

Turbo: Effective Caching in Differentially-Private Databases

Kelly Kostopoulou^{1*}, Pierre Tholoniati^{1*}, Asaf Cidon¹, Roxana Geambasu¹, and Mathias Lécuyer²
¹Columbia University and ²University of British Columbia

Abstract

Differentially-private (DP) databases allow for privacy-preserving analytics over sensitive datasets or data streams. In these systems, *user privacy* is a limited resource that must be conserved with each query. We propose *Turbo*, a novel, state-of-the-art caching layer for linear query workloads over DP databases. Turbo builds upon private multiplicative weights (PMW), a DP mechanism that is powerful in theory but ineffective in practice, and transforms it into a highly-effective caching mechanism, *PMW-Bypass*, that uses prior query results obtained through an external DP mechanism to train a PMW to answer arbitrary future linear queries accurately and “for free” from a privacy perspective. Our experiments on public Covid and CitiBike datasets show that Turbo with PMW-Bypass conserves 1.7 – 15.9× more budget compared to vanilla PMW and simpler cache designs, a significant improvement. Moreover, Turbo provides support for range query workloads, such as timeseries or streams, where opportunities exist to further conserve privacy budget through DP parallel composition and warm-starting of PMW state. Our work provides a theoretical foundation and general system design for effective caching in DP databases.

CCS Concepts: • Security and privacy → Data anonymization and sanitization.

1 Introduction

ABC collects lots of user data from its digital products to analyze trends, improve existing products, and develop new ones. To protect user privacy, the company uses a restricted interface that removes personally identifiable information and only allows queries over aggregated data from multiple users. Internal analysts use interactive tools like Tableau to examine static datasets and run jobs to calculate aggregate metrics over data streams. Some of these metrics are shared with external partners for product integrations. However, due to data reconstruction attacks on similar “anonymized” and “aggregated” data from other sources, including the US Census Bureau [29] and Aircloak [17], the CEO has decided to pause external

aggregate releases and severely limit the number of analysts with access to user data statistics until the company can find a more rigorous privacy solution.

The preceding scenario, while fictitious, is representative of what often occurs in industry and government, leading to obstacles to data analysis or incomplete privacy solutions. In 2007, Netflix withdrew “anonymized” movie rating data and canceled a competition due to de-anonymization attacks [52]. In 2008, genotyping aggregate information from a clinical study led to the revelation of participants’ membership in the diagnosed group, prompting the National Institutes of Health to advise against the public release of statistics from clinical studies [4]. In 2021, New York City excluded demographic information from datasets released from their CitiBike bike rental service, which could reveal sensitive user data [3]. The city’s new, more restrained data release not only remains susceptible to privacy attacks but also prevents analyses of how demographic groups use the service.

Differential privacy (DP) provides a rigorous solution to the problem of protecting user privacy while analyzing and sharing statistical aggregates over a database. DP guarantees that analysts cannot confidently learn anything about any individual in the database that they could not learn if the individual were not in the database. Industry and government have started to deploy DP for various use cases [22], including publishing trends in Google searches related to Covid [10], sharing LinkedIn user engagement statistics with outside marketers [53], enabling analyst access to Uber mobility data while protecting against insider attacks [36], and releasing the US Census’ 2020 redistricting data [5]. To facilitate the application of DP, industry has developed a suite of systems, ranging from specialized designs like the US Census TopDown [5] and LinkedIn Audience Engagements [53] to more general DP SQL systems, like GoogleDP [62], Uber Chorus [36], and Tumult Analytics [12].

DP systems face a significant challenge that hinders their wider adoption: they struggle to handle large workloads of queries while maintaining a reasonable privacy guarantee. This is known as the “running out of privacy budget” problem and affects any system, whether DP or not, that aims to release multiple statistics from a sensitive dataset. A seminal paper by Dinur and Nissim [23] proved that releasing too many accurate linear statistics from a dataset fundamentally enables its reconstruction, setting a lower bound on the necessary error in queries to prevent such reconstruction. Successful reconstructions of the US Census 2010 data [29] and Aircloak’s data [17] from the aggregate statistics released by these entities exemplify this fundamental limitation. DP, while not

*These authors contributed equally to this work

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SOSP ’23, October 23–26, 2023, Koblenz, Germany

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0229-7/23/10.

<https://doi.org/10.1145/3600006.3613174>

immune to this limitation, provides a means of bounding the reconstruction risk. DP randomizes the output of a query to limit the influence of individual entries in the dataset on the result. Each new DP query increases this limit, consuming part of a *global privacy budget* that must not be exceeded, lest individual entries become vulnerable to reconstruction.

Recent work proposed treating the global privacy budget as a *system resource* that must be managed and conserved, similar to traditional resources like CPU [45]. When computation is expensive, *caching* is a go-to solution: it uses past results to save CPU on future computations. Caches are ubiquitous in all computing systems – from the processor to operating systems and databases – enabling scaling to much larger workloads than would otherwise be afforded with fixed resources. In this paper, we thus ask: *How should caching work in DP systems to significantly increase the number of queries they can support under a privacy guarantee?* While DP theory has explored algorithms to reuse past query results to save privacy budget in future queries, there is no general DP caching system that is effective in common practical settings.

We propose *Turbo*, the first general and effective caching layer for DP SQL databases that boosts the number of linear queries (such as sums, averages, counts) that can be answered accurately under a fixed, global DP guarantee. In addition to incorporating a traditional *exact-match cache* that saves past DP query results and reuses them if the same query reappears, Turbo builds upon a powerful theoretical construct, known as *private multiplicative weights (PMW)* [31], that leverages past DP query results to learn a histogram representation of the dataset that can go on to answer *arbitrary* future linear queries for free once it has converged. While PMW has compelling convergence guarantees in theory, we find it ineffective in practice, being overrun even by an exact-match cache.

We make three main contributions to PMW design to boost its effectiveness and applicability. First, we develop *PMW-Bypass*, a variant of PMW that bypasses it during the privacy-expensive learning phase of its histogram, and switches to it once it has converged to reap its free-query benefits. This change requires a new mechanism for updating the histogram despite bypassing the PMW, plus new theory to justify its convergence. The PMW-Bypass technique is highly effective, significantly outperforming both the exact-match cache and vanilla PMW in the number of queries it can support. Second, we optimize our mechanisms for workloads of range queries that do not access the entire database. These types of queries are typical in timeseries databases and data streams. For such workloads, we organize the cache as a tree of multiple PMW-Bypass objects and demonstrate that this approach outperforms alternative designs. Third, for streaming workloads, we develop warm-starting procedures for tree-structured PMW-Bypass histograms, resulting in faster convergence.

We formally analyze each of our techniques, focusing on privacy, per-query accuracy, and convergence speed. Each technique represents a contribution on its own and can be

used separately, or, as we do in Turbo, as part of the first *general, effective, and accurate DP-SQL caching design*. We prototype Turbo on TimescaleDB, a timeseries database, and use Redis to store caching state. We evaluate Turbo on workloads based on Covid and CitiBike datasets. We show that Turbo significantly improves the number of linear queries that can be answered with less than 5% error (w.h.p.) under a global (10, 0)-DP guarantee, compared to not having a cache and alternative cache designs. Our approach outperforms the best-performing baseline in each workload by 1.7 to 15.9 times, and even more significantly compared to vanilla PMW and systems with no cache at all (such as most existing DP systems). These results demonstrate that our Turbo cache design is both general and effective in boosting workloads in DP SQL databases and streams, making it a promising solution for companies like ABC that seek an effective DP SQL system to address their user data analysis and sharing concerns. We make Turbo available open-source at <https://github.com/columbia/turbo>, part of a broader set of infrastructure systems we are developing for DP, all described here: <https://systems.cs.columbia.edu/dp-infrastructure>.

2 Background

Threat model. We consider a threat model known as *centralized differential privacy*: one or more untrusted analysts query a dataset or stream through a restricted, aggregate-only interface implemented by a trusted database engine of which Turbo is a trusted component. The goal of the database and Turbo is to provide accurate answers to the analysts’ queries without compromising the privacy of individual users in the database. The two main adversarial goals that an analyst may have are membership inference and data reconstruction. Membership inference is when the adversary wants to determine whether a known data point is present in the dataset. Data reconstruction involves reconstructing unknown data points from a known subset of the dataset. To achieve their goals, the adversary can use composition attacks to single out contributions from individuals, collude with other analysts to coordinate their queries, link anonymized records to public datasets, and access arbitrary auxiliary information *except for* timing side-channel information. Previous research demonstrated attacks under this threat model [17, 21, 28, 29, 33, 52].

Differential privacy (DP). DP [25] randomizes aggregate queries over a dataset to prevent membership inference and data reconstruction [24, 61]. DP randomization (a.k.a. noise) ensures that the probability of observing a specific result is stable to a change in one datapoint (e.g., if user x is removed or replaced in the dataset, the distribution over results remains similar). More formally, a query Q is (ϵ, δ) -DP if, for any two datasets D and D' that differ by one datapoint, and for any result subset S we have: $\mathbb{P}(Q(D) \in S) \leq e^\epsilon \mathbb{P}(Q(D') \in S) + \delta$. ϵ quantifies the privacy loss due to releasing the DP query’s result (higher means less privacy), while δ can be interpreted as a failure probability and is set to a small value.

Two common mechanisms to enforce DP are the Laplace and Gaussian mechanisms. They add noise from an appropriately scaled Laplace/Gaussian distribution to the true query result, and return the noisy result. As an example, for counting queries and a database of size n , adding noise from $\text{Laplace}(0, 1/n\epsilon)$, ensures $(\epsilon, 0)$ -DP (a.k.a. pure DP); adding noise from $\text{Gaussian}(0, \sqrt{2\ln(1.25/\delta)}/n\epsilon)$ ensures (ϵ, δ) -DP. The accuracy for such queries can be controlled probabilistically by converting it into the (ϵ, δ) parameters.

Answering multiple queries on the same data fundamentally degrades privacy [23]. DP quantifies this over a sequence of DP queries using the *composition property*, which in its basic form states that releasing two (ϵ_1, δ_1) -DP and (ϵ_2, δ_2) -DP queries is $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -DP. When queries access disjoint data subsets, their composition is $(\max(\epsilon_1, \epsilon_2), \max(\delta_1, \delta_2))$ -DP and is called *parallel composition*. Using composition, one can enforce a global (ϵ_G, δ_G) -DP guarantee over a workload, with each DP query “consuming” part of a *global privacy budget* that is defined upfront as a system parameter [54].

Good values of the global privacy budget in interactive DP SQL systems remain subject for debate [34], but generally, an ideal value for strong theoretical guarantees is $\epsilon_G = 0.1$, while $\epsilon_G = 1$ are considered acceptable. Larger values are often considered vacuous semantically, since individuals’ privacy risk grows with e^{ϵ_G} . In this paper, we aim to achieve values of $\epsilon_G = 1$ or smaller over a query workload.

