

Public-Key Encryption with Efficient Amortized Updates

Nishanth Chandran^{1*}, Rafail Ostrovsky^{2**}, and William E. Skeith III³

¹ Department of Computer Science, University of California, Los Angeles.

² Departments of Computer Science & Mathematics, UCLA

Email: rafail@cs.ucla.edu

³ Department of Computer Science, City College, CUNY.

Abstract. Searching and modifying public-key encrypted data has received a lot of attention in recent literature. In this paper we re-visit this important topic and achieve improved amortized bounds including resolving a prominent open question posed by Boneh et al. [3].

First, we consider the following much simpler to state problem: A server holds a copy of Alice’s database that has been encrypted under Alice’s public key. Alice would like to allow other users in the system to replace a bit of their choice in the server’s database by communicating directly with the server, despite other users not having Alice’s private key. However, Alice requires that the server should not know which bit was modified. Additionally, she requires that the modification protocol should have “small” communication complexity (sub-linear in the database size). This task is referred to as private database modification, and is a central tool in building a more general protocol for modifying and searching over public-key encrypted data. Boneh et al. [3] first considered the problem and gave a protocol to modify 1 bit of an N -bit database with communication complexity $\mathcal{O}(\sqrt{N})$. Naturally, one can ask if we can improve upon this. Indeed, the recent work of Gentry [9] shows that under lattice assumptions, better asymptotic communication complexity is possible. However, current algebraic techniques based on any singly homomorphic encryption, or bilinear maps (which includes for example, all known cryptosystems based on factoring and discrete logs) cannot achieve communication better than $\mathcal{O}(\sqrt{N})$ (see [17]). In this paper we study the problem of improving the communication complexity for modifying L bits of an N -bit database. Our main result is a *black-box* construction of a private database modification protocol to modify L bits of an N -bit database, using a protocol for modifying 1 bit. Our protocol has communication complexity $\tilde{\mathcal{O}}(N^\beta L^{(1+\alpha)(1-\beta)})$, where $0 < \alpha < 1$ can be an arbitrary constant and $N^\beta, 0 < \beta < 1$ (for constant β) is the communication complexity of a protocol for modifying 1 bit of an N -bit database. We stress that our amortized protocol improves the communication complexity in all cases when the single bit modification

* Supported in part by NSF grants 0716835, 0716389, 0830803, 0916574.

** Supported in part by IBM Faculty Award, Xerox Innovation Group Award, the Okawa Foundation Award, Intel, Teradata, NSF grants 0716835, 0716389, 0830803, 0916574 and U.C. MICRO grant.

protocol uses any known cryptosystem based on factoring or discrete logs.

In addition to our general reduction, we show how to realize an implementation of our amortized protocol under the subgroup decision problem [2]. (We remark that in contrast with recent work of Lipmaa [16] on the same topic, our database size *does not grow* with every update, and stays exactly the same size.)

As sample corollaries to our main result, we obtain the following:

- First, we apply our private database modification protocol to answer the main open question of [3]. More specifically, we construct a public-key encryption scheme supporting PIR queries that allows every message to have a non-constant number of keywords associated with it, which is secure under the subgroup decision problem.
- Second, we show that one can apply our techniques to obtain more efficient communication complexity when parties wish to increment or decrement multiple cryptographic counters (formalized by Katz et al. [15]).

We believe that “public-key encrypted” amortized database modification is an important cryptographic primitive in its own right and will be useful in other applications.

1 Introduction

The problem of private database modification was first studied in the context of public-key encryption supporting private information retrieval (PIR) queries by Boneh et al. [3]. The private database modification protocol of [3] requires communication complexity $\mathcal{O}(\sqrt{N})$ to modify (i.e., change a 0 bit into a 1 and vice-versa) one bit of an N -bit database. Furthermore Ostrovsky and Skeith showed in [17] that using currently known algebraic techniques (which will not increase the database size after an update) with singly homomorphic encryption or bilinear maps, one cannot obtain better communication complexity to modify a single bit. Hence, we turn to the question of modifying multiple bits of the database. Using repeated application of the protocol from [3], one can obtain a private database modification protocol to modify L bits with communication complexity $\mathcal{O}(L\sqrt{N})$.

Lipmaa, in [16], also considered the question of amortizing the communication complexity of private database modification. However, his protocol has a significant drawback. In [16], the size of the database increases with every update made and after only $\mathcal{O}(\frac{\sqrt[4]{N}}{\log^4 N})$ bits have been updated in the database, Alice (the owner of the database) needs to download the entire database and re-send a new encrypted database, for the protocol to have efficient communication complexity thereafter. We shall see a little later that in applications of the private database modification protocol, this drawback is significant.

1.1 Main Result

Let N^β (for constant $0 < \beta < 1$) be the communication complexity of a private database modification protocol for modifying 1 bit of an N -bit database. Our

main contribution in this paper is a *black-box* construction of an amortized protocol (from a protocol that modifies 1 bit) with communication complexity $\tilde{O}(N^\beta L^{(1+\alpha)(1-\beta)})$ when modifying L bits of the database (where $0 < \alpha < 1$ is an arbitrary constant), *without ever increasing the size of the database, regardless of the number of updates*. Although we believe the amortized protocol for database modification to be of independent interest, we also describe two results that we obtain as corollaries of our main result.

1.2 Applications

Our first application is an answer to the main open question of [3] (resolving the main drawback of their solution.) Recall that in [3], a private database modification protocol was used to construct a public-key encryption scheme supporting PIR queries. An illustrative example of this concept is that of web-based email. Suppose that Alice stores her email on the server of a storage provider Bob (as is the case for a Yahoo! or Hotmail email account, for example). Bob must provide Alice with the ability to collect, retrieve, search and delete emails but at the same time learn nothing about the contents of the email nor the search criteria used by Alice. For example, a user might send an encrypted email to Alice (via Bob) that is marked with the keyword “Urgent”. Bob should store this email without knowing that a user sent an email marked as urgent. Later, Alice should have the ability to retrieve all email messages marked with “Urgent” without Bob knowing what search criteria Alice used. The goal was to obtain communication efficient protocols for this task. This goal was accomplished by making use of a protocol for private database modification in order to mark emails as containing certain keywords. However, in order to keep the communication complexity of the protocol sub-linear, the authors of [3] put constraints on the number of keywords that could be attached to a single message. In particular, this number was forced to be a constant. We can directly apply our batch update protocol to remove this constraint, allowing for non-constant numbers of keywords to be associated to a single message (which may often be the case for large messages). Note, that if we were to use a modification protocol in which the size of the database increased with every update (such as the one in [16]), then Alice needs to frequently download her entire email and send an updated database back to Bob, so that further executions of the modification protocol can have efficient communication complexity. This means that if for example, Alice does not check her mail for a period of time, then users will no longer be able to mark messages with keywords with efficient communication complexity. This defeats the entire purpose of having an email system supporting oblivious collect, retrieve, search and delete queries. Alice could simply achieve all these queries in an oblivious manner when she downloads the entire database. Furthermore, [16] requires that message senders be aware of a certain aspect of the state of the email database (in particular, the number of layers of encryption) before they send a message and perform updates to the database to mark the message with keywords. This would appear to force an interactive protocol for message sending (even for the

case of updating a bit with a semi-honest server), which seems undesirable and need not be the case (as shown in [3]).

