# Faster Private Release of Marginals on Small Databases* 

Karthekeyan Chandrasekaran<br>Justin Thaler Jonathan Ullman Andrew Wan

July 29, 2018


#### Abstract

We study the problem of answering $k$-way marginal queries on a database $D \in\left(\{0,1\}^{d}\right)^{n}$, while preserving differential privacy. The answer to a $k$-way marginal query is the fraction of the database's records $x \in\{0,1\}^{d}$ with a given value in each of a given set of up to $k$ columns. Marginal queries enable a rich class of statistical analyses on a dataset, and designing efficient algorithms for privately answering marginal queries has been identified as an important open problem in private data analysis. For any $k$, we give a differentially private online algorithm that runs in time $$
\text { poly }\left(n, \min \left\{\exp \left(d^{1-\Omega(1 / \sqrt{k})}\right), \exp \left(d / \log ^{0.99} d\right)\right\}\right)
$$ per query and answers any (adaptively chosen) sequence of poly $(n) k$-way marginal queries with error at most $\pm 0.01$ on every query, provided $n \gtrsim d^{0.51}$. To the best of our knowledge, this is the first algorithm capable of privately answering marginal queries with a non-trivial worst-case accuracy guarantee for databases containing poly $(d, k)$ records in time $\exp (o(d))$. Our algorithm runs the private multiplicative weights algorithm (Hardt and Rothblum, FOCS '10) on a new approximate polynomial representation of the database.

We derive our representation for the database by approximating the OR function restricted to low Hamming weight inputs using low-degree polynomials with coefficients of bounded $L_{1}$ norm. In doing so, we show new upper and lower bounds on the degree of such polynomials, which may be of independent approximation-theoretic interest. First, we construct a polynomial that approximates the $d$-variate $O R$ function on inputs of Hamming weight at most $k$ such that the degree of the polynomial is at most $d^{1-\Omega(1 / \sqrt{k})}$ and the $L_{1}$-norm of its coefficient vector is $d^{0.01}$. Then we show the following lower bound that exhibits the tightness of our approach: for any $k=o(\log d)$, any polynomial whose coefficient vector has $L_{1}$-norm poly $(d)$ that pointwise approximates the $d$-variate OR function on all inputs of Hamming weight at most $k$ must have degree $d^{1-O(1 / \sqrt{k})}$.


[^0]
## 1 Introduction

Consider a database $D \in\left(\{0,1\}^{d}\right)^{n}$ in which each of the $n(=|D|)$ rows corresponds to an individual's record, and each record consists of $d$ binary attributes. The goal of privacy-preserving data analysis is to enable rich statistical analyses on the database while protecting the privacy of the individuals. In this work, we seek to achieve differential privacy [DMNS06, which guarantees that no individual's data has a significant influence on the information released about the database.

One of the most important classes of statistics on a dataset is its marginals. A marginal query is specified by a set $S \subseteq[d]$ and a pattern $t \in\{0,1\}^{|S|}$. The query asks, "What fraction of the individual records in $D$ has each of the attributes $j \in S$ set to $t_{j}$ ?" A major open problem in privacy-preserving data analysis is to efficiently release a differentially private summary of the database that enables analysts to answer each of the $3^{d}$ marginal queries. A natural subclass of marginals are $k$-way marginals, the subset of marginals specified by sets $S \subseteq[d]$ such that $|S| \leq k$.

Privately answering marginal queries is a special case of the more general problem of privately answering counting queries on the database, which are queries of the form, "What fraction of individual records in $D$ satisfy some property $q$ ?" Early work in differential privacy DN03, BDMN05, DMNS06 showed how to privately answer any set of counting queries $\mathcal{Q}$ approximately by perturbing the answers with appropriately calibrated noise, ensuring good accuracy (say, within $\pm .01$ of the true answer) provided $|D| \gtrsim|\mathcal{Q}|^{1 / 2}$.

However, in many settings data is difficult or expensive to obtain, and the requirement that $|D| \gtrsim|\mathcal{Q}|^{1 / 2}$ is too restrictive. For instance, if the query set $\mathcal{Q}$ includes all $k$-way marginal queries then $|\mathcal{Q}| \geq d^{\Theta(k)}$, and it may be impractical to collect enough data to ensure $|D| \gtrsim|\mathcal{Q}|^{1 / 2}$, even for moderate values of $k$. Fortunately, a remarkable line of work initiated by Blum et al. [BLR08] and continuing with [DNR ${ }^{+} 09$, DRV10, RR10, HR10, HLM12, GRU12, JT12, has shown how to privately release approximate answers to any set of counting queries, even when $|\mathcal{Q}|$ is exponentially larger than $|D|$. For example, the online private multiplicative weights algorithm of Hardt and Rothblum HR10 gives accurate answers to any (possibly adaptively chosen) sequence of queries $\mathcal{Q}$ provided $|D| \gtrsim \sqrt{d} \log |\mathcal{Q}|$. Hence, if the sequence consists of all $k$-way marginal queries, then the algorithm will give accurate answers provided $|D| \gtrsim k \sqrt{d}$. Unfortunately, all of these algorithms have running time at least $2^{d}$ per query, even in the simplest setting where $\mathcal{Q}$ is the set of 2 -way marginals.

Given this state of affairs, it is natural to seek efficient algorithms capable of privately releasing approximate answers to marginal queries even when $|D| \ll d^{k}$. The most efficient algorithm known for this problem, due to Thaler, Ullman, and Vadhan [TUV12] (building on the work of Hardt, Rothblum, and Servedio HRS12]) runs in time $d^{O(\sqrt{k})}$ and releases a summary from which an analyst can compute the answer to any $k$-way marginal query in time $d^{O(\sqrt{k})}$.

Even though $|D|$ can be much smaller than $|\mathcal{Q}|^{1 / 2}$, a major drawback of this algorithm and other efficient algorithms for releasing marginals (e.g. [GHRU11, CKKL12, HRS12, FK13, DNT13]) is that the database still must be significantly larger than $\tilde{\Theta}(k \sqrt{d})$, which we know would suffice for inefficient algorithms. Recent experimental work of Hardt, Ligett, and McSherry HLM12] demonstrates that for some databases of interest, even the $2^{d}$-time private multiplicative weights algorithm is practical, and also shows that more efficient algorithms based on adding independent noise do not provide good accuracy for these databases. Motivated by these findings, we believe that an important approach to designing practical algorithms is to achieve a minimum database size comparable to that of private multiplicative weights, and seek to optimize the running time of the algorithm as much as possible. In this paper we give the first algorithms for privately answering marginal queries for this parameter regime.

| Reference | Running Time per Query | Database Size |
| :---: | :---: | :---: |
| DN03, DN04,,BDMN05, DMNS06] | $O(1)$ | $\widetilde{O}\left(d^{k / 2}\right)$ |
| [HR10,GRU12 | $2^{O(d)}$ | $\widetilde{O}(k \sqrt{d})$ |
| [HRS12,TUV12] | $d^{O(\sqrt{k})}$ | $d^{O(\sqrt{k})}$ |
| [DNT13] | $d^{O(k)}$ | $\widetilde{O}\left(d^{[k / 2] / 2}\right)$ |
| This paper | $2^{O\left(d / \log ^{0.99} d\right)}$ | $k d^{0.5+o(1)}$ |
| This paper | $2^{d^{1-\Omega(1 / \sqrt{k})}}$ | $\widetilde{O}\left(k d^{.51}\right)$ |

Table 1: Summary of prior results on differentially private release of $k$-way marginals with error $\pm 0.01$ on every marginal. Note that the running time ignores dependence on the database size, privacy parameters, and the time required to evaluate the query non-privately.

### 1.1 Our Results

In this paper we give faster algorithms for privately answering marginal queries on databases of size $\widetilde{O}\left(d^{0.51} / \varepsilon\right)$, which is nearly the smallest a database can be while admitting any differentially private approximation to marginal queries [BUV13].

Theorem 1.1. There exists a constant $C>0$ such that for every $k, d, n \in \mathbb{N}, k \leq d$, and every $\varepsilon, \delta>0$, there is an $(\varepsilon, \delta)$-differentially private online algorithm that, on input a database $D \in$ $\left(\{0,1\}^{d}\right)^{n}$, runs in time

$$
\text { poly }\left(n, \min \left\{\exp \left(d^{1-1 / C \sqrt{k}}\right), \exp \left(d / \log ^{0.99} d\right)\right\}\right)
$$

per query and answers any sequence $\mathcal{Q}$ of (possibly adaptively chosen) $k$-way marginal queries on $D$ up to an additive error of at most $\pm 0.01$ on every query with probability at least 0.99 , provided that $n \geq C d^{0.51} \log |\mathcal{Q}| \log (1 / \delta) / \varepsilon$.

When $k$ is much smaller than $d$, it may be useful to view our algorithm as an offline algorithm for releasing answers to all $k$-way marginal queries. This offline algorithm can be obtained simply by requesting answers to each of the $d^{\Theta(k)}$ distinct $k$-way marginal queries from the online mechanism. In this case we obtain the following corollary.

Corollary 1.2. There exists a constant $C>0$ such that for every $k, d, n \in \mathbb{N}, k=O(d / \log d)$, and every $\varepsilon, \delta>0$, there is an $(\varepsilon, \delta)$-differentially private offline algorithm that, on input a database $D \in\left(\{0,1\}^{d}\right)^{n}$, runs in time

$$
\text { poly }\left(n, \min \left\{\exp \left(d^{1-1 / C \sqrt{k}}\right), \exp \left(d / \log ^{0.99} d\right)\right\}\right)
$$

and, with probability at least 0.99, releases answers to every $k$-way marginal query on $D$ up to an additive error of at most $\pm 0.01$, provided that $n \geq C k d^{0.51} \log (1 / \delta) / \varepsilon$.

Here $\binom{d}{\leq k}:=\sum_{i=0}^{k}\binom{d}{i}$, and the number of $k$-way marginals on $\{0,1\}^{d}$ is bounded by a polynomial in this quantity. See Table 1 for a comparison of relevant results on privately answering marginal queries.

## Remarks.

1. When $k=\Omega\left(\log ^{2} d\right)$, the minimum database size requirement can be improved to $n \geq$ $C k d^{0.5+o(1)} \log (1 / \delta) / \varepsilon$, but we have stated the theorems with a weaker bound for simplicity. (Here $C>0$ is a universal constant and the $o(1)$ is with respect to $d$.)
2. Our algorithm can be modified so that instead of releasing approximate answers to each $k$ way marginal explicitly, it releases a summary of the database of size $\widetilde{O}\left(k d^{0.01}\right)$ from which an analyst can compute an approximate answer to any $k$-way marginal in time $\widetilde{O}\left(k d^{1.01}\right)$.

A key ingredient in our algorithm is a new approximate representation of the database using polynomial approximations to the $d$-variate OR function restricted to inputs of Hamming weight at most $k$. For any such polynomial, the degree determines the runtime of our algorithm, while the $L_{1}$-weight of the coefficient vector determines the minimum required database size. Although low-degree low $L_{1}$-weight polynomial approximations to the OR function have been studied in the context of approximation theory and learning theory [STT12], our setting requires an approximation only over a restricted subset of the inputs. When the polynomial needs to approximate the OR function only on a subset of the inputs, is it possible to reduce the degree and $L 1$-weight (in comparison to [STT12]) of the polynomial?

Our main technical contribution addresses this variant of the polynomial approximation problem. We believe that our construction of such polynomials (Theorem (1.3) as well as the lower bound (Theorem 1.4) could be of independent approximation-theoretic interest. The following theorem shows a construction of polynomials that achieve better degree and $L_{1}$-weight in comparison to STT12] for small values of $k$. Let $\mathrm{OR}_{d}:\{-1,1\}^{d} \rightarrow\{-1,1\}$ denote the OR function on $d$ variables with the convention that -1 is TRUE, and for any vector $x \in\{-1,1\}^{d}$, let $|x|$ denote the number of coordinates of $x$ equal to -1 .

Theorem 1.3. Let $k \in[d]$. For some constant $C>0$, there exists a polynomial $p$ such that

$$
\begin{equation*}
\left|p(x)-O R_{d}(x)\right| \leq 1 / 400 \text { for every } x \in\{-1,1\}^{d}:|x| \leq k, \tag{i}
\end{equation*}
$$

(ii) the $L_{1}$-weight of the coefficient vector of $p$ is at most $d^{0.01}$, and
(iii) the degree of $p$ is at most

$$
\min \left\{d^{1-\frac{1}{C \sqrt{k}}}, \frac{d}{\log ^{0.995} d}\right\} .
$$

The degree bound of $d / \log ^{0.995} d$ in the above theorem follows directly from techniques developed in [STT12], while the degree bound of $d^{1-\frac{1}{C \sqrt{k}}}$ requires additional insight. We also show a lower bound to exhibit the tightness of our construction.

Theorem 1.4. Let $k=o(\log d)$, and let $p$ be a real d-variate polynomial satisfying $\left|p(x)-O R_{d}(x)\right| \leq$ $1 / 6$ for all $x \in\{-1,1\}^{d}$ with $|x| \leq k$. If the $L_{1}$-weight of the coefficient vector of $p$ is $d^{O(1)}$, then the degree of $p$ is at least $d^{1-O(1 / \sqrt{k})}$.

