# Polylogarithmic Private Approximations and Efficient Matching 

Piotr Indyk ${ }^{1}$ and David Woodruff ${ }^{1,2}$<br>${ }^{1}$ MIT CSAIL<br>\{indyk, dpwood\}@mit.edu<br>${ }^{2}$ Tsinghua University


#### Abstract

In 12] a private approximation of a function $f$ is defined to be another function $F$ that approximates $f$ in the usual sense, but does not reveal any information about $x$ other than what can be deduced from $f(x)$. We give the first two-party private approximation of the $l_{2}$ distance with polylogarithmic communication. This, in particular, resolves the main open question of 12 .

We then look at the private near neighbor problem in which Alice has a query point in $\{0,1\}^{d}$ and Bob a set of $n$ points in $\{0,1\}^{d}$, and Alice should privately learn the point closest to her query. We improve upon existing protocols, resolving open questions of [13, 10]. Then, we relax the problem by defining the private approximate near neighbor problem, which requires introducing a notion of secure computation of approximations for functions that return sets of points rather than values. For this problem we give several protocols with sublinear communication.


Keywords: private approximations, secure multiparty computation, nearest neighbor, communication complexity.

## 1 Introduction

Recent years witnessed the explosive growth of the amount of available data. Large data sets, such as transaction data, the web and web access logs, or network traffic data, are in abundance. Much of the data is stored or made accessible in a distributed fashion. This neccessitates the development of efficient protocols that compute or approximate functions over such data (e.g. see [2]).

At the same time, the availability of this data has raised significant privacy concerns. It became apparent that one needs cryptographic techniques in order to control data access and prevent potential misuse. In principle, this task can be achieved using the general results of secure function evaluation (SFE) [33, 18. However, in most cases the resulting private protocols are much less efficient than their non-private counterparts. Moreover, SFE applies only to algorithms that compute functions exactly, while for most massive data sets problems, only

[^0]efficient approximation algorithms are known or are possible. Indeed, while it is true that SFE can be used to privately implement any efficient algorithm, it is of little use applying it to an approximation algorithm when the approximation leaks more information about the input than the solution itself.

In a pioneering paper [12], the authors introduced a framework for secure computation of approximations. They also proposed an $\tilde{O}(\sqrt{n})$-communication ${ }^{2}$ two-party protocol for approximating the Hamming distance between two binary vectors. This improves over the linear complexity of computing the distance exactly via SFE, but still does not achieve the polylogarithmic efficiency of a non-private protocol of [25]. Improving the aforementioned bound was one of the main problems left open in [12].

In this paper we provide several new results for secure computation of approximations. Our first result is an $\tilde{O}(1)$-communication protocol for approximating the Euclidean $\left(\ell_{2}\right)$ distance between two vectors. This, in particular, solves the open problem of [12]. Since distance computation is a basic geometric primitive, we believe that our result could lead to other algorithms for secure approximations. Indeed, in [1] the authors show how to approximate the $\ell_{2}$ distance using small space and/or short amount of communication, initiating a rich body of work on streaming algorithms.

In the second part of the paper, we look at secure computation of a near neighbor for a query point $q$ (held by Alice) among $n$ data points $P$ (held by Bob) in $\{0,1\}^{d}$. We improve upon known results [10, 13] for this problem under various distance metrics, including $\ell_{2}$, set difference, and Hamming distance over arbitrary alphabets. Our techniques also result in better communication for the all-near neighbors problem, where Alice holds $n$ different query points, resolving an open question of [13], and yield a binary inner product protocol with communication $d+O(k)$ in the common random string model.

| Complexity | Problem | Prior work | SFE |
| :---: | :---: | :---: | :---: |
| $\tilde{O}(n+d)$ | Near neighbor under $l_{2}$, <br> Hamming over $\{0,1\}^{d}$, Set difference | $[10$ | $\tilde{O}(n d)$ |
| $\tilde{O}(d U+n)$ | Near neighbor under distances <br> $f(a, b)=\sum_{i=1}^{d} f_{i}\left(a_{i}, b_{i}\right), a_{i}, b_{i} \in[U]$ | $[10]$ | $\tilde{O}(n d \log U)$ |
| $\lceil\log d\rceil d+O(k)$ | Hamming distance | $[14]$ | $O(k d)$ |
| $\tilde{O}\left(n d^{2}+n^{2}\right)$ | All-near neighbors | $[13]$ | $\tilde{O}\left(n^{2} d\right)$ |

However, all of our protocols for the near neighbor problem have the drawback of needing $\Omega(n)$ bits of communication, though the dependence on $d$ is often optimal. Thus, we focus on what we term the approximate near neighbor problem. For this we introduce a new definition of secure computation of approximations for functions that return points (or sets of points) rather than values.

[^1]Approximate privacy. Let $P_{t}(q)$ be the set of points in $P$ within distance $t$ from $q$. In the c-approximate near neighbor problem, the protocol is required to report a point in $P_{c r}(q)$, as long as $P_{r}(q)$ is nonempty. We say that a protocol solving this problem is $c^{\prime}$-private (or just private if $c^{\prime}=c$ ) if Bob learns nothing, while Alice learns nothing except what can be deduced from the set $P_{c^{\prime} r}(q)$. In our paper we always set $c^{\prime}=c$.

We believe this to be a natural definition of privacy in the context of the approximate near neighbor problem. First, observe that if we insist that Alice learns only the set $P_{r}$ (as opposed to $P_{c r}$ ), then the problem degenerates to the exact near neighbor problem. Indeed, even though the definition of correctness allows the protocol to output a point $p \in P_{c r}-P_{r}$, in general Alice cannot simulate this protocol given only the set $P_{r}$. Thus, in order to make use of the flexibility provided by the approximate definition of the problem, it seems necessary to relax the definition of privacy as well.

Second, the above relaxation of privacy appears natural in the context of applications of near neighbor algorithms. In most situations, the distance function is only a heuristic approximation of the dis-similarity between objects, and there is no clear rationale for a sharp barrier between objects that can or cannot be revealed (still, it is important that the information leak is limited). Our model formalizes this intuition, and our algorithmic results shows that it is possible to exploit the model to obtain more efficient algorithms.

Specifically, within this framework, we give a $c$-approximate near neighbor protocol with communication $\tilde{O}\left(n^{1 / 2}+d\right)$ for any constant $c>1$. The protocol is based on dimensionality reduction technique of [25]. We show how the dependence on $d$ can be made polylogarithmic if Alice just wants a coordinate of a point in $P_{c r}$. We also give a protocol based on locality-sensitive hashing (LSH) [23], with communication $\tilde{O}\left(n^{1 / 2+1 /(2 c)}+d\right)$, but significantly less work (though still polynomial).

Finally, proceeding along the lines of [20], we say the protocol leaks $b$ bits of information if it can be simulated given $b$ extra bits which may depend arbitrarily on the input. With this definition, we give a protocol with $\tilde{O}\left(n^{1 / 3}+d\right)$ communication leaking only $k$ bits, where $k$ is a security parameter.

General vs specific solutions. As described above, this paper offers solutions to specific computational problems. In principle, a general "compiler-like" approach (as in [33, [18]) would be preferable. However, it appears unlikely that a compiler approach can be developed in the context of approximate problems. Indeed, there is no general method that, for a given problem, generates an efficient approximation algorithm (even ignoring the privacy issue). This implies that a compiler would have to start from a particular approximation to a given function. Unfortunately, as mentioned earlier, such approximation itself can leak too much information.

This argument leads us to believe that, in context of approximate algorithms, designing efficient private solutions to specific problems is the only possible approach.

## 2 Preliminaries

Background on homomorphic encryption, oblivious transfer (OT), and secure function evaluation (SFE) can be found in appendix A. We write negl $(k, n)$ to denote an arbitrary negligible function of $k, n$, that is a function which shrinks faster than any inverse polynomial in $n, k$.

We assume both parties are computationally bounded and semi-honest, meaning they follow the protocol but may keep message histories in an attempt to learn more than is prescribed. In [18, 7, 29], it is shown how to transform a semihonest protocol into a protocol secure in the malicious model. Further, [29] does this at a communication blowup of at most a small factor of poly $(k)$. Therefore, we assume parties are semi-honest in the remainder of the paper.

We briefly review the semi-honest model, referring the reader to [17, 26] for more details. Let $f:\{0,1\}^{*} \times\{0,1\}^{*} \rightarrow\{0,1\}^{*} \times\{0,1\}^{*}$ be a function, the first element denoted $f_{1}\left(x_{1}, x_{2}\right)$ and the second $f_{2}\left(x_{1}, x_{2}\right)$. Let $\pi$ be a two-party protocol for computing $f$. The views of players $P_{1}$ and $P_{2}$ during an execution of $\pi\left(x_{1}, x_{2}\right)$, denoted $\operatorname{View}_{1}^{\pi}\left(x_{1}, x_{2}\right)$ and $\operatorname{View}_{2}^{\pi}\left(x_{1}, x_{2}\right)$ respectively, are:
$\operatorname{View}_{1}^{\pi}\left(x_{1}, x_{2}\right)=\left(x_{1}, r_{1}, m_{1,1}, \ldots, m_{1, t}\right), \operatorname{View}_{2}^{\pi}\left(x_{1}, x_{2}\right)=\left(x_{2}, r_{2}, m_{2,1}, \ldots, m_{2, t}\right)$,
where $r_{i}$ is the random input and $m_{i, j}$ the messages received by player $i$ respectively. The outputs of $P_{1}$ and $P_{2}$ during an execution of $\pi\left(x_{1}, x_{2}\right)$ are denoted output ${ }_{1}^{\pi}\left(x_{1}, x_{2}\right)$ and output ${ }_{2}^{\pi}\left(x_{1}, x_{2}\right)$. We define output ${ }^{\pi}\left(x_{1}, x_{2}\right)$ to be (output ${ }_{1}^{\pi}\left(x_{1}, x_{2}\right)$, output ${ }_{2}^{\pi}\left(x_{1}, x_{2}\right)$ ). We say that $\pi$ privately computes a function $f$ if there exist PPT algorithms $S_{1}, S_{2}$ for which for $i \in\{1,2\}$ we have the following indistinguishability

$$
\left\{S_{i}\left(x_{i}, f_{i}\left(x_{1}, x_{2}\right)\right), f\left(x_{1}, x_{2}\right)\right\} \stackrel{\mathrm{c}}{\equiv}\left\{\operatorname{View}_{i}^{\pi}\left(x_{1}, x_{2}\right), \text { output }^{\pi}\left(x_{1}, x_{2}\right)\right\} .
$$

This simplifies to $\left\{S_{i}\left(x_{i}, f_{i}\left(x_{1}, x_{2}\right)\right)\right\} \stackrel{\text { c }}{\equiv}\left\{\operatorname{View}_{i}^{\pi}\left(x_{1}, x_{2}\right)\right\}$ if either $f_{1}\left(x_{1} x_{2}\right)=$ $f_{2}\left(x_{1}, x_{2}\right)$ or if $f\left(x_{1}, x_{2}\right)$ is deterministic or equals a specific value with probability $1-\operatorname{negl}(k, n)$, for $k$ a security parameter.