A second use of our main result is a protocol for amortized updates of cryptographic counters. Cryptographic counters, formalized by [15], allow a group of participants to increment and decrement a public-key encrypted cryptographic representation of a hidden numeric value, stored on a server, in such a way that the server can not observe the value of the counter. The value of the counter can then be determined only by someone with a private key. Cryptographic counters can be used in several electronic voting protocols [5, 1, 6, 7, 18, 8]. One can imagine a situation in which parties would like to modify not one cryptographic counter but several such counters. For example, there could be a total of N counters and every party could wish to modify at most L of them. This could be seen in situations where every voter must vote for at most L out of N candidates and possibly also specify a ranking of the L selected candidates. We show how to obtain a communication efficient protocol for batch cryptographic counters (better than the batch protocol that is obtained through the trivial repetition of existing protocols). Once again, we cannot use a modification protocol in which the size of the database increases with every update. This is because of the following reason. After every party updates the counters (or casts his or her votes), the party holding the private key, must obtain the value of the counter and send a new “updated” (re-encrypted) value to the server. Clearly, privacy of the protocol would be lost here.

1.3 Our Techniques and High-Level Outline of Our Constructions

Our starting point are the techniques from Ishai, Kushilevitz, Ostrovsky and Sahai [13] on batch codes. Before we explain the main ideas of our construction (and in particular why batch codes do not apply directly) we need to give short background on batch codes.

Recall that batch codes of [13] are used for encoding an N -bit database on M different servers such that a user could read L bits of the N -bit database, by reading at most t bits from every server. The goal is to minimize both the total storage of the M servers as well as minimize t . At a high level, Ishai et al., make D copies of every bit and store each copy on a different server. To decide the server on which the r^{th} copy of a bit should be stored, they use an expander graph for the encoding.

We note that we cannot apply the construction of Ishai et al., in our setting: in the setting of Ishai et al., the problem is to read a bit of the (static) database, and reading any of the D copies of the bit gives the correct value. However, while writing a bit of the database, modifying one copy out of D copies does not yield a correct solution and modifying all eliminates all the efficiency savings.

The setting in our main construction is as follows. Users wish to change L 0-bits of the database on the server, into 1-bits without the server knowing which L bits were changed to 1. Additionally, Alice (the owner of the database) wishes to change L 1-bits of the database on the server, into 0-bits without the server knowing which L bits were changed to 0. We wish to obtain an amortized

communication complexity for both these tasks. For solving this, we consider other encodings of a bit of the database. One fruitful approach is to encode every bit of the database through D bits such that the *majority* of the D bits decodes to the bit in the database. The D bits that encode a bit will again be on M different “virtual” databases. However, now in order to modify a bit in the database, one needs to modify a majority of the copies of the bit. Unfortunately, using a generic expander (as in [13]) the encoding *does not* allow us to enjoy the property that a user reads the same number of bits from every virtual database. However, it turns out that the careful use of *lossless* expanders for our encoding achieves the desired savings. This requires us to prove that for every set of L bits, one can modify a majority of each of the bits in the encoding by modifying the same (small) number of bits on each virtual database. The solution that we then obtain gives us an amortization on the communication complexity when we modify a 0 bit into a 1 as well as vice-versa. Finally, we remark that the recent work of Gentry on fully homomorphic encryption [9] could indeed be used to achieve better asymptotic communication complexity under lattice assumptions. However, as shown in [17], current algebraic techniques based on any singly homomorphic encryption, or bilinear maps (which includes for example, all known cryptosystems based on factoring and discrete logs) cannot achieve communication better than $\mathcal{O}(\sqrt{N})$. Furthermore, we note that the techniques for amortizing communication given here are abstract and apply to any protocol which privately updates single bits of a database. We stress that our amortized protocol improves the communication complexity in all cases when the single bit modification protocol uses any known cryptosystem based on factoring or discrete logs.

Organization of the paper. We begin with a brief description of the private database modification protocol of [3] in Section 2. In Section 3, we describe our main result, the private database modification protocol for modifying L bits with amortized communication complexity. In Section 4, we show how to use the amortized private database modification protocol to obtain a public-key encryption scheme supporting PIR queries with non-constant numbers of keywords associated with each message. In Appendix B, we show how to apply the amortized modification protocol to get an amortization of cryptographic counters.

2 Background: Private Database Modification with Sub-linear Communication

Consider the following problem. A server is holding a database of Alice’s, which has been encrypted under her public key. Alice would like to allow her friends to use her public key to update a bit (of their choice) in the database by communicating directly with the server. For concreteness, typically what will be meant by “update” or “modify” is translation by a non-identity element in a group (as in the [3] implementation). In some instances where the updates are made by

Alice herself (who has knowledge of pieces of the database contents) the updates can be designed to explicitly write any values of her choice. However all that is required in most cases (where the parties know nothing of the database contents) is that the new value be different than the previous (which is always the case for non-identity element translation in a group). Alice requires the following conditions:

1. The details of each modification are hidden from the server. That is, without the private key, each transaction for an update is computationally indistinguishable from any another.
2. Her friends need only a “small” amount of communication (sub-linear in the database size) in order to perform an update of a single bit.

For a database of size N , we’ll call a protocol for privately updating subsets of L bits by $\text{Update}(N, L)$, and protocols for updating a single bit will accordingly be denoted by $\text{Update}(N, 1)$. The first construction of an $\text{Update}(N, 1)$ protocol which satisfied the above requirements was developed in [3]. The [3] protocol made use of a homomorphic cryptosystem that allows computation of polynomials of total degree 2 on ciphertexts (due to Boneh et al. [2]). That protocol has communication complexity $\mathcal{O}(\sqrt{N})$ for updating a single bit.

Given only an $\text{Update}(N, 1)$ protocol, an $\text{Update}(N, L)$ protocol can be constructed simply by running $\text{Update}(N, 1)$ L times sequentially, which will of course come at a $\mathcal{O}(L\sqrt{N})$ cost in communication complexity. Hence, if $\Omega(\sqrt{N})$ updates are to be made at once, the total communication becomes $\Omega(N)$, making the scheme no better than various trivial privacy-preserving solutions⁴.

