

Fast Oblivious Storage

RADU SION, Stony Brook University
 PETER WILLIAMS, Stony Brook University

We introduce three practicality-critical components for Oblivious RAM (ORAM): (i) parallelism, (ii) deamortization, and (iii) fork-consistency defense against malicious adversaries.

In (i), instead of waiting for the completion of ongoing client-server transactions, client threads can now engage a server in parallel *without loss of privacy*. This critical piece is missing from existing ORAMs, which can not allow multiple clients threads to operate simultaneously without revealing intra- and inter-query correlations and thus incurring privacy leaks. And since ORAMs often require many communication rounds, this significantly and unnecessarily constrains throughput. The mechanisms introduced here eliminate this constraint, allowing overall throughput to be bound by server bandwidth only, and thus to increase by an order of magnitude.

Further, in (ii) new de-amortization techniques bring the extremely impractical worst case ORAM query cost (e.g., having to periodically wait for tens or hundreds of hours for ORAM reshuffles to complete before being able to execute the next query) in line with the average cost. Both (i) and (ii) are shown to be fundamental to any practical ORAM. Finally, (iii) extensions providing fork consistency against an actively malicious adversary are then presented.

A high performance, fully functional implementation (PD-ORAM) was designed, built and analyzed. It performs multiple queries per second on a 1TB+ database across 50ms latency links, with unamortized, bound query latencies. Based on PD-ORAM, *PrivateFS* an oblivious file system that enables access to remote storage, while keeping both the file contents and client *access patterns* secret was built and deployed on Linux as a userspace file system. For a single-CPU server throughput can exceed 300Kbps with almost optimal bandwidth utilization.

Categories and Subject Descriptors: D.0 [Software]: General; E.3 [Data Encryption]

Additional Key Words and Phrases: Access Privacy, Cloud Computing, Oblivious RAM

ACM Reference Format:

Sion, R., and Williams, P., 2012. Fast Oblivious Storage ACM Trans. Info. Syst. Sec. 15, 4, Article 39 (March 2013), 28 pages.

DOI = 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

1. INTRODUCTION

This paper introduces PD-ORAM (“Parallel De-amortized Oblivious RAM”), a collection of new techniques, applicable to a large class of existing Oblivious RAMs (ORAM), to improve their performance and practical relevance by eliminating critical bottlenecks and drawbacks. PD-ORAM is then used as the block-level building block for *PrivateFS*, a parallel, oblivious file system.

This paper is a vastly extended version of [Williams et al. 2012].

This work was supported in part by the National Science Foundation under grants 1161541, 0937833, 0845192, 0803197, 0708025.

Author’s addresses: R. Sion, Computer Science Department, Stony Brook University, Stony Brook, NY 11794-4400, sion@cs.stonybrook.edu; P. Williams, MIT Lincoln Laboratory, 244 Wood Street Lexington, MA 02421, petertw@gmail.com (work performed while at Stony Brook University).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2013 ACM 1094-9224/2013/03-ART39 \$15.00

DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

In today's "cloud-rich" environments outsourced storage architectures become more prevalent than ever. In these systems, data is often geographically distributed, stored on potentially vulnerable remote servers or transferred across untrusted networks and Remote servers are maintained by third party storage vendors. This is because the total cost of storage management is 5–10 times higher than the initial acquisition costs [Gartner, Inc. 1999]. However, most cloud storage service vendors do not provide strong assurances of data confidentiality and integrity.

Privacy guarantees of such services are at best declarative and often subject customers to unreasonable fine-print clauses – e.g., allowing the server operator (and thus malicious attackers gaining access to its systems) to use customer behavior for commercial profiling, or governmental surveillance purposes [CNN 2006].

To protect data stored in such an untrusted server model, security systems should offer users assurances of data confidentiality and access privacy.

As a first line of defense, to ensure confidentiality, all data and associated meta-data can be encrypted at the client side using non-malleable encryption, before being stored on the server. The data remains encrypted throughout its lifetime on the server and is decrypted by the client upon retrieval.

Encryption provides important privacy guarantees at low cost. It however, is only a first step as significant information is still leaked through the access pattern of encrypted data. Even when outsourced data is encrypted, the sequence of *locations* read and written to storage reveals information about the user and the data.

Consider a database management system running on top of an untrusted, encrypted file system. The file system, in satisfying requests for the transaction processor, learns semantic information about the transactions through the sequence of disk blocks accessed. If, for example, an alphabetical, encrypted keyword index is updated as an encrypted record is inserted, it can learn what keywords are present in the new record, *based only on the locations updated within the encrypted index*.

The Oblivious RAM (ORAM) [Goldreich and Ostrovsky 1996] model is one of the prominent proposed approaches for ensuring client access pattern privacy (and confidentiality) in the case of a single-owner in full control over the data. And, since the introduction of the first ORAM, approaches to increase query throughput have been sought and discovered. However, practical constructions (and thus practical oblivious file systems) have remained elusive.

One main reason for this is the fact that existing ORAMs were faced with significant drawbacks which severely limit their practicality: average time complexity lies in the hundreds to thousands of seconds per single data access even for small databases; worst case time complexity can be in the hundreds of hours; finally, malicious servers can often alter the results returned to the client without being detected.

The high average cost is a result of existing ORAMs' significant number of round-trips, typically $O(\log(n))$, per query (with some notable exceptions; see Section 2). On wide-area networks, with mid to high latency, this imposes strict limits on the query response times as well as throughput.

To mitigate this issue, one first contribution of this work is to introduce support for oblivious parallel data accesses. By supporting querying in parallel from multiple clients, PD-ORAM eliminates the effect of query response time on throughput, while presenting an opportunity for new multi-client scenarios.

The parallelization technique consists of a round-trip-optimal and wait-optimal protocol to transform a single-client "period-based" ORAM into a multi-client parallel ORAM. A "period-based" ORAM is defined in Section 5.1 as an ORAM that operates on small batches of queries, with its data structure not sensitive to the order of reads or modifications within a particular batch.

Further, to control the impractical worst-case behavior of existing ORAMs (the result of periodically required database reshuffling), we introduce a new de-amortization construction (converting an algorithm with an amortized bound into one with a worst-case bound) to process queries simultaneously with database re-shuffling. Re-shuffling is an essential and extremely costly task which, in amortized solutions, completely blocks the server for extended time periods after a certain number of queries. In PD-ORAM and other de-amortized solutions, instead of suspending queries to reorder data periodically, the server continuously re-shuffles the database in the background, in loose synchronization with querying, to guarantee minimal client latencies for both the average and worst case. Since reshuffle costs typically dwarf online query costs, and most existing ORAMs claim practicality only in the *average* (and not worst-case) query cost, this de-amortization is critical in real deployments.

The de-amortization technique operates on any “level-based” amortized cost ORAMs to bring the worst-case query cost close to the average query cost. “Level-based” ORAMs are defined in Section 5.3 as ORAMs in which querying consists of oblivious searches down a pyramid data structure. A large number of ORAM constructions fall under both these definitions, making the techniques widely applicable.

A fully functional PD-ORAM implementation is developed and benchmarked at multiple queries per second on a terabyte database, the highest throughput to date on a medium-latency link. While the mechanisms described here can be directly applied to a large number of different ORAM techniques, PD-ORAM is based on the ORAM described in [Williams et al. 2008], but with de-amortized level construction, support for parallel queries, and a new, simpler, Bloom filter (BF) construction.

PD-ORAM is then deployed and benchmarked in Linux to build *PrivateFS*, a first oblivious file system which enables clients to read and write without revealing the *access pattern* (which files are accessed, and any correlation between accesses).

The contributions of this work are multi-fold:

- A parallel secure ORAM querying protocol with its abstraction.
- Techniques for general ORAM de-amortization.
- Logarithmic-storage construction of a BF ORAM.
- Implementation and performance analysis of a parallel, de-amortized ORAM.
- Implementation and analysis of *PrivateFS*, the first oblivious file system.

2. RELATED WORK

Oblivious RAM, introduced by Goldreich and Ostrovsky [Goldreich and Ostrovsky 1996] (“GO-ORAM”) is a primitive that provides access pattern privacy to a *single client* (or software process) accessing a remote database (or RAM). GO-ORAM requires only logarithmic storage at the client; the amortized communication and computational complexities of this construction are $O(\log^3 n)$. Due to a large hidden constant factor, the GO-ORAM authors offer an alternate solution with computational complexity of $O(\log^4 n)$, that would be more efficient for all currently plausible database sizes. A discussion of GO-ORAM and recent extensions follows.

2.1. ORAM Overview

In ORAM, the database is considered to be a set of n semantically-secure encrypted blocks (with an *ORAM key* held by the client) and supported operations are $\text{read}(id)$, and $\text{write}(id, newvalue)$. In GO-ORAM for example, data is organized into $\log_2(n)$ *levels*, as a pyramid. Level i consists of up to 2^i blocks; each block is assigned to one of the 2^i buckets (originally $\log_4(n)$ levels sized 4^i ; levels sized 2^i are used here for simplicity) at this level as determined by a hash function. Due to hash collisions each bucket may contain from 0 to $k \ln n$ blocks.

GO-ORAM Reads. To obtain the value of block id , a client must perform a read query in a manner that maintains two invariants: (i) it never reveals which level the desired block is at, and (ii) it never looks twice in the same spot for the same block. To maintain (i), the client always scans a single bucket in every level, starting at the top (Level 0, 1 bucket) and working down. The hash function informs the client of the candidate bucket at each level, which the client then scans. *Once the client has found the desired block, the client still proceeds to each lower level, scanning random buckets instead of those indicated by their hash function.* For (ii), once all levels have been queried, the client re-encrypts the query result with the ORAM key and a different nonce (so it looks different to the server) and places it in the *top* level. This ensures that when it repeats a search for this block, it will locate the block immediately (in a different location), and the rest of the search pattern is randomized. The top level quickly fills up; how to dump the top level into the one below is described later.

ORAM Writes. Writes are performed identically to reads in terms of the data traversal pattern, with the exception that the new value is inserted into the top level at the end. Inserts are performed identically to writes, since no old value will be discovered in the query phase. Semantic security properties of the re-encryption function ensure the server is unable to distinguish between reads, writes, and inserts, since the access patterns are indistinguishable.

Level Overflow. Once a level is full, it is emptied into the level below. This second level is then re-encrypted and re-ordered, according to a new hash function. Thus, accesses to this new *generation* of the second level will henceforth be completely independent of any previous accesses. Each level overflows once the level above it has been emptied twice. Any re-ordering must be performed obliviously: once complete, the adversary must be unable to make any correlation between the old block locations and the new locations. A sorting network is used to re-order the blocks.

To enforce invariant (i), note also that all buckets must contain the same number of blocks. For example, if the bucket scanned at a particular level has no blocks in it, then the adversary would be able to determine that the desired block was *not* at that level. Therefore, each re-order process fills all partially empty buckets to the top with *fake* blocks. Recall that since every block is encrypted with a semantically secure encryption function, the adversary cannot distinguish between fake and real blocks.

2.2. Private Information Retrieval

Another set of existing mechanisms handle access pattern privacy (but *not data confidentiality*) in the presence of *multiple clients*. Private Information Retrieval (PIR) [Chor et al. 1995] protocols aim to allow (arbitrary, multiple) clients to retrieve information from public or private databases, without revealing to the database servers which records are retrieved.