We need a standard composition theorem [17 concerning private subprotocols. An oracle-aided protocol (see [26]) is a protocol augmented with a pair of oracle tapes for each party and oracle-call steps. In an oracle-call step parties write to their oracle tape and the oracle responds to the requesting parties. An oracle-aided protocol uses the oracle-functionality $f=\left(f_{1}, f_{2}\right)$ if the oracle responds to query $x, y$ with $\left(f_{1}(x, y), f_{2}(x, y)\right)$, where $f_{1}, f_{2}$ denote first and second party's output respectively. An oracle-aided protocol privately reduces $g$ to $f$ if it privately computes $g$ when using oracle-functionality $f$.

Theorem 1. [17] If a function $g$ is privately reducible to a function $f$, then the protocol $g^{\prime}$ derived from $g$ by replacing oracle calls to $f$ with a protocol for privately computing $f$, privately computes $g$.

We now define the functional privacy of an approximation as in [12]. For our approximation protocols we will have $f_{1}(x, y)=f_{2}(x, y)=f(x, y)$.

Definition 1. Let $f(x, y)$ be a function, and let $\hat{f}(x, y)$ be a randomized function. Then $\hat{f}(x, y)$ is functionally private for $f$ if there is an efficient simulator $S$ s.t. for every $x, y$, we have $\hat{f}(x, y) \stackrel{c}{\equiv} S(f(x, y))$.

A private approximation of $f$ privately computes a randomized function $\hat{f}$ that is functionally private for $f$.

Finally, we need the notion of a protocol for securely evaluating a circuit with $R O M$. In this setting, the $i$ th party has a table $R_{i} \in\left(\{0,1\}^{r}\right)^{s}$ defined by his inputs. The circuit, in addition to the usual gates, is equipped with lookup gates which on inputs $(i, j)$, output $R_{i}[j]$.

Theorem 2. [29] If $C$ is a circuit with $R O M$, then it can be securely computed with $\tilde{O}(|C| T(r, s))$ communication, where $T(r, s)$ is the communication of 1-out-of-s OT on words of size $r$.

## 3 Private $\ell_{2}$ Approximation

Here we give a private approximation of the $\ell_{2}$ distance. Alice is given a vector $a \in[M]^{n}$, and Bob a vector $b \in[M]^{n}$. Note that $\|a-b\|^{2} \leq T_{\max } \stackrel{\text { def }}{=} n M^{2}$. In addition, parameters $\epsilon, \delta$ and $k$ are specified. For simplicity, we assume that $k=\Omega(\log (n M))$. The goal is for both parties to compute an estimate $E$ such that $\left|E-\|x\|^{2}\right| \leq \epsilon\|x\|^{2}$ with probability at least $1-\delta$, for $x \stackrel{\text { def }}{=} a-b$. Further, we want $E$ to be a private approximation of $\|x\|$, as defined in section 2, As discussed there, wlog we assume the parties are semi-honest. We set the parameter $B=\Theta(k)$; this notation means $B=c k$ for a large enough constant $c$ independent from $k, n, M, \delta, \epsilon$. In our protocol we make the following cryptographic assumptions.

1. There exists a PRG $G$ stretching $\operatorname{poly} \log (n)$ bits to $n$ bits secure against poly $(n)$-sized circuits.
2. There exists an OT scheme for communicating 1 of $n$ bits with communication polylog $(n)$.

At the end of the section we discuss the necessity and plausibility of these assumptions. Our protocol relies on the following fact and corollary.

Fact 3. 27] Let $A$ be a random $n \times n$ orthonormal matrix (i.e., $A$ is picked from a distribution defined by the Haar measure). Then there is $c>0$ such that for any $x \in \Re^{n}$, any $i=1, \ldots, n$, and any $t>1$,

$$
\operatorname{Pr}\left[\left|(A x)_{i}\right| \geq \frac{\|x\|}{\sqrt{n}} t\right] \leq e^{-c t^{2}}
$$

Corollary 1. Suppose we sample $A$ as in Fact 3 but instead generate our randomness from $G$, rounding its entries to the nearest multiple of $2^{-\Theta(B)}$. Then,

$$
\begin{aligned}
\forall & x \in[M]^{n}, \\
& \operatorname{Pr}\left[\left(1-2^{-B}\right)\|x\|^{2} \leq\|A x\|^{2} \leq\|x\|^{2} \text { and } \forall_{i}(A x)_{i}^{2}<\frac{\|x\|^{2}}{n} B\right]>1-\operatorname{neg}(k, n)
\end{aligned}
$$

Proof. If there were an infinite sequence of $x \in[M]^{n}$ for which this did not hold, a circuit with $x$ hardwired would contradict the pseudorandomness of $G$.

Protocol Overview: Before describing our protocol, it is instructive to look at some natural approaches and why they fail. We start with the easier case of approximating the Hamming distance, and suppose the parties share a common random string. Consider the following non-private protocol of [25] discussed in [12: Alice and Bob agree upon a random $O(\log n) \times n$ binary matrix $R$ where the $i$ th row consists of $n$ i.i.d. Bernoulli $\left(\beta^{i}\right)$ entries, where $\beta$ is a constant depending on $\epsilon$. Alice and Bob exchange $R a, R b$, and compute $R(a-b)=R x$. Then $\|x\|$ can be approximated by observing that $\operatorname{Pr}\left[(R a)_{i}=(R b)_{i}\right] \approx 1 / 2$ if $\|x\| \gg \beta^{-i}$, and $\operatorname{Pr}\left[(R a)_{i}=(R b)_{i}\right] \approx 1$ if $\|x\| \ll \beta^{-i}$. Let the output be $E$. The communication is $O(\log n)$, but it is not private since both parties learn $R x$. Indeed, as mentioned in [12], if $a=0$ and $b=e_{i}$, then $R x$ equals the $i$ th column of $R$, which cannot be simulated without knowing $i$.

However, given only $\|x\|$, it is possible to simulate $E$. Therefore, as pointed out in [12, one natural approach to try to achieve privacy is to run an SFE with inputs $R a, R b$, and output $E$. But this also fails, since knowing $E$ together with the randomness $R$ may reveal additional information about the inputs. If $E$ is a deterministic function of $R a, R b$, and if $a=0$ and $b=e_{i}$, Alice may be able to find $i$ from $a$ and $R$.

In [12], two private protocols which each have $\Omega(n)$ communication for a worst-case choice of inputs, were cleverly combined to overcome these problems and to achieve $\tilde{O}(\sqrt{n})$ communication. The first protocol, High-Distance Estimator, works when $\|x\|>\sqrt{n}$. The idea is for the parties to obliviously sample random coordinates of $x$, and use these to estimate $\|x\|$. Since the sampling is oblivious, the views depend only on $\|x\|$, and since it is random, the estimate is good provided we take $\tilde{O}(\sqrt{n})$ samples.

The second protocol, Low-Distance Estimator, works when $\|x\| \leq \sqrt{n}$. Roughly, the idea is for the parties to perfectly hash their vectors into $\tilde{O}(\sqrt{n})$ buckets so that at most one coordinate $j$ for which $a_{j} \neq b_{j}$ lies in any given bucket. The parties then run an SFE with their buckets as input, which can compute $\|x\|$ exactly by counting the number of buckets which differ.

Our protocol breaks this $O(\sqrt{n})$ communication barrier as follows. First, Alice and Bob agree upon a random orthonormal matrix $A$ in $\mathbb{R}^{n \times n}$, and compute $A a$ and $A b$. The point of this step is to uniformly spread the mass of the difference vector $x$ over the $n$ coordinates, as per Fact 3, while preserving the length. Since we plan to sample random coordinates of $A x$ to estimate $\|x\|$, it is crucial to spread out the mass of $\|x\|$, as otherwise we could not for instance, distinguish $x=0$ from $x=e_{i}$. The matrix multiplication can be seen as an analogue to the perfect hashing in Low-Distance Estimator, and the coordinate sampling as an analogue to that in High-Distance Estimator.

To estimate $\|x\|$ from the samples, we need to be careful of a few things. First, the parties should not learn the sampled values $(A x)_{j}$, since these can reveal too much information. Indeed, if $a=0$, then $(A x)_{j}=(A b)_{j}$, which is not private.

To this end, the parties run a secure circuit with ROM (see section 2) $A a$ and $A b$, which privately obtains the samples.

Second, we need the circuit's output distribution $E$ to depend only on $\|x\|$. It is not enough for $\mathbf{E}[E]=\|x\|^{2}$, since a polynomial number of samples from $E$ may reveal non-simulatable information about $x$ based on $E$ 's higher moments. To this end, the circuit uses the $(A x)_{j}$ to independently generate r.v.s $z_{j}$ from a Bernoulli distribution with success probability depending only on $\|x\|$. Hence, $z_{j}$ depends only on $\|x\|$.

Third, we need to ensure that the $z_{j}$ contain enough information to approximate $\|x\|$. We do this by maintaining a loop variable $T$ which at any point in time is guaranteed to be an upper bound on $\|x\|^{2}$ with overwhelming probability. Using Corollary 1 , for all $j$ it holds that $q \stackrel{\text { def }}{=} n(A x)_{j}^{2} /(T B) \leq 1$ for a parameter $B$, so we can generate the $z_{j}$ from a $\operatorname{Bernoulli}(q)$ distribution. Since $T$ is halved in each iteration, for some iteration $\mathbf{E}\left[\sum_{j} z_{j}\right]$ will be large enough to ensure that $E$ is tightly concentrated.