In this work, we present an oblivious database modification protocol that amortizes the communication complexity of modifying L bits of the database. Our $\text{Update}(N, L)$ protocol is obtained using any protocol to modify a single bit, in a black-box manner (for example the $\text{Update}(N, 1)$ protocol from [3]).

3 Private Database Modification with Batches

In this section, we describe how one can amortize the communication complexity when running a private database modification protocol to modify L bits. In other words, let $\text{Update}(N, 1)$ denote any protocol for private database modification to modify 1 bit of an N -bit database and let the communication complexity of $\text{Update}(N, 1)$ be denoted by $C_{N,1} = N^\beta$, for constant $0 < \beta < 1$. Let $\text{Update}(N, L)$ denote a private database modification protocol to modify L bits of an N -bit database and let the communication complexity of $\text{Update}(N, L)$ be denoted by $C_{N,L}$. We construct a protocol for $\text{Update}(N, L)$, such that

⁴ E.g., the database could be encrypted under any homomorphic scheme and then Alice’s friend could simply request the entire encrypted database, make the updates using homomorphic encryption, re-randomize by translating with encryptions of the identity, and send the resulting database back. A non-interactive version may also be obtained just by communicating an $\Omega(N)$ -length vector of ciphertext in a homomorphic scheme.

$C_{N,L} = \tilde{O}(N^\beta L^{(1+\alpha)(1-\beta)})$, where $0 < \alpha < 1$ can be an arbitrary constant. We note that, we can pick α such that $(1 + \alpha)(1 - \beta) < 1$ and this ensures that the communication complexity of our protocol $< LC_{N,1}$ for sufficiently large values of L .

We first begin with the description of our security game for privacy. We assume a semi-honest adversary \mathcal{A} that runs in probabilistic polynomial time (PPT). Informally, a semi-honest adversarial server should not have any knowledge about which L bits a user modified in the database (changed from 0 to 1). The security game when Alice modifies L 1-bits in the database to 0-bits is exactly the same (except that the protocol requires Alice to know the secret key of the encryption scheme being used). The interface for $\text{Update}(N, L)$ is described in Section 2. The security game for privacy is defined through the two experiments (0 and 1) below. Let W_b denote the probability with which \mathcal{A} outputs 1 in Experiment b for $b = 0, 1$.

1. In both experiments,
 - (a) The challenger picks the public key of the encryption scheme pk and sends it to \mathcal{A} .
 - (b) \mathcal{A} sends an N -bit string denoting the database to the challenger.
 - (c) The challenger encrypts these N bits using pk and sends it to \mathcal{A} .
 - (d) \mathcal{A} picks 2 sets S_0 and $S_1 \subseteq [N]$, such that $|S_0| = |S_1| = L$, where at every index in S_0 and S_1 , the database contains a 0. Also, for every index in S_0 and S_1 , \mathcal{A} specifies if the bit at that index must be changed to a 1 or not. Note that \mathcal{A} can also choose to not change any bit to a 1; this corresponds to the case when \mathcal{A} does not change any bit in the database.
2. In Experiment b , the challenger runs $\text{Update}(N, L)$ with \mathcal{A} using S_b as input.
3. \mathcal{A} outputs a bit 0 or 1.

Definition 1 *We say that $\text{Update}(N, L)$ is L -private, if for all semi-honest PPT adversaries \mathcal{A} , we have $|W_0 - W_1|$ is negligible.*

Let $\text{Update}^*(N, 1)$ denote a private database modification protocol in which a user runs the modification protocol but does not modify any bit of the database. For example, $\text{Update}^*(N, 1)$ could be defined just as $\text{Update}(N, 1)$ of [3], but replacing the encryptions of characteristic vectors with encryptions of 0-vectors.

We first begin with some background on expanders in §3.1. In §3.2, we describe our main construction.

3.1 Expander Graphs

Expanders are graphs that are sparse but highly connected. Expanders have had several applications in computer science (see for example the survey of [12]). We define expanders and lossless expanders below and refer the reader to [4, 12, 20, 11] for further details.

Definition 2 A bipartite multi-graph with $N = 2^n$ left vertices and $M = 2^m$ right vertices, where every left vertex has degree $D = 2^d$, can be specified by a function $\Gamma : [N] \times [D] \rightarrow [M]$, where $\Gamma(u, r)$ denotes the r^{th} neighbor of vertex $u \in [N]$. For a set $S \subseteq [N]$, $\Gamma(S)$ denotes the set of neighbors of S . That is, $\Gamma(S) = \{\Gamma(x, y) : x \in S, y \in [D]\}$. Let $|\Gamma(S)|$ denote the size of the set $\Gamma(S)$.

Definition 3 A bipartite graph $\Gamma : [N] \times [D] \rightarrow [M]$ is a (L, A) expander, if for every set $S \subseteq [N]$, with $|S| = L$, we have $|\Gamma(S)| \geq A \cdot L$. Γ is a $(\leq L_{\max}, A)$ expander if it is a (L, A) expander for every $L \leq L_{\max}$. An expander is unbalanced if $M \ll N$.

Definition 4 A $(\leq L_{\max}, A)$ expander $\Gamma : [N] \times [D] \rightarrow [M]$ is a (L_{\max}, ϵ) lossless expander if $A = (1 - \epsilon)D$.

3.2 Main Construction

Our solution, uses techniques from the work of Ishai et al. [13] on batch codes and their applications. Ishai et al. considered the problem of encoding an N -bit database on M different servers such that a user could read L bits of the N -bit database, by reading at most t bits from every server. The goal is to minimize the total storage of the M servers as well as minimize t .

The idea in Ishai et al. is as follows. Make D copies of every bit in the N -bit database. The parameters D and M are picked such that $\Gamma : [N] \times [D] \rightarrow [M]$ is a $(\leq L_{\max}, A)$ expander for some $A > 1$. Now, the ND bits are distributed among the M servers according to the expander graph. In other words, the r^{th} copy of bit $i \in [N]$ of the database is stored in database $\Gamma(i, r)$. Now one can show that to read any L bits of the N -bit database (with $L \leq L_{\max}$), one only needs to read at most 1 bit from each of the M servers. So, by reading 1 bit from each of the M servers, t is minimized. The bound on the total storage of the M servers is obtained through the expansion property of Γ , thus satisfying the other required property.

Note that [13], do not consider the problem of modifying bits of a database. The encoding in [13] works because in order to read a bit from the N -bit database, one only needs to read any copy of that bit. The encoding does not directly apply in our setting as modifying 1 bit out of the D bits that encode a bit does not result in a correct modification.

