

Authentic Data Publication over the Internet*

Premkumar Devanbu, Michael Gertz, Charles Martel
Department of Computer Science
University of California
Davis, CA 95616 USA
{devanbu|gertz|martel}@cs.ucdavis.edu

Stuart G. Stubblebine
Stubblebine Consulting, LLC
8 Wayne Blvd
Madison, NJ 07940
stuart@stubblebine.com

Abstract

Integrity critical databases, such as financial information used in high-value decisions, are frequently published over the Internet. Publishers of such data must satisfy the integrity, authenticity, and non-repudiation requirements of clients. Providing this protection over public data networks is an expensive proposition. This is, in part, due to the difficulty of building and running secure systems. In practice, large systems can not be verified to be secure and are frequently penetrated. The negative consequences of a system intrusion at the publisher can be severe. The problem is further complicated by data and server replication to satisfy availability and scalability requirements.

To our knowledge this work is the first of its kind to give general approaches for reducing the trust required of publishers of large databases. To do this, we separate the roles of data owner and data publisher. With a few digital signatures on the part of the owner and no trust required of a publisher, we give techniques based on Merkle hash trees that publishers can use to provide authenticity and non-repudiation of the answer to database queries posed by a client. This is done *without* requiring a key to be held in an on-line system, thus reducing the impact of system penetrations. By reducing the trust required of the publisher, our solution is a step towards the publication of large databases in a scalable manner.

1 Introduction

Consider a financial markets database, used by businesses to make high-value decisions. Examples include financial data about a range of investment vehicles such as stocks, bonds and mortgage-backed securities. These databases might be resold to “financial portals”, which republish the information over insecure public networks. This data will be queried at high rates, for example, by client investment tools. We focus our attention on data which changes infrequently. Examples of relatively static information include financial earnings reports, market history such as individual trades, offers and bids, intra- and inter-day highs and volume, and sometimes financial ratings. We assume extremely high query ratios, with millions of queries executed daily. The data needs to be delivered promptly, reliably and accurately. One approach to this problem is to digitally sign the results to a query. Here, $\sigma_{K_O^{-1}}()$ represents the signature using the private signature key of the data owner, K_O^{-1} .

1. *Client* \longrightarrow *Owner* : *Query*
2. *Owner* \longrightarrow *Client*: *Data, Timestamp, $\sigma_{K_O^{-1}}(\text{Data, Timestamp, Query})$*

*We gratefully acknowledge support from the NSF ITR Program, Grant No. 0085961.

The *client* may thus be assured of the correctness of the answer, since it is in response to the query, the data is adequately recent, and it is signed by the owner (and thus satisfies the non-repudiation requirements). However, there are several issues here: first the *owner* may not be willing or able to provide a reliable and efficient database service to query over this data at the scale and the rates required by clients and republishers of the information. Second, even if the owner is willing and able to provide a service, the owner needs to maintain a high-level of physical and computer security required to defend against attacks. This has to be done to protect the signing key, which must be resident in the database server at all times to sign outgoing data. However, in practice, most large software systems have vulnerabilities. Some of these can be “patched”, (albeit usually only partially) with a great deal of skill, effort, time and expense. Cryptographic techniques like threshold cryptography or using hardware devices to protect the private key help but do not solve the systems vulnerability problem and are generally too expensive to implement in the application domains mentioned above.

A more scalable approach to this problem is to use trusted third-party publishers of the data, in conjunction with a key management mechanism, which allows certification of the signing keys of the publishers to speak for the owner of the data. The database (DB) and updates are provided securely to the publishers.

Owner \longrightarrow *Publishers* $P_i, i = 1 \dots n$: $DB, Timestamp, \sigma_{K_O^{-1}}(DB, Timestamp)$

This is followed by the clients accessing the publishers in a query response style similar to the previous protocol description, with publishers signing answers using $K_{P_i}^{-1}, i = 1 \dots n$.

