Private Information Storage

(Extended Abstract)

Rafail Ostrovsky*
Bellcore

Victor Shoup†
Bellcore, IBM

May 1997

Abstract

This paper deals with the problem of efficiently and privately storing and retrieving information that is distributively maintained in several databases that do not communicate with one another. The goal is to minimize the communication complexity while maintaining privacy (i.e., so that individual databases do not get any information about the data or the nature of the users' queries). The question of private retrieval from multiple databases was introduced in a very nice paper of Chor, Goldreich, Kushilevitz and Sudan (FOCS '95), but the question whether it is possible to perform *both reading and writing* in a communication-efficient manner remained open. In this paper, we answer this question in the affirmative, and show that efficient read/write schemes are indeed possible. In fact, we show a general information-theoretic reduction from reading and writing to any read-only scheme that preserves the communication complexity of the read scheme to within a poly-logarithmic factor (in the size of the database), thus establishing that read/write schemes could be implemented as efficiently (up to poly-log factors) as read-only schemes. Additionally, we consider the question of both reading and writing in the computational security setting.

*Bell Communications Research, MCC-1C365B, Morristown, NJ 07960-6438, USA. Email: `rafail@bellcore.com`

†Work done while at Bellcore. Current addr: IBM Research–Zurich, Säumerstr. 4, 8803 Rüschlikon, Switzerland. Email: `sho@zurich.ibm.com`

1 Introduction

The setting

In this paper, we address the issue of privacy for the database maintenance problem. However, before we address this issue, let us first state what the setting is without the privacy considerations. Suppose a collection of "users" wish to maintain a "database", where the "database" (in its simplest form) is just an n-bit array. The users wish to be able to execute both "read(address)" and "write(address,value)" operations. In this setting, we measure the amount of communication which must be sent between some "user" and the "database". If we do not care about privacy, then both read and write operations clearly take $O(\log n)$ bits of communication only.

Now, let us address privacy. Suppose the "users" do not trust the database administrator and wish to conceal from the "database" all their data and the nature of their queries. That is, they wish to conceal from the database administrator not only the contents of the data, but also the access pattern, i.e., which particular addresses of the database are being read from or written to. This could be done in two different settings, as explained below:

SINGLE DATABASE [G-87, Ost-90, GO-96]: In a single database scheme there is only one database from which users wish to hide information, while still using it to store data. In this setting, all participating users agree on an *encryption/decryption* scheme which they use to encrypt every value stored in the database. In addition, however, they must also hide from the database which particular addresses are being accessed. One simple way out of this is to read and re-encrypt the entire database for every read/write operation—then clearly the database has no clue as to which particular location is really read and/or modified. What are the drawbacks of this scheme? We list some of them: (1) the communication complexity is huge (proportional to the size of the database times the security parameter); (2) users must maintain small amount of secret information (private key); and (3) the security achieved is only computational. The first drawback (i.e., communication complexity) was resolved in [Ost-90, GO-96], where they show how to construct a communication-efficient scheme, where communication for each read or write operation is poly-logarithmic in the size of the database times the security parameter of the underlying cryptographic scheme. Their scheme, however, still suffers from the second and third drawbacks. A suggestion made in the paper of Chor, Goldreich, Kushilevitz

and Sudan [CGKS-95] is to use multiple databases that do not talk to one another:

MULTIPLE DATABASES [CGKS-95]: The second way to proceed is to distribute the database, so that the composite database is implemented as a number of non-communicating, constituent databases. The advantage of this setting is that the users need not keep any state information, and the security achieved is information-theoretic. This is the approach we follow in this paper. The case of reading-only schemes was considered in [CGKS-95, Amb-96, CG-97]; here we address the question of both writing and reading.

A TRIVIAL EXAMPLE: Before we proceed let us present a simple (but communication-inefficient) example of a private read/write scheme on n bits, to demonstrate that information-theoretic security is indeed easy to attain, as long as we do not care about communication complexity. The idea is to use the most rudimentary form of "secret sharing," representing each bit of the database as two random bits, or "shares," and placing each share in a separate constituent database. We then use the idea of scanning the entire database, similar to the single-database approach of hiding the access pattern. Here are the details. The two constituent databases physically maintain bit vectors D_1 and D_2, respectively, both of length n. At any point in time, the composite database logically maintains the bit vector $D = D_1 \oplus D_2$, i.e., the bit-wise exclusive-or of D_1 and D_2. To read a single bit, the user simply asks for copies of D_1 and D_2. To write a single bit, the user does the following: (1) asks for copies of D_1 and D_2; (2) computes D; (3) changes the desired bit of D; (4) picks a random bit vector D_1 and computes D_2, (subject to $D = D_1 \oplus D_2$); (5) gives the new D_1 and D_2 back to the respective databases. The communication complexity of this scheme is $\Theta(n)$. One of the objectives of this paper is to present multiple-database (information-theoretically) private read/write schemes with much lower communication complexity. Before we state our main results, let us make the problem more precise.

Model and Problem statement

We suppose that the *composite database* "logically" maintains a bit vector of length n. The composite database is implemented as some number k of non-communicating *constituent databases*, each of which is responsible for "physically" maintaining its own bit vector. These latter bit vectors need

bear no particular relation with each other or with that of the composite database.

A user may perform a read operation, $\mathsf{read}(i)$, for a given address $0 \leq i < n$, obtaining the bit stored at location i, or a write operation $\mathsf{write}(i, b)$, for a given address $0 \leq i < n$ and bit $b \in \{0, 1\}$, setting location i to b. When a user performs a read or write operation on the composite database, she engages in a protocol with the k constituent databases. After engaging in one or more such protocols, each constituent database has its own *view*, consisting of the messages it has received from the user(s), along with any random bits it may have generated while executing the protocol.