We note that our algorithmic approach for designing efficient private data release algorithms would work equally well if we have any small set of functions whose low- $L_{1}$-weight linear combinations approximate disjunctions restricted to inputs of Hamming weight at most $k$. Our lower bound limits the applicability of our approach if we choose to use low-degree monomials as the set of functions. We observe that this also rules out several natural candidates that can themselves be computed exactly by a low-weight polynomial of low-degree (e.g., the set of small-width conjunctions). There is some additional evidence from prior work that low-degree monomials may be the optimal choice: if we only care about the size of the set of functions used to approximate disjunctions on inputs of Hamming weight at most $k$, then prior work shows that low-degree monomials are indeed optimal [She11 (see also Section 5 in the full version of [TUV12]). It remains an interesting open question to determine whether this optimality still holds when we restrict the $L_{1}$ weight of the linear combinations used in the approximations to be poly $(d)$.

### 1.2 Techniques

For notational convenience, we focus on monotone $k$-way disjunction queries. However, our results extend straightforwardly to general non-monotone $k$-way marginal queries via simple transformations on the database and queries. A monotone $k$-way disjunction is specified by a set $S \subseteq[d]$ of size $k$ and asks what fraction of records in $D$ have at least one of the attributes in $S$ set to 1 .

Following the approach introduced by Gupta et al. [GHRU11] and developed into a general theory in [HRS12], we view the problem of releasing answers to conjunction queries as a learning problem. That is, we view the database as specifying a function $f_{D}:\{-1,1\}^{d} \rightarrow[0,1]$, in which each input vector $s \in\{-1,1\}^{d}$ is interpreted as the indicator vector of a set $S \subseteq\{1, \ldots, d\}$, with $s_{i}=-1$ iff $i \in S$, and $f_{D}(s)$ equals the evaluation of the conjunction query specified by $S$ on the database $D$. Then, our goal is to privately learn to approximate the function $f_{D}$; this is accomplished in HRS12] by approximating $f_{D}$ succinctly with polynomials and learning the polynomial privately. Polynomial approximation is central to our approach as well, as we explain below.

We begin with a description of how the parameters of the online learning algorithm determine the parameters of the online differentially private learning algorithm. We consider the "IDC framework" GRU12] - which captures the private multiplicative weights algorithm HR10 among others RR10, GRU12, JT12 - for deriving differentially private online algorithms from any online learning algorithm that may not necessarily be privacy preserving.

Informally, an online learning algorithm is one that takes a (possibly adaptively chosen) sequence of inputs $s_{1}, s_{2}, \ldots$ and returns answers $a_{1}, a_{2}, \ldots$ to each, representing "guesses" about the values $f_{D}\left(s_{1}\right), f_{D}\left(s_{2}\right), \ldots$ for the unknown function $f_{D}$. After making each guess $a_{i}$, the learner is given some information about the value of $f_{D}\left(s_{i}\right)$. The quantities of interest are the running time required by the online learner to produce each guess $a_{i}$ and the number of "mistakes" made by the learner, which is the number of rounds $i$ in which $a_{i}$ is "far" from $f_{D}\left(s_{i}\right)$. Ultimately, for the differentially private algorithm derived in the IDC framework, the notion of far will correspond to the accuracy, the per query running time will essentially be equal to the running time of the online learning algorithm, and the minimum database size required by the private algorithm will be proportional to the square root of the number of mistakes.

We next describe the well-known technique of deriving faster online learning algorithms that commit fewer mistakes using polynomial approximations to the target function. Indeed, it is wellknown that if $f_{D}$ can be approximated to high accuracy by a $d$-variate polynomial $p_{D}:\{-1,1\}^{d} \rightarrow \mathbb{R}$ of degree $t$ and $L_{1}$-weight at most $W$, where the weight is defined to be the sum of the absolute values of the coefficients, then there is an online learning algorithm that runs in time poly $\left.\binom{d}{\leq t}\right)$ and makes $O\left(W^{2} d\right)$ mistakes. Thus, if $t \ll d$, the running time of such an online learning algorithm will be significantly less than $2^{d}$ and the number of mistakes (and thus the minimum database size of the resulting private algorithm) will only blow up by actor of $W$.

Consequently, our goal boils down to constructing the best possible polynomial representation $p_{D}$ for any database $D$ - one with low-degree, low- $L_{1}$-weight such that $\left|p_{D}(s)-f_{D}(s)\right|$ is small for all vectors $s \in\{-1,1\}^{d}$ corresponding to monotone $k$-way disjunction queries. To accomplish this goal, it is sufficient to construct a low-degree, low- $L_{1}$-weight polynomial that can approximate the $d$-variate OR function on inputs of Hamming weight at most $k$ (i.e., those that have -1 in at most $k$ indices). Such problems are well-studied in the approximation-theory literature, however our variant requires polynomials to be accurate only on a restricted subset of inputs. In fact, the existence of a polynomial with degree $d / \log ^{0.99} d$ and $L_{1}$-weight $d^{o(1)}$ that approximates the $d$-variate OR function on all inputs follows from the work of Servedio et al. [STT12]. We improve these bounds for small values of $k$ by constructing an approximating polynomial that has degree
$d^{1-\Omega(1 / \sqrt{k})}$ and $L_{1}$-weight $d^{0.01}$.
We also prove a new approximation-theoretic lower bound for polynomials that seek to approximate a target function for a restricted subset of inputs. Specifically, we show that for any $k=o(\log d)$, any polynomial $p$ of weight $\operatorname{poly}(d)$ that satisfies $|p(s)-\mathrm{OR}(s)| \leq 1 / 6$ for all inputs $s \in\{-1,1\}^{d}$ of Hamming weight at most $k$ must have degree $d^{1-O(1 / \sqrt{k})}$. We prove our lower bound by expressing the problem of constructing such a low-weight, low-degree polynomial $p$ as a linear program, and exhibiting an explicit solution to the dual of this linear program. Our proof is inspired by recent work of Sherstov [She09, She11, She12b] and Bun-Thaler [BT13].

### 1.3 Related Work

Other Results on Privately Releasing Marginals. In work subsequent to our result, Dwork et al. DNT13] show how to privately release marginals in a very different parameter regime. Their algorithm is faster than ours, running in time $\operatorname{poly}\left(\binom{d}{\leq k}\right)$, and has better dependence on the error parameter. However, their algorithm requires that the database size is $\widetilde{\Omega}\left(d^{\lceil k / 2\rceil / 2}\right)$ for answering with error $\pm 0.01$. This size is comparable to the optimal $\widetilde{\Omega}(k \sqrt{d})$ only when $k \leq 2$. In contrast, our algorithm has nearly-optimal minimum database size for every choice of $k$.

While we have focused on accurately answering every $k$-way marginal query, or more generally every query in a sequence of marginal queries, several other works have considered more relaxed notions of accuracy. These works show how to efficiently release a summary of the database from which an analyst can efficiently compute an approximate answer to marginal queries, with the guarantee that the average error of a marginal query is at most .01 , when the query is chosen from a particular distribution. In particular, Feldman and Kothari [FK13] achieve small average error over the uniform distribution with running time and database size $\tilde{O}\left(d^{2}\right)$; Gupta et al. GHRU11] achieve small average error over any product distribution with running time and minimum database size poly $(d)$; finally Hardt et al. HRS12 show how to achieve small average error over arbitrary distributions with running time and minimum database size $2^{\tilde{O}\left(d^{1 / 3}\right)}$. All of these results are based on the approach of learning the function $f_{D}$.

Several works have also considered information theoretic bounds on the minimum database size required to answer $k$-way marginals. Kasiviswanathan et al. KRSU10] showed that $|D| \geq$ $\min \left\{1 / \alpha^{2}, d^{k / 2} / \alpha\right\}$ is necessary to answer all $k$-way marginals with error $\pm \alpha$. De [De12] extended this result to hold even when accuracy $\pm \alpha$ can be violated for a constant fraction of $k$-way marginals. In our regime, where $\alpha=\Omega(1)$, their results do not give a non-trivial lower bound. In forthcoming work, Bun, Ullman, and Vadhan BUV13] have proven a lower bound of $|D| \geq \tilde{\Omega}(k \sqrt{d})$, which is nearly optimal for $\alpha=\Omega(1)$.

Relationship with Hardness Results for Differential Privacy. Ullman Ull13] (building on the results of Dwork et al. $\left.\mathrm{DNR}^{+} 09\right]$, showed that any $2^{o(d)}$-time differentially private algorithm that answers arbitrary counting queries can only give accurate answers if $|D| \gtrsim|\mathcal{Q}|^{1 / 2}$, assuming the existence of exponentially hard one-way functions. Our algorithms have running time $2^{o(d)}$ and are accurate when $|D| \ll|\mathcal{Q}|^{1 / 2}$, and thus show a separation between answering marginal queries and answering arbitrary counting queries.

When viewed as an offline algorithm for answering all $k$-way marginals, our algorithm will return a list of values containing answers to each $k$-way marginal query. It would in some cases be more attractive if we could return a synthetic database, which is a new database $\widehat{D} \in\left(\{0,1\}^{d}\right)^{\widehat{n}}$ whose rows are "fake", but such that $\widehat{D}$ approximately preserves many of the statistical properties of the database $D$ (e.g., all the marginals). Some of the previous work on counting query release
has provided synthetic data $\mathrm{BCD}^{+} 07, \mathrm{BLR} 08, \mathrm{DNR}^{+} 09, \mathrm{DRV10}, \mathrm{HLM} 12$.
Unfortunately, Ullman and Vadhan [UV11] (building on [DNR ${ }^{+} 09$ ]) have shown that no differentially private sanitizer with running time poly $(d)$ can take a database $D \in\left(\{0,1\}^{d}\right)^{n}$ and output a private synthetic database $\widehat{D}$, all of whose 2 -way marginals are approximately equal to those of $D$, assuming the existence of one-way functions. They also showed that under certain strong cryptographic assumptions, there is no differentially private sanitizer with running time $2^{d^{1-\Omega(1)}}$ can output a private synthetic database, all of whose 2 -way marginals are approximately equal to those of $D$. Our algorithms indeed achieve this running time and accuracy guarantee when releasing $k$-way marginals for constant $k$, and thus it may be inherent that our algorithms do not generate synthetic data.

Relationship with Results in Approximation Theory. Servedio et al. [STT12] focused on developing low-weight, low-degree polynomial threshold functions (PTFs) for decision lists, motivated by applications in computational learning theory. As an intermediate step in their PTF constructions, they constructed low- $L_{1}$-weight, low-degree polynomials that approximate the OR function on all Boolean inputs. Our construction of lower-weight, lower-degree polynomials that approximate the OR function on low Hamming weight inputs is inspired by and builds on Servedio et al.'s construction of approximations that are accurate on all Boolean inputs.

The proof of our lower bound is inspired by recent work that has established new approximate degree lower bounds via the construction of dual solutions to certain linear programs. In particular, Sherstov [She09] showed that approximate degree and PTF degree behave roughly multiplicatively under function composition, while Bun and Thaler [BT13] gave a refinement of Sherstov's method in order to resolve the approximate degree of the two-level AND-OR tree, and also gave an explicit dual witness for the approximate degree of any symmetric Boolean function. We extend these lower bounds along two directions: (1) we show degree lower bounds that take into account the $L_{1}$-weight of the coefficient vector of the approximating polynomial, and (2) our lower bounds hold even when we only require the approximation to be accurate on inputs of low Hamming weight, while prior work only considered approximations that are accurate on all Boolean inputs.

Some prior work has studied the degree of polynomials that point-wise approximate partial Boolean functions She12b, She12a. Here, a function $f: Y \rightarrow \mathbb{R}$ is said to be partial if its domain $Y$ is a strict subset of $\{-1,1\}^{d}$, and a polynomial $p$ is said to $\epsilon$-approximate $f$ if

1. $|f(x)-p(x)| \leq \epsilon$ for all $x \in Y$, and
2. $|p(x)| \leq 1+\epsilon$ for all $x \in\{-1,1\}^{d} \backslash Y$.

In contrast, our lower bounds apply even in the absence of Condition 2, i.e., when $p(x)$ is allowed to take arbitrary values on inputs in $\{-1,1\}^{d} \backslash Y$.

Finally, while our motivation is private data release, our approximation theoretic results are similar in spirit to recent work of Long and Servedio [LS13, who are motivated by applications in computational learning theory. Long and Servedio consider halfspaces $h$ defined on inputs of small Hamming weight, and (using different techniques very different from ours) give upper and lower bounds on the weight of these halfspaces when represented as linear threshold functions.

Organization. In Section 3, we describe our private online algorithm and show that it yields the claimed accuracy given the existence of sufficiently low-weight polynomials that approximate the $d$-variate OR function on inputs of low Hamming weight. The results of this section are a combination of known techniques in differential privacy [RR10, HR10, GRU12] and learning theory
(see e.g., [KS04]). Readers familiar with these literatures may prefer to skip Section 3 on first reading. In Section 4, we give our polynomial approximations to the OR function, both on low Hamming weight Boolean inputs and on all Boolean inputs. Finally, in Section 5, we state and prove our lower bounds for polynomial approximations to the OR function on restricted inputs.

## 2 Preliminaries

### 2.1 Differentially Private Sanitizers

Let a database $D \in\left(\{0,1\}^{d}\right)^{n}$ be a collection of $n$ rows $x^{(1)}, \ldots, x^{(n)}$ from a data universe $\{0,1\}^{d}$. We say that two databases $D, D^{\prime} \in\left(\{0,1\}^{d}\right)^{n}$ are adjacent if they differ only on a single row, and we denote this by $D \sim D^{\prime}$.

