Efficient Consistency Proofs for Generalized Queries on a Committed Database*

Rafail Ostrovsky[1], Charles Rackoff[2], and Adam Smith[3]

[1] UCLA Dept. of Computer Science, Los Angeles, CA, USA. rafail@cs.ucla.edu
[2] University of Toronto, Toronto, Ontario, Canada. rackoff@cs.toronto.edu
[3] MIT Computer Science and AI Lab, Cambridge, MA, USA. asmith@csail.mit.edu

Abstract. A *consistent query protocol* (CQP) allows a database owner to publish a very short string c which *commits* her and everybody else to a particular database D, so that any copy of the database can later be used to answer queries and give short proofs that the answers are consistent with the commitment c. Here *commits* means that there is at most one database D that anybody can find (in polynomial time) which is consistent with c. (Unlike in some previous work, this strong guarantee holds even for owners who try to cheat while creating c.) Efficient CQPs for membership and one-dimensional range queries are known [4, 11, 16]: given a query pair $a, b \in \mathbb{R}$, the server answers with all the keys in the database which lie in the interval $[a, b]$ and a proof that the answer is correct.

This paper explores CQPs for more general types of databases. We put forward a general technique for constructing CQPs for any type of query, assuming the existence of a data structure/algorithm with certain inherent robustness properties that we define (called a *data robust algorithm*). We illustrate our technique by constructing an efficient protocol for *orthogonal range queries*, where the database keys are points in $\mathbb{R}^d$ and a query asks for all keys in a rectangle $[a_1, b_1] \times \ldots \times [a_d, b_d]$. Our data-robust algorithm is within a $O(\log N)$ factor of the best known standard data structure (a range tree, due to Bentley [2]).

We modify our protocol so that it is also *private*, that is, the proofs leak no information about the database beyond the query answers. We show a generic modification to ensure privacy based on zero-knowledge proofs, and also give a new, more efficient protocol tailored to hash trees.

1 Introduction

Informally, a *consistent query* protocol (CQP) allows a database owner to publish a short string c which *commits* her to a particular database D, so that she can later answer queries and give short proofs that her answers are consistent with D. Here *commits* means that she cannot change her mind about D — there is at most one database she can find (in polynomial time) which is consistent with c (e.g. c could be a secure hash of D). Similarly, she can only find valid proofs for query answers which are consistent

* Preliminary work done during the summer of 2000 when all authors were visiting/working at Telcordia Technologies. Preliminary version appeared as MIT LCS Technical Report TR-887, Feb. 2003 [20]. Work of the first author at UCLA is partially supported by a gift from Teradata.

with D. The challenge is to make both the commitment and the proofs of consistency as short and simple as possible.

One may also require *privacy* – that is, the proofs of consistency should not leak any information on the database beyond the query answers. Privacy is important, for example, in settings in which query answers are sold individually, or in which the database contains personal data. Adding this requirement to a CQP brings it much closer to the traditional cryptographic notion of a commitment scheme.

Below, we discuss relevant related work and then describe our results in detail.

Related Work We discuss the related work in the context of cryptographic commitment protocols. These have been studied extensively, and part of our contribution is to tie them in to an algorithmic point of view. A commitment protocol allows Alice to put a value a in a virtual envelope and hand it to Bob. Bob learns nothing about the value (*hiding*), but Alice can later open the envelope, without being able to reveal a different value a' (*binding*).

Commitment Schemes for Large Datasets. The notion of commitment has been generalized considerably to allow revealing only partial information about the committed data, using very little communication. Merkle [17] proposed the following protocol for committing to a list of N values $a_1, ..., a_N$: Pick a collision-resistant hash-function[4] H (say from $2k$ bits to k bits), pair up inputs $(a_1, a_2), \ldots, (a_{N-1}, a_N)$ and apply H to each pair. Now, pair up the resulting hash values and repeat this process, constructing a binary tree of hash values, until you get to a single root of length k. If the root of the tree is published (or sent to Bob by Alice), the entire collection of values is now committed to, though not necessarily hidden—we discuss hiding further below. To reveal any particular value a_i, Alice can reveal a path from the root to a_i together with all the siblings along the path. This requires only $k \log N$ bits. This idea has many cryptographic applications, including efficient signature schemes [17, 5], efficient zero-knowledge arguments [10, 1] and computationally sound proofs [15].

Recently Buldas, Laud and Lipmaa [3], Kilian [11] and Micali and Rabin [16] independently generalized this idea to allow committing to a *set* of values. The server produces a short commitment to her set of $(\mathsf{key}, \mathsf{value})$ pairs which is made public. When a client makes a *membership query* (i.e. "do you have an entry with key x?"), the server returns the answer along with a short proof of consistency. (We call a scheme for this task a CQP for membership queries.) A very similar data structure (again, a Merkle tree) also allows one to also answer one-dimensional *range queries*, e.g. "What keys lie between x and y?" [4, 11, 16]. Merkle trees were subsequently modified to allow efficient updates by changing the structure to resemble a skip list [12]. Our work generalizes these ideas to more complex queries and data structures, and provides rigorous proofs of security.