Formally, we model the above setting in a way similar to the multi-prover interactive proof setting of [BGKW-88]. That is, we model k constituent databases and the user as $k+1$ interactive Turing machines $DB_1, \ldots, DB_k$ and U of [GMR-85] defined as follows:

- Let $DB_1, DB_2, \ldots, DB_k$ are interactive Turing machines. Each DB_i has distinct read-only input tape, a work-tape, and two communication tapes: a write-only communication tape for sending messages to U and read-only communication tape for received messages from U.

- U is an interactive probabilistic polynomial time Turing machine with input tape, a work tape, a read-once random tape, a write-only output tape, and k read-only and k write-only communication tapes, each read/write communication tape shared with each each D_i.

The *protocol* for k databases consists of the the probabilistic polynomial time *initialization* algorithm *Init*, $DB_1, \ldots, DB_k$ and U. *Init* takes as an input an n-bit vector (the initial value of the database) and outputs k-tuple of strings to be used as inputs for each corresponding DB_i. The syntax of the input for U is as follows: it consists of sequence of "read(i)" or "write(v, i)" *instructions*, where $1 \leq i \leq n$ and $v \in \{0, 1\}$. We insist that U operates as follows:

- it reads the next single instruction from its input tape;

- it then engages in the interactive protocol with all the databases of polynomial (in n) number of *rounds*, where a round consists of $2k$ messages, consisting of k "questions" to all the databases and their k "responses";

- U must then output a single bit on its output tape, (which in case of a write instruction could be arbitrary, but in case of read instruction corresponds to the bit "retrieved" from the database);
- U must then completely erase its work tape, and only then read the next instruction on its input tape.

Correctness means that for any n-bit string which corresponds to the initial value of the database, and for any coin-flips of the initialization process, and for any coin-flips of U, the output produced by the user is consistent with the usual read/write sequence.

We define the *view* of the constituent database DB_i to be the value of its input tape and the sequence of messages written on its read/write communication tapes.

Privacy means that for every length n and any initial value of the *Init* algorithm (of length n) and for any sequence of *instructions* (i.e., read or write operations of U) the probability distribution (over the coin-tosses of the initialization algorithm and coin-tosses of U) of each constituent database's view is independent of the initialization vector (of the *Init* algorithm,) of addresses in the read and write instructions, and of the value of the data in the write operations.

We will call such a scheme a *k-database private read/write scheme (on n bits)*. For a given scheme, its *communication complexity* is the total number of bits transmitted during one execution of either the read or write protocols. Throughout this paper, communication complexity will be expressed as a function $R(k, n)$ of k and n, and measured in terms of worst-case (as opposed to amortized or average-case) behavior.

VARIANTS OF THE PROBLEM: There are several variants of the above setting:

- **Number of Constituent Databases.** In [CGKS-95, Amb-96, CG-97] as well as in the current paper, we examine the dependence between the number of constituent databases k and the total communication complexity, between the user and all the databases.
- **Identical vs. Distinct Databases.** In the original model of [CGKS-95], there was no distinction made between a constituent database and a global database. That is, since the objective there was to hide the reading pattern only, each constituent database simply held an identical copy of the actual database. In case of writing,

where we wish to hide the data as well, we allow constituent databases to hold "shares" [S-79] of the actual database. Thus, in our setting we allow (and utilize the fact) that constituent databases need not be identical.

- **Rounds.** In the setting of [CGKS-95, Amb-96, CG-97], a single round of interaction between the user and the constituent databases (where the user sends a single message to each constituent database and gets an answer back) is used. In our solutions, we use multiple rounds.

- **Privacy.** The definition of privacy could be relaxed to the computational setting. In this setting, we model constituent databases as probabilistic polynomial time Turing machines, and rely on complexity assumptions. Hence, we only require that the views of the constituent databases will be *computationally indistinguishable* instead of identical.

- **Isolated vs. Active Security [G-96]:** In the model described above, constituent databases are not allowed to communicate with one another and are only allowed to communicate with the user. We call such a model an "isolated" model. A stronger type of an adversary, suggested by Oded Goldreich [G-96], is the one where databases are also allowed to "hire users" and query all other databases, pretending to be legitimate users. Indeed, in such an "active" model, any constituent database can find out all the data of the actual database. One can still insist, however, that if afterwards, a "legitimate" user reads a bit, then it is still remains hidden as to which address she read, and if the constituent databases do not again re-read the database (by pretending to be users again), then the subsequent write operation should also remain hidden (i.e., both the address and the value of the updated data should remain hidden). We remark that the private reading schemes of [CGKS-95, Amb-96] (where all constituent databases are identical) are resistant against this stronger type of an adversary. However, if constituent databases keep different "shares" of the actual database, and for the case of writing, more care must be exercised against this stronger adversary.

- **Computational Efficiency.** In addition to communication complexity, we also consider computational complexity needed to execute the protocol by both the user and the constituent databases.

Our Results

We consider both information-theoretic and computational setting.

INFORMATION-THEORETIC SETTING.

As mentioned above, our work builds on that of [CGKS-95]. They consider a restricted version of our problem where users are only allowed to perform read operations (further efficiency improvement was suggested by [Amb-96].) All of their schemes have the additional property that each constituent database simply maintains identical copies of the composite database. We will call such a scheme a *private read scheme*.

We give a general reduction, showing how to build private read/write schemes from private read schemes, with only a modest increase in both the number of databases required and in communication complexity.

Theorem 1 For any $k \geq 2$, if there is a k-database private read scheme on n bits with communication complexity $R(k, n)$, then there is a $(k+1)$-database private read/write scheme on n bits with communication complexity

$$O(R(k, nk) \cdot k \cdot (\log n)^3).$$

Each constituent database maintains $O(nk)$ bits.