Let $\mathcal{A}:\left(\{0,1\}^{d}\right)^{n} \rightarrow \mathcal{R}$ be an algorithm that takes a database as input and outputs some data structure in $\mathcal{R}$. We are interested in algorithms that satisfy differential privacy.

Definition 2.1 (Differential Privacy DMNS06). An algorithm $\mathcal{A}:\left(\{0,1\}^{d}\right)^{n} \rightarrow \mathcal{R}$ is $(\varepsilon, \delta)-$ differentially private if for every two adjacent databases $D \sim D^{\prime} \in\left(\{0,1\}^{d}\right)^{n}$ and every subset $S \subseteq \mathcal{R}$,

$$
\operatorname{Pr}[\mathcal{A}(D) \in S] \leq e^{\varepsilon} \operatorname{Pr}\left[\mathcal{A}\left(D^{\prime}\right) \in S\right]+\delta .
$$

Since a sanitizer that always outputs $\perp$ satisfies Definition 2.1, we focus on sanitizers that are accurate. In particular, we are interested in sanitizers that give accurate answers to counting queries. A counting query is defined by a boolean predicate $q:\{0,1\}^{d} \rightarrow\{0,1\}$. Abusing notation, we define the evaluation of the query $q$ on a database $D \in\left(\{0,1\}^{d}\right)^{n}$ to be $q(D)=\frac{1}{n} \sum_{i=1}^{n} q\left(x^{(i)}\right)$. Note that the value of a counting query is in $[0,1]$. We use $\mathcal{Q}$ to denote a set of counting queries.

For the purposes of this work, we assume that the range of $\mathcal{A}$ is simply $\mathbb{R}^{|\mathcal{Q}|}$. That is, $\mathcal{A}$ outputs a list of real numbers representing answers to each of the specified queries.

Definition 2.2 (Accuracy). The output of $\mathcal{A}(D), a=\left(a_{q}\right)_{q \in \mathcal{Q}}$, is $\alpha$-accurate for the query set $\mathcal{Q}$ if

$$
\forall q \in \mathcal{Q},\left|a_{q}-q(D)\right| \leq \alpha
$$

A sanitizer is $(\alpha, \beta)$-accurate for the query set $\mathcal{Q}$ if for every database $D, \mathcal{A}(D)$ outputs $a$ such that with probability at least $1-\beta$, the output $a$ is $\alpha$-accurate for $\mathcal{Q}$, where the probability is taken over the coins of $\mathcal{A}$.

We remark that the definition of both differential privacy and ( $\alpha, \beta$ )-accuracy extend straightforwardly to the online setting. Here the algorithm receives a sequence of $\ell$ (possibly adaptively chosen) queries from $\mathcal{Q}$ and must give an answer to each before seeing the rest of the sequence. Here we require that with probability at least $1-\beta$, every answer output by the algorithm is within $\pm \alpha$ of the true answer on $D$. See e.g., HR10 for an elaborate treatment of the online setting.

### 2.2 Query Function Families

Given a set of queries of interest, $\mathcal{Q}$ (e.g., all marginal queries), we think of the database $D$ as specifying a function $f_{D}$ mapping queries $q$ to their answers $q(D)$. We now describe this transformation more formally:

Definition 2.3 ( $\mathcal{Q}$-Function Family). Let $\mathcal{Q}=\left\{q_{y}\right\}_{y \in Y_{\mathcal{Q}} \subseteq\{-1,1\}^{m}}$ be a set of counting queries on a data universe $\{0,1\}^{d}$, where each query is indexed by an $m$-bit string. We define the index set of $\mathcal{Q}$ to be the set $Y_{\mathcal{Q}}=\left\{y \in\{-1,1\}^{m} \mid q_{y} \in \mathcal{Q}\right\}$.

We define the $\mathcal{Q}$-function family $\mathcal{F}_{\mathcal{Q}}=\left\{f_{x}:\{-1,1\}^{m} \rightarrow[0,1]\right\}_{x \in\{0,1\}^{d}}$ as follows: For every possible database row $x \in\{0,1\}^{d}$, the function $f_{x}:\{-1,1\}^{m} \rightarrow[0,1]$ is defined as $f_{x}(y)=q_{y}(x)$. Given a database $D \in\left(\{0,1\}^{d}\right)^{n}$ we define the function $f_{\mathcal{Q}, D}:\{-1,1\}^{m} \rightarrow[0,1]$ where $f_{\mathcal{Q}, D}(y)=$ $\frac{1}{n} \sum_{i=1}^{n} f_{x^{(i)}}(y)$. When $\mathcal{Q}$ is clear from context we will drop the subscript $\mathcal{Q}$ and simply write $f_{x}$, $f_{D}$, and $\mathcal{F}$.

When $\mathcal{Q}$ is the set of all monotone $k$-way disjunctions on a database $D \in\left(\{0,1\}^{d}\right)^{n}$, the queries are defined by sets $S \subseteq[d],|S| \leq k$. In this case, we represent each query by the $d$-bit $-1 / 1$ indicator vector $y_{S}$ of the set $S$, where $y_{S}(i)=-1$ if and only if $i \in S$. Thus, $y_{S}$ has at most $k$ entries that are -1 . Hence, we can take $m=d$ and $Y_{\mathcal{Q}}=\left\{y \in\{-1,1\}^{d} \mid \sum_{j=1}^{d} \mathbf{1}_{\left\{y_{i}=-1\right\}} \leq k\right\}$.

### 2.3 Low-Weight Polynomial Approximations

Given an $m$-variate real polynomial $p:\{-1,1\}^{m} \rightarrow \mathbb{R}$,

$$
p(y)=\sum_{S \subseteq[m]} c_{S} \cdot \prod_{i \in S} y_{i}
$$

we define the degree, weight $w(\cdot)$ and non-constant weight $w^{*}(\cdot)$ of the polynomial as follows:

$$
\begin{aligned}
\operatorname{deg}(p) & :=\max \left\{|S|: S \subseteq[m], c_{S} \neq 0\right\}, \\
w(p) & :=\sum_{S \subseteq[m]}\left|c_{S}\right|, \text { and } \\
w^{*}(p) & :=\sum_{S \subseteq[m], S \neq \emptyset}\left|c_{S}\right| .
\end{aligned}
$$

We use $\binom{[m]}{\leq t}$ to denote $\left\{S \subseteq[m]||S| \leq t\}\right.$ and $\binom{m}{\leq t}=\left|\binom{[m]}{\leq t}\right|=\sum_{j=0}^{t}\binom{m}{j}$.
We will attempt to approximate the functions $f_{x}:\{-1,1\}^{m} \rightarrow\{0,1\}$ on all the indices in $Y_{\mathcal{Q}}$ by a family of polynomials with low degree and low weight. Formally and more generally:

Definition 2.4 (Restricted Approximation by Polynomials). Given a function $f: Y \rightarrow \mathbb{R}$, where $Y \subseteq \mathbb{R}^{m}$, and a subset $Y^{\prime} \subseteq Y$, we denote the restriction of $f$ to $Y^{\prime}$ by $\left.f\right|_{Y^{\prime}}$. Given an $m$-variate real polynomial $p$, we say that $p$ is a $\gamma$-approximation to the restriction $\left.f\right|_{Y^{\prime}}$, if $|f(y)-p(y)| \leq \gamma$ $\forall y \in Y^{\prime}$. Notice there is no restriction whatsoever placed on $p(y)$ for $y \in Y \backslash Y^{\prime}$.

Given a family of $m$-variate functions $\mathcal{F}=\left\{f_{x}: Y \rightarrow \mathbb{R}\right\}_{x \in\{0,1\}^{d}}$, where $Y \subseteq \mathbb{R}^{m}$, a set $Y^{\prime} \subseteq Y$ we use $\left.\mathcal{F}\right|_{Y^{\prime}}=\left\{\left.f_{x}\right|_{Y^{\prime}}\right\}_{x \in\{0,1\}^{d}}$ to denote the family of restricted functions. Given a family $\mathcal{P}$ of $m$-variate real polynomials, we say that the family $\mathcal{P}$ is a $\gamma$-approximation to the family $\left.\mathcal{F}\right|_{Y^{\prime}}$ if for every $x \in\{0,1\}^{d}$, there exists $p_{x} \in \mathcal{P}$ that is a $\gamma$-approximation to $\left.f_{x}\right|_{Y^{\prime}}$.

Let $H_{m, k}=\left\{x \in\{-1,1\}^{m}: \sum_{i=1}^{m}\left(1-x_{i}\right) / 2 \leq k\right\}$ denote the set of inputs of Hamming weight at most $k$. We view the $d$ variate OR function, $\operatorname{OR}_{d}$ as mapping inputs from $\{-1,1\}^{d}$ to $\{-1,1\}$, with the convention that -1 is TRUE and 1 is FALSE. Let $\mathcal{P}_{t, W}(m)$ denote the family of all $m$-variate real polynomials of degree $t$ and weight $W$. For the upper bound, we will show that for certain small values of $t$ and $W$, the family $\mathcal{P}_{t, W}(d)$ is a $\gamma$-approximation to the family of all disjunctions restricted to $H_{d, k}$.
Fact 2.5. If $\mathcal{Q}$ is the set of all monotone $k$-way disjunctions on a database $D \in\left(\{0,1\}^{d}\right)^{n}, \mathcal{F}$ is its function family, and $Y=H_{d, k}$ is its index set, then $\mathcal{P}_{t, W}(d)$ is a $\gamma$-approximation to the restriction $\left.\mathcal{F}\right|_{Y}$ if and only if there is a degree $t$ polynomial of weight $O(W)$ that $\gamma$-approximates $\left.O R_{d}\right|_{H_{d, k}}$.

The fact follows easily by observing that for any $x \in\{0,1\}^{d}, y \in\{-1,1\}^{d}$,

$$
f_{x}(y)=\bigvee_{i \in x} \mathbf{1}_{\left\{y_{i}=-1\right\}}=\frac{1-\mathrm{OR}_{d}\left(y_{1}^{x_{1}}, \ldots, y_{d}^{x_{d}}\right)}{2} .
$$

For the lower bound, we will show that any collection of polynomials with small weight that is a $\gamma$-approximation to the family of disjunctions restricted to $H_{m, k}$ should have large degree. We need the following definitions:

Definition 2.6 (Approximate Degree). Given a function $f: Y \rightarrow \mathbb{R}$, where $Y \subseteq \mathbb{R}^{m}$, the $\gamma$ approximate degree of $f$ is

$$
\operatorname{deg}_{\gamma}(f):=\min \{d: \exists \text { real polynomial } p \text { that is a } \gamma \text {-approximation to } f, \operatorname{deg}(p)=d\} .
$$

Analogously, the $(\gamma, W)$-approximate degree of $f$ is
$\operatorname{deg}_{(\gamma, W)}(f):=\min \{d: \exists$ real polynomial $p$ that is a $\gamma$-approximation to $f, \operatorname{deg}(p)=d, w(p) \leq W\}$.
It is clear that $\operatorname{deg}_{\gamma}(f)=\operatorname{deg}_{(\gamma, \infty)}(f)$.
We let $w^{*}(f, t)$ denote the degree-t non-constant margin weight of $f$, defined to be:

$$
w^{*}(f, t):=\min \left\{w^{*}(p): \exists \text { real polynomial } p \text { s.t. } \operatorname{deg}(p) \leq t, f(y) p(y) \geq 1 \forall y \in Y\right\}
$$

The above definitions extend naturally to the restricted function $\left.f\right|_{Y^{\prime}}$.
Our definition of non-constant margin weight is closely related to the well-studied notion of the degree- $t$ polynomial threshold function (PTF) weight of $f$ (see e.g., She11), which is defined as $\min _{p} w(p)$, where the minimum is taken over all degree- $t$ polynomials $p$ with integer coefficients, such that $f(x)=\operatorname{sign}(p(x))$ for all $x \in\{-1,1\}^{d}$. Often, when studying PTF weight, the requirement that $p$ have integer coefficients is used only to ensure that $p$ has non-trivial margin, i.e. that $|p(x)| \geq 1$ for all $x \in\{-1,1\}^{d}$; this is precisely the requirement captured in our definition of nonconstant margin weight. We choose to work with margin weight because it is a cleaner quantity to analyze using linear programming duality; PTF weight can also be studied using LP duality, but the integrality constraints on the coefficients of $p$ introduces an integrality gap that causes some loss in the analysis (see e.g., Sherstov [She11, Theorem 3.4] and Klauck [Kla11, Section 4.3]).

## 3 From Low-Weight Approximations to Private Data Release

In this section we show that low-weight polynomial approximations imply data release algorithms that provide approximate answers even on small databases. The main goal of this section is to prove the following theorem.