At a high level, our protocol for private database modification to modify L_{\max} bits of an N -bit database is as follows. We encode every bit of the database through D bits. The majority value of these D bits decodes to the original bit in the database. The resulting ND bits from the encoding are distributed into M “virtual” databases according to a (L_{\max}, ϵ) lossless expander graph Γ . Let the number of bits in each of the M virtual databases be denoted by a_1, a_2, \dots, a_M .

We will then show that to modify a majority of each of the bits in any set of L_{\max} bits of the N -bit database, one only needs to modify at most 1 bit from each of the M virtual databases. One can modify 1 bit from each of the M virtual databases using $\text{Update}(a_i, 1)$ for all $1 \leq i \leq M$. The bound on the

communication complexity of the protocol will be obtained through the lossless expansion property of Γ .

While reading a bit from the N -bit database, one reads all D bits that encode this bit from the M virtual databases and takes the majority value. We first describe how to create the virtual databases.

Creating Virtual Databases. Consider a (L_{max}, ϵ) lossless expander $\Gamma : [N] \times [D] \rightarrow [M]$ as defined in §3.1. Let $L = 2^l$ and $L_{max} = 2^{l_{max}}$.

Every node $u \in [N]$ represents a bit in the database and the $D = 2^d$ neighbors of the node u are the encoded bits of u . For bit $u \in [N]$, the r^{th} bit of the encoding of u is present in database $\Gamma(u, r)$. To read the value of bit $u \in [N]$, one reads all D bits of the encoding of u and takes the majority of these values as the value of u . Hence, note that to modify bit u , one has to modify a majority of the bits that encode u . Below, we show that this can be done by modifying at most 1 bit in every virtual database.

By the lossless expansion property, first note that for all sets $S \subseteq [N]$ with $|S| = L \leq L_{max}$, we have $|\Gamma(S)| \geq (1 - \epsilon)LD$. To modify bits from a set S , we show that there is a strategy to modify at least $(1 - 2\epsilon)D$ bits of the encodings of all the bits in S by modifying at most 1 bit in each of the virtual databases.

Lemma 1 *Let Γ be a lossless expander as above. Then for every subset $S \subseteq [N]$ where $|S| = L \leq L_{max}$, the number of nodes $v \in [M]$ that have exactly one neighbor in S (v is then called a unique neighbor node with respect to S) is at least $(1 - 2\epsilon)LD$.*

Proof. Let $S \subseteq [N]$ where $|S| = L \leq L_{max}$. Let x_1 be the number of nodes $v \in [M]$ that have exactly one neighbor in S and let x_2 be the number of nodes $v \in [M]$ that have at least 2 neighbors in S . Now, $|\Gamma(S)| = x_1 + x_2 \geq (1 - \epsilon)LD$ (from the property of lossless expansion). Assume for contradiction that $x_1 < (1 - 2\epsilon)LD$. That is, let $x_1 = (1 - 2\epsilon - \delta)LD$ for some $\delta > 0$. This means, that $x_2 = |\Gamma(S)| - x_1 \geq (\epsilon + \delta)LD$. From counting the edges that originate out of S , we have $LD \geq x_1 + 2x_2 \geq (1 + \delta)LD$ which cannot be true for $\delta > 0$ and is a contradiction. Hence, the lemma. \square

Lemma 2 *Fix any set $S \subseteq [N]$ where $|S| = L \leq L_{max}$. Let $g_S(v)$ for all $v \in [M]$, be a function such that $g_S(v) = \text{NIL}$ or $g_S(v) = u$ such that $u \in S$ and there exists $r \in [D]$ such that $\Gamma(u, r) = v$. In other words, $g_S(v)$ is either NIL or a neighbor of v in S . Let $h_S(u) = |g_S^{-1}(u)|$ for all $u \in S$. That is, $h_S(u)$ is the number of $v \in [M]$ such that $g_S(v) = u$. There exists a polynomial time computable function $g_S(\cdot)$, such that $h_S(u) \geq (1 - 2\epsilon)D$ for all $u \in S$, and furthermore the function $g_S(\cdot)$ can be constructed in polynomial time for any $S \subseteq [N]$.*

Proof. We shall construct $g_S(\cdot)$ as follows:

1. Let $S' = S$. A node $u \in S'$ is satisfied if $h_S(u) \geq (1 - 2\epsilon)D$.

2. For every node $v \in [M]$, let $g_S(v) = u$ if u is the only neighbor of v in S' . Let H denote the set of nodes in S' that are satisfied.
3. Set $S' = S - H$. If S' is not empty, repeat Step 2, otherwise halt, setting $g_S(v)$, for all unassigned nodes v , to NIL.

We will prove that at every iteration of the algorithm, at least one node in S' is satisfied. This means the algorithm will halt in time $\mathcal{O}(L)$. In this case, every node $u \in S$ is satisfied and hence $h_S(u) \geq (1 - 2\epsilon)D$ for all $u \in S$. Let $|S'| = L'$. We will show that $|H| > 0$. Let $|H| = h$. Let l be the number of unique neighbor nodes with respect to S' in $[M]$. We have that $l \geq (1 - 2\epsilon)L'D$ (By Lemma 1).

Now, consider a *satisfied* node $u \in S'$. The number of unique neighbor nodes with respect to S' in $[M]$ that have their unique neighbor as u can be at most D , as the degree of every node in $[N]$ is at most D .

Consider a node $u \in S'$ that is not satisfied. The number of unique neighbor nodes with respect to S' in $[M]$ that have their unique neighbor as u is strictly less than $(1 - 2\epsilon)D$. Otherwise, u would be satisfied. (This is because at no stage of the algorithm did we assign $g_S(v)$ to be u when v was also a neighbor of a node $u' \in S$ that was not already satisfied.)

Hence we have $l < hD + (L' - h)(1 - 2\epsilon)D$. Since $l \geq (1 - 2\epsilon)L'D$, we have that $hD + (L' - h)(1 - 2\epsilon)D > (1 - 2\epsilon)L'D$, which means $h > 0$. \square

We note that the above proof is similar in flavor to the proof of error correction in linear time encodable/decodable expander codes (Refer [19, 4] for further details.). Our protocol uses the specific lossless expander explicit construction from [11]. We pick ϵ , such that $1 - 2\epsilon > \frac{1}{2}$. In other words, $\epsilon < \frac{1}{4}$. We state the theorem below.

Theorem 1 [11] *For all constants $\alpha > 0$, every $N \in \mathbb{N}$, $L_{max} \leq N$, and $\epsilon > 0$, there is an explicit (L_{max}, ϵ) lossless expander $\Gamma : [N] \times [D] \rightarrow [M]$ with degree $D = \mathcal{O}((\log N)(\log L_{max})/\epsilon)^{1+1/\alpha}$ and $M \leq D^2 \cdot L_{max}^{1+\alpha}$. Moreover D is a power of 2.*

Protocol Description. Let $\text{Update}(N, 1)$ denote any private database modification protocol for modifying 1 bit of an N -bit database. We now describe our black-box construction of private database modification protocol $\text{Update}(N, L_{max})$ from $\text{Update}(N, 1)$.

