cdot \sigma)$   $\triangleright s$  is stability factor of  $sv$  wrt its parent
11:    if  $ans = \text{FALSE}$ , return FALSE.
12:     $B(sv) += \sigma$ . Return TRUE.
13:  end if
14: end procedure

```

LEMMA 4.2. Let $\mathbb{C} \sim \mathbb{C}'$. For any $sv \in SV$ with non-empty $Q(sv)$, the following holds:

$$\begin{aligned}
 & \prod_{(q, s, \sigma, v) \in Q(sv)} P(q(E(s), \sigma) = v) \\
 & \leq \exp(B(sv) \times |E(sv) \oplus E'(sv)|) \times \prod_{(q, s, \sigma, v) \in Q'(sv)} P(q(E'(s), \sigma) = v).
 \end{aligned} \tag{1}$$

PROOF. Proof by induction on a reverse topological order of the transformation graph.

Base case: Consider a single sv at the end of the topological order (therefore it has no children). If $Q(sv)$ is empty, then it holds trivially. Assume non-empty. Consider any $(q, s, \sigma, v) \in Q(sv)$. Since sv has no children, then $s = sv$. Furthermore, because the only budget requests that apply to sv are from direct queries, we have (according to Algorithm 2), $B(sv) = \sum_{(q, s, \sigma, v) \in Q(sv)} \sigma$. Since we assume that any query operator satisfies ϵ -differential privacy with respect to its source input, we have $P(q(E(s), \sigma) = v) \leq P(q(E'(s), \sigma) = v) \times \exp(\sigma \times |E(s) \oplus E'(s)|)$. Substituting sv for s and taking the product over all terms in $Q(sv)$, we get Equation (1).

Inductive case: Assume Equation (1) holds for all nodes later in the topological order. Therefore, it holds for any child c of sv . We can combine the inequalities for each child into the following inequality over all children:

$$\begin{aligned}
 & \prod_{c \in \text{children}(sv)} \prod_{(q, s, \sigma, v) \in Q(c)} P(q(E(s), \sigma) = v) \\
 & \leq \prod_{c \in \text{children}(sv)} \exp(B(c) \times |E(c) \oplus E'(c)|) \times \prod_{(q, s, \sigma, v) \in Q(c)} P(q(E'(s), \sigma) = v) \\
 & = \exp\left(\sum_{c \in \text{children}(sv)} B(c) \times |E(c) \oplus E'(c)|\right) \\
 & \quad \times \prod_{c \in \text{children}(sv)} \prod_{(q, s, \sigma, v) \in Q(c)} P(q(E'(s), \sigma) = v).
 \end{aligned}$$

There are two cases, depending what type of table variable sv is.

First, consider the case when sv is *not* a special partition variable. We know by transformation stability that $|E(c) \oplus E'(c)| \leq s \times |E(sv) \oplus E'(sv)|$, where s is the stability factor for the transformation. In addition, $\sum_c B(c) \times s \leq B(sv)$, because, according to Algorithm 2, every time a request of σ is made to child c , a request of $s \times \sigma$ is made to sv . Therefore,

$$\begin{aligned} \sum_{c \in \text{children}(sv)} B(c) \times |E(c) \oplus E'(c)| &\leq \sum_{c \in \text{children}(sv)} B(c) \times s \times |E(sv) \oplus E'(sv)| \\ &\leq B(sv) \times |E(sv) \oplus E'(sv)|. \end{aligned}$$

Furthermore, observe that each term in $(q, s, \sigma, v) \in Q(c)$ also appears in $Q(sv)$. In addition, $Q(sv)$ includes any queries on sv directly (and we know from an argument similar to the base case that Equation (1) holds for these queries). Therefore, Equation (1) holds on sv .

Now, consider the case where sv is the special partition variable. Let $m = \max_c B(c)$. We get the following:

$$\begin{aligned} \sum_{c \in \text{children}(sv)} B(c) \times |E(c) \oplus E'(c)| &\leq \sum_{c \in \text{children}(sv)} m \times |E(c) \oplus E'(c)| \\ &= m \times \sum_{c \in \text{children}(sv)} |E(c) \oplus E'(c)| = m \times |E(sv) \oplus E'(sv)| \\ &= B(sv) \times |E(sv) \oplus E'(sv)|. \end{aligned}$$

The second-to-last line follows from the fact that sv is partition transformation. The last line follows from how $B(sv)$ is updated according Algorithm 2. $\square$

Main Proof. We use $\mathbb{C}_0(T, \epsilon_{tot}, P_0)$ to denote the initial configuration in which the protected kernel has been initialized with **Init**(T, ϵ_{tot}) and the client state is initialized to P_0 . We use the notation $\mathbb{C}_0(T, \epsilon_{tot}, P_0) \xrightarrow[t]{p} \mathbb{C}$ to mean that starting in $\mathbb{C}_0$ after t operations, the probability of being in configuration $\mathbb{C}$ is p .

THEOREM 4.3. *If $T \sim_1 T'$ and $\mathbb{C}_0(T, \epsilon_{tot}, P_0) \xrightarrow[t]{p} \mathbb{C}$ such that $B(sv_{root}) = \epsilon$ in $\mathbb{C}$, then $\epsilon \leq \epsilon_{tot}$ and there exists $\mathbb{C}'$ such that $\mathbb{C}_0(T', \epsilon_{tot}, P_0) \xrightarrow[t]{q} \mathbb{C}'$, where $\mathbb{C} \sim \mathbb{C}'$ and $p \leq q \cdot \exp(\epsilon)$.*

Theorem 4.1 follows as a corollary from Theorem 4.3.

PROOF. Proof by induction on t .

Base case: $t = 0$. This implies that $p = q = 1$, $\epsilon = 0$, and $\mathbb{C} = \mathbb{C}_0(T, \epsilon_{tot}, P_0)$ and $\mathbb{C}' = \mathbb{C}_0(T', \epsilon_{tot}, P_0)$. It follows that $\mathbb{C} \sim \mathbb{C}'$, because we are given that $T \sim_1 T'$ and the rest of the claim follows.

Inductive case: Assume the claim holds for t , we will show it holds for $t + 1$. Let $\mathbb{C}_1$ be any configuration such that $\mathbb{C}_0(T, \epsilon_{tot}, P_0) \xrightarrow[t]{p_1} \mathbb{C}_1$ where in $\mathbb{C}_1$, we have $B(sv_{root}) = \epsilon_1$.

The inductive hypothesis tells us that $\epsilon_1 \leq \epsilon_{tot}$ and that there exists a $\mathbb{C}'_1$ such that $\mathbb{C}_0(T', \epsilon_{tot}, P_0) \xrightarrow[t]{q_1} \mathbb{C}'_1$ and $\mathbb{C}_1 \sim \mathbb{C}'_1$ and $p_1 \leq q_1 \times \exp(\epsilon_1)$.

Because $\mathbb{C}_1 \sim \mathbb{C}'_1$, it follows that the client is in the same state and so the next operation request from the client will be the same in $\mathbb{C}_1$ and $\mathbb{C}'_1$. The proof requires a case analysis based on the nature of the operator. We omit analysis of transformation operators or operators that are purely on the client side as those cases are straightforward: Essentially, we must show that the appropriate bookkeeping is performed by the protected kernel. We focus on the case where the operator is a query operator.

Transform		Partition selection		Query selection	
TV	T-Vectorize	PA	AHPpartition	SI	Identity
TW	T-Where	PG	Grid	ST	Total
TPR	T-Project	PD	Dawa	SP	Privelet
TP	V-SplitByPartition	PW	Workload-based	SH2	H2
TR	V-ReduceByPartition	PS	Stripe(attr)	SHB	HB
Inference		PM	Marginal(attr)	SG	Greedy-H
		PU	UniformPartition	SU	UniformGrid
		Query		SA	AdaptiveGrids
				SQ	Quadtree
LS	Least squares	LM	Vector Laplace	SHD	HDMM
NLS	Nneg Least squares			SS	Stripe(attr)
MW	Mult Weights			SW	Worst-approx
HR	Thresholding			SPB	PrivBayes select

Fig. 1. The operators currently implemented in ϵ KTELO. Private operators are red, Private→Public operators are orange, and Public operators are green.

For a query operator, there are two cases: (a) running out of budget and (b) executing a query. For the first case, by the inductive hypothesis $\mathbb{C}_1 \sim \mathbb{C}'_1$ and therefore if executing Algorithm 2 yields False on the protected kernel state in $\mathbb{C}_1$, it will also do so on the protected kernel state in $\mathbb{C}'_1$. For the second case, suppose query q is executed on source sv with noise σ and answer v is obtained. The protected kernel adds the corresponding entry to the query history Q . Let $\mathbb{C}$ denote the resulting state. Let $\mathbb{C}'$ correspond to extending $\mathbb{C}'_1$ in a similar way. Thus $\mathbb{C} \sim \mathbb{C}'$.

It remains to show two things. First, letting $B(sv_{root}) = \epsilon$, we must show that $\epsilon \leq \epsilon_{tot}$. This follows from Algorithm 2, which does not permit $B(sv_{root})$ to exceed ϵ_{tot} . Second, we must bound the probabilities. Suppose that the probability of this query answer in $\mathbb{C}$ is p_2 and the probability of this answer on $\mathbb{C}'$ is q_2 . It remains to show that $p_1 \cdot p_2 \leq \exp(\epsilon) \cdot q_1 \cdot q_2$. For this we rely on Lemma 4.2 applied to sv_{root} with the observations that the product of probabilities bounded in Lemma 4.2 corresponds to the probabilities in $p_1 \cdot p_2$ that do not trivially equal 1 and that $|E(sv_{root}) \oplus E'(sv_{root})| = 1$. $\square$

5 OPERATORS AND OPERATOR CLASSES

We now describe in detail the operators and operator classes in ϵ KTELO. A full list of operators is shown in Figure 1 where they are arranged into classes and color-coded by type (Private, Private→Public, or Public). Along with descriptions of the operator classes we explain their role in plans and prove supporting properties.

5.1 Transformation Operators

Transformation operators take as input a data source variable (either a table or a vector) and output a transformed data source (again, either a table or vector). Transformation operators modify the data held in the kernel, without returning answers. So while they do not expend privacy budget, they can affect the privacy analysis through their *stability* (Section 3). Every transformation in ϵ KTELO has a well-established stability.

Table Transformations. ϵ KTELO supports table transformations PROJECT and WHERE, each with stability 1. The definitions of the operators are nearly identical to those described in PINQ [35] and are not repeated here. As ϵ KTELO currently handles programs that use linear queries on single tables, the JOIN operator is not yet supported.

Vectorization. All of the plans in εKTELO start with table transformations and typically transform the resulting table into a vector using T-VECTORIZE (and all later operations happen on vectors). The T-VECTORIZE operator is a transformation operator that takes as input a table T and outputs a vector $\mathbf{x}$ that has as many cells as the number of elements in the table's domain (recall the discussion of domain Section 3). Each cell in $\mathbf{x}$ represents the number of records in the table that correspond to the domain element encoded by the cell. T-VECTORIZE is a 1-stable transformation.

The vectorize operation can significantly impact the performance of the code, especially in high-dimensional cases, as we represent one cell per element in the domain. For this reason we allow table transformations (WHERE and PROJECT) to reduce the domain size before running T-VECTORIZE. One of the primary reasons for working with the vector representation is to allow for inference operators downstream. Once in vector form, data can be further transformed as described next.

Vector Transformations. εKTELO supports transformations on vector data sources. Each vector transformation takes as input a vector $\mathbf{x}$ and a matrix $\mathbf{M}$ and produces a vector $\mathbf{x}' = \mathbf{M}\mathbf{x}$. The linearity of vector transformations is an important feature that is leveraged by downstream inference operators. The stability of vector transformations is equal to the largest L_1 column norm of $\mathbf{M}$.

The V-REDUCEBYPARTITION operator is a 1-stable vector transformation operator that reduces the dimensionality of the data vector $\mathbf{x}$ by eliminating cells from $\mathbf{x}$ or grouping together cells in $\mathbf{x}$. Such transformations are useful to (a) filter out parts of the domain that are uninteresting for the analyst, (b) reduce the size of the $\mathbf{x}$ vector so that algorithm performance can be improved, and (c) reduce the number of cells in $\mathbf{x}$ so that the amount of noise added by measurement operators is reduced.

V-REDUCEBYPARTITION takes as input a partition defining a grouping of the cells in the $\mathbf{x}$. It can be carried out by representing the partition as a $(p \times n)$ matrix $\mathbf{P}$ where n is the number of cells in $\mathbf{x}$, p is the number of groups in the partition, and $P_{ij} = 1$ if cell j in $\mathbf{x}$ is mapped to group i and 0 otherwise.

The V-SPLITBYPARTITION operator is the vector analogue of the tabular SPLITBYPARTITION operator. It takes as input a partition and splits the data vector $\mathbf{x}$ into k vectors, $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(k)}$, each representing a disjoint subset of the original domain. This operator allows us to create different subplans for disjoint parts of the domain. This is a 1-stable vector transform. (Note: V-SPLITBYPARTITION can be expressed as k linear transforms with matrices that select the appropriate elements of the domain for each partition.)

5.2 Query Operators

Query operators are responsible for computing noisy answers to queries on a data source. Since answers are returned, query operators necessarily expend privacy budget. Query operators take a data source variable and ϵ as input.

For tables, the NOISYCOUNT operator takes as input a table D and ϵ and returns $|D| + \eta$, where η is drawn from the Laplace distribution with scale $1/\epsilon$. For vectors, the VECTOR LAPLACE operator takes as input a vector $\mathbf{x}$, epsilon, and a set of linear counting queries $\mathbf{M}$ represented in matrix form. Let $\mathbf{M}$ be a matrix of size $(m \times n)$. VECTOR LAPLACE returns $\mathbf{M}\mathbf{x} + \frac{\sigma(\mathbf{M})}{\epsilon}\mathbf{b}$, where $\mathbf{b}$ is a vector of m independently drawn Laplace random variables with scale 1 and $\sigma(\mathbf{M})$ is the maximum L_1 norm of the columns of $\mathbf{M}$.

For both query operators, it is easy to show they satisfy ϵ -differential privacy with respect to their data source input [32, 35]. Note, however, in the case the source is derived from other data sources through transformation operators, the total privacy loss could be higher. The cumulative privacy loss depends on the stability of the transformations and is tracked by the protected kernel.

5.3 Query Selection Operators

Since each query operation consumes privacy budget, the plan author must be judicious about what queries are being asked. Recent privacy work has shown that if the plan author's goal is to answer a *workload* of queries, then simply asking these queries directly can lead to sub-optimal accuracy (e.g., when workload queries ask about overlapping regions of the domain). Instead, higher accuracy can be achieved by designing a *query strategy*, a collection of queries whose answers can be used to reconstruct answers to the workload. This approach was formalized by the matrix mechanism [32] and has been a key idea in many algorithms [12, 23, 29, 31, 33, 44, 51]. Among these the recent HDMM algorithm is notable, because it uses an optimization-based approach to find the query strategy that most-effectively answers the workload. HDMM effectively scales to multi-dimensional domains and offers state-of-the-art utility on many workloads [33].