Private multiplicative weights (PMW). PMW is a DP mechanism to answer online linear queries with bounded error [31]. We defer detailed description of PMW, plus an example illustrating its functioning, to §4 and only give here an overview. PMW maintains an approximation of the dataset in the form of a *histogram*: estimated counts of how many times any possible data point appears in the dataset. When a query arrives, PMW estimates an answer using the histogram and computes the *error of this estimate* against the real data in a DP way, using a DP mechanism called *sparse vector (SV)* [26] (described shortly). If the estimate’s error is low, it is returned to the analyst, consuming no privacy budget (i.e., the query is answered “for free”). If the estimate’s error is large, then PMW executes the DP query on the data with the Laplace/Gaussian mechanism, consuming privacy budget as needed. It returns the DP result and also uses it to update the histogram for more accurate estimates to future queries.

An additional cost in using PMW comes from the SV, a well-known DP mechanism that can be used to test the error of a sequence of query estimates against the ground truth with DP guarantees and limited privacy budget consumption [26]. We refer the reader to textbook descriptions of SV for detailed functioning [26] and provide here only an overview of its semantics. SV is a stateful mechanism that receives queries and estimates for their results one by one, and assesses the error between these estimates and the ground-truth query results. While the estimates have error below a preset threshold with

high probability, SV returns success and consumes *zero privacy*. However, as soon as SV detects a large-error estimate, it requires a *reset*, which is a privacy-expensive operation that re-initializes state within the SV to continue the assessments. In common SV implementations, a reset costs as much as $3\times$ the privacy budget of executing one DP query on the data.

The theoretical vision of PMW is as follows. Under a stream of queries, PMW first goes through a “training” phase, where its histogram is inaccurate, requiring frequent SV resets and consuming budget. Failed estimation attempts update the histogram with low-error results obtained by running the DP query. Once the histogram becomes sufficiently accurate, the SV tests consistently pass, thereby ameliorating the initial training cost. Theoretical analyses provide a compelling *worst-case convergence* guarantee for the histogram, determining a worst-case number of updates required to train a histogram that can answer *any future linear query* with low error [32]. However, no one has examined whether this worst-case bound is practical and if PMW outperforms natural baselines, such as an exact-match cache.

3 Turbo Overview

Turbo is a caching layer that can be integrated into a DP SQL engine, significantly increasing the number of linear queries that can be executed under a fixed, global (ϵ_G, δ_G) -DP guarantee. We focus on *linear queries* like sums, averages, and counts (defined in §4), which are widely used in interactive analytics and constitute the class of queries supported by approximate databases such as BlinkDB [6]. These queries enable powerful forms of caching like PMW, and also allow for accuracy guarantees, which are important when doing approximate analytics, as one does on a DP database.

3.1 Design Goals

In designing Turbo, we were guided by several goals:

- (G1) *Guarantee privacy*: Turbo must satisfy (ϵ_G, δ_G) -DP.
- (G2) *Guarantee accuracy*: Turbo must ensure (α, β) -accuracy for each query, defined for $\alpha > 0, \beta \in (0, 1)$ as follows: if R' and R are the returned and true results, then $|R' - R| \leq \alpha$ with $(1 - \beta)$ probability. If β is small, a result is α -accurate *w.h.p.* (with high probability).
- (G3) and (G4) *Provide worst-case convergence guarantees but optimize for empirical convergence*: We aim to maintain PMW’s theoretical convergence (G3), but we prioritize for *empirical convergence* speed, a new metric that measures, on a workload, the number of updates needed to answer most queries for free (G4).
- (G5) *Improve privacy budget consumption*: We aim for *significant improvements* in privacy budget consumption compared to both not having a cache and having an exact-match cache or a vanilla PMW.
- (G6) *Support multiple use cases*: Turbo should benefit multiple important workload types, including static and streaming databases, and queries that arrive over time.

(G7) *Easy to configure*: Turbo should include few knobs with fairly stable performance.

(G1) and (G2) are strict requirements. (G3) and (G4) are driven by our belief that DP systems should not only possess meaningful theoretical properties but also be optimized for practice. (G5) is our main objective. (G6) requires further attention, given shortly. (G7) is driven by the limited guidance from PMW literature on parameter tuning. PMW meets goals (G1-G3) but falls significantly short for (G4-G7). Turbo achieves all goals; we provide theoretical analyses for (G1-G3) in §4 and empirical evaluations for (G4-G7) in §5.

3.2 Use Cases

The DP literature is fragmented, with different algorithms developed for different use cases. We seek to create a *general system* that supports multiple settings, highlighting three here: (1) **Non-partitioned databases** are the most common use case in DP. A group of untrusted analysts issue queries over time against a static database, and the database owner wishes to enforce a global DP guarantee. Turbo should allow a larger workload of queries compared to existing approaches.

(2) and (3) **Partitioned databases** are less frequently investigated in DP theory literature, but important to distinguish in practice [50, 57]. When queries tend to access different data ranges, it is worth partitioning the data and accounting for consumed privacy budget in each partition separately through DP’s parallel composition. This lowers privacy budget consumption in each partition and permits more non- or partially-overlapping queries against the database. This kind of workload is inherent in *timeseries* and *streaming databases*, where analysts typically query the data by *windows of time*, such as how many new Covid cases occurred in the week after a certain event, or what is the average age of positive people over the past week. We distinguish two cases:

(2) **Partitioned static database**, where the database is static and partitioned by an attribute that tends to be accessed in ranges, such as time, age, or geo-location. All partitions are available at the beginning. Queries arrive over time and most are assumed to run on some range of interest, which can involve one or more partitions. Turbo should provide significant benefit not only compared to the baseline caching techniques, but also compared to not having partitioning.

(3) **Partitioned streaming database**, where the database is partitioned by time and partitions arrive over time. In such workloads, queries tend to run *continuously* as new data becomes available. Hence, new partitions see a similar query workload as preceding partitions. Turbo should take advantage of this similarity to further conserve privacy.

For all three use cases, we aim to support *online workloads* of queries that are not all known upfront. As §7 reviews, most works on optimizing global privacy budget consumption operate in the *offline setting*, where all queries are known upfront. For that setting, algorithms are known to answer all queries simultaneously with optimal use of privacy budget. However,

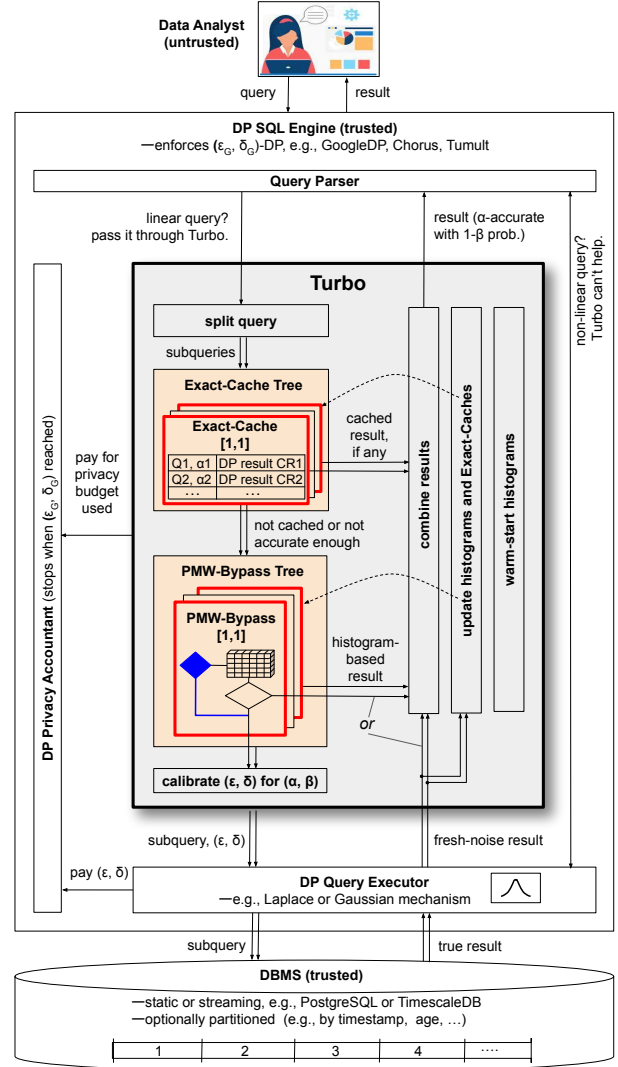


Fig. 1. Turbo architecture.

this setting is unrealistic for real use cases, where analysts adapt their queries based on previous results, or issue new queries for different analyses. In such cases, which correspond to the *online setting*, we require adaptive algorithms that accurately answer queries on-the-fly. Turbo does this by making effective use of PMW, as we next describe.

3.3 Turbo Architecture

Fig. 1 shows the Turbo architecture. It is a caching layer that can be added to a DP SQL engine, like GoogleDP [62], Uber Chorus [36], or Tumult Analytics [12], to boost the number of linear queries that can be answered accurately under a fixed global DP guarantee. The filled components indicate our additions to the DP SQL engine, while the transparent components are standard in DP SQL engines. Here is how a typical DP SQL engine works *without Turbo*. Analysts issue queries against the engine, which is trusted to enforce a global (ϵ_G, δ_G) -DP guarantee. The engine executes the queries using a DP query executor, which adds noise to query results with

the Laplace/Gaussian mechanism and consumes a part of the global privacy budget. A budget accountant tracks the consumed budget; when it runs out, the DP SQL engine either stops responding to new queries (as do Chorus and Tumult Analytics) or sacrifices privacy by “resetting” the budget (as does LinkedIn Audience Insights). We assume the former.

Turbo intercepts the queries before they go into the DP query executor and performs a very proactive form of caching for them, reusing prior results as much as possible to avoid consuming privacy budget for new queries. Turbo’s architecture is organized in two types of components: *caching objects* (denoted in light-orange background in Fig. 1) and *functional components* that act upon them (denoted in grey background).

Caching objects. Turbo maintains several types of caching objects. First, the *Exact-Cache* stores previous queries and their DP results, allowing for direct retrieval of the result without consuming any privacy budget when the same query is seen again on the same database version. Second, the *PMW-Bypass* is an improved version of PMW that reduces privacy budget consumption during the training phase of its histogram (§4.3). Given a query, PMW-Bypass uses an effective heuristic to judge whether the histogram is sufficiently trained to answer the query accurately; if so, it uses it, thereby spending no budget. Critically, PMW-Bypass includes a mechanism to *externally update* the histogram even when bypassing it, to continue training it for future, free-budget queries.

Turbo aims to enable parallel composition for workloads that benefit from it, such as timeseries or streaming workloads, by supporting database partitioning. In theory, partitions could be defined by attributes with public values that are typically queried by range, such as time, age, or geo-location. In this paper, we will focus on partitioning by time. Turbo uses a *tree-structured PMW-Bypass* caching object, consisting of multiple histograms organized in a binary tree, to support linear range queries over these partitions effectively (§4.4). This approach conserves more privacy budget and enables larger workloads to be run when queries access only subsets of the partitions, compared to alternative methods.