1. Create M smaller databases according to lossless expander Γ from Theorem 1 and encode the bits of the database into the M smaller databases as described earlier. Let size of database $v \in [M]$ be denoted by a_v .
2. To modify a set $S \subseteq [N]$ of bits of the database with $|S| = L_{max}$, create $g_S(v)$ as described in Lemma 2.
3. Run $\text{Update}(a_v, 1)$ to modify bit $g_S(v)$ in database v for all databases $v \in [M]$. If $g_S(v) = \text{NIL}$, then run $\text{Update}^*(a_v, 1)$ with database v .

Protocol Correctness and Security. Let $\text{Update}(N, 1)$ denote a protocol for privately modifying 1 bit of an N -bit database. The correctness of the protocol $\text{Update}(N, L_{\max})$ follows trivially from Lemma 2 and from the correctness of the $\text{Update}(N, 1)$ protocol. The security is proven below.

Theorem 2 *If $\text{Update}(N, 1)$ is 1-private, then $\text{Update}(N, L_{\max})$ is L_{\max} -private.*

Proof. Lemma 2 shows that the number of bits we modify in each of the M virtual databases is independent of the subset of L_{\max} bits we wished to modify in the original database. In particular, in each virtual database, we either modify 1 bit by running $\text{Update}(a_v, 1)$ or do not modify any bits by running $\text{Update}^*(a_v, 1)$. Now, since $\text{Update}(a_v, 1)$ is 1-private, no adversary can distinguish between the case when we run $\text{Update}(a_v, 1)$ and modify a bit and when we run $\text{Update}^*(a_v, 1)$. Hence, it follows by a simple hybrid argument that $\text{Update}(N, L_{\max})$ is L_{\max} -private. \square

Communication Complexity. We now analyze the communication complexity of protocol $\text{Update}(N, L_{\max})$. Let the communication complexity of $\text{Update}(N, 1)$ be $C_{N,1} = N^\beta$ for some constant $0 < \beta < 1$. Note that if a_v is the number of bits in the v^{th} smaller database, then the communication complexity of $\text{Update}(N, L_{\max})$ is $C_{N,L_{\max}} = \sum_{i=1}^M C_{a_i,1}$. We have $C_{a_i,1} = a_i^\beta$. We also have $\sum_{i=1}^M a_i = ND$. Now, Hölder's inequality, states the following:

Let $1 \leq p, q \leq \infty$ with $\frac{1}{p} + \frac{1}{q} = 1$. Let n be a positive integer. Then,

$$\sum_{i=1}^n |x_i y_i| \leq \left(\sum_{i=1}^n |x_i|^p \right)^{\frac{1}{p}} \left(\sum_{i=1}^n |y_i|^q \right)^{\frac{1}{q}}$$

for all $(x_1, x_2, \dots, x_n), (y_1, y_2, \dots, y_n) \in \mathbb{R}^n$. In this inequality, let $n = M$, $x_i = a_i^\beta$ for all $1 \leq i \leq M$, $y_i = 1$ for all $1 \leq i \leq M$, $p = \frac{1}{\beta}$ and $q = \frac{1}{1-\beta}$. Now, by Hölder's inequality, we get

$$\sum_{i=1}^M |a_i^\beta| \leq \left(\sum_{i=1}^M |a_i| \right)^\beta \left(\sum_{i=1}^M 1 \right)^{1-\beta}$$

Now, since $\sum_{i=1}^M a_i = ND$, it follows that $C_{N,L_{\max}} \leq (ND)^\beta M^{1-\beta}$. Next, setting the parameters according to Theorem 1, we get the communication complexity to be $\mathcal{O}(N^\beta L_{\max}^{(1+\alpha)(1-\beta)} \left(\frac{\log N \log L_{\max}}{\epsilon} \right)^{(2-\beta)(1+\frac{1}{\alpha})})$, where $0 < \alpha < 1$ is an arbitrary constant. Now, we can pick $\alpha < \frac{\beta}{1-\beta}$, giving us $(1+\alpha)(1-\beta) < 1$. Hence, we get $C_{N,L_{\max}} = \tilde{\mathcal{O}}(N^\beta L_{\max}^{(1+\alpha)(1-\beta)})$. We note that if we use the protocol of [3] for $\text{Update}(N, 1)$, we have $\beta = \frac{1}{2}$ and the communication complexity of $\text{Update}(N, L_{\max})$ is $\tilde{\mathcal{O}}(\sqrt{NL_{\max}^{1+\alpha}})$.

Maintaining Consistencies over Different Values of L_{max} . Note that if we run $\text{Update}(N, L_{max})$ when we wish to modify L bits in the database with $L < L_{max}$, then the communication complexity is not optimal as the communication complexity depends only on L_{max} and not on L . For example, if we have $L_{max} = \mathcal{O}(\sqrt{N})$ and $C_{N,1} = \sqrt{N}$, then running $\text{Update}(N, L_{max})$ when we want to modify $L = \mathcal{O}(\sqrt[4]{N})$ bits, will not be optimal. $\text{Update}(N, L_{max})$ will then have communication complexity $\mathcal{O}(D^{3/2} \sqrt[4]{N^{3+\alpha}})$, which is more than the communication complexity when running $\text{Update}(N, 1)$, $\mathcal{O}(\sqrt[4]{N})$ times.

Now, if we use different lossless expanders for the encoding depending on the number of bits we wish to modify, then the encoding of each bit of the original database will not be consistent. More specifically, since the degree of the graphs are not uniform, we may not modify “enough” copies of a particular bit.

To overcome this difficulty, we pick $W = \mathcal{O}((\log N)^2/\epsilon)^{1+1/\alpha}$ bits to encode every bit and store them. Now, for all values of $L_{max} \leq N$, the corresponding value of D is $\leq W$. When we wish to modify L_{max} bits of the original database, we use the corresponding lossless expander with degree D . We repeat this protocol $\lceil \frac{W}{D} \rceil$ times using a different (disjoint) set of D bits of the encoding of every bit in each iteration. Now, since in each execution we modify at least $(1 - 2\epsilon)D$ bits of the encoding of a bit, in total we will modify at least $(1 - 2\epsilon)W$ bits of the encoding of every bit that we modify and hence the decoding of majority still works. Furthermore, note that we do not increase the communication complexity of the protocol.