In contrast to ORAM, PIR targets scenarios involving *multiple, unrelated clients* querying a *public database*. And, while the data in most cases is *not* confidential, PIR does not assume nor usually require client-side credentials management, in contrast to ORAM which, in principle assumes all accessing clients have access to the same credentials and the data is to be protected from the server. For this reason, encryption techniques such as used in ORAM often don't apply. Moreover, unlike ORAM, PIR typically only provides read privacy, not update privacy. Notable exceptions include [Lipmaa and Zhang 2010] and [Boneh et al. 2007].

Practicality of PIR. In initial results, Chor et.al [Chor et al. 1995] proved that in an information theoretic setting, any single-server solution requires $\Omega(n)$ bits of communication. PIR schemes with only sub-linear communication overheads, such as [Chor et al. 1995], require multiple non-communicating servers to hold replicated copies of

the data. Relaxing the information theoretic guarantee allows for single-server solutions with better complexities [Gasarch a; b].

For single-server PIR, initial results [Sion and Carbunar 2007], seemed to suggest it is extremely difficult to provide PIR implementations that are faster than the trivial case (downloading the entire public database to the client). Recent work [Olumofin and Goldberg 2011] however, has shown that multi-server schemes as well as more recent lattice-based mechanisms are faster than downloading the entire database.

PIR vs. ORAM. Unfortunately, this still leaves PIR several orders of magnitude behind ORAM mechanisms – which are faster at the expense of requiring client-side key management. For example, for the 1TB database considered here, over a dedicated 100Mbps link, in the absolute ideal case, ignoring any other constants, round-trip times and systems issues, a lattice-based single-server PIR query would still take upwards of 3 minutes (2 orders of magnitude slower than PD-ORAM).

When considering client query encoding, network round-trips, server disk seeks¹ as well as the cost of the extremely high server CPU utilization, the ORAM-PIR difference can grow beyond 3 orders of magnitude.

2.3. Oblivious Transfer (OT)

OT [Naor and Pinkas 1999; Rabin 2005; Ishai and Kushilevitz 1997] is a related cryptographic primitive in which clients request a number of database records from a server with full privacy. Further, clients are prevented from learning any additional database entries. In contrast to ORAM and PIR, OT does not aim to achieve any efficiency targets, e.g., such as sublinear communication in the database size etc.

Oblivious Transfer with Access Control. Camenish et al. [Camenisch et al. 2007] study the problem of performing k sequential oblivious transfers (OT) between a client and a server storing N values. The work makes the case that previous solutions tolerate selective failures. A selective failure occurs when the server may force the following behavior in the i th round (for any $i=1..k$): the round should fail if the client requests item j (of the N items) and succeed otherwise. The paper introduces security definitions to include the selective failure problem and then propose two protocols to solve the problem under the new definitions.

Coull et al. [Coull et al. 2009] propose an access control oblivious transfer problem. Specifically, the server wants to enforce access control policies on oblivious transfers performed on the data stored: The client should only access fields for which it has the credentials. However, the server should not learn which credentials the client has used and which items it accesses.

Note that OT does not consider (by definition) the problem of writing to the data.

2.4. Anonymity

The goal of anonymous communication is to hide not the data being accessed, but the identity of the party accessing this data. In many cases this provides suitable privacy. Anonymous communication approaches are less complete, however; efficient anonymous communication is able to provide only weak security guarantees, relying on the adversary's inability to obtain any sort of global network view. A discussion of state of the art anonymous routing approaches can be found in [Johnson et al. 2011]. Moreover, often anonymous communication is not sufficient. Combined with external knowledge, the access patterns can compromise anonymity, or reveal sensitive information in themselves.

¹Possibly catastrophically dominant, as our results on deploying the Goodrich et al. theoretically-attractive randomized shell sorting network [Goodrich 2010] have shown (see Section 2.6)

2.5. Recent ORAM Developments

Starting with [Williams and Sion 2008] researchers have sought to improve the overhead from the polylogarithmic performance of the original ORAM. Williams et al. in [Williams et al. 2008] introduced a faster ORAM variant which also features correctness guarantees, with computational complexity costs and storage overheads of only $O(\log n \log \log n)$ (amortized per-query), under the assumption of $O(\sqrt{n})$ temporary client memory. In their work, the assumed client storage is used to speed up the reshuffle process by taking advantage of the predictable nature of a merge sort on uniform random data.

Recently, a new approach to speed up ORAM was revealed by Pinkas et al. [Pinkas and Reinman 2010], showing the applicability of the Cuckoo hash construction from [Pagh and Rodler 2004]. Unfortunately, this was shown to leak access privacy information [Goodrich and Mitzenmacher 2011]. A similar, but secure, approach, allowing efficient item lookup while hiding success was then developed [Goodrich and Mitzenmacher 2011]. This approach has found continued utility in other solutions [Goodrich et al. 2011a; Kushilevitz et al. 2012].

Researchers have long recognized the utility of constructions with efficient worst cases; the first de-amortized construction followed shortly after the introduction of Oblivious RAM [Ostrovsky and Shoup 1997]. More recent solutions have also featured de-amortized constructions [Shi et al. 2011; Kushilevitz et al. 2012; Goodrich et al. 2011a; Boneh et al. 2011]. These de-amortized solutions are mostly based on the same core idea as we describe here: constructing future levels in the background, while still querying copies of current levels. One exceptional ORAM [Shi et al. 2011] is naturally un-amortized (rather than de-amortized), performing a well-defined ORAM update on every query. These previous solutions do not apply directly to a Bloom filter-based ORAM; this de-amortization requires maintaining a delete log to delay level updates until after the corresponding shuffle.

A promising recursive construction technique is introduced in [Stefanov et al. 2012] under the assumption of $O(n \log n)$ reliable client storage. Using this storage, it promises to reduce the level construction cost while requiring only a constant number of online round trips. The clear drawback here is the assumed $O(n \log n)$ client storage—enough to keep track of the positions of all items, instead of querying recursively for them as in the previously described ORAMs. The authors present the $O(n \log n)$ client storage assumption as not necessarily unreasonable, since a large block size means that the client only needs a fraction of the outsourced storage. Their alternative construction, requiring only $O(\sqrt{n})$ storage, recursively uses a $\log n$ -round-trip ORAM to store this position map, but now incurs $\log n$ round trips per query.

Another notable alternative to the constructions of $\log(n)$ round trips is found in [Wang et al. 2006]; Ding et al. build an Oblivious RAM requiring only a constant number of online round trips. The main idea is to extend Goldreich's \sqrt{n} -solution [Goldreich and Ostrovsky 1996] to store the recent query cache in client memory. The drawback is the significantly higher shuffle cost, since the entire database must be scanned once the cache is filled (e.g. after a period of \sqrt{n} queries). Moreover, the amortized shuffle costs in Oblivious RAMs easily dominate the online costs both in theory and in practice [Pinkas and Reinman 2010]. Boneh et al. examine a similar construction [Boneh et al. 2011], and formalize the security model.

Another constant-round-trip solution, based on the same core caching idea as [Wang et al. 2006] and [Boneh et al. 2011], is introduced in [Goodrich et al. 2011b], but the single-level format prevents de-amortization. The constant-round-trip claim depends on the assumption of $M = n^{1/u}$ client memory for a constant u . Nevertheless, impor-

tant notions regarding optimal use of local memory in performing efficient oblivious external-memory sorts are introduced.

In addition to providing de-amortization for the construction in [Williams et al. 2008], PD-ORAM introduces a *general* approach to de-amortization, based on rearranging levels to safely allow querying during level re-construction. Different constraints of this base ORAM require new techniques, but the core idea remains similar to related de-amortization work. This de-amortization approach applies to any ORAM with a logarithmic number of levels. This includes a Bloom filter-based ORAM [Williams et al. 2008], the original Goldreich-Ostrovsky logarithmic ORAM [Goldreich and Ostrovsky 1996], and recent cuckoo-hash based solutions [Goodrich and Mitzenmacher 2011], though we remind the reader that prior work [Ostrovsky and Shoup 1997; Goodrich et al. 2011a] has already de-amortized the latter two solutions.

A multi-client ORAM is introduced in [Goodrich et al. 2012]. Because the client state (aside from the secret key) is stored on the server, and scanned on every access, clients can take turns performing accesses. PD-ORAM takes this notion one step further: not only are the clients “stateless,” but accesses are actually performed in parallel. That is, clients begin future accesses while other clients are still processing previous ones.

In this paper, we choose to instantiate our Bloom-filter based ORAM [Williams et al. 2008]. Regarding the choice to use Bloom filter as an underlying ORAM, we are aware of two solutions with better than the Bloom filter ORAM’s $O(\log^2 n)$ complexity: one provided by Stefanov et al [Stefanov et al. 2012] and one provided by Goodrich and Mitzenmacher [Goodrich and Mitzenmacher 2011]. Neither solution, however, can provide detection of misbehavior by an actively malicious adversary, which we deem to be an important property. Moreover, the techniques we introduce apply to many other existing ORAMs.

2.6. When Reality Hits

PD-ORAM is unique in providing a tangible de-amortized implementation. De-amortized constructions do not always lead readily to an actual implementation. Instructions such as “perform a chunk of work sized x ”, or “run the shuffle in the background while querying”, are fine for establishing existence proofs, but can be devastatingly inappropriate in achieving an actual prototype due to large hidden constants, of little help when building an implementation. New techniques that address these hidden complications of de-amortization are presented in Section 5.4.

A randomized shell sorting network identified by Goodrich et al. [Goodrich 2010] was subsequently employed to construct an Oblivious RAM [Goodrich and Mitzenmacher 2011; Pinkas and Reinman 2010]. The authors of [Goodrich and Mitzenmacher 2011] claim logarithmic overhead, in a specific model calculating message complexity and based on some assumptions of block sizes and message sizes, rather than the traditional computational or communication complexity models. Since it uses the randomized shell sort to perform oblivious sorts of the data, it has a lower bound on par with existing work.

This sort procedure was used in the first design of PD-ORAM, but was found to result in too many disk seeks to make it usable on a even a medium-sized database. The randomized nature of the shell sort guarantees that the order of item access appears non sequential and random, making efficient use of rotational hard disks difficult. The amortized cost of constructing the bottom level in a terabyte database with an acceptable failure rate is in the range of hundreds of disk seeks per query, already putting the implementation outside of the targeted performance goals.

To understand this in more detail, consider the construction of the largest level of a 1 TB database. This level contains at least 0.5 TB of data. For 10 KB blocks, this

translates into 50 million blocks. The randomized shell sort makes $6k \log_2 n$ random passes across the database, incurring a total of $6kn \log_2 n$ seeks every n queries, where k influences the sort failure probability. For $n = 5 \times 10^7$, $\log_2 n = 23$, translating to $138 \times k$ disk seeks per query for the largest level alone.

For $k = 4$ as suggested in the original paper [Goodrich 2010], this amortizes to 550+ disk seeks per query. Even for high-speed, low-latency disks with 6ms seek times, this becomes at least 3.3 seconds/query (in addition to any/all other significant network and CPU overheads)!

Multiple Disks. Of course this can be mitigated by using multiple disks at additional cost, in which case, the number of seeks for each disk is reduced (linearly in the added cost) as they may occur in parallel. However, increasing performance linearly in the added cost is not surprising nor desirable, and an efficient base-case construction should be found instead.

Unfortunately, ORAM imposes a unique sorting requirement that is difficult to satisfy using the randomized shell sort. This requirement derives from the fact that, to maintain privacy, the sort must succeed with overwhelming probability. Moreover, all sorts must be indistinguishable, eliminating the possibility of retrying in the case of failure. This is because observation of a sort failure translates into an advantage at distinguishing the permutation from random, which translates into a privacy leak.