Presumably, the market for the useful data in *DB* will motivate other *publishers* to provide this service, unburdening *owner* of the need to do so. The *owner* simply needs to sign the data once and distribute it to the *publishers*. As the demand increases, more such *publishers* will emerge, thus making this approach inherently more scalable. However, this approach also suffers from the problem and expense of maintaining a secure system accessible from the Internet. Furthermore, a *client* might worry that a *publisher* engages in deception. The *client* has to find a *publisher* that she can trust, such as a trusted brand-name belonging to a large corporation. The *client* would also have to trust the key-management infrastructure that allowed the keys of the *publishers* to be certified. In addition, she would have to believe that the *publisher* is both competent and careful with site administration and physical access to the database. Particularly, the *client* might well worry that the private signing-key of the *publisher* would have to be resident at the database server, and is therefore vulnerable to attacks. To get a *client* to trust him to provide really valuable data, the *publisher* would have to adopt careful and stringent administrative policies, which would be more expensive for him (and thus also for the *client*). The abovementioned need for trusted publishers would increase the reliance on brand-names, which would also limit the tendency of market competition to reduce costs. Doug Tygar lists this as an important open problem in e-commerce. He asks [26]: “How can we protect re-sold or re-distributed information...?”. In this paper, we present an approach to this problem.

Our Approach. We propose a new approach to certifying responses to queries, by which a completely *untrusted publisher* can provide a *verification object* \mathcal{VO} to a *client*, along with the answer to the query. This \mathcal{VO} is generated by the *publisher*, and provides an independent means of verifying that the answer to a query is correct. The verification object is based on a small number of *digests* (Σ s) that are distributed periodically to the clients by the data *owner*. The digests are bottom-up hashes computed recursively over tree type indexes for the entire set of records in

the *owner*'s database, signed with K_O^{-1} . Answers to queries are various combinations of subsets of these records. Given a query, the *publisher* computes the answer. To show that an answer is correct the *publisher* constructs a \mathcal{VO} using the same tree structures that were used by *owner* to compute the digests. This \mathcal{VO} validates an answer set by providing a hard to forge "proof" which links the answer to the appropriate Σ , which was already signed by the *owner*. Our techniques are founded on cryptographic assumptions regarding the security of hash functions and public-key crypto-systems. Our approach has several features:

1. With the exception of the security at the *client* host, a *client* needs only trust the key of the *owner*; in addition, the *owner* only needs to distribute Σ during update cycles. Hence, the private signing key need not be resident in a "online" machine, and can be better protected, e.g., it can be ensconced in a hardware token that is used only to sign hash digests during updates [27].
2. *Clients* need not trust the *publishers*, nor their keys.
3. For all the techniques we describe, a \mathcal{VO} is always linear in the size of the answer to a query.
4. The \mathcal{VO} guarantees that the answer is correct: it contains all the requested records and no superfluous records.
5. If the *publisher* evaluates queries using the same data structures used by *owner* to compute Σ , *publisher*'s overhead to compute the \mathcal{VO} is relatively minor.

An incorrect answer and \mathcal{VO} will almost always be rejected by the *client*; it is infeasible for *publisher* to forge a correct \mathcal{VO} for an incorrect answer to a *client*'s query.

We offer solutions for authenticity in a specific context. We do not directly address the orthogonal issue of diverse types of access control policies [14], which might restrict queries and updates, nor issues of confidentiality and privacy. Our techniques focus on data managed in relational databases, and we handle only some important classes of queries against relational databases. Our techniques do involve the construction of some complex data structures, although the cost of this can be amortized over more efficient query processing. Some of the approaches are similar to view materialization, e.g., [12, 24], which can enable data warehouses to provide efficient query processing.

Structure of the Paper. In Section 2, we begin with a discussion of the problem space, as we have framed it, and briefly set the context for the preliminary results presented in this paper, and open issues that still remain. Next, after a brief background on relational databases, we describe the essence of our extension of the work of Naor & Nissim [20] to authentic data publication. In Section 4, we give our basic approach. In Section 5, we present some novel ongoing research for using multi-dimensional verification objects. More pragmatic issues and related work are discussed in Section 6. We conclude the paper in Section 7.