Theorem 2 For any $k \geq 2$, if there is a k-database private read scheme on n bits with communication complexity $R(k, n)$, then there is a $2k$-database private read/write scheme on n bits with communication complexity

$$O(R(k, n) \cdot (\log n)^3).$$

Each constituent database maintains $O(n)$ bits.

REMARKS:

- For both theorems above, the communication complexity bounds are independent of the order of reading and writing. In fact, it is strait-forward to hide even whether we perform reading or writing, by always performing both read and write and just re-writing the same value in case of an actual read.

- Notice that the difference when going from $(k+1)$ databases to $(2k)$ databases in our two results is a k multiplicative factor in the communication complexity and the size of each constituent database. For a number of databases between $k+1$ and $2k$ simple tradeoffs can be achieved.
- For reading schemes where privately reading a contiguous block of l bits has smaller communication complexity them reading l single bits, our reductions are more efficient, and the poly-log exponent can be further reduced. Moreover, the savings on reading schemes for blocks could be used to achieve savings in our schemes when reading/writing blocks of bits.
- Using secret sharing [S-79], our results could be adopted to the case where coalitions of databases are allowed to communicate. Moreover, we can adopt our solution to the malicious case, in the sense of certifying (with overwhelming probability) if data has been tampered with, as long as there exists at least one non-corrupted database.
- Our reductions "preserve" active security, in the sense that if the underlying read-only scheme is secure against such an active attack (which *does* hold in the case of [CGKS-95, Amb-96]) then our resulting read/write scheme is also secure in this stronger active security model.
- In case of writing, one can consider the number of bits not to be fixed, but grow as a function of time. In this case, we get poly-log overhead results as well. For example, we get an analog of information-theoretically secure Oblivious RAM simulation (see [G-87, Ost-90, GO-96]), where t steps of the original program can be simulated in an oblivious manner using $O(R(k,t) \cdot (\log t)^3)$ overhead per step using k databases. Moreover, all the solutions for the Oblivious RAM model of [GO-96] are *amortized*, where as all our solutions are not.

Combining our general reductions with the private read schemes of [CGKS-95] and of [Amb-96], we obtain the following corollary:

Corollary 3 To store n bits of data, we have

- a three-database private read/write scheme with active security and communication complexity $O(n^{1/3} \cdot (\log n)^3)$;

- for all constants $k \geq 2$, a $(k+1)$-database private read/write scheme with active security and communication complexity $O(n^{1/(2k-1)} \cdot (\log n)^3)$;
- an $O(\log n)$-database private read/write scheme with active security and communication complexity $O((\log n)^5 \log\log n)$.

Notice that we show a general reduction from reading and writing to reading, with only poly-logarithmic overhead. Hence, due to the general nature of our reduction, any improvement in the efficiency of reading schemes would yield a more efficient reading and writing scheme as well. We remark that prior to the current paper no trivial (i.e., sub-linear) bounds for private information storage were known.

COMPUTATIONAL SETTING.

Chor and Gilboa [CG-97] show how in the computational setting, one can keep two *identical* databases and in one-round perform private reading with $O(n^\epsilon)$ communication complexity.

We consider a weaker model, where we allow constituent databases to keep *different* data and allow a multi-round scheme. In this setting, Gene Itkis (private communication by Goldreich [G-96]) has shown a four-database scheme which achieves poly-logarithmic *amortized* overhead for both reading and writing. We extend his result, and show how to achieve poly-logarithmic overhead *without* amortization and with just two constituent databases:

Theorem 4 Suppose one-way trapdoor permutations exist, and let g be a security parameter. Then to store n bits of data, we have a two-database computationally-private read/write scheme with active security and communication complexity $O(g^{O(1)} \cdot (\log n)^{O(1)})$.

In fact, in order to show the above result, we exhibit how to make all the results in the Oblivious RAM simulation paper of [GO-96] non-amortized, which we believe is of interest in its own right.

Comparison with Previous Work

The general approach of distributing information to maintain privacy has been used in many situations, including the previous work on reading from

a distributed database [CGKS-95, Amb-96, CG-97], the U. S. Government's Clipper Chip proposal [U.S.-93], Micali's fair-cryptosystems [M-92], secret-sharing schemes [S-79], and instance-hiding schemes [RAD-78, AFK-89, BF-90, BFKL-90].

Closely related to the private storage problem is the *oblivious RAM simulation problem*, studied in [G-87, Ost-90, GO-96], and indeed the techniques we employ here build on those used to solve the oblivious RAM simulation problem. So, let us mention that setting here. The problem is to simulate a random-access machine (RAM) with another so that the memory contents and access patterns of the latter machine are independent of the input. In the oblivious RAM simulation problem, the central processing unit (CPU) plays the role of the user, and the main memory plays the role of the database. It is perhaps worth pointing out the technical differences between these two problems: (1) unlike the main memory of the RAM, the databases are distributed; (2) for our general reduction, we require information-theoretic privacy, whereas in the oblivious RAM simulation problem, complexity-theoretic assumptions are used and only computational privacy is achieved (either that, or access to a random oracle is required); (3) unlike the CPU, the user does not maintain any state; (4) whereas the techniques of [Ost-90, GO-96] yield bounds on the amortized communication complexity, our techniques yield worst-case bounds.

The problem of performing private database queries with multiple databases that do not interact with one other was studied in two other settings: in instance hiding schemes of [BF-90, BFKL-90] and on private database queries of [CGKS-95, Amb-96, CG-97]. In both of these settings, the contents of the database is static, and each constituent database maintains exact copies of the database. Our work shows how to support dynamic databases, while maintaining privacy. One technical difference is that the constituent databases in our schemes must maintain somewhat larger amounts of non-identical data.