We note that our protocols for modifying $L < L_{max}$ bits of the N -bit database, do not attempt to hide the value of L_{max} . Note that any protocol for modifying bits of a database, will reveal an upper bound on the number of bits that are modified. This is because, if we wish to hide the value of L_{max} , then such a protocol must be indistinguishable from a protocol where a user modifies all bits of the N -bit database and this protocol, by an information theoretic argument, must have communication complexity $\Omega(N)$. We assume that the bound, L_{max} , on the number of bits that we wish to modify in the database is public. Note that for optimal communication complexity L must equal L_{max} .

An Important Remark. In the context of the remote mail storage example, consider a scenario in which Alice, the holder of the secret key wishes to update some part of her own database (perhaps she would like to delete a recently read message). In such a situation, the contents of the portion of the database to be modified are known. Say, Alice wishes to modify L bits which maybe either 0 or 1. Alice will know the contents of the database, including that of the encoding (she can learn this through an efficient PIR protocol). Now, the protocol described above gives us an amortization on the communication complexity both while marking messages with keywords as well as when removing marked keywords (when deleting a message). This is because, each bit is encoded through D bits and the majority of the D bits decode to the bit in the N -bit database. Now, irrespective of whether we are changing a 0 bit into a 1, or vice-versa, it will always suffice to modify at most a majority of the D copies of the bit.

Other Encodings. In addition to the majority encoding, other encodings can also prove useful in this context. Rather than using a lossless expander and an encoding where D bits encode a 0 or a 1 through the majority, one could encode every bit through D bits where the sum (modulo 2) of the D bits determines the encoded bit. Now, in order to modify a bit (either from 0 to 1 or vice-versa), one only needs to modify any single bit out of the D bits and can then apply the results of [13] on batch codes more directly. This approach may be of particular use when it is desirable for a party without the secret key to modify 1-bits back to 0-bits (although this is not of utility in the main application of remote mail storage).

4 Applications of Batch Protocols for Database Modification to [3]

The protocol of [3] applies to a scenario that models a somewhat ideal internet-based email service: all email messages are encrypted under the user’s public key, yet the user can still perform the common tasks of searching for and retrieving messages via keywords, erasing messages, etc., without revealing any information to the service provider about the messages nor the keywords being searched for. Furthermore, this can be done with “small” (i.e. sub-linear) communication.

The protocols will typically involve a message sender, receiver, and a storage provider. We’ll use the following notational conventions to represent the various parties: \mathcal{X} will refer to a message sending party; \mathcal{Y} will refer to the message receiving party (owner of the private key); \mathcal{S} will refer to the server/storage provider.

The protocol of [3] accomplished the basic task outlined above, but in order to maintain sub-linear communication complexity as well as to preserve the correctness of the protocol, several limitations were enforced. The most prominent conditions needed were as follows:

1. The number of messages associated to a single keyword must be bounded by a constant.
2. The number of keywords in use must be proportional to the number of messages.
3. The number of keywords associated to a particular message must be bounded by a constant.

We still enforce conditions 1 and 2 (which apply for the same technical reasons regarding correctness) however using batch protocols for private database modification, we show how to relax the third condition and allow non-constant numbers of keywords to be associated with a single message. Clearly the protocol of [3] cannot have this capability for a keyword set of size $\Omega(\sqrt{N})$: The expected number of bits one is required to update would similarly be $\Omega(\sqrt{N})$, and \sqrt{N} executions of `Update($N, 1$)` from [3] will yield $\Omega(N)$ communication complexity for sending this single message, violating the requirement of maintaining sub-linear communication.

Protocol Description

The details of the protocol are fairly straightforward. Let $\mathcal{K}, \mathcal{E}, \mathcal{D}$ represent the key generation, encryption and decryption algorithms, respectively, of a public key cryptosystem that allows for the evaluation of polynomials of total degree 2 on ciphertext (e.g., [2]). Adopting the notation of [3], we'll denote the maximum number of keywords that can be associated to a single message by θ . The protocol of [3] requires that this value in fact be constant. We will make no such assumption on θ and demonstrate a protocol that satisfies the same definitions of correctness and of privacy.

For brevity, we direct the reader to [3] for formal definitions of correctness and privacy for such a protocol, and instead provide an intuitive summary here. Roughly speaking, for integers λ, θ and a message database of N messages, a public-key storage with keyword search is said to be (N, λ, θ) -correct if the protocol for retrieval of messages by keywords yields the appropriate results after any sequence of executions of the message sending protocol, given that not more than θ keywords are associated to a single message. The work of [3] presents a public-key storage with keyword search that is (N, λ, θ) -correct where θ is a constant, independent of N , and independent of the message size. Below, we extend this protocol to maintain communication efficiency in the case of non-constant θ . In order to simplify the description, we will present the protocol at a high level and refer the reader to the work of [3] for details when needed. The protocol consists of the following three algorithms.

KeyGen(s) — Run the key generation algorithm \mathcal{K} of the underlying cryptosystem to produce a public and private pair of keys.

Send $_{\mathcal{X}, \mathcal{S}}$ (M, W) — Sender \mathcal{X} wants to send message M marked with the set of keywords W to \mathcal{Y} via \mathcal{S} . \mathcal{X} encrypts M and the keywords and then proceeds as in [3] in order to update the keyword-message association structure. However, rather than repeatedly applying **Update**($N, 1$), \mathcal{X} will use the **Update**(N, θ) protocol described in §3.2 to efficiently perform the updates as a batch. Note that in order to mark a message as having a only single keyword, \mathcal{X} is required to update $\Omega(\log^2 N)$ bits of the Bloom filter structure that holds the keyword-message associations.

Retrieve $_{\mathcal{Y}, \mathcal{S}}$ (w) — \mathcal{Y} wishes to retrieve all messages associated with the keyword w and optionally erase them from the server. This protocol consists of steps similar to [3] in order to decrypt the locations of matching messages and subsequently download and decrypt. However in the case where message erasure is also performed, we will have \mathcal{Y} execute **Update**(N, θ) from §3.2 with \mathcal{S} , as opposed to repeated usage of **Update**($N, 1$) (as found in [3]) which allows us to handle non-constant numbers of keywords to be associated with a single message.

Theorem 3 *The Public-Key Storage with Keyword Search from the preceding construction is (N, λ, θ) -correct according to the definition of [3].*

Remark: The proof of the above theorem follows in much the same way as that of [3]. However, there is a need for one remark on this subject. Recall that in [3]

it was required that only a constant number of messages were associated to a particular keyword, since fixed-length buffers were needed to represent sets. As mentioned, we have adopted this same requirement for much the same reason. However, we would like to note that in a practical implementation of our protocol this may be harder to achieve since the increased number of keywords per message will naturally lead to more messages being associated with a particular keyword. That is, if θ is large, it will generally not be possible to have *every* message associated to θ keywords without exceeding λ messages associated to some particular keyword. None the less, we emphasize that our protocol conforms to the very same definitions for correctness as that of [3]; it is just that the antecedent will perhaps not come as easily.