While in other applications it may suffice to repeat the sort until it succeeds, when applied to ORAM, the sort parameter k must be chosen to guarantee success with overwhelming probability. In [Goodrich 2010] the failure rate is determined experimentally; it is not obvious from the proofs what k is necessary to obtain a given acceptable error rate on database sized n . While the paper does prove that the probability of failure is negligible in k , which is sufficient for a construction existence proof, it is unclear what parameters can be chosen for a practical implementation.

It has been observed elsewhere that this sort is significantly less efficient than non-oblivious sorts. In particular, [Pinkas and Reinman 2010] opts to use a standard sort on the client for those levels that fit in memory. PD-ORAM uses the asymptotically equivalent merge sort, described in [Williams et al. 2011], that runs faster (requiring $\log_2 n$ sequential passes instead of $6k \log_2 n$ random passes), but requires $k\sqrt{n \ln n}$ blocks of local client memory.

Further, recently it was shown [Pinkas and Reinman 2010] that without a careful choice of Bloom filter parameters, the construction of [Williams et al. 2008] is susceptible to a leak via false positives in the Bloom filter construction. Avoiding this leak requires a number of hash functions proportional to the security parameter k . PD-ORAM uses this construction, choosing k to bound the false positive rate at 2^{-64} .

3. MODEL

An ORAM setup consists of three types of parties: ORAM Clients, who issue read and write queries, the limited-storage ORAM Instances, who satisfy these queries for their client while maintaining privacy, and the ORAM Server, who has plentiful storage and is willing to help the ORAM Instances, but is deemed untrustworthy (Figure 1).

The semantics used here correspond to existing ORAMs (read and write of fixed-size blocks is supported), with the addition of concurrency and multiple “Instances”. Because of the introduction of concurrency, it is also *important to offer an atomic record-level test-and-set instruction* to ORAM clients. Analogous to the CPU primitive, it updates a record and returns its old value as a single atomic operation.

ORAM Client: a party who is authorized to issue reads and writes to the ORAM Interface. Data is accessed in “blocks”, a term used to denote a fix sized record. “Block” is used instead of “word” to convey target applications broader than memory access (in-

cluding file system and database outsourcing). Block IDs are arbitrary bit sequences. There are multiple clients. Each ORAM client has access to the following interface provided by a corresponding *ORAM Instance*: $\text{read}(id): val$; $\text{write}(id, val)$; $\text{test-and-set}(id, val): val$.

Accesses are serialized in the order the ORAM server receives them, which guarantees each client sees a serialized view of its own requests. Among multiple clients, however, the ordering is not guaranteed to correspond with the time-ordering of incoming client requests.

ORAM Instance: a trusted party providing an ORAM interface to ORAM Clients. Instances are stateless but have access to the *ORAM key* (a secret shared among instances, enabling access to the ORAM) and address of the server.

The Instance-to-Server protocol details are implementation-specific (and typically optimized to the instance to minimize network traffic and the number of round trips).

ORAM Server: the untrusted party providing the storage backend, filling requests from the instance. The ORAM Instances communicate only with the server (not with each other). This corresponds to the idea that it is unreasonable to require clients to be aware and participate in every access of each other. Moreover, this reflects the notion that the adversary may have a complete view of the network.

The server is assumed throughout the bulk of this paper to be honest but curious; however, Section 5.6 details inexpensive adjustments that ensure “fork consistency” against even an actively malicious adversary. Defined in [Li et al. 2004], fork consistency acknowledges that a malicious server can present different versions of the database to different clients, e.g., by not including updates from some of the other clients, but guarantees that each “fork” view is self-consistent. Moreover, clients will detect this behavior if they ever communicate with each other, or if the server ever attempts to rejoin the views. The server is allowed to know which client issues queries and when queries are issued.

Implementations are assumed to be free of timing vulnerabilities.

Communication between ORAM Clients and Instances is assumed secured, e.g., with IPC access controls if they are on the same machine, or with SSL otherwise. Communication between ORAM Instances and Servers is also secured, e.g., with SSL.

Notation. c = number of parallel clients, i = a level within a pyramid-based ORAM (for the smallest, “top” level, $i = 0$), k = security parameter, n = database size, in blocks, \cong_c is equivalence modulo c .

4. PARALLEL QUERIES: A FIRST PASS

We now examine how to query existing ORAMs in parallel. The idea is to start with an ORAM on which it is safe to run non-repeating *unique* queries (targeting different records in the underlying database) simultaneously, and to build from this an ORAM which can also safely run *colliding* queries (targeting the same underlying data record).

For illustration, consider the GO-ORAM construction discussed above. We can augment it to handle parallelism for sets of unique queries as follows.

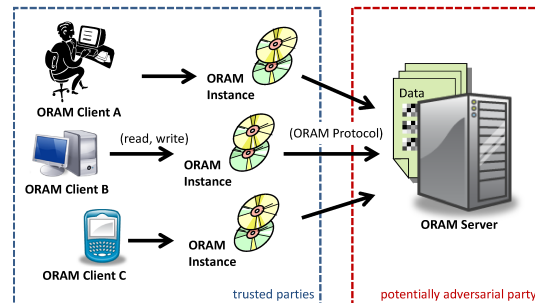


Fig. 1. Overview: ORAM Clients access data obliviously through a simple ORAM Instance interface. The stateless ORAM Instances use an (untrusted) ORAM Server to store and retrieve the data obliviously.

- First, consider that as a client query searches across the database for a particular item, its accesses are by design indistinguishable from random.
- Second, by ORAM construction, different unique queries touch independent sets of database locations (because their coin-flips are independent).
- Now consider a set of multiple queries. Any of its reorderings results in accessing the same locations (provided each query gets the same coin flips), albeit in a different order.
- Thus, intuitively, it is safe for clients to submit unique queries simultaneously: the server sees an identical transcript independent of the underlying queries.

The above privacy intuition only holds for non-repeating unique sets of queries, of course. Clients simultaneously querying the same item reveal this to the server, as their accessed locations are substantially similar. This raises the interesting question of how to guarantee query uniqueness over arbitrary incoming client query patterns, without revealing any inter-query collisions to the server.

Since the model requires the ORAM Instances to communicate only via the ORAM Server (capturing an adversary with a complete view of the network), one idea is to have Instances synchronize with each other via a server-hosted data structure, in a way that prevents overlapping equivalent queries while still guaranteeing indistinguishability of all query patterns.

Strawman. The simplest approach to ensuring uniqueness is for each new client to examine ongoing queries. Finding an intersection, the client can then wait for ongoing queries to complete before trying again. This is fundamentally insecure, however, since the server learns when a query is being repeated by a later client, since the later client waits for the previous client. Further, correcting this by making all clients wait for all previous queries to complete before initiating their own query defeats the goal of parallelism.

Alphaman. To fix these issues, the server will help clients maintain an encrypted query log. Scanning this query log allows an Instance to identify simultaneously ongoing requests. In the case of overlap, a random unique query is executed instead. Once its query (real or random) is completed, the Instance reports its result back to the shared cache/result log. It then searches in the log for the result of the simultaneously ongoing overlapping query it had identified (if any). This guarantees that in either case the Instance gets to learn the result it seeks.

The following sections provide more detail of this approach, identifying the conditions of applicability to existing ORAMs, and showing compatibility with de-amortization techniques.

5. ABSTRACTIONS AND SOLUTIONS

This section presents generic approaches to parallelization and de-amortization, after identifying necessary properties of the underlying ORAM Instance building blocks. The idea is to use pieces of existing ORAMs as building blocks; (Figure 2).

5.1. Parallelism Abstraction: “Period-based”

The main idea is to run queries simultaneously *between* reshuffles. We are limited by the size of the top level: this is the number of appends that can be performed before a shuffle is required. Thus, we designate the maximum parallelism *and* the size of the top level as c . Subsequent levels will be sized $c2^i$ for this reason.

DEFINITION 1. A **period-based stateless ORAM Instance** is an ORAM that performs a series of c queries between each shuffle. The transcripts of unique queries within a period are independent of their order. Previously executed queries are scanned from a server-stored cache sized c , triggering a fake lookup instead.

Procedure Query-Period(id, val)

```

 $queryid \leftarrow queryid + 1$ 
 $coinflips \leftarrow PRNG()$ 
if  $id \in cache$  then
   $result = cache\{id\}$ 
  QueryORAM( $fake, queryid, \perp,$ 
     $coinflips, secretkey$ )
else
   $result \leftarrow$  QueryORAM( $id, queryid,$ 
     $val, coinflips, secretkey$ )
end
 $cache \leftarrow cache \cup (id, result)$ 
if  $gen \simeq_c 0$  then
  AppendAndRebuild( $cache$ )
   $cache \leftarrow \emptyset$ 
end
return  $result$ 

```

The following operations define a period-based stateless ORAM (which in turn invoke its underlying ORAM Server methods QueryORAM($id, queryid, val$), and AppendAndRebuild(new items)): **read**(id): return Query-Period(id, \perp), **write**(id, val): Query-Period(id, val), and **test-and-set**(id, val): return Query-Period(id, val).

Insight. The goal of this abstraction is to capture the fact that the underlying ORAM already supports up to c simultaneous *unique* queries. Specifically, the ORAM runs in periods of queries over which the transcripts of unique queries are independent (e.g., the original ORAM [Goldreich and Ostrovsky 1996] satisfies this). For a given period between reshuffling (which lasts several queries, until the top level overflows), and choosing the random number

generator coin flips for two unique queries ahead of time, the transcript for each query is the same regardless of the order they are run. The transcripts of two identical queries, however, are inter-dependent, since the first-to-execute query necessarily searches farther down in the database than the second query, as the item is moved up to the top level once the first query completes. The second-to-execute query finds the item immediately at the top, and thus the remainder of the search is random.

5.2. Parallelism Construction

The motivation behind this construction is to minimize waiting, while guaranteeing serializability. The basic assumption about the underlying ORAM is that its data structure supports simultaneous querying for unique blocks. Then, the *ultimate purpose of this protocol is to make sure all the simultaneous queries request different items*. The challenge is to juggle this uniqueness requirement (in the presence of colliding queries from different clients) with the requirement that the server not learn any inter-query correlation.

One solution is to obviously guarantee uniqueness of queries using an append log in combination with a results log. Clients will still need to wait for previous queries to complete before outputting a result, but now there is no requirement of blocking on other clients during query execution. This increases throughput significantly (Figure 3):

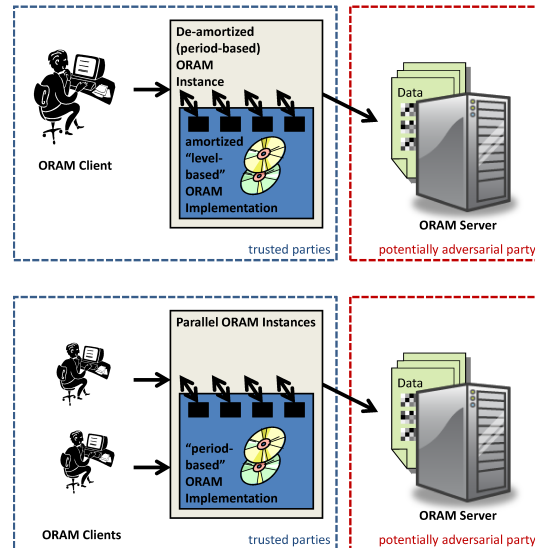


Fig. 2. Using existing ORAMs as building blocks. Top: building a de-amortized from an amortized level-based ORAM. Bottom: building a parallel ORAM from a period-based ORAM.