Theorem 3.1. Given $\alpha, \beta, \varepsilon, \delta>0$, and a family of linear queries $\mathcal{Q}$ with index set $Y \subseteq\{-1,1\}^{m}$. Suppose for some $t \leq m, W>0$, the family of polynomials $\mathcal{P}_{t, W}(m)(\alpha / 4)$-approximates the function family $\left.\mathcal{F}_{\mathcal{Q}}\right|_{Y}$. Then there exists an $(\epsilon, \delta)$-differentially private online algorithm that is $(4 \alpha, \beta)$-accurate for any sequence of $\ell$ (possibly adaptively chosen) queries from $\mathcal{Q}$ on a database $D \in\left(\{0,1\}^{d}\right)^{n}$, provided

$$
n \geq \frac{128 W \log (\ell / \beta) \log (4 / \delta)}{\alpha^{2} \varepsilon} \sqrt{\log \left(2\binom{m}{\leq t}+1\right)}
$$

The private algorithm has running time poly $\left(n,\binom{m}{\leq t}, \log W, \log (1 / \alpha), \log (1 / \beta), \log (1 / \varepsilon), \log (1 / \delta)\right)$.

We note that the theorem can be assembled from known techniques in the design and analysis of differentially private algorithms and online learning algorithms. We include the proof of the theorem here for the sake of completeness, as to our knowledge they do not explicitly appear in the privacy literature.

We construct and analyze the algorithm in two steps. First, we use standard arguments to show that the non-private multiplicative weights algorithm can be used to construct a suitable online learning algorithm for $f_{\mathcal{Q}, D}$ whenever $f_{\mathcal{Q}, D}$ can be approximated by a low-weight, low-degree polynomial. Here, a suitable online learning algorithm is one that fits into the IDC framework of Gupta et al. GRU12. We then apply the generic conversion from IDCs to differentially private online algorithms RR10, HR10, GRU12 to obtain our algorithm.

### 3.1 IDCs

We start by providing the relevant background on the iterative database construction framework. An IDC will maintain a sequence of functions $f_{D}^{(1)}, f_{D}^{(2)}, \ldots$ that give increasingly good approximations to the $f_{\mathcal{Q}, D}$. In our case, these functions will be low-degree polynomials. Moreover, the mechanism produces the next approximation in the sequence by considering only one query $y^{(t)}$ that "distinguishes" the real database in the sense that $\left|f^{(t)}\left(y^{(t)}\right)-f_{D}\left(y^{(t)}\right)\right|$ is large.

Definition 3.2 (IDC RR10, HR10, GRU12]). Let $\mathcal{Q}=\left\{q_{y}\right\}_{y \in Y_{\mathcal{Q}} \subseteq\{-1,1\}^{m}}$ be a family of counting queries indexed by $m$-bit strings. Let $\mathbf{U}$ be an algorithm mapping a function $f: Y_{\mathcal{Q}} \rightarrow \mathbb{R}$, a query $y \in Y_{\mathcal{Q}}$, and a real number $a$ to a new function $f^{\prime}$. Let $D \in\left(\{0,1\}^{d}\right)^{n}$ be a database and $\alpha>0$ be a parameter. Consider the following game with an adversary. Let $f_{D}^{(1)}$ be some function. In each round $t=1,2, \ldots$ :

1. The adversary chooses a query $y^{(t)} \in Y_{\mathcal{Q}}$ (possibly depending on $f_{D}^{(t)}$ ).
2. If $\left|f_{D}^{(t)}\left(y^{(t)}\right)-f_{D}\left(y^{(t)}\right)\right|>\alpha$, then we say that the algorithm has made a mistake.
3. If the algorithm made a mistake, then it receives a value $a^{(t)} \in \mathbb{R}$ such that $\mid f_{D}^{(t)}\left(y^{(t)}\right)-$ $f_{D}\left(y^{(t)}\right) \mid \leq \alpha / 2$ and computes a new function $f_{D}^{(t+1)}=\mathbf{U}\left(f_{D}^{(t)}, y^{(t)}, a^{(t)}\right)$. Otherwise let $f_{D}^{(t+1)}=f_{D}^{(t)}$.

If the number of rounds $t$ in which the algorithm makes a mistake is at most $B$ for every adversary, then $\mathbf{U}$ is a iterative database construction for $\mathcal{Q}$ with mistake bound $B$.

Theorem 3.3 (Variant of [GRU12]). For any $\alpha>0$, and any family of queries $\mathcal{Q}$, if there is an iterative database construction for $\mathcal{Q}$ with mistake bound $B$, then there is an $(\varepsilon, \delta)$-differentially private online algorithm that is $(4 \alpha, \beta)$-accurate for any sequence of $\ell$ (possibly adaptively chosen) queries from $\mathcal{Q}$ on a database $D \in\left(\{0,1\}^{d}\right)^{n}$, so long as

$$
n \geq \frac{32 \sqrt{B} \log (\ell / \beta) \log (4 / \delta)}{\alpha \varepsilon}
$$

Moreover, if the iterative database construction, $\mathbf{U}$, runs in time $T_{\mathbf{U}}$, then the private algorithm has running time $\operatorname{poly}\left(n, T_{\mathbf{U}}, \log (1 / \alpha), \log (1 / \beta), \log (1 / \varepsilon), \log (1 / \delta)\right)$ per query.

The IDC we will use is specified in Algorithm (1. The IDC will use approximations $f_{D}^{(t)}$ in the form of low-degree polynomials of low-weight, and thus we need to specify how to represent such a function. Specifically, we will represent a polynomial $p$ as a vector $\bar{p}$ of length $2(\underset{\leq t}{m})+1$ with only
non-negative entries. For each coefficient $S \in\binom{[m]}{\leq t}$, the vector will have two components $\bar{p}_{S}, \bar{p}_{\neg S}$. (Recall that $\binom{[m]}{\leq t}:=\{S \subseteq[m]| | S \mid \leq t\}$.) Intuitively these two entries represent the positive part and negative part of the coefficient $c_{S}$ of $p$. There will also be an additional entry $\bar{p}_{0}$ that is used to ensure that the $L_{1}$-norm of the vector is exactly 1 . Given a polynomial $p \in \mathcal{P}_{t, W}$ with coefficients $\left(c_{S}\right)$, we can construct this vector by setting

$$
\bar{p}_{S}=\frac{\max \left\{0, c_{S}\right\}}{W} \quad \bar{p}_{\neg S}=\frac{\max \left\{0,-c_{S}\right\}}{W}
$$

and choosing $\bar{p}_{0}$ so that $\|\bar{p}\|_{1}=1$. Observe that $\bar{p}_{0}$ can always be set appropriately since the weight of $p$ is at most $W$.

Similarly, we want to associate queries with vectors so that we can replace the evaluation of the polynomial $p$ on a query $y$ with the inner product $W\langle\bar{p}, \bar{y}\rangle$. We can do so by defining the vector $\bar{y}$ of length $2\binom{m}{\leq t}+1$ such tha $\bar{y}_{0}=0, \bar{y}_{S}=\prod_{i \in S} y_{i}$ and $\bar{y}_{\neg S}=-\prod_{i \in S} y_{i}$.

Fact 3.4. For every $m$-variate polynomial $p$ of degree at most $t$ and weight at most $W$, and every query $y \in\{-1,1\}^{m}, W\langle\bar{p}, \bar{y}\rangle=p(y)$.

```
Algorithm 1 The Multiplicative Weights Algorithm for Low-Weight Polynomials.
\(\mathbf{U}_{\alpha}^{M W}\left(\bar{p}^{(t)}, y^{(t)}, a^{(t)}\right):\)
    Let \(\eta \leftarrow \alpha / 4 W\).
    If: \(p^{(t)}=\emptyset\) then: output
\[
\bar{p}^{(t)}=\frac{1}{2\binom{m}{\leq t}+1}(1, \ldots, 1)
\]
```

(representing the constant 0 polynomial).
Else if: $a^{(t)}<W\left\langle\bar{p}^{(t)}, \bar{y}^{(t)}\right\rangle$
Let $\bar{r}^{(t)}=\bar{y}^{(t)}$
Else:
Let $\bar{r}^{(t)}=-\bar{y}^{(t)}$


$$
\begin{gathered}
\bar{p}_{I}^{(t+1)} \leftarrow \exp \left(-\eta \bar{r}_{I}^{(t)}\right) \cdot \bar{p}_{I}^{(t)} \\
\bar{p}^{(t+1)} \leftarrow \frac{\bar{p}^{(t+1)}}{\left\|\bar{p}^{(t+1)}\right\|_{1}}
\end{gathered}
$$

Output $\bar{p}^{(t+1)}$.
We summarize the properties of the multiplicative weights algorithm in the following theorem:
Theorem 3.5. For any $\alpha>0$, and any family of linear queries $\mathcal{Q}$ if $\mathcal{P}_{t, W}(\alpha / 4)$-approximates the restriction $\left.\mathcal{F}\right|_{Y}$ then Algorithm 1 is an iterative database construction for $\mathcal{Q}$ with mistake bound $B$ for

$$
B=B(W, m, t, \alpha)=\frac{16 W^{2} \log \left(2\binom{m}{\leq t}+1\right)}{\alpha^{2}} .
$$

Moreover, $\mathbf{U}$ runs in time $\operatorname{poly}\left(\binom{m}{\leq t}, \log W, \log (1 / \alpha)\right)$.

Proof. Let $D \in\left(\{0,1\}^{d}\right)^{n}$ be any database. For every round $t$ in which $\mathbf{U}$ makes a mistake, we consider the tuple $\left(\bar{p}^{(t)}, y^{(t)}, \widehat{a}^{(t)}\right)$ representing the information used to update the approximation in round $t$. In order to bound the number of mistakes, it will be sufficient to show that after $B \leq 16 W^{2} \log \left(2\binom{m}{\leq t}+1\right) / \alpha^{2}$, the vector $\bar{p}^{(t)}$ is such that

$$
\forall y \in Y_{\mathcal{Q}},\left|W\left\langle\bar{p}^{(t)}, \bar{y}\right\rangle-f_{D}(y)\right| \leq \alpha
$$

That is, after making $B$ mistakes $\bar{p}^{(t)}$ represents a polynomial that approximates $f_{D}$ on every query, and thus there can be no more than $B$ makes.

First, we note that there always exists a polynomial $p_{D} \in \mathcal{P}_{t, W}$ such that

$$
\begin{equation*}
\forall y \in Y_{\mathcal{Q}},\left|p_{D}(y)-f_{D}(y)\right| \leq \frac{\alpha}{4} \tag{1}
\end{equation*}
$$

The assumption of our theorem is that for every $x^{(i)} \in D$, there exists $p_{x^{(i)}} \in \mathcal{P}_{t, W}$ such that

$$
\forall y \in Y_{\mathcal{Q}},\left|p_{x^{(i)}}(y)-f_{x^{(i)}}(y)\right| \leq \frac{\alpha}{4} .
$$

Thus, since $f_{D}=\frac{1}{n} \sum_{i=1}^{n} f_{x^{(i)}}$, the polynomial $p_{D}=\frac{1}{n} \sum_{i=1}^{n} p_{x^{(i)}}$ will satisfy (11). Note that $p_{D} \in \mathcal{P}_{t, W}$, thus if we represent $p_{D}$ as a vector,

$$
\forall y \in Y_{\mathcal{Q}},\left|W\left\langle\bar{p}_{D}, \bar{y}\right\rangle-f_{D}(y)\right| \leq \frac{\alpha}{4}
$$

Given the existence of $\bar{p}_{D}$, we will define a potential function capturing how far $\bar{p}^{(t)}$ is from $\bar{p}_{D}$. Specifically, we define

$$
\Psi_{t}:=K L\left(\bar{p}_{D} \| \bar{p}^{(t)}\right)=\sum_{\left.I \in\left\{0,\binom{[m]}{\leq t},\right\urcorner\binom{[m]}{\leq t}\right\}} \bar{p}_{D, I} \log \left(\frac{\bar{p}_{D, I}}{\bar{p}_{I}^{(t)}}\right)
$$

to be the KL divergence between $\bar{p}_{D}$ and the current approximation $\bar{p}^{(t)}$. Note that the sum iterates over all $2\binom{m}{\leq t}+1$ indices in $\bar{p}$. We have the following fact about KL divergence.
Fact 3.6. For all $t: \Psi_{t} \geq 0$, and $\Psi_{0} \leq \log (2(\underset{\leq t}{m})+1)$.
We will argue that after each mistake the potential drops by at least $\alpha^{2} / 16 W^{2}$. Note that the potential only changes in rounds where a mistake was made. Because the potential begins at $\log (2(\underset{\leq t}{m})+1)$, and must always be non-negative, we know that there can be at most $B(\alpha) \leq 16 W^{2} \log \left(2\binom{m}{\leq t}+1\right) / \alpha^{2}$ mistakes before the algorithm outputs a (vector representation of) a polynomial that approximates $f_{D}$ on $Y_{\mathcal{Q}}$.

The following lemma is standard in the analysis of multiplicative-weights based algorithms.