A TECHNICAL REMARK: why can't we do it much easier?

At a first glance, the following (incorrect) argument seem to give a much stronger result in a trivial manner. Take the Oblivious RAM solution of [Ost-90, GO-96] and implement it using two databases: one to represent the RAM memory and another one to represent the CPU memory. This seems to trivially give us a solution with poly-logarithmic amortized overhead, using just two databases. However, this does not work. The reason

is that [Ost-90, GO-96] use in an essential way pseudo-random functions of [GGM-86] in order to implement a random oracle in their construction. Thus, the guarantees are only computational by the nature of the [Ost-90, GO-96] construction. Moreover, even if we are willing to opt for the computational security only and are willing to settle for an amortized solution, this does not work if we require Active Security (see "variants of the problem" section of the introduction). The main technical contribution of our information theoretic reduction is to show that if we are given an information-theoretically secure reading scheme, then we can implement information-theoretically secure reading and writing in an efficient manner without any need for pseudo-random functions or random oracles. Moreover, in the computational case, we show how Active Security can be achieved, while still maintaining only two databases.

Overview of the paper

In §2 we discuss some elementary solutions as a means of illustrating several techniques used in our general reduction. These elementary solutions are asymptotically inferior to our general reduction (but are useful to explain our general reduction). Then in §3 we prove Theorem 2, and show how to modify this proof to obtain Theorem 1. Finally, in §4 we discuss the computational setting and show how to obtain Theorem 4.

2 Elementary methods

Elementary linear solution—another way to look at it

Recall that the method presented in the introduction uses two databases and represents each bit as the exclusive-or of two bits in two different databases. This allows users to hide from each database the value of each bit that is being stored. Thus, one can think of this operation as "encrypting" data so that each database sees the access pattern (just scanning the database from left to right) but does not see the actual values of the bits, and hence does not know which bit was re-written. This technique allows us to implement *semi-private writes*, where the access pattern is visible to each constituent database but the value being written as not visible.

Separating writing from reading

In this subsection we show how the above elementary scheme for writing could be augmented to have efficient reading with four databases. Recall that [CGKS-95] show that with two databases *which contain identical n bits of data*, it is possible to privately read a bit with $O(n^{1/3})$ communication complexity. Notice, however, that in the two-database scheme presented in the previous subsection, the two databases contain different data, since every bit is represented as the exclusive-or of two corresponding bits from two databases. The idea is very simple: using four databases, maintain two identical copies of each database-pair of the previous subsection. Now, writing still takes $O(n)$ steps, since we still must re-write the entire database (and in fact maintain two copies), but reading can be done in $O(n^{1/3})$ steps just by reading the appropriate bit from both identical pairs of databases using twice the reading scheme of [CGKS-95] and then just computing the exclusive-or of these two bits.

An Elementary Sub-linear Scheme

In this subsection, we present an elementary 8-database private read/write scheme with communication complexity $O(n^{1/2})$ for writing and $O(n^{1/3})$ for reading. The idea will be an extension of the solution of the previous section, but with more efficient writing. Here is the overall strategy. We assume we have 4 databases that already support private reading but non-private writing. Using these 4 databases, we show how to implement private writing as well. Then we replace each of the 4 constituent databases with 2 identical copies of ordinary databases, and apply a result of

Assume we have 4 databases, D_{st} $(s,t \in \{0,1\})$, each of which supports private reading and non-private writing. As in the simple two-database scheme discussed in the introduction, we will split each bit of the database in shares, but this time four shares. That is, at any point in time, each bit in the composite database is represented as the exclusive-or of the four corresponding bits in the constituent databases. We now show how to privately *toggle* a particular bit in the database. Let $d = \lceil n^{1/2} \rceil$. For any address $i \in \{0, \ldots, n-1\}$, we can write

$$i = jd + k \qquad (0 \leq j < d,\ 0 \leq k < d).$$

To toggle bit i, the user generates two random bit-vectors $v, w \in \{0,1\}^d$. To

each component database D_{st}, the user sends vectors $v', w' \in \{0,1\}^d$ where

$$v'_l = \begin{cases} v_l & \text{if } l \neq j, \\ v_l \oplus s & \text{if } l = j, \end{cases}$$

for $0 \leq l < d$, and

$$w'_m = \begin{cases} w_m & \text{if } m \neq k, \\ w_m \oplus t & \text{if } m = k, \end{cases}$$

for $0 \leq m < d$. Upon receiving vectors v', w', the database D_{st} toggles all bits whose address is of the form $ld + m$, where $v'_l = 1$ and $w'_m = 1$.

Consider the effect of this operation on an arbitrary bit in the database whose address is $ld + m$:

$$\begin{aligned} D(ld+m) &= D_{00}(ld+m) \\ &\oplus D_{01}(ld+m) \oplus D_{10}(ld+m) \oplus D_{11}(ld+m). \end{aligned} \tag{1}$$

Case 1. If $l = j$ and $m = k$, then exactly 1 term in (1) are toggled, effectively toggling the sum.
Case 2. If $l \neq j$ and $m \neq k$, then either none or all of the terms in (1) are toggled, leaving the sum unchanged.
Case 3. If $l \neq j$ or $m \neq k$, but not both, then either 0 or 2 of the terms in (1) are toggled, again leaving the sum unchanged.