We now describe the protocol in detail. Set $\ell=\Theta(B)\left(1 / \epsilon^{2} \log (n M) \log (1 / \delta)+\right.$ $k)$. In the following, if $q>1$, then the distribution $\operatorname{Bernoulli}(q)$ means Bernoulli(1).

## $\ell_{2}$-Approx $(a, b)$ :

1. Alice, Bob exchange a seed of $G$ and generate $A$ as in Corollary 1
2. Set $T=T_{\text {max }}$
3. Repeat:
(a) $\left\{\right.$ Assertion: $\left.\|x\|^{2} \leq T\right\}$
(b) A secure circuit with ROM $A a, A b$ computes the following

- Generate random $i_{1}, \ldots, i_{\ell}$ and compute $(A x)_{i_{1}}^{2}, \ldots(A x)_{i_{\ell}}^{2}$
- Generate $\left\{z_{j}\right\}_{j \in[\ell]}$ from i.i.d. Bernoulli $\left(\frac{n(A x)_{i_{j}}^{2}}{T B}\right)$ distributions
(c) $\mathrm{T}=\mathrm{T} / 2$

4. Until $\sum_{i} z_{i} \geq \frac{\ell}{4 B}$ or $T<1$
5. Output $E=\frac{2 T B}{l} \sum_{i} z_{i}$ as an estimate of $\|x\|^{2}$

Note that the protocol can be implemented in $O(1)$ rounds by parallelizing the secure circuit invocations.

Lemma 1. The probability that assertion 3a holds in every iteration of step 3 is $1-\operatorname{neg}(k, n)$. Moreover, if $\|x\|^{2} \neq 0$, then when the algorithm exits, with probability $1-\operatorname{neg}(k, n)$ it holds that $\mathbf{E}\left[\sum_{j} z_{j}\right]=\Theta(\ell / B)$.

Proof. By Corollary [1 $\operatorname{Pr}_{A}\left[\left(1-2^{-B}\right)\|x\|^{2} \leq\|A x\|^{2} \leq\|x\|^{2}\right.$ and $\forall_{i}(A x)_{i}^{2}<$ $\left.\frac{\|x\|^{2}}{n} B\right]=1-\operatorname{neg}(k, n)$, so we may condition on this event occurring. If $\|x\|^{2}=0$, then $A x=0$, and thus $\operatorname{Pr}[E=0]=1$.

Otherwise, $\|x\|^{2} \geq 1$. Consider the smallest $j$ for which $T_{\max } / 2^{j}<\|x\|^{2}$. We show for $T=T_{\max } / 2^{j-1} \geq\|x\|^{2} \geq 1$ that $\operatorname{Pr}\left[\sum_{j} z_{j}<\ell /(4 B)\right]=\operatorname{neg}(k, n)$. The assertion holds at the beginning of the $j$ th iteration by our choice of $T$. Thus,
$n(A x)_{i}^{2} \leq T B$ for all $i \in[n]$ by the properties of $A$. So for all $j, \operatorname{Pr}\left[z_{j}=1\right]=$ $\frac{\|A x\|^{2}}{T B} \geq\left(1-2^{-B}\right) /(2 B)$, and thus $\mathbf{E}\left[\sum_{j} z_{j}\right] \geq \ell /(3 B)$. By a Chernoff bound, $\operatorname{Pr}\left[\sum_{j} z_{j}<\ell /(4 B)\right]=\operatorname{neg}(k, n)$, so if ever $T=T_{\max } / 2^{j-1}$, then this is the last iteration with overwhelming probability.

Note that the second part of the lemma follows from standard Chernoff bounds. Indeed, if $\|x\|^{2} \neq 0$, then we have shown with overwhelming probability that in some iteration, $T \geq 1$ and $\sum_{i} z_{i} \geq \ell / 4 B$, so we may condition on the event that the algorithm exits in such an iteration. But for a certain constant in the big-Oh notation, one can show (by Chernoff and union bounds) that the probability $\sum_{i} z_{i} \geq \ell / 4 B$ when $\mathbf{E}\left[\sum_{j} z_{j}\right]=O(\ell / B)$ is negligible. On the other hand, once $\mathbf{E}\left[\sum_{j} z_{j}\right] \geq \ell / 3 B$, we have shown that $\sum_{i} z_{i} \geq \ell / 4 B$ with overwhelming probability. Thus we exit with $\mathbf{E}\left[\sum_{j} z_{j}\right]=\Theta(\ell / B)$.

Correctness. We claim $\operatorname{Pr}\left[\left|E-\|x\|^{2}\right| \leq \epsilon\right] \geq 1-\delta$. Since $A x=0$ if $\|x\|^{2}=$ 0 , we have that $E=0$ in this case, and the claim is immediate. So suppose $\|x\|^{2} \neq 0$. By Lemma when the algorithm exits, with probability $1-\operatorname{neg}(k, n)$, $\mathbf{E}\left[\sum_{i} z_{i}\right]=\Theta(\ell / B)$, so we assume this event occurs. By a Chernoff and a union bound over iterations, we may assume that whenever $\mathbf{E}\left[\sum_{i} z_{i}\right]=\Theta(\ell / B)$,

$$
\operatorname{Pr}\left[\left|\sum_{i} z_{i}-E\left[\sum_{i} z_{i}\right]\right| \geq \frac{\epsilon}{2} E\left[\sum_{i} z_{i}\right]\right] \leq e^{-\Theta\left(\epsilon^{2} \frac{\ell}{B}\right)}<\frac{\delta}{2}
$$

Thus, this holds when the algorithm exits. By Lemma 1 assertion 3a holds, so that $\ell\left(1-2^{-B}\right)\|x\|^{2} \leq T B \cdot \mathbf{E}\left[\sum_{i} z_{i}\right] \leq \ell\|x\|^{2}$. Setting $E=\frac{2 T B}{\ell} \sum_{i} z_{i}$ (recall that $T$ is halved in step 3 c ) then shows that $\operatorname{Pr}\left[\left|E-\|x\|^{2}\right| \geq \epsilon\|x\|^{2}\right] \leq \delta$.

Privacy. We replace the secure circuit with ROM in step 3b of $\ell_{2}$-Approx with an oracle (see section(2). We construct a single simulator Sim, which given $\Delta \stackrel{\text { def }}{=}\|x\|^{2}$, satisfies $\operatorname{Sim}(\Delta) \stackrel{\mathrm{c}}{\equiv} \operatorname{View}_{\mathrm{A}}^{\pi}(\mathrm{a}, \mathrm{b})$ and $\operatorname{Sim}(\Delta) \stackrel{\mathrm{c}}{\equiv} \operatorname{View}_{\mathrm{B}}^{\pi}(\mathrm{a}, \mathrm{b})$, where $\operatorname{View}_{A}^{\pi}(\mathrm{a}, \mathrm{b})$, $\operatorname{View}_{B}^{\pi}(\mathrm{a}, \mathrm{b})$ are Alice, Bob's real views respectively. This, in particular, implies functional privacy. It will follow that $\ell_{2}$-Approx is a private approximation of $\Delta$.
$\underline{\operatorname{Sim}(\Delta): ~}$

1. Generate a random seed of $G$
2. Set $T=T_{\max }$
3. Repeat:
(a) Generate $\left\{z_{j}\right\}_{j \in[\ell]}$ from i.i.d. Bernoulli $\left(\frac{\Delta}{T B}\right)$ distributions
(b) $T=T / 2$
4. Until $\sum_{i} z_{i} \geq \frac{\ell}{4 B}$ or $T<1$
5. Output $E=\frac{2 T B}{l} \sum_{i} z_{i}$

With probability $1-\operatorname{neg}(k, n)$, the matrix $A$ satisfies the property in Corollary 1 so we assume this event occurs. In each iteration, the random variables $z_{j}$ are
independent in both the simulation and the protocol. Further, the probabilities that $z_{j}=1$ in the simulated and real views differ only by a multiplicative factor of $\left(1-2^{-B}\right)$ as long as $T \geq \Delta$. But the probability that, in either view, we encounter $T<\Delta$ is $\operatorname{neg}(k, n)$.
Complexity. Given our cryptographic assumptions, we use $\tilde{O}(1)$ communication and $O(1)$ rounds.

Remark 1. Our cryptographic assumptions are fairly standard, and similar to the ones in [12]. There the authors make the weaker assumptions that PRGs stretching $n^{\gamma}$ bits to $n$ bits and OT with $n^{\gamma}$ communication exist for any constant $\gamma$. In fact, the latter implies the former [21, 15]. If we were to instead use these assumptions, our communication would be $O\left(n^{\gamma}\right)$, still greatly improving upon the $O\left(n^{1 / 2+\gamma}\right)$ communication of [12]. A candidate OT scheme satisfying our assumptions can be based on the $\Phi$-Hiding Assumption [6], and can be derived by applying the PIR to OT transformation of 30 to the scheme in that paper.

Remark 2. For the special case of Hamming distance, we have an alternative protocol based on the following idea. Roughly, both parties apply the perfect hashing of the Low-Distance Estimator protocol of 12 for a logarithmic number of levels $j$, where the $j$ th level contains $\tilde{O}\left(2^{j}\right)$ buckets. To overcome the $\tilde{O}(\sqrt{n})$ barrier of [12], instead of exchanging the buckets, the set of buckets is randomly and obliviously sampled. From the samples, an estimate of $\Delta(a, b)$ is output. For some $j, 2^{j} \approx \Delta(a, b)$, so the estimate will be tightly concentrated, and for reasons similar to $\ell_{2}$-Approx, will be simulatable. We omit the details, but note that two advantages of this alternative protocol are that the time complexity will be $\tilde{O}(n)$ instead of $\tilde{O}\left(n^{2}\right)$, and that we don't need the PRG $G$, as we may use $k$-wise independence for the hashing.

## 4 Private Near Neighbor and c-Approximate Near Neighbor Problems

We consider the case in which Alice has a point $q$, and Bob a set of $n$ points $P$.