## Lemma 3.7.

$$
\Psi_{t}-\Psi_{t+1} \geq \eta\left(\left\langle\bar{p}^{(t)}, r^{(t)}\right\rangle-\left\langle\bar{p}_{D}, r^{(t)}\right\rangle\right)-\eta^{2}
$$

Proof.

$$
\begin{aligned}
\Psi_{t}-\Psi_{t+1} & =\sum_{I \in\left\{0,\binom{[m]}{\leq t}, \neg\left(\left[\begin{array}{c}
{[m]} \\
\leq t)
\end{array}\right)\right\}\right.} \bar{p}_{D, I} \log \left(\frac{\bar{p}_{I}^{(t+1)}}{\bar{p}_{I}^{(t)}}\right) \\
& =-\eta\left\langle\bar{p}_{D}, r^{(t)}\right\rangle-\log \left(\sum_{I \in\left\{0,\binom{[m]}{\leq t}, \neg\left(\left[\begin{array}{l}
{[m]} \\
\leq t
\end{array}\right)\right\}\right.} \exp \left(-\eta r_{I}^{(t)}\right) \bar{p}_{I}^{(t)}\right) \\
& \geq-\eta\left\langle\bar{p}_{D}, r^{(t)}\right\rangle-\log \left(\sum_{\substack{[m] \\
I \in\left\{0,\left(\begin{array}{c}
{[m] \\
\leq t}
\end{array}\right), \neg\left(\left[\begin{array}{l}
{[m] \\
\leq t}
\end{array}\right)\right\}\right.}} \bar{p}_{I}^{(t)}\left(1+\eta^{2}-\eta r_{I}^{(t)}\right)\right) \\
& \geq \eta\left(\left\langle\bar{p}^{(t)}, r^{(t)}\right\rangle-\left\langle\bar{p}_{D}, r^{(t)}\right\rangle\right)-\eta^{2}
\end{aligned}
$$

The rest of the proof now follows easily. By the conditions of an iterative database construction algorithm, $\left|\widehat{a}^{(t)}-f_{D}\left(y^{(t)}\right)\right| \leq \alpha / 2$. Hence, for each $t$ such that $\left|W\left\langle\bar{p}^{(t)}, y^{(t)}\right\rangle-f_{D}\left(y^{(t)}\right)\right| \geq \alpha$, we also have that $W\left\langle\bar{p}^{(t)}, y^{(t)}\right\rangle>f_{D}\left(y^{(t)}\right)$ if and only if $W\left\langle\bar{p}^{(t)}, y^{(t)}\right\rangle>\widehat{a}^{(t)}$.

In particular, if $r^{(t)}=y^{(t)}$, then $W\left\langle\bar{p}^{(t)}, y^{(t)}\right\rangle-W\left\langle\bar{p}_{D}, y^{(t)}\right\rangle \geq \alpha / 2$. Similarly, if $r^{(t)}=-y^{(t)}$, then $W\left\langle\bar{p}_{D}, y^{(t)}\right\rangle-W\left\langle\bar{p}^{(t)}, y^{(t)}\right\rangle \geq \alpha$. Here we have utilized the fact that $\left|p_{D}(y)-f_{D}(y)\right| \leq \alpha / 4$. Therefore, by Lemma 3.7 and the fact that $\eta=\alpha / 4 W$ :

$$
\Psi_{t}-\Psi_{t+1} \geq \frac{\alpha}{4 W}\left(\left\langle\bar{p}^{(t)}, r^{(t)}\right\rangle-\left\langle\bar{p}_{D}, r^{(t)}\right\rangle\right)-\frac{\alpha^{2}}{16 W^{2}} \geq \frac{\alpha}{4 W}\left(\frac{\alpha}{2 W}\right)-\frac{\alpha^{2}}{16 W^{2}}=\frac{\alpha^{2}}{16 W^{2}}
$$

Theorem 3.1 follows immediately from Theorems 3.3 and 3.5

## 4 Upper Bounds

Fact 2.5 and Theorem 3.1 show that in order to develop a differentially private mechanism that can release all $k$-way marginals of a database, it is sufficient to construct a low-weight polynomial that approximates the $\mathrm{OR}_{d}$, on all Boolean inputs of Hamming weight at most $k$. This is the purpose to which we now turn.

The $\mathrm{OR}_{d}$ function is easily seen to have an exact polynomial representation of constant weight and degree $d$ (see Fact 4.3 below); however, an approximation with smaller degree may be achieved at the expense of larger weight. The best known weight-degree tradeoff, implicit in the work of Servedio et al. STT12, can be stated as follows: there exists a polynomial $p$ of degree $t$ and weight $(d \log (1 / \gamma) / t)^{\left(d(\log 1 / \gamma)^{2} / t\right)}$ that $\gamma$-approximates the function $\mathrm{OR}_{d}$ on all Boolean inputs, for every $t$ larger than $\sqrt{d} \log (1 / \gamma)$. Setting the degree $t$ to be $O\left(d / \log ^{0.99} d\right)$ yields a polynomial of weight at most $d^{0.01}$ that approximates the $\mathrm{OR}_{d}$ function on all Boolean inputs to any desired constant accuracy. On the other hand, Lemma 8 of STT12] can be shown to imply that any polynomial of weight $W$ that $1 / 3$-approximates the $\mathrm{OR}_{d}$ function requires degree $\Omega(d / \log W)$, essentially matching the $O\left(d / \log { }^{-99} d\right)$ upper bound of Servedio et al. when $W=d^{\Omega(1)}$.

However, in order to privately release $k$-way marginals, we have shown that it suffices to construct polynomials that are accurate only on inputs of low Hamming weight. In this section, we give
a construction that achieves significantly improved weight degree trade-offs in this setting. In the next section, we demonstrate the tightness of our construction by proving matching lower bounds.

We construct our approximations by decomposing the $d$-variate OR function into an OR of OR's, which is the same approach taken by Servedio et al. [STT12]. Here, the outer OR has fan-in $m$ and the inner OR has fan-in $d / m$, where the subsequent analysis will determine the appropriate choice of $m$. In order to obtain an approximation that is accurate on all Boolean inputs, Servedio et al. approximate the outer OR using a transformation of the Chebyshev polynomials of degree $\sqrt{m}$, and compute each of the inner OR's exactly.

For $k \ll \log ^{2} d$, we are able to substantially reduce the degree of the approximating polynomial, relative to the construction of Servedio et al., by leveraging the fact that we are interested in approximations that are accurate only on inputs of Hamming weight at most $k$. Specifically, we are able to approximate the outer OR function using a polynomial of degree only $\sqrt{k}$ rather than $\sqrt{m}$, and argue that the weight of the resulting polynomial is still bounded by a polynomial in $d$.

We now proceed to prove the main lemmas. For the sake of intuition, we begin with weightdegree tradeoffs in the simpler setting in which we are concerned with approximating the $\mathrm{OR}_{d}$ function over all Boolean inputs. The following lemma, proved below for completeness, is implicit in the work of [STT12].

Lemma 4.1. For every $\gamma>0$ and $m \in[d]$, there is a polynomial of degree $t=O((d / \sqrt{m}) \log (1 / \gamma))$ and weight $W=m^{O(\sqrt{m} \log (1 / \gamma))}$ that $\gamma$-approximates the $O R_{d}$ function.

Our main contribution in this section is the following lemma that gives an improved polynomial approximation to the $\mathrm{OR}_{d}$ function restricted to inputs of low Hamming weight.

Lemma 4.2. For every $\gamma>0, k<d$ and $m \in[d] \backslash[k]$, there is a polynomial of degree $t=$ $O(d \sqrt{k} \log (1 / \gamma) / m)$ and weight $W=m^{O(\sqrt{k} \log (1 / \gamma))}$ that $\gamma$-approximates the $O R_{d}$ function restricted to inputs of Hamming weight at most $k$.

For any constant $\gamma$, one may take $m=d^{O(1 / \sqrt{k})}$ in the lemma (here the choice of constant depends on the constants in Fact 4.4) and obtain a polynomial of degree $d^{1-\Omega(1 / \sqrt{k})}$ and weight $d^{0.01}$.

Our constructions use the following basic facts.
Fact 4.3. The real polynomial $p_{d}:\{-1,1\}^{d} \rightarrow \mathbb{R}$

$$
p_{d}(x)=2\left(\sum_{S \subseteq[d]} 2^{-d} \prod_{i \in S} x_{i}\right)-1=2 \prod_{i=1}^{d}\left(\frac{1+x_{i}}{2}\right)-1
$$

computes $O R_{d}(x)$ and has weight $w\left(p_{d}\right) \leq 3$.
Fact 4.4. [see e.g., TUV12]] For every $k \in \mathbb{N}$ and $\gamma>0$, there exists a univariate real polynomial $p=\sum_{i=0}^{t_{k}} c_{i} x^{i}$ of degree $t_{k}$ such that

1. $t_{k}=O(\sqrt{k} \log (1 / \gamma))$,
2. for every $i \in\left[t_{k}\right],\left|c_{i}\right| \leq 2^{O(\sqrt{k} \log (1 / \gamma))}$,
3. $p(0)=0$, and
4. for every $x \in[2 k],|p(x)-1| \leq \gamma / 2$.

Proof of Lemma 4.1. We can compute $\mathrm{OR}_{d}(y)$ as a disjunction of disjunctions by partitioning the inputs $y_{1}, \ldots, y_{d}$ into blocks of size $d / m$ and computing:

$$
\operatorname{OR}_{m}\left(\operatorname{OR}_{d / m}\left(y_{1}, \ldots, y_{d / m}\right), \ldots, \operatorname{OR}_{d / m}\left(y_{d-d / m+1}, \ldots, y_{d}\right)\right)
$$

In order to approximately compute $\mathrm{OR}_{d}(y)$, we compute the inner disjunctions exactly using the polynomial $p_{d / m}$ given in Fact 4.3 and approximate the outer disjunction using the polynomial from Fact 4.4, Let

$$
Z(y)=p_{d / m}\left(y_{1}, \ldots, y_{d / m}\right)+\cdots+p_{d / m}\left(y_{d-d / m+1}, \ldots, y_{d}\right) .
$$

Setting $k=m$ in Fact 4.4, let $q_{m}$ be the resulting polynomial of degree $O(\sqrt{m} \log (1 / \gamma))$ and weight $O\left(m^{\sqrt{m} \log (1 / \gamma)}\right)$. Our final polynomial is

$$
1-2 q_{m}(m-Z(y)) .
$$

Note that $m-Z(y)$ takes values in $\{0, \ldots, m\}$ and is 0 exactly when all inputs $y_{1}, \ldots, y_{d}$ are FALSE. It follows that our final polynomial indeed approximates $\mathrm{OR}_{d}$ to additive error $\gamma$ on all Boolean inputs.

We bound the degree and weight of this polynomial in $y$. By Fact 4.3, the inner disjunctions are computed exactly using degree $d / m$ and weight at most 3 . Hence, the total degree is $O(\sqrt{m} \log (1 / \gamma) \cdot d / m)$. To bound the weight, we observe that the outer polynomial $q_{m}(\cdot)$ has at most $T=m^{O(\sqrt{m} \log (1 / \gamma))}$ terms where each one has degree at most $D_{\text {outer }}=O(\sqrt{m} \log (1 / \gamma))$ and coefficients of absolute value at most $c_{\text {outer }}=2^{O(\sqrt{m} \log (1 / \gamma))}$. Expanding the polynomials for $Z(y)$, the weight of each term incurs a multiplicative factor of $c_{\text {inner }} \leq 3^{D_{\text {outer }}}=3^{O(\sqrt{m} \log 1 / \gamma)}$ so the total weight is at most $c_{\text {inner }} \cdot c_{\text {outer }} \cdot T=m^{O(\sqrt{m} \log 1 / \gamma)}$.

Proof of Lemma 4.2. Again we partition the inputs $y_{1}, \ldots, y_{d}$ into blocks of size $d / m$ and view the disjunction as:

$$
\mathrm{OR}_{m}\left(\operatorname{OR}_{d / m}\left(y_{1}, \ldots, y_{d / m}\right), \ldots, \mathrm{OR}_{d / m}\left(y_{d-d / m+1}, \ldots, y_{d}\right)\right)
$$

Once again, we compute the inner disjunctions exactly using the polynomial from Fact 4.3, Let

$$
Z(y)=p_{d / m}\left(y_{1}, \ldots, y_{d / m}\right)+\cdots+p_{d / m}\left(y_{d-d / m+1}, \ldots, y_{d}\right) .
$$

If the input $y$ has Hamming weight at most $k$, then $Z(y)$ also takes values in $\{m, \ldots, m-2 k\}$. Thus, we may approximate the outer disjunction using a polynomial of degree $O(\sqrt{k} \log (1 / \gamma))$ from Fact 4.4. Our final polynomial is:

$$
1-2 q_{k}(m-Z) .
$$

The bound on degree and weight may be obtained as in the previous lemma.

### 4.1 Proof of upper bound theorems

We first present the proof of Theorem [1.3.
Proof of Theorem 1.3. Taking $m=O\left((\log d / \log \log d)^{2}\right)$ in Lemma 4.1 and $m=d^{O(1 / \sqrt{k})}$ in Lemma 4.2, it follows that for some constant $C>0$, the $d$-variate disjunction restricted to $H_{d, k}$ is (1/400)-approximated by a $d$-variate real polynomial of degree $t$ and weight $W$ where

$$
t=\min \left\{d^{1-\frac{1}{C \sqrt{k}}}, \frac{d}{\log ^{0.995} d}\right\} \text { and } W=d^{0.01}
$$

Proof of Theorem [1.1. By Theorem [1.3, we have a polynomial $p$ that ( $1 / 400$ )-approximates the function $\left.O R_{d}\right|_{H_{d, k}}$. Moreover, $p$ has weight $W<=d^{0.01}$ and degree

$$
t \leq \min \left\{d^{1-\frac{1}{C \sqrt{k}}}, \frac{d}{\log ^{0.995} d}\right\}
$$

Thus, by Fact 2.5, we have a family of polynomials $\mathcal{P}_{t, W}(d)$ that ( $1 / 400$ )-approximates the function family $\left.\mathcal{F}_{\mathcal{Q}}\right|_{H_{d, k}}$. We have the ingredients needed to apply Theorem 3.1. Taking $\alpha=1 / 100, \beta=$ $1 / 100$ gives the conclusion.