2 The Design Space

To begin with, we clarify the problem setting outlined above, and map out the potential design space of solutions. Our general setting can be described thus:

1. The *owner* populates the relations in the database, does some pre-computation on the database to construct data structures that support query-processing (such as B⁺-trees), and computes digests (i.e., Σ s) of these data structures.
2. The *owner* distributes these Σ s securely to *clients* and the database to the *publishers* (e.g., using a private key).
3. A *client* sends a query to a *publisher*. The *publisher* computes the answer q , a verification object \mathcal{VO} for q , and sends both back to the *client*.
4. The *client* verifies the correctness and completeness of the answer q by recomputing Σ using q and the \mathcal{VO} .

Given the general problem of publishing data securely using third parties, different approaches are possible, which lead to different computing and storage costs for the *clients*, *publishers*, and *owners*. One extreme approach, for example, would be for the *owner* to simply deliver the entire database to all *clients* and let them do their own query processing. This approach entails huge data transmission costs and requires a great deal of storage and processing by the *clients*. The other extreme is for the *owner* to pre-compute and sign a whole range of answers to possible queries. The *publishers* would simply cache the queries along with the applicable signatures from the *owner*, and return the signatures with the answers to the queries. From the *client's* perspective, this approach is attractive: each pre-computed answer comes with a *constant*-length verification object (a signature) directly from the *owner*. However, this approach is not practical in general: there are simply too many possible queries *clients* might want to ask.

Our goal is a compromise design: one that does not require pre-computation of arbitrary query answers, nor shipments of the entire database to *clients*. Given a particular database, the *client* can choose from a potentially infinite set of queries (although constrained by the query patterns implemented in the applications at the client sites). We adopt the position that requires a certain amount of effort from all parties: the *owner*, the *publisher*, and the *client*. In all our suggested techniques, the *owner* has to do work $O(n \log^{d-1} n)$ time and space in the size n of the database (where d is a small constant denoting roughly the number of operations, such as selections and projections, in the query). In particular, simple selections over a particular attribute of a single relation only require $O(n)$ preprocessing by the *owner* (where n is the number of tuples in the relation). However, this work is mostly towards the construction of index structures, which is amortized over repeated, efficient query processing. These index structures can either be transmitted to the *publishers*, or (exactly and precisely) recomputed by them. Our query processing algorithms are similar in performance to those of standard databases. The construction of the \mathcal{VO} s are a small constant overhead over the standard query processing algorithms. Finally, and perhaps most importantly, the size of the \mathcal{VO} grows linearly with the size of the answer set, and poly-logarithmically with the size of the database. The verification step itself takes time linear in the size of the \mathcal{VO} .

This view suggests that there may be many other perspectives on this problem of authentic data publication. Differing assumptions on the levels of storage and computational effort expected of the *owner*, *publishers* and *clients* may lead to different viable solutions to this problem. We believe that (in addition to the techniques presented below) many useful “operating points” await discovery.

3 Preliminaries

In this section, we will discuss the basic notions, definitions and concepts necessary for the approach presented in this paper. In Section 3.1, we will present the basic notions underlying relational databases and queries formulated in relational algebra. In Section 3.2, we will discuss the computation and usage of Merkle Hash Trees.

3.1 Relational Databases

The data model underlying our approach is the relational data model (see, e.g., [8, 25]). That is, we assume that the data owner and publishers manage the data using a relational database management system (DBMS). The basic structure underlying the relational data model is the relation. A *relation schema* $R\langle A_1, A_2, \dots, A_n \rangle$ consists of a relation name R and an ordered set of attribute names $\langle A_1, A_2, \dots, A_n \rangle$, also denoted by $schema(R)$. Each attribute A_i is defined on a domain D_i . An *extension* of a relation schema with arity n (also called *relation*, for short) is a finite subset of the Cartesian product $D_1 \times \dots \times D_n$. The extension of a relation schema R is denoted by r . The value of a tuple $t \in r$ for an attribute A_i is denoted by $t.A_i$. We assume that with each relation schema R a set $pk(R) \subseteq \{A_1, \dots, A_n\}$ is associated that designates the primary key. The number of tuples in a relation r is called the *cardinality* of the relation, denoted by $|r|$. The extension of the relation schemas at a particular point in time is called a *database instance* (or *database*).