We again leave the formal definitions of privacy for the sender and receiver to [3] for the purposes of brevity. Roughly, they state that the sending and receiving protocols do not reveal any information about the messages nor the keyword associations to a computationally bounded adversary. This is phrased via a standard indistinguishability condition.

Theorem 4 *Assuming CPA-security of the underlying cryptosystem, the Public-Key Storage with Keyword Search from the above construction is sender-private as well as receiver-private, according to the definitions of [3].*

Proof. This follows almost immediately from the privacy of the $\text{Update}(N, 1)$ protocol, Theorem 2 and the analogous theorem from [3]. \square

References

1. Josh C Benaloh and Moti Yung. Distributing the power of a government to enhance the privacy of voters. In *PODC '86: Proceedings of the fifth annual ACM symposium on Principles of distributed computing*, pages 52–62. ACM, 1986.
2. Dan Boneh, Eu-Jin Goh, and Kobbi Nissim. Evaluating 2-DNF formulas on ciphertexts. In *TCC'05*, pages 325–341, 2005.
3. Dan Boneh, Eyal Kushilevitz, Rafail Ostrovsky, and William E. Skeith. Public key encryption that allows PIR queries. In *CRYPTO'07*, pages 50–67, 2007.
4. Michael R. Capalbo, Omer Reingold, Salil P. Vadhan, and Avi Wigderson. Randomness conductors and constant-degree lossless expanders. In *IEEE Conference on Computational Complexity*, page 15, 2002.
5. Josh D. Cohen and Michael J. Fischer. A robust and verifiable cryptographically secure election scheme. *Symposium on Foundations of Computer Science*, pages 372–382, 1985.
6. Ronald Cramer, Matthew K. Franklin, Berry Schoenmakers, and Moti Yung. Multi-authority secret-ballot elections with linear work. In *EUROCRYPT'96*, pages 72–83, 1996.
7. Ronald Cramer, Rosario Gennaro, and Berry Schoenmakers. A secure and optimally efficient multi-authority election scheme. In *EUROCRYPT'97*, pages 103–118, 1997.
8. Ivan Damgard and Mads Jurik. Efficient protocols based on probabilistic encryption using composite degree residue classes, 2000.

9. Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.
10. Oded Goldreich. *Foundations of Cryptography: Volume 2, Basic Applications*. Cambridge University Press, New York, NY, USA, 2004.
11. Venkatesan Guruswami, Christopher Umans, and Salil P. Vadhan. Unbalanced expanders and randomness extractors from Parvaresh-Vardy codes. In *IEEE Conference on Computational Complexity*, pages 96–108, 2007.
12. Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc.*, 43:439–561, 2006.
13. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Batch codes and their applications. In *STOC'04*, pages 262–271, 2004.
14. Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography with constant computational overhead. In *STOC'08*, pages 433–442, 2008.
15. Jonathan Katz, Steven Myers, and Rafail Ostrovsky. Cryptographic counters and applications to electronic voting. In *EUROCRYPT'01*, pages 78–92, 2001.
16. Helger Lipmaa. Private branching programs: On communication-efficient crypto-computing. Cryptology ePrint Archive, Report 2008/107, 2008. <http://eprint.iacr.org/2008/107>.
17. Rafail Ostrovsky and William E. Skeith. Communication complexity in algebraic two-party protocols. In *CRYPTO'08*, pages 379–396, 2008.
18. Berry Schoenmakers. A simple publicly verifiable secret sharing scheme and its application to electronic. In *CRYPTO'99*, pages 148–164, 1999.
19. Michael Sipser and Daniel A. Spielman. Expander codes. *IEEE Transactions on Information Theory*, 42(6):1710–1722, 1996.
20. Amnon Ta-Shma, Christopher Umans, and David Zuckerman. Lossless condensers, unbalanced expanders, and extractors. *Combinatorica*, 27(2):213–240, 2007.
21. Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *FOCS'82*, pages 160–164, 1982.

A Appendix

A.1 Privacy Game for Batch Cryptographic Counters

The privacy for batch cryptographic counters with respect to server T (modified from [15]) is given through the following two experiments 0 and 1. Let the batch cryptographic counter protocol with N counters, where every party modifies at most L counters, be denoted by $\text{Counter}(N, L)$. Let W_b denote the probability with which \mathcal{A} outputs 1 in Experiment b for $b = 0, 1$.

1. In both experiments,
 - (a) The challenger picks the public key of the encryption scheme pk and sends it to \mathcal{A} . \mathcal{A} plays the role of the server T here.
 - (b) \mathcal{A} sends N values denoting the initial values of the N cryptographic counters to the challenger.
 - (c) The challenger encrypts these N values using pk and sends it to \mathcal{A} .
 - (d) \mathcal{A} picks 2 sets S_0 and $S_1 \subseteq [N]$, such that $|S_0| = |S_1| = L$. Also, for every index in S_0 and S_1 , \mathcal{A} denotes the value by which the counter in the index must be incremented or decremented (\mathcal{A} may also specify that this counter not be changed).

2. In Experiment b , the challenger runs $\text{Counter}(N, L)$ with \mathcal{A} using S_b as input.
3. \mathcal{A} outputs a bit 0 or 1.

Definition 5 We say that $\text{Counter}(N, L)$ is L -private with respect to server T (run by \mathcal{A}), if for all semi-honest PPT adversaries \mathcal{A} , we have $|W_0 - W_1|$ is negligible.

The privacy of $\text{Transfer}(x, y)$ with respect to party P and server T is defined via standard definitions of two-party computation which guarantees that P does not learn anything other than the final value of the counter and that T does not learn anything.

B Application of Batch Protocols for Database Modification to Cryptographic Counters

Cryptographic counters, formalized by Katz et al. [15], allow a group of participants to increment and decrement a cryptographic representation of a hidden numeric value privately. The value of the counter can then be determined by a specific party only. More formally, there are a set of R parties, $\{P_1, P_2, \dots, P_R\}$. These parties wish to increment and decrement the value of a specific counter C which is stored by a party T , assumed to be semi-honest. After they have incremented/decremented the counter C , T must reveal the value of the counter to a specific party denoted by P . T is semi-honest and is trusted not to collude with P . At the same time, parties wish only the output of the counter to be revealed to P . One can implement this protocol in the following way. Let P pick a public/private key pair (pk_P, sk_P) of an additively homomorphic encryption scheme over \mathbb{Z}_n with n larger than the maximum value of the counter. Let P_i hold input x_i . Let $\mathcal{E}(pk, m)$ denote the encryption of message m with public key pk . Now, P_i sends $\mathcal{E}(pk_P, x_i)$ to T . Using the additive homomorphic property, T computes $\mathcal{E}(pk_P, \sum_{i=1}^R x_i)$ and sends the result to P who can then compute $\sum_{i=1}^R x_i$.