Remark 1 in the Introduction follows from using a slightly different choice of $m$ in Lemma 4.1, namely $m=O\left(\log ^{2} d / \log ^{3} \log d\right)$.

To obtain the summary of the database promised in Remark 2, we request an answer to each of the $k$-way marginal queries $B(1 / 400)$ times. Doing so, will ensure that we obtain a maximal database update sequence, and it was argued in Section 2.1 that the polynomial resulting from any maximal database update sequence accurately answers every $k$-way marginal query. Finally, we obtain a compact summary by randomly choosing $\tilde{O}\left(k d^{0.01}\right)$ samples from the normalized coefficient vector of this polynomial to obtain a new sparse polynomial that accurately answers every $k$-way marginal query (see e.g. [BS92]). Our compact summary is this final sparse polynomial.

## 5 Lower Bounds

In this section, we address the general problem of approximating a block-composed function $G=$ $F(\ldots, f(),. \ldots)$, where $F:\{-1,1\}^{k} \rightarrow\{-1,1\}, f: Y \rightarrow\{-1,1\}, Y \subseteq \mathbb{R}^{d / k}$ over inputs restricted to a set $\mathcal{Y} \subseteq Y^{k}$ using low-weight polynomials. We give a lower bound on the minimum degree of such polynomials. In our main application, $G$ will be $\mathrm{OR}_{d}$, and $\mathcal{Y}$ will be the set of all $d$-dimensional Boolean vectors of Hamming weight at most $k$.

Our proof technique is inspired by the composition theorem lower bounds of She09, Theorem 3.1], where it is shown that the $\gamma$-approximate degree of the composed function $G$ is at least the product of the $\gamma$-approximate degree of the outer function and the PTF degree of the inner function. Our main contribution is a generalization of such a composition theorem along two directions: (1) we show degree lower bounds that take into account the L1-norm of the coefficient vector of the approximating polynomial, and (2) our lower bounds hold even when we require the approximation to be accurate only on inputs of low Hamming weight, while prior work only considered approximations that are accurate on all Boolean inputs.

Our main theorem is stated below. In parsing the statement of the theorem, it may be helpful to think of $G=\mathrm{OR}_{d}, \mathcal{Y}=H_{d, k}$, the set of all $d$-dimensional Boolean vectors of Hamming weight at most $k, f=\mathrm{OR}_{d / k}, F=\mathrm{OR}_{k}, Y=\{-1,1\}^{d / k}$, and $H=H_{d / k, 1}$. This will be the setting of interest in our main application of the theorem.

Theorem 5.1. Let $Y \subset \mathbb{R}^{d / k}$ be a finite set and $\gamma>0$. Given $f: Y \rightarrow\{-1,1\}$ and $F:\{-1,1\}^{k} \rightarrow$ $\{-1,1\}$ such that $\operatorname{deg}_{2 \gamma}(F)=D$, let $G: Y^{k} \rightarrow\{-1,1\}$ denote the composed function defined by $G\left(Y_{1}, \ldots, Y_{k}\right)=F\left(f\left(Y_{1}\right), \ldots, f\left(Y_{k}\right)\right)$. Let $\mathcal{Y} \subseteq Y^{k}$. Suppose there exists $H \subseteq Y$ such that for every $\left(Y_{1}, \ldots, Y_{k}\right) \in Y^{k} \backslash \mathcal{Y}$ there exists $i \in[k]$ such that $Y_{i} \in Y \backslash H$. Then, for every $t \in \mathbb{Z}_{+}$,

$$
\operatorname{deg}_{(\gamma, W)}\left(\left.G\right|_{\mathcal{Y}}\right) \geq \frac{1}{2} t D \text { for every } W \leq \gamma 2^{-k} w^{*}\left(\left.f\right|_{H}, t\right)^{\frac{D}{2}}
$$

We will derive the following corollary from Theorem [5.1. Theorem 1.4 follows immediately from Corollary 5.2 by considering any $k=o(\log d)$.

Corollary 5.2. Let $k \in[d]$. Then, there exists a universal constant $C>0$ such that

$$
\operatorname{deg}_{(1 / 6, W)}\left(\left.O R_{d}\right|_{H_{d, k}}\right)=\Omega\left(\frac{d}{\sqrt{k}} \cdot \frac{W^{-\frac{1}{C \sqrt{k}}}}{2^{\sqrt{k} / C}}\right) \text { for every } W \leq \frac{1}{6.2^{k}}\left(\frac{d}{k}\right)^{C \sqrt{k}} .
$$

Intuition underlying our proof technique. Recall that our upper bound in Section 4 worked as follows. We viewed $\mathrm{OR}_{d}$ as an "OR of ORs", and we approximated the outer OR with a polynomial $p$ of degree deg $_{\text {outer }}$ chosen to be as small as possible, and composed $p$ with a lowweight but high-degree polynomial computing each inner OR. We needed to make sure the weight $W_{\text {inner }}$ of the inner polynomials was very low, because the composition step potentially blows the weight up to roughly $W_{\text {inner }}^{\text {deger }^{\prime}}$. As a result, the inner polynomials had to have very high degree, to keep their weight low.

Intuitively, we construct a dual solution to a certain linear program that captures the intuition that any low-weight, low-degree polynomial approximation to $\mathrm{OR}_{d}$ must look something like our primal solution, composing a low-degree approximation to an "outer" OR with low-weight approximations to inner ORs. Moreover, our dual solution formalizes the intuition that the composition step must result in a massive blowup in weight, from $W_{\text {inner }}$ to roughly $W_{\text {inner }}^{\mathrm{deg}_{\text {outer }}}$.

In more detail, our dual construction works by writing $\mathrm{OR}_{d}$ as an OR of ORs, where the outer OR is over $k$ variables, and each inner ORs is over $d / k$ variables. We obtain our dual solution by carefully combining a dual witness $\Gamma$ to the high approximate degree of the outer OR, with a dual witness $\psi$ to the fact that any low-degree polynomial with margin at least 1 for each inner OR, must have "large" weight, even if the polynomial must satisfy the margin constraint only on inputs of Hamming weight 0 or 1 . This latter condition, that $\psi$ must witness high non-constant margin weight even if restricted to inputs of Hamming weight 0 or 1 , is essential to ensuring that our combined dual witness does not place any "mass" on irrelevant inputs, i.e. those of Hamming weight larger than $k$.

### 5.1 Duality Theorems

In the rest of the section, we let $\chi_{S}(x)=\prod_{i \in S} x_{i}$ for any given set $S \subseteq[d]$. The question of existence of a weight $W$ polynomial with small degree that $\gamma$-approximates a given function can be expressed as a feasibility problem for a linear program. Now, in order to show the non-existence of such a polynomial, it is sufficient to show infeasibility of the linear program. By duality, this is equivalent to demonstrating the existence of a solution to the corresponding dual program. We begin by summarizing the duality theorems that will be useful in exhibiting this witness.
Theorem 5.3 (Duality Theorem for $(\gamma, W)$-approximate degree). Fix $\gamma \geq 0$ and let $f: Y \rightarrow$ $\{-1,1\}$ be given for some finite set $Y \subset \mathbb{R}^{d}$. Then, $\operatorname{deg}_{(\gamma, W)}(f) \geq t+1$ if and only if there exists a function $\Psi: Y \rightarrow \mathbb{R}$ such that

1. $\sum_{y \in Y}|\Psi(y)|=1$,
2. $\sum_{y \in Y} \Psi(y) f(y)-W \cdot\left|\sum_{y \in Y} \Psi(y) \chi_{S}(y)\right|>\gamma$ for every $S \subseteq[d],|S| \leq t$.

Proof. By definition, $\operatorname{deg}_{(\gamma, W)}(f) \leq t$ if and only if $\exists\left(\lambda_{S}\right)_{S \subseteq[d]|, S| \leq t}$ :

$$
\begin{aligned}
\sum_{S \subseteq[d]],|S| \leq t}\left|\lambda_{S}\right| & \leq W, \text { and } \\
\left|f(y)-\sum_{S \subseteq[d]],|S| \leq t} \lambda_{S} \chi_{S}(y)\right| & \leq \gamma \forall y \in Y .
\end{aligned}
$$

By Farkas' lemma, $\operatorname{deg}_{(\gamma, W)}(f) \leq t$ if and only if $\nexists \Psi: Y \rightarrow \mathbb{R}$ such that

$$
\frac{1}{W} \sum_{y \in Y}(f(y) \Psi(y)-\gamma|\Psi(y)|)>\left|\sum_{y \in Y} \chi_{S}(y) \Psi(y)\right| \forall S \subseteq[d],|S| \leq t
$$

The dual witness that we construct to prove Theorem 5.1 is obtained by combining a dual witness for the large non-constant margin weight of the inner function with a dual witness for the large approximate degree for the outer function. The duality conditions for these are given below. The proof of the duality condition for the case of $\gamma$-approximate degree is well-known, so we omit the proof for brevity.

Theorem 5.4 (Duality Theorem for $\gamma$-approximate degree). [She11, Š08, BT13] Fix $\gamma \geq 0$ and let $f: Y \rightarrow\{-1,1\}$ be given, where $Y \subset \mathbb{R}^{d}$ is a finite set. Then, $\operatorname{deg}_{\gamma}(f) \geq t+1$ if and only if there exists a function $\Gamma: Y \rightarrow \mathbb{R}$ such that

1. $\sum_{y \in Y}|\Gamma(y)|=1$,
2. $\sum_{y \in Y} \Gamma(y) p(y)=0$ for every polynomial $p$ of degree at most $t$, and
3. $\sum_{y \in Y} \Gamma(y) f(y)>\gamma$.

Theorem 5.5 (Duality Theorem for non-constant margin weight). Let $Y \subset \mathbb{R}^{d}$ be a finite set, let $f: Y \rightarrow\{1,-1\}$ be a given function and $w>0$. The non-constant margin weight $w^{*}(f, t) \geq w$ if and only if there exists a distribution $\mu: Y \rightarrow[0,1]$ such that

1. $\sum_{y \in Y} \mu(y) f(y)=0$
2. $\left|\sum_{y \in Y} \mu(y) f(y) \chi_{S}(y)\right| \leq \frac{1}{w}$ for every $S \subseteq[d],|S| \leq t$.

Proof. Let $\mathcal{S}=\{S \subseteq[d]:|S| \leq t\}, \overline{\mathcal{S}}=\mathcal{S} \backslash\{\emptyset\}$. By definition, $w^{*}(f, t)$ is expressed by the following linear program:

$$
\begin{gathered}
\min \sum_{S \in \overline{\mathcal{S}}}\left|\lambda_{S}\right| \\
f(y) \sum_{S \in \mathcal{S}} \lambda_{S} \chi_{S}(y) \geq 1 \forall y \in Y .
\end{gathered}
$$

The above linear program can be restated as follows:

$$
\begin{aligned}
\min \sum_{S \in \overline{\mathcal{S}}} \alpha_{S} & \\
\alpha_{S}+\lambda_{S} & \geq 0 \forall S \in \overline{\mathcal{S}}, \\
\alpha_{S}-\lambda_{S} & \geq 0 \forall S \in \overline{\mathcal{S}}, \\
f(y) \sum_{S \in \mathcal{S}} \lambda_{S} \chi_{S}(y) & \geq 1 \forall y \in Y, \text { and } \\
\alpha_{S} & \geq 0 \forall S \in \overline{\mathcal{S}} .
\end{aligned}
$$

The dual program is expressed below:

$$
\begin{aligned}
\max \sum_{y} \mu(y) & \\
u_{1}(S)+u_{2}(S) & \leq 1 \forall S \in \overline{\mathcal{S}}, \\
\sum_{y \in Y} \mu(y) f(y) \chi_{S}(y)+u_{1}(S)-u_{2}(S) & =0 \forall S \in \overline{\mathcal{S}}, \\
\sum_{y \in Y} \mu(y) f(y) & =0 \\
\mu(y) \geq 0 \forall y \in Y, u_{1}(S), u_{2}(S) & \geq 0 \forall S \in \overline{\mathcal{S}} .
\end{aligned}
$$

By standard manipulations, the above dual program is equivalent to

$$
\begin{aligned}
\max \sum_{y} \mu(y) & \\
\left|\sum_{y \in Y} \mu(y) \chi_{S}(y) f(y)\right| & \leq 1 \forall S \in \overline{\mathcal{S}} \\
\sum_{y \in Y} \mu(y) f(y) & =0, \\
\mu(y) & \geq 0 \forall y \in Y
\end{aligned}
$$

Finally, given a distribution $\mu^{\prime}$ satisfying the hypothesis of the theorem, one can obtain a dual solution $\mu$ to show that $w^{*}(f, t) \geq w$ by taking $w^{-1}=\max _{S \in \mathcal{S}}\left|\sum_{y \in Y} \mu^{\prime}(y) \chi_{S}(y) f(y)\right|$ and setting $\mu(y)=w \mu^{\prime}(y) \forall y \in Y$. In the other direction, if $w^{*}(f, t) \geq w$, then we have a dual solution $\mu$ satisfying the above dual program such that $\sum_{y \in Y} \mu(y)=w^{*}(f, t)$. By setting $\mu^{\prime}(y)=$ $\mu(y) / w^{*}(f, t) \forall y \in Y$, we obtain the desired distribution.