A query selection operator is distinguished by its output type: a set of linear counting queries $\mathbf{M}$ represented in matrix form (i.e., the matrix input to the VECTOR LAPLACE operator described above). As Figure 1 indicates, ϵ KTELO supports a large number of query selection operators, most of which are extracted from algorithms proposed in the literature. While these operators agree in terms of their output, they vary in terms of their input: Some employ fixed strategies that depend only on the size of $\mathbf{x}$ (e.g., IDENTITY and PREFIX in Algorithm 1), some adapt to the workload (e.g., GREEDY-H), some depend on prior measurements (e.g., ADAPTIVEGRIDS), and so on.

Most query selection operators only rely on non-private information (domain size, workload) and therefore are of **Public** type. But there are a few that consult the private data and thus have the **Private**→**Public** type. For example, WORST-APPROX is an operator that picks the query from a workload that is the worst approximated by a current estimate of the data. Such an operator is used by iterative algorithms like MWEM [21]. Another is PRIVBAYES SELECT, an operator that privately constructs a Bayes net over the attributes of the data source and then returns a matrix corresponding to the sufficient statistics for fitting the parameters of the Bayes net. This was used as a subroutine in PrivBayes [55].

5.4 Partition Selection Operators

Partition selection operators compute a matrix $\mathbf{P}$ that can serve as the input to the V-REDUCEBYPARTITION and V-SPLITBYPARTITION operators described earlier. Of course, the matrix $\mathbf{P}$ must be appropriately structured to be a valid partition of $\mathbf{x}$.

This is an important operator class, since much of recent innovation into state-of-the-art algorithms for answering histograms and range queries has used partitions to either reduce the domain size of the data vector by grouping together cells with similar counts or split the data vector into smaller vectors and leverage the parallel composition of differential privacy to process each subset of the domain independently. ϵ KTELO includes partition selection operators AHPARTITION and DAWA, which are subroutines from the AHP [57] and DAWA [29] algorithms, respectively. Both of these operators are data adaptive and hence are **Private**→**Public**. We also introduce new partition selection operators, WORKLOAD-BASED and STRIPE, described in Sections 8.1 and 9.2, respectively.

5.5 Inference Operators

An inference operator derives new estimates to queries based on the history of transformations and query answers. Inference operators never use the input data directly and hence are **Public**. Plans typically terminate with a call to an inference operator to estimate a final set of query answers reflecting all available information computed during execution of the plan. Some plans may also perform inference as the plan executes.

Ideally, an inference method should: (i) properly account for measurements with unequal noise; (ii) support inference over incomplete measurements (in which derived answers are not completely

determined by available measurements); (iii) should incorporate all available information (including a prior or constraint on the input dataset); and, last, (iv) inference should efficiently scale to large domains. Many versions of inference have been considered in the literature [7, 21, 23, 28, 30, 41, 44, 50, 57] but none meet all of the objectives above. εKTELO currently supports multiple inference methods, in part to support algorithms from past work and in part to offer necessary tradeoffs among the properties above.

All the inference operators supported in εKTELO take as input a set of queries, represented as a matrix $\mathbf{M}$, and noisy answers to these queries, denoted $\mathbf{y}$. The output of inference is a data vector $\hat{\mathbf{x}}$ that best fits the noisy answers—i.e., an $\hat{\mathbf{x}}$ such that $\mathbf{M}\hat{\mathbf{x}} \approx \mathbf{y}$. The estimated $\hat{\mathbf{x}}$ can then be used to derive an estimate of any linear query $\mathbf{q}$ by computing $\mathbf{q} \cdot \hat{\mathbf{x}}$. The inference operator may optionally take as input a set of weights, one per query (row) in $\mathbf{M}$ to account for queries with different noise scales.

εKTELO supports two variants of least squares inference, the most widely used form of inference in the current literature [23, 30, 44]. εKTELO extends these methods and formulates them as general operators, allowing us to replicate past algorithms and consider new forms of inference that support constraints. The first variant solves a classical least squares problem:

Definition 5.1 (Ordinary Least Squares (LS)).

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{M}\mathbf{x} - \mathbf{y}\|_2. \quad (2)$$

Our second variant imposes a non-negativity constraint on $\hat{\mathbf{x}}$.

Definition 5.2 (Non-negative Least Squares (NNLS)). Given scaled query matrix $\mathbf{M}$ and answer vector $\mathbf{y}$, the non-negative least squares estimate of $\mathbf{x}$ is

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \geq 0} \|\mathbf{M}\mathbf{x} - \mathbf{y}\|_2. \quad (3)$$

These inference methods can also support some forms of prior information, particularly if it can be represented as a linear query. For example, if the total number of records in the input table is publicly known, or other special queries have publicly available answers, then they can be added as “noisy” answers with negligible noise scale and they will naturally be incorporated into the inference process and the derivation of new query estimates.

We also support an inference method based on a multiplicative weights update rule, which is used in the MWEM [21] algorithm. This inference algorithm is closely related to the principle of maximum entropy and is especially effective when one has measured an incomplete set of queries.

Defining inference under vector transformations. Recall that in the discussion above we describe inference as operating on a single vector $\mathbf{x}$ with a corresponding query matrix $\mathbf{M}$. However, plans can include an arbitrary combination of vector transformations, followed by query operators, resulting in a collection of query answers defined over various vector representations of the data. εKTELO handles this by taking advantage of the structure of vector transformations and query operators, both of which perform linear transformations, therefore making it possible to map measured queries back on to the original domain (i.e., a vector produced by the VECTORIZE operation) and perform inference there. This allows for the most complete form of inference but other alternatives are conceivable, for example by performing inference locally on transformed vectors and combining inferred queries. This might have efficiency advantages, but would likely sacrifice accuracy, and is left for future investigation.

Inference: impact on accuracy. Because inference is an operator in εKTELO algorithm authors are encouraged to use inference consistently, using all available measurements, even if they are measured in different parts of a plan. In contrast, some existing algorithms use inference in an ad hoc manner, performing inference on one set of measurements separately from another set of

measurements. As we show below, for unbiased plans, this is always sub-optimal and ϵ_{KTELO} helps to relieve the algorithm designer of the complexity of integrating measured information properly. The following theorem follows the intuition that any unbiased noisy measurement provides information about the true data that can lower error, in expectation.

THEOREM 5.3. *Given any (full rank) matrix $\mathbf{M}$ of linear measurements and any linear query $\mathbf{q}$, the expected error of $\mathbf{q}$ is never higher if we include additional linear measurements using least squares inference.*

PROOF. Assuming all measurements have variance 1, the expected error on a query $\mathbf{q}$ is $\text{Error}_{\mathbf{M}}(\mathbf{q}) = \mathbf{q}(\mathbf{M}^T \mathbf{M})^{-1} \mathbf{q}^T$ [30]. In general, the variance of the measurements depends on the privacy budget and the sensitivity of $\mathbf{M}$, but they can always be scaled to have variance 1. If we augment $\mathbf{M}$ with a new linear query $\mathbf{b}$, then it becomes $\mathbf{M}' = \begin{bmatrix} \mathbf{M} \\ \mathbf{b} \end{bmatrix}$. We can write $\mathbf{M}'^T \mathbf{M}' = \mathbf{M}^T \mathbf{M} + \mathbf{b}^T \mathbf{b}$, where $\mathbf{b}^T \mathbf{b}$ is the outer product. Using the Sherman-Morrison formula [48], we see that

$$(\mathbf{M}'^T \mathbf{M}')^{-1} = (\mathbf{M}^T \mathbf{M})^{-1} - \frac{1}{1 + \mathbf{b}(\mathbf{M}^T \mathbf{M})^{-1} \mathbf{b}^T} (\mathbf{M}^T \mathbf{M})^{-1} \mathbf{b}^T \mathbf{b} (\mathbf{M}^T \mathbf{M})^{-1}.$$

Since $(\mathbf{M}^T \mathbf{M})^{-1}$ is positive-definite, $\mathbf{b}(\mathbf{M}^T \mathbf{M})^{-1} \mathbf{b}^T \geq 0$ and the fraction is just some positive constant c . With some algebraic manipulation, we arrive at the following expression:

$$\text{Error}_{\mathbf{M}'}(\mathbf{q}) = \text{Error}_{\mathbf{M}}(\mathbf{q}) - c \mathbf{q}(\mathbf{M}^T \mathbf{M})^{-1} \mathbf{b}^T \mathbf{b} (\mathbf{M}^T \mathbf{M})^{-1} \mathbf{q}^T.$$

Letting $v = \mathbf{q}(\mathbf{M}^T \mathbf{M})^{-1} \mathbf{b}^T$, we can write that as $\text{Error}_{\mathbf{M}'}(\mathbf{q}) = \text{Error}_{\mathbf{M}}(\mathbf{q}) - cv^2$, since $(\mathbf{M}^T \mathbf{M})^{-1}$ is symmetric. Clearly cv^2 is non-negative, so $\text{Error}_{\mathbf{M}'}(\mathbf{q}) \leq \text{Error}_{\mathbf{M}}(\mathbf{q})$. This completes the proof. $\square$

6 EXPRESSING KNOWN ALGORITHMS

As stated in Section 2, ϵ_{KTELO} represents differentially private algorithms as plans composed over a rich library of operators and supports not only simple linear sequences but also plans with iteration, recursion, and branching. To highlight the expressiveness of ϵ_{KTELO} , we re-implemented state-of-the-art algorithms as ϵ_{KTELO} plans. Once the necessary operators are implemented, the plan definition for an existing algorithm is typically a few lines of code for combining operators and managing parameters. We performed extensive testing to confirm that reimplementations in ϵ_{KTELO} of existing algorithms provide statistically equivalent outputs.

We examined 12 differentially private algorithms for answering low-dimensional counting queries that were deemed competitive² in a recent benchmark study [22] and one new algorithm published after the benchmark study [33]. PLAN #1-13 in Figure 2 abstract their ϵ_{KTELO} implementations as *plan signatures* where operators are represented using colored abbreviations.

6.1 Re-implementing Existing Algorithms

The algorithms are listed roughly in the order in which they were proposed in the literature and reflect the evolution of increasingly complex algorithmic techniques. The simplest Algorithm, IDENTITY [13], is a natural application of the Laplace mechanism. It simply measures each component of the data vector. Algorithms 2 through 5 reflect the evolution of more sophisticated measurements selection, targeted toward specific workloads. Many of these techniques were originally designed to support range queries (a small subclass of linear queries) over one- or two-dimensional (2E) data. PRIVELET [51] uses a Haar wavelet as its measurements, which allows for sensitivity that grows logarithmically with the domain size yet allows accurate reconstruction of any range query. The HIERARCHICAL (H2) technique uses measurements that form a binary tree over the domain,

²This is the subset of algorithms that offered the best accuracy for at least one of the input settings of the benchmark.

ID	Cite	Algorithm name	Plan signature
1	Dwork et al. 2006	Identity	SI LM
2	Xiao et al. 2010	Privelet	SP LM LS
3	Hay et al. 2010	Hierarchical (H2)	SH2 LM LS
4	Qardaji et al. 2013	Hierarchical Opt (HB)	SHB LM LS
5	Li et al. 2014	Greedy-H	SG LM LS
6	-	Uniform	ST LM LS
7	Hardt et al. 2012	MWEM	I:(SW LM MW)
8	Zhang et al. 2014	AHP	PA TR SI LM LS
9	Li et al. 2014	DAWA	PD TR SG LM LS
10	Cormode et al. 2012	Quadtree	SQ LM LS
11	Qardaji et al. 2013	UniformGrid	SU LM LS
12	Qardaji et al. 2013	AdaptiveGrid	SU LM LS PU TP[SA LM] LS
13	McKenna et al. 2018	HDMM	SHD LM LS
14	NEW	DAWA-Striped	PS TP[PD TR SG LM] LS
15	NEW	HB-Striped	PS TP[SHB LM] LS
16	NEW	HB-Striped_kron	SS LM LS
17	NEW	PrivBayesLS	SPB LM LS
18	NEW	MWEM variant b	I:(SW SH2 LM MW)
19	NEW	MWEM variant c	I:(SW LM NLS)
20	NEW	MWEM variant d	I:(SW SH2 LM NLS)

Fig. 2. The high-level signatures of plans implemented in εKTELO (referenced by ID). All plans begin with a vectorize transformation, omitted for readability. We also omit parameters of operators, including ϵ budget shares. $I(subplan)$ refers to iteration of a subplan, and $TP[subplan]$ means that *subplan* is executed on each partition produced by TP.

achieving effects similar to the wavelet measurements. QUADTREE [12] is the two-dimensional realization of the hierarchical structures. All the algorithms above follow similar design idioms, allowing us to implement them using operator sequences of the same pattern: *Query selection*, *Query*, and *Inference*.

All of the algorithms above are *data-independent*, with constant error rates for any input dataset. More recent algorithms are *data-dependent*, displaying different error rates on different inputs, often because the algorithmic techniques are adapting to features of the data to lower error. The simplest data-dependent algorithm is UNIFORM, which simply estimates the total number of records in the input and assumes uniformity across the data vector. This simple algorithm also follows the simple pattern.

A more complex example is the Multiplicative-Weights Exponential Mechanism (MWEM) [21], which takes a workload of linear queries as input and runs several rounds of estimation, measuring one workload query in each round, and using the multiplicative update rule to revise its estimate of the data vector. In each round, the Exponential Mechanism is used to select the workload query that is most poorly approximated using the current data vector estimate.

In Algorithm 3, we rewrite the algorithm with abstracted subroutines and find out that this algorithm can be represented as several iterations of the *Query selection*, *Query*, and *Inference* sequence. In Figure 2, the iteration inherent to PLAN #7 (MWEM) is shown with $I : (..)$.

Other data-dependent algorithms exploit partitioning, in which components of the data vector are merged and estimated only in their entirety, which uniformity assumption imposed within the regions. The DAWA [29] and AHP [57] algorithms have custom partition selection methods

ALGORITHM 3: MWEM

```

1:  $D \leftarrow \text{PROTECTED}(\text{source\_uri})$  ▷ Init
2:  $\hat{x} \leftarrow \text{T-VECTORIZE}(D)$  ▷ Transform
3: for  $i = 1 : T$  do
4:    $M \leftarrow \text{WORSTAPPROX}(\hat{x}, \epsilon/2T)$  ▷ Query Selection
5:    $y \leftarrow \text{VECTORLAPLACE}(M, \epsilon/2T)$  ▷ Query
6:    $\hat{x} \leftarrow \text{MULTWEIGHTS}(M, y)$  ▷ Inference
7: end for
8: return  $\hat{x}$  ▷ Output

```

that consume part of the privacy budget to identify approximately uniform partition blocks. The partition selection methods work by finding a grouping of the bins in a vector and are the key innovations of the algorithms. We encapsulate these subroutines as new operators in our framework (in the cases above, we added a partition selection operators DAWAPARTITION (PD) in PLAN #9 and AHPPARTITION (PA) in PLAN #8).

UNIFORMGRID and ADAPTIVEGRID [43] focus on 2D data and both end up with partitioned sets of measurements forming a grid over a 2D domain. UNIFORMGRID imposes a static grid, while ADAPTIVEGRID uses an initial round of measurements to adjust the coarseness of the grid, avoiding estimation of small sparse regions.

6.2 Re-implementation Strategies

The process of re-implementing in ϵ KTELO this seemingly diverse set of algorithms consisted of breaking the algorithms down into key subroutines and translating them into operators. To summarize, the translation strategy typically falls into one of three categories.