Functional components. When Turbo receives a linear query through the DP SQL engine’s query parser, it applies its caching objects to the query. If the database is partitioned, Turbo splits the query into multiple sub-queries based on the available tree-structured caching objects. Each sub-query is first passed through an Exact-Cache, and if the result is not found, it is forwarded to a PMW-Bypass, which selects whether to execute it on the histogram or through direct Laplace/Gaussian. For sub-queries that can leverage histograms, the answer is supplied directly without execution or budget consumption. For sub-queries that require execution with Laplace/Gaussian, the (ϵ, δ) parameters for the mechanism are computed based on the (α, β) accuracy parameters, using the “calibrate (ϵ, δ) for (α, β) ” functional component in Fig. 1. Then, each sub-query and its privacy parameters are passed to the DP query executor for execution.

Turbo combines all sub-query results obtained from the caching objects to form the final result, ensuring that it is within α of the true result with probability $1 - \beta$ (functional component “combine results”). New results computed with fresh noise are used to update the caching objects (functional component “update histograms and Exact-Caches”). Additionally, Turbo includes cache management functionality, such as “warm-start of histograms,” which reuses trained histograms from previous partitions to warm-start new histograms when a new partition is created (§4.5).

4 Detailed Design

We next detail the novel caching objects and mechanisms in Turbo, using different use cases from §3.2 to illustrate each concept. We describe PMW-Bypass in the static, non-partitioned database, then introduce partitioning for the tree-structured PMW-Bypass, followed by the addition of streaming to discuss warm-start procedures. We focus on the Laplace mechanism and basic composition, thus only discussing pure $(\epsilon, 0)$ -DP and ignoring δ . We also assume β is small enough for Turbo results to count as α -accurate w.h.p. Appendix A.6 extends all our theoretical results to (ϵ, δ) -DP, non-zero β , the Gaussian mechanism, and Rényi composition; in theory, all these should help to further conserve privacy budget, so we speculate they will be important for practice, but we leave their implementation and evaluation for future work.

4.1 Notation

Our algorithms require some notation. Given a data domain \mathcal{X} , a database x with n rows can be represented as a histogram $h \in \mathbb{N}^{\mathcal{X}}$ as follows: for any data point $v \in \mathcal{X}$, $h(v)$ denotes the number of rows in x whose value is v . $h(v)$ is the *bin* corresponding to value v in the histogram. We denote $N = |\mathcal{X}|$ the size of the data domain and n the size of the database. When \mathcal{X} has the form $\{0, 1\}^d$, we call d the data domain dimension. Example: a database with 3 binary attributes has domain $\mathcal{X} = \{0, 1\}^3$ of dimension $d = 3$ and size $N = 8$; $h(0, 0, 1)$ is the number of rows that are equal to $(0, 0, 1)$. §4.2 exemplifies a database, its dimensions, and its histogram.

We define *linear queries* as SQL queries that can be transformed or broken into the following form:

```
SELECT AVG(*) FROM ( SELECT q(A, B, C, ...) FROM Table ),
```

where q takes d arguments (one for each attribute of `Table`, denoted A, B, C, \dots) and outputs a value in $[0, 1]$. When q has values in $\{0, 1\}$, a query returns the *fraction of rows satisfying predicate* q . To get raw counts, we multiply by n , which we assume is public information. PMW (and hence Turbo) is designed to support only linear queries. Examples of *non-linear* queries are: maximum, minimum, percentiles, top-k.

4.2 Running Example

Fig. 2 gives a running example inspired by our evaluation Covid dataset. Analysts run queries against a database consisting of Covid test results over time. Fig. 2(a) shows a simplified version of the database, with only three attributes:

(a) “Covid” table:

Time (T)	Positive (P)	Age Bracket (A)
02/01/21	0	0 (1-17)
02/01/21	1	1 (18-49)
02/02/21	1	2 (50-64)
02/02/21	1	3 (65+)
... say n=100 total rows ...		

(b) Previously executed queries:

Q1: SELECT COUNT(*) FROM Covid
WHERE Positive=1

Q2: SELECT COUNT(*) FROM Covid
WHERE AgeBracket=0

(c) Histogram state after executing Q1, then Q2:

	A = 0	A = 1	A = 2	A = 3
P = 0	h(0,0): 12.5->18.3->8 (real: 13) c: 1	h(0,1): 12.5->18.3->21.7 (real: 27) c: 0	h(0,2): 12.5->18.3->21.7 (real: 15) c: 0	h(0,3): 12.5->18.3->21.7 (real: 25) c: 0
P = 1	h(1,0): 12.5->6.7->2.9 (real: 3) c: 2	h(1,1): 12.5->6.7->8 (real: 5) c: 1	h(1,2): 12.5->6.7->8 (real: 8) c: 1	h(1,3): 12.5->6.7->8 (real: 4) c: 1

Format: h(p,a): default-bin-value->value-after-Q1->value-after-Q2 (current value)
real: real value of the histogram bin (no DP, included as reference for h(v))
c: number of purposeful updates to the histogram bin

(d) Next query to execute:

Q3: SELECT COUNT(*) FROM Covid WHERE Positive=1 AND AgeBracket=0

Fig. 2. Running example. (a) Simplified Covid tests dataset with $n = 100$ rows and data domain size $N = 8$ for the two non-time attributes, test outcome P and subject’s age bracket A . (b) Two queries that were previously run. (c) State of the histogram as queries are executed. (d) Next query to run.

the test’s date, T ; the outcome, P , which can be 0 or 1 for negative/positive; and subject’s age bracket, A , with one of four values as in the figure. The database could be either static or actively streaming in new test data. Initially, we assume it static and ignore the T attribute. Our example database has $n = 100$ rows and data domain size $N = 8$ for P and A .

Fig. 2(b) shows two queries that were previously executed. While queries in Turbo return the *fraction* of entries satisfying a predicate, for simplicity we show raw counts. $Q1$ requests the positivity rate and $Q2$ the fraction of tested minors. Fig. 2(c) illustrates the histogram representation corresponding to the dataset, as estimated by the PMW algorithm, whose execution we discuss shortly. Fig. 2(d) shows the next query that will be executed, $Q3$, requesting the fraction of positive minors. $Q3$ is not identical to either $Q1$ or $Q2$, but it is *correlated* with both, as it accesses data that overlaps with both queries. Thus, while neither $Q1$ ’s nor $Q2$ ’s DP results can be used to directly answer $Q3$, intuitively, they both should help. That is the insight that PMW (and PMW-Bypass) exploits through its query-by-query build-up of a DP histogram representation of the database that becomes increasingly accurate in bins that are accessed by more queries.

Fig. 2(c) shows the state of the histogram after executing $Q1$ and $Q2$ but before executing $Q3$. Each bin in the histogram stores an *estimation* of the number of rows equal to (p, a) . This is the $h(p, a)$ field in the figure, for which we show the sequence of values it has taken following updates due to $Q1$ and $Q2$. Initially, $h(p, a)$ in all bins is set assuming a uniform distribution over $P \times A$; in this case the initial value was $n/N = 12.5$. The figure also shows the real (non-private) count for each bin (denoted *real*), which is *not* part of the histogram, but we include it as a reference. As queries are executed, $h(p, a)$ values are updated with DP results, depending on which bins are accessed. $Q1$ and $Q2$ have already been executed, and both are assumed to have resorted to the

Laplace mechanism, so they both contributed DP results to specific bins (we specify the update algorithm later when discussing Alg. 1). $Q1$ accessed, and hence updated, data in the $P = 1$ bins (the bottom row of the histogram). $Q2$ did so in the $A = 0$ bins (the left column of the histogram). Through a renormalization step, these queries have also changed the other bins, though not necessarily in a query-informed way. The c variable in each bin shows the number of queries that have purposely updated that bin. We can see that estimates in the $c > 0$ bins are a bit more accurate compared to those in the $c = 0$ bins. The only bin that has been updated twice is $(P = 1, A = 0)$, as it lies at the intersection of both queries; that bin has diverged from its neighboring, singly-updated bins and is getting closer to its true value. (Bin $(P = 1, A = 2)$, updated only once, is even more accurate purely by chance.)

Our last query, $Q3$, which accesses $(P = 1, A = 0)$, may be able to leverage its estimation “for free,” assuming the estimation’s error is within α w.h.p. Assessing that the error is within α – privately, and without consuming privacy budget if it is – is the purview of the SV mechanism incorporated in a PMW. The catch is that the SV consumes privacy budget, in copious amounts, if this test fails. This is what makes vanilla PMW impractical, a problem that we address next.

4.3 PMW-Bypass

PMW-Bypass addresses practical inefficiencies of PMW, which we illustrate with simple demonstration.

Demo experiment. Using a four-attribute Covid dataset with domain size 128 (so a bit larger than in our running example), we generate a query pool of over 34K unique queries by taking all possible combinations of values over the four attributes. From this pool, we sample uniformly with replacement 35K queries to form a workload; there is therefore some identical repetition of queries but not much. This workload is not necessarily realistic, but it should be an *ideal showcase* for PMW: there are many unique queries relative to the small data domain size (giving the PMW ample chance to train), and while most queries are unique, they tend to overlap in the data they touch (giving the PMW ample chance to reuse information from previous queries). We evaluate the cumulative privacy budget spent as queries are executed, comparing the case where we execute them through PMW vs. directly with Laplace, with and without an exact-match cache.

Fig. 3 shows the results. As expected for this workload, the PMW works, as it converges after roughly the first 10K queries and consumes very little budget afterwards.

However, before con-




Fig. 3. Demo experiment.

verging, the PMW consumes enormous budget. In contrast,

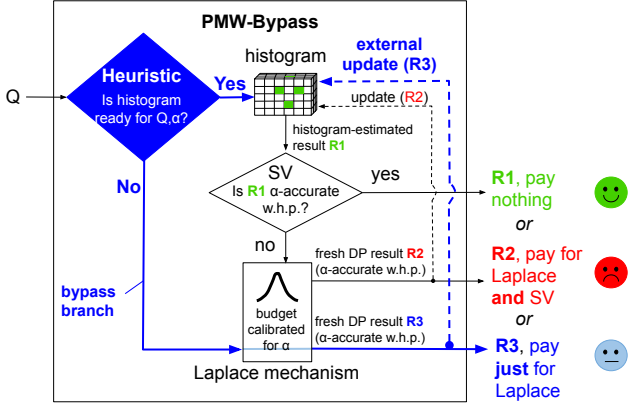


Fig. 4. PMW-Bypass. New components over vanilla PMW are in blue/bold. direct execution through Laplace grows linearly, but more slowly compared to PMW’s beginning. The PMW eventually becomes better than Laplace, but only after $\approx 27K$ queries.