Queries against a database are formulated in SQL [19, 9]. Such queries are typically translated by the DBMS query processing engine into expressions of the relational algebra for the purpose of query optimization and execution. In this paper, we are mainly concerned with providing verification objects for query results where the queries are formulated as expressions of relational algebra. Those queries are either simple, containing at most one basic operator, or complex, containing a composition of basic operators. The basic operators of the relational algebra are as follows (with R, S being relation schemas):

- Selection (σ): $\sigma_P(r) := \{t \mid t \in r \text{ and } P(t)\}$ where r is a relation (name), P is a condition of the form $A_i \Theta c$ with $A_i \in schema(R)$, $c \in D_i$, and $\Theta \in \{=, \neq, <, >, \leq, \geq\}$.
- Projection (π): $\pi_{A_k, \dots, A_l}(r) := \{\langle t.A_k, \dots, t.A_l \rangle \mid t \in r\}$.
- Cartesian Product (\times): $r \times s := \{tq \mid t \in r \text{ and } q \in s\}$.
- Union Operator (\cup): $r \cup s := \{t \mid t \in s \text{ or } t \in r\}$. R and S must be union compatible.
- Set Difference Operator ($-$): $r - s := \{t \mid t \in r \text{ and } t \notin s\}$. R and S must be union compatible.

Additional operators of the relational algebra, which are typically used in complex queries, are *natural join or equi-join* \bowtie , *condition join or theta-join* \bowtie_C (with C being a condition on join attributes), and *set-intersection* \cap . All these operators can be derived from the above five basic operators.

3.2 Merkle Hash Trees

We describe the computation of a Merkle Hash Tree [18] for a given relation r with relation schema $R = \langle A_1, \dots, A_n \rangle$. For this, assume that $\mathcal{A} = \langle A_i, \dots, A_k \rangle$ is a list of attributes from $schema(R)$.

The Merkle Hash Tree computed is denoted by $MHT(r, \mathcal{A})$. A Merkle Hash Tree is a binary tree with $|r|$ leaves and a hash value $H(i)$ associated with each node i of the tree.

1. First, compute the *tuple hash* h_t for each tuple $t \in r$ (\parallel denotes a separator between arguments):

$$h_t(t) = h(h(t.A_1) \parallel \dots \parallel h(t.A_n))$$

The tuple hash (by the collision resistance of the hash function) functions as a “nearly unique” tuple identifier (for a hash-length of 128 bits, probability of collisions approaches 2^{-128}).

The use of the tuple hash in calculations often enables us to keep tuples themselves confidential, while using the tuple hash for authenticity.

2. Next, compute the Merkle hash tree for relation r . For this, assume that r is sorted by the values of $\mathcal{A} \subseteq \text{schema}(R)$. Given two tuples t, t' in r , the value of t is less than the value of t' , denoted $t.\mathcal{A} < t'.\mathcal{A}$ if for some $l, i \leq l \leq k, t.A_l < t'.A_l$ and $t.A_j = t'.A_j$ for $i \leq j < l$.

We compute the value $H(i)$ associated with node i as follows:

Leaf-nodes : $H(i) = h_t(t_i), i = 1 \dots |r|, 0$ otherwise

Internal-nodes : $H(i) = h(H(L) \parallel H(R))$ where L,R are the left and right children of node i

We note that the “root hash”, denoted $H_{r,\mathcal{A}}$, is a digest of all the data in the Merkle hash tree.

This construction is illustrated in Figure 1. H_i is the tuple hash value of t_i ($i = 1, 2, 3, 4$); e.g., H_{34} is the hash value for tuples 3 and 4. We note that this construction easily generalizes to a higher branching factor greater 2, such as in a B^+ -tree; however, for our presentation here, we primarily use binary trees. In this case, constructing a \mathcal{VO} is a very minor overhead over the query evaluation process itself.