Protocols with a Trusted Committer. There is substantial work on *authenticated data structures* [18], which allow one to guarantee the consistency of many replicated copies of a database. That work tackles a different problem from ours, since it assumes that the commitment phase is always performed honestly. As with ordinary commitments, assuming a trusted committer allows for simpler, more efficient solutions than are known

[4] A hash function family $H_\kappa(\cdot)$ is *collision-resistant* if no poly-time algorithm given κ can find a pair of inputs that map to the same output for a randomly chosen key κ (see Section 2).

in our (general) setting; the generic construction in this paper can be viewed as a more robust version of the generic constructions of authenticated data structures [18, 13, 8]. For discussions of the dangers of assuming a trusted committer, see [3, 12].

Privacy for Committed Databases. Micali, Rabin and Kilian [14] show how to prove consistency of answers to membership queries while also hiding information about unanswered queries. They require that consistency proofs leak nothing about the database except the query answer—not even the size of the database. (They call the primitive a *zero-knowledge set*.) They give an efficient protocol based on the DDH assumption, with proof length $O(k \log M)$ where M is an upper bound on the set size (k is the output length of the hash function). Our techniques achieve this result with $poly(k)$ communication under more general assumptions and for more general types of queries. Subsequent to our work, [9] achieved the results of [14] based on general assumptions.

Our Contributions This paper considers CQPs for types of queries beyond simple membership and range queries. We give a general framework for designing such protocols based on query algorithms with a certain robustness property, and illustrate our paradigm for *orthogonal range queries*, constructing protocols with an $O(k \log N)$ overhead over the fastest known standard query alogrithms. We also show how to make the protocols *private* without too much loss of efficiency.

A general paradigm for CQP*s.* We introduce *data-robust algorithms* (DRAs). These are search algorithms (paired with data structures) which are robust against corruptions of the data by an unbounded, *malicious* adversary: for any input—essentially, an arbitrary string— the algorithm will answer all queries consistently with one (valid) database.

Assuming the existence of collision-resistant hash functions, any DRA which accesses memory via pointers can be transformed into a consistent query protocol whose (non-interactive) consistency proofs have length at most $O(kT)$, where k is the output size of the hash function and T is the running time of the DRA.

CQP *for Orthogonal Range Queries.* We present a consistent query protocol scheme that allows efficient orthogonal range queries in d dimensions. That is, the database consists of tuples $(\mathsf{key}_1, ..., \mathsf{key}_d, \mathsf{value})$, a query consists of d intervals $[a_1, b_1], \ldots, [a_d, b_d]$, and an answer is the set of all database elements whose keys lie inside the corresponding hypercube. The server not only proves that it has provided all the points in the database which match the query, but also that no others exist.

Our consistency proofs have size $O(k(m+1)\log^d N)$, where N is the database size, k is the security parameter, and m is the number of keys in the database satisfying the query (the computation required is $O((m+1)\log^d N)$ hash evaluations). For range queries on a single key, our construction reduces essentially to that of [4, 16, 11].

Our protocol is obtained by first constructing a DRA based on range trees, a classic data structure due to Bentley [2]. Existing algorithms (in particular, the authenticated data structures of [13]) do not suffice, as inconsistencies in the data structure can lead to inconsistent query answers. Instead, we show how local checks can be used to ensure that all queries are answered consistently with a single database. For d-dimensional queries, the query time is $O((m+1)\log^d N)$, where m is the number of hits for the query and N is the number of keys in the database. This is within $\log N$ of the best known (non-robust) data structure.

Achieving Privacy Efficiently. Consistent query protocols will, in general, leak information about the database beyond the answer to the query. It is possible to add privacy to any CQP using generic techniques: one can replace the proof of consistency π with a zero-knowledge proof of knowledge of π. Surprisingly, this leads to schemes with better asymptotic communication complexity, namely $O(poly(k))$. This generic transformation can hide the size of the database, as in [14].

However, the use of NP reductions and probabilistically checkable proofs in generic constructions means that the advantages only appear for extremely large datasets. We give a simpler zero-knowledge protocol tailored to Merkle trees, which does not hide the size of the database. The crux of that protocol is to avoid NP reductions when proving zero-knowledge statements about values of the hash function, and so the result is called an *explicit-hash Merkle tree*. As a sample application, we show how this protocol can be used to add privacy to one-dimensional range trees.

Organization. Section 2 formally defines CQPs. Section 3 explains data-robust algorithms, and the transformation from DRAs to CQPs. Section 4 gives our DRA for orthogonal range queries. Section 5 discusses techniques for making CQPs private. Due to lack of space, all proofs are deferred to the full version.