### 4.1 Private Near Neighbor Problem

Suppose for some integer $U$, Alice has $q \in[U]^{d}$, Bob has $P=p_{1}, \ldots, p_{n} \in[U]^{d}$, and Alice should learn $\min _{i} f\left(q, p_{i}\right)$, where $f$ is some distance function. In [10] protocols for $\ell_{1}, \ell_{2}$, Hamming distance over $U$-ary alphabets, set difference, and arbitrary distance functions $f(a, b)=\sum_{i=1}^{d} f_{i}\left(a_{i}, b_{i}\right)$ were proposed, using an untrusted third party. We improve the communication of these protocols and remove the third party using homomorphic encryption to implement polynomial evaluation as in [13], and various hashing tricks.

In [13], the authors consider the private all-near neighbors problem in which Alice has $n$ queries $q_{1}, \ldots, q_{n} \in[U]^{d}$ and wants all $p_{i}$ for which $\Delta\left(p_{i}, q_{j}\right) \leq t<d$ for some $j$ and parameter $t$. Our techniques improve the $\tilde{O}\left(n^{2} d\right)$ communication of a generic SFE and the $\tilde{O}\left(n\binom{d}{t}\right)$ communication of [13] for this problem to
$\tilde{O}\left(n d^{2}+n^{2}\right)$. Finally, in the common random string model we achieve $\lceil\log d\rceil+$ $O(k)$ communication for the (exact) Hamming distance, and an inner product protocol with $d+O(k)$ communication.

For the details of our schemes, see the full version of our paper [24]. We do not focus on them here since they still suffer from an $\Omega(n)$ communication cost. We instead focus on how to privately approximate these problems.

### 4.2 Private $\boldsymbol{c}$-Approximate Near Neighbor Problem

Suppose $q \in\{0,1\}^{d}$ and $p_{i} \in\{0,1\}^{d}$ for all $i$. Let $P_{t}=\{p \in P \mid \Delta(p, q) \leq t\}$, and $c>1$ be a constant.
Definition 2. A c-approximate NN protocol is correct if when $P_{r} \neq \emptyset$, Alice outputs a point $f(q, P) \in P_{c r}$ with probability $1-2^{-\Omega(k)}$. It is private if in the computational sense, Bob learns nothing, while Alice learns nothing except what follows from $P_{c r}$. Formally, Alice's privacy is implied by an efficient simulator Sim for which $\langle q, P, f(q, P)\rangle \stackrel{\text { c }}{\equiv}\left\langle q, P, \operatorname{Sim}\left(1^{n}, P_{c r}, q\right)\right\rangle$ for $\operatorname{poly}(d, n, k)$-time machines. Following [20], we say the protocol leaks $b$ bits of information if there is a deterministic "hint" function $h:\{0,1\}^{(n+1) d} \rightarrow\{0,1\}^{b}$ such that the distributions $\langle q, P, f(q, P)\rangle$ and $\left\langle q, P, \operatorname{Sim}\left(1^{n}, P_{c r}, q, h(P, q)\right)\right\rangle$ are indistinguishable. As motivated in section 1] we believe these to be natural extensions of private approximations in [12, 20] from values to sets of values.

We give a private $c$-approximate NN protocol with communication $\tilde{O}(\sqrt{n}+d)$ and a $c$-approximate NN protocol with communication $\tilde{O}\left(n^{1 / 3}+d\right)$ which leaks $k$ bits of information. Both protocols are based on dimensionality reduction in the hypercube [25]. There it is shown that for an $O(\log n) \times d$ matrix $A$ with entries i.i.d. Bernoulli $(1 / d)$, there is an $\tau=\tau(r, c r)$ such that for all $p, q \in\{0,1\}^{d}$, the following event holds with probability at least $1-1 / \operatorname{poly}(n)$
If $\Delta(p, q) \leq r$, then $\Delta(A p, A q) \leq \tau$, and if $\Delta(p, q) \geq c r$, then $\Delta(A p, A q)>\tau$.
Here, arithmetic occurs in $\mathbb{Z}_{2}$. We use this idea in the following helper protocol $\operatorname{DimReduce}(\tau, B, q, P)$. Let $A$ be a random matrix as described above. Let $S=$ $\{p \in P \mid \Delta(A p, A q) \leq \tau\}$. If $|S|>B$, replace $S$ with the lexicographically first $B$ elements of $S$. DimReduce outputs random shares of $S$.

DimReduce $(\tau, B, q, P)$ :

1. Bob performs the following computation

- Generate a matrix $A$ as above, and initialize $L$ to an empty list.
- For each $v \in\{0,1\}^{O(\log n)}$, let $L(v)$ be the first $B p_{i}$ for which $\Delta\left(A p_{i}, v\right) \leq \tau$.

2. A secure circuit with ROM $L$ and input $(q, A)$ executes:

- Compute Aq.
- Lookup $A q$ in $L$ to obtain $S$. If $|S|<B$, pad $S$ so that all $S$ have the same length.
- Output random shares $\left(S^{1}, S^{2}\right)$ of $S$ so that $S=S^{1} \oplus S^{2}$.

It is an easy exercise to show the correctness and privacy of DimReduce.
Remark 3. As stated, the communication is $\tilde{O}(d B)$. The dependence on $d$ can be improved to $\tilde{O}(d+B)$ using homomorphic encryption. Roughly, Alice sends $E\left(q_{1}\right), \ldots, E\left(q_{d}\right)$ to Bob, who sets $L(v)$ to be the first $B$ different $E\left(\Delta\left(p_{i}, q\right)\right)$ for which $\Delta\left(A p_{i}, v\right) \leq \tau$. Note that $E\left(\Delta\left(p_{i}, q\right)\right)$ is efficiently computable, and has size $\tilde{O}(1) \ll d$.
It will be useful to define the following event $\mathcal{H}\left(r_{1}, r_{2}, P\right)$ with $r_{1}<r_{2}$. Suppose we run DimReduce independently $k$ times with matrices $A_{i}$. Then $\mathcal{H}\left(r_{1}, r_{2}, P\right)$ is the event that at least $k / 2$ different $i$ satisfy

$$
\forall p \in P_{r_{1}}, \Delta\left(A_{i} p, A_{i} q\right) \leq \tau\left(r_{1}, r_{2}\right) \text { and } \forall p \in P \backslash P_{r_{2}}, \Delta\left(A_{i} p, A_{i} q\right)>\tau\left(r_{1}, r_{2}\right)
$$

The next lemma follows from the properties of the $A_{i}$ and Chernoff bounds:
Lemma 2. $\operatorname{Pr}\left[\mathcal{H}\left(r_{1}, r_{2}, P\right)\right]=1-2^{-\Omega(k)}$.

## 4.3 c-Approximate NN Protocol

Protocol Overview: Our protocol is based on the following intuition. When $\left|P_{c r}\right|$ is large, a simple solution is to run a secure function evaluation with Alice's point $q$ as input, together with a random sample $P^{\prime}$ of roughly a $k /\left|P_{c r}\right|$ fraction of Bob's points $P$. The circuit returns a random point of $P^{\prime} \cap P_{c r}$, which is non-empty with overwhelming probability. The communication is $\tilde{O}\left(n /\left|P_{c r}\right|\right)$.

On the other hand, when $\left|P_{c r}\right|$ is small, if for $k$ independent trials Alice and Bob run DimReduce $\left(\tau(r, c r),\left|P_{c r}\right|, q, P\right)$, then with overwhelming probability $P_{r} \subseteq \cup_{i} S_{i}$, where $S_{i}$ denotes the (randomly shared) output in the $i$ th execution. A secure function evaluation can then take in the random shares of the $S_{i}$ and output a random point of $P_{r}$. The communication of this scheme is $\tilde{O}\left(\left|P_{c r}\right|\right)$.

Our protocol combines these two protocols to achieve $\tilde{O}(\sqrt{n})$ communication, by sampling roughly an $n^{-1 / 2}$ fraction of Bob's points in the first protocol, and by invoking DimReduce with parameter $B=\tilde{O}(\sqrt{n})$ in the second protocol. This approach is similar in spirit to the "high distance / low distance" approach used to privately approximate the Hamming distance in 12 .
$c$-Approx (q, P):

1. Set $B=\tilde{O}(\sqrt{n})$.
2. Independently run $\operatorname{DimReduce}(\tau(r, c r), B, q, P) k$ times, generating shares $\left(S_{i}^{1}, S_{i}^{2}\right)$.
3. Bob finds a random subset $P^{\prime}$ of $P$ of size $B$.
4. On inputs $q, S_{i}^{1}, S_{i}^{2}, P^{\prime}$, a secure circuit executes:

- Compute $S_{i}=S_{i}^{1} \oplus S_{i}^{2}$ for all $i$.
- Let $f(q, P)$ be a random point from $P_{c r} \cap P^{\prime} \neq \emptyset$ if it is non-empty,
- Else let $f(q, P)$ be a random point from $P_{r} \cap \cup_{i} S_{i}$ if it is non-empty, otherwise set $f(q, P)=\emptyset$.
- Output ( $f(q, P)$, null).

Using the ideas in Remark 3, the communication is $\tilde{O}(d+B)$, since the SFE has size $\tilde{O}(B)$. Let $\mathcal{F}$ be the event that $P^{\prime} \cap P_{c r} \neq \emptyset$, and put $\mathcal{H}=\mathcal{H}(r, c r, P)$.

Correctness. Suppose $P_{r}$ is nonempty. The probability $s$ of correctness is just the probability we don't output $\emptyset$. Thus $s \geq \operatorname{Pr}[\mathcal{F}]+\operatorname{Pr}[\neg \mathcal{F}] \operatorname{Pr}[f(q, P) \neq \emptyset \mid \neg \mathcal{F}]$.
Case $\left|P_{c r}\right| \geq \sqrt{n}$ : For sufficiently large $B$, we have $s \geq \operatorname{Pr}[\mathcal{F}]=1-2^{-\Omega(k)}$.
Case $\left|P_{c r}\right|<\sqrt{n}$ : It suffices to show $\operatorname{Pr}[f(q, P) \neq \emptyset \mid \neg \mathcal{F}]=1-2^{-\Omega(k)}$. But this probability is at least $\operatorname{Pr}[f(q, P) \neq \emptyset \mid \mathcal{H}, \neg \mathcal{F}] \operatorname{Pr}[\mathcal{H}]$, and if $\mathcal{H}$ occurs, then $f(q, P) \neq \emptyset$. By Lemma $2 \operatorname{Pr}[\mathcal{H}]=1-2^{-\Omega(k)}$.