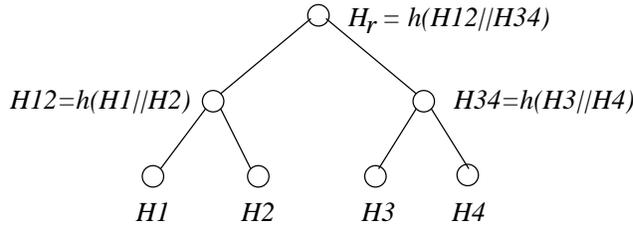


Figure 1: Computation of a Merkle hash tree

Note that by the cryptographic assumption of a collision-resistant hash function, if the correct value of the *parent* is known to the client, the publisher cannot feasibly forge the value of the hash siblings. Our entire approach flows from the signed, correct value of the root of a Merkle tree, just as in the work of Naor & Nissim [20].

Definition 1 (Hash Path) Let i be a leaf node in $MHT(r, \mathcal{A})$ corresponding to a tuple $t_i \in r$. The nodes necessary to compute the *hash path* up to the root hash is denoted as $path(t_i)$. Such a hash path always has the length $\lceil \log_2(|r|) \rceil$ and comprises $2 * \lceil \log_2(|r|) \rceil - 1$ nodes where exactly two nodes are leaf nodes. Of these, only $\lceil \log_2(|r|) \rceil + 1$ need be provided to recompute the value at the root. Hash paths can also be provided for non-leaf nodes. \square

The $\lceil \log_2(|r|) \rceil + 1$ nodes in $path(t_i)$ constitute the verification object \mathcal{VO} showing that t_i is actually in the relation with hash value $H_{r,\mathcal{A}}$ at the root node; the *owner's* signature on the root node certifies its authenticity. Indeed any interior node within the hash tree can be authenticated by giving a path to the root.

Definition 2 (Boundaries) For a given non-empty contiguous sequence $q = \langle t_i, \dots, t_j \rangle$ of leaf nodes in a Merkle Hash Tree $MHT(r, \mathcal{A})$, there are two special leaf nodes $LUB(q)$ and $GLB(q)$ that describe the *lowest upper* and *greatest lower bound* values, respectively, of q and are defined as follows:

- (1) $GLB(q) := \{t \mid t \in r \wedge t.\mathcal{A} < t_i.\mathcal{A} \wedge (\neg \exists t' \in r : t'.\mathcal{A} > t.\mathcal{A} \wedge t'.\mathcal{A} < t_i.\mathcal{A})\}$
- (2) $LUB(q) := \{t \mid t \in r \wedge t.\mathcal{A} > t_j.\mathcal{A} \wedge (\neg \exists t' \in r : t'.\mathcal{A} < t.\mathcal{A} \wedge t'.\mathcal{A} > t_j.\mathcal{A})\}$ □

We assume that both $GLB(q)$ and $LUB(q)$ are singletons. This can easily be accomplished by adding $pk(R)$ to the list \mathcal{A} of attributes by which the leaves in $MHT(r, \mathcal{A})$ are ordered.

Definition 3 (Lowest Common Ancestor) For a given non-empty contiguous sequence $q = \langle t_i, \dots, t_j \rangle$ of leaf nodes in a Merkle Hash Tree $MHT(r, \mathcal{A})$, the lowest common ancestor $LCA(q)$ for q in $MHT(r, \mathcal{A})$ is defined as the root of the minimal subtree in $MHT(r, \mathcal{A})$ that has all tuples in q as leaf nodes. □

This situation is illustrated in Figure 2. Given $LCA(q)$, one can show a hash path $path(LCA(q))$ to the authenticated root hash value. After this is done, (shorter) hash paths from each tuple to $LCA(q)$ can provide evidence of membership of q in the entire tree. This is also useful to build a \mathcal{VO} showing that two nodes occur consecutively in the tree.