2 Definitions

A function $f(k)$ is *negligible* in a parameter k if $f(k) \in O(\frac{1}{k^c})$ for all integers $c > 0$. Assigning the (possibly randomized) output of algorithm A on input x to variable y is denoted by $y \leftarrow A(x)$. An important component is collision-resistant hash functions (CRHF). This is a family of length-reducing functions (say from $3k$ bits to k bits) such that given a randomly chosen function h from the family, it is computationally infeasible to find a collision, i.e. $x \neq y$ with $h(x) = h(y)$.

Consistent Query Protocols A query structure is a triple $(\mathcal{D}, \mathcal{Q}, Q)$ where $\mathcal{D}$ is a set of *valid* databases, $\mathcal{Q}$ is a set of possible queries, and Q is a rule which associates an answer $a_{q,D} = Q(q, D)$ with every query/database pair $q \in \mathcal{Q}, D \in \mathcal{D}$.

In a CQP, there is a server who, given a database, produces a commitment which is made public. Clients then send queries to the server, who provides the query answer along with a proof of consistency of the commitment. There may also be a public random string to be provided by a trusted third party. Though we formulate our definitions in that context, our constructions mostly do not require the third party.

Syntactically, a query protocol consists several probabilistic poly-time (PPT) algorithms: (1) a server setup algorithm $\mathcal{S}_s$, which takes the database D, a security parameter 1^k and any public randomness σ, and outputs the commitment c and some internal state information *state*; (2) an answering algorithm $\mathcal{S}_a$ which takes *state* and a query q and returns an answer-proof pair (a, π); (3) a client verification algorithm which takes a triple (c, q, a, π) and outputs "accept" or "reject;" (4) an algorithm Σ for sampling the public random string.

Definition 1. *A query protocol is* consistent *if it is complete and sound:*

- *Completeness: For every valid database D and query q, if setup is performed correctly then with overwhelming probability, $\mathcal{S}_a$ outputs both the correct answer and a proof which is accepted by $\mathcal{C}$. Formally, for all $q \in \mathcal{Q}$ and for all $D \in \mathcal{D}$,*

$$\Pr[\sigma \leftarrow \Sigma(1^k); (c, \text{state}) \leftarrow \mathcal{S}_s(\sigma, D); (a, \pi) \leftarrow \mathcal{S}_a(q, \text{state}) : \\ \mathcal{C}(\sigma, c, q, a, \pi) = \text{“accept” and } a = Q(q, D)] \geq 1 - negl(k)$$

- *(Computational) Soundness: For every non-uniform PPT adversary $\tilde{\mathcal{S}}$: run $\tilde{\mathcal{S}}$ to obtain a commitment c along with a list of triples (q_i, a_i, π_i). We say $\tilde{\mathcal{S}}$* acts consistently *if there exists $D \in \mathcal{D}$ such that $a_i = Q(q_i, D)$ for all i for which π_i is a valid proof. The protocol is* sound *if all PPT adversaries $\tilde{\mathcal{S}}$ act consistently. Formally:*

$$\Pr[\sigma \leftarrow \Sigma(1^k); \big(c, (q_1, a_1, \pi_1), \ldots, (q_t, a_t, \pi_t)\big) \leftarrow \tilde{\mathcal{S}}; b_i \leftarrow \mathcal{C}(\sigma, c, q_i, a_i, \pi_i) : \\ \exists \tilde{D} \text{ such that } (a_i = Q(q_i, \tilde{D}) \text{ or } b_i = 0) \text{ for all } i] \geq 1 - negl(k)$$

Privacy Informally, we require that an adversarial client interacting with an (honest) server learn no more information from the answer/proof pairs he receives than what he gets from the answers alone. specifically, a simulator who has access only to the query answers should be able to give believable-looking proofs of consistency. The definition comes from [11, 16, 14], though we use a cleaner formulation due to [9].

Definition 2 (Computational privacy). *A consistent query protocol for $(\mathcal{D}, \mathcal{Q}, Q)$ is private if there exists a PPT simulator* Sim*, such that for every non-uniform PPT adversary $\tilde{\mathcal{C}}$, the outputs of the following experiments are computationally indistinguishable:*

$\sigma \leftarrow \Sigma(1^k)$,	$\sigma', c', state_{\mathsf{Sim}} \leftarrow \mathsf{Sim}(1^k)$,
$(D, state_{\tilde{\mathcal{C}}}) \leftarrow \tilde{C}(\sigma)$,	$(D, state_{\tilde{\mathcal{C}}}) \leftarrow \tilde{C}(\sigma')$,
$(c, state) \leftarrow \mathcal{S}_s(\sigma, D)$,	
Output $z \leftarrow \tilde{\mathcal{C}}^{\mathcal{S}_a(\cdot, state)}(c, state_{\tilde{\mathcal{C}}})$	*Output* $z \leftarrow \tilde{\mathcal{C}}^{\mathsf{Sim}(\cdot, state_{\mathsf{Sim}}, Q(\cdot, D))}(c', state_{\tilde{\mathcal{C}}})$

Here $\tilde{\mathcal{C}}^{\mathcal{O}(\cdot)}$ denotes running $\tilde{\mathcal{C}}$ with oracle access to $\mathcal{O}$. The simulator Sim *has access to a query oracle $Q(\cdot, D)$, but* asks only queries which are asked to Sim by $\tilde{\mathcal{C}}$.