- (1) The Client issues a query to its ORAM Instance.

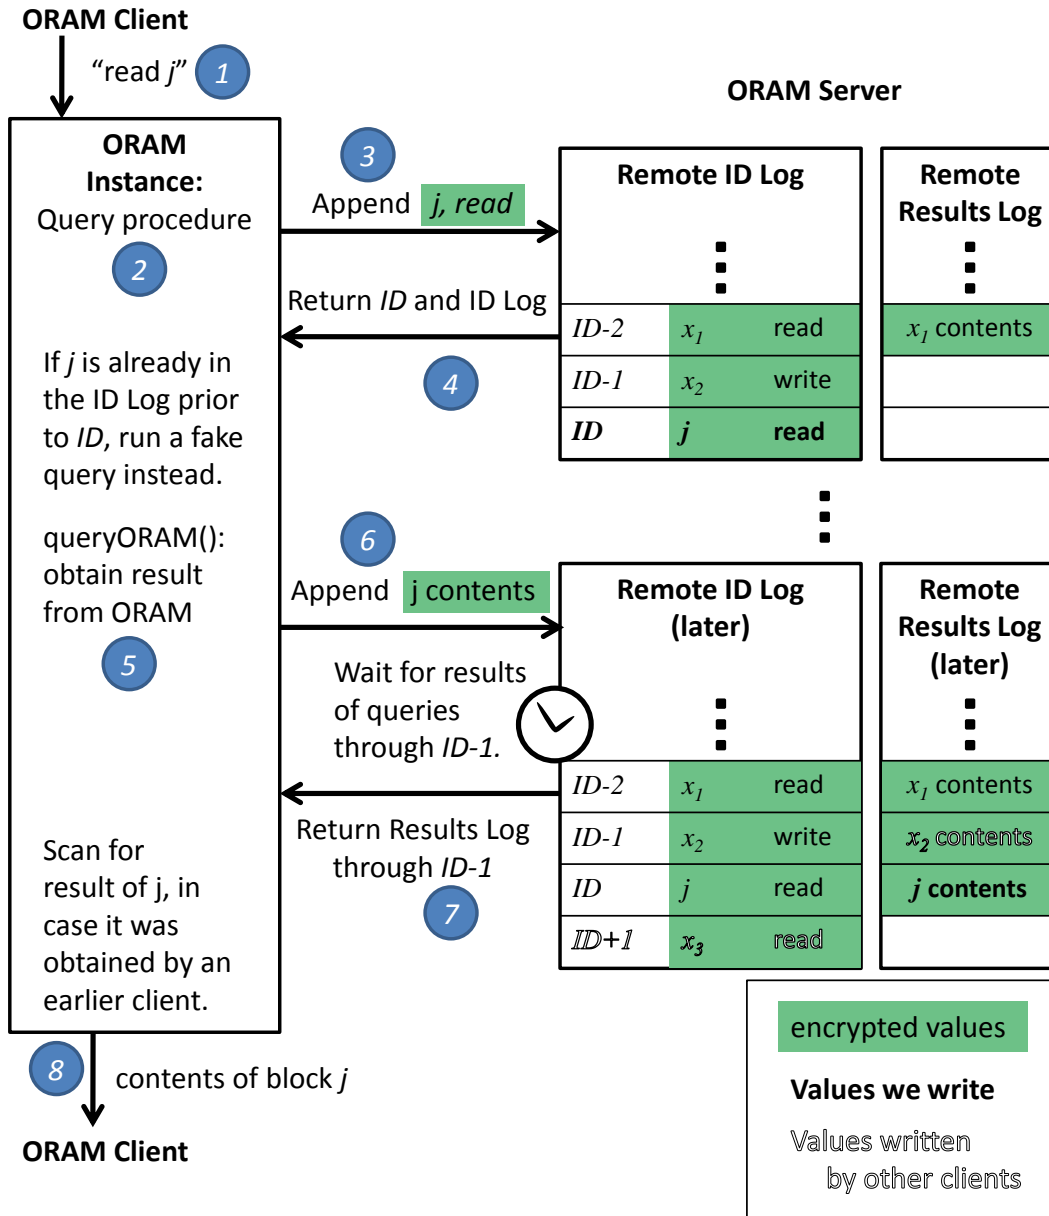


Fig. 3. Parallel query protocol

- (2) The ORAM Instance generates the request, consisting of the block ID, and a bit indicating whether this is a read or a write/insert. Test-and-set operations are identified as writes.
- (3) The request is encrypted and sent to the server.
- (4) The server appends this to the encrypted query log (analogous to the top level), and returns a sequential query ID, together with the query log, containing all queries since the top level was last emptied.

- (5) The ORAM Instance interactively queries the underlying ORAM. If the query was already in the query log (i.e., from another running client), the instance runs a dummy query instead.
- (6) The ORAM Instance sends its result, encrypted, back to the server, which appends it to the query result log.
- (7) The ORAM Instance reads the results log up to its own entry (only interested in previous clients' results), in case the current query is accessing a block that was previously read or written. This is the only step that requires waiting for the earlier queries to complete.
- (8) The ORAM Instance returns to the ORAM Client the result (obtained from the database, or the result log).

ORAM Instances wait for the entire result log prior to that query's registered location to complete before returning to the Client. However, in many cases the result will already be obtained. If it is assumed that the adversary cannot observe at what point the ORAM Instance returns to the client, then it is safe for the ORAM Instance to return its value sooner.

Server network traffic over a period of c parallel client queries is quadratic in c , since each Instance needs to be aware of what each simultaneous Client is doing². The number of parallel clients that optimizes total throughput is thus a function of network bandwidth, latency, and database size.

On the one hand, for c clients, and a database of n blocks, the sequence of $\log_2(n) + 3$ round trips per query imposes a network-latency based maximum query throughput:

$$\frac{c}{((\log_2 n) + 3) \times \text{latency}}$$

On the other hand, the cost of supporting multiple clients (quadratic in c over c queries; linear in c per query), and the online data transfer cost of $\log_2 n$ blocks, impose a server bandwidth based maximum throughput:

$$\frac{\text{bandwidth}}{((c - 1)/2 + \log_2 n) \times \text{blocksize}}$$

To find the optimal number of clients c for a given configuration the lower of these upper bounds needs to be maximized. This relationship is plotted in Figure 4 for a representative setting.

²To disseminate this information in a more network-efficient manner, e.g., by only requesting the log entry the client is interested in (instead of the entire log) a PIR protocol may be deployed. This results in a tradeoff between server bandwidth and server CPU time. This tradeoff is mostly unfavorable [Sion and Carbunar 2007], yet recent results [Olumofin and Goldberg 2011] indicate that it may become feasible in the future.




Fig. 4. Upper bound on online query rate of a 100GB database (blocksize = 10^4 and $n = 10^7$) and assuming a 100Mbit (12.5 MB/sec) network link. The plot is shown for various network round trip times, from 0 ms through 75 ms. Single clients, incurring $\log_2 n$ round trips per query, are tightly bound by the round trip time. Adding more parallel clients increases this throughput linearly, up to the point where bandwidth limits from the query log traffic take over.

After c queries execute, the query log is converted into the ORAM top level. Queries are thus applied to the database, in the ORAM-sense (Figure 5). This is where the requirement of a “period-based” Instance is necessary.

In the process, the shuffler consolidates the query log. A single block may be accessed multiple times by different queries, but is placed only once back into the database. The value chosen for a given block is the last write, if there is one, or otherwise the first read. Recall that subsequent reads are associated with fake results.

Properties. We now introduce *optimality* and *query privacy*.

THEOREM 1. *In any model in which the server can associate all visible read/write data accesses corresponding to a given query, hiding access patterns in a “non-simultaneous ORAM” requires waiting for the results of all previous queries (“wait-optimality”).*

PROOF. Take any “non-simultaneous” ORAM that requires serial execution of queries for the same underlying block. In such an ORAM, repeating the query before the previous query has completed would result in overlap between visible data accesses, and would allow an adversary to correlate these queries.

Thus, instead of attempting to repeat a previous query, clients must obtain the result from it once that query has finished executing.

Finally, to not reveal which previous query they are waiting for (if any) and make all query sequences indistinguishable to the server, even in a parallel ORAM, clients will still need to wait for *all* prior queries, as if their query depended on each of these previous queries. \square

Theorem 2 shows that running parallel queries in this model is safe. The intuition is as follows. In the parallel case, client behavior from the perspective of the server is identical for each new query instance regardless of whether and how often the same equivalent query appears earlier in the log. For every query, the server only sees a semantically secure encrypted *append* operation to the query log, a *query* to the underlying ORAM, and a *scan* of the results log up to that point.

The only difference now is that transcripts of the different queries are interleaved, but otherwise contain the same accesses as when executed in a traditional ORAM. This is so because in the parallel case, if a client queries for a block corresponding to a simultaneously ongoing query, the client’s ORAM Instance will instead issue a fake query—which is what it would have done anyway in the traditional ORAM had it found the query result at the top level. Thus, from the server’s point of view, the transcripts contain the same (random looking) accesses.

Further, query transcripts are independent of their query and, without knowledge of the secret ORAM key, indistinguishable from random, as required by Definition 18 of traditional ORAM [Goldreich and Ostrovsky 1996]. Then, an advantage at distinguish-

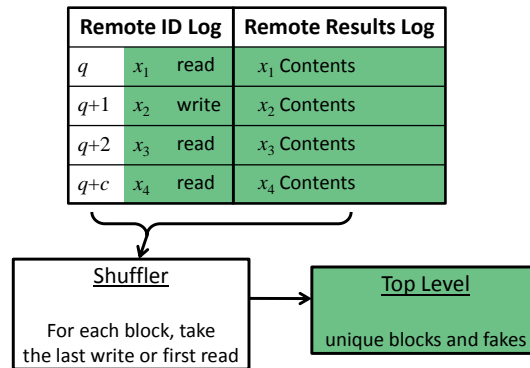


Fig. 5. Reconciliation of the query log into the new top level. After a period of c queries, the query log is shuffled and becomes the top level. Regarding the query log hosted set of ids, consider that an ORAM block can only be read once during this period. Thus, the correct value for a given block is either the last write of this period, if there is one, or otherwise the first read.

ing the new transcripts translates into an equivalent advantage at distinguishing the underlying ORAMs if parallelism were not enabled.

THEOREM 2. *The existence of an adversary with non-negligible advantage at violating query privacy (as in Definition 18 [Goldreich and Ostrovsky 1996]) in a parallel ORAM implies existence of an adversary with non-negligible advantage at violating the privacy of the underlying single-client ORAM (“query privacy inheritance”).*

PROOF. Take an adversary A with advantage ϵ at correlating queries in a parallel ORAM based on an underlying ORAM O . We now construct an adversary B with equivalent advantage at correlating queries in O .

B simulates adversary A on every query. Since the interleaving information is publicly known, it includes this information in the transcripts it uses to simulate A . It appends random information to the query log contents for A . These query log contents give A no additional non-negligible advantage; otherwise a distinguisher could be built distinguishing the semantically secure encryption function output from random.

B then outputs the guesses and requests provided by A , and obtains the same advantage as A . \square

5.3. “Level-based amortized ORAM”

This section introduces a generic de-amortization technique, together with the abstractions necessary to deploy it to de-amortize a large class of ORAMs and reduce their worst case query cost to the average case.

While several other ORAMs have recently introduced de-amortization (reducing the worst case query cost to average case), as described in Section 2, a generic technique is described here, with the abstractions necessary to make use of it.