### 5.2 Proof of Theorem 5.1

Our approach to exhibiting a dual witness as per Theorem 5.3 is to build a dual witness by appropriately combining the dual witnesses for the "hardness" of the inner and outer functions. Our method of combining the dual witnesses is inspired by the technique of [She09, Theorem 3.7].

Proof of Theorem 5.1. Let $w=w^{*}\left(\left.f\right|_{H}, t\right)$. We will exhibit a dual witness function $\Psi: \mathcal{Y} \rightarrow \mathbb{R}$ corresponding to Theorem 5.3 for the specified choice of degree and weight. For $y \in Y^{k}$, let $Y_{i}=\left(y_{(i-1)(d / k)+1}, \ldots, y_{i d / k}\right)$. By Theorem [5.5, we know that there exists a distribution $\mu: H \rightarrow \mathbb{R}$ such that

$$
\begin{align*}
\sum_{y \in H} \mu(y) f(y) & =0  \tag{2}\\
\left|\sum_{y \in H} \mu(y) f(y) \chi_{S}(y)\right| & \leq \frac{1}{w} \forall S \subseteq\left[\frac{d}{k}\right],|S| \leq t \tag{3}
\end{align*}
$$

We set $\mu(y)=0$ for $y \in Y \backslash H$.

Since $\operatorname{deg}_{2 \gamma}(F)=D$, by Theorem [5.4, we know that there exists a function $\Gamma:\{-1,1\}^{k} \rightarrow \mathbb{R}$ such that

$$
\begin{align*}
\sum_{x \in\{-1,1\}^{k}}|\Gamma(x)| & =1  \tag{4}\\
\sum_{x \in\{-1,1\}^{k}} \Gamma(x) p(x) & =0 \text { for every polynomial } p \text { of degree at most } D, \text { and }  \tag{5}\\
\sum_{x \in\{-1,1\}^{k}} \Gamma(x) F(x) & >2 \gamma . \tag{6}
\end{align*}
$$

Consider the function $\Psi: Y^{k} \rightarrow \mathbb{R}$ defined as $\Psi(y)=2^{k} \Gamma\left(f\left(Y_{1}\right), \ldots, f\left(Y_{k}\right)\right) \prod_{i=1}^{k} \mu\left(Y_{i}\right)$. By the hypothesis of the theorem, we know that if $\left(Y_{1}, \ldots, Y_{k}\right) \in Y^{k} \backslash \mathcal{Y}$, then there exists $i \in[k]$ such that $Y_{i} \in Y \backslash H$ and hence $\mu\left(Y_{i}\right)=0$ and therefore $\Psi\left(Y_{1}, \ldots, Y_{k}\right)=0$.
1.

$$
\begin{aligned}
\sum_{y \in \mathcal{Y}}|\Psi(y)| & =\sum_{y \in \mathcal{Y}} 2^{k}\left|\Gamma\left(f\left(Y_{1}\right), \ldots, f\left(Y_{k}\right)\right)\right| \prod_{i=1}^{k} \mu\left(Y_{i}\right) \\
& =2^{k} \mathbb{E}_{y \sim \Phi}\left(\left|\Gamma\left(f\left(Y_{1}\right), \ldots, f\left(Y_{k}\right)\right)\right|\right)
\end{aligned}
$$

where $y \sim \Phi$ denotes $y$ chosen from the product distribution $\Phi: Y^{k} \rightarrow[0,1]$ defined by $\Phi(y)=\prod_{i \in[k]} \mu\left(Y_{i}\right)$. Since $\sum_{y \in Y} \mu(y) f(y)=0$, it follows that if $Y_{i}$ is chosen with probability $\mu\left(Y_{i}\right)$, then $f\left(Y_{i}\right)$ is uniformly distributed in $\{-1,1\}$. Consequently,

$$
\sum_{y \in \mathcal{Y}}|\Psi(y)|=2^{k} \mathbb{E}_{z \sim U}\{-1,1\}^{k}\left(\left|\Gamma\left(z_{1}, \ldots, z_{k}\right)\right|\right)=1
$$

The last equality is by using (4).
2. By the same reasoning as above, it follows from (6) that

$$
\sum_{y \in \mathcal{Y}} \Psi(y) G(y)=\sum_{z \in\{-1,1\}^{k}} \Gamma(z) F(z)>2 \gamma
$$

3. Fix a subset $S \subseteq[d]$ of size at most $t D / 2$. Let $S_{i}=S \cap\{(i-1)(d / k)+1, \ldots, i d / k\}$ for each $i \in[k]$. Consequently, $\chi_{S}(y)=\prod_{i=1}^{k} \chi_{S_{i}}\left(Y_{i}\right)$.
Now using the Fourier coefficients $\hat{\Gamma}(T)$ of the function $\Gamma$, we can express

$$
\Gamma\left(z_{1}, \ldots, z_{k}\right)=\sum_{T \subseteq[k]} \hat{\Gamma}(T) \prod_{i \in T} z_{i}=\sum_{\substack{T \subseteq[k],|T| \geq D}} \hat{\Gamma}(T) \prod_{i \in T} z_{i}
$$

since $\hat{\Gamma}(T)=0$ if $|T|<D$ by (5). Hence,

$$
\Psi(y)=2^{k} \sum_{\substack{T \subseteq[k],|T| \geq D}} \hat{\Gamma}(T) \prod_{i \in T} f\left(Y_{i}\right) \mu\left(Y_{i}\right) \cdot \prod_{i \in[k] \backslash T} \mu\left(Y_{i}\right)
$$

Therefore, $\sum_{y \in \mathcal{Y}} \Psi(y) \chi_{S}(y)$

$$
\begin{aligned}
& =\sum_{y \in \mathcal{Y}} \Psi(y) \prod_{i \in[k]} \chi_{S_{i}}\left(Y_{i}\right) \\
& =2^{k} \sum_{y \in \mathcal{Y}}\left(\sum_{\substack{T \subseteq[k],|T| \geq D}} \hat{\Gamma}(T) \prod_{i \in T} f\left(Y_{i}\right) \mu\left(Y_{i}\right) \cdot \prod_{i \in[k] \backslash T} \mu\left(Y_{i}\right)\right) \prod_{i \in[k]} \chi_{S_{i}}\left(Y_{i}\right) \\
& =2^{k} \sum_{\substack{T \subseteq[k],|T| \geq D}} \hat{\Gamma}(T) \sum_{y \in \mathcal{Y}}\left(\prod_{i \in T} f\left(Y_{i}\right) \mu\left(Y_{i}\right) \cdot \prod_{i \in[k] \backslash T} \mu\left(Y_{i}\right) \prod_{i \in[k]} \chi_{S_{i}}\left(Y_{i}\right)\right) \\
& =2^{k} \sum_{\substack{T \subseteq[k],|T| \geq D}} \hat{\Gamma}(T) \sum_{Y_{1}, \ldots, Y_{k} \in H}\left(\prod_{i \in T} f\left(Y_{i}\right) \mu\left(Y_{i}\right) \chi_{S_{i}}\left(Y_{i}\right) \cdot \prod_{i \in[k] \backslash T} \mu\left(Y_{i}\right) \chi_{S_{i}}\left(Y_{i}\right)\right) .
\end{aligned}
$$

Rearranging, we have $\sum_{y \in \mathcal{Y}} \Psi(y) \chi_{S}(y)=$

$$
\begin{equation*}
2^{k} \sum_{\substack{T \subseteq[k],|T| \geq D}} \hat{\Gamma}(T) \prod_{i \in T}\left(\sum_{Y_{i} \in H} f\left(Y_{i}\right) \mu\left(Y_{i}\right) \chi_{S_{i}}\left(Y_{i}\right)\right) \prod_{i \in[k] \backslash T}\left(\sum_{Y_{i} \in H} \mu\left(Y_{i}\right) \chi_{S_{i}}\left(Y_{i}\right)\right) . \tag{7}
\end{equation*}
$$

Now, we will bound each product term in the outer sum by $w^{-D / 2}$. We first observe that for every $i \in[k]$,

$$
\sum_{x \in H} \mu(x) \chi_{S_{i}}(x) \leq \sum_{x \in H} \mu(x)=1
$$

If $\left|S_{i}\right| \leq t$, by (3)

$$
\left|\sum_{x \in H} f(x) \mu(x) \chi_{S_{i}}(x)\right| \leq \frac{1}{w} .
$$

If $\left|S_{i}\right|>t$, then

$$
\left|\sum_{x \in H} f(x) \mu(x) \chi_{S_{i}}(x)\right| \leq \sum_{x \in H} \mu(x)=1
$$

Since $\sum_{i=1}^{k}\left|S_{i}\right| \leq t D / 2$, it follows that $\left|S_{i}\right| \leq t$ for more than $k-D / 2$ indices $i \in[k]$. Thus, for each $T \subseteq[k]$ such that $|T| \geq D$, there are at least $D / 2$ indices $i \in T$ such that $\left|S_{i}\right| \leq t$. Hence,

$$
\left|\sum_{y \in \mathcal{Y}} \Psi(y) \chi_{S}(y)\right| \leq 2^{k} w^{-\frac{D}{2}} \sum_{T \subseteq[k],|T| \geq D}|\hat{\Gamma}(T)| \leq 2^{k} w^{-\frac{D}{2}}
$$

Here, the last inequality is because $|\hat{\Gamma}(T)| \leq 2^{-k}$ from (4).
From 1, 2 and 3, we have

$$
\sum_{y \in \mathcal{Y}} \Psi(y) G(y)-W \max _{S \subseteq[d],|S| \leq \frac{t D}{2}}\left|\sum_{y \in \mathcal{Y}} \Psi(y) \chi_{S}(y)\right|>\gamma
$$

if $W \leq \gamma 2^{-k} w^{D / 2}$.
We now derive Corollary 5.2. We need the following theorems on the approximate degree and the non-constant margin weight of the $\mathrm{OR}_{d}$ function.
Theorem 5.6 (Approximate degree of $\left.\mathrm{OR}_{d}\right)$. Pat92] $\operatorname{deg}_{1 / 3}\left(O R_{d}\right)=\Theta(\sqrt{d})$.
Lemma 5.7 (Non-constant margin weight of $\left.\mathrm{OR}_{d}\right) \cdot w^{*}\left(\left.O R_{d}\right|_{H_{d, 1}}, t\right) \geq d / t$.
Proof. The function

$$
\mu(x)=\left\{\begin{aligned}
1 / 2 & \text { if } x=(1, \ldots, 1), \\
1 / 2 d & \text { if } x \in H_{d, 1} \backslash\{(1, \ldots, 1)\} .
\end{aligned}\right.
$$

acts as the dual witness in Theorem 5.5.
Proof of Corollary 5.2. We use Theorem 5.1 in the following setting. Let $Y=\{-1,1\}^{d / k}$, the inner function $f: Y \rightarrow\{-1,1\}$ be $\mathrm{OR}_{d / k}$ and the outer function $F:\{-1,1\}^{k} \rightarrow\{-1,1\}$ be $\mathrm{OR}_{k}$, $\mathcal{Y}=H_{d, k}$ and $H=H_{d / k, 1}$. By a simple counting argument, if $\left(Y_{1}, \ldots, Y_{k}\right) \in\{-1,1\}^{d} \backslash H_{d, k}$, then there exists $i \in[k]$ such that $Y_{i} \in\{-1,1\}^{d / k} \backslash H_{d / k, 1}$. Further, by Theorem 5.6, we know that $\operatorname{deg}_{1 / 3}(F)=\Theta(\sqrt{k})$ and by Claim 5.7, we know that $w^{*}\left(\left.f\right|_{H}, t\right) \geq d / k t$. Therefore, by Theorem 5.1. we have that, for every $t \in \mathbb{Z}_{+}$,

$$
\operatorname{deg}_{1 / 6, W}\left(\left.O R_{d}\right|_{H_{d, k}}\right)=\Omega(t \sqrt{k}) \text { for every } W \leq \frac{1}{6} 2^{-k}\left(\frac{d}{k t}\right)^{C \sqrt{k}} .
$$

We obtain the conclusion by taking $t=\left\lfloor(d / k)\left(6 W 2^{k}\right)^{-1 / C \sqrt{k}}\right\rfloor$. Since $W \leq(1 / 6) 2^{-k}(d / k)^{C \sqrt{k}}$, it follows that $t \geq 1$ and hence is a valid choice for $t$ in applying Theorem 5.1.

Comparison to [STT12]. As described in the beginning of Section 4, Lemma 8 of the work of Servedio et al. [STT12] can be shown to imply that any polynomial $p$ of weight $W$ that $1 / 3$ approximates the $\mathrm{OR}_{d}$ function on all Boolean inputs requires degree $\Omega(d / \log W)$ The proof in [TT12] relies on a Markov-type inequality that bounds the derivative of a univariate polynomial in terms of its degree and the size of its coefficients. The proof of this Markov-type inequality is non-constructive and relies on complex analysis.