Hiding Set Size. In general, a private protocol should not leak the size of the database [14]. Nonetheless, for the sake of efficiency we will sometimes leak a *polynomial* upper bound T on the database size, and call the corresponding protocols *size-T-private* [11]. This can be reflected in the definition by giving the simulator an upper bound T on the size of D as an additional input. One recovers the original definition by letting T be exponential, e.g. $T = 2^k$.

Interactive proofs. The definitions extend to a model where consistency proofs are interactive (although the access of the simulator to the adversarial client is more tricky).

3 Data-robust algorithms and consistent query protocols

In this section, we describe a general framework for obtaining secure consistent query protocols, based on designing efficient algorithms which are “data-robust”. Assuming the availability of a collision-resistant hash function, we show that any such algorithm which accesses its input by “following” pointers can be transformed into a consistent query protocol whose (non-interactive) consistency proofs have complexity at most proportional to the complexity of the algorithm.

Data-robust algorithms Suppose a programmer records a database on disk in some kind of static data structure which allows efficient queries. Such data structures are often augmented with redundant information, for example to allow searching on two different fields. If the data structure later becomes corrupted, then subsequent queries to the structure might be mutually inconsistent: for example, if entries are sorted on two fields, some entry might appear in one of the two structures but not the other. A data-robust algorithm prevents such inconsistencies.

Suppose we have a query structure $(\mathcal{D}, \mathcal{Q}, Q)$. A data-robust algorithm (DRA) for these consists of two polynomial-time[5] algorithms (T, A): First, a setup transformation $T : D \to \{0,1\}^*$ which takes a database D and makes it into a static data structure (i.e. a bit string) $S = T(D)$ which is maintained in memory. Second, a query algorithm A which takes a query $q \in \mathcal{Q}$ and an arbitrary "structure" $\tilde{S} \in \{0,1\}^*$ and returns an answer. The structure $\tilde{S}$ needn't be the output of T for any valid database D.

Definition 3. *The algorithms* (T, A) *form a* data-robust algorithm *for* $(\mathcal{D}, \mathcal{Q}, Q)$ *if:*

- **Termination** *A terminates in polynomial time on* all *input pairs* $(q, \tilde{S})$*, even when* $\tilde{S}$ *is not an output from* T*.*
- **Soundness** *There exists a function* $T^* : \{0,1\}^* \to \mathcal{D}$ *such that* for all *inputs* $\tilde{S}$*, the database* $D = T^*(\tilde{S})$ *satisfies* $A(q, \tilde{S}) = Q(q, D)$ *for all queries* q*.*
  (There is no need to give an algorithm for T^**; we only need it to be well-defined.)*
- **Completeness** *For all* $D \in \mathcal{D}$*, we have* $T^*(T(D)) = D$*.*
  (That is, on input q *and* $T(D)$*, the algorithm* A *returns the correct answer* $Q(q, D)$*.)*

We only allow A *read* access to the data structure (although the algorithm may use separate space of its own). Moreover, A is *stateless*: it shouldn't have to remember any information between invocations.

The running time of A. There is a naive solution to the problem of designing a DRA: A could scan the corrupted structure $\tilde{S}$ in its entirety, decide which database D this corresponds to, and answer queries with respect to D. The problem, of course, is that this requires at least linear time *on every query* (recall that A is stateless). Hence the task of designing robust algorithms is most interesting when there are natural *sub-linear* time algorithms; the goal is then to maintain efficiency while also achieving robustness. In our setting, efficiency means the running-time of the algorithm A on *correct* inputs, in either a RAM or pointer-based model. On incorrect inputs, an adversarially-chosen structure could, in general, make A waste time proportional to the size of the structure $\tilde{S}$; the termination condition above restricts the adversary from doing too much damage (such as setting up an infinite loop, etc).

Constructing consistent query protocols from DRAs Given a DRA which works in a pointer-based memory model, we can obtain a cryptographically secure consistent query protocol of similar efficiency. Informally, a DRA is pointer-based if it operates by following pointer in a directed acyclic graph with a single source (see the full version for details). Most common search algorithms fit into this model.

[5] We assume for simplicity that the algorithms are deterministic; this is not strictly necessary.