Moreover, if instead of always executing with Laplace, we trivially cached the results in an exact-match cache for future reuse if the same query reappeared – a rare event in this workload – then the PMW would *never* become notably better than this simple baseline! This happens for a workload that should be ideal for PMW. §5 shows that for other workloads, less favorable for PMW but more realistic, the outcome persists: *PMWs underperform even the simplest baselines in practice.*

We propose *PMW-Bypass*, a re-design for PMWs that releases their power and makes them *very effective*. We make multiple changes to PMWs, but the main one involves *bypassing* the PMW while it is training (and hence expensive) and instead executing directly with Laplace (which is less expensive). Importantly, we do this while still updating the histogram with the Laplace results so that eventually the PMW becomes good enough to switch to it and reap its zero-privacy query benefits. The PMW-Bypass line in Fig. 3 shows just how effective this design is in our demo experiment: PMW-Bypass follows the low, direct-Laplace curve instead the PMW’s up until the histogram converges, after which it follows the flat shape of PMW’s convergence line. In this experiment, as well as in others in §5, the outcome is the same: *our changes make PMWs very effective.* We thus believe that PMW-Bypass should replace PMW in most settings where the latter is studied, not just in our system’s design.

PMW-Bypass. Fig. 4 shows the functionality of PMW-Bypass, with the main changes shown in blue and bold. Without our changes, a vanilla PMW works as follows. Given a query Q , PMW first estimates its result using the histogram ($R1$) and then uses the SV protocol to test whether it is α -accurate w.h.p. The test involves comparing $R1$ to the *exact result* of the query executed on the database. If a noisy version of the absolute error between the two is within a threshold comfortably far from α , then $R1$ is considered accurate w.h.p. and outputted directly. This is the good case, because the query need not consume *any privacy*. The bad case is when the SV test fails. First, the query must be executed directly through

Laplace, giving a result $R2$, whose release costs privacy. But beyond that, the SV must be *reset*, which consumes privacy. In total, if the Laplace execution costs ϵ , then releasing $R2$ costs $4 * \epsilon$! This is what causes the extreme privacy consumption during the training phase for vanilla PMW, when the SV test mostly fails. Still, in theory, after paying handsomely for this histogram “miss,” $R2$ can be used to update the histogram (the arrow denoted “update ($R2$)” in Fig. 4), in hopes that future correlated queries “hit” in the histogram.

PMW-Bypass adds three components to PMW: (1) a *heuristic* that assesses whether the histogram is likely ready to answer Q with the desired accuracy; (2) a *bypass branch*, taken if the histogram is deemed not ready and direct query execution with Laplace instead of going through (and likely failing) the SV test; and (3) an *external update* procedure that updates the histogram with the bypass branch result. Given Q , PMW-Bypass first consults the heuristic, which only inspects the histogram, so its use is free. Two cases arise:

Case 1: If the heuristic says the histogram is ready to answer Q with α -accuracy w.h.p., then the PMW is used, $R1$ is generated, and the SV is invoked to test $R1$ ’s actual accuracy. If the heuristic’s assessment was correct, then this test will succeed, and hence the free, $R1$ output branch will be taken. Of course, no heuristic that lacks access to the raw data can guarantee that $R1$ will be accurate enough, so if the heuristic was actually wrong, then the SV test will fail and the expensive $R2$ path is taken. Thus, a key design question is whether there exist heuristics good enough to make PMW-Bypass effective. We discuss heuristic designs below, but the gist is that simple and easily tunable heuristics work well, enabling the significant privacy budget savings in Fig. 3.

Case 2: If the heuristic says the histogram is not ready to answer Q with α -accuracy w.h.p., then the bypass branch is taken and Laplace is invoked directly, giving result $R3$. Now, PMW-Bypass must pay for Laplace, but because it bypassed the PMW, it does not risk an expensive SV reset. A key design question here is whether we can still reuse $R3$ to update the histogram, even though we did not, in fact, consult the SV to ensure that the histogram is truly insufficiently trained for Q . We prove that performing the same kind of update as the PMW would do, from outside the protocol, would break its theoretical convergence guarantee. Thus, for PMW-Bypass, we design an *external update* procedure that *can* be used to update the histogram with $R3$ while preserving the PMW’s worst-case convergence, albeit at slower speed.

Heuristic ISHISTOGRAMREADY. One option to assess if a histogram is ready to answer a query accurately is to check if it has received at least C updates, for some global threshold C . However, this approach is often imprecise as it fails to detect histogram regions that might still be untrained. Thus, we use a separate threshold value per bin, raising the question of how to configure all these thresholds. To keep configuration easy (goal **(G6)**), we use an *adaptive per-bin threshold*. For each bin, we initialize its threshold C with a value C_0 and increment

Algorithm 1 PMW-Bypass algorithm.

```
1: Cfg.: PRIVACYACCOUNTANT, HEURISTIC, accuracy params  $(\alpha, \beta)$ ,  
   histogram convergence params  $lr, \tau$ , database DATA with  $n$  rows.  
2: function UPDATE( $h, q, s$ )  
3:   Update estimated values:  $\forall v \in \mathcal{X}, g(v) \leftarrow h(v)e^{s \cdot q(v)}$   
4:   Renormalize:  $\forall v \in \mathcal{X}, h(v) \leftarrow g(v) / \sum_{w \in \mathcal{X}} g(w)$   
5:   return  $h$   
6: function CALIBRATEBUDGET( $\alpha, \beta$ )  
7:   return  $\frac{4 \ln(1/\beta)}{n\alpha}$   
8: Initialize histogram  $h$  to uniform distribution on  $\mathcal{X}$   
9:  $\epsilon \leftarrow \text{CALIBRATEBUDGET}(\alpha, \beta)$   
10: PRIVACYACCOUNTANT.PAY( $3 \cdot \epsilon$ ) // Pay to initialize first SV  
11: while PRIVACYACCOUNTANT.HASBUDGET() do  
12:    $\hat{\alpha} \leftarrow \alpha/2 + \text{Lap}(1/\epsilon n)$  // SV reset  
13:    $SV \leftarrow \text{NOTCONSUMED}$   
14:   while  $SV == \text{NOTCONSUMED}$  do  
15:     Receive next query  $q$   
16:     if HEURISTIC.ISHISTOGRAMREADY( $h, q, \alpha, \beta$ ) then  
17:       // Regular PMW branch:  
18:       if  $|q(\text{DATA}) - q(h)| + \text{Lap}(1/\epsilon n) < \hat{\alpha}$  then // SV test  
19:         Output  $R1 = q(h)$  // →R1, pay nothing  
20:       else  
21:         PRIVACYACCOUNTANT.PAY( $4 * \epsilon$ ) // →R2, pay for  
22:         Output  $R2 = q(\text{DATA}) + \text{Lap}(1/\epsilon n)$  // Laplace, SV  
23:         // Update histogram (R2):  
24:          $s \leftarrow \begin{cases} lr & \text{if } R2 > q(h) \\ -lr & \text{if } R2 < q(h) \end{cases}$   
25:          $h \leftarrow \text{UPDATE}(h, q, s)$   
26:          $SV \leftarrow \text{CONSUMED}$  // force SV reset  
27:         HEURISTIC.PENALIZE( $q, h$ )  
28:       else  
29:         // Bypass branch:  
30:         PRIVACYACCOUNTANT.PAY( $\epsilon$ ) // →R3, pay for  
31:         Output  $R3 = q(\text{DATA}) + \text{Lap}(1/\epsilon n)$  // Laplace  
32:         // External update of histogram (R3):  
33:          $s \leftarrow \begin{cases} lr & \text{if } R3 > q(h) + \tau\alpha \\ -lr & \text{if } R3 < q(h) - \tau\alpha \\ 0 & \text{otherwise // no updates if we're not confident!} \end{cases}$   
34:          $h \leftarrow \text{UPDATE}(h, q, s)$ 
```

C by an additive step S_0 every time the heuristic errs (i.e., predicts it is ready when it is in fact not ready for that query). While the threshold is too small, the heuristic gets penalized until it reaches a threshold high enough to avoid mistakes. For queries that span multiple bins, we only penalize the least-updated bins to prevent a single, inaccurate bin from setting back the histogram from queries using accurate bins only. With these thresholds, we only configure initial parameters C_0 and S_0 , which we find experimentally easy to do (§5.2).

External updates. While we want to bypass the PMW when the histogram is not “ready” for a query, we still want to update the histogram with the result from the Laplace execution (R3); otherwise, the histogram will never get trained. That is the purpose of our external updates (lines 33-34 in Alg. 1). They follow a similar structure as a regular PMW update (lines 24-25 in Alg. 1), with a key difference. In vanilla PMW, the histogram is updated with the result $R2$ from Laplace

only when the SV test fails. In that case, PMW updates the relevant bins in one direction or another, depending on the sign of the error $R2 - q(h)$. For example, if the histogram is underestimating the true answer, then $R2$ will likely be higher than the histogram-based result, so we should increase the value of the bins (case $R2 > q(h)$ of line 24 in Alg. 1).

In PMW-Bypass, external updates are performed not just when the authoritative SV test finds the histogram estimation inaccurate, but also when our heuristic predicts it to be inaccurate even though it may actually be accurate. In the latter case, performing external updates in the same way as PMW updates would add bias into the histogram and forfeit its convergence guarantee. To prevent this, in PMW-Bypass, external updates are executed only when we are quite confident, based on the direct-Laplace result $R3$, that the histogram overestimates or underestimates the true result. Line 33 shows the change: the term $\tau\alpha$ is a *safety margin* that we add to the comparison between the histogram’s estimation and $R3$, to be confident that the estimation is wrong and the update warranted. This lets us prove worst-case convergence akin to PMW. Finally, like regular PMW updates, external updates reuse the already DP result $R3$, hence they do not consume any additional privacy budget beyond what was already consumed to generate $R3$.

Learning rate. In addition to the bypass option, we make another key change to PMW design for practicality. When updating a bin, we increase or decrease the bin’s value based on a learning rate parameter, lr , which determines the size of the update step taken (line 3 in Alg. 1). Prior PMW works fix learning rates that minimize theoretical convergence time, typically $\alpha/8$ [58]. However, our experiments show that larger values of lr can lead to much faster convergence, as dozens of updates may be needed to move a bin from its uniform prior to an accurate estimation. However, increasing lr beyond a certain point can impede convergence, as the updates become too coarse. Taking cue from deep learning, PMW-Bypass uses a scheduler to adjust lr over time. We start with a high lr and progressively reduce it as the histogram converges.

Guarantees. **(G1) Privacy:** PMW-Bypass preserves ϵ_G -DP across the queries it executes (Thm. A.1). **(G2) Accuracy:** PMW-Bypass is α -accurate with $1 - \beta$ probability for each query (Thm. A.3). This property stems from how we calibrate Laplace budget ϵ to α and β . This is function CALIBRATEBUDGET in Alg. 1 (lines 6-7). For n datapoints, setting $\epsilon = \frac{4 \ln(1/\beta)}{n\alpha}$ ensures that each query is answered with error at most α with probability $1 - \beta$. **(G3) Worst-case convergence:** If $lr/\alpha < \tau \leq 1/2$, then w.h.p. PMW-Bypass needs to perform at most $\frac{\ln |X|}{lr(\tau\alpha - lr)/2}$ updates (Thm. A.4). PMW-Bypass’s worst-case convergence is thus similar to PMW’s, but roughly $1/2\tau$ times slower. §5.2 confirms this empirically.