Consider a “level-based amortized ORAM” which divides the database into levels, with the more recently accessed items stored in the smaller top levels. Retrieval requires searching within a set of specific locations in each level from the top down, and, once found, lifting the result back up into the top level. The next time the same query is repeated, the block is discovered at the top, which causes the ORAM Instance to instead access a random set of locations at each level below. The top level overflows periodically into the lower levels, and a level is guaranteed to get re-shuffled

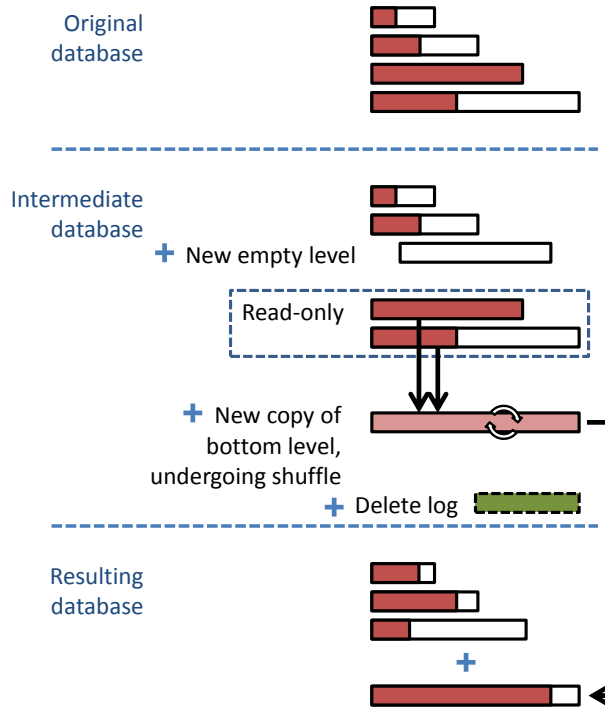


Fig. 6. Background construction of a single level. Top = initial database; middle = database during construction of new bottom level; bottom = resulting database. In this scenario, the third level is full, so it needs to be combined with the fourth level. Read-only copies of these two levels are made, and can be queried while the combination is occurring. Simultaneously, overflows from the top level are placed into a replacement third level. Queries access all five levels at this time. Once the bottom level construction is complete, querying resumes with this new bottom level, replacing the two read-only levels (containing the same items).

before a given query is repeated at that level. Thus each query looks like only an selection of random, independent locations in each level.

DEFINITION 2. A **level-based amortized ORAM Instance** is an ORAM that searches levels recursively, appending the result back to the first level. Privacy results from the property that an item is sought at a particular level no more than once between two consecutive shuffles of that level.

De-amortization techniques need to deal efficiently with the level constructions resulting from overflow of the top levels. Their goal is to arrange the levels such that they can be queried while the items are simultaneously being inserted and re-shuffled into new levels. That is, instead of suspending querying to wait for shuffling to proceed, a new level must be available as soon as it is needed.

Deamortized ORAMs expose the following Client interface (implemented by also calling the underlying ORAM's `Try-Level(id, level)`, `Fake-Try(level)`, `Top-Append(Item)`, and `Combine(levels)` operations): `read(id)`: return `Query(id, ⊥)`, `write(id, val)`: `Query(id, val)`, and `test-and-set(id, val)`: return `Query(id, val)`, where `Query()` is one of `Query-Amortized()` (original ORAM) or `Query-Deamortized()` (defined below).

Intuition. The goal of this definition is to identify those ORAMs that recursively query a data structure divided into multiple levels. These ORAMs guarantee privacy by enforcing that blocks are never sought in the same location twice. That is, after seeking a block once at a given level, that level will be “re-constructed” before the block is sought there again.

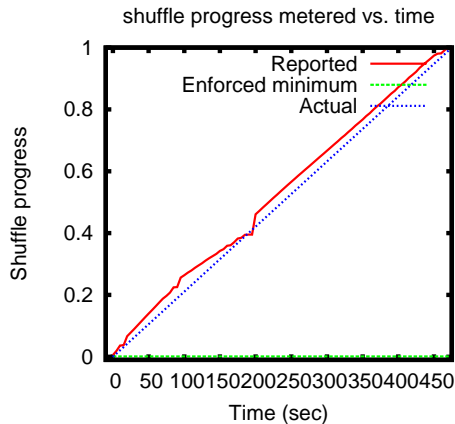


Fig. 7. Level shuffle progress: the linear nature of the plot indicates the extrapolation based on partial shuffling is well done.

Procedure Query-Amortized(id, val)

```

success  $\leftarrow \perp$ 
for  $i = 1$  to NumLevels do
  if success then
    | Fake-Try( $id$ )
  else
    | (res, success) = Try-Level( $id, i$ )
  end
end
if  $val \neq \perp$  then
  | Top-Append( $val$ )
else
  | Top-Append(res)
end
Rebuild()
return res

```

Procedure Rebuild

```

for  $i = 1$  to NumLevels do
  if Overflow( $i$ ) then
    | Level  $i + 1 \leftarrow$  Combine( $i, i + 1$ )
    | Level  $i \leftarrow \emptyset$ 
  end
end
end

```

Level-based ORAMs move each retrieved block up to the top level of the ORAM, from where it slowly makes its way back down the data structure as the lower levels are re-constructed. Level-based amortized ORAMs include [Goldreich and Ostrovsky 1996; Williams et al. 2011; Williams et al. 2008; Goodrich and Mitzenmacher 2011; Pinkas and Reinman 2010].

5.4. De-amortization Construction

We now construct a level-based de-amortized ORAM from a level-based amortized ORAM. The main idea is to provide pre-emptive shuffling. Rather than waiting for querying to complete before shuffling a level, its transformation into a new level begins as soon as a level is constructed, and right as its querying begins.

To allow this, we duplicate a level into two copies: a read-only variant that is used in the querying process, and a writable variant which is dynamically updated into the new generation of this level. The read-only copy is discarded at the end of the period, since it is no longer needed for either queries or level construction.

Level De-amortization. Consider the de-amortization of a single level. In a traditional ORAM, a level is reconstructed by combining into it the above level that has filled up and now overflows (Section 2.1). This necessarily stops the query process for its duration.

To de-amortize this, when beginning its construction, instead of pausing queries and waiting for the construction to finish, querying can continue via the read-only level copy while a new generation is produced into the writable variant. Critically, during this process, existing levels can overflow into a fresh, empty, replacement for this level. (Figure 6).

Delete Log. The second change is to delay level updates until the end of the shuffle. Some ORAMs avoid the complexity of reconciling multiple versions of an item (and in the process, reduce storage overheads) by deleting blocks from the levels where they are found. In the de-amortized construction, this is not possible, since the queried level copies are now read-only. Instead, these changes are appended to an update log. Items marked for deletion in this log are now deleted by the server before the *next* shuffle of the level. Not all ORAMs modify the pyramid structure during queries; some [Goldreich and Ostrovsky 1996] do not need such modification. This de-amortization protocol requires that those changes that are made can be applied *after* the level has been reconstructed (instead of before).

Details. A level at height i , containing $m = c2^i$ items, is queried m times. At the end of the m th query, the shuffle will have completed, and the remaining items from this level are now in a new level.

Each level is a data set, completely specified by its height i , the sequentially increasing “generation” j at that height, and a one-bit marker for the odd generations. The two levels at an even generation j , denoted by $i.j$ and $i.j*$, are combined to produce a level at the next height $i + 1$, generation $j/2$. Those levels at a height i with an odd generation j are reshuffled to produce level $i.j + 1*$.

The levels currently reshuffled are the ones queried – at any one time there are either one or two active levels (queried and being shuffled) at each given height, for $i \leq \log_2 n$. Figure 7 compares the live extrapolated progress estimates used in de-




Fig. 8. Sample unmodified, amortized, construction for reference. Observe that querying halts periodically while levels undergo construction.

amortized level shuffling, to the actual level construction time. Figure 9 illustrates the shuffling and querying schedule resulting from de-amortization.

The above construction allows de-amortization of the construction of all the levels except the top level. De-amortization of the top level appears to be possible using a rotating query log, but is left as future work (since its impact is minimal as the top level is very small). Figure 8 shows a sample *amortized* construction (before de-amortization).

If based on an underlying stateless period-based amortized ORAM, the result is a stateless period-based de-amortized ORAM, suitable for parallelization.

THEOREM 3. *A polynomially bounded adversary has no non-negligible advantage at guessing the access pattern based on observing the transcripts of a de-amortized BF-based ORAM.*

PROOF. The output of the underlying BF ORAM level construction process is shown in previous work to be a randomly ordered set of stores (ID,value pairs). Both the ID and value are opaque; that is, the server cannot distinguish them from random. The only other time the server sees this opaque ID is when the ID is retrieved later on. The same opaque ID is never retrieved twice. BF locations are also accessed in a manner uncorrelable to the access pattern, as the locations corresponding to IDs are chosen with a pseudo-random number generator. \square

THEOREM 4. *As long as the underlying level-based ORAM destroys all inter-item correlations on shuffle, then existence of an adversary with non-negligible advantage at violating query privacy in a de-amortized ORAM implies existence of an adversary with non-negligible advantage at violating the privacy of this underlying level-based ORAM (“query privacy inheritance”).*

PROOF. A de-amortized ORAM operates equivalently to the underlying ORAM, with two differences. First, the level access structure is different—there are up to twice as many levels—but the privacy-preserving property that no item is requested twice under the same (opaque) identifier is the same. Second, we don’t delete items until after the shuffle (instead of before the shuffle), but this provides no additional information as long as the level construction process is also opaque.

This is the only new information given to the adversary. Privacy is shown by reducing to a property of ideal permutations: uncovering a portion of a permutation does not reveal anything about the remaining portion. Since the underlying level-based ORAM selects the level permutation according to an ideal secret permutation, this follows trivially: all arrangements of the elements in the “still secret” portion are equally likely. It constitutes a secret permutation in itself.