Proposition 1. *(Informally) Let (T, A) be a* DRA *for query structure $(\mathcal{D}, \mathcal{Q}, Q)$ which fits into the pointer-based framework described above. Suppose that on inputs q and $T(D)$ (correctly formed), the algorithm A examines $b(q, D)$ memory blocks and a total of $s(q, D)$ bits of memory, using $t(q, D)$ time steps. Assuming the availability of a public collision-resistant hash function, there exists a consistent query protocol for $(\mathcal{D}, \mathcal{Q}, Q)$ which has proof length $s(q, D) + kb(q, D)$ on query q. The server's computation on each query is $O(s(q, D) + t(q, D) + kb(q, D))$.*

To get a consistent query protocol from a DRA, we essentially build a Merkle tree (or graph, in fact) which mimics the structure of the data, replacing pointers with hashes of the values they point to. The client runs the query algorithm starting from hash of the unique source in the graph (that hash value is the public commitment). When the query algorithm needs to follow a pointer, the server merely provides the corresponding pre-image of the hash value.

4 Orthogonal Range Queries

In the case of join queries, a database D is a set of key/value pairs (entries) where each key is a point in $\mathbb{R}^d$, and each query is a rectangle $[a_1, b_1] \times \cdots \times [a_d, b_d]$. Note that these are also often called *(orthogonal) range queries*, and we shall adopt this terminology here for consistency with the computational geometry literature. For concreteness, we consider the two-dimensional case; the construction naturally extends to higher dimensions. In two dimensions, each query q is a rectangle $[a_1, b_1] \times [a_2, b_2]$. The query answer $Q(q, D)$ is a list of all the entries in D whose key $(\mathsf{xkey}, \mathsf{ykey})$ lies in q.

4.1 A data-robust algorithm for range queries

Various data structures for efficient orthogonal range queries exist (see [7] for a survey). The most efficient (non-robust) solutions have query time $O((m + 1) \log^{d-1} N)$ for d-dimensional queries. We recall the construction of *multi-dimensional range trees* (due to Bentley [2]), and show how they can be queried robustly. The query time of the robust algorithm is $O((m + 1) \log^d N)$. It is an interesting open question to find a robust algorithm which does as well as the best non-robust algorithms.

One-dimensional range trees Multidimensional range trees are built recursively from one-dimensional range trees (denoted 1-DRT), which were also used by [4, 16, 11]. In a 1-DRT, $(\mathsf{key}, \mathsf{value})$ pairs are stored in sorted order as the leaves of a (minimum-height) binary tree. An internal node n stores the minimum (denoted a_n) and maximum (denoted b_n) keys which appear in the subtree rooted at n. For a leaf l, we take $a_l = b_l$ to be the value of the key_l key stored at l. Additionally, leaves store the value value_l associated to key_l.

Setup. Given a database $D = \{(\mathsf{key}_1, \mathsf{value}_1), \ldots, (\mathsf{key}_N, \mathsf{value}_N)\}$, the setup transformation T_{1DRT} constructs a minimum-height tree based on the sorted keys. All the intervals $[a_n, b_n]$ can be computed using a single post-order traversal.

Algorithm 1. $A_{\mathsf{1DRT}}(\ [a,b],\ n,\ )$
Input: a target range $[a,b]$, a node n in a (possibly misformed) 1-DRT.
Output: a set of $(\mathsf{key},\mathsf{value})$ pairs.

1. **if** n is not properly formed (i.e. does not contain the correct number of fields) **then** return $\emptyset$
2. **if** n is a leaf: **if** $a_n = b_n = \mathsf{key}_n$ and $\mathsf{key}_n \in [a,b]$, **then** return $\{(\mathsf{key}_n, \mathsf{value}_n)\}$ **else** return $\emptyset$
3. **if** n is an internal node:
   - $l \leftarrow \mathsf{left}_n$, $r \leftarrow \mathsf{right}_n$
   - **if** $a_n = a_l \leq b_l < a_r \leq b_r = b_n$ **then** return $A_{\mathsf{1DRT}}\ (\ [a,b],\ l) \cup A_{\mathsf{1DRT}}\ (\ [a,b],\ r)$
   - **else** return $\emptyset$

Robust queries. It is easy to see that a 1-DRT allows efficient range queries when it is correctly formed (given the root n of a tree and a target interval $[a,b]$, descend recursively to those children whose intervals overlap with $[a,b]$). However, in our setting we must also ensure that the queries return consistent answers even when the data structure is corrupted. The data structure we will use is exactly the one above. To ensure robustness we will modify the querying algorithm to check for inconsistencies.

Assume that we are given a *rooted* graph where all nodes n have an associated interval $[a_n, b_n]$, and all nodes have outdegree either 0 or 2. A *leaf* l is any node with outdegree 0. A leaf is additionally assumed to have to extra fields key_l and value_l. Consider the following definitions:

Definition 4. *A node n is* consistent *if its interval agrees with those of its children. That is, if the children are l and r respectively, then the node is consistent if $a_n = a_l \leq b_l < a_r \leq b_r = b_n$. Moreover, we should have $a_n = b_n$ for a node if and only if it is a leaf.*

A path from the root to a node is consistent *if n is consistent and all nodes on the path to the root are also consistent.*

Definition 5. *A leaf l in a* 1-DRT *is* valid *if there is a consistent path from the root to l.*

In order to query a (possibly misformed) 1-DRT in a robust manner, we will ensure that the query algorithm A returns *exactly* the set of valid leaves whose keys lie in the target range. Thus for any string $\tilde{S}$, the database $T^*(\tilde{S})$ consists of the data at all the valid leaves one finds when $\tilde{S}$ is considered as the binary encoding of a graph.