Privacy. Note that Bob gets no output, so Alice's privacy follows from the composition of of DimReduce and the secure circuit protocol of step 5. Similarly, if we can construct a simulator $\operatorname{Sim}$ with inputs $1^{n}, P_{c r}, q$ so that the distributions $\langle q, P, f(q, P)\rangle$ and $\left\langle q, P, \operatorname{Sim}\left(1^{n}, P_{c r}, q\right)\right\rangle$ are statistically close, Bob's privacy will follow by that of DimReduce and the secure circuit protocol of step 5.

## $\underline{\operatorname{Sim}\left(1^{n}, P_{c r}, q\right)}:$

1. Set $B=\tilde{O}\left(n^{1 / 2}\right)$.
2. With probability $1-\binom{n-\left|P_{c r}\right|}{B}\binom{n}{B}^{-1}$, output a random element of $P_{c r}$,
3. Else output a random element of $P_{r}$.

Let $X$ denote the output of $\operatorname{Sim}\left(1^{n}, P_{c r}, q\right)$. It suffices to show that for each $p \in$ $P,|\operatorname{Pr}[f(q, P)=p]-\operatorname{Pr}[X=p]|=2^{-\Omega(k)}$, since this also implies $\mid \operatorname{Pr}[f(q, P)=$ $\emptyset]-\operatorname{Pr}[X=\emptyset] \mid=2^{-\Omega(k)}$. We have

$$
\begin{aligned}
\operatorname{Pr}[f(q, P)=p] & =\operatorname{Pr}[f(q, P)=p, \mathcal{F}]+\operatorname{Pr}[f(q, P)=p, \neg \mathcal{F}] \\
& =\operatorname{Pr}[f(q, P)=p, \mathcal{F}]+\operatorname{Pr}[f(q, P)=p, \neg \mathcal{F} \mid \mathcal{H}] \pm 2^{-\Omega(k)} \\
& =\operatorname{Pr}[\mathcal{F}]\left|P_{c r}\right|^{-1}+\operatorname{Pr}[\neg \mathcal{F}] \operatorname{Pr}[f(q, P)=p \mid \mathcal{H}, \neg \mathcal{F}] \pm 2^{-\Omega(k)}
\end{aligned}
$$

where we have used Lemma 2 Since $\operatorname{Pr}[\mathcal{F}]=1-\binom{n-\left|P_{c r}\right|}{B}\binom{n}{B}^{-1}$, we have

$$
\begin{gathered}
|\operatorname{Pr}[f(q, P)=p]-\operatorname{Pr}[X=p]| \leq \\
\left.\operatorname{Pr}[\neg \mathcal{F}]\left|\operatorname{Pr}[f(q, P)=p \mid \mathcal{H}, \neg \mathcal{F}]-\delta\left(p \in P_{r}\right)\right| P_{r}\right|^{-1} \mid+2^{-\Omega(k)}
\end{gathered}
$$

If $\left|P_{c r}\right| \geq \sqrt{n}$, then $\operatorname{Pr}[\neg \mathcal{F}]=2^{-\Omega(k)}$. If $\left|P_{c r}\right|<\sqrt{n}$, then $\operatorname{Pr}[f(q, P)=p \mid$ $\mathcal{H}, \neg \mathcal{F}]=\delta\left(p \in P_{r}\right)\left|P_{r}\right|^{-1}$.

Reducing the dependence on $d$ : The way the current problem is stated, there is an $\Omega(d)$ lower bound. We now sketch how, if Alice just wants to learn some coordinate of an element of $P_{c r}$, this dependence can be made polylogarithmic. The idea is to perform an approximation to the Hamming distance instead of using the $E\left(\Delta\left(p_{i}, q\right)\right)$ in the current protocol (see, e.g., DimReduce, and the following remark). The approximation we use is that given in [25], namely, the
parties will agree upon random matrices $A_{i}$ for some subset of $i$ in $[n]$, and from the $A_{i} p_{i}$ and $A_{i} q$ will determine $(1 \pm \epsilon)$ approximations to the $\Delta\left(p_{i}, q\right)$ with probability $1-2^{-k}$. We don't need private approximations since the parties will not learn these values, but rather, they will input the $A_{i} p_{i}, A_{i} q$ into a secure circuit which makes decisions based on these approximations.

More precisely, Bob samples $B$ of his vectors $p_{i}$, and in parallel agrees upon $B$ matrices $A_{i}$ and feeds the $A_{i} p_{i}$ into a secure circuit. Alice feeds in the $A_{i} q$. Let $c \geq 1+8 \epsilon$. The circuit looks for an approximation of at most $r(1+6 \epsilon)$. If such a value exists, the circuit gives Alice the corresponding index. Observe that if $\left|P_{r(1+4 \epsilon)}\right|>\sqrt{n}$, then with probability $1-2^{-k}$ an index is returned to an element in $P_{c r}$, and that this distribution is simulatable. So assume $\left|P_{r(1+4 \epsilon)}\right| \leq \sqrt{n}$.

The parties proceed by running a variant of $\operatorname{DimReduce}(\tau(r, r(1+4 \epsilon)), B, q, P)$, with the important difference being that the output no longer consists of shares of the $E\left(\Delta\left(p_{i}, q\right)\right)$. Instead, for each entry $L(v)$, Bob pretends he is running the approximation of [25] with Alice's point $q$. That is, the parties agree on $B$ different matrices $A_{i}$ and Bob computes $A_{i} p$ for each $p \in L(v)$. A secure circuit obtains these products, and computes the approximations. It outputs an index to a random element with approximation at most $r(1+2 \epsilon)$. If $P_{r}$ is nonempty, such an index will exist with probability $1-2^{-k}$. Also, the probability that an index to an element outside of $P_{r(1+4 \epsilon)}$ is returned is less than $2^{-k}$, and so the distribution of the index returned is simulatable.

Finally, given the index of some element in $P_{c r}$, the parties perform OT and Alice obtains the desired coordinate, The communication is $\tilde{O}(\sqrt{n}+\operatorname{polylog}(d))$.

Locality-sensitive hashing (LSH): We also have a similar protocol based on LSH, which only achieves $\tilde{O}\left(n^{1 / 2+1 /(2 c)}+d\right)$ communication, but has much smaller time complexity (though still polynomial). More precisely, the work of the LSH scheme is $n^{O(1)}$, whereas the work of $c$-Approx is $n^{O\left(1 /(c-1)^{2}\right)}$, which is polynomial only for constant $c$. See Appendix B for the details.

## $4.4 \quad c$-Approximate NN Protocol Leaking $k$ Bits

Protocol Overview: We consider three balls $P_{r} \subseteq P_{b r} \subseteq P_{c r}$, where $c-b, b-1 \in$ $\Theta(1)$. We start by trying to use dimensionality reduction to separate $P_{r}$ from $P \backslash P_{b r}$, and to output a random point of $P_{r}$. If this fails, we try to sample and output a random point of $P_{c r}$. If this also fails, then it will likely hold that $n^{1 / 3} \leq\left|P_{b r}\right| \leq\left|P_{c r}\right| \leq n^{2 / 3}$. We then sample down the pointset $P$ by a factor of $n^{-1 / 3}$, obtaining $\tilde{P}$ with survivors $\tilde{P}_{b r}, \tilde{P}_{c r}$ of $P_{b r}, P_{c r}$ respectively. It will now likely hold that we can use dimensionality reduction to separate $\tilde{P}_{b r}$ from $\tilde{P} \backslash \tilde{P}_{c r}$ to obtain and output a random point of $\tilde{P}_{b r}$. The hint function will encode the probability, to the nearest multiple of $2^{-k}$, that the first dimensionality reduction fails, which may be a non-negligible function of $P \backslash P_{c r}$. This hint will be enough to simulate the entire protocol.

The protocol can be implemented in polynomial time with communication $\tilde{O}(B+d)=\tilde{O}\left(n^{1 / 3}+d\right)$.

## $c$-ApproxWithHelp (q, P):

1. Set $B=\tilde{O}\left(n^{1 / 3}\right)$.
2. Independently run $\operatorname{DimReduce}(\tau(r, b r), B, q, P) k$ times, generating shares $\left(S_{i}^{1}, S_{i}^{2}\right)$.
3. Bob finds random subsets $P^{\prime}, \tilde{P}$ of $P$ of respective sizes $B$ and $n^{2 / 3}$.
4. Independently run $\operatorname{DimReduce}(\tau(b r, c r), B, q, \tilde{P}) k$ times, generating shares $\left(\tilde{S}_{i}^{1}, \tilde{S}_{i}^{2}\right)$.
5. On inputs $q, S_{i}^{1}, S_{i}^{2}, P^{\prime}, \tilde{S}_{i}^{1}, \tilde{S}_{i}^{2}$, a secure circuit executes:

- Compute $S_{i}=S_{i}^{1} \oplus S_{i}^{2}$ and $\tilde{S}_{i}=\tilde{S}_{i}^{1} \oplus \tilde{S}_{i}^{2}$ for all $i$.
- If for most $i,\left|S_{i}\right|<B$, let $f(q, P)$ be a random point in $P_{r} \cap \cup_{i} S_{i}$, or set it to $\emptyset$ if it is empty.
- Else if $P_{c r} \cap P^{\prime} \neq \emptyset$, let $f(q, P)$ be a random point in $P_{c r} \cap P^{\prime}$.
- Else let $f(q, P)$ be a random point in $P_{b r} \cap \cup_{i} \tilde{S}_{i}$ if it is non-empty, otherwise set $f(q, P)=\emptyset$.
- Output ( $f(q, P)$, null).