4.4 Tree-Structured PMW-Bypass

We now switch to the partitioned-database use cases, focusing on time-based partitions, as in timeseries databases, whether static or dynamic. Rather than accessing the entire




Fig. 5. Example of tree-structured histograms.

database, analysts tend to query specific time windows, such as requesting the Covid positivity rate over the past week, or the fraction of minors diagnosed with Covid in the two weeks following school reopening. This allows the opportunity to leverage DP’s parallel composition: the database is partitioned by time (say a week’s data goes in one partition), and privacy budget is consumed at the partition level. Queries can run at finer or coarser granularity, but they will consume privacy against the partition(s) containing the requested data. With this approach, a system can answer more queries under a fixed global (ϵ_G, δ_G) -DP guarantee compared to not partitioning [40, 45, 50, 56]. We implement support for partitioning and parallel composition in Turbo through a new caching object called a *tree-structured PMW-Bypass*.

Example. Fig. 5 shows an extension of the running example in §4.2, with the database partitioned by week. Denote n_i the size of each partition. A new query, Q , asks for the positivity rate over the past three weeks. How should we structure the histograms we maintain to best answer this query? One option would be to maintain *one histogram per partition* (i.e., just the leaves in the figure). To resolve Q , we query the histograms for weeks 2, 3, 4. Assume the query results in an update. Then, we need to update histograms, computing the answer with DP within our α error tolerance. Updating histograms for weeks 2, 3, and 4 requires querying the result for each of them with parallel composition. Given that Laplace $(1/n\epsilon)$ has standard deviation $\sqrt{2}/n\epsilon$, for week 4 for instance, we need noise scaled to $1/n_4\epsilon$. Thus, we consume a fairly large ϵ for an accurate query to compensate for the smaller n_4 . Another option would be to use *one histogram per range* (i.e. set of contiguous partitions), but that involves maintaining a large state that grows quadratically in the number of partitions.

Instead, our approach is to maintain a *binary-tree-structured set of histograms*, as shown in Fig. 5. For each partition, but also for a binary tree growing from the partitions, we maintain a separate histogram. To resolve Q , we split the query into two sub-queries, one running on the histogram for week 2 ([2,2]) and the other running on the histogram for the range week 3 to week 4 ([3,4]). That last sub-query would then run on a larger dataset of size $n_3 + n_4$, requiring a smaller budget consumption to reach the target accuracy.

Design. Fig. 6 shows our design. Given a query Q , we split it into sub-queries based on the histogram tree, applying the min-cuts algorithm to find the smallest set of nodes in the tree

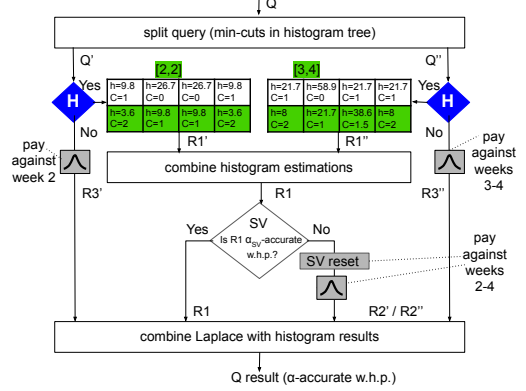


Fig. 6. Tree-structured PMW-Bypass.

that covers the requested partitions. In our example, this gives two sub-queries, Q' and Q'' , running on histograms [2,2] and [3,4], respectively. For each sub-query, we use our heuristic to decide whether to use the histogram or invoke Laplace directly. If both histograms are “ready,” we compute their estimations and combine them into one result, which we test with an SV against an accuracy goal. In our example, there are only two sub-queries, but in general there can be more, some of which will use Laplace while others use histograms. We adjust the SV’s accuracy target to an $(\alpha_{SV}, \beta_{SV})$ calibrated to the aggregation that we will need to do among the results of these different mechanisms. We pay for any Laplace’s and SV resets against the queried data partitions and finally combine Laplace results with histogram-based results. Each subquery updates the corresponding histograms of the tree (details in Alg. 2) and increments c for updated nodes.

Guarantees. (G1) *Privacy* and (G2) *accuracy* are unchanged (Thm. A.5, A.6). (G3) *Worst-case convergence*: For T partitions, if $lr/\alpha < \tau \leq 1/2$, then w.h.p. we perform at most $\frac{2T(\lceil \log T \rceil + 1) \ln |\mathcal{X}|}{\eta(\tau\alpha - \eta)/2}$ updates (Thm. A.8).

4.5 Histogram Warm-Start

An opportunity exists in streams to warm-start histograms from previously trained ones to converge faster. Prior work on PMW initialization [44] only justifies using a public dataset close to the private dataset to learn a more informed initial value for histogram bins than a uniform prior. We prove that warm-starting a histogram by copying an entire, trained histogram preserves the worst-case convergence. In Turbo, we use two procedures: for new leaf histograms, we copy the previous partition’s leaf node; for non-leaf histograms, we take the average of children histograms. We also initialize the per-bin thresholds and update counters of each node.

Guarantees. (G1) *Privacy* and (G2) *accuracy* guarantees are unchanged. (G3) *Worst-case convergence*: If there exists $\lambda \geq 1$ such that the initial histogram h_0 in Alg. 1 satisfies $\forall x \in \mathcal{X}, h_0(x) \geq \frac{1}{\lambda|\mathcal{X}|}$, then we show that each PMW-Bypass converges, albeit at a slower pace (Thm. A.9). The same properties hold for the tree.

5 Evaluation

We prototype Turbo using TimescaleDB as the underlying database and Redis for storing the histograms and exact-cache. Using two public timeseries datasets – Covid and CitiBike – we evaluate Turbo in the three use cases from §3.2. Each use case lets us do *system-wide evaluation*, answering the critical question: *Does Turbo significantly improve privacy budget consumption compared to reasonable baselines for each use case?* This corresponds to evaluating our §3.1 design goals (G5) and (G6). In addition, each setting lets us evaluate a different set of caching objects and mechanisms:

(1) Non-partitioned database: We configure Turbo with a single PMW-Bypass and Exact-Cache, letting us evaluate the PMW-Bypass object, including its empirical convergence and the impact of its heuristic and learning rate parameters.

(2) Partitioned static database: We partition the datasets by time (one partition per week) and configure Turbo with the tree-structured PMW-Bypass and Exact-Cache. This lets us evaluate the tree-structured cache.

(3) Partitioned streaming database: We configure Turbo with the tree-structured PMW-Bypass, Exact-Cache, and histogram warm-up, letting us evaluate warm-up.

As highlighting, our results show that PMW-Bypass unleashes the power of PMW, enhancing privacy budget consumption for linear queries well beyond the conventional approach of using an exact-match cache (goal (G5)). Moreover, Turbo as a whole seamlessly applies to multiple settings, with its novel tree-structured PMW-Bypass structure scoring significant benefit for timeseries workloads where database can be partitioned to leverage parallel composition (goal (G6)). Configuration of our objects and mechanisms is straightforward (goal (G7)), and we tune them based on empirical convergence rather than theoretical convergence, boosting their practical effectiveness (goal (G4)). Finally, we provide a basic runtime and memory evaluation, which shows that while Turbo performs reasonably for our datasets, further research is needed for larger-domain data.

5.1 Methodology

For each dataset, we create query workloads by (1) generating a pool of linear queries and (2) sampling queries from this pool based on a Zipfian distribution. Covid uses a completely synthetic query pool. CitiBike uses a pool based on real-user queries from prior CitiBike analyses. We use the former as a microbenchmark, the latter as a macrobenchmark.

Covid. Dataset: We take a California dataset of Covid-19 tests from 2020 that provides daily *aggregate information* of the number of Covid tests and their positivity rates for various demographic groups defined by age \times gender \times ethnicity. We combine this data with US Census data to generate a synthetic dataset that contains $n = 50,426,600$ per-person test records, each with the date and four attributes: positivity, age, gender, and ethnicity. These attributes have domain sizes of 2, 4, 2 and 8, respectively, so the dataset domain size is $N = 128$.

The dataset spans 50 weeks, so in partitioned use cases we have up to 50 partitions. *Query pool:* We create a synthetic and rich pool of correlated queries comprising all possible count queries that can be posed on Covid. This gives 34,425 unique queries, plenty for us to microbenchmark Turbo.

CitiBike. Dataset: We take a dataset of NYC bike rentals from 2018-2019, which includes information about individual rides, such as start/end date, start/end geo-location, and renter’s gender and age. The original data is too granular with 4,000 geo-locations and 100 ages, making it impractical for PMWs. Since all the real-user analyses we found consider the data at coarser granularity (e.g. broader locations and age brackets), we group geo-locations into ten neighborhoods and ages into four brackets. This yields a dataset with $n = 21,096,261$ records, domain size $N = 604,800$, and spanning 50 weeks. *Query pool:* We collect a set of pre-existing CitiBike analyses created by various individuals and made available on Public Tableau [2]. An example is here [1]. We extract 30 distinct queries, most containing ‘GROUP BY’ statements that we decompose into multiple *primitive queries* that can interact with Turbo histograms. This gives us a pool of 2,485 queries, which is smaller than Covid’s but more realistic and suitable as a macrobenchmark.

Workload generation. As is customary in caching literature [8, 18, 63], we use a Zipfian distribution to control the skewness of query distribution, which affects hit rates in the exact-match cache. From a pool of Q queries, a query of type $x \in [1, Q]$ is sampled with probability $\propto x^{-k_{\text{zipf}}}$, where $k_{\text{zipf}} \geq 0$ is the parameter that controls skewness. We evaluate with several k_{zipf} values but report only results for $k_{\text{zipf}} = 0$ (uniform) and $k_{\text{zipf}} = 1$ (skewed) for Covid. For CitiBike, we evaluate only for $k_{\text{zipf}} = 0$ to avoid reducing the small query pool further with skewed sampling. For streaming, queries arrive online with arrival times following a Poisson process; they request a window of certain size over recent timestamps. **Metrics.** • *Average cumulative budget:* the average budget consumed across all partitions. • *Systems metrics:* traditional runtime, process RAM. • *Empirical convergence:* We periodically evaluate the quality of Turbo’s histogram by running a validation workload sampled from the same query pool. We measure the accuracy of the histogram as the fraction of queries that are answered with error $\geq \alpha/2$ by the histogram. We define *empirical convergence* as the number of histogram updates necessary to reach 90% validation accuracy.

Default parameters. Unless stated otherwise, we use the following parameter values: privacy ($\epsilon_G = 10, \delta_G = 0$); accuracy ($\alpha = 0.05, \beta = 0.001$); for Covid: {learning rate lr starts from 0.25 and decays to 0.025, heuristic ($C_0 = 100, S_0 = 5$), external updates $\tau = 0.05$ }; for CitiBike: {learning rate $lr = 0.5$, heuristic ($C_0 = 5, S_0 = 1$), external updates $\tau = 0.01$ }.