Now, suppose there are N such counters C_1, C_2, \dots, C_N that the parties wish to update. Furthermore, assume that each user updates no more than L of these N counters. No user P_i wishes to reveal to anyone, which of the counters he/she modified. An example of this situation would be a voting protocol which has several candidates (N candidates). Voters have to select L out of these N candidates and rank them. Candidates are then selected according to a weighted sum of their votes.

Now, let $x_i[1], \dots, x_i[N]$ denote the inputs (or weighted votes) held by P_i (only at most L out of these values are non-zero). Using the solution described above, P_i can send $\mathcal{E}(pk_P, x_i[1]), \dots, \mathcal{E}(pk_P, x_i[N])$ to T . Using the additive homomorphic property, T can compute $\mathcal{E}(pk_P, \sum_{i=1}^R x_i[1]), \dots, \mathcal{E}(pk_P, \sum_{i=1}^R x_i[N])$ and send the result to P . However, this protocol has communication complexity $\mathcal{O}(N)$ for every user P_i .

Let a cryptographic counter protocol between a user and server T , where there are N counters and every party modifies at most L out of the N counters, be denoted by $\text{Counter}(N, L)$. Let the protocol to transfer an encrypted

value of the final counter value from the server T and user P be denoted by $\text{Transfer}(x, y)$, where x and y are inputs of T and P respectively. The privacy game for $\text{Counter}(N, L)$ with respect to server T and the privacy game for $\text{Transfer}(x, y)$ with respect to party P and server T is given in Appendix A.1. We do not focus on the other security requirements such as universal verifiability and robustness ([15]). Universal verifiability informally means that any party (including third parties) can be convinced that all votes were cast correctly and that the tally was made correctly. Robustness informally means that the final output can be computed even in the presence of a few faulty parties.

We now describe below two solutions for $\text{Counter}(N, L)$ that have communication complexity $\mathcal{O}(\sqrt{L^{1+\alpha}N})$ poly-log N .

B.1 Protocol Using Lossless Unbalanced Expanders

This protocol, uses the private database modification protocol for modifying L bits of an N -bit database from §3.2. We first note that the additively homomorphic encryption scheme of [2] can be used to encrypt messages from a polynomially large message space (of size n). Choose n such that n is greater than the maximum value of the counter. Now, since the encryption scheme of [2] is additively homomorphic, we can use this scheme in order to encrypt the value of the counter. Every user P_i can update the counter by sending an encryption of their input x_i to T .

Now, we describe below how we can amortize the communication complexity of this protocol. We use the (L_{max}, ϵ) lossless expander Γ from Theorem 1. The protocol requires the number of parties R to be less than $\frac{1}{4\epsilon}$. The protocol $\text{Counter}(N, L)$ is described below.

1. Encode every counter through D different counters. To decode, the simple majority value of all values held in these D counters is the value of the counter. Initially these D counters all hold an encryption of 0 under P 's public key according to the encryption scheme of [2].
2. Now, using the protocol from §3.2, each party P_i modifies $(1 - 2\epsilon)$ copies of each of the L counters that he/she wishes to update.

$\text{Transfer}(x, y)$: For each counter, T runs an efficient two-party computation [21, 10, 14] with P to compute the simple majority value present in these counters and returns the value to P . T 's input x is all the D encrypted values of a counter and P 's input y is the secret key of the homomorphic encryption scheme.

Protocol Correctness, Security and Communication Complexity. We have $R < \frac{1}{4\epsilon}$. We note that since each party modifies at least $(1 - 2\epsilon)D$ copies of every counter, after the first modification to a counter, at least $(1 - 2\epsilon)D$ copies of every counter hold the correct value. Now, after the second modification to the counter at least $(1 - 4\epsilon)D$ copies of the counter hold the correct value and so on. Since $R < \frac{1}{4\epsilon}$, after all parties have modified the counters, a majority of the counters still hold the correct value of the counter and hence when evaluating the simple majority of the value held in the counter (via the two-party computation protocol between P and T), the output obtained will be correct.

Theorem 5 $\text{Counter}(N, L)$ is L -private with respect to T according to the definition given in §A.1.

Proof. This theorem follows from Theorem 2, that guarantees that T cannot tell which of the L counters were modified. \square

Theorem 6 $\text{Transfer}(x, y)$ is private with respect to P and T according to the definition given in §A.1.

Proof. This theorem follows from the security of the two-party computation protocol (that guarantees that P learns only the output value of the counter) and that T does not learn the value of the counter. \square

The communication complexity of $\text{Counter}(N, L)$ for every user is the same as that in §3.2, that is $\mathcal{O}(\sqrt{NL}^{1+\alpha} \text{poly-log } N)$ for some constant $0 < \alpha < 1$. Since, the server can run an efficient two-party computation protocol with P , with inputs of size $\mathcal{O}(\text{poly-log } N)$ to compute the majority value of each counter, the communication complexity of $\text{Transfer}(x, y)$ is $\mathcal{O}(N \text{poly} D)$ which is $\mathcal{O}(N \text{poly-log } N)$.

B.2 Protocol Using Unbalanced Expanders

We present a protocol for batch counters in which there is no restriction on the number of parties R . $\text{Counter}(N, L)$ is describe below.

1. Encode every counter as D different counters. To decode, compute the sum of all these counters to obtain the value of the counter. Initially these D counters all hold an encryption of 0 under P 's public key according to [2].
2. To modify a counter, each party P_i modifies any 1 copy of each of the L counters as follows:
 - (a) Following work from Ishai et al. [13], it follows that one can modify 1 out of the D bits that encode every bit in a set of L_{\max} bits of the N -bit database by modifying at most 1 bit in each of the M virtual databases. Using the explicit unbalanced expander from Guruswami et al. [11], one can obtain the same communication complexity as in the protocol described in §3.2. Here, we note that we do not require the expander to be lossless, but only that it is unbalanced.
 - (b) P_i modifies one of the D different counters encoding every counter that P_i wishes to modify. This modifies the value of the encoding as well (as the encoding is simply a sum of all the counters).
 - (c) Again, in order to use the protocol with different values of L_{\max} , we store more copies of each bit and use the same solution as described earlier in §3.2.

$\text{Transfer}(x, y)$: For each counter, the server (using the additive homomorphism property of [2]) computes the decoding of the counter and returns the encrypted value of the counter to P .

Protocol Correctness, Security and Communication Complexity. The correctness and privacy of the protocol is easy to show. The communication complexity of each user is the same as that in the protocol for database modification.