This results in an ordering and labeling indistinguishable from random. \square

Procedure Query-Deamortized(id,val)

```

success  $\leftarrow \perp$ 
gen = queryid/c
for  $i = 1$  to NumLevels do
  if success then
    | Fake-Try(id, i.gen)
  else
    | (res,success)  $\leftarrow$  Try-Level(id, i.gen)
  end
  if  $gen \cong_2 0$  then
    | if success then
      | | Fake-Try(id, i.gen.*)
    | else
      | | (res,success)
      | |  $\leftarrow$  Try-Level(id, i.gen.*)
    | end
  end
  gen  $\leftarrow \lfloor (gen - 1)/2 \rfloor$ 
end
if val  $\neq \perp$  then
  | Top-Append(val)
else
  | Top-Append(res)
end
Some-Rebuild()
return res

```

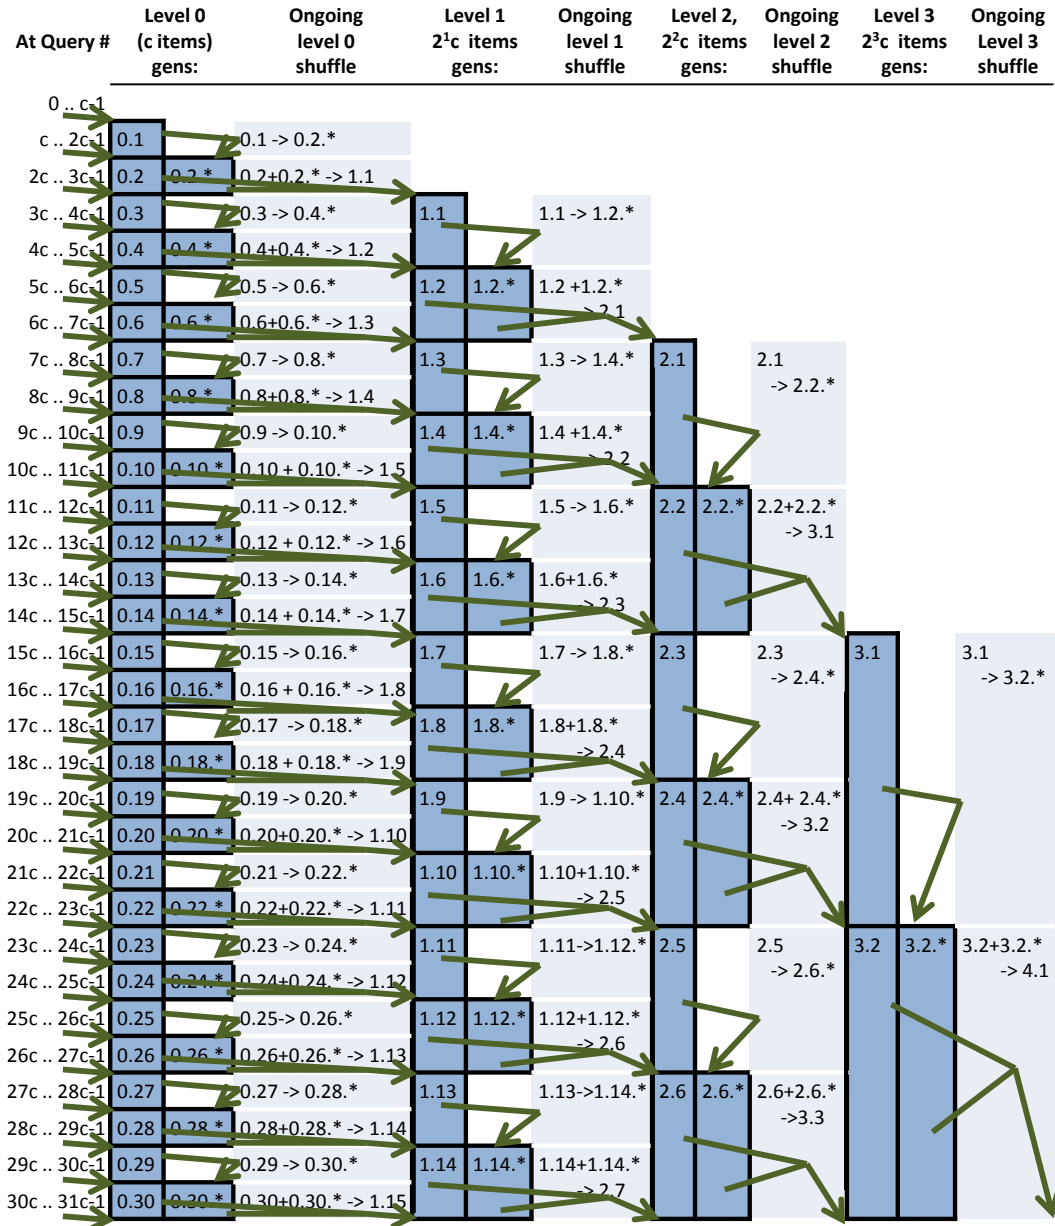


Fig. 9. Shuffle schedule: Query numbers are listed in the first column. For a given query, the dark outlined boxes overlapping that row specify which levels and generations to access (from left to right). The light boxes identify the simultaneously occurring reshuffle.

Storage overhead. Since the client storage required to shuffle level i is less than or equal to the space required to shuffle the next larger level, $i + 1$, de-amortization requires an extra factor of client storage upper-bound by $\log_2 n$ (since all $\log_2 n$ levels are being shuffled at once).

Communication overhead. This construction requires querying up to twice as many levels simultaneously. However, since it is known in advance that the item won't show up in both levels, this querying can be done in the same number of round trips.

Shuffling overhead. The amount of work performed per query is roughly similar. Additional overhead may stem from shuffling being done *before* level items are removed. But since only half of the items would be removed anyway, for ORAMs where the level i construction cost is $sm \log_2 m$ for constant s and $m = 2^i$, this keeps the new overhead for shuffling within a factor of $\frac{s(2m \log_2 2m)}{s(m \log_2 m)} = 2 + \frac{1}{\log_2 m} < 3$.

5.5. BF ORAM with logarithmic memory

We now detail the final piece of the construction: the underlying base ORAM mechanism. For illustration we take our ORAM [Williams et al. 2008] that separated each ORAM pyramid level into two data structures: a hash table of items, indexed by unique random IDs, and an encrypted BF to check whether an item is stored at a given level. Querying this ORAM proceeds analogously to the standard pyramidal ORAM model—starting at the top, searching downward until the item is found, querying randomly from then forward. However, instead of scanning a bucket sized $O(\log n)$ at each level, the encrypted BF is checked, requiring only a single item retrieval from each level.

At reshuffle time, in [Williams et al. 2008], BFs for each level were built securely by the client using a somewhat complex procedure based on an oblivious scramble of a list representation of positions to set, followed by bucket sort. This allowed costs under $O(\log^2 n)$, requiring $k\sqrt{n \ln n}$ client storage.

We now introduce a technique resulting in an $O(\log^2 n)$ BF-based ORAM and requiring only *logarithmic* client storage. Instead of using the “bucket sort” method that builds \sqrt{n} -sized chunks of the encrypted BF, we construct, then sort, a server-stored list of the BF segments and indexes:

First, in a single pass over the encrypted items of the level undergoing construction:

- build an encrypted, server-stored, list of the positions that need to be set in the encrypted BF
- divide the BF into segments of size $O(\log n)$
- append one encrypted segment identifier for each possible segment to this list, padded so the server cannot distinguish segment identifiers from BF positions.

Second, obviously sort this entire list [Williams et al. 2008; Goodrich 2010], with segments interspersed among the positions, such that they immediately follow the positions they belong in.

Third, in a single pass over this sorted list:

- output one BF segment for every list element: for each segment identifier, output an encrypted segment, with the appropriate positions set. Recall that these positions immediately preceded this segment identifier in the sorted list.
- for each position encountered in this list, output a dummy empty segment.

Fourth, obviously sort the resulting list by segment identifier, with the dummy segments at the end.

Fifth, truncate off the dummy segments from this list to yield the encrypted BF.

The advantage of this approach over [Williams et al. 2008] is that using a logarithmic-space $O(m \log m)$ oblivious sort [Goodrich 2010] provides a logarithmic-space construction still running in $O(\log^2 n)$ time. As discussed, employing this as a de-amortized ORAM brings the client storage requirement to $O(\log^2 n)$.

This construction process is secure, provided (a) the oblivious sort is private, (b) the server cannot distinguish an encrypted dummy segment from an encrypted BF segment, and (c) the server cannot distinguish an encrypted segment identifier from an encrypted BF position. If these three conditions are satisfied, then every run of this

process over a set of m items appears identical to the server, regardless of the positions in the BF being set.

The PD-ORAM implementation assumes $k\sqrt{n \ln n}$ client storage to employ the Oblivious Merge Sort (instead of the disk-seek intensive less efficient randomized shell sort).

5.6. Security against malicious adversaries

Ensuring security against a malicious adversary first requires an underlying ORAM providing these guarantees. Fortunately, the BF-based ORAM of [Williams et al. 2008] provides such a mechanism. Second, maintaining security in a parallel setting requires clients to test that they all see a consistent view. This is achieved using a hash tree over the set of all previous queries. Whenever a client performs the top level shuffle, it is responsible for updating this hash tree and attaching the root value with a MAC to the new query log. This entails hashing the current query log, appending it as a new leaf node to the hash tree, and recomputing hash values along the path from this new node to the top.

Second, whenever a client performs a query, it performs one additional operation: it verifies that the last query it performed is included in the hash tree, whose root corresponds to the value attached to this query log. This is done by verifying all nodes in the hash tree adjacent to the path from the root to its last request.

These operations will not detect a forking / split universe attack, unless out-of-band communication is available, or clients perform periodic accesses to ensure they are still operating in the same view as other clients. However, this solution has the property that once the server forks the universe into multiple views, it cannot rejoin the views without being detected by the clients (by hash value disagreement).

6. EXPERIMENTS AND ANALYSIS

Amortized Measurements. A significant challenge in measuring the performance of any amortized system is ensuring the trial captures the average performance, not just peak performance. This is complicated by the requirement of running trials for periods that are too short to encompass the full period over which the amortization is performed. For example, at even several queries per second, the reconstruction of the lowest level of a terabyte database would require significant time, of the order of weeks or longer.

De-amortizing the level construction provides the opportunity of improving measurement accuracy. The challenge remains, however, of ensuring the de-amortized background shuffle proceeds proportionally to the query rate: the de-amortization is perfect when the new level construction is completed at the instant it is needed by a query. Inaccuracies in this rate synchronization will affect the measured results, since measured query throughput of a short period might be higher or lower than the sustainable rate.

To avoid this effect, PD-ORAM maintains *progress meters* for level construction, allowing queries to proceed when every level is proportionally constructed. The level constructions are also suspended when a level gets too far ahead of the current query. This keeps querying and level construction smooth, minimizing worst case latency (Figure 10).




Fig. 10. Individual query latencies + running averages of a single client. With perfect de-amortization, all queries would require the same amount of time.

Proper de-amortization: Theory vs. Reality. Performing proper de-amortization proved a non-trivial systems challenge. Research solutions, such as [Goodrich et al. 2011a], and PD-ORAM (Section 5.4) express de-amortization in terms such as “perform the proportional amount of work required”, or “perform the next $O(f(x))$ accesses.” While these terms suffice for proving existence of a de-amortized construction, programming models do not typically provide this type of abstract control. PD-ORAM achieves this control by metering progress over the construction of individual levels. Since the level construction involves different types of computation across the client and server, accurate progress metering required splitting level construction into tasks whose progress can be reported over time. Moreover, this metering uses experimentally determined values to identify what portion of the level construction corresponds to which subtasks.

Figure 7 illustrates the result of informing the progress meter with experimentally determined values. Strict de-amortization is employed: construction is paused until the next level is proportionally constructed, to ensure the correct portion of resources are allocated to construction of each level. Moreover, individual queries are also paused until the correct proportion of each level is constructed. This ensures that the next level is ready to query as soon as it is needed, eliminating waiting.

As an aside, suspending the level construction when it outpaces the queries proved critical on the larger database sizes. The sheer number of requests being sent from the ORAM Instance to the ORAM Server for construction tended to starve the requests of actual queries (much fewer in number), causing the query rate to drop quickly as more levels were introduced. This behavior was corrected by forcing level construction to remain proportional to query progress: this keeps the individual query rates much closer to average.

Since it is impractical to repeatedly running trials over the entire (up to 1TB) measured epoch, the database for these trials is first constructed non-obliviously on the server via a specially designed module. The items are inserted randomly so that the final result mirrors an oblivious construction (as would occur from a sequence of write queries).

6.1. Setup

PD-ORAM is written in Java. Clients are distributed evenly across several Intel Xeon machines, each with two quad-core 3.16GHz Xeon X5460’s. The Server runs on a single Quad-Core Intel i7-2600K Sandy Bridge 3.4GHz (3.8GHz Turbo Boost) LGA 95W 1155 CPU, with 16GB DDR3 1600 (PC3 12800) SD-RAM and seven HITACHI Deskstar 2TB 7200 RPM SATA 3.0Gb/s 3.5” SATA disks (RAID0 / LVM).