The first translation strategy was to identify specific implementations of common differentially private operations and replace them with a single unified general-purpose operator in ϵ KTELO. For instance, the Laplace mechanism (LM), which adds noise drawn from the Laplace distribution, appears in *every one* of the 13 algorithms. Noise addition can be implemented in a number of ways (e.g., calling a function in the numpy.random package, taking the difference of exponential random variables, etc.). In ϵ KTELO, all these plans call the same VECTOR LAPLACE operator with a single unified sensitivity calculation.

Another less obvious example of this translation is for subroutines that infer an estimate of x using noisy query answers. With the exception of IDENTITY and MWEM, each of the algorithms uses instances of least squares inference, often customized to the structure of the noisy query answers. For instance, PRIVELET uses Haar wavelet reconstruction, hierarchical strategies like HB and DAWA use a tree-based implementation of inference, and others like UNIFORM and AHP use uniform expansion. We replaced each of these custom inference methods with a single general-purpose least squares inference operator (LS operators in Figure 2). It would still be possible to implement a specialized inference operator in ϵ KTELO that exploited particular properties of a query set, but, given the efficient inference methods described in Section 7.6, we did not find this to be beneficial.

Our second translation strategy was to identify higher-level patterns that reflect design idioms that exist across multiple algorithms. In these cases, we replace one or more subroutines in the original code with a sequence of operators that capture this idiom. As shown earlier, PLAN #2, 3, 4, 5, 6, 10, 11, and 13 all consist of the operator sequence: Query selection, Query (LM), and Inference (LS), differing only in Query selection method. For other algorithms, this idiom reappears as a subroutine, as in PLAN #8 (AHP) and PLAN #9 (DAWA).

Finally, we were left with distinct subroutines of algorithms that represented key intellectual advances in the differential privacy literature. We encapsulate such subroutines as new operators (e.g., PA in PLAN #8 (AHP) and PD in PLAN #9 (DAWA)) in the framework.

6.3 Benefits

We highlight the benefits of re-implementing known algorithms in εKTELO.

Code reuse. Once reformulated in εKTELO, nearly all algorithms use the VECTOR LAPLACE operator and least squares inference. This means that any improvements to either of these operators will be inherited by all the plans. We show such an example in Section 7.6.

Reduced privacy verification effort. Code reuse also reduces the number of critical operators that must be carefully vetted. The operators that require careful vetting are ones that consume the privacy budget, which are the Private→Public operators in Figure 1. These are: VECTOR LAPLACE, the partition selection operators for both DAWA [29] and AHP [57], a query selection operator used by PrivBayes [55], and a query selection operator used by the MWEM [21] algorithm that privately derives the worst-currently approximated workload query. In contrast, for the DPBench code base, the entire code has to be vetted to audit the use and management of the privacy budget. The end result is that verifying the privacy of an algorithm requires checking fewer lines of code. For example, to verify the QuadTree algorithm in the DPBench codebase requires checking 163 lines of code. However, with εKTELO, this only requires vetting the 30-line VECTOR LAPLACE operator. (Furthermore, by vetting just this one operator, we have effectively vetted 10 of the 18 algorithms in Figure 2, since the only privacy sensitive operator these algorithms use is VECTOR LAPLACE.). When we consider all of the DPBench algorithms in Figure 2, algorithms 1–12, verifying the DPBench implementation requires checking a total of 1,837 lines of code while vetting all the privacy-critical operators in εKTELO requires checking 517 lines of code.

Transparency. As noted above, εKTELO plans make explicit the typical patterns that result in accurate differentially private algorithms. Moreover, εKTELO plans help clarify the distinctive ingredients of state-of-the-art algorithms. For instance, DAWA and AHP (PLAN #9 and PLAN #8, respectively, in Figure 2) have the same structure but differ only in two operators: partition selection and query selection.

7 IMPLEMENTATION: EFFICIENT MATRIX SUPPORT

Matrices and operations on matrices are central to the implementation of εKTELO operators but can become a performance bottleneck. In this section, we describe a set of specialized matrix representation techniques, based on the implicit definition of matrices, which allows for greater scalability as the size of the data vector grows.

We review next the types of matrix objects in εKTELO and then, in Section 7.2, the different ways matrices can be represented including implicit matrices. In Section 7.3, we decompose the common matrix operations in εKTELO into a small set of primitive operations that every implicit matrix should support. We then describe in Section 7.4 a general matrix type from which implicit matrices can be built, and use that matrix type, in Section 7.5, to implement common query selection operators in εKTELO. We conclude with Section 7.6 describing the implementation of inference using implicit matrices.

7.1 Matrix Types and Their Operations

Recall that matrices are used to represent three different objects in εKTELO: sets of workload queries, sets of measurement queries, and partitions of the domain. In all cases, the matrices

Table 1. Types of Matrix Objects in ϵ KTELO (Workload, Measurement, Partition) and the Key Computations Performed in Plans along with the Primitive Methods Required to Support Each Computation

Key computations by matrix type		Primitive methods
Workload matrix, W		
Query evaluation	$\mathbf{W}\hat{\mathbf{x}}$	Matrix-vector product
L_1 Sensitivity	$\ \mathbf{W}\ _1$	Abs, Transpose, Matrix-vector product
L_2 Sensitivity	$\ \mathbf{W}\ _2$	Sqr, Transpose, Matrix-vector product
Gram Matrix	$\mathbf{W}^T\mathbf{W}$	Transpose, Matrix multiplication
Row indexing	$\mathbf{w}_i$	Transpose, Matrix-vector product
Measurement matrix, M		
Query evaluation	$\mathbf{M}\mathbf{x}$	Matrix-vector product
L_1 Sensitivity	$\ \mathbf{M}\ _1$	Abs, Transpose, Matrix-vector product
L_2 Sensitivity	$\ \mathbf{M}\ _2$	Sqr, Transpose, Matrix-vector product
Inference (LS)	$\arg \min_{\mathbf{x}} \ \mathbf{M}\mathbf{x} - \mathbf{y}\ _2$	Transpose, Matrix-vector product
Inference (NNLS)	$\arg \min_{\mathbf{x} \geq 0} \ \mathbf{M}\mathbf{x} - \mathbf{y}\ _2$	Transpose, Matrix-vector product
Inference (MW)	$\hat{\mathbf{x}}^{(k+1)} \propto \hat{\mathbf{x}}^{(k)} \odot \exp(\mathbf{g}/N)$ $\mathbf{g} = 0.5\mathbf{M}^T(\mathbf{M}\hat{\mathbf{x}}^{(k)} - \mathbf{y})$	Transpose, Matrix-vector product
Partition matrix, P		
Reduce workload	$\mathbf{W}' = \mathbf{W}\mathbf{P}^+$	Transpose, Matrix multiplication
Reduce data vector	$\mathbf{x}' = \mathbf{P}\mathbf{x}$	Matrix-vector product
Expand workload	$\mathbf{W} = \mathbf{W}'\mathbf{P}$	Matrix multiplication
Expand data vector	$\mathbf{x} = \mathbf{P}^+\mathbf{x}'$	Transpose, Matrix-vector product

contain one column for each element of a corresponding data vector. In the case of both workload and measurement matrices, rows represent linear queries. A partition matrix describes a linear transformation that can be applied to a data vector or to a workload; one row describes a set of elements of the domain that will be combined after the transformation.

The key computations on each matrix type are shown in Table 1 (in the left two columns). Workload matrices and measurement matrices both represent sets of queries and so they share similar computations (such as query evaluation on a data vector and calculation of sensitivity) however we only do inference on measurement matrices. Partition matrices are used to reduce and expand both workload matrices and data vectors.

In common plans, the number of rows in a workload or measurement measurement matrix can be as large or larger than n , the number of elements in the data vector. Partitions have at most n rows, but may still be large. For plans operating on large data vectors, where n approaches the size of memory, these matrices, in standard form, are infeasible to represent in memory and operate on. To address this, ϵ KTELO provides flexible and efficient matrix capabilities that can be used for the efficient implementation of operators.

7.2 Matrix Representations: Dense, Sparse, and Implicit

The matrix class in ϵ KTELO supports matrices using a combination of the dense, sparse, and implicit matrices. These representations differ in their space utilization, their generality, and the efficiency of the matrix operations they support.

A dense $m \times n$ matrix is the standard representation that stores mn values. Obviously, any matrix can be represented in this manner and all operations in Table 1 are supported. A sparse matrix stores only non-zero elements of a matrix. Any matrix can be represented in sparse form, but its efficiency depends on the number of nonzero entries. Where $nnz(\mathbf{A})$ denotes the number of nonzero

elements in matrix A , if $nnz(A) \approx mn$, and then sparse matrices do not offer any benefit and may even be more expensive to represent than the dense representation. However, if $nnz(A) \ll mn$, then there may be significant improvements to performance in using this representation.

An implicit matrix is a virtual representation of a matrix that may not explicitly store all (or even any) of the elements of the matrix. Because it is a virtual object, it must define appropriate methods so that computations with the implicit matrix produce correct results. While not all matrices allow for efficient implicit representations, we have found that many of the matrices used in εKTELO operators have a structure that can be exploited for efficient implicit representation. Note that implicit matrix representations are lossless: They do not approximate some dense matrix but represent it exactly. Therefore an implicit matrix can always be materialized in sparse or dense form, although the goal is to perform computations without materialization.

As an example of an implicit matrix, recall the Prefix workload, an encoding of an empirical CDF, which was used in the example plan (Algorithm 1) of Section 2:

Example 7.1 (The Prefix Workload: Dense, Sparse, and Implicit). In dense form, the prefix workload is defined as a lower-triangular matrix containing 1's. If $n = 5$, then we have:

$$\mathbf{W}_{\text{pre}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 \end{bmatrix}.$$

A sparse representation of $\mathbf{W}_{\text{pre}}$ would store (a list of) only the nonzero elements of this matrix, but the space complexity of both dense and sparse representations remains $O(n^2)$. In addition, the time complexity of computing matrix-vector products using the dense or sparse representation is $O(n^2)$. However, the Prefix matrix can be completely specified by a single parameter, n , which is the only state stored for the implicit version of $\mathbf{W}_{\text{pre}}$. Further, we can evaluate the matrix-vector product $y = \mathbf{W}_{\text{pre}}\mathbf{x}$ using a simple one-pass algorithm over $\mathbf{x}$: First, compute $y_1 = x_1$. Then, for $k = 2 \dots n$, compute $y_k = y_{k-1} + x_k$. Therefore, by representing the Prefix workload implicitly we can achieve $O(1)$ space complexity and $O(n)$ time complexity for computing matrix-vector products.

7.3 Computing with Implicit Matrices

Most implicit matrices require very little internal state to be stored. The main challenge is therefore to insure that all necessary computations involving an implicit matrix can be carried out efficiently, hopefully without falling back to materialization of the dense form of the matrix. Before defining additional implicit matrix constructions, we review the key computations εKTELO matrix objects must support.

A careful examination of the operators currently implemented in εKTELO resulted in the list of key computations in Table 1, where the left two columns describe operations on matrices that commonly occur in plans. Importantly, these plan-level matrix computations can be decomposed into just five fundamental matrix methods, which we call *primitive methods*: matrix-vector product, transpose, matrix multiplication, element-wise absolute value (*abs*), and element-wise square (*sqr*). Matrix-vector product takes as input a vector and returns a vector. The multiplication of two implicit matrices returns a new implicit matrix, as does transpose, which returns another implicit matrix that represents the transpose linear transformation. Element-wise *sqr* and *abs* both return new implicit matrices.

Below we review the key computations on εKTELO matrix objects in Table 1, how they can be decomposed into primitive methods, and implementation considerations. Our goal will then be to construct implicit matrices that efficiently support the primitive methods.

Query evaluation and data reduction. In ϵ KTELO plans, matrix-vector multiplication is used by workload and measurement matrices for query evaluation and by partition matrices for reduction of the data vector.

Sensitivity. The sensitivity of measurement and workload matrices can be computed using a combination of primitive methods (abs, sqr, transpose, and matrix-vector product). For a matrix $\mathbf{M}$, we compute the maximum column sum of $\text{abs}(\mathbf{M})$ (for L_1 sensitivity) or $\text{sqr}(\mathbf{M})$ (for L_2 sensitivity), which can be done by doing a transpose-matrix-vector product with $\mathbf{1}$ —the vector of all 1's:

$$\|\mathbf{M}\|_1 = \max(\text{abs}(\mathbf{M})^T \mathbf{1}) \quad \|\mathbf{M}\|_2 = \sqrt{\max(\text{sqr}(\mathbf{M})^T \mathbf{1})}.$$

Inference. The most common form of inference in ϵ KTELO plans is based on least squares. While classical solutions to the least squares problem involve matrix decompositions and computation of the pseudo-inverse, we will show that iterative methods lead to much greater scalability in combination with our implicit matrix representations. Least squares, non-negative least squares, and multiplicative-weights inference can all be implemented using just the matrix-vector product and transpose primitive methods. We will discuss iterative inference in Section 7.6.

Gram Matrix. Some workload-adaptive mechanisms like GreedyH and HDMM require the (materialized) gram matrix of the workload. For a workload $\mathbf{W}$, the gram matrix is $\mathbf{W}^T \mathbf{W}$. This computation can be implemented in terms of transpose, matrix multiplication, and materialize. For extremely large workloads with special structure, where $\mathbf{W}^T \mathbf{W}$ is much smaller than $\mathbf{W}$ (like the set of all range queries), a more efficient version can be implemented directly that avoids using the primitive methods.

Partition reduction and expansion. Partition matrices need to be able to reduce the data vector and workload. They also have to be able to do the reverse expansion operations. As we show in the proof of Prop. 8.3, because of the special structure of partition matrices, the pseudo-inverse of any matrix $\mathbf{P}$ can be computed as the product of $\mathbf{P}^T$ and a diagonal matrix $\mathbf{D}$. Thus, partition matrices simply need the three primitive methods: matrix-vector product, transpose, and matrix-multiplication.

Row Indexing. The MWEM algorithm and its variants (described in Section 9.1) use the worst-approximated query selection operator, which requires row indexing, or materialization of the i th row of a matrix. This can be implemented in terms of the primitive methods transpose and matrix-vector product as follows: $\mathbf{w}_i = \mathbf{W}^T \mathbf{e}_i$, where $\mathbf{e}_i$ is the i th column of an identity matrix.

Materialize. If a plan requires working with matrices in a manner not supported by the interface of our implicit matrices, then the matrix can always be materialized, at which point standard implementations of matrix methods can be employed. As in row indexing, materialization can be performed by a sequence of matrix vector products with the columns of an identity matrix, i.e., $\mathbf{Ae}_i$ for $i = 1, \dots, n$.

7.4 Generalized Matrix Construction

ϵ KTELO contains a matrix class, denoted *EMatrix*, that generalizes dense, sparse, and implicit matrices, supporting flexible matrix construction using a small set of specially designed *core matrices* that may then be combined with combining operations (union, product, and Kronecker product). This provides a flexible and extensible mechanism for constructing a wide range of matrices. The

Table 2. Comparison of Core Implicit Matrices to Their Corresponding Sparse and Dense Representations in Terms of Space Usage and Time Complexity of a Matrix-vector Product

Core Matrix	Implicit		Dense Space/Time	Sparse Space/Time
	Space Usage	Time (mat-vec)		
<i>Identity</i>	$O(1)$	$O(n)$	$O(n^2)$	$O(n)$
<i>Ones</i>	$O(1)$	$O(m + n)$	$O(mn)$	$O(mn)$
<i>Prefix</i>	$O(1)$	$O(n)$	$O(n^2)$	$O(n^2)$
<i>Suffix</i>	$O(1)$	$O(n)$	$O(n^2)$	$O(n^2)$
<i>Wavelet</i>	$O(1)$	$O(n \log n)$	$O(n^2)$	$O(n \log n)$

For sparse and dense matrices, the time complexity is the same as the space complexity.