To prove correctness and privacy, we introduce some notation. Let $\mathcal{E}_{1}$ be the event that the majority of the $\left|S_{i}\right|$ are less than $B$, and $\mathcal{E}_{2}$ the event that $P_{r} \subseteq \cup_{i} S_{i}$. Let $\mathcal{F}$ be the event that $P^{\prime} \cap P_{c r} \neq \emptyset$. Let $\mathcal{G}_{1}$ be the event that $1 \leq \tilde{P}_{b r} \leq \tilde{P}_{c r} \leq B$ and $\mathcal{G}_{2}$ the event that $\tilde{P}_{b r} \subseteq \cup_{i} \tilde{S}_{i}$. Finally, let $\mathcal{H}_{1}=\mathcal{H}(r, b r, P)$ and $\mathcal{H}_{2}=\mathcal{H}(b r, c r, \tilde{P})$. Note that $\operatorname{Pr}\left[\mathcal{H}_{1}\right], \operatorname{Pr}\left[\mathcal{H}_{2}\right]$ are $1-2^{-\Omega(k)}$ by Lemma 2. We need two lemmas:

Lemma 3. $\operatorname{Pr}\left[\mathcal{E}_{2} \mid \mathcal{E}_{1}\right]=1-2^{-\Omega(k)}$.
Proof. If $\mathcal{H}_{1}$ and $\mathcal{E}_{1}$ occur, then there is an $i$ for which $P_{r} \subseteq S_{i}$, so $\mathcal{E}_{2}$ occurs.
Lemma 4. $\operatorname{Pr}\left[\mathcal{G}_{2} \mid \mathcal{G}_{1}\right]=1-2^{-\Omega(k)}$.
Proof. If $\mathcal{H}_{2}$ and $\mathcal{E}_{2}$ occur, then the majority of the $\tilde{S}_{i}$ contain $\tilde{P}_{b r}$, so $\mathcal{G}_{2}$ occurs.
Correctness. We may assume $P_{r} \neq \emptyset$. The probability $s$ of correctness is just the probability the algorithm doesn't return $\emptyset$. Since $\mathcal{F}, \mathcal{E}_{1}$, and $\mathcal{G}_{1}$ are independent,

$$
s \geq \operatorname{Pr}\left[\mathcal{E}_{1}\right] \operatorname{Pr}\left[\mathcal{E}_{2} \mid \mathcal{E}_{1}\right]+\operatorname{Pr}\left[\neg \mathcal{E}_{1}\right]\left(\operatorname{Pr}[\mathcal{F}]+\operatorname{Pr}[\neg \mathcal{F}] \operatorname{Pr}\left[\mathcal{G}_{1}\right] \operatorname{Pr}\left[\mathcal{G}_{2} \mid \mathcal{G}_{1}\right]\right)
$$

Case $\left|P_{b r}\right|<B: \mathcal{H}_{1}$ implies $\mathcal{E}_{1}$ since $\left|P_{b r}\right|<B$, and using Lemma 3, $s \geq$ $\operatorname{Pr}\left[\mathcal{E}_{1}\right] \operatorname{Pr}\left[\mathcal{E}_{2} \mid \mathcal{E}_{1}\right]=1-2^{-\Omega(k)}$.
Case $\left|P_{b r}\right| \geq B$ : Since $\operatorname{Pr}\left[\mathcal{E}_{2} \mid \mathcal{E}_{1}\right]=1-2^{-\Omega(k)}$ by Lemma 3 we just need to show that $\operatorname{Pr}[\mathcal{F}]+\operatorname{Pr}[\neg \mathcal{F}] \operatorname{Pr}\left[\mathcal{G}_{1}\right] \operatorname{Pr}\left[\mathcal{G}_{2} \mid \mathcal{G}_{1}\right]=1-2^{-\Omega(k)}$. If $\left|P_{c r}\right|>n^{2 / 3}$, it suffices to show $\operatorname{Pr}[\mathcal{F}]=1-2^{-\Omega(k)}$. This holds for large enough $B=$ $\tilde{O}\left(n^{1 / 3}\right)$. Otherwise, if $\left|P_{c r}\right| \leq n^{2 / 3}$, then it suffices to show $\operatorname{Pr}\left[\mathcal{G}_{1}\right] \operatorname{Pr}\left[\mathcal{G}_{2} \mid \mathcal{G}_{1}\right]=$ $1-2^{-\Omega(k)}$. By assumption, $B \leq\left|P_{b r}\right| \leq\left|P_{c r}\right| \leq n^{2 / 3}$. Therefore, for large
enough $B, \operatorname{Pr}\left[\mathcal{G}_{1}\right]=1-2^{-\Omega(k)}$, and thus by Lemma [4, $\operatorname{Pr}\left[\mathcal{G}_{1}\right] \operatorname{Pr}\left[\mathcal{G}_{2} \mid \mathcal{G}_{1}\right]=$ $1-2^{-\Omega(k)}$.

Privacy. Note that Bob gets no output, so Alice's privacy follows from the composition of DimReduce and the secure circuit protocol of step 5. Similarly, if we can construct a simulator Sim with inputs $1^{n}, P_{c r}, q, h\left(P_{c r}, q\right)$ so that the distributions $\langle q, P, f(q, P)\rangle$ and $\left\langle q, P, \operatorname{Sim}\left(1^{n}, P_{c r}, q, h\left(P_{c r}, q\right)\right)\right\rangle$ are statistically close, Bob's privacy will follow by that of DimReduce and the secure circuit of step 5.

We define the hint function $h\left(P_{c r}, q\right)$ to output the nearest multiple of $2^{-k}$ to $\operatorname{Pr}\left[\mathcal{E}_{1}\right]$. In the analysis we may assume that $\operatorname{Sim}$ knows $\operatorname{Pr}\left[\mathcal{E}_{1}\right]$ exactly, since its output distribution in this case will be statistically close to its real output distribution.

## $\underline{\operatorname{Sim}\left(1^{n}, P_{c r}, q, \operatorname{Pr}\left[\mathcal{E}_{1}\right]\right):}$

1. Set $B=\tilde{O}\left(n^{1 / 3}\right)$.
2. With probabiity $\operatorname{Pr}\left[\mathcal{E}_{1}\right]$, output a random element of $P_{r}$, or $\emptyset$ if $P_{r}=\emptyset$.
3. Else with probability $1-\binom{n-\left|P_{c r r}\right|}{B}\binom{n}{B}^{-1}$, output a random element of $P_{c r}$,
4. Else output a random element of $P_{b r}$.

Let $X$ denote the output of $\operatorname{Sim}\left(1^{n}, P_{c r}, q, \operatorname{Pr}\left[\mathcal{E}_{1}\right]\right)$. It suffices to show that for each $p \in P$,

$$
|\operatorname{Pr}[f(q, P)=p]-\operatorname{Pr}[X=p]|=2^{-\Omega(k)}
$$

since then we have $|\operatorname{Pr}[f(q, P)=\emptyset]-\operatorname{Pr}[X=\emptyset]|=2^{-\Omega(k)}$. Using the independence of $\mathcal{F}, \mathcal{E}_{1}, \mathcal{G}_{1}$, and Lemmas 3, 4, we bound $\operatorname{Pr}[f(q, P)=p]$ as follows

$$
\begin{aligned}
& \operatorname{Pr}[f(q, P)=p]=\operatorname{Pr}\left[\mathcal{E}_{1}, f(q, P)=p\right]+\operatorname{Pr}\left[\neg \mathcal{E}_{1}, f(q, P)=p\right] \\
= & \operatorname{Pr}\left[\mathcal{E}_{1}\right] \operatorname{Pr}\left[f(q, P)=p \mid \mathcal{E}_{2} \mathcal{E}_{1}\right] \pm 2^{-\Omega(k)}+\operatorname{Pr}\left[\neg \mathcal{E}_{1}\right] \operatorname{Pr}[\mathcal{F}] \operatorname{Pr}\left[f(q, P)=p \mid F, \neg \mathcal{E}_{1}\right] \\
+ & \operatorname{Pr}\left[\neg \mathcal{E}_{1}\right] \operatorname{Pr}[\neg \mathcal{F}] \operatorname{Pr}\left[f(q, P)=p \mid \neg F, \neg \mathcal{E}_{1}\right] \\
= & \operatorname{Pr}\left[\mathcal{E}_{1}\right]\left|P_{r}\right|^{-1} \delta\left(p \in P_{r}\right) \pm 2^{-\Omega(k)}+\operatorname{Pr}\left[\neg \mathcal{E}_{1}\right] \operatorname{Pr}[\mathcal{F}]\left|P_{c r}\right|^{-1} \\
+ & \operatorname{Pr}\left[\neg \mathcal{E}_{1}\right] \operatorname{Pr}[\neg \mathcal{F}] \operatorname{Pr}\left[\mathcal{G}_{1}\right] \operatorname{Pr}\left[f(q, P)=p \mid \mathcal{G}_{1} \mathcal{G}_{2} \neg \mathcal{F} \neg \mathcal{E}_{1}\right] \pm 2^{-\Omega(k)} \\
+ & \operatorname{Pr}\left[\neg \mathcal{E}_{1}\right] \operatorname{Pr}[\neg \mathcal{F}] \operatorname{Pr}\left[\neg \mathcal{G}_{1}\right] \operatorname{Pr}\left[f(q, P)=p \mid \neg \mathcal{G}_{1} \neg \mathcal{F} \neg \mathcal{E}_{1}\right] \\
= & \operatorname{Pr}\left[\mathcal{E}_{1}\right]\left|P_{r}\right|^{-1} \delta\left(p \in P_{r}\right)+\operatorname{Pr}\left[\neg \mathcal{E}_{1}\right] \operatorname{Pr}[\mathcal{F}]\left|P_{c r}\right|^{-1} \\
+ & \operatorname{Pr}\left[\neg \mathcal{E}_{1}\right] \operatorname{Pr}[\neg \mathcal{F}] \operatorname{Pr}\left[\mathcal{G}_{1}\right]\left|P_{b r}\right|^{-1} \delta\left(p \in P_{b r}\right) \\
+ & \operatorname{Pr}\left[\neg \mathcal{E}_{1}\right] \operatorname{Pr}[\neg \mathcal{F}] \operatorname{Pr}\left[\neg \mathcal{G}_{1}\right] \operatorname{Pr}\left[f(q, P)=p \mid \neg \mathcal{E}_{1} \neg \mathcal{F} \neg \mathcal{G}_{1}\right] \pm 2^{-\Omega(k)} .
\end{aligned}
$$