Here, we observe that our dual witness construction used to prove Corollary 5.2 also yields a general lower bound on the tradeoffs achievable between the weight and degree of the approximating polynomial $p$, even when we require $p$ to be accurate only on inputs of Hamming weight at most $O(\log W)$ (see Theorem 5.8). The methods of Servedio et al. do not yield any non-trivial lower bound on the degree in this setting. We also believe our proof technique is of interest in comparison to the methods of Servedio et al. as it is constructive (exhibiting an explicit dual witness for the lower bound) and avoids the use of complex analysis.
Theorem 5.8. Any polynomial $p$ of weight $W$ that $1 / 6$-approximates the $O R_{d}$ function on all Boolean inputs requires degree $d / 2 O(\sqrt{\log W})$.
Proof. As $p$ is accurate on the entire Boolen hypercube, it is accurate on inputs of Hamming weight at most $\log W$. The theorem follows by setting $k=\log W$ in the statement of Corollary 5.2,

[^1]
## 6 Discussion

We gave a differentially private online algorithm for answering $k$-way marginal queries that runs in time $2^{o(d)}$ per query, and guarantees accurate answers for databases of size poly $(d, k)$. More precisely, we showed that if there exists a polynomial of degree $t$ and weight $W$ approximating the $d$-variate OR function on Boolean inputs of Hamming weight at most $k$, then a variant of the private multiplicative weights algorithm can answer $k$-way marginal queries in time roughly $\binom{d}{t}$ per query and guarantee accurate answers on databases of size roughly $W \sqrt{d}$. To this end, we gave a new construction showing the existence of polynomial approximations to the OR function on inputs of low Hamming weight. Specifically, we showed that polynomials of weight $d^{0.01}$ and degree $d^{1-\Omega(1 / \sqrt{k})}$ exist.

Practical Considerations. Our algorithm for answering $k$-way marginals is essentially the same as in HR10 except for using a different set of base functions (specifically, the set of all low-degree parities), which leads to an efficiency gain. We note that our algorithm degrades smoothly to the private multiplicative weights algorithm as the degree of the promised polynomial approximation increases, and never gives a worse running time. This behavior suggests that our algorithm may lead to practical improvements even for relatively small values of $d$, for which the asymptotic analysis does not apply. In such cases one might use an alternative but similar analysis that shows the existence of a polynomial of degree $k d^{1-c / k}$ and weight $d^{c}$ (for any $0<c<k$ ) that exactly computes the $d$-variate OR function on inputs of Hamming weight at most $k$. Such a polynomial may be obtained as in our construction, by breaking the $d$-variate OR function into an OR of ORs, and using a degree $k$ polynomial defined via polynomial interpolation, instead of Chebyshev polynomials, to approximate the outer OR on inputs of Hamming weight at most $k$. This variant does achieve asymptotic properties not shared by the algorithm of Theorem 1.1, owing to the fact that it uses exact rather than approximate representations of the database: for any constant $k$, this variant also runs in time $\exp \left(d^{1-\Omega(1)}\right)$, and achieves worst-case additive error $o(1)$ for sufficiently large databases, i.e., for $n \gtrsim \tilde{O}\left(d^{c}\right)$.

Relationship with [NTZ13, DNT13]. As we mentioned in the introduction, in subsequent work, Dwork et al. DNT13 show how to privately release marginals in a very different parameter regime from what we consider here. Although their algorithm is quite different from ours, there are some important similarities. Their algorithm is based on a recent algorithm of Nikolov et. al. [NTZ13] for answering arbitrary counting queries. This algorithm proceeds as follows. First, it adds noise $O\left(|\mathcal{Q}|^{1 / 2} /|D|\right)$ to the answer to every query to obtain a vector of noisy answers. The noise is sufficient to ensure differential privacy. Second, it "projects" the vector of noisy answers onto the convex body $K$ consisting of all vectors of answers that are consistent with some real database. Surprisingly, the projection step will improve the accuracy. In fact, Nikolov et al. [NTZ13] show that the projected answers will be accurate even for very small databases. However current best algorithms for projecting on the body $K$ require time $2^{O(d)}$ for an arbitrary set of queries.

Dwork et al. [DNT13] show that for the set of $k$-way marginals there is a convex body $L$ such that the projection into $L$ can be computed in time $\operatorname{poly}\left(\binom{d}{\leq k}\right)$ and $L$ approximates $K$ well enough to achieve accuracy on smaller databases than would be achievable with independent noise (however, these databases still have size $\left.d^{\Omega(k)}\right)$. Our approximation of $f_{D}$ by low-weight polynomials of degree $t=o(d)$ can be shown to imply a polytope $L^{\prime}$ with $2^{o(d)}$ vertices-which is sufficient to imply that projection can be computed in time $2^{o(d)}$ - that approximates $K$ well enough to achieve accuracy on databases of nearly optimal size (i.e., size roughly $k d^{0.51}$ ). Thus our approximation-theoretic
approach is relevant for understanding the capabilities and limitations of algorithms in the Nikolov et al. [NTZ13] framework.

Future Directions. Our lower bounds show that our polynomial approximation to the $\mathrm{OR}_{d}$ function on inputs of Hamming weight $k$ is essentially the best possible; in particular, we cannot hope to substantially improve the running time on poly $(d, k)$ size databases by giving approximating polynomials with better weight and degree bounds. This also rules out several natural candidates that can themselves be computed exactly by a low-weight polynomial of low-degree (e.g., the set of small-width conjunctions).

However, we do not know if it is possible to do better by using different feature spaces (other than the set of all low-degree monomials) to approximate all disjunctions over $d$ variables. There is some additional evidence from prior work that low-degree monomials may be the optimal choice: if the parameter of significance is the size of the set of functions used to approximate disjunctions on inputs of Hamming weight at most $k$, then low-degree monomials are indeed optimal She11] (see also Section 5 in the full version of TUV12]). It would be interesting to determine whether this optimality still holds when we restrict the $L_{1}$ weight of the linear combinations used in the approximations to be poly $(d)$.

Acknowledgments. We thank Salil Vadhan for helpful discussions about this work.

## References

$\left[\mathrm{BCD}^{+} 07\right]$ Boaz Barak, Kamalika Chaudhuri, Cynthia Dwork, Satyen Kale, Frank McSherry, and Kunal Talwar. Privacy, accuracy, and consistency too: a holistic solution to contingency table release. In PODS, pages 273-282, 2007.
[BDMN05] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: the sulq framework. In PODS, pages 128-138, 2005.
[BLR08] Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to noninteractive database privacy. In STOC, pages 609-618, 2008.
[BS92] Jehoshua Bruck and Roman Smolensky. Polynomial Threshold Functions, $A C^{0}$ Functions, and Spectral Norms. SIAM J. Comput., 21(1):33-42, 1992.
[BT13] Mark Bun and Justin Thaler. Dual lower bounds for approximate degree and markovbernstein inequalities. In Proceedings of 40th International Colloquium on Automata, Languages, and Programming, pages 303-314, 2013.
[BUV13] Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the true price of differential privacy. Manuscript, 2013.
[CKKL12] Mahdi Cheraghchi, Adam Klivans, Pravesh Kothari, and Homin K. Lee. Submodular functions are noise stable. In Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 1586-1592, 2012.
[De12] Anindya De. Lower bounds in differential privacy. In TCC, pages 321-338, 2012.
[DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In TCC '06, pages 265-284, 2006.
[DN03] Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In PODS, pages 202-210, 2003.
[DN04] Cynthia Dwork and Kobbi Nissim. Privacy-preserving datamining on vertically partitioned databases. In CRYPTO, pages 528-544, 2004.
[DNR $\left.{ }^{+} 09\right]$ Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil P. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In STOC '09, pages 381-390, 2009.
[DNT13] Cynthia Dwork, Aleksandar Nikolov, and Kunal Talwar. Efficient algorithms for privately releasing marginals via convex relaxations. CoRR, abs/1308.1385, 2013.
[DRV10] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and differential privacy. In FOCS, pages 51-60, 2010.
[FK13] Vitaly Feldman and Pravesh Kothari. Learning coverage functions. Manuscript, 2013.
[GHRU11] Anupam Gupta, Moritz Hardt, Aaron Roth, and Jonathan Ullman. Privately releasing conjunctions and the statistical query barrier. In STOC '11, pages 803-812, 2011.
[GRU12] Anupam Gupta, Aaron Roth, and Jonathan Ullman. Iterative constructions and private data release. In TCC, pages 339-356, 2012.
[HLM12] Moritz Hardt, Katrina Ligett, and Frank McSherry. A simple and practical algorithm for differentially private data release. NIPS '12, 2012.
[HR10] Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacypreserving data analysis. In FOCS, pages 61-70, 2010.
[HRS12] Moritz Hardt, Guy N. Rothblum, and Rocco A. Servedio. Private data release via learning thresholds. In SODA, pages 168-187, 2012.
[JT12] Prateek Jain and Abhradeep Thakurta. Mirror descent based database privacy. In APPROX-RANDOM, pages 579-590, 2012.
[Kla11] Hartmut Klauck. On arthur merlin games in communication complexity. In Proceedings of the 26th Annual Conference on Computational Complexity (CCC), CCC, pages 189199, 2011.
[KRSU10] Shiva Prasad Kasiviswanathan, Mark Rudelson, Adam Smith, and Jonathan Ullman. The price of privately releasing contingency tables and the spectra of random matrices with correlated rows. In STOC, pages 775-784, 2010.
[KS04] Adam R. Klivans and Rocco A. Servedio. Toward attribute efficient learning of decision lists and parities. In John Shawe-Taylor and Yoram Singer, editors, COLT, volume 3120 of Lecture Notes in Computer Science, pages 224-238. Springer, 2004.
[LS13] Philip M. Long and Rocco A. Servedio. Low-weight halfspaces for sparse boolean vectors. In ITCS, pages 21-36, 2013.
[NTZ13] Aleksandar Nikolov, Kunal Talwar, and Li Zhang. The geometry of differential privacy: the sparse and approximate cases. In STOC, pages 351-360, 2013.
[Pat92] R. Paturi. On the degree of polynomials that approximate symmetric boolean functions. In Proceedings of the twenty-fourth annual ACM Symposium on Theory of Computing, pages 468-474, 1992.
[RR10] Aaron Roth and Tim Roughgarden. Interactive privacy via the median mechanism. In STOC '10, pages 765-774, 2010.
[She09] Alexander A. Sherstov. The intersection of two halfspaces has high threshold degree. In Proceedings of the 2009 50th Annual IEEE Symposium on Foundations of Computer Science, FOCS '09, pages 343-362, 2009.
[She11] A. A. Sherstov. The pattern matrix method. SIAM Journal on Computing, 40(6):19692000, 2011.
[She12a] Alexander A. Sherstov. The multiparty communication complexity of set disjointness. In STOC, pages 525-548, 2012.
[She12b] Alexander A. Sherstov. Strong direct product theorems for quantum communication and query complexity. SIAM J. Comput., 41(5):1122-1165, 2012.
[STT12] Rocco Servedio, Li-Yang Tan, and Justin Thaler. Attribute-efficient learning and weight-degree tradeoffs for polynomial threshold functions. In 25th Annual Conference on Computational Learning Theory (COLT), JMLR Workshop and Conference Proceeding, volume 23, pages 14.1-14.19, 2012.
[TUV12] Justin Thaler, Jonathan Ullman, and Salil P. Vadhan. Faster algorithms for privately releasing marginals. In ICALP, pages 810-821, 2012.
[Ull13] Jonathan Ullman. Answering $\mathrm{n}^{2+\mathrm{o}(1)}$ counting queries with differential privacy is hard. In STOC, pages 361-370, 2013.
[UV11] Jonathan Ullman and Salil P. Vadhan. PCPs and the hardness of generating private synthetic data. In TCC '11, pages 400-416, 2011.
[Š08] Robert Špalek. A dual polynomial for OR. CoRR, abs/0803.4516, 2008.


[^0]:    *Harvard University, School of Engineering and Applied Sciences. Email: \{karthe, jthaler, jullman, atw12\}@seas.harvard.edu. Karthekeyan Chandrasekaran is supported by Simons Fellowship. Justin Thaler is supported by an NSF Graduate Research Fellowship and NSF grants CNS-1011840 and CCF-0915922. Jonathan Ullman is supported by NSF grant CNS- 1237235 and a Siebel Scholarship. Andrew Wan is supported by NSF grant CCF0964401 and NSFC grant 61250110218.

[^1]:    ${ }^{1}$ More precisely, [STT12, Lemma 8] as stated shows that if the coefficients of a univariate polynomial $P$ each have absolute value at most $W$, and $1 / 2 \leq \max _{x \in[0,1]}|P(x)| \leq R$, then $\max _{x \in[0,1]}\left|P^{\prime}(x)\right|=O(\operatorname{deg}(P) \cdot R \cdot(\log W+$ $\log \operatorname{deg}(P))$ ), where $P^{\prime}(x)$ denotes the derivative of $P$ at $x$, and $\operatorname{deg}(P)$ denotes the degree of $P$. By inspection of the proof, it is easily seen that if the $L_{1}$-norm of the coefficients of $P$ is bounded by $W$, then the following slightly stronger conclusion holds: $\max _{x \in[0,1]}\left|P^{\prime}(x)\right|=O(\operatorname{deg}(P) \cdot R \cdot \log W)$. When combined with the symmetrization argument of [STT12], this stronger conclusion implies: any polynomial $p$ of weight $W$ that $1 / 3$-approximates the $\mathrm{OR}_{d}$ function on all Boolean inputs requires degree $\Omega(d / \log W)$.