following grammar describes the construction of *EMatrix* instances:

$$\begin{aligned}
 \text{CoreMatrix} &= \text{Identity} \mid \text{Ones} \mid \text{Prefix} \mid \text{Suffix} \mid \text{Wavelet} \\
 \text{EMatrix} &= \text{DenseMatrix} \mid \text{SparseMatrix} \mid \text{CoreMatrix} \\
 \text{EMatrix} &= \text{Union}(\text{EMatrix}, \text{EMatrix}) \\
 \text{EMatrix} &= \text{Product}(\text{EMatrix}, \text{EMatrix}) \\
 \text{EMatrix} &= \text{Kronecker}(\text{EMatrix}, \text{EMatrix}).
 \end{aligned}$$

Unless an *EMatrix* is defined as a single *SparseMatrix* or *DenseMatrix*, we consider it implicit, since it is not fully materialized.

Core matrices. The *CoreMatrix* type forms the basic building block for *EMatrix* and each is defined implicitly. The following are custom core matrices we designed to support εKTELO operators:

- **Identity:** Identity is the simplest building block. It is defined as the matrix I having the property that $I\mathbf{v} = \mathbf{v}$ for all vectors $\mathbf{v}$. Thus, the implementation of matrix-vector product is trivial. Similarly, transpose, abs, and sqr are simple no-ops.
- **Ones:** Ones is the $m \times n$ matrix of all ones. Matrix-vector products can be efficiently computed by summing up the entries of the input vector, and constructing a m -length vector with that value. The transpose is a $n \times m$ Ones matrix, and abs and sqr are simple no-ops. **Total** is a special case of the Ones matrix where $m = 1$.
- **Prefix and Suffix:** The description of prefix and the algorithm for efficiently computing matrix-vector products is given in Example 7.1. The transpose of Prefix is Suffix, and abs and sqr are simple no-ops.
- **Wavelet:** Wavelet is the Haar wavelet transform. Efficient algorithms exist for evaluating matrix-vector products implicitly with the Haar wavelet [51]. The transpose has a similar form, but abs and sqr, if needed, must materialize the matrix. However, for this matrix sensitivity may be computed directly, without going through abs and sqr.

The primitive methods described above have very simple and efficient implementations for the core matrices. In Table 2, we report the space utilization for each core matrix, along with the time complexity for one of the most important primitive methods (matrix-vector product). We compare the complexity of the core implicit matrices with their standard dense and sparse representations, showing significant reductions in space usage—up to a factor n^2 . We observe a similar reduction in time complexity of matrix-vector products by up to a factor of n .

Composing matrices. Core matrices and arbitrary sparse or dense matrices can be combined using a Union, Product (including with a constant), and Kronecker Product to form new matrices that are implicit (or partially implicit).

If matrix $\mathbf{M}_1$ and $\mathbf{M}_2$ each represent queries, then $\text{Union}(\mathbf{M}_1, \mathbf{M}_2)$ is a matrix that represents the union of the queries of $\mathbf{M}_1$ and $\mathbf{M}_2$. It is useful for building complex workloads and measurement matrices, and it also important in plans to bring together all measured queries for global inference. Product is less frequently needed but is used for multiplying partition matrices with workload and measurement matrices.

Kronecker product is especially useful for constructing workload and measurement matrices over multi-dimensional domains. Suppose our input is a relation $R(A, B)$, we vectorize its projection of $\pi_A(R)$ to get data vector $\mathbf{x}_A$, and we define a set of queries of interest as matrix $\mathbf{M}_A$. If we similarly form a matrix of queries $\mathbf{M}_B$ over the vectorization of $\pi_B(R)$, then $\text{Kronecker}(\mathbf{M}_A, \mathbf{M}_B)$ (denoted $\mathbf{M}_A \otimes \mathbf{M}_B$ in matrix equations) is a matrix that encodes a new set of queries over both attributes A and B and it contains $q_{a_i} \wedge q_{b_j}$ for each q_{a_i} in $\mathbf{M}_A$ and each q_{b_j} in $\mathbf{M}_B$, i.e., it contains the conjunctive combination of all pairs of queries drawn from $\mathbf{M}_A$ and $\mathbf{M}_B$.

The formal definition of the Kronecker product is as follows.

Definition 7.2. The Kronecker product $\mathbf{A} \otimes \mathbf{B}$ between a $m_A \times n_A$ matrix $\mathbf{A}$ and a $m_B \times n_B$ matrix $\mathbf{B}$ is a $m_A m_B \times n_A n_B$ matrix defined as:

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{11}\mathbf{B} & \dots & a_{1n_A}\mathbf{B} \\ \vdots & \ddots & \vdots \\ a_{m_A 1}\mathbf{B} & \dots & a_{m_A n_A}\mathbf{B} \end{bmatrix}.$$

Kronecker products were first used within the context of a privacy mechanism by McKenna et al. [33].

We present these as binary operations on implicit matrices, but because they are associative, they can also be applied to a collection of k sub-matrices. For example, for $k = 3$ we may write $\text{Union}(\mathbf{A}, \mathbf{B}, \mathbf{C})$ as shorthand notation for $\text{Union}(\mathbf{A}, \text{Union}(\mathbf{B}, \mathbf{C}))$.

Example 7.3. Suppose our input relation is $R(\text{age}, \text{income}, \text{marital-status})$, where age and income are discretized into 100 bins, and marital-status is a categorical attribute with seven possible values, resulting in a data vector of size 70,000. We want to accurately answer range queries on age and income, broken down by various marital statuses. Thus, we may construct the following workload using tools from above:

$$\mathbf{W} = \text{Kronecker}(\text{Prefix}, \\ \text{Prefix}, \\ \text{Union}(\text{Total}, \text{Identity}, \text{Dense})),$$

where *Dense* is a 2×7 query matrix with two queries that aggregate over the marital status attribute into two groups: “married” and “unmarried.” Using Table 3, we see that the only storage required to represent $\mathbf{W}$ is the 2×7 *Dense* matrix and metadata for the *Prefix*, *Total*, and *Identity* matrices. In contrast, the sparse and dense representation of $\mathbf{W}$ would require about 8 GB and 56 GB, respectively.

Supporting the primitive methods. Core, sparse, and dense matrices have native support for the primitive methods discussed previously. When a primitive method is invoked for an *EMatrix* that results from one or more of the combining operations, the work is delegated to the constituent sub-matrices, and thus the matrices formed by composition inherit the performance characteristics

Table 3. Space and Time Complexity of Composed Matrices in Terms of the Complexity of Sub-matrices

Matrix	Operation	Space Usage	Time (mat-vec)
$Union(A, B)$	$\begin{bmatrix} A \\ B \end{bmatrix}$	$Space(A) + Space(B)$	$Time(A) + Time(B)$
$Product(A, B)$	AB	$Space(A) + Space(B)$	$Time(A) + Time(B)$
$Kronecker(A, B)$	$A \otimes B$	$Space(A) + Space(B)$	$n_B Time(A) + m_A Time(B)$

of the sub-matrices. In particular, the key primitive methods can be implemented efficiently on matrices formed from unions, products, and Kronecker products.

The key performance characteristics of composed matrices are summarized in Table 3.

7.5 Matrix Constructions for εKTELO Operators

Using the generalized matrix construction described above, we can re-implement many of the existing query selection and partition operators currently in εKTELO. We will show in Section 10 that the use of implicit matrices leads to significant improvements in efficiency and scalability, as well as the reduction of space consumption.

Query selection operators based on range queries. A notable class of query matrices that we can efficiently represent using the tools from above is an arbitrary collection of range queries. This type of workload has been extensively studied in the literature and many of the query selection operators in εKTELO are specially designed sets of range queries including H2, Hb, QuadTree, UniformGrid, and AdaptiveGrid.

Recall that a single range query over a one-dimensional domain can be specified by a pair of indices (i, j) , and a workload of range queries can be represented as a list of these pairs. This suggests we can store any range query workload using only $O(m)$ space, where m is the number of queries. Furthermore, one way to evaluate matrix-vector products is by iterating through each query one by one and evaluating $\sum_{k=i}^j x_k$. The time complexity of this approach is $O(mn)$ in the worst case, which is equivalent to the sparse and dense representations.

Our general matrix construction allows us to do even better by exploiting the fact that any range query can be expressed as the difference of two prefix queries. Thus the matrix can be represented as $Product(Sparse, Prefix)$, where $Sparse$ is a $m \times n$ sparse matrix with two non-zero entries per row. An illustrating example is shown below:

Example 7.4 (Range Queries). A collection of four range queries over a domain of size five, represented implicitly as the product of a $Sparse$ matrix and the $Prefix$ matrix (displayed here in dense form for illustration purposes):

$$\begin{bmatrix} 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} -1 & 0 & 0 & 1 & 0 \\ 0 & 0 & -1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \\ -1 & 1 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 \end{bmatrix}.$$

Using this construction, we can evaluate matrix-vector products in $O(n + m)$ time, which is a substantial improvement over the other representations. The range query construction can be naturally extended to multi-dimensional domains by replacing $Prefix$ with $Kronecker(Prefix, \dots, Prefix)$ and replacing $Sparse$ with a sparse matrix with up to 2^d nonzero entries per row, where d is the number of dimensions of the domain.

A special case of this range query construction is hierarchical and grid-based matrices used by H2, Hb, and QuadTree. These matrices always have an Identity matrix, and while these matrices can be efficiently represented using this construction, it is actually slightly better to observe that they all contain Identity as a sub-matrix, and instead represent them as $Union(Identity, Product(Sparse, Prefix))$. We use this representation in our empirical evaluation.

Note that even though *Product* does not natively support *abs* and *sqr*, for the case of range queries, or more generally any matrix with binary values, *abs* and *sqr* are simple no-ops.

Representing marginals. A common task for multi-dimensional data analysis is computing the marginals of a dataset. Marginals may be used both as part of workloads and measurement matrices. They can be efficiently represented using the tools from above, as demonstrated in Example 7.5.

Example 7.5 (Marginals). Any marginal can be represented as a Kronecker product of *Identity* and *Total* building blocks. For example, the two-way marginal that sums out the second attribute can be encoded as:

$$W_{13} = Kronecker(Identity, Total, Identity).$$

Further, an arbitrary collection of marginals can be encoded as *Union* of these Kronecker products. All two-way marginals is

$$W_{2way} = Union(Kronecker(Identity, Identity, Total), \\ Kronecker(Identity, Total, Identity), \\ Kronecker(Total, Identity, Identity)).$$

Partition operators. The matrices used by partition operators are represented simply as *Sparse* matrices. While an implicit definition is possible, it would not offer any improvement in space or time over the sparse representation.

7.6 Implementing Inference

Inference is a fundamental operator that can improve error with no cost to privacy and, accordingly, we saw that it appeared in virtually every algorithm re-implemented in ϵ KTELO (as shown in Figure 2). But inference can be a costly operation. Recall that the input to inference is a measurement matrix, denoted by M , containing m queries defined over a data vector of size n , and the list of noisy answers y . The least squares solution (Equation (2)) is given by the solution to the normal equations $M^T M \hat{x} = M^T y$. Assuming $M^T M$ is invertible, then the solution is unique and can be expressed as $\hat{x} = (M^T M)^{-1} M^T y$. Often explicit matrix inversion is avoided, in favor of suitable factorizations of M (e.g., QR or SVD). However, the time complexity of such “direct” methods is still generally cubic in the domain size when $m = O(n)$. In practice we have found that the runtime of such direct methods is unacceptable when n is greater than about 5,000.

Algorithms in prior work [23, 43, 44, 51] have used least squares inference on large domains by restricting the selection of queries, namely to those representing a set of hierarchical queries. This allows for inference in time linear in the domain size, avoiding the explicit matrix representation of the queries. We avoid this approach in ϵ KTELO, because it means that a custom inference method may be required for each query selection operation, and because it limits the measurement sets that can be used. In addition, hierarchical methods only work for least squares but not least squares with non-negativity constraints.

An alternative approach to least squares inference is to use an iterative gradient-based method, which solves the normal equations by repeatedly computing matrix-vector products Mv and $M^T v$ until convergence. The time complexity of these methods is $O(kn^2)$ for dense matrix representations where k is the number of iterations. In experiments, we use a well-known iterative method,

LSMR [18]. Empirically, we observe LSMR to converge in far fewer than n iterations when $\mathbf{M}$ is well-conditioned, which is the case as long as the queries are not taken with vastly different noise scales, and thus we expect $k \ll n$.

In the original version of this work [54], we demonstrated a significant performance benefit from the combination of iterative solution methods with sparse matrix representations. This benefit is amplified when the underlying matrix representation is implicit. Letting $Time(\mathbf{M})$ denote the time complexity of evaluating a matrix-vector product with $\mathbf{M}$, the time complexity of least squares inference is $O(k \cdot Time(\mathbf{M}))$, where k is the number of iterations, as before. As shown in Table 2 and Table 3, $Time(\mathbf{M})$ is often $O(n)$, resulting in a very favorable $O(kn)$ time complexity for inference.

Iterative approaches, using implicit matrices, are well suited to the other inference methods in εKTELO: least squares with non-negativity constraints (Equation (3)) and multiplicative weights. For the former, we use the limited memory BFGS algorithm with bound constraints [10]. The time complexity of this algorithm is the same as LSMR, although the number of iterations needed for convergence may be different, and there is a constant factor overhead for storing the low-rank approximation to the inverse Hessian matrix. The multiplicative weights inference algorithm is defined in an iterative manner and requires the same primitive methods as ordinary least squares and non-negative least squares.

The combination of our general matrix construction techniques with iterative inference result in flexible inference capabilities for plan authors. With relative freedom, they can construct measurement matrices or combine measurements from parts of a plan and apply a single generic inference operator, which will run efficiently. In Section 10, we show that using iterative least squares on implicitly represented matrices, we can scale inference to domains consisting of hundreds of millions of cells while staying within modest runtime bounds, well beyond what is possible with sparse or dense representations.

8 WORKLOAD - BASED PARTITION SELECTION

In many cases, the goal of a differentially private algorithm (and its corresponding εKTELO plan) is to answer a given workload of queries, $\mathbf{W}$, defined in terms of a data vector $\mathbf{x}$. We describe next a method for reducing the representation of the $\mathbf{x}$ vector to precisely the elements required to correctly answer the workload queries. This is a new partition selection operator, called *workload-based partition selection*, which can be used as input to a V-REDUCEBYPARTITION transformation.

We define the partition below and prove that, under reasonable assumptions, using such domain reduction can never hurt accuracy. We provide an algorithm for computing the partition, which can be executed using implicit workload representations. Later, in Section 10.3, we will show empirically that using this partition in plans can offer significant improvement in both runtime and error.