5.2 Use Case (1): Non-partitioned Database

System-wide evaluation. Question 1: *In a non-partitioned database, does Turbo significantly improve privacy budget*




Fig. 7. Non-partitioned database: (a-c) system-wide evaluation (Question 1); (d) empirical convergence for PMW-Bypass vs. PMW (Question 2). (a-c) Turbo, instantiated with one PMW-Bypass and Exact-Cache, significantly improves budget consumption compared to both baselines. (d) Uses Covid $k_{\text{zipf}} = 1$. PMW-Bypass has similar empirical convergence to PMW, and both converge faster with much larger lr than anticipated by worst-case convergence.

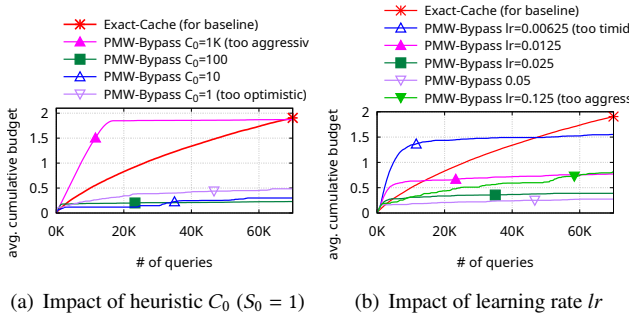


Fig. 8. Impact of parameters (Question 3). Uses Covid $k_{\text{zipf}} = 1$. Being too optimistic or pessimistic about the histogram’s state (a), or too aggressive or timid in learning from each update (b), give poor performance.

consumption compared to vanilla PMW and a simple Exact-Cache? Fig. 7(a)-7(c) show the cumulative privacy budget used by three workloads as they progress to 70K queries. Two workloads correspond to Covid, one uniform ($k_{\text{zipf}} = 0$) and one skewed ($k_{\text{zipf}} = 1$), and one uniform workload for CitiBike. Turbo surpasses both baselines across all three workloads. The improvement is enormous when compared to vanilla PMW: 15.9 – 37.4 \times ! PMW’s convergence is rapid but consumes lots of privacy; Turbo uses little privacy during training and then executes queries for free. Compared to just an Exact-Cache, the improvement is less dramatic but still significant. The greatest improvement over Exact-Cache is seen in the uniform Covid workload: 16.7 \times (Fig. 7(a)). Here, queries are relatively unique, resulting in low hit rate for the Exact-Cache. That hit rate is higher for the skewed workload (Fig. 7(b)), leaving less room for improvement for Turbo: 9.7 \times better than Exact-Cache. For CitiBike (Fig. 7(c)), the query pool is much smaller ($< 2.5K$ queries), resulting in many exact repetitions in a large workload, even if uniform. Nevertheless, Turbo gives a 1.7 \times improvement over Exact-Cache. And in this workload, Turbo outperforms PMW by 37.4 \times (omitted from figure for visualization reasons). Overall, then, Turbo significantly reduces privacy budget consumption in non-partitioned databases, achieving 1.7 – 15.9 \times improvement over the best baseline for each workload (goal **(G5)**).

PMW-Bypass evaluation. Using Covid $k_{\text{zipf}} = 1$, we microbenchmark PMW-Bypass to understand the behavior of this key Turbo component. *Question 2: Does PMW-Bypass*

converge similarly to PMW in practice? Through theoretical analysis, we have shown that PMW-Bypass achieves similar worst-case convergence to PMW, albeit at slower speed (§4.3). Fig. 7(d) compares the *empirical convergence* (defined in §5.1) of PMW-Bypass vs. PMW, as a function of the learning rate lr . We make three observations, two of which agree with theory, and the last differs. First, the results confirm the theory that (1) PMW-Bypass and PMW converge similarly, but (2) for “good” values of lr , vanilla PMW converges slightly faster: e.g., for $lr = 0.025$, PMW-Bypass converges after 1853 updates, while PMW after 944. Second, as theory suggests, very large values of lr (e.g., $lr \geq 0.4$) impede convergence in practice. Third, although theoretically, $lr = \alpha/8 = 0.00625$ is optimal for worst-case convergence, and it is commonly hard-coded in PMW protocols [58], we find that empirically, larger values of lr (e.g., $lr = 0.05$, which is 8 \times larger) give much faster convergence. This is true for both PMW and PMW-Bypass, and across all our workloads. This justifies the need to adapt and tune mechanisms based on not only theoretical but also empirical behavior (goal **(G4)**).

Question 3: How do PMW-Bypass heuristic, learning rate, and external update parameters impact consumed budget? We experimented with all parameters and found that the two most impactful are (a) C_0 , the initial threshold for the number of updates each bin involved in a query must have received to use the histogram, and (b) the learning rate. Fig. 8 shows their effects. *Heuristic C_0 (Fig. 8(a)):* Higher C_0 results in a more pessimistic assessment of histogram readiness. If it’s too pessimistic ($C_0 = 1K$), PMW is never used, so we follow a direct Laplace. If it’s too optimistic ($C_0 = 1$), errors occur too often, and the histogram’s training overpays. $C_0 = 100$ is a good value for this workload. *Learning rate lr (Fig. 8(b)):* Higher lr leads to more aggressive learning from each update. Both too aggressive ($lr = 0.125$) and too timid ($lr = 0.00625$) learning slow down convergence. Good values hover around $lr = 0.025$. Overall, only a few parameters affect performance, and even for those, performance is relatively stable around good values, making them easy to tune (goal **(G7)**).

Question 4: How does Turbo’s adaptive, per-bin heuristic compare to alternatives? We experimented with three alternative ISHISTOGRAMREADY designs that forgo either (1) the




Fig. 9. Partitioned static database: system-wide evaluation (Question 5). Turbo is instantiated with tree-structured PMW-Bypass and Exact-Cache. Turbo significantly improves budget consumption compared to both a single Exact-Cache and a tree-structured set of Exact-Caches.

per-bin granular thresholds, or (2) the adaptivity property, or (3) both. We make two observations. First, the *coarse-grained heuristics* consume more privacy budget than the fine-grained heuristics, especially on more skewed workloads, such as $k_{zipf} = 1.5$, which have less diversity so they tend to train histogram bins less uniformly. For example, a coarse-grained heuristic that uses a histogram-level count of the number of updates, with a threshold C_0 to determine when the histogram is ready to receive *any* query, consumes at best 0.7 global privacy budget on a Covid workload with $k_{zipf} = 1.5$; this is achieved when C_0 is optimally configured to a value of 2070 updates. In contrast, a fine-grained heuristic, which uses a per-bin update count with a threshold C_0 for each bin, consumes at best 0.44 global privacy budget, achieved when C_0 is set to 160 updates. Second, the *adaptive heuristics* consume similar budget as the optimally-configured, non-adaptive ones, but the former are much easier to configure, as they offer stable performance around wide ranges of the C_0 parameter. For example, when C_0 varies in range $[20, 200]$, the non-adaptive per-bin heuristic’s budget consumption varies in range $[0.44, 0.81]$ for the $k_{zipf} = 1.5$ workload, and in range $[0.31, 0.76]$ for $k_{zipf} = 1$ workload. In contrast, Turbo’s adaptive, per-bin heuristic’s budget consumption varies in much tighter ranges under the same circumstances: $[0.44, 0.52]$ and $[0.28, 0.48]$ for the $k_{zipf} = 1.5$ and $k_{zipf} = 1$ workload, respectively. Thus, Turbo’s heuristic is the best of these options.

5.3 Use Case (2): Partitioned Static Database

System-wide evaluation. *Question 5: In a partitioned static database, does Turbo significantly improve privacy budget consumption, compared to a single Exact-Cache and a tree-structured set of Exact-Caches?* We divide each database into 50 partitions and select a random contiguous window of 1 to 50 partitions for each query. We adjust the (C_0, S_0) heuristic parameters to $(50, 1)$ for Covid and $(1, 1)$ for CitiBike. Fig. 9(a)-9(c) show the average budget consumed per partition up to 300K queries. Compared to the static case, Turbo can now support more queries under $\epsilon_G = 10$ thanks to parallel composition: each query only consumes privacy from the accessed partitions. Turbo further divides privacy budget consumption by $1.9 - 4.7\times$ compared to the best-performing baseline for each workload, demonstrating its effectiveness as a caching strategy for the static partitioned use case.

Tree structure evaluation. *Question 6: When does the tree structure for histograms outperform a flat structure that maintains one histogram per partition?* We vary the average size of the windows requested by queries from 1 to 50 partitions based on a Gaussian distribution with std-dev 5. We find the tree structure for histograms is beneficial when queries tend to request more partitions (25 partitions or more). Because the tree structure maintains more histograms than the flat structure, it fragments the query workload more, resulting in fewer histogram updates per histogram and more use of direct-Laplace. The tree’s advantage in combining fewer results makes up for this privacy overhead caused by histogram maintenance when queries tend to request larger windows of partitions, while the linear structure is more justified when queries tend to request smaller windows of partitions.

5.4 Use Case (3): Partitioned Streaming Database

System-wide evaluation. *Question 7: In streaming databases partitioned by time, does Turbo significantly improve privacy budget consumption compared to baselines? Does warm-start help?* Fig. 10(a)-10(c) show Turbo’s budget consumption compared to the baselines. The experiments simulate a streaming database, where partitions arrive over time and queries request the latest P partitions, with P chosen uniformly at random between 1 and the number of available partitions. Turbo outperforms both baselines significantly for all workloads, particularly when warm-start is enabled. Without warm-start, Turbo improves performance by $1.5 - 3.5\times$ at the end of the workload. With warm-start, Turbo gives $1.9 - 5.4\times$ improvement over the best baseline for each workload, showing its effectiveness for the streaming use case. When there is a large variety of unique queries the tree-structured Exact-Cache has a significantly better hit-rate than the Exact-Cache baseline and performs better (Fig. 10(a)). In Fig. 10(b) and 10(c) the query pool is considerably smaller. Both baselines have a good enough hit-rate while the tree-structured Exact-Cache needs to consume more privacy budget to compensate for the aggregation error which makes it perform worse. This concludes our evaluation across use cases (goal **(G6)**).