All the machines share a gigabit switch. Network latency is shaped by forcing server threads to sleep for the desired round trip duration upon receiving a request. This allows simulation of link latency without capping link bandwidth.

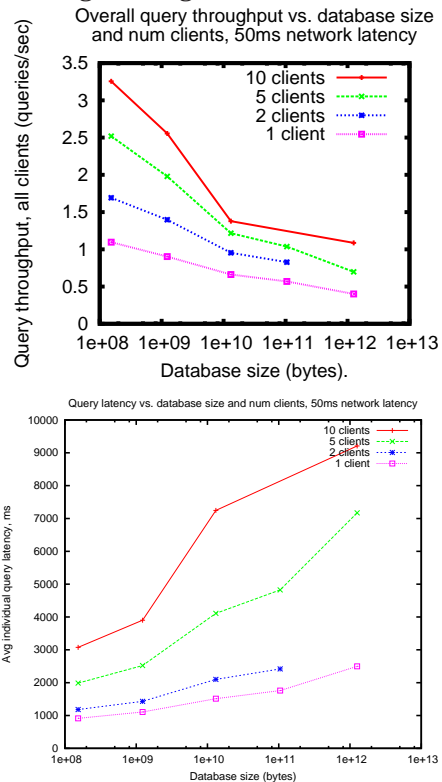


Fig. 11. (a) Throughput vs. data size for varying number of clients (x-axis is log scale). (b) Latency vs. database size. Each point is sampled over 3000 queries.

The implementation uses a BF with 8 hash functions, and 2400 bits of space per item which allows an efficient construction within the false positive rate of 2^{-64} per lookup. The resulting BF constitutes roughly 25% of the total size of the database records.

Optimization. Rather than optimizing the BF size required to obtain this error rate by using a larger number of hashes, as suggested by Pinkas et al. [Pinkas and Reinman 2010], PD-ORAM uses larger BFs with fewer hashes, to minimize item lookup disk seeks while obtaining the same error rate.

6.2. Experiments

One main goal of the experiments is to understand the interaction between network parameters and the parallel nature of PD-ORAM.

Size + clients vs. query throughput. Figure 11 plots the effect of database size and client parallelization on overall query throughput. Fresh databases were used for all trials to prevent dependency of the measurements on the order of the trials, except for the 1TB trials, where this proved impractical. Even though individual query latency increases with higher resource contention, the benefit of parallelization are obvious: significantly higher overall throughputs.

Clients + network latency vs. performance. Figure 12 plots the effect of parallel clients and network latency on overall query throughput for a fixed database size. The premise of this measurement is that parallelization becomes more important as network latency increases.

De-amortization optimality. Figure 10 plots the observed latency of individual queries vs. time on a growing database. With perfect de-amortization, all queries would require the same amount of time. Most queries take around 1200ms; a fixed lower limit is imposed by the network latency. The bands at 2600ms and 3100ms reflect the construction of the top level, which is not de-amortized. The slight downward slope suggests the progress estimation at the beginning of the shuffle is conservative.

The points in the top left, in particular, suggest slight inefficiencies. However, even those points are much improved over the worst case costs without de-amortization (and background shuffling).

Overall, this enforced de-amortization leads to a lower sampled query rate than the one obtained in [Williams et al. 2008] – the values measured here reflect sustainable rather than the peak query rates.

Progress metering. To validate the accuracy of the progress metering, Figure 7 shows the reported construction progress of a single level as sampled every 5 seconds. Strict de-amortization and querying is disabled for this trial, to avoid cool-down periods when construction has progressed farther than is needed, and to ensure measurement of its progress only.

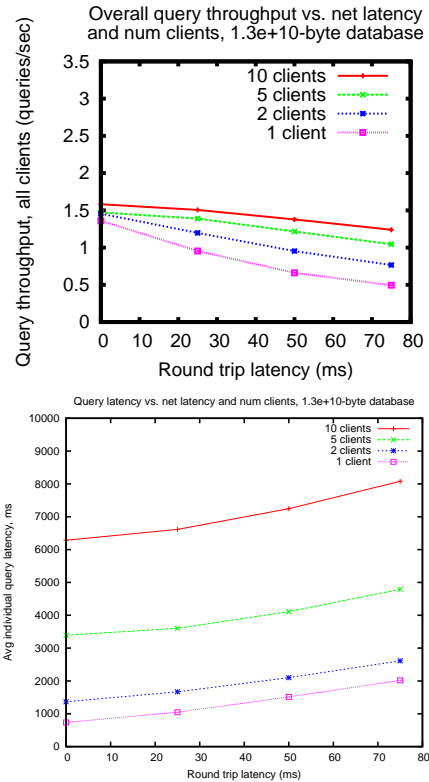


Fig. 12. (a) Throughput vs. network latency for varying number of clients. (b) Query latency vs. network latency. Each point is sampled over 3000 queries.

6.3. Impact of disk latency

The experiments were repeated (for database sizes up to 300GB) in a different setup, in which the server was run on dual 3.16Ghz Xeon X5460 quad-core CPUs and *six 0.4 TB 15K RPM SCSI (hardware RAID0)* disks.

This configuration surprisingly outperformed the setup above by a factor of 2x in most trials. The primary advantage is the superior seek time on the server disks, so the markedly different results suggest that server disk seek costs play an important role in overall performance. This was somewhat surprising, since the level construction mechanisms were designed specifically to minimize disk seeks (with the hash table insertion being the only random-access operation during level construction, requiring an average of 2 random writes per insert). The rest of the level construction simply requires reading from one or two sequential buffers, and writing out sequentially to one or two.

The culprit is most likely the de-amortization process, which constructs different levels in parallel, and in effect randomizes disk access patterns. While individual level construction is mostly limited by sequential disk throughput, running many of these processes in parallel across the same file system results in disk seeks even in the sequential access regions, resulting in a much lower overall disk throughput.

SSDs. Several software and hardware solutions present themselves. Better data placement would split data in a more efficient manner across the available disks (instead of using a RAID configuration), to allow the sequential nature of each level construction process to transfer to sequential disk access. Obtaining optimal throughput in this manner would require a relatively large number of disks. Further, the use of more expensive (but quickly dropping in cost) low-latency solid state disks (SSDs) would be a simple hardware solution to eliminate this performance bottleneck. It remains to be seen however, whether the sustained random write performance degradation plaguing current SSDs does not constitute a bigger bottleneck in itself, as preliminary throughput experiments on several recent 128GB Samsung SSDs with 2011 firmware updates seem to suggest.

7. PRIVATEFS: AN OBLIVIOUS FILE SYSTEM

ORAM lends itself naturally to the creation of a block device. Due to existing results' impractical performance overhead this has not been previously possible. A Linux-based deployment of PD-ORAM is used here to design and build *PrivateFS*, a fully-functional oblivious network file system in which files can be accessed on a remote server with computational access privacy and data confidentiality.

An initial implementation was built on top of the Linux Network Block Device (NBD) driver, which is the simplest and most natural approach, since PD-ORAM already provides a block interface. However, NBD supports only serial, synchronous requests. To take advantage of the parallel nature of PD-ORAM, *PrivateFS* is instead built on FUSE (Filesystem in Userspace [Henk and Szeredi 2012]).

A second attempt used *ext2fuse*, a FUSE-based ext2 implementation [Scholl 2009], by rerouting block access through PD-ORAM. However, thread safety difficulties prevented us from modifying it to support parallel writes or reads. Additionally, because of its nature as a block device file system, *ext2fuse* requires mechanisms for allocating blocks for files, such as block groups, free block bitmaps and indirect file block pointers inside inodes. These mechanisms are not all thread-safe and pose a challenge to synchronize. Moreover,

File System	Write speed	Read speed
nbd (ext2)	14.80KB/s	10.45KB/s
ext2fuse	7.11KB/s	9.62KB/s
PrivateFS	83KB/s	61.69KB/s

Fig. 13. Performance in a zero latency environment for a small 10MB filesystem using the two initial trial implementations, and the *PrivateFS* implementation.

locking the code using synchronization primitives would not result in a sufficient degree of parallelization.

Instead, we implemented a full file system in FUSE. It fully leverages PD-ORAM parallelism and also takes advantage of the non-contiguous block labeling in a way that block-device file systems cannot.

Following the Linux file system model, in *PrivateFS* files are represented by inodes. Directories are inodes containing a list of directory entries; each directory entry is the name of a file or subdirectory along with its inode number. Inodes are numbered using 256-bit values and are mapped directly to ORAM blocks, such that inode x is stored in ORAM block x . Inodes hold metadata such as type, size and permissions. *PrivateFS* and (this instance of) PD-ORAM use 256-bit block identifiers and 4096-byte blocks.

Because the ORAM provides random access to 256-bit addressable blocks, a block can be allocated simply by generating a random 256-bit number. We take advantage of this in two ways. First, to read or write the i^{th} block of file with inode number x , the pair (x, i) is hashed with the collision-resistant SHA256 hash, yielding the 256-bit ORAM block ID for that file block. Second, when a new file is created, a 256-bit inode number is randomly generated, as opposed to maintaining and synchronizing access to an inode counter.

Our design eliminates the complexity of contiguous block device file systems and minimizes the need for locking when writing or reading files. As opposed to *ext2fuse*, *PrivateFS* does not incur the overhead of maintaining free block or inode bitmaps, grouping blocks into block groups, or traversing indirect block pointers to read files. The potential drawback is that sequential blocks of a given file will not be stored contiguously in the file system. However, this is harmless when using an ORAM, since there is no notion of sequential block numbers (which would compromise access privacy).

PrivateFS employs exclusive locks when reading and writing directories. In addition, an LRU cache is implemented to quickly retrieve an inode's data given its inode number and also for file path to inode number translation, which helps avoid long directory traversals (and associated locking).

PrivateFS communicates with the ORAM server by means of a proxy (written in Java), which receives block requests from the file system and satisfies them using parallel connections to the ORAM server. This design choice affords us a higher degree of modularity, enabling us to connect *PrivateFS* to other ORAM schemes in the future.

We benchmarked *PrivateFS* along with our previous file system attempts, using a parallel workload writing ten, 512KB files simultaneously to the file system, then checking their integrity (also done simultaneously) using the sha256sum utility. For the single-client NBD implementation, throughput ranged between 10.45KB/s and

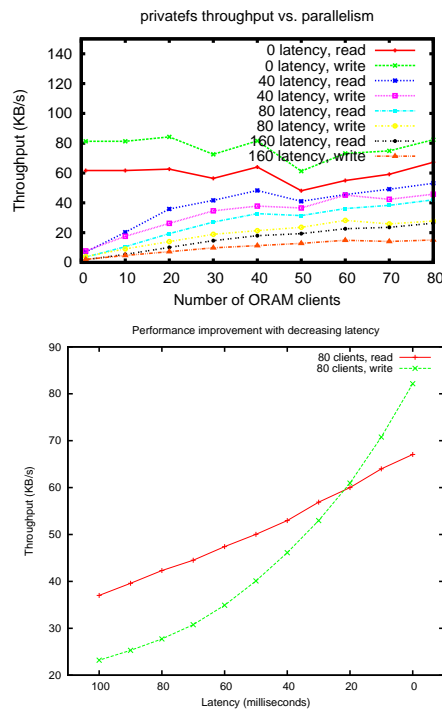


Fig. 14. (a) r/w performance vs. the number of ORAM Clients, for various network latencies. The zero latency environment is mostly unaffected by the degree of parallelism. For higher latencies, throughput grows with the number of parallel clients that offset the latency. (b) Varying latency for 80 clients in (10MB database).