8.1 The Workload-based Partition and Its Properties

For a workload $\mathbf{W}$ of linear queries described on data vector $\mathbf{x}$, it is often possible to define a reduction of $\mathbf{x}$, to a smaller $\mathbf{x}'$, and appropriately transform the workload to $\mathbf{W}'$, so that all workload query answers are preserved, i.e., $\mathbf{W}\mathbf{x} = \mathbf{W}'\mathbf{x}'$. Intuitively, such a reduction is possible when a set of elements of $\mathbf{x}$ is not distinguished by the workload: Each linear query in the workload either ignores it or treats it in precisely the same way. In that case, that portion of the domain need not be represented by multiple cells but instead by a single cell in a reduced data vector. It is in this sense that the reduction is lossless with respect to the workload. Following this intuition, the domain reduction can be computed from the matrix representation $\mathbf{W}$ of the workload by finding groups of identical columns: Elements of these groups will be merged in $\mathbf{W}$ to get $\mathbf{W}'$ while the corresponding cells in $\mathbf{x}$ are summed.

Example 8.1. Consider a table with schema Census(age, sex, salary). If the workload consists of queries Q1(salary ≤ 100K, sex = M) and Q2(salary > 100K, sex = F), then the workload only requires a data vector consisting of 2 cells. If the workload consists of all one-way marginals, then no workload-based data reduction is possible.

Note that calculating this partition only requires knowledge of the workload and is therefore done in the unprotected client space (and does not consume the privacy budget). The partition is then input to a V-REDUCEBYPARTITION transformation operator carried out by the protected kernel and its stability is one.

The new workload-based partition selection operator can be formalized in terms of a linear matrix operator, as follows.

Definition 8.2 (Workload-based Partition Selection). Let $\mathbf{w}_1, \dots, \mathbf{w}_n$ denote the columns of $\mathbf{W}$ and let $\mathbf{u}_1, \dots, \mathbf{u}_p$ denote those that are unique. For $h(\mathbf{u}) = \{j \mid \mathbf{w}_j = \mathbf{u}\}$, define the transformation matrix $\mathbf{P} \in \mathbb{R}^{p \times n}$ to have $P_{ij} = 1$ if $j \in h(\mathbf{u}_i)$ and $P_{ij} = 0$ otherwise. The reverse transformation is the pseudo-inverse $\mathbf{P}^+ \in \mathbb{R}^{n \times p}$.

The matrix $\mathbf{P}$ defines a partition of the data, which can be passed to V-REDUCEBYPARTITION to transform the data vector, and $\mathbf{P}^+$ can be used to transform the workload accordingly. When $\mathbf{P}$ is passed to V-REDUCEBYPARTITION, the operator produces a new data vector $\mathbf{x}' = \mathbf{P}\mathbf{x}$, where x'_i is the sum of entries in $\mathbf{x}$ that belong to i th group of $\mathbf{P}$. When viewed as an operation on the workload, $\mathbf{P}^+$ merges duplicate columns by taking the row-wise average for each group. This is formalized as follows.

PROPOSITION 8.3 (PROPERTIES: WORKLOAD-BASED REDUCTION). *Given transform matrix $\mathbf{P}$ and its pseudo-inverse $\mathbf{P}^+$, the following hold:*

- $\mathbf{x}' = \mathbf{P}\mathbf{x}$ is the reduced data vector;
- $\mathbf{W}' = \mathbf{W}\mathbf{P}^+$ is the workload matrix, represented over $\mathbf{x}'$;
- The transformation is lossless: $\mathbf{W}\mathbf{x} = \mathbf{W}'\mathbf{x}'$.

PROOF. First, note that $\mathbf{P}^+ = \mathbf{P}^T \mathbf{D}^{-1}$ where $\mathbf{D}$ is the $p \times p$ diagonal matrix with $D_{ii} = |h(\mathbf{u}_i)|$ for h defined in Definition 8.2. Since $\mathbf{P}$ has linearly independent rows, $\mathbf{P}^+ = \mathbf{P}^T (\mathbf{P}\mathbf{P}^T)^{-1}$ and $\mathbf{P}\mathbf{P}^T = \mathbf{D}$, because $h(\mathbf{u}_i)$ and $h(\mathbf{u}_j)$ are disjoint for $i \neq j$. By definition of $\mathbf{P}$, we see that $x'_i = \sum_{j \in h(\mathbf{u}_i)} x_j$ for $1 \leq i \leq p$. Similarly, the i th column of $\mathbf{W}'$ is given by $\mathbf{w}'_i = \frac{1}{|h(\mathbf{u}_i)|} \sum_{j \in h(\mathbf{u}_i)} \mathbf{w}_j$. Since $\mathbf{w}_j = \mathbf{u}_i$ when $j \in h(\mathbf{u}_i)$, we have $\mathbf{w}'_i = \mathbf{u}_i$, which shows that $\mathbf{W}'$ is just $\mathbf{W}$ with the duplicate columns removed. Using these definitions, we show that the transformation is lossless:

$$\mathbf{W}\mathbf{x} = \sum_{i=1}^n \mathbf{w}_i x_i = \sum_{i=1}^p \mathbf{u}_i \sum_{j \in h(\mathbf{u}_i)} x_j = \sum_{i=1}^p \mathbf{w}'_i x'_i = \mathbf{W}'\mathbf{x}'. \quad \square$$

As noted in Example 8.1, not all workloads allow for reduction (in some cases, the $\mathbf{P}$ matrix computed above is the identity). But others may allow a significant reduction, which improves the efficiency of subsequent operators. Less obvious is that workload-based data reduction would impact accuracy. In fact, many query selection methods from existing work depend implicitly on the representation of the data in vector form, and these approaches may be improved by domain reduction. In Section 8.4, we measure the impact of this transform on accuracy and efficiency.

We show next that this reduction does not hurt accuracy: For any selected set of measurement queries, their reduction will provide lower error after transformation.

THEOREM 8.4. *Given a workload $\mathbf{W}$ and data vector $\mathbf{x}$, let $\mathbf{M}$ be any query matrix that answers $\mathbf{W}$. Then if $\mathbf{q}' = \mathbf{q}\mathbf{P}^+$ is a reduced query and $\mathbf{M}' = \mathbf{M}\mathbf{P}^+$ is the query matrix on the reduced domain, $\text{Error}_{\mathbf{q}'}(\mathbf{M}') \leq \text{Error}_{\mathbf{q}}(\mathbf{M})$ for all $\mathbf{q} \in \mathbf{W}$.*

PROOF. We use the definition of squared error from Reference [32], which shows that $\forall \mathbf{q} \in \mathbf{W}$, $Error_{\mathbf{q}}(\mathbf{M}) \propto \|\mathbf{M}\|_1^2 \|\mathbf{qM}^+\|_2^2$ as long as $\mathbf{M}$ supports $\mathbf{W}$. Let $\hat{\mathbf{m}}_j$ and $\mathbf{m}_j$ denote the j th column of $\hat{\mathbf{M}}$ and $\mathbf{M}$ respectively. First, we show that $\|\hat{\mathbf{M}}\|_1 \leq \|\mathbf{M}\|_1$:

$$\begin{aligned} \|\hat{\mathbf{M}}\|_1 &= \max_{1 \leq i \leq p} \|\hat{\mathbf{m}}_i\|_1 \\ &= \max_{1 \leq i \leq p} \left\| \frac{1}{|h(\mathbf{u}_i)|} \sum_{j \in h(\mathbf{u}_i)} \mathbf{m}_j \right\|_1 \\ &\leq \max_{1 \leq i \leq p} \frac{1}{|h(\mathbf{u}_i)|} \sum_{j \in h(\mathbf{u}_i)} \|\mathbf{m}_j\|_1 \\ &\leq \max_{1 \leq i \leq p} \max_{j \in h(\mathbf{u}_i)} \|\mathbf{m}_j\|_1 \\ &= \max_{1 \leq i \leq n} \|\mathbf{m}_i\|_1 \\ &= \|\mathbf{M}\|_1, \end{aligned}$$

where $h(\mathbf{u})$ and $\mathbf{u}_i$ are defined in Definition 8.2. Now we show that $\|\mathbf{qM}^+\|_2 \leq \|\hat{\mathbf{q}}\hat{\mathbf{M}}^+\|_2$. Observe that it is possible to write $\mathbf{q}$ as a linear combination of the rows of $\mathbf{M}$, since $\mathbf{M}$ supports $\mathbf{W}$. That is, there exists a $\mathbf{z}$ satisfying $\mathbf{zM} = \mathbf{q}$. In general, there may be infinitely many solutions to this linear system, but $\mathbf{z} = \mathbf{qM}^+$ is the *minimum-norm* solution [40]. On the reduced domain, we also know there exists a $\hat{\mathbf{z}}$ satisfying $\hat{\mathbf{z}}\hat{\mathbf{M}} = \hat{\mathbf{q}}$, or equivalently $\hat{\mathbf{z}}\hat{\mathbf{M}}\mathbf{P}^+ = \hat{\mathbf{q}}\mathbf{P}^+$. By making the substitution $\mathbf{zM} = \mathbf{q}$, it is easy to see that $\hat{\mathbf{z}} = \mathbf{z}$ is one solution to this linear system. The minimum norm solution to this linear system is $\hat{\mathbf{z}} = \hat{\mathbf{q}}\hat{\mathbf{M}}^+$, which implies $\|\hat{\mathbf{z}}\|_2 \leq \|\mathbf{z}\|_2$. This shows that $\|\hat{\mathbf{q}}\hat{\mathbf{M}}^+\|_2 \leq \|\mathbf{qM}^+\|_2$, and it immediately follows that $Error_{\hat{\mathbf{q}}}(\hat{\mathbf{M}}) \leq Error_{\mathbf{q}}(\mathbf{M})$ as desired. $\square$

8.2 Computing the Partition

The computation of the partition $\mathbf{P}$ in Definition 8.2 is conceptually straightforward: It simply requires finding the unique columns of $\mathbf{W}$ and grouping them. Finding the unique columns of $\mathbf{W}$ exactly by inspecting the entries of $\mathbf{W}$ requires an explicit matrix in dense form or materializing an implicit matrix. Algorithm 4 provides an efficient method for finding the column groupings that does not require a explicit matrix representation, relying instead only on the primitive methods of transpose and matrix-vector product. This approach is highly scalable and is compatible with the implicit matrix representations of the workload discussed in Section 7.

By grouping the elements of $\mathbf{h}$ (line 6), we recover the column groupings of $\mathbf{W}$, because if $\mathbf{w}_i = \mathbf{w}_j$, then $h_i = h_j$, and if $\mathbf{w}_i \neq \mathbf{w}_j$, then $P(h_i = h_j) = 0$, since h_i and h_j are continuous random variables. While algorithm Algorithm 4 is a randomized algorithm, it returns the correct result almost surely. The probability of incorrectly grouping two different columns is approximately 10^{-16} with a 64-bit floating point representation, but if needed, then we can repeat the procedure k times until the probability of failure ($\sim 10^{-16k}$) is vanishingly small.

9 CASE STUDIES: εKTELO IN ACTION

In this section, we put εKTELO into action by developing new algorithms. First, we show that it is easy to re-design and improve existing algorithms by combining operators in new ways. In particular, we develop a variant of the MWEM algorithm with significantly improved accuracy. Then, we use εKTELO to tackle two practical use-cases, constructing new plans that offer state-of-the-art accuracy. We evaluate all of the proposed plan in Section 10.

ALGORITHM 4: An algorithm for workload-based data reduction

```

1: procedure COMPUTE REDUCTION MATRIX( $W$ )
2:   Input:  $m \times n$  matrix  $W$ 
3:   Output:  $p \times n$  matrix  $P$  where  $p \leq n$ 
4:   set  $\mathbf{v}$  = vector of  $m$  samples from Uniform(0, 1)  $\triangleright 1 \times m$ 
5:   compute  $\mathbf{h} = W^T \mathbf{v}$   $\triangleright 1 \times n$ 
6:   let  $G = g_1, \dots, g_p$  be groups of common values in  $\mathbf{h}$ 
7:   initialize matrix  $P$  with zeros  $\triangleright p \times n$ 
8:   for  $g_i$  in  $G$  do
9:     set row  $i$  of  $P$  to 1 in each position of  $g_i$ 
10:  end for
11:  return  $P$ 
12: end procedure

```

9.1 Recombination of Operators to Improve MWEM

Using ϵ KTELO, we design new variants of the well-known *Multiplicative Weights Exponential Mechanism* (MWEM) [21] algorithm. MWEM repeatedly derives the worst-approximated workload query with respect to its current estimate of the data, then measures the selected query, and uses the multiplicative weights update rule to refine its estimate, often along with any past measurements taken. This repeats a number of times, determined by an input parameter.

When viewed as a plan in ϵ KTELO, a deficiency of MWEM becomes apparent. Its query selection operator selects a *single* query to measure, whereas most query selection operators select a set of queries such that the queries in the set measure disjoint partitions of the data. By the parallel composition property of differential privacy, measuring the entire set has the same privacy cost as asking any single query from the set. This means that MWEM could be measuring more than a single query per round (with no additional consumption of the privacy budget). To exploit this opportunity, we designed an augmented query selection operator that adds to the worst-approximated query by attempting to build a binary hierarchical set of queries over the rounds of the algorithm. In round one, it adds any unit length queries that do not intersect with the selected query. In round two, it adds length two queries, and so on.

Adding more measurements to MWEM has an undesirable side effect on runtime, however. Because it measures a much larger number of queries across rounds of the algorithm and the runtime of multiplicative weights inference scales with the number of measured queries, inference can be considerably slower. Thus, we also use replace it with a version of least squares with a non-negativity constraint (NNLS) and incorporate a high-confidence estimate of the total that is assumed by MWEM.

In total, we consider three MWEM variants: an alternative query selection operator (PLAN #18), an alternative inference operator (PLAN #19), and the addition of both alternative operators (PLAN #20). These are shown in Figure 2 and evaluated in Section 10.

9.2 Census Case Study

The U.S. Census Bureau collects data about U.S. citizens and releases a wide variety of tabulations describing the demographic properties of individuals. We consider a subset of the (publicly released) March 2000 Current Population Survey. The data report on 49,436 heads-of-household describing their income, age (in years), race, marital status, and gender. We divide Income into 5,000 uniform ranges from (0, 750000), age in five uniform ranges from (0, 100), and there are seven, four, and two possible values for status, race, and gender.

We author differentially private plans for answering a workload of queries similar to Census tabulations. This is challenging, because the data domain is large and involves multiple dimensions. The workloads we consider are (a) the Identity workload (or counts on the full domain of 1.4M cells), (b) a workload of all two-way marginals (age $\times$ gender, race $\times$ status, and so on), and (c) a workload suggested by U.S. Census Bureau staff: Prefix(Income), which consists of all counting queries of the form (income $\in (0, i_{high})$, age = a , marital = m , race = r , gender = g), where $(0, i_{high})$ is an income range, and a, m, r, g may be values from their respective domains or $\langle any \rangle$.

There are few existing algorithms suitable for this task. We were unable to run the DAWA [29] algorithm directly on such a large domain. In addition, it was designed for 1d- and 2d- inputs. One of the few algorithms designed to scale to high dimensions is PrivBayes [55]. While not a workload-adaptive algorithm, PrivBayes generates synthetic data that can support the census workloads above. We use PrivBayes as a baseline, and we use εKTELO to construct three new plans composed of operators in our library. The proposed plans are HB-STRIPED (PLAN #15), DAWA-STRIPED (PLAN #14), and PRIVBAYESLS (PLAN #17). The first two “striped” plans showcase the ability to adapt lower-dimensional techniques to a higher-dimensional problem avoiding scalability issues. The third plan considers improving on PrivBayes by changing its inference step.