5.5 Runtime and Memory Evaluation

Question 8: What are Turbo’s runtime and memory bottlenecks? We evaluate Turbo’s runtime and memory consumption to identify areas of improvement. Fig. 10(d) shows the average runtime of Turbo’s main execution paths in a non-partitioned database. The Exact-Cache hit path is the cheapest




Fig. 10. (a-c) Partitioned streaming database: system-wide consumed budget (Question 7); (d) PMW-Bypass runtime in non-partitioned setting (Question 8). (a-c) Turbo is instantiated with tree-structured PMW-Bypass and Exact-Cache, with and without warm-start. (d) Uses Covid $k_{zipf} = 1$ and one Exact-Cache and PMW-Bypass. Shows execution runtime for different execution paths. Most expensive is when the SV test fails.

and the other paths are more expensive. Histogram operations are the bottlenecks in CitiBike due to the larger domain size (N), while query execution in TimescaleDB is the bottleneck in Covid due to the larger database size (n). The R1 path is similar across the two datasets because their distinct bottlenecks compensate. Failing the SV check (output path R2) is the costliest path for both datasets due to the extra operations needed to update the heuristic’s per-bin thresholds. We also conduct an experiment in the partitioned streaming case and find the same bottlenecks: TimescaleDB for Covid, histogram operations for CitiBike. Finally, we report Turbo’s memory consumption in the streaming case with 50 partitions: 5.21MB for Covid and 1.43GB for CitiBike. For context, the raw datasets occupy on disk 600MB and 795MB, respectively. Thus, Turbo’s memory overhead is significant and it is caused by the PMWs. The next section discusses this limitation and proposes potential directions to address it.

6 Discussion

We discuss several of Turbo’s strengths and weaknesses. Turbo provides benefits when queries overlap in the data they access, i.e., new queries access histogram bins that have been accessed by past queries. The functions computed atop these bins can differ among queries (e.g., the new query can compute an average while all the past ones computed count fractions). If there is no data overlap in the queries, then Turbo does not give any benefit and comes with memory/computational costs. This is typical for caching systems: they only help if the workload has some level of locality.

A key strength in Turbo is its support for dynamic workloads, both new queries and new data arriving in the system. First, Turbo adapts seamlessly to changing queries. In the worst case, the new queries will access completely “untrained” regions within a histogram. Our heuristic will detect this and trigger a new cycle of external updates. In more moderate cases, the workload will touch a mix of “trained” and “untrained” regions. This will yield a mix of hits and misses in the heuristic, and Turbo will use just the right amount of privacy budget to adapt to these slower workload changes. Second, thanks to histogram warm-start, Turbo adapts to new data partitions arriving into the system with minimal privacy budget consumption: as new partitions arrive, their histograms are initialized from past ones and then fine-tuned for the new data

by a few external updates. This way, the new histograms will quickly start serving query answers for free, conserving privacy budget. Still, there is a limitation: while we support new data arriving in the system, we do not support updates on past data; such updates would result in our heuristics predicting less accurately when the histogram can answer a query, and thus in more expensive SV failures.

By far, Turbo’s biggest limitation is the memory consumed to maintain the PMW histograms. Each histogram is a RedisAI vector whose size grows with data domain size N , i.e., exponentially in data domain dimension d (N and d are defined in Section 4.1). With T partitions and k queries, Turbo maintains a binary tree of such histograms, which means it stores $\approx 2TN$ scalar values. By comparison, the Tree Exact-Cache baseline stores at most $\log(T)k$ scalar values, a much lower memory consumption. This impacts not only the scale of the datasets that can be handled with Turbo, but also the runtime performance of Turbo-mediated queries. Indeed, as shown in the preceding section, histogram operations for CitiBike are the bottleneck in runtime due to the relatively high domain size. Some techniques have previously been proposed to address this rather fundamental challenge for PMW [30]. However, for even larger-scale deployments, we believe that it will be worth considering PMW alternatives that may not offer as compelling convergence guarantees as PMW but which are much more lightweight. One example may be the relaxed adaptive projection (RAP) [9], which builds a lightweight representation of the dataset by learning a small subset of representative data points using gradient-descent. One would have to be willing to forfeit the theoretical convergence guarantees to use this mechanism, and to develop an adaptive version of RAP to support realistic systems settings involving dynamic workloads and data. Even so, some of the core concepts we have proposed in this paper may transfer to this new design, including passing RAP-based estimations through an SV to ensure result accuracy while incorporating a heuristic-based bypass to avoid expensive failures in the SV.

Finally, we touch on several potential vulnerabilities. First, an adversary may craft queries that consume budget by generating cache misses. The convergence proofs in §A.5 provide a bound on how much such queries can affect budget consumption when a straightforward cutoff parameter is configured

upfront. Second, response time can be a side-channel, which we leave out of scope but should be addressed in the future. Third, n , the number of elements in the database (or in each partition), is considered public knowledge. This can leak information and should be addressed by consuming some of the budget to compute n privately, as done in [40].

7 Related Work

This paper presents the first design, implementation, and evaluation for a *general, effective, and accurate DP-caching system for interactive DP-SQL systems*. In computer systems, caching is a heavily-explored topic, with numerous algorithms and implementations [11, 16, 64], some pervasively used in processors, operating systems, databases, and more. However, traditional forms of caching differ significantly from DP caching, justifying the need for a specialized approach for DP. The primary purposes of traditional caching are to conserve CPU and to improve throughput and latency; for these purposes, existing caches can be readily reused in DP systems. However, DP caching aims to conserve privacy budget, which requires a new design to be truly effective. For example, layering Redis on a DP database to cache query results would save CPU, but for privacy it would be equivalent to the “Exact-Cache” baseline that our evaluation shows is less effective than Turbo. This paper thus builds upon general traditional caching concepts – such as the two-layer design, the principle of generality in supporting multiple workloads – but develops a cache specialized in conserving DP budget.

To our knowledge, no existing DP system incorporates such a specialized caching system. Most DP systems do not incorporate caching capabilities at all [7, 12, 37, 50, 53, 55]; [62] explicitly leaves the design of an effective DP cache for future work. Some DP systems incorporate what amounts to an Exact-Cache by deterministically generating the same noise upon the arrival of the same query. Three systems consider more sophisticated mechanisms for DP result reuse: PrivateSQL [38], Chorus [36], and CacheDP [48]. But the result reuse components in these systems suffer from such significant limitations that they cannot be considered general and effective caching designs. **PrivateSQL** [38] takes a batch of “representative” offline queries and precomputes a private synopsis that answers them all. If new queries arrive (online), PrivateSQL uses the synopsis to answer them in a best-effort way, without accuracy guarantees. It does not learn on-the-fly from them, so it is unsuited for online workloads and does not support data streams. **Chorus** [36] provides a trivialized implementation of MWEM, a variant of PMW, however the implementation only works for databases with a *single attribute*. The paper does not evaluate the MWEM-based implementation, nor integrates it as a caching layer. **CacheDP** [48] is an interactive DP query engine and has a built-in DP cache that answers queries using the Matrix Mechanism [43]. Our experience with the CacheDP code suggests that it is not a general, effective, or accurate caching

layer for DP databases. First, CacheDP’s implementation only scales to a few attributes and does not support parallel composition on data partitions; this suggests that it is not general enough to support a variety of workloads. Second, the “Tree Exact-Cache” baseline with which we compare in evaluation matches, to our understanding, the CacheDP design while scaling to the higher-dimension datasets and streaming workloads we evaluate against. Our evaluation shows Turbo more effective than Tree Exact-Cache.

While DP caching are under-explored in systems, the topic of optimizing global privacy budget for a query workload is heavily explored in theory. Approaches include generating synthetic datasets or histograms that can answer certain classes of queries, such as linear queries, with accuracy guarantees and no further privacy consumption [9, 13, 30, 31, 44, 60]; and optimizing privacy consumption over a batch of queries by adapting the noise distribution to properties of the queries [42, 43, 49]. Apart from PMW [31], all these methods operate in the offline setting, where queries are known upfront. This setting is unrealistic, as discussed in §3.2.

All of the theory works cited above, including PMW, suffer from another limitation: they operate on static datasets and do not support new data arriving into the system. PMWG [20] is an extension of PMW for dynamic “growing” databases, but operates in a setting where all queries request the *entire database*. This precludes the use of parallel composition for queries that access less than the entire database, such as queries over windows of time. Other algorithms focus on continuously releasing specific statistics over a stream, such as the streaming counter [35] that inspired our tree structure, and extensions to top-k and histogram queries [14]. These works do not support arbitrary linear queries, and they answer all predefined queries at every time step while we only pay budget for queries that are actually posed by analysts.

8 Conclusion

Turbo is a caching layer for differentially-private databases that increases the number of linear queries that can be answered accurately with a fixed privacy guarantee. It employs a PMW, which learns a histogram representation of the dataset from prior query results and can answer future linear queries at no additional privacy cost once it has converged. To enhance the practical effectiveness of PMWs, we bypass them during the privacy-expensive training phase and only switch to them once they are ready. This transforms PMWs from ineffective to very effective compared to simpler cache designs.

9 Acknowledgments

We thank our shepherd Andreas Haeberlen and the anonymous reviewers. We thank Junfeng Yang for feedback throughout the project. We thank Heeyun Kim, Hailey Onweller, Sally Wang, Yucheng Wu, Chris Yoon for indirect contributions to prototype and evaluation. The work was supported by NSF EEC-2133516, CNS-2104292, Sloan, Microsoft, and Google faculty fellowships, and Google Cloud credits.