From the above discussion, it is clear that this operation has the effect of toggling bit i in the database. Moreover, each constituent database receives two random bit-vectors that are independent of i. Thus, a write operation (a read followed by a toggle) is private. The communication complexity is $O(n^{1/2})$. To complete the discussion, we observe that each of the four constituent databases, which support private reading, can be implemented using a pair of identical, ordinary databases. Using a result of [CGKS-95], a private query can then be implemented with communication complexity $O(n^{1/3})$. Putting all of this together, we get an 8-database scheme where private read operations have a communication complexity of $O(n^{1/3})$, and private write operations have a communication complexity of $O(n^{1/2})$.

Remark: Michael Fischer independently discovered this 8-database method also (communicated by Oded Goldreich [G-96].)

Remark: The above method could be naturally extended to higher dimensions, similar to [CGKS-95] approach for constant k. However, our general reductions in the next two sections yield asymptotically better results, and thus we do not present this simple extension.

3 Proof of Theorem 2

In this section, we present the proof of theorem 2. That is, we assume that we have a k-database private read scheme on n-bits with communication complexity $R(k,n)$ and we show how to construct a $2k$-database private read/write scheme with communication complexity $O(R(k,n)\cdot(\log n)^3)$. Each constituent database will hold $O(n)$ bits.

The reduction will proceed in two steps. First, we will assume that we have a scheme that supports private reading and *semi-private* writing — that is, where reading can be done in a completely private way, but writing can be done while hiding (from the databases) the value that is begin written but not the address in the memory where it is written to. Second, we will show how to implement the private read/semi-private write scheme using just a private read scheme.

Part 1: using a private read/semi-private write scheme

We make use of a variant of the memory-hierarchy idea used in [Ost-90, GO-96] for the oblivious RAM simulation problem. However, there are several obstacles that we must overcome:

- in the oblivious RAM simulation solution, the user (i.e. CPU) accesses a random oracle (or pseudo-random functions), whereas in our case the user is allowed to flip coins, but she does not has access to a random oracle;
- in the oblivious RAM simulation solution, the user has local storage, whereas in our case the user is completely memoryless from one read/write operation to the next;
- the solution presented in [Ost-90, GO-96] is amortized while here we do not allow any amortization.

Offsetting these difficulties is our assumption that we already have a private read/semi-private write scheme at our disposal: all we have to do is hide the access pattern of the write operations.

The data structure

The data structure is a kind of "memory hierarchy." There are $m+2$ "levels," where $m = \lfloor \log_2(n/\log n) \rfloor$; the levels are numbered $0, 1, \ldots, m+1$. Intuitively, data stored at lower-numbered levels is more recent than at

higher-numbered levels; as data ages, it is gradually moved from lower-numbered levels to higher-numbered levels.

We will need a counter *ctr* that keeps track of the number of write operations that have been performed. This counter is maintained modulo 2^{m+1}, and is initially 0.

For $0 \leq l \leq m$, level l is structured as follows. There are three vectors of length 2^l, each component of which contains an address/data pair (i, b), where $0 \leq i < n$, and $b \in \{0, 1\}$.

At any instant, one of the vectors is assigned the role of a "buffer," one the role of "primary data vector," and the other the role of "secondary data vector." The roles of the vectors will change over time; we need to maintain a constant amount of state information at this level to keep track of the current assignment.

For each of the three vectors we maintain a length variable, which ranges between 0 and 2^l. These variables denote the current effective length of the corresponding vector. Initially, these length variables are 0.

The components of the three vectors are always sorted in order of increasing addresses.

For each of the two data vectors, we will need two pointer variables, which range between 0 and 2^l. These are initially zero. (These variables are needed to move data from one level to the next, as will be explained below.)

The last level, level $m + 1$, is simply a bit vector of length n. This is initialized to zero (or to any desired default initial value).

Performing a private read

To read the contents of address i, we do the following. For $l = 0, \ldots, m$, we perform a binary search first on the primary data vector at level l, and then on the secondary data vector in at level l. The first place that we find (i, b), we take b as the current value stored at location i. If we do not find address i at levels $0, \ldots, m$, we then we simply obtain the current value by reading the bit vector at level $m + 1$.

Since we are assuming the underlying database supports private reading, we only need to ensure that the number of read operations we perform is always the same. This is easily accomplished by performing an equal number of "dummy" reads for all the levels, even if we already have the value.

Performing a private write

Suppose we want to write b into location i.

- First, we insert (i, b) into the buffer at level 0 (as will be seen, this buffer is always empty just before the insertion).
- Now, for $l = 0, 1, \ldots, m-1$, we do the following:

    - If $\lfloor ctr/2^l \rfloor$ is odd, we perform two steps of the merge-sort algorithm on the data vectors. With each merge-sort step, we do the following. Using the pointer variables as indices into the two data vectors, we compare the two corresponding addresses. If the addresses are different, we copy the corresponding component (i, b) of the smaller address to the next level, and increment the corresponding pointer variable. If the addresses are identical, we copy the component (i, b) from the primary data vector into the next level, and increment both pointer variables (this gives precedence to the data in the primary data vector).
    - To copy (i, b) to level $l+1$, we simply insert (i, b) into the buffer vector at level $l+1$ at the next available slot in the buffer (as will be seen, this buffer will never overflow).

- Level m requires special treatment. If $\lfloor ctr/2^m \rfloor$ is odd, then we update $\lceil n/2^m \rceil$ (which is $O(\log n)$) successive entries in the vector at level $m+1$. This is done using a similar merge-sort technique as above.
- As the data vectors at each level may not be completely full, we have to perform an appropriate number of "dummy reads" and also "dummy writes" to the appropriate locations.
- After performing the above steps, we increment ctr modulo 2^{m+1}.
- Now we go back through levels $l = 0, 1, \ldots, m$, and at each level, if $ctr \equiv 0 \bmod 2^l$, we do the following:

    If $ctr/2^l$ is odd,

        switch the roles of the buffer and primary data vectors, and clear the pointer variables;

otherwise (i.e. $ctr/2^l$ is even),

make the primary and secondary data vectors empty (by clearing the corresponding length variables) and then switch the roles of the buffer and secondary data vectors.