Both HB-STRIPED and DAWA-STRIPED use the same plan structure: First, they partition the full domain, then they execute subplans to select measurements for each partition, and, last, given the measurement answers, they perform inference on the full domain and answer the workload queries. The partitioning of the initial step is done as follows: Given a high-dimensional dataset with N attributes and an attribute A of that domain, our partitions are parallel “stripes” of that domain for each fixed value of the rest of the $N - 1$ attributes, so that the measurements are essentially the one-dimensional histograms resulting from each stripe. In the case of HB-STRIPED (fully described in Algorithm 5), the subplan executed on each partition is the HB algorithm [44], which builds an optimized hierarchical set of queries, while in the case of the DAWA-STRIPED the subplan is DAWA algorithm [29]. Note that while the data-independent nature of the HB subplan means that all the measurements from each stripe are the same, that is not the case with DAWA, which potentially selects different measurement queries for each stripe, depending on the local vector it sees. For our experiments, the attribute chosen was Income, and for DAWA-STRIPED we set the DAWA parameter ρ to 0.25.

ALGORITHM 5: HB-STRIPED

```

1:  $D \leftarrow \text{PROTECTED}(\text{source\_uri})$  ▷ Init
2:  $\mathbf{x} \leftarrow \text{T-VECTORIZE}(D)$  ▷ Transform
3:  $\mathbf{R} \leftarrow \text{StripePartition}(\text{Attr})$  ▷ Partition Selection
4:  $\mathbf{x}_R \leftarrow \text{V-SPLITBYPARTITION}(\mathbf{x}, \mathbf{R})$ 
5:  $\mathbf{M} \leftarrow \emptyset$ 
6:  $\mathbf{y} \leftarrow \emptyset$ 
7: for  $\mathbf{x}' \in \mathbf{x}_R$  do
8:    $\mathbf{M} \leftarrow \mathbf{M} \cup \text{HB}(\mathbf{x}')$  ▷ Query Selection
9:    $\mathbf{y} \leftarrow \mathbf{y} \cup \text{VECLAPLACE}(\mathbf{x}', \mathbf{M}, \epsilon)$  ▷ Query
10: end for
11:  $\hat{\mathbf{x}} \leftarrow \text{LS}(\mathbf{M}, \mathbf{y})$ 
12: return  $\hat{\mathbf{x}}$  ▷ Output

```

Selection in HB-STRIPED’s subplans are data independent, unlike in DAWA-STRIPED, so the exact same set of measurements will be selected on each partition. As introduced in Section 7.4, this set of measurements can be represented compactly as a Kronecker product. So we introduce

a new selection operator $\text{STRIPE}(\text{ATTR})$ where a global measurement is composed by constructing the Kronecker product of HB measurements on the stripe dimension and Identity matrices on other dimensions. HB-STRIPED_KRON (PLAN #16) is a sequence starting with the new SS selection operator, followed by Laplace measurement and LS inference. The complete plan is shown in Algorithm 6. This non-iterative alternative implementation is more efficient, and we compare the efficiency and scalability of the two implementations in Section 10.1.1.

ALGORITHM 6: HB-STRIPED_KRON

```

1:  $D \leftarrow \text{PROTECTED}(\text{source\_uri})$  ▷ Init
2:  $\mathbf{x} \leftarrow \text{T-VECTORIZE}(D)$  ▷ Transform
3:  $\mathbf{M} \leftarrow \text{STRIPESELECT}(\text{Attr})$  ▷ Query Selection
4:  $\mathbf{y} \leftarrow \text{VecLAPLACE}(\mathbf{x}, \mathbf{M}, \epsilon_3)$  ▷ Query
5:  $\hat{\mathbf{x}} \leftarrow \text{LS}(\mathbf{M}, \mathbf{y})$ 
6: return  $\hat{\mathbf{x}}$  ▷ Output

```

Our final plan is a variant of PRIVBAYES in which we replace the original inference method with least squares, retaining the original PRIVBAYES query selection and query steps. We call this algorithm PRIVBAYESLS, and it is fully described in Algorithm 7.

ALGORITHM 7: PRIVBAYESLS

```

1:  $D \leftarrow \text{PROTECTED}(\text{source\_uri})$  ▷ Init
2:  $\mathbf{x} \leftarrow \text{T-VECTORIZE}(D)$  ▷ Transform
3:  $\mathbf{M} \leftarrow \text{PBSELECT}(\mathbf{x}, \epsilon_2)$  ▷ Query Selection
4:  $\mathbf{y} \leftarrow \text{VecLAPLACE}(\mathbf{x}, \mathbf{M}, \epsilon_3)$  ▷ Query
5:  $\hat{\mathbf{x}} \leftarrow \text{LS}(\mathbf{M}, \mathbf{y})$  ▷ Inference
6: return  $\mathbf{W} \cdot \hat{\mathbf{x}}$  ▷ Output

```

We evaluate the error incurred by these plans in Section 10.2.2 and show that the best of our plans outperforms the state-of-the-art PRIVBAYES by at least 10 $\times$ in terms of error.

9.3 Naive Bayes Case Study

We also demonstrate how ϵKTELO can be used for constructing a Naive Bayes classifier. To learn a NaiveBayes classifier that predicts a binary label attribute Y using predictor variables $(X_1, \dots, X_k)$ requires computing $2k+1$ 1d histograms: a histogram on Y , histogram on each X_i conditioned on each value on Y . We design ϵKTELO plans to compute this workload of $2k+1$ histograms and use them to fit the classifier under the Multinomial statistical model [26].

We develop two new plans and compare them to two plans that correspond to algorithms considered in prior work. WORKLOAD represents the $2k+1$ histograms as a matrix and uses VECTOR LAPLACE to estimate the histogram counts. This corresponds to a technique proposed in the literature [11]. The other baseline is IDENTITY (Plan 1): It estimates all point queries in the contingency table defined by the attributes, adds noise to it, and marginalizes the noisy contingency table to compute the histograms.

The first new plan is WORKLOADLS, which runs WORKLOAD followed by a least squares inference operator, which for this specific workload would make all histograms have consistent totals. Our second plan is called SELECTLS (fully described in Algorithm 8) and selects a different algorithm (subplan) for estimating each of the histograms. SELECTLS first runs $2k+1$ domain reductions to compute $2k+1$ vectors, one for each histogram. Then, for each vector, SELECTLS uses a conditional statement to select between two subplans: If the vector size is less than 80, then IDENTITY is chosen,

ALGORITHM 8: SELECTLS

```

1:  $D \leftarrow \text{PROTECTED}(\text{source\_uri})$  ▷ Init
2:  $\mathbf{x} \leftarrow \text{T-VECTORIZE}(D)$  ▷ Transform
3:  $\mathbf{R} \leftarrow \text{MARGREDUCTION}(\mathbf{x}, \text{Att})$  ▷ Partition Selection
4:  $\mathbf{M} \leftarrow \emptyset, \mathbf{y} \leftarrow \emptyset$ 
5: for  $i = 1 : k$  do ▷ Iterate over Dimensions
6:    $\mathbf{x}' \leftarrow \text{V-REDUCEBYPARTITION}(\mathbf{x}, \mathbf{R}_i)$ 
7:   if  $\text{DomainSize}_i > 80$  then
8:      $\mathbf{R}' \leftarrow \text{RDAWA}(\mathbf{x}', \epsilon_1/k)$  ▷ Partition Selection
9:      $\mathbf{x}'_R \leftarrow \text{V-REDUCEBYPARTITION}(\mathbf{x}, \mathbf{R}')$ 
10:     $\mathbf{M} \leftarrow \text{GREEDYH}(\mathbf{x}'_R)$  ▷ Query Selection
11:     $\mathbf{y} \leftarrow \mathbf{y} \cup \text{VECLAPLACE}(\mathbf{x}'_R, \mathbf{M}, \epsilon_2/k)$  ▷ Query
12:   else
13:      $\mathbf{M} \leftarrow \text{MU IDENTITY}(\mathbf{x}')$  ▷ Query Selection
14:      $\mathbf{y} \leftarrow \mathbf{y} \cup \text{VECLAPLACE}(\mathbf{x}', \mathbf{M}, \epsilon/k)$  ▷ Query
15:   end if
16:    $\mathbf{x} \leftarrow \text{V-REDUCEBYPARTITION}(\mathbf{x}', \mathbf{R}_i)$  ▷ Domain Expansion
17: end for
18:  $\hat{\mathbf{x}} \leftarrow \text{LS}(\mathbf{M}, \mathbf{y})$ 
19: return  $\hat{\mathbf{x}}$  ▷ Output

```

else a subplan that runs DAWA partition selection followed by IDENTITY is chosen. We combine the answers from all subplans and use least squares inference jointly on all measurements. The inputs to the inference operator are the noisy answers and the workload of effective queries on the full domain. In Section 10.2.3, we show that our new plans not only outperform existing plans but also approach the accuracy of the non-private classifier in some cases.

10 EXPERIMENTAL EVALUATION

Our prototype implementation of ϵ KTELO, including all algorithms and variants used below, consists of 7.9k lines of code: Twenty-five percent is the framework itself, 46% consist of operator implementations, 14% consist of definitions of plans used in our experiments and the remaining 15% are tests and examples provided for the users. We conduct experimental evaluation on an Amazon EC2 2xLarge instance with 32 GB of RAM running Ubuntu 14.04.

In this section: We first demonstrate the performance improvement of implicit matrices by comparing efficiency and scalability of plans and inference operators with previous implementations using dense or sparse matrices from Reference [54]. Then we report the results of using ϵ KTELO in the case studies of Section 9. Finally, we evaluate the impact of workload-based domain reduction introduced in Section 8.1.

10.1 Performance Improvements of Implicit Matrices

As discussed in Section 7, most ϵ KTELO operators involve performing operations on matrices, which can be implemented using several different representations, including the dense and sparse representations from Reference [54] as well as the novel implicit representations described earlier. All of them are lossless representations of the underlying matrix, so the choice of implementation does not influence plan accuracy. However, it could impact the efficiency and scalability.

In this section, we compare these alternative implementations. We first evaluate how the choice of matrix implementation impacts scalability and efficiency of plans. Then we have a focused experiment on a key matrix operation: inference.

Fig. 3. Plan execution time with different implementations of measurement matrices with the Identity workload. Implicit represents plans implemented with the implicit matrices introduced in Section 7. Dense and sparse represent previous implementations [54], and basic sparse is an alternative to Kronecker product (a type of implicit matrix) materialized as sparse matrices. Results show the new implicit representation can increase scalability by a factor up to 1,000 $\times$.

10.1.1 Scalability and Efficiency of Plans. To understand the impact of different physical implementations for plans in Figure 2, we compare the runtime of plans using implicit measurement matrices with previous implementations from Reference [54], which can use either dense or sparse matrices. For the plans HDMM and HB-STRIPED_KRON, which contain operators that were not supported in the previous implementation, we make the comparison by converting implicit matrices to their sparse and dense representations. We measure the average end-to-end execution time over five random trials for three implementations of each plan along increasing domain sizes. We stop any execution when it runs for more than 1,000 s.

Figure 3(a) shows runtime for low-dimensional plans. All plans are applied on two-dimensional domain square domains, except for DAWA and GREEDY-H, which are designed for one-dimensional domains. PLAN #19 MWEM VARIANT B is omitted, because it timed out even at the smallest domain

Fig. 4. For a given computation time, the proposed iterative and implicit inference methods permit scaling to data vector sizes as much as 1000× larger than previous techniques using direct approaches and dense matrices. Dense and sparse implementations are from Reference [54] and tree-based is from Reference [23].

we tested here. The results show that for most plans, the implicit implementation has the best scalability. Also, looking at a fixed domain size, the implicit representation usually leads to faster runtime than its dense and sparse counterparts.

The performance improvement is most pronounced with plans HB, QUADTREE, and UNIFORM-GRID, where implicit representation can scale to domains larger by a factor of 1000×. These algorithms construct hierarchical/grid-based measurement matrices and can be represented as *Range Queries*, a special instance of the implicit matrices. As discussed in Section 7.5, this representation is compact and supports faster matrix-vector products. There are few cases where the difference between implementations is less pronounced. DAWA and GREEDY-H share the same special selection subroutine that needs to materialize the matrix. ADAPTIVEGRID has a plan that requires iterating through a potentially large number of partitions, and this step appears to dominate the runtime.

Results for high-dimensional plans are shown in Figure 3(b). For the first three plans, sparse and implicit representations exhibit similar performance and scale to domains at least 10× larger than using dense. For the last plan, HB-STRIPED_KRON, recall from Section 9.2 that this plan is an alternative way of expressing the same algorithm as the HB-STRIPED plan, but instead of partitioning the data, it uses Kronecker products to express queries compactly in terms of submatrices. By comparing adjacent figures, we can see the approach based on Kronecker products allows plans to scale to at least 10× larger domains across implementations of the submatrices. As another comparison point that illustrates the benefits of Kronecker products, in the last figure, we include “Basic sparse,” an alternative implementation of the HB-STRIPED_KRON plan, where the query Kronecker product matrix is replaced with a materialized sparse matrix over the full domain.

10.1.2 Scalability of Inference. Inference is one of the most computation-intensive operators in εKTELO, especially for large domains resulting from multi-dimensional data. Next, we show the impact of implementation choices on the scalability of inference. Figure 4 shows the computation time for running our main inference operators (LS and NNLS) as a function of data vector size.

Recall that the methods described in Section 7.6 provide efficiency improvements by using iterative solution strategies (*iterative* instead of *direct* in the figure) and exploiting sparsity in the measurement matrix (*sparse* or *implicit* as opposed to *dense* in the figure). For this experiment, we fix the measured query set to consist of binary hierarchical measurements [23]. Figure 4 shows that using sparse matrices and iterative methods allow inference to scale to data vectors consisting of millions of counts on a single machine in less than a minute. The use of implicit matrices

Table 4. For Three New Algorithms, (b), (c), and (d), the Multiplicative Factors by Which Error Is Improved Presented as (min, mean, max) over Datasets

	MWEM Variants		ERROR IMPROVEMENT			RUNTIME
	Query Selection	Inference	min	mean	max	
(a)	worst-approx	MW	1	1	1	1
(b)	worst-approx + H2	MW	1.03	2.80	7.93	354.9
(c)	worst-approx	NNLS, known total	0.78	1.08	1.54	1.0
(d)	worst-approx + H2	NNLS, known total	0.89	2.64	8.13	9.0

For runtime, the mean is shown, normalized to the runtime of standard MWEM (1D, $n = 4096$, $W = \text{RandomRange}(1000)$, $\epsilon = 0.1$).

permits additional scale-up for both LEASTSQUARES and NNLS. We also compare against the inference method introduced by Hay et al., denoted “Tree-based” in the figure. It is an algorithm that is logically equivalent to LEASTSQUARES but specialized for hierarchically structured measurements. The general-purpose LEASTSQUARES implementation is able to scale to much larger domains.