References

- [1] Citibike tableau story. <https://public.tableau.com/app/profile/james.jeffrey/viz/CitiBikeRideAnalyzer/CitiBikeRdeAnalyzer>. Accessed: 2023-04-13.
- [2] Tableau. <https://public.tableau.com/app/discover>. Accessed: 2023-04-13.
- [3] Citibike system data. <https://www.citibikenyc.com/system-data>, 2018.
- [4] NOT-OD-17-110: Request for Comments: Proposal to Update Data Management of Genomic Summary Results Under the NIH Genomic Data Sharing Policy, Apr. 2023. [Online; accessed 17. Apr. 2023].
- [5] J. M. Abowd, R. Ashmead, R. Cumings-Menon, S. Garfinkel, M. Heineck, C. Heiss, R. Johns, D. Kifer, P. Leclerc, A. Machanavajhala, et al. The 2020 census disclosure avoidance system topdown algorithm. *Harvard Data Science Review*, (Special Issue 2), 2022.
- [6] S. Agarwal, B. Mozafari, A. Panda, H. Milner, S. Madden, and I. Stoica. BlinkDB: queries with bounded errors and bounded response times on very large data. In *Proceedings of the 8th ACM European conference on computer systems*, pages 29–42, 2013.
- [7] K. Amin, J. Gillenwater, M. Joseph, A. Kulesza, and S. Vassilvitskii. Plume: Differential Privacy at Scale. *arXiv*, Jan. 2022.
- [8] B. Atikoglu, Y. Xu, E. Frachtenberg, S. Jiang, and M. Paleczny. Workload analysis of a large-scale key-value store. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE joint international conference on Measurement and Modeling of Computer Systems*, pages 53–64, 2012.
- [9] S. Aydöre, W. Brown, M. Kearns, K. Kenthapadi, L. Melis, A. Roth, and A. A. Siva. Differentially private query release through adaptive projection. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 457–467. PMLR, 2021.
- [10] S. Bavadekar, A. Dai, J. Davis, D. Desfontaines, I. Eckstein, K. Everett, A. Fabrikant, G. Flores, E. Gabrilovich, K. Gadepalli, S. Glass, R. Huang, C. Kamath, D. Kraft, A. Kumok, H. Marfatia, Y. Mayer, B. Miller, A. Pearce, I. M. Perera, V. Ramachandran, K. Raman, T. Roessler, I. Shafraan, T. Shekel, C. Stanton, J. Stimes, M. Sun, G. Wellenius, and M. Zoghi. Google COVID-19 Search Trends Symptoms Dataset: Anonymization Process Description (version 1.0). *arXiv*, Sept. 2020.
- [11] N. Beckmann, H. Chen, and A. Cidon. Lhd: Improving cache hit rate by maximizing hit density. In *Proceedings of the 15th USENIX Conference on Networked Systems Design and Implementation, NSDI'18*, page 389–403, USA, 2018. USENIX Association.
- [12] S. Berghel, P. Bohannon, D. Desfontaines, C. Estes, S. Haney, L. Hartman, M. Hay, A. Machanavajhala, T. Magerlein, G. Miklau, A. Pai, W. Sexton, and R. Shrestha. Tumult analytics: a robust, easy-to-use, scalable, and expressive framework for differential privacy. *CoRR*, abs/2212.04133, 2022.
- [13] A. Blum, K. Ligett, and A. Roth. A learning theory approach to non-interactive database privacy. In C. Dwork, editor, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 609–618. ACM, 2008.
- [14] A. R. Cardoso and R. Rogers. Differentially private histograms under continual observation: Streaming selection into the unknown. In G. Camps-Valls, F. J. R. Ruiz, and I. Valera, editors, *International Conference on Artificial Intelligence and Statistics, AISTATS 2022, 28-30 March 2022, Virtual Event*, volume 151 of *Proceedings of Machine Learning Research*, pages 2397–2419. PMLR, 2022.
- [15] T.-H. H. Chan, E. Shi, and D. Song. Private and continual release of statistics. *ACM Trans. Inf. Syst. Secur.*, 14(3), nov 2011.
- [16] H. Che, y. Tung, and Z. Wang. Hierarchical web caching systems: Modeling, design and experimental results. *Selected Areas in Communications, IEEE Journal on*, 20:1305 – 1314, 10 2002.
- [17] A. Cohen and K. Nissim. Linear program reconstruction in practice. *Journal of Privacy and Confidentiality*, 2020.
- [18] B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears. Benchmarking cloud serving systems with YCSB. In *Proceedings of the 1st ACM symposium on Cloud computing*, pages 143–154, 2010.
- [19] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Oct. 2001.
- [20] R. Cummings, S. Krehbiel, K. A. Lai, and U. Tantipongpipat. Differential privacy for growing databases. *CoRR*, abs/1803.06416, 2018.
- [21] Y.-A. De Montjoye, C. A. Hidalgo, M. Verleyesen, and V. D. Blondel. Unique in the crowd: The privacy bounds of human mobility. *Scientific reports*, 2013.
- [22] D. Desfontaines. Real world DP use-cases. <https://desfontain.es/privacy/real-world-differential-privacy.html>. Accessed: 2023-04-13.
- [23] I. Dinur and K. Nissim. Revealing information while preserving privacy. 2003.
- [24] J. Dong, A. Roth, and W. J. Su. Gaussian differential privacy. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 2022.
- [25] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proc. of the Theory of Cryptography Conference (TCC)*, 2006.
- [26] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [27] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 2014.
- [28] S. R. Ganta, S. Kasiviswanathan, and A. Smith. Composition attacks and auxiliary information in data privacy. 2008.
- [29] S. Garfinkel, J. M. Abowd, and C. Martindale. Understanding database reconstruction attacks on public data. *Communications of the ACM*, 2019.
- [30] M. Hardt, K. Ligett, and F. Mcsherry. A simple and practical algorithm for differentially private data release. In F. Pereira, C. Burges, L. Bottou, and K. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012.
- [31] M. Hardt and G. N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 61–70, 2010.
- [32] M. Hardt and G. N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *Symposium on Foundations of Computer Science*, 2010.
- [33] N. Homer, S. Szeling, M. Redman, D. Duggan, W. Tembe, J. Muehling, J. V. Pearson, D. A. Stephan, S. F. Nelson, and D. W. Craig. Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays. *PLoS Genetics*, 2008.
- [34] J. Hsu, M. Gaboardi, A. Haeberlen, S. Khanna, A. Narayan, B. Pierce, and A. Roth. Differential privacy: An economic method for choosing epsilon. volume 2014, 07 2014.
- [35] T.-H. Hubert Chan, E. Shi, and D. Song. Private and continual release of statistics. In S. Abramsky, C. Gavoille, C. Kirchner, F. Meyer auf der Heide, and P. G. Spirakis, editors, *Automata, Languages and Programming*, pages 405–417, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- [36] N. Johnson, J. P. Near, J. M. Hellerstein, and D. Song. Chorus: a programming framework for building scalable differential privacy mechanisms. In *2020 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 535–551, 2020.
- [37] N. M. Johnson, J. P. Near, and D. Song. Towards practical differential privacy for SQL queries. *Proc. VLDB Endow.*, 11(5):526–539, 2018.
- [38] I. Kotsogiannis, Y. Tao, X. He, M. Fanaeepour, A. Machanavajhala, M. Hay, and G. Miklau. Privatesql: A differentially private sql query engine. *Proc. VLDB Endow.*, 12(11):1371–1384, jul 2019.

- [39] M. Lécuyer. Practical Privacy Filters and Odometers with Rényi Differential Privacy and Applications to Differentially Private Deep Learning. In *arXiv*, 2021.
- [40] M. Lécuyer, R. Spahn, K. Vodrahalli, R. Geambasu, and D. Hsu. Privacy Accounting and Quality Control in the Sage Differentially Private ML Platform. In *Proc. of the ACM Symposium on Operating Systems Principles (SOSP)*, 2019.
- [41] M. Lecuyer, R. Spahn, K. Vodrahalli, R. Geambasu, and D. Hsu. Privacy accounting and quality control in the sage differentially private ML platform. Online Supplements (also available on <https://arxiv.org/abs/1909.01502>), 2019.
- [42] C. Li, M. Hay, G. Miklau, and Y. Wang. A data- and workload-aware query answering algorithm for range queries under differential privacy. *Proc. VLDB Endow.*, 7(5):341–352, 2014.
- [43] C. Li, G. Miklau, M. Hay, A. McGregor, and V. Rastogi. The matrix mechanism: optimizing linear counting queries under differential privacy. *VLDB J.*, 24(6):757–781, 2015.
- [44] T. Liu, G. Vietri, T. Steinke, J. Ullman, and S. Wu. Leveraging public data for practical private query release. In M. Meila and T. Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 6968–6977. PMLR, 18–24 Jul 2021.
- [45] T. Luo, M. Pan, P. Tholoniati, A. Cidon, R. Geambasu, and M. Lécuyer. Privacy budget scheduling. In *15th USENIX Symposium on Operating Systems Design and Implementation (OSDI 21)*, pages 55–74. USENIX Association, July 2021.
- [46] M. Lyu, D. Su, and N. Li. Understanding the sparse vector technique for differential privacy. *Proc. VLDB Endow.*, 10(6):637–648, 2017.
- [47] X. Lyu. Composition theorems for interactive differential privacy. *CoRR*, abs/2207.09397, 2022.
- [48] M. Mazmudar, T. Humphries, J. Liu, M. Rafuse, and X. He. Cache me if you can: Accuracy-aware inference engine for differentially private data exploration. *Proc. VLDB Endow.*, 16(4):574–586, 2022.
- [49] R. McKenna, G. Miklau, M. Hay, and A. Machanavajjhala. Optimizing error of high-dimensional statistical queries under differential privacy. *Proc. VLDB Endow.*, 11(10):1206–1219, 2018.
- [50] F. McSherry. Privacy integrated queries: an extensible platform for privacy-preserving data analysis. In U. Çetintemel, S. B. Zdonik, D. Kossmann, and N. Tatbul, editors, *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2009, Providence, Rhode Island, USA, June 29 - July 2, 2009*, pages 19–30. ACM, 2009.
- [51] I. Mironov. Rényi Differential Privacy. In *Computer Security Foundations Symposium (CSF)*, 2017.
- [52] A. Narayanan and V. Shmatikov. Robust de-anonymization of large sparse datasets. In *Proc. of IEEE Symposium on Security and Privacy (S&P)*, 2008.
- [53] R. Rogers, S. Subramaniam, S. Peng, D. Durfee, S. Lee, S. K. Kancha, S. Sahay, and P. Ahammad. LinkedIn’s audience engagements api: A privacy preserving data analytics system at scale. *arXiv preprint arXiv:2002.05839*, 2020.
- [54] R. M. Rogers, A. Roth, J. Ullman, and S. Vadhan. Privacy odometers and filters: Pay-as-you-go composition. 2016.
- [55] E. Roth, H. Zhang, A. Haeberlen, and B. C. Pierce. Orchard: Differentially private analytics at scale. In *14th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2020, Virtual Event, November 4-6, 2020*, pages 1065–1081. USENIX Association, 2020.
- [56] I. Roy, S. T. Setty, A. Kilzer, V. Shmatikov, and E. Witchel. Airavat: Security and privacy for MapReduce. In *Proc. of the USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2010.
- [57] J. Smith, H. J. Asghar, G. Gioiosa, S. Mrabet, S. Gaspers, and P. Tyler. Making the most of parallel composition in differential privacy. *Proc. Priv. Enhancing Technol.*, 2022(1):253–273, 2022.
- [58] S. Vadhan. The Complexity of Differential Privacy. In *Tutorials on the Foundations of Cryptography*, pages 347–450. Springer, Cham, Switzerland, Apr. 2017.
- [59] S. P. Vadhan and W. Zhang. Concurrent composition theorems for all standard variants of differential privacy. *CoRR*, abs/2207.08335, 2022.
- [60] G. Vietri, G. Tian, M. Bun, T. Steinke, and Z. S. Wu. New oracle-efficient algorithms for private synthetic data release. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 9765–9774. PMLR, 2020.
- [61] L. Wasserman and S. Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 2010.
- [62] R. J. Wilson, C. Y. Zhang, W. Lam, D. Desfontaines, D. Simmons-Marengo, and B. Gipson. Differentially private sql with bounded user contribution. *Proceedings on Privacy Enhancing Technologies*, 2020(2):230–250, 2020.
- [63] J. Yang, Y. Yue, and K. Rashmi. A large scale analysis of hundreds of in-memory cache clusters at Twitter. In *Proceedings of the 14th USENIX Conference on Operating Systems Design and Implementation*, pages 191–208, 2020.
- [64] J. Yang, Y. Yue, and K. V. Rashmi. A large-scale analysis of hundreds of in-memory key-value cache clusters at twitter. *ACM Trans. Storage*, 17(3), aug 2021.
- [65] Y. Zhu and Y.-X. Wang. Improving Sparse Vector Technique with Rényi Differential Privacy. *Advances in Neural Information Processing Systems*, 33:20249–20258, 2020.