On the other hand, since $\operatorname{Pr}[\mathcal{F}]=1-\binom{n-\left|P_{c r}\right|}{B}\binom{n}{B}^{-1}$, then $\operatorname{Pr}[X=p]$ is
$\operatorname{Pr}\left[\mathcal{E}_{1}\right]\left|P_{r}\right|^{-1} \delta\left(p \in P_{r}\right)+\operatorname{Pr}\left[\neg \mathcal{E}_{1}\right] \operatorname{Pr}[\mathcal{F}]\left|P_{c r}\right|^{-1}+\operatorname{Pr}\left[\neg \mathcal{E}_{1}\right] \operatorname{Pr}[\neg \mathcal{F}]\left|P_{b r}\right|^{-1} \delta\left(p \in P_{b r}\right)$,
so that,

$$
\begin{gathered}
|\operatorname{Pr}[f(q, P)=p]-\operatorname{Pr}[X=p]| \leq \\
\operatorname{Pr}\left[\neg \mathcal{E}_{1}\right] \operatorname{Pr}[\neg \mathcal{F}] \operatorname{Pr}\left[\neg \mathcal{G}_{1}\right] \operatorname{Pr}\left[f(q, P)=p \mid \neg \mathcal{E}_{1} \neg \mathcal{F} \neg \mathcal{G}_{1}\right]+2^{-\Omega(k)} .
\end{gathered}
$$

If $\left|P_{b r}\right|<B, \operatorname{Pr}\left[\neg \mathcal{E}_{1}\right]=2^{-\Omega(k)}$. If $\left|P_{c r}\right| \geq n^{2 / 3}$, $\operatorname{Pr}[\neg \mathcal{F}]=2^{-\Omega(k)}$. Otherwise $B \leq\left|P_{b r}\right| \leq\left|P_{c r}\right| \leq n^{2 / 3}$, and as shown for correctness, $\operatorname{Pr}\left[\neg \mathcal{G}_{1}\right]=2^{-\Omega(k)}$, which shows that $|\operatorname{Pr}[f(q, P)=p]-\operatorname{Pr}[X=p]|=2^{-\Omega(k)}$.

## Acknowledgments

The second author would like to thank Andrew Yao for support and hospitality while visiting Tsinghua University.

## References

[1] N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments, STOC, 1996.
[2] K. Bharat and A. Broder. Estimating the relative size and overlap of public web search engines, WWW, 1998.
[3] A. Beimel, Y. Ishai, T. Malkin. Reducing the servers computation in private information retrieval: PIR with preprocessing, CRYPTO, 2000.
[4] J. D. C. Benaloh, Verifiable secret-ballot elections. PhD thesis, Yale University, 1987.
[5] C. Cachin, J. Camenisch, J. Kilian and J. Müller. One-round secure computation and secure autonomous mobile agents, ICALP, 2000.
[6] C. Cachin, S. Micali and M. Stadler. Computationally private information retrieval with polylogarithmic communication, Eurocrypt, 1999.
[7] R. Canetti, Y. Lindell, R. Ostrovsky, and A. Sahai. Universally composable twoparty computation, STOC, 2002.
[8] B. Chor, N. Gilboa and M. Naor, Private information retrieval by keywords, Technical Report CS0917, Department of Computer Science, Technion, 1997.
[9] B. Chor, O. Goldreich, E. Kushilevitz and M. Sudan. Private information retrieval, FOCS, 1995.
[10] W. Du and M. J. Attalah. Protocols for secure remote database access with approximate matching, CCS - Workshop on Security and Privacy in E-commerce, 2000.
[11] S. Even, O. Goldreich and A. Lempel. A randomized protocol for signing contracts, Communications of the ACM, 1985.
[12] J. Feigenbaum, Y. Ishai, T. Malkin, K. Nissim, M. Strauss, and R. Wright. Secure multiparty computation of approximations, ICALP 2001.
[13] M. Freedman, K. Nissim and B. Pinkas. Efficient private matching and set intersection, Eurocrypt, 2004.
[14] B. Goethals, S. Laur, H. Lipmaa, and T. Mielikainen. On secure scalar product computation for privacy-preserving data mining, ICISC, 2004.
[15] J. Hastad, R. Impagliazzo, L. A. Levin, and M. Luby. Construction of a pseudorandom generator from any one-way function, Technical Report TR-91-068, International Computer Science Institute, 1991.
[16] Y. Gertner, Y. Ishai, E. Kushilevitz and T. Malkin. Protecting data privacy in private information retrieval schemes, STOC, 1998.
[17] O. Goldreich. Secure multi-party computation, 1998. Available at http://philby.ucsd.edu/
[18] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game, STOC, 1987.
[19] S. Goldwasser and S. Micali. Probabilistic encryption, JCSS, 1984.
[20] S. Halevi, R. Krauthgamer, E. Kushilevitz, and K. Nissim. Private approximation of NP-hard functions, STOC 2001.
[21] R. Impagliazzo and M. Luby. One-way functions are essential for complexity-based cryptography, FOCS, 1989.
[22] P. Indyk. High-dimensional computational geometry. PhD Thesis, Stanford University, 2000.
[23] P. Indyk and R. Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality, STOC, 1998.
[24] P. Indyk and D. Woodruff. Polylogarithmic private approximations and efficient matching, ECCC, Technical Report TR05-117, 2005.
[25] E. Kushilevitz, R. Ostrovsky and Y. Rabani. Efficient search for approximate nearest neighbor in high dimensional spaces, STOC, 1998.
[26] Y. Lindell and B. Pinkas. Privacy preserving data mining, Crypto, 2000.
[27] V.D. Milman and G. Schechtman. Asymptotic Theory of Finite Dimensional Normed Spaces. Lecture Notes in Mathematics, 1200, Springer Verlag, 1986.
[28] D. Naccache and J. Stern. A new public key cryptosystem, Eurocrypt, 1997.
[29] M. Naor and K. Nissim. Communication complexity and secure function evaluation, STOC, 2001.
[30] M. Naor and B. Pinkas. Oblivious transfer and polynomial evaluation, STOC, 1999.
[31] P. Paillier. Public-key cryptosystems based on composite degree residuosity classes, Eucrocrypt, 1999.
[32] M. Rabin. How to exchange secrets by oblivious transfer. Technical report TR81, Aiken Computation Lab, 1981.
[33] A. C. Yao. Protocols for secure computations, FOCS, 1982.

## A Cryptographic Tools

Homomorphic Encryption. An encryption scheme, $E:\left(G_{1},+\right) \rightarrow\left(G_{2}, \cdot\right)$ is homomorphic if for all $a, b \in G_{1}, E(a+b)=E(a) \cdot E(b)$. For more background on this primitive see, for example, [19, 28]. We will make use of the Paillier homomorphic encryption scheme 31.

Oblivious Transfer and SPIR. Oblivious transfer is equivalent to the notion of symmetrically-private information retrieval (SPIR), where the latter usually refers to communication-efficient implementations of the former. SPIR was introduced in [16. With each invocation of a SPIR protocol a user learns exactly one bit of a binary database while giving the server no information about which bit was learned. We rely on single-server SPIR schemes in our protocols. Such schemes necessarily offer computational, rather than unconditional, security 9 . Applying the transformation of 30 to the PIR scheme of [6] give SPIR constructions with $\tilde{O}(n)$ server work and $\tilde{O}(1)$ communication.

One issue is that in some of our schemes, we actually perform OT on records rather than on bits. It is a simple matter to convert a binary OT scheme into an OT scheme on records by running $r$ invocations of the binary scheme in parallel, where $r$ is the record size. This gives us a 1-round, $\tilde{O}(r)$ communication, $\tilde{O}(n r)$ server work OT protocol on records of size $r$. The dependence on $r$ can be improved using techniques of [8].

Secure Function Evaluation. In [18, 33] it is shown how two parties holdings inputs $x$ and $y$ can privately evaluate any circuit $C$ with communication $O(k(|C|+|x|+|y|))$, where $k$ is a security parameter. In [5] it is shown how to do this in one round for the semi-honest case we consider. The time complexity is the same as the communication. We use such protocols as black boxes in our protocols.

## B Private c-Approximate NN Based on Locality Sensitive Hashing

We give an alternative private $c$-approximate NN protocol, with slightly more communication than that in section 4.2, but less work (though still polynomial). It is based on locality sensitive hashing (LSH) [23]. The fact we need is that there is a family of functions $\mathcal{G}:\{0,1\}^{d} \rightarrow\{0,1\}^{\tilde{O}(1)}$ such that each $g \in \mathcal{G}$ has description size $\tilde{O}(1)$, and $\mathcal{G}$ is such that for all $p, q \in\{0,1\}^{d}$,

$$
\operatorname{Pr}_{g \in \mathcal{G}}[g(p)=g(q)]=\Theta\left(n^{-\Delta(p, q) / c r}\right)
$$

Recall that Alice has a point $q \in\{0,1\}^{d}$ and Bob has $n$ points $P \subseteq\{0,1\}^{d}$. For correctness, Alice should learn a point of $P_{c r}$ provided $P_{r} \neq \emptyset$. For privacy, her view should be simulatable given only $P_{c r}$.

Our protocol is similar to that in section4.2. When $\left|P_{c r}\right|$ is large, one can run a secure function evaluation with Alice's point $q$ as input, together with a random sample $P^{\prime}$ of roughly a $k /\left|P_{c r}\right|$ fraction of Bob's points $P$. The circuit returns a random point of $P_{\tilde{O}}^{\prime} \cap P_{c r}$ which is non-empty with probabiity $1-2^{-\Omega(k)}$. The communication is $\tilde{O}\left(n /\left|P_{c r}\right|\right)$.

On the other hand, when $\left|P_{c r}\right|$ is small, if Alice and Bob exchange functions $g_{i}$ independently $\tilde{O}\left(n^{1 / c}\right)$ times, then with overwhelming probability $P_{r} \subseteq \cup_{i} S_{i}$, where $S_{i}$ denotes the subset of Bob's points $p$ with $g_{i}(p)=g_{i}(q)$. Using a secure ciruit with ROM, we can obtain these sets $S_{i}$, and output a random point of $P_{r}$. The communication is $\tilde{O}\left(n^{1 / c}\left|P_{c r}\right|\right)$.