10.2 Case Studies

10.2.1 MWEM: Improved Query Selection and Inference. We evaluate the three new plans described in Section 9.1 that were variants of ϵ KTELO plan for the MWEM [21] algorithm. Recall that the variants were achieved by replacing key operators in the MWEM plan. These algorithms are data-dependent algorithms so we evaluate them over a diverse collection of 10 datasets taken from DPBench [22]. The results are shown in Table 4.

The performance of the first variant, line (b), shows that the augmenting query selection with H2 can significantly improve error: by a factor of 2.8 on average (over various input datasets) and by as much as a factor of 7.9. (Error and runtime measures are normalized to the values for the original MWEM; min/mean/max error values represent variation across datasets.) Unfortunately, this operator substitution has a considerable impact on performance: The added queries slow down by a factor of more than 300. But combining augmented query selection with NNLS inference, line (d), reduces runtime significantly: It is still slower than the original MWEM algorithm, but by only a factor of 9. Using the original MWEM query selection with NNLS inference, line (c), has largely equivalent error and runtime to the original MWEM. Thus, the performance gains of NNLS inference over MW appear to be most pronounced when the number of measured queries is large.

10.2.2 Census Data Analysis. In this section, we compare the ϵ KTELO plans proposed in Section 9.2, measuring their effectiveness in computing workloads inspired by Census tabulations. We compare our three new plans PRIVBAYESLS, Hb-STRIPED, and DAWA-STRIPED with baseline algorithms IDENTITY (PLAN #1 in Figure 2) and PRIVBAYES, our ϵ KTELO reimplement of a state-of-the-art algorithm for high-dimensional data [55].

Table 5 presents the results for each workload. We use scaled, per-query L2 error to measure accuracy. First, we find that PRIVBAYES performs worse than IDENTITY on all workloads. Interestingly, on Identity and two-way marginal workloads, it is improved by our new plan PRIVBAYESLS that replaces its inference step with least squares. PRIVBAYES may be more suitable to input data with higher correlations between the attributes. Second, our striped plans Hb-STRIPED and DAWA-STRIPED offer significant improvements in error. DAWA-STRIPED is the best performer: The data-dependent nature of DAWA exploits uniform regions in the partitioned data vectors. This shows the benefit from ϵ KTELO in allowing algorithm idioms designed for lower-dimensional data to be adapted to high-dimensional problems.

Table 5. Results on Census Data; Domain size 1,400,000, and scale of error is indicated under each workload.

Algorithm	Workload		
	Identity (1e-9)	Two-way Marg. (1e-7)	Prefix (Income) (1e-7)
IDENTITY	24.18	12.04	18.97
PRIVBAYES	76.93	65.31	28.70
PRIVBAYESLS	5.86	13.29	36.81
HB-STRIPED	70.31	21.91	4.13
DAWA-STRIPED	3.43	1.96	2.50

Fig. 5. New εKTELO plans WORKLOADLS and SELECTLS result in NaiveBayes classifiers with lower error than plans that correspond to algorithms from prior work and approach the accuracy of a non-private classifier for various ϵ values.

10.2.3 Naive Bayes Classification. We evaluate the performance of the Naive Bayes classifier on *Credit Default* [52], a credit card clients dataset that we use to predict whether a client will default on their payment or not. The data consists of 30k tuples and 24 attributes from which one is the target binary variable “Default” and the rest are the predictive variables. We used the predictive variables $X_3 - X_6$ for a total combined domain size of 17,248.

In our experiments, we measure the average area under the curve (AUC) of the receiver operating characteristic curve across a 10-fold cross validation test. The AUC measures the probability that a randomly chosen positive instance will be ranked higher than a randomly chosen negative instance. We repeat this process 10 times (for a total of 100 unique testing/training splits) to account for the randomness of the differentially private algorithms and report the {25, 50, 75}-percentiles of the average AUC. As a baseline we show the majority classifier, which always predicts the majority class of the training data and also show the unperturbed classifier as an upper bound for the utility of our algorithms.

In Figure 5, we report our findings: each group of bars corresponds to a different ϵ value and each bar shows the median value of the AUC for an algorithm. For each DP algorithm, we also plot the error bars at the 25 and 75 percentiles. The dotted line is plotted at 0.5067 and shows the AUC of the majority classifier. The continuous red line is the performance of the non-private classifier (Unperturbed). For larger ϵ values, we see that our plans significantly outperform the baseline and reach AUC levels close to the unperturbed. As ϵ decreases, the quality of the private classifiers degrades and for $\epsilon = 10^{-3}$ the noise added to the empirical distributions drowns the signal and the AUC of the private classifiers reach 0.5, which is the performance of a random classifier. Our

Table 6. Runtime (sec) and Error Improvements Resulting from Workload-based Domain Reduction

Algorithm	Original Domain		Reduced Domain		Factor Improved	
	Error/Runtime		Error/Runtime		Error/Runtime	
AHP	1.68e-5	777.10	1.30e-5	145.00	1.29	5.36
DAWA	1.06e-5	0.23	1.07e-5	0.25	0.99	0.92
IDENTITY	4.74e-5	0.66	1.64e-5	0.90	2.89	0.73
HB	3.20e-5	0.05	2.38e-5	0.08	1.34	0.62

W = RandomRange, small ranges. Original domain size: AHP (128,128), DAWA 4096, Identity (256,256), HB 4096.

plan WORKLOADLS is essentially the algorithm in Reference [12] with an extra inference operator. This shows that the addition of an extra operator to a previous solution significantly increases its performance.

10.3 Workload-driven Data Reduction

Next, we evaluate the impact of workload-driven data reduction, as described in Section 8.1. For selected algorithms, Table 6 shows that performing workload-driven data reduction improves *both* error and runtime, almost universally.

The biggest improvement in error (a factor of 2.89) is witnessed for the IDENTITY algorithm. Without workload-driven reduction, groups of elements of the domain are estimated independently even though the workload only uses the total of the group. After reduction, the sum of the group of elements is estimated and will have lower variance than the sum of independent measurements.

The biggest improvement in runtime occurs for the AHP algorithm. This algorithm has an expensive clustering step, performed on each element of the data vector. Workload-driven reduction reduces the cost of this step, since it is performed on a smaller data vector. It also tends to improve error, because higher-quality clusters are found on the reduced data representation.

10.4 Summary of Findings

The experiments evaluate the accuracy, scalability, and efficiency of ϵ KTELO. The study of scalability and efficiency found that implicit matrix representation can lead to huge performance gains, increasing scalability by a factor of 1000 $\times$ in some cases. The generalized inference implementation scales well and outperforms specialized algorithms. The case studies show that ϵ KTELO can lead to more accurate algorithms with relatively little effort from the programmer: The MWEM algorithm can be improved significantly by replacing a few key operators; for the Census and Naive Bayes case studies, ϵ KTELO can be used to design novel algorithms from existing building blocks, offering state-of-the-art error rates. Finally, the evaluation shows the workload-driven data reduction improves accuracy and runtime, almost universally, so that it can be added to all workload-based plans with little cost and significant potential for gains.

11 RELATED WORK

ϵ KTELO was first described by the authors in a 2018 conference paper [54]. The open source codebase for ϵ KTELO is publicly available [6], and the original experimental results [54] have passed the SIGMOD reproducibility review. This manuscript extends the conference paper [54] by providing a unified and improved approach to the efficient representation of matrix objects. The new matrix representations are fundamental to ϵ KTELO plans, as they are used to represent workload queries, measurement queries, and partitions. This innovation impacts many aspects of ϵ KTELO, including

inference and workload-based partition selection. Existing plans were re-implemented using these new matrix techniques to measure their impact and were demonstrated to allow ϵ KTELO plans to scale to far larger data vectors than previously possible.

The implicit matrix representation furthers the notion that there can be a separation between the matrix as a logical representation and its physical implementation. Implementation choices can be transparent to plan authors and open up interesting directions for developing highly optimized implementations.

A number of languages and programming frameworks have been proposed to make it easier for users to write private programs [15, 35, 42, 46]. The *Privacy Integrated Queries* (PINQ) platform began this line of work and is an important foundation for ϵ KTELO. We use the fundamentals of PINQ to ensure that plans implemented in ϵ KTELO are differentially private. In particular, we adapt and extend a formal model of a subset of PINQ features, called Featherweight PINQ [15], to show that plans written using ϵ KTELO operators satisfy differential privacy. Our extension adds support for the partition operator, a valuable operator for designing complex plans.

Additionally, there is a growing literature on formal verification tools that prove that an algorithm satisfies differential privacy [9, 19, 53]. For instance, LightDP [53] is a simple imperative language in which differentially private programs can be written. LightDP allows for verification of sophisticated differentially private algorithms with little manual effort. LightDP's goal is orthogonal to that of ϵ KTELO: It simplifies proofs of privacy, while ϵ KTELO's goal is to simplify the design of algorithms that achieve high accuracy. Nevertheless, an interesting future direction would be to implement ϵ KTELO operators in LightDP to simplify both problems of verifying privacy and achieving high utility.

Concurrently with Reference [54], Kellaris et al. [25] observed that algorithms for single-dimensional histogram tasks share subroutines that perform common functions. The authors compare a number of existing algorithms along with new variants formed by combining subroutines, empirically evaluating tradeoffs between accuracy and efficiency. The focus is exclusively one-dimensional tasks.

The use of inference in differentially private algorithm design is not new [8, 23, 50] and is used in various guises throughout recent work [7, 12, 28, 29, 32, 33, 41, 51, 57]. Proserpio et al. [41] propose a general-purpose inference engine based on MCMC that leverages properties of its operators to offset the otherwise high time/space cost of this form of inference. Our work is complementary in that we focus on a different kind of inference (based on least squares) in part because it is used, often implicitly, in many published techniques. A deeper investigation of alternative inference strategies is a compelling research direction.

Our use of implicit matrices was inspired by their use in Reference [33], where Kronecker products were used extensively to represent high-dimensional workloads and measurements. ϵ KTELO implicit matrices extend and generalize those matrix constructions. Techniques that use measurements based on wavelets [8] (for range query workloads) and Fourier basis queries [51] (for marginals) are examples in which measurement and inference is performed without materialization of a matrix, and so they could be seen as implicit methods. Hierarchical query sets [23, 44] also admit inference methods that do not require materialization of a matrix. Each of these examples relies on the special structure of the query sets they use to achieve scalability. Our approach in ϵ KTELO is more flexible: It allows the plan author to focus on *what* to measure rather than how to measure it and how to perform inference efficiently.

ϵ KTELO provides an execution framework for privacy algorithms but does not perform plan-level optimization. The matrix mechanism [32] (and more recently the high-dimensional matrix mechanism [33]) formulates an optimization problem that corresponds to query selection in ϵ KTELO. The mechanism then estimates the selected queries and applies least squares inference. This can be

seen as a kind of optimization but in a limited plan space that admits only data-independent plans. Recent work [26] examines the problem of algorithm selection—selecting the best algorithm for a given private dataset and task—and proposes a meta-algorithm, Pythia, capable of choosing among a set of “black box” algorithms. In contrast, ϵ KTELO takes a “white box” approach, decomposing existing algorithms into modular operators and allowing plan authors to design new algorithms. Pythia could be adapted to automatically select operators in ϵ KTELO and, in fact, Pythia could be implemented as an ϵ KTELO plan.

As noted above, our efficiency and scalability efforts have so far been focused on a centralized setting where ϵ KTELO plans are executed on a single machine and data vectors fit in memory. As such, ϵ KTELO currently makes use Python’s pandas and SciPy modules for relational and matrix processing, respectively. In the future, it may be beneficial for ϵ KTELO to use a distributed data processing platform such as Apache’s Hive [1], Accumulo [3], or Spark [4], or to consider the innovations of academic projects such as Weld [38], LaraDB [24], SPOOF [16], and Samsara [47], many of which support parallelized matrix operations.

However, these platforms do not provide easy scalability solutions for ϵ KTELO and are not substitutes for exploiting the special matrix structure present in ϵ KTELO operators, which will be necessary and beneficial in any execution platform that is adopted. For example, no computing platform is so efficient that it obviates the need to reduce a matrix from 56 GB to less than 100 bytes, as in Example 7.3. Please see Section 12.2 for further discussion of scalability.

12 DISCUSSION AND LIMITATIONS

In this section, we discuss current limitations of ϵ KTELO and opportunities for future work that will extend its capabilities.

12.1 More Complex Queries and Tasks

ϵ KTELO is extensible and we hope to expand the classes of tasks that can be supported. ϵ KTELO is currently focused on statistical queries on a single table, or those tasks, such as naive Bayes classification, which can be directly supported by such statistics. Answering more expressive aggregate queries, for example, those expressible as SQL queries over multi-relational schemas, is possible using the general framework of ϵ KTELO but will require a number of extensions.

In follow-up work a subset of the authors of the current paper have proposed an architecture for a differentially private SQL query-answering system [27]. This architecture builds on the fundamental structure of ϵ KTELO but requires a number of extensions. First, it requires support for specifying and enforcing more complex privacy policies over multiple input relations. Second, it requires extended relational transformations, including joins and a truncation operator, which can reduce the sensitivity of downstream computations. Third, more advanced stability calculations must accompany these new transformations. Using these extensions, the approach outlined in Reference [27] would generate views over the base relations, each generated by an ϵ KTELO plan consisting of transformation and measurement operators. The final challenge then becomes inference, which would require extensions to achieve global consistency.

Another possible extension for ϵ KTELO is adding support for aggregate functions over continuous and numerical domains. Queries like SUM, STDDEV, and MEAN can be used as building blocks for a wider class of algorithms (e.g., the naive Bayes classifier of Reference [49] for continuous variables). To support those queries, ϵ KTELO will need to know the domain of each attribute involved in these aggregations to correctly compute their respective sensitivities. For instance, given attribute A , with upper bound a_{max} , ϵ KTELO can correctly compute the sensitivity of SUM(A) as a_{max} . Similarly and given domain information, ϵ KTELO can provide upper bounds for the sensitivity of STDDEV, and MEAN. Last, note that support for user-defined queries for which sensitivity

computation is non-trivial is still viable, as long as these queries are accompanied with a vetted sensitivity estimate.

12.2 Scaling to Higher Dimensions

Many of the ϵ KTELO plans described in this article operate on a single materialized data vector, and the current version of ϵ KTELO always keeps this data vector in memory. Recall that a data vector can grow as large as the product of the domain sizes of the attributes in the input table. We have proposed a number of innovations that allow even the most expensive operators (like inference) to scale to vectors as large as 10^8 . But it is nevertheless clear that for high-dimensional data (or even data with a modest number of dimensions, but large attribute domains) plans attempting to fully materialize the data vector will be infeasible.

ϵ KTELO can still be used to define and execute plans in the high-dimensional case. The principal difference from most plans explored in this article is that the plans will operate on *multiple* data vectors, each defined over a subset of the attributes of the input table. Specifically, such a plan would begin by applying multiple PROJECT transformations to the input table, to perform a relational projection onto a subset of the attributes. Then each would be followed by T-VECTORIZE, ultimately resulting in a collection of data vectors. A plan could then apply Query Selection and Query operators to each vector independently, resulting in sets of differentially private estimates, each over a projection of the original data.