The following lemma proves that one-dimensional range trees, along with the algorithm A_{1DRT}, form a DRA for range queries.

Lemma 1. *The algorithm A_{1DRT} will return exactly the set of valid leaves whose keys are in the target range. In the worst case, the adversary can force the queries to take time $O(s)$ where s is the total size of the data structure. Conversely, given a collection of N entries there is a tree such that the running time of the algorithm is $O((m+1)\log N)$, where m is the number of points in the target range. This tree can be computed in time $O(N \log N)$ and takes $O(N)$ space to store.*

Algorithm 2. $A_{\mathsf{2DRT}}(\ [a^{(x)}, b^{(x)}] \times [a^{(y)}, b^{(y)}],\ n)$
Input: a target range $[a^{(x)}, b^{(x)}] \times [a^{(y)}, b^{(y)}]$, a node n in a 2-DRT.
Output: a set of (xkey, ykey, value) triples.

1. **if** n is not properly formed (i.e. does not contain the correct number of fields), **then** return $\emptyset$.
2. Check for consistency (if check fails, return $\emptyset$):
   - **if** n is a leaf **then** check $a_n = b_n = \mathsf{key}_n$
   - **if** n is an internal node, **then** check $a_n = a_{\mathsf{left}_n} \le b_{\mathsf{left}_n} < a_{\mathsf{right}_n} \le b_{\mathsf{right}_n} = b_n$
3. (a) **if** $[a_n, b_n] \cap [a^{(x)}, b^{(x)}] = \emptyset$ **then** return $\emptyset$
   (b) **if** $[a_n, b_n] \subseteq [a^{(x)}, b^{(x)}]$ **then**
   - $B \leftarrow A_{\mathsf{1DRT}}(\ [a^{(y)}, b^{(y)}],\ \mathsf{tree}_n)$
   - Remove elements of B for which $\mathsf{xkey} \notin [a_n, b_n]$
   - **if** n is an internal node:
     For each point p in B, check that p is 2-valid in either left_n or right_n.
     If the check fails, remove p from B.
   - Return B

   (c) Otherwise
   - $B \leftarrow A_{\mathsf{2DRT}}\Big(\ ([a^{(x)}, b^{(x)}] \cap [a_{\mathsf{left}_n}, b_{\mathsf{left}_n}]) \times [a^{(y)}, b^{(y)}],\ \mathsf{left}_n\Big)$
     $\cup\ A_{\mathsf{2DRT}}\Big(\ ([a^{(x)}, b^{(x)}] \cap [a_{\mathsf{right}_n}, b_{\mathsf{right}_n}]) \times [a^{(y)}, b^{(y)}],\ \mathsf{right}_n\Big)$
   - Remove elements of B which are not valid leaves of tree_n.
   - Return B

Two-dimensional range trees Here, the database is a collection of triples (xkey, ykey, value), where the pairs (xkey, ykey) are all distinct (they need not differ in both components). The data structure, a two-dimensional range tree (denoted 2-DRT), is an augmented version of the one above. The skeleton is a 1-DRT (called the *primary* tree), which is constructed using the xkey's of the data as its key values. Each node in the primary tree has an attached 1-DRT called its *secondary* tree:

- Each leaf l of the primary tree (which corresponds to a single xkey value $a_l = b_l$) stores all entries with that xkey value. They are stored in the 1-DRT tree_l which is constructed using ykey's as its key values.
- Each internal node n (which corresponds to an interval $[a_n, b_n]$ of xkey's) stores a 1-DRT tree_n containing all entries with xkey's in $[a_n, b_n]$. Again, this "secondary" tree is organized by ykey's.

The setup algorithm T_{2DRT} creates a 2-DRT given a database by first sorting the data on the key xkey, creating a *primary* tree for those keys, and creating a secondary tree based on the ykey for each of nodes in the primary tree. In a 2-DRT, each point is stored d times, where d is its depth in the primary tree. Hence, the total storage can be made $O(N \log N)$ by choosing minimum-height trees.

Searching in a 2-DRT. The natural recursive algorithm for range queries in this structure takes time $O(\log^2 N)$ [7]: Given a target range $[a^{(x)}, b^{(x)}] \times [a^{(y)}, b^{(y)}]$ and an

internal node n, there are three cases: if $[a^{(x)}, b^{(x)}] \cap [a_n, b_n] = \emptyset$, then there is nothing to do; if $[a^{(x)}, b^{(x)}] \supseteq [a_n, b_n]$, then perform a search on the second-level tree attached to n using the target range $[a^{(y)}, b^{(y)}]$; otherwise, recursively explore n's two children.