An illustration

Before proceeding with the analysis, we illustrate the data-movement at a level l, where $0 \leq l < m$. Starting with $ctr = 0$, we divide sequences write operations into *periods* consisting of 2^l write operations, and *cycles* consisting of two periods. We illustrate the first four periods.

Period 0. During period 0, the buffer gets filled with some data (at most 2^l items), call it A. So during period 0, the vectors at level l look like this:

buffer: A	primary: –	secondary: –

At the end of period 0, we swap the buffer and the primary vector pointers:

primary: A	buffer: –	secondary: –

Period 1. During period 1, the buffer gets filled with data, call it B. Also during this period, we merge and copy to the next level the contents of the primary and secondary vectors. Right now, the secondary vector is empty so this just has the effect copying A to the next level. During period 1, the situation looks like this:

primary: A	buffer: B	secondary: –

At the end of period 1, we clear the primary and secondary vectors, and then swap the buffer and secondary vector pointers:

primary: –	secondary: B	buffer: –

Period 2. During period 2, the buffer is filled with data, call it C:

primary: –	secondary: B	buffer: C

At the end of period 2, we swap the buffer and primary vector pointers:

buffer: –	secondary: B	primary: C

Period 3. During period 3, the buffer is filled with data, call it D. Also during this period, B and C are merged and copied to the next level. Notice that merge-sort is running at "double speed," so there is enough time to copy all of B and C during this period. The situation looks like this:

buffer: D	secondary: B	primary: C

At the end of period 3, we clear the secondary and primary vectors, and swap the buffer and the secondary vector pointers:

secondary: D	buffer: –	primary: –

Finally, note that at each *cycle* of length 2^{l+1} the total "flow of data" into level l and out of level l is "balanced". We now proceed with the analysis.

Analysis

We first make some observations about the movement of data from one level to the next. Consider level l, where $0 \leq l \leq m$. The actions performed at this level cycle are cyclical, repeating themselves once every 2^{l+1} write operations. Let us say that a cycle begins when $ctr \equiv 0 \bmod 2^{l+1}$. As above, we divide each cycle into a first period (when $\lfloor ctr/2^l \rfloor$ is even) and a second period (when $\lfloor ctr/2^l \rfloor$ is odd).

We make several claims:

(1) At the beginning of a cycle, the buffer and primary data vectors at this level are empty.

(2) During one cycle, at most 2^{l+1} address/data pairs are copied to level $l+1$; in particular, these address/data pairs are the merged contents of the two data vectors at the start of the second period of the cycle.

(3) The buffer at this level never overflows.

Claim (1) is certainly true at the beginning of execution, when $ctr = 0$. Moreover, at the end of every cycle, we clear the two data vectors and then switch the roles of the buffer and secondary data vectors, so at this point the buffer and primary data vectors are again empty. That proves (1).

We now prove (2). During the first period of the cycle, the buffer gets filled with data from the previous level. As we will argue below, the buffer does not overflow, but for now, assume that any insertion into the buffer that would cause an overflow is simply discarded. At the end of the first period of the cycle, the buffer has been filled (perhaps only partially) and we switch the roles of the buffer and primary data vector. So now the buffer is empty, and the primary data vector contains the contents of what was the buffer. During the second period of this cycle, the buffer again gets filled with data from the previous level. Also during the second period of the cycle, the entire contents of both data vectors is merged and copied to the next level. To see that everything is copied, suppose first that $l < m$. Note that both data vectors together contain at most 2^{l+1} items, and every write operation we copy *two* items; therefore, after the 2^l write operations of the second period of the cycle, *all* of the items from the two data vectors have been copied. The case $l = m$ is also straightforward to analyze; we omit the details. That proves (2).

Claim (3) for level l follows immediately from claim (1) at level l and claim (2) it level $l - 1$.

It is also straightforward to see that the value returned by a read from location i is equal to the last value written to location i. To see this, consider what happens at the end of a cycle at level l. We make the two data vectors empty at this point, so we have to show that no data is lost. However, by claim (2), all of the contents of the two data vectors has been transferred to the next level. If $l = m$, then the bit vector has been properly updated, and there is nothing more to show. If $l < m$, then the end of the cycle at level l is either the end of the first or second period of the cycle at level $l + 1$. In the former case, the buffer at level $l + 1$ (which now contains the contents of the old data vectors at level l) becomes the primary data vector at level $l + 1$; thus, the data is available at level $l + 1$, and takes precedence over the other data at level $l + 1$ (which is in the secondary data vector). In the latter case, the buffer at level $l + 1$ becomes the secondary data vector at level $l + 1$, but the primary data vector at level $l + 1$ becomes empty; thus, the data is available at level $l + 1$, and is the *only* data available at this level.

That these algorithms attain privacy has already been argued. During a read operation, we only need to ensure that the number of reads on the underlying database is constant. This we do by performing the appropriate number of dummy reads. During a write operation, the locations that are read from and written to on the underlying database depend only on the value of ctr, provided care is taken to perform dummy writes as necessary; in

particular, whenever as we perform merge-sort steps at one level, we should always write something to successive locations of the buffer at the next level, even if the data vectors at the current level have already been exhausted.

It is straightforward to see that during a read operation, $O((\log n)^3)$ read operations are performed on the underlying database, and during a write operation, $O((\log n)^2)$ read and write operations are performed on the underlying database.