In fact, plans of precisely this form were designed in response to the recent *Differential Privacy Synthetic Data Challenge* [5] created by the U.S. National Institute of Standards and Technology (NIST). The challenge required competitors to generate differentially private synthetic data from a source dataset with 98 attributes. The resulting data vector would have approximately 5×10^{205} elements. Not only is this infeasible to represent, but, given that the input data contained just 660 thousand records, measurements offering reasonable utility under differential privacy will tend to be those that aggregate over many dimensions. Among the top four solutions to the final round of the challenge, none measured more than a three-attribute projection. The winning solution constructed approximately 200 such low-dimensional projections, performed query selection, and then post-processing.

The true challenge is indeed post-processing many independently estimated noisy sets of queries. The methods proposed in this article, while substantially expanding the efficiency of inference, cannot support global inference over many data vectors, because they rely on full materialization to relate the measured queries to one another. Some existing techniques have proposed methods for post-processing on extremely large domains by restricting the structure of the noisy queries [21, 45, 55], but they are not general solutions to the problem that can be effectively used within different types of plans. A more general solution to the problem was recently proposed that uses graphical models as a representation tool for the high-dimensional data distribution [34]. Such a technology could be used within ϵ KTELO plans, allowing them to effectively combine evidence about multiple data vectors derived from the same data source.

In summary, while we believe ϵ KTELO is a promising platform for plans involving high-dimensional data, further research and development is required to fully incorporate scalable post-processing methods, as well as to reconsider query selection operators for plans that involve many projections over a single dataset. We note that recent research into (non-private) summarization of high-dimensional data may have a role to play here [37, 39].

12.3 Automated Optimization

The ϵ KTELO system is analogous to a query executor in a relational system: It allows plans to be specified and executed, but it does not create plans in response to a given query or task. Some

operators perform *operator-level* optimization (e.g., the HDMM and Greedy-H query selection operators both automatically adapt to a given input workload). However, a compelling future goal for ϵ KTELO is automatic *plan-level* optimization.

Much like a relational optimizer, we envision adding a component that can explore the plan space implicitly defined by the collection of operators implemented in ϵ KTELO. But optimization here has dual objectives (both accuracy and efficiency must be considered), and, in addition, it may be important to accommodate user-defined accuracy metrics. Further, the accuracy of some plans depends on the input data and may incur a privacy cost if it is used naively by the system during optimization. Last, in classical optimization, the set of logically equivalent plans is well defined, but for private algorithms, which are randomized, an appropriate definition of plan equivalence requires further investigation.

13 CONCLUSIONS

We have described the design and implementation of ϵ KTELO: an extensible programming framework and system for defining and executing differentially private algorithms. Many state-of-the-art differentially private algorithms can be specified as plans consisting of sequences of operators, increasing code reuse and facilitating more transparent algorithm comparisons. Algorithms implemented in ϵ KTELO are often faster and scale to larger domains by leveraging ϵ KTELO's compact internal data representations, based on implicit matrices. Using ϵ KTELO, we designed new algorithms that outperform the state of the art in accuracy on tasks requiring answers to sets of linear queries.

By allowing plan authors to focus on the simpler problem of designing a plan, and shifting the burden of implementing Privacy-critical and accuracy-critical operators to privacy engineers, we hope ϵ KTELO will drive the wider adoption of differential privacy.

REFERENCES

- [1] 2010. Apache Hive. Retrieved from <https://hive.apache.org>.
- [2] 2010. OnTheMap. Retrieved from <https://onthemap.ces.census.gov/>.
- [3] 2012. Apache Accumulo. Retrieved from <https://accumulo.apache.org>.
- [4] 2014. Apache Spark. Retrieved from <https://spark.apache.org>.
- [5] 2018. 2018 Differential Privacy Synthetic Data Challenge. Retrieved from <https://www.nist.gov/communications-technology-laboratory/pscr/funding-opportunities/open-innovation-prize-challenges-1>.
- [6] 2018. Ektelo source code. Retrieved from <https://github.com/ektelo/ektelo>.
- [7] Gergely Ács, Claude Castelluccia, and Rui Chen. 2012. Differentially private histogram publishing through lossy compression. In *Proceedings of the IEEE International Conference on Data Mining (ICDM'12)*. 1–10.
- [8] Boaz Barak, Kamalika Chaudhuri, Cynthia Dwork, Satyen Kale, Frank McSherry, and Kunal Talwar. 2007. Privacy, accuracy, and consistency too: A holistic solution to contingency table release. In *Proceedings of the Twenty-Sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. 273–282.
- [9] Gilles Barthe, Gian Pietro Farina, Marco Gaboardi, Emilio Jesus Gallego Arias, Andy Gordon, Justin Hsu, and Pierre-Yves Strub. 2016. Differentially private bayesian programming. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS'16)*. 68–79.
- [10] Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. 1995. A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.* 16, 5 (1995), 1190–1208.
- [11] Graham Cormode. 2011. Personal privacy vs population privacy: Learning to attack anonymization. In *Proceedings of the Conference on Knowledge Discovery and Data Mining (KDD'11)*.
- [12] Graham Cormode, Magda Procopiuc, Entong Shen, Divesh Srivastava, and Ting Yu. 2012. Differentially private spatial decompositions. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE'12)*. 20–31.
- [13] Cynthia Dwork, Frank McSherry Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Theory of Cryptography Conference (TCC'06)*. 265–284.
- [14] Cynthia Dwork and Aaron Roth. 2014. *The Algorithmic Foundations of Differential Privacy*. Foundations and Trends in Theoretical Computer Science.
- [15] Hamid Ebadi and David Sands. 2016. Featherweight PINQ. *Journal of Privacy and Confidentiality* 7, 2 (2016), 159–164.

- [16] Tarek Elgamal, Shangyu Luo, Matthias Boehm, Alexandre V. Evfimievski, Shirish Tatikonda, Berthold Reinwald, and Prithviraj Sen. 2017. SPOOF: Sum-product optimization and operator fusion for large-scale machine learning. In *Proceedings of the Conference on Innovative Data Systems Research (CIDR'17)*.
- [17] Úlfar Erlingsson, Vasył Pihur, and Aleksandra Korolova. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS'14)*.
- [18] David Chin-Lung Fong and Michael Saunders. 2011. LSMR: An iterative algorithm for sparse least-squares problems. *SIAM J. Sci. Comput.* 33, 5 (Oct. 2011), 2950–2971.
- [19] Marco Gaboardi, Andreas Haeberlen, Justin Hsu, Arjun Narayan, and Benjamin C. Pierce. 2013. Linear dependent types for differential privacy. In *Proceedings of the ACM SIGPLAN Symposium on Principles of Programming Languages (POPL'13)*. 357–370.
- [20] Samuel Haney, Ashwin Machanavajjhala, John Abowd, Matthew Graham, Mark Kutzbach, and Lars Vilhuber. 2017. Utility cost of formal privacy for releasing national employer-employee statistics. In *Proceedings of the ACM Conference on Management of Data (SIGMOD'17)*.
- [21] Moritz Hardt, Katrina Ligett, and Frank McSherry. 2012. A simple and practical algorithm for differentially private data release. In *Proceedings of the Conference on Neural Information Processing Systems (NIPS'12)*.
- [22] Michael Hay, Ashwin Machanavajjhala, Gerome Miklau, Yan Chen, and Dan Zhang. 2016. Principled evaluation of differentially private algorithms using DPBench. In *Proceedings of the ACM Conference on Management of Data (SIGMOD'16)*.
- [23] Michael Hay, Vibhor Rastogi, Gerome Miklau, and Dan Suciu. 2010. Boosting the accuracy of differentially private histograms through consistency. *Proc. VLDB* 3, 1–2 (2010), 1021–1032.
- [24] Dylan Hutchison, Bill Howe, and Dan Suciu. 2017. LaraDB: A minimalist kernel for linear and relational algebra computation. In *Proceedings of the 4th ACM SIGMOD Workshop on Algorithms and Systems for MapReduce and Beyond*.
- [25] Georgios Kellaris, Stavros Papadopoulos, and Dimitris Papadias. 2015. Differentially private histograms for range-sum queries: A modular approach. *arXiv* (2015).
- [26] Ios Kotsogiannis, Ashwin Machanavajjhala, Michael Hay, and Gerome Miklau. 2017. Pythia: Data dependent differentially private algorithm selection. In *Proceedings of the ACM Conference on Management of Data (SIGMOD'17)*.
- [27] Ios Kotsogiannis, Yuchao Tao, Ashwin Machanavajjhala, Gerome Miklau, and Michael Hay. 2019. Architecting a differentially private SQL engine. In *Proceedings of the Conference on Innovative Data Systems Research (CIDR'19)*.
- [28] Jaewoo Lee, Yue Wang, and Daniel Kifer. 2015. Maximum likelihood postprocessing for differential privacy under consistency constraints. In *Proceedings of the Conference on Knowledge Discovery and Data Mining (KDD'15)*.
- [29] Chao Li, Michael Hay, and Gerome Miklau. 2014. A data- and workload-aware algorithm for range queries under differential privacy. *Proc. VLDB* 7, 5 (2014), 341–352.
- [30] Chao Li, Michael Hay, Vibhor Rastogi, Gerome Miklau, and Andrew McGregor. 2010. Optimizing linear counting queries under differential privacy. In *Proceedings of the Twenty-Ninth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*. 123–134.
- [31] Chao Li and Gerome Miklau. 2012. An adaptive mechanism for accurate query answering under differential privacy. *Proc. VLDB* 5, 6 (2012), 514–525.
- [32] Chao Li, Gerome Miklau, Michael Hay, Andrew McGregor, and Vibhor Rastogi. 2015. The matrix mechanism: Optimizing linear counting queries under differential privacy. *VLDB J.* 24, 6 (2015), 757–781.
- [33] Ryan McKenna, Gerome Miklau, Michael Hay, and Ashwin Machanavajjhala. 2018. Optimizing error of high-dimensional statistical queries under differential privacy. *Proc. VLDB Endow.* 11, 10 (2018), 1206–1219.
- [34] Ryan McKenna, Daniel Sheldon, and Gerome Miklau. 2019. Graphical-model based estimation and inference for differential privacy. In *Proceedings of the International Conference on Machine Learning (ICML'19)*.
- [35] Frank D. McSherry. 2009. Privacy integrated queries: An extensible platform for privacy-preserving data analysis. In *Proceedings of the ACM Conference on Management of Data (SIGMOD'09)*. 19–30.
- [36] Ilya Mironov. 2012. On significance of the least significant bits for differential privacy. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS'12)*.
- [37] Laurel Orr, Magdalena Balazinska, and Dan Suciu. 2017. Probabilistic database summarization for interactive data exploration. *Proc. VLDB Endow.* 10, 10 (Jun. 2017), 1154–1165. DOI: <https://doi.org/10.14778/3115404.3115419>
- [38] Shoumik Palkar, James J. Thomas, Anil Shanbhag, Deepak Narayanan, Holger Pirk, Malte Schwarzkopf, Saman Amarasinghe, and Matei Zaharia. 2017. Weld: A common runtime for high performance data analytics. In *Proceedings of the Conference on Innovative Data Systems Research (CIDR'17)*.
- [39] Yongjoo Park, Ahmad Shahab Tajik, Michael Cafarella, and Barzan Mozafari. 2017. Database learning: Toward a database that becomes smarter every time. In *Proceedings of the ACM Conference on Management of Data (SIGMOD'17)*. 587–602. DOI: <https://doi.org/10.1145/3035918.3064013>
- [40] R. Penrose. 1956. On best approximate solutions of linear matrix equations. *Math. Proc. Camb. Philos. Soc.* 52, 1 (1956), 17–19. DOI: <https://doi.org/10.1017/S0305004100030929>

- [41] Davide Proserpio, Sharon Goldberg, and Frank McSherry. 2012. A workflow for differentially-private graph synthesis. In *Workshop on Online Social Networks*.
- [42] Davide Proserpio, Sharon Goldberg, and Frank McSherry. 2014. Calibrating data to sensitivity in private data analysis: A platform for differentially-private analysis of weighted datasets. *Proc. VLDB Endow.* 7, 8 (Apr. 2014), 637–648.
- [43] Wahbeh Qardaji, Weining Yang, and Ninghui Li. 2013. Differentially private grids for geospatial data. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE'13)*. IEEE, 757–768.
- [44] Wahbeh Qardaji, Weining Yang, and Ninghui Li. 2013. Understanding hierarchical methods for differentially private histograms. *Proc. VLDB* 6, 14 (2013).
- [45] Wahbeh Qardaji, Weining Yang, and Ninghui Li. 2014. Privview: Practical differentially private release of marginal contingency tables. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*. ACM, 1435–1446.
- [46] Indrajit Roy, Srinath T. V. Setty, Ann Kilzer, Vitaly Shmatikov, and Emmett Witchel. 2010. Airavat: Security and privacy for MapReduce. In *Proceedings of the USENIX Symposium on Networked Systems Design and Implementation (NSDI'10)*.
- [47] Sebastian Schelter, Andrew Palumbo, Shannon Quinn, Suneel Marthi, and Andrew Musselman. 2016. Samsara: Declarative machine learning on distributed dataflow systems. In *Proceedings of the Neural Information Processing Systems (NeurIPS'16)*.
- [48] Jack Sherman and Winifred J. Morrison. 1950. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *Ann. Math. Stat.* 21, 1 (3 1950), 124–127. DOI: <https://doi.org/10.1214/aoms/1177729893>
- [49] J. Vaidya, B. Shafiq, A. Basu, and Y. Hong. 2013. Differentially private naive bayes classification. In *Proceedings of the 2013 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI'13) and Intelligent Agent Technologies (IAT'13)*, Vol. 1. 571–576. DOI: <https://doi.org/10.1109/WI-IAT.2013.80>
- [50] Oliver Williams and Frank McSherry. 2010. Probabilistic inference and differential privacy. *Proceedings of the Conference on Neural Information Processing Systems (NIPS'10)*. 2451–2459.
- [51] Xiaokui Xiao, Guozhang Wang, and Johannes Gehrke. 2010. Differential privacy via wavelet transforms. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE'10)*. 225–236.
- [52] I-Cheng Yeh and Che hui Lien. 2009. The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert Syst. Appl.* 36, 2 (2009), 2473–2480.
- [53] Danfeng Zhang and Daniel Kifer. 2017. LightDP: Towards automating differential privacy proofs. In *Proceedings of the ACM SIGPLAN Symposium on Principles of Programming Languages (POPL'17)*. 888–901.
- [54] Dan Zhang, Ryan McKenna, Ios Kotsogiannis, Michael Hay, Ashwin Machanavajjhala, and Gerome Miklau. 2018. EKTELO: A framework for defining differentially-private computations. In *Proceedings of the 2018 International Conference on Management of Data (SIGMOD'18)*. ACM, New York, NY, 115–130. DOI: <https://doi.org/10.1145/3183713.3196921>
- [55] Jun Zhang, Graham Cormode, Cecilia M. Procopiuc, Divesh Srivastava, and Xiaokui Xiao. 2017. PrivBayes: Private data release via Bayesian networks. *ACM Trans. Database Syst.* 42, 4 (2017).
- [56] Jun Zhang, Xiaokui Xiao, and Xing Xie. 2016. PrivTree: A differentially private algorithm for hierarchical decompositions. In *Proceedings of the ACM Conference on Management of Data (SIGMOD'16)*.
- [57] Xiaojian Zhang, Rui Chen, Jianliang Xu, Xiaofeng Meng, and Yingtao Xie. 2014. Towards accurate histogram publication under differential privacy. In *Proceedings of the SIAM International Conference on Data Mining (SDM'14)*.

Received December 2018; revised June 2019; accepted September 2019