Based on the natural query algorithm, we can construct a DRA A_{2DRT} by adding the following checks:

- All queries made to the 1-D trees (both primary and secondary) are made robustly following Algorithm 1 (A_{1DRT}), i.e. checking consistency of each explored node.
- For every point which is retrieved in the query, make sure it is present and valid in all the secondary 1-D trees which are on the path to the root (in the primary tree).

Definition 6. *A point* $p = (\mathsf{xkey}, \mathsf{ykey}, \mathsf{value})$ *in a (corrupted)* 2-DRT *is* 2-valid *if*

1. *p appears at a valid leaf in the secondary* 1-DRT tree_l *belonging to a* leaf *l of the primary tree with key value* $\mathsf{xkey} = a_l = b_l$.
2. *For every (primary) node n on the path to l from the root of the primary tree, n is consistent and p is a valid leaf in the (one-dimensional) tree* tree_n.

For robust range queries, we obtain Algorithm 2 (A_{2DRT}). As before, the idea is to return only those points which are 2-valid. Thus, for an arbitrary string $\tilde{S}$, the induced database $T^*_{\mathsf{2DRT}}(\tilde{S})$ is the collection of all 2-valid points in the graph represented by $\tilde{S}$. The following lemma shows that the algorithms $(T_{\mathsf{2DRT}}, A_{\mathsf{2DRT}})$ form a DRA for two-dimensional range queries with query complexity $O((m+1)\log^2 N)$ (where m is the number of points in the target range).

Lemma 2. *Algorithm 2 (A_{2DRT}) will return exactly the set of 2-valid points which are in the target range. On arbitrary inputs, A_{2DRT} terminates in worst-case time $O(L)$, where L is the total size of the data structure.*

Conversely, given a collection of N entries there is a tree such that the running time of the algorithm A_{2DRT} is $O((m+1)\log^2 N)$, where m is the number of points in the target range. This tree can be computed in time $O(N \log^2 N)$ and takes $O(N \log N)$ space to store.

One can use similar ideas to make robust range queries on d-dimensional keys, where $d \geq 2$. The structure is built recursively, as in the 2-D case. Although the algorithm is polylogarithmic for any fixed dimension, the exponent increases:

Lemma 3. *There exists a* DRA *for d dimensional range queries such that queries run in time $O((m+1)\log^d N)$, and the data structure requires $O(N \log^d N)$ preprocessing and $O(N \log^{d-1} N)$ storage.*

Using the generic transformation of the previous section, we obtain:

Theorem 1 (Two dimensions). *Assuming the existence of collision-resistant hash functions, there is a consistent query protocol for two-dimensional range queries with commitment size k and non-interactive consistency proofs of length at most $O(k(m+1)\log^2 N)$, where m is the number of keys in the query range, and k is the security parameter (output size of the hash function).*

For higher dimensions, our construction yields proofs of length $O(k(m+1)\log^d N)$.

5 Privacy for Consistent Query Protocols

One can construct private CQPs (Definition 2) with good asymptotic complexity using generic techniques, as follows: Universal arguments [1] allow one to (interactively) give a zero-knowledge proof of knowledge of an NP statement of arbitrary polynomial length, using only a fixed, $poly(k)$ number of bits of communication. This allows one to handle arbitrary query structures (as long as answering queries takes at most polynomial time). It also hides the set size of the database as in [14], since the universal argument leaks only a super-polynomial bound on the length of the statement being proven.

The generic technique can be made slightly more efficient by starting from a (non-private), efficient CQP, and replacing each proof of consistency π with a zero-knowledge argument of knowledge of π. With a public random string, one can also use non-interactive zero-knowledge proofs. This approach will typically leak some bound on the size N of the database. One can avoid that leakage if the original proofs take time and communication $poly(\log N)$, as with membership and orthogonal range queries. Replacing N with the upper bound 2^k, we once again again get $poly(k)$ communication. (A different proof of the result for membership queries can be found in [9].)

Theorem 2. *(a) Assume that there exists a collision-resistant hash family. For any query structure with polynomial complexity, there exists a* private CQP *with a constant number of rounds of interaction and $poly(k)$ communication.*

(b) Given a public random string, any CQP *with proofs of length $\ell(N)$ can be made size-N-private with no additional interaction at a $poly(k\ \ell(N))$ multiplicative cost in communication, assuming non-interactive zero-knowledge proof systems exist.*

Although the asymptotics are good, the use of generic NP reductions and probabilistically checkable proofs in [1] means that the advantages only appear for extremely large datasets. We therefore construct simpler protocols tailored to Merkle trees.

Explicit-Hash Merkle trees. The Merkle tree commitment scheme leaks information about the committed values, since a collision-resistant function cannot hide all information about its input. At first glance, this seems easy to resolve: one can replace the values a_i at the leaves of the tree with hiding commitments $C(a_i)$. However, there is often additional structure to the values $a_1, ..., a_N$. In CQPs for range queries, they are stored in sorted order. Revealing the path to a particular value then reveals its rank in the data set. The problem gets even more complex when we want to reveal a subset of the values, as we have to hide not only whether paths go left or right at each branching in the tree, but whether or not different paths overlap.