Our protocol balances these approaches to achieve $\tilde{O}\left(n^{1 / 2+1 /(2 c)}\right)$ communication.

There are a few technicalities dodged by this intuition. First, even though the parties exchange $\tilde{O}\left(n^{1 / c}\right)$ different $g_{i}$, and can thus guarantee that each $p$ is in some $S_{i}$ with probability $1-2^{-\Omega(k)}$, it may be that whenever $p \in S_{i}$, many points from $P \backslash P_{c r}$ also land in $S_{i}$, so that $S_{i}$ is very large. Even though we only expect $\left|P \backslash P_{c r}\right| O(1 / n)=O(1)$ points from $P \backslash P_{c r}$ in $S_{i}$, since $\operatorname{Pr}\left[p \in S_{i}\right]=\Theta\left(n^{-1 / c}\right)$
is small, $p$ may only be in $S_{i}$ when $S_{i}$ is large. Because the size of the $S_{i}$ affects the communication of our protocol, we cannot always afford for the ROM to receive the whole $S_{i}$ (sometimes we will truncate it). However, in the analysis, we show that the average $S_{i}$ is small, and this will be enough to get by with low communication.

Second, we need to extend the notion of a lookup gate given in section 2 Instead of just mapping inputs $(i, j)$ to output $R_{i}[j]$, the $j$ th entry in the $i$ th party's ROM, we also allow $j$ to be a key, so that the output is the record in $R_{i}$ keyed by $j$. This can be done efficiently using [8], and Theorem 2 is unchanged, assuming the length of the keys is $\tilde{O}(1)$.

## LSH $(q, P)$ :

1. Set $B=\tilde{O}\left(n^{1 / 2+1 /(2 c)}\right)$ and $C=\tilde{O}\left(n^{1 / c}\right)$.
2. Bob finds a random subset $P^{\prime}$ of $P$ of size $B$.
3. For $i=1$ to $k$,
(a) Alice and Bob agree upon $C$ random $g_{i, j} \in \mathcal{G}$.
(b) Bob creates a ROM $L$ with entries $L(v)$ containing $p$ s.t. $g(p)=v$.
(c) A secure circuit with ROM $L$ on input ( $q,\left\{g_{i, j}\right\}$ ) executes:

- Compute $v_{i, j}=g_{i, j}(q)$ for each $j$.
- Lookup the $L\left(v_{i, j}\right)$ one by one for the different $v_{i, j}$ until the communication exceeds $d B$. If it is less, pad it to $d B$.
- Output shares $S_{i}^{1}, S_{i}^{2}$ so that $S_{i}^{1} \oplus S_{i}^{2}$ is the set of sets $L\left(v_{j}\right)$.

4. A secure circuit on inputs $P^{\prime}, S_{i}^{1}, S_{i}^{2}$ executes:

- Compute the set $S_{i}=S_{i}^{1} \oplus S_{i}^{2}=\cup_{j} L\left(v_{j}\right)$ for all $i$.
- Let $f(q, P)$ be random in $P_{c r} \cap P^{\prime}$ if it is non-empty.
- Else let $f(q, P)$ be random in $P_{r} \cap \cup_{i} S_{i}$ if it is non-empty, else set $f(q, P)=\emptyset$.
- Output $(f(q, P)$, null $)$.

The communication is $\tilde{O}(d B)$. By using homomorphic encryption, one can reduce the dependence on $d$, as per remark 3. Let $\mathcal{E}$ be the event that $P_{r} \subseteq \cup_{i} S_{i}$, and let $\mathcal{F}$ be the event that $P_{c r} \cap P^{\prime}$ is non-empty.

Correctness. Suppose $P_{r} \neq \emptyset$. The probability $s$ of correctness is just the probability we don't output $\emptyset$. Thus $s \geq \operatorname{Pr}[\mathcal{F}]+\operatorname{Pr}[\neg \mathcal{F}] \operatorname{Pr}[f(q, P) \neq \emptyset \mid \neg \mathcal{F}]$.

Case $\left|P_{c r}\right| \geq n^{1 / 2-1 /(2 c)}$ : For sufficiently large $B$, we have $s \geq \operatorname{Pr}[\mathcal{F}]=1-2^{-\Omega(k)}$. Case $\left|P_{c r}\right|<n^{1 / 2-1 /(2 c)}$ : It is enough to show $\operatorname{Pr}[f(q, P) \neq \emptyset \mid \neg \mathcal{F}]=1-2^{-\Omega(k)}$. Fix $i$. Put $Y=\sum_{j}\left|L\left(v_{i, j}\right)\right|$, where $\left|L\left(v_{i, j}\right)\right|$ denotes the number of points in $L\left(v_{i, j}\right)$. The expected number of points in $P \backslash P_{c r}$ that are in $L\left(v_{i, j}\right)$ is at most $n \cdot O(1 / n)=O(1)$. Since $\left|P_{c r}\right|<n^{1 / 2-1 /(2 c)}, \mathbf{E}\left[L\left(v_{i, j}\right)\right]<n^{1 / 2-1 /(2 c)}+O(1)$. Thus $\mathbf{E}[Y] \leq B / 3$ for large enough $B$, so $\operatorname{Pr}[Y>B] \leq 1 / 3$ by Markov's inequality. Thus, with probability $1-2^{-\Omega(k)}$, for at least half of the $i, S_{i}$ is not truncated in step 3c. Moreover, for large enough $B$, any $i$, and any $p \in P_{r}$,
$\operatorname{Pr}\left[p \in S_{i}\right]=1-2^{-\Omega(k)}$ for large enough $C$. By a few union bounds then, $\operatorname{Pr}\left[P_{r} \subseteq \cup_{i} S_{i}\right]=\operatorname{Pr}[\mathcal{E}]=1-2^{-\Omega(k)}$. Thus,

$$
\begin{aligned}
\operatorname{Pr}[f(q, P) \neq \emptyset \mid \neg \mathcal{F}] & \geq \operatorname{Pr}[f(q, P) \neq \emptyset, \mathcal{E} \mid \neg \mathcal{F}] \\
& =\operatorname{Pr}[f(q, P) \neq \emptyset \mid \mathcal{E}, \neg \mathcal{F}] \operatorname{Pr}[\mathcal{E}] \\
& \geq 1-2^{-\Omega(k)}
\end{aligned}
$$

Privacy. Note that Bob gets no output, so Alice's privacy follows from that of the secure circuit protocol. We construct a simulator $\operatorname{Sim}\left(1^{n}, P_{c r}, q\right)$ so that the distributions $\langle q, P, f(q, P)\rangle$ and $\left\langle q, P, \operatorname{Sim}\left(1^{n}, P_{c r}, q\right)\right\rangle$ are statistically close. Bob's privacy then follows by the composition with the secure circuit protocol.

## $\boldsymbol{\operatorname { S i m }}\left(1^{n}, P_{c r}, q\right):$

1. Set $B=\tilde{O}\left(n^{1 / 2+1 /(2 c)}\right)$.
2. With probabiity $1-\binom{n-\left|P_{c r}\right|}{B}\binom{n}{B}^{-1}$, output a random element of $P_{c r}$.
3. Else output a random element of $P_{r}$.

Let $X$ denote the output of $\operatorname{Sim}\left(1^{n}, P_{c r}, q\right)$. It suffices to show that for each $p \in$ $P,|\operatorname{Pr}[f(q, P)=p]-\operatorname{Pr}[X=p]|=2^{-\Omega(k)}$, since this also implies $\mid \operatorname{Pr}[f(q, P)=$ $\emptyset]-\operatorname{Pr}[X=\emptyset] \mid=2^{-\Omega(k)}$. We have

$$
\begin{aligned}
\operatorname{Pr}[f(q, P)=p] & =\operatorname{Pr}[f(q, P)=p, \mathcal{F}]+\operatorname{Pr}[f(q, P)=p, \neg \mathcal{F}] \\
& =\operatorname{Pr}[\mathcal{F}]\left|P_{c r}\right|^{-1}+\operatorname{Pr}[f(q, P)=p, \neg \mathcal{F}]
\end{aligned}
$$

Note that $\operatorname{Pr}[\mathcal{F}]=1-\binom{n-\left|P_{c r}\right|}{B}\binom{n}{B}^{-1}$. Therefore,

$$
|\operatorname{Pr}[f(q, P)=p]-\operatorname{Pr}[X=p]|=\left.\operatorname{Pr}[\neg \mathcal{F}]\left|\operatorname{Pr}[f(q, P)=p \mid \neg \mathcal{F}]-\delta\left(p \in P_{r}\right)\right| P_{r}\right|^{-1} \mid
$$

If $\left|P_{c r}\right| \geq n^{1 / 2-1 /(2 c)}$, this is $2^{-\Omega(k)}$, since then $\operatorname{Pr}[\neg \mathcal{F}]=2^{-\Omega(k)}$. Otherwise, $\left|P_{c r}\right|<n^{1 / 2-1 /(2 c)}$, and as shown in the proof of correctness, we have $\operatorname{Pr}[\mathcal{E}]=$ $\operatorname{Pr}\left[P_{r} \subseteq \cup_{i} S_{i}\right]=1-2^{-\Omega(k)}$. Thus

$$
\begin{aligned}
\operatorname{Pr}[f(q, P) & =p \mid \neg \mathcal{F}] \\
& =\operatorname{Pr}[f(q, P)=p \mid \mathcal{E}, \neg \mathcal{F}] \operatorname{Pr}[\mathcal{E}] \pm 2^{-\Omega(k)} \\
& =\delta\left(p \in P_{r}\right)\left|P_{r}\right|^{-1} \pm 2^{-\Omega(k)}
\end{aligned}
$$

which completes the proof.


[^0]:    ${ }^{1}$ A rare exception is the result of [29, who show how to obtain private and communication-efficient versions of non-private protocols, as long as the communication cost is logarithmic.

[^1]:    ${ }^{2}$ We write $f=\tilde{O}(g)$ if $f(n, k)=O\left(g(n, k) \log ^{O(1)}(n) \operatorname{poly}(k)\right)$, where $k$ is a security parameter.