Part 2: implementing a private read/semi-private write scheme

The approach here is essentially the same used in the two-database scheme in the introduction. We utilize a k-database private read scheme. Each bit in the database is split into two random bits, or "shares," whose exclusive-or is the value of the bit. The database is partitioned into two components, and each share is stored in one component. Each component is then distributed and replicated k times, and the k-database scheme for private reads is then used for reading bits in one component. This gives rise to a $2k$-database private read/semi-private write scheme whose communication complexity is bounded by a $O((logn)^3))$ times that of the underlying k-database private read scheme. The size of each of the $2k$ databases is $O(n)$.

Finally, it is easy to see that if the underlying read scheme is secure against active adversary, then so is our reduction: the case of reading consists of a serious of calls to reading of the underlying scheme, and in the case of writing, the new values always enter in a fixed manner into the database.

4 Proof of Theorem 1

In this section, we show how to modify the proof of Theorem 2 to obtain Theorem 1. The only change is in Part 2 of the construction: implementing a private read/semi-private write scheme. The approach presented in the previous section requires $2k$ databases. Here, we show how to do this using only $k+1$ databases. Before we show the general construction, let us the three database scheme, which achieves communication complexity $O(n^{1/3} \cdot (\log n)^3)$:

- We represent each bit b in the original database as as a random three-bit vector $\{b_1, b_2, b_3\}$ subject to the constraint that $b = b_1 \oplus b_2 \oplus b_3$.

- We distribute these three bits among the three constituent databases as follows: $DB_1 \leftarrow \{b_2, b_3\}$; $DB_2 \leftarrow \{b_1, b_3\}$; $DB_3 \leftarrow \{b_1, b_2\}$. We do so for every bit.
- Notice that the original bit b is still hidden from each constituent database. On the other hand every bit b_i appears in *two* different databases, so we can use the [CGKS-95] reading scheme with $O(n^{1/3})$ communication complexity.

We now generalize this in a strait-forward manner as follows. Each bit in the private read/semi-private write database is split into $k+1$ random bits, or "shares," whose exclusive-or is the value of the bit. Number the constituent databases and shares 1 through $k+1$. Then share i is given to all databases except database i. Thus, each database contains only k of the shares, which keeps the value of the bit private. To read a bit from the database, one needs to obtain all $k+1$ shares. To obtain share i, one uses a k-database private read scheme on all databases other than database i. This gives rise to a $(k+1)$-database private read/semi-private scheme whose communication complexity is $O(k)$ times that of the underlying k-database scheme for private queries. Notice also that the sizes of each of the $k+1$ databases is $O(nk)$.

The above construction, combined with the construction in the previous section, proves Theorem 1.

5 Computational case

Our starting point is the Oblivious RAM simulation of [Ost-90, GO-96], which consists of a protected CPU and an encrypted memory. We also remark that their scheme is *tamper-proof* (for definitions, see [GO-96].) Unfortunately, their scheme is *amortized.*

Our first step is to get rid of the amortization in the manner similar to our information-theoretic reduction of section §3 as follows. First it is easy to note that except for the smallest level, the oblivious re-hash operations of [GO-96] always move data from the previous (i.e., smaller) level, which is available beforehand, and is only changed during such move. Thus, the random function can be picked in advance, and the process of oblivious re-hash can be done slowly at each level, building in advance the next level, and then just doing "copy" as follows: we keep *two* versions of each level, one completely constructed "current version" and one "under construction" version,

which is being slowly built from "current version" and the previous level "current version", analogous to the information-theoretic solution presented in section §3. When the construction of the level is complete, we just switch pointers which is the "current version" and which is the "under construction" version, thus doing the copy in the non-amortized sense, by spreading the cost uniformly. Thus, we are able to realize the Oblivious RAM and the software protection results of [GO-96] in a non-amortized sense.

Now, if we keep one database to store the contents of the CPU (which a user reads before she begins) and another database to represent the memory, then combined with the above, this gives us a two-database computationally-secure scheme with poly-logarithmic cost, but only with *Isolated* security. That is, (as pointed out by Goldreich [G-96]) in case if the database that represents the memory is malicious and is allowed to hire a "user" and accesses the value of the "CPU" database component (i.e., if we consider *active security* – see introduction), then the database that represents the memory learns the value of the key of the pseudo-random function stored in the CPU, and now the access pattern to the memory component is no longer oblivious.

In the computational setting, the way around this problem (of active security) was suggested by Gene Itkis (communicated by [G-96]): Itkis suggested to use three databases to represent the CPU, and the fourth database to represent the memory, where the key new observation (due to Itkis) is that the three DB's that hold "shares" of the state of the CPU can engage in a multi-party secure function evaluation (communicating through the user) in order to evaluate the pseudo-random function (stored in a distributed fashion among the three databases that represent the CPU) and that the evaluation of the pseudo-random function and other CPU operations can be represented as a small (poly-logarithmic times the security parameter) circuit, hence yielding a 4-database scheme with poly-logarithmic (in his case amortized) overhead with active security. Finally, we note that his solution could be reduced to just two databases, where both databases keep "shares" of the state of the CPU, and additionally one of the databases also keeps the contents of the Oblivious RAM memory. The main reason why we can allow one of the constituent databases to keep *both* the "share" of the CPU and the Oblivious RAM memory and still show that the view of this constituent database is computationally indistinguishable for all executions is that the Oblivious RAM memory component is kept in an encrypted (and tamper-resistant) form (see [GO-96]), according to a distributed (be-

tween both databases) private-key stored in the CPU. For every step of the CPU computation, both databases execute secure *two-party* function evaluation of [Y-82, GMW-87] which can be implemented based on any one-way trapdoor permutation family (again communicating through the user) in order to both update their shares and output re-encrypted value stored in a tamper-resistant way in Oblivious RAM memory component. Thus, the key observation is that the database that holds both the share of the CPU state and Oblivious RAM memory, holds both of them in an encrypted and tamper-resistant manner. Hence, it can not modify the memory component in any way, accept via the CPU access, which, in turn, requires secure function evaluation and thus can not be modified without both databases. Thus, the proof reduces to the software protection proof of [GO-96] and we are done.