14.80KB/s. For ext2fuse, throughput ranged between 7.11KB/s and 9.62KB/s. In contrast, the performance results for the full implementation indicate major speedups (Figure 13).

Further benchmarks (Figure 14) indicate that *PrivateFS* is benefiting from the high degree of parallelism in high-latency environments. We ran five trials for each point, varying the number of ORAM clients used by the proxy to satisfy file system block requests and performance increased proportionally with the degree of parallelism. The average over the five trials is plotted.

Reads are less expensive than writes in the 40ms and 160ms trials, because the individual writes are performed synchronously, while the reads can be parallel, resulting in a higher overall degree of parallelism. On the other hand, the low latency in the 0ms trial prevents this parallelism from having an impact. The reads are, in turn, more expensive, due to inefficiencies in the de-amortization method resulting in slightly slower queries at later points in the process.

We also wanted to analyze the behavior of the datastore as more data was written to it (Figure 15). We used the widely used iozone benchmark [Capps 2012] to test the write throughput as we wrote more data to the file system. We started with an empty file-system and repeatedly ran a parallel write throughput test using iozone. The IOzone test writes ten 1MB files and then rewrites them.

We compared *PrivateFS* to NFSv3 and ext4 using the IOzone workload. The *PrivateFS* tests wrote ten, 1MB files concurrently, while the NFS and ext4 tests wrote ten 1GB files concurrently. We decided to use larger file sizes when performing the tests for NFS and ext4 in order to minimize the impact of caching. As expected, throughput hovered higher around 57MB/s for NFSv3 and 138MB/s for ext4.

Thus, in this simple setup, *PrivateFS* features a modest throughput when compared to unsecured file systems. This is the inherent cost of achieving privacy.




Fig. 15. IOzone write performance vs. increasing file system written data. Write throughput is slowly decreasing, as expected due to the inherent slowdown in the (increasingly larger) ORAM.

8. DISCUSSION ON USING MAKE-UP ON THE PIG

A few notes are in order to outline how performance can be scaled with increasing resources thrown at the problem.

On Using SSDs and multiple disks to decrease latency. Deploying (multiple, Sections 2.6, 6.3) zero-latency media server-side would significantly impact performance. This is straightforward and relatively uninteresting research-wise. For example, deploying SSDs would shift bottlenecks and immediately increase throughput by an order of magnitude or more.

On Using Larger Block Sizes to Pretend Higher Throughput (while wasting network bandwidth). *PrivateFS* has been benchmarked with 4KB blocks. One can straightforwardly increase their size by considering larger blocks and obtain an (artificially inflated) “higher throughput” up to the maximum available bandwidth. E.g., going from 4KB to 128KB blocks could increase “throughput” by another order of magnitude or more at the expense of wasted bandwidth. Similarly, we felt this is not interesting research-wise and should be decided on an application-specific basis.

On Using Millions of CPUs to Claim Better “Performance”. Finally, to eliminate bottlenecks, large amounts of server-side resources can be thrown at the problem to speed up things even further. We posit that this is also not interesting – what ulti-

mately counts is the usable bang for the buck achieved, i.e., in this case the throughput (at some fixed block-size) achieved per compute cycle spent server-side. Otherwise, results become meaningless – e.g., if instead of the test setup above we were to deploy a large number of compute cores, performance would increase almost linearly up to network saturation. With careful fine-tuning throughputs of hundreds of Mbps can be achieved without any changes in the base protocol. However, from a security point of view or research-wise in general this is not interesting but should be pursued in industrial R&D.

9. CONCLUSION

This paper includes mechanisms for secure parallel querying of existing ORAMs to increase throughput, a generalization of ORAM de-amortization, and implementation of an efficient ORAM based on these techniques, performing a transaction per second on a terabyte database in an average-latency network (a first). An implementation of *PrivateFS*, the first oblivious networked file system, is provided.

Ongoing work includes support for real time file system sharing between remote clients and finding alternate approaches to increasing throughput, such as constant-round-trip ORAMs [Williams and Sion 2012] and de-amortization of the top level.

ACKNOWLEDGMENTS

The authors would like to thank Alin Tomescu for the FUSE implementation work.

REFERENCES

- BONEH, D., KUSHILEVITZ, E., OSTROVSKY, R., AND SKEITH, III., W. E. 2007. Public key encryption that allows PIR queries. In *Proceedings of the 27th annual international cryptology conference on Advances in cryptology*. CRYPTO'07. Springer-Verlag, Berlin, Heidelberg, 50–67.
- BONEH, D., MAZIÉRES, D., AND POPA, R. A. 2011. Remote oblivious storage: Making Oblivious RAM practical. Tech. rep., MIT. MIT-CSAIL-TR-2011-018 March 30, 2011.
- CAMENISCH, J., NEVEN, G., AND SHELAT, A. 2007. Simulatable adaptive oblivious transfer. In *EUROCRYPT '07: Proceedings of the 26th annual international conference on Advances in Cryptology*.
- CAPPS, D. 2012. IOzone Filesystem Benchmark. Online at <http://www.iozone.org/>.
- CHOR, B., GOLDREICH, O., KUSHILEVITZ, E., AND SUDAN, M. 1995. Private information retrieval. In *IEEE Symposium on Foundations of Computer Science*. 41–50.
- CNN. 2006. Feds seek Google records in porn probe. Online at <http://www.cnn.com>.
- COULL, S., GREEN, M., AND HOHENBERGER, S. 2009. Controlling access to an oblivious database using stateful anonymous credentials. In *International Conference on Practice and Theory in Public Key Cryptography (PKC)*.
- GARTNER, INC. 1999. Server Storage and RAID Worldwide. Tech. rep., Gartner Group/Dataquest. www.gartner.com.
- GASARCH, W. A WebPage on Private Information Retrieval. Online at <http://www.cs.umd.edu/~gasarch/pir/pir.html>.
- GASARCH, W. A survey on private information retrieval. Online at <http://citeseer.ifi.unizh.ch/gasarch04survey.html>.
- GOLDREICH, O. AND OSTROVSKY, R. May 1996. Software protection and simulation on Oblivious RAMs. *Journal of the ACM* 45, 431–473.
- GOODRICH, M. AND MITZENMACHER, M. 2011. MapReduce Parallel Cuckoo Hashing and Oblivious RAM Simulations. In *ICALP*.
- GOODRICH, M., MITZENMACHER, M., OHRIMENKO, O., AND TAMASSIA, R. 2011a. Oblivious RAM Simulation with Efficient Worst-Case Access Overhead. In *ACM Cloud Computing Security Workshop (CCSW)*.
- GOODRICH, M. T. 2010. Randomized shellsort: A simple oblivious sorting algorithm. In *SODA*.
- GOODRICH, M. T., MITZENMACHER, M., OHRIMENKO, O., AND TAMASSIA, R. 2011b. Oblivious storage with low I/O overhead. *CoRR abs/1110.1851*.
- GOODRICH, M. T., MITZENMACHER, M., OHRIMENKO, O., AND TAMASSIA, R. 2012. Privacy-preserving group data access via stateless oblivious RAM simulation. In *SODA*.

- HENK, C. AND SZEREDI, M. 2012. FUSE: Filesystem in Userspace. Online at <http://sourceforge.net/projects/fuse>.
- ISHAI, Y. AND KUSHILEVITZ, E. 1997. Private simultaneous messages protocols with applications. *Theory of Computing Systems, Israel Symposium on the 0*, 174.
- JOHNSON, A., SYVERSON, P., DINGLEDINE, R., AND MATHEWSON, N. 2011. Trust-based anonymous communication: Adversary models and routing algorithms. In *ACM Computer and Communications Security Conference (CCS 2011)*.
- KUSHILEVITZ, E., LU, S., AND OSTROVSKY, R. 2012. On the (in)security of hash-based oblivious RAM and a new balancing scheme. In *SODA*, Y. Rabani, Ed. SIAM, 143–156.
- LI, J., KROHN, M., MAZIÈRES, D., AND SHASHA, D. 2004. Secure untrusted data repository (SUNDR). In *OSDI'04: Proceedings of the 6th conference on Symposium on Operating Systems Design & Implementation*. USENIX Association, Berkeley, CA, USA, 9–9.
- LIPMAA, H. AND ZHANG, B. 2010. Two new efficient PIR-writing protocols. In *Proceedings of the 8th international conference on Applied cryptography and network security*. ACNS'10. Springer-Verlag, Berlin, Heidelberg, 438–455.
- NAOR, M. AND PINKAS, B. 1999. Oblivious transfer with adaptive queries. In *Proceedings of the 19th Annual International Cryptology Conference on Advances in Cryptology*. CRYPTO '99. Springer-Verlag, London, UK, 573–590.
- OLUMOFIN, F. AND GOLDBERG, I. 2011. Revisiting the computational practicality of private information retrieval. In *In Financial Cryptography and Data Security '11*.
- OSTROVSKY, R. AND SHOUP, V. 1997. Private information storage (extend abstract). In *IN PROCEEDINGS OF STOC*. ACM Press, 294–303.
- PAGH, R. AND RODLER, F. F. 2004. Cuckoo hashing. *J. Algorithms* 51, 122–144.
- PINKAS, B. AND REINMAN, T. 2010. Oblivious RAM revisited. In *CRYPTO*. 502–519.
- RABIN, M. O. 2005. How to exchange secrets with oblivious transfer. Cryptology ePrint Archive, Report 2005/187. <http://eprint.iacr.org/>.
- SCHOLL, T. 2009. ext2fuse: ext2 filesystem in userspace. Online at <http://sourceforge.net/projects/ext2fuse/>.
- SHI, E., CHAN, T.-H. H., STEFANOV, E., AND LI, M. 2011. Oblivious RAM with $o(\log n)^3$ worst-case cost. In *ASIACRYPT*. 197–214.
- SION, R. AND CARBUNAR, B. 2007. On the computational practicality of private information retrieval. In *Proceedings of the Network and Distributed Systems Security (NDSS) Symposium*.
- STEFANOV, E., SHI, E., AND SONG, D. 2012. Towards Practical Oblivious RAM. In *Proceedings of the Network and Distributed System Security (NDSS) Symposium*.
- WANG, S., DING, X., DENG, R. H., AND BAO, F. 2006. Private information retrieval using trusted hardware. In *Proceedings of the European Symposium on Research in Computer Security ESORICS*. 49–64.
- WILLIAMS, P. AND SION, R. 2008. Usable PIR. In *Proceedings of the 2008 Network and Distributed System Security (NDSS) Symposium*.
- WILLIAMS, P. AND SION, R. 2012. Round-Optimal Access Privacy on Outsourced Storage. In *ACM Conference on Computer and Communication Security CCS*.
- WILLIAMS, P., SION, R., AND CARBUNAR, B. 2008. Building castles out of mud: practical access pattern privacy and correctness on untrusted storage. In *ACM Conference on Computer and Communications Security*. 139–148.
- WILLIAMS, P., SION, R., AND SOTAKOVA, M. 2011. Practical oblivious outsourced storage. *ACM Trans. Inf. Syst. Secur.* 14, 20:1–20:28.
- WILLIAMS, P., SION, R., AND TOMESCU, A. 2012. PrivateFS: A Parallel Oblivious File System. In *ACM Conference on Computer and Communication Security CCS*.

Received August 2012; revised November 2012; accepted January 2013