When one attempts to solve the problem using generic zero-knowledge proofs, the main bottleneck lies in proving that $y = H(x)$, given commitments $C(x)$ and $C(y)$—the circuit complexity of the statement is too high. The challenge, then, is to provide zero-knowledge proofs that a set $a'_1, ..., a'_t$ is a subset of the committed values, without going through oblivious evaluation of such complicated circuits. We present a modification of Merkle trees where one reveals all hash-function input-output pairs explicitly, yet retains privacy. We call our construction an *Explicit-Hash Merkle Tree*.

Lemma 4. *Assuming the existence of collision-resistant hash families and homomorphic perfectly-hiding commitment schemes,* explicit-hash Merkle trees *allow proving (in*

zero-knowledge) the consistency of t paths (of length $d = \log N$) using $O(d \cdot t^2 \cdot k^2)$ bits of communication, where k is the security parameter. The protocol uses 5 rounds of interaction. It can be reduced to a single message in the random oracle model.

To illustrate, we apply this idea to the for one-dimensional range queries. The main drawback of the resulting protcol is that the server needs to maintains state between invocations; we denote by t the number of previous queries.

Theorem 3. *There exists an efficient,* size-N-private *consistent query protocol for 1-D range queries. For the t-th query to the server, we obtain proofs of size $O((t + m) \cdot s \cdot k^2 \cdot \log N)$, where s is the maximum length of the keys used for the data, and m is the total number of points returned on range queries made so far. The protocol uses 5 rounds of interaction and requires no common random string. The protocol can be made non-interactive in the random oracle model.*

Acknowledgements. We thank Leo Reyzin and Silvio Micali for helpful discussions.

References

1. B. Barak and O. Goldreich. Universal Arguments. In *Proc. Complexity (CCC) 2002*.
2. J. L. Bentley. Multidimensional divide-and-conquer. *Comm. ACM*, 23:214–229, 1980.
3. A Buldas, P. Laud and H. Lipmaa. Eliminating Counterevidence with Applications to Accountable Certificate Management. *J. Computer Security*, 2002. (Originally in *CCS 2000*.)
4. A. Buldas, M. Roos, J. Willemson. Undeniable Replies to Database Queries. In *DBIS 2002*.
5. I. B. Damgård, T. P. Pedersen, and B. Pfitzmann. On the existence of statistically hiding bit commitment schemes and fail-stop signatures. In *CRYPTO '93*, pp. 22–26.
6. A. De Santis and G. Persiano Zero-Knowledge Proofs of Knowledge Without Interaction (Extended Abstract). In *Proc. of FOCS 1992*, pp. 427-436.
7. J. Goodman and J. O'Rourke, editors. *Handbook of Discrete and Computational Geometry*. CRC Press, 1997.
8. M. T. Goodrich, R. Tamassia, N. Triandopoulos and R. Cohen. Authenticated Data Structures for Graph and Geometric Searching. In *Proc. RSA Conference, Cryptographers' Track*, 2003.
9. A. Healy, A. Lysyanskaya, T. Malkin, L. Reyzin. Zero-Knowledge Sets from General Assumptions. Manuscript, March 2004.
10. J. Kilian. A note on efficient zero-knowledge proofs and arguments. In *24th STOC*, 1992.
11. J. Kilian. Efficiently committing to databases. Technical report, NEC Research, 1998.
12. P. Maniatis and M. Baker. Authenticated Append-only Skip Lists. ArXiv e-print cs.CR/0302010, February, 2003.
13. C. Martel, G. Nuckolls, M. Gertz, P. Devanbu, A. Kwong, S. Stubblebine. A General Model for Authentic Data Publication. Manuscript, 2003.
14. S. Micali, M. Rabin and J. Kilian. Zero-Knowledge Sets. In *Proc. FOCS 2003*.
15. S. Micali. Computationally Sound Proofs. *SIAM J. Computing*, 30(4):1253–1298, 2000.
16. S. Micali and M. Rabin. Accessing personal data while preserving privacy. Talk announcement (1997), and personal communication with M. Rabin (1999).
17. R. Merkle A digital signature based on a conventional encryption function. In *CRYPTO '87*, pp. 369–378, 1988.
18. M. Naor and K. Nissim. Certificate Revocation and Certificate Update. In *7th USENIX Security Symposium*, 1998.
19. M. Naor, M. Yung. Universal One-Way Hash Functions and their Cryptographic Applications. In *21st STOC*, 1989.
20. R. Ostrovsky, C. Rackoff, A. Smith. Efficient Consistency Proofs on a Committed Database MIT LCS Technical Report TR-887. Feb 2003. See http://www.lcs.mit.edu/publications