Note that since we use two-party secure function evaluation of [Y-82, GMW-87], we need stronger cryptographic assumptions (i.e., the existence of trapdoor permutations instead of general one-way functions needed for four databases). We also remark that the above scheme also achieves sub-linear (in the size of the database) *computational* efficiency, where a single read/write operation requires only $g^{O(1)}$ (where g is the security parameter of the underlying trapdoor permutation) times poly-logarithmic (in the size of the database) computational steps by all the constituent databases and the user.

6 Conclusion

We have given several constructions for distributed databases that support private reading and writing — both in the information-theoretic and in the computational setting.

In the information-theoretic setting we have shown that the communication complexity of private reading and writing is within a poly-logarithmic factor of private reading. Many of the extensions in [CGKS-95] (such as privacy against coalitions and efficient, private access to blocks of data) apply to our constructions as well. One of our schemes achieves a communication complexity of $O(n^{1/3}(\log n)^3)$ using just three databases. An open question is whether there exists a two-database information-theoretic private read/write scheme with sub-linear communication complexity.

We also addressed the computational case, where assuming that databases are allowed to keep different data, and we allow a multi-round schemes, we have shown an efficient (both in communication and computational complexity) two-database read/write scheme based on any one-way trapdoor permutation family. This should be contrasted with the reading scheme of Chor and Gilboa [CG-97], where they only assume the existence of general one-way functions, keep identical databases and use only a single round of interaction. Clearly, achieving optimal performance with minimal assumptions, with the minimal number of databases and with minimal number of rounds would be interesting both for read-only and read/write computational schemes.

Acknowledgments

We thank Oded Goldreich and Eyal Kushilevitz for many helpful discussions.

References

[AFK-89] M. Abadai, J. Feigenbaum, and J. Kilian. On hiding information from an oracle. *JCSS* 39(1):21–50, 1989.

[N-89] N. Adam and J. Wortmann. Security control methods for statistical databases: a comparative study. *ACM Computing Surveys* 21(4):515–555, 1989.

[Amb-96] A. Ambainis. Upper bound on the communication complexity of private information retrieval. On-line version published in *Theory of Cryptography Library*, **http://theory.lcs.mit.edu/~tcryptol**, May 1996.

[BF-90] D. Beaver and J. Feigenbaum. Hiding instances in multi-oracle queries. In *Proc. of 7th STACS, Springer-Verlag LNCS, Vol. 415*, pp. 37–48, 1990.

[BFKL-90] D. Beaver, J. Feigenbaum, J. Kilian, and P. Rogaway. Security with low communication overhead. In *Advances in Cryptology—Proc. Crypto '90*, 1990.

[BGKW-88] M. Ben-or, S. Goldwasser, J. Kilian and A. Wigderson. Multi prover interactive proofs: How to remove intractability. STOC 88.

[B-79] G. R. Blakley. Safeguarding cryptographic keys. In *Proc. NCC AFIPS*, pp. 313–317, 1979.

[CG-97] B. Chor and N. Gilboa. Computationally Private Information Retrieval In this proceedings – *STOC '97*.

[CGKS-95] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. In *Proc. 36th Annual IEEE Symp. Foundations Comp. Sci.*, pp. 41–50, 1995.

[G-96] O. Goldreich. Personal communication, June of 1996.

[RAD-78] R. L. Rivest, L. Adleman, and M. L. Dertouzos. On data banks and privacy homomorphisms. In *Foundations of Secure Computation* (eds. R. DeMillo, D. Dobkin, A. Jones, and R. Lipton). Academic Press, 1978.

[G-87] O. Goldreich. Towards a theory of software protection and simulation by oblivious RAMs. In *Proc. 19th Annual ACM Symp. Theory Comp.*, 1987.

[GMR-85] S. Goldwasser, S. Micali and C. Rackoff, The Knowledge Complexity of Interactive Proof-Systems, *SIAM J. Comput.* 18 (1989), pp. 186-208; (also in STOC 85, pp. 291-304.)

[GMW-87] O. Goldreich, S. Micali, and A. Wigderson. "How to Play Any Mental Game". Proc. of 19th STOC, pp. 218-229, 1987.

[GO-96] O. Goldreich and R. Ostrovsky. Software protection and simulation by oblivious RAMs. *JACM*, 1996.

[GGM-86] Goldreich, O., S. Goldwasser, and S. Micali, "How To Construct Random Functions," *Journal of the Association for Computing Machinery*, Vol. 33, No. 4 (October 1986), 792-807.

[M-92] S. Micali. Fair public-key cryptosystems. In *Advances in Cryptology—Proc. Crypto '92*, pp. 113–138, 1992.

[Ost-90] R. Ostrovsky. Software protection and simulation on oblivious RAMs. M.I.T. Ph. D. Thesis in Computer Science, June 1992. Preliminary version in *Proc. 22nd Annual ACM Symp. Theory Comp.*, 1990.

[S-79] A. Shamir. How to Share a Secret. *CACM* 22:612–613, 1979.

[U.S.-93] A proposed federal information processing standard for an escrowed encryption standard. Federal Register, July 30, 1993.

[Y-82] Yao, A.C., "Theory and Applications of Trapdoor Functions", *23rd FOCS*, 1982, pp. 80-91.