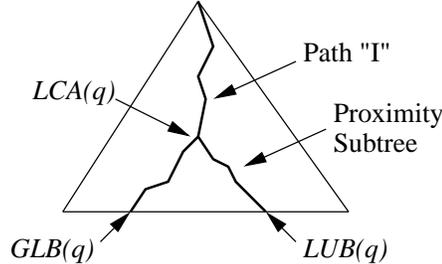


Figure 2: A Merkle tree with a contiguous subrange q , a least common ancestor $LCA(q)$, and upper and lower bounds. Note the verifiable hash path “I” from $LCA(q)$ to the root, and the proximity subtrees (thick lines) for the “near miss” tuples for $LUB(q)$ and $GLB(q)$ which show that q is complete.

Definition 4 (Proximity Subtree) Consider a consecutive pair of tuples (leaf nodes) s, t in $MHT(r, \mathcal{A})$, and their lowest common ancestor, $LCA(\langle s, t \rangle)$. This node, along with the two paths showing that s (respectively, t) is the rightmost (leftmost) element in the left (right) subtree of $LCA(\langle s, t \rangle)$ constitute the “proximity subtree” of s and t , denoted by $ptree(s, t)$. □

Proximity subtrees are used in boundary cases, with $GLBs$ and $LUBs$, i.e., to show a “near-miss” tuple that occurs just outside the answer set lies next to the extremal tuple in the answer set. In this case, it is important to note that by construction, we just need to reveal the relevant attribute

value in the “near-miss” to show that it is indeed a near miss; with just the hash of the other attributes, the tuple hash, and the rest of the proximity tree can be exhibited.

We finally define important properties of the answer set q returned by *publisher*. For this, we assume that *owner* can use a database system to process queries from the *client* in the same fashion as done by the *publisher*.

Definition 5 Assume a query Q issued by a client. Let q_{pub} and q_{owner} denote the query result computed at the data publisher and data owner site, respectively.

q_{pub} is said to be an *inclusive* answer to Q iff $\forall t: t \in q_{pub} \Rightarrow t \in q_{owner}$ holds.

q_{pub} is said to be a *complete* answer to Q iff $\forall t: t \in q_{owner} \Rightarrow t \in q_{pub}$ holds. □

In the following, we also sometimes use the term *correct answer* to refer to the combination of both an inclusive and complete answer.

4 Base level Relational Queries

In this section we outline the computation of \mathcal{VO} s for answers to queries. We illustrate the basic idea behind our approach for selection and projection queries in Section 4.1 and 4.2, respectively. Slightly more complex types of queries (join queries) and set operators are discussed in Sections 4.3 and 4.4.

4.1 Selections

Assume a selection query of the form $\sigma_{A_i \Theta c}(r)$, $c \in D_i$ which determines a result set $q \subseteq r$. Furthermore, assume that the Merkle Hash Tree $MHT(r, A_i)$ has been constructed. For each possible comparison predicate $\Theta \in \{=, \neq, <, >\}$, we show how the *publisher* can construct the \mathcal{VO} , with which the *client* can verify the inclusiveness and completeness of the query answer q . Again, we emphasize that in all the following cases, *if* the Merkle hash tree is constructed by *owner* and *publisher* over the same index structures used for querying, the overhead for constructing the \mathcal{VO} is minor. We first consider the cases for the comparison predicate $\Theta \equiv =$.

Case 1. If $A_i = pk(R)$ and $q \neq \{\}$, then the \mathcal{VO} is just $path(t)$ where $t = q$ is the only tuple that satisfies the selection condition. In this case, the size of the \mathcal{VO} is $O(\log_2 |r|)$.

Case 2. If $A_i = pk(R)$ and $q = \{\}$, then we have to show that no tuple exists that satisfies the selection. For this, we have to provide paths to the two tuples that would “surround” the non-existing tuple. The two tuples are determined by $GLB(q')$ and $LUB(q')$ with $q' = c$. Determining $path(GLB(q'))$ and $path(LUB(q'))$ requires searching the two associated tuples in the leaf nodes of $MHT(r, A_i)$. The proximity subtree $ptree(GLB(q'), LUB(q'))$ provides the required evidence that the answer set is empty. The size of the \mathcal{VO} again is $O(\log_2 |r|)$.

Case 3. A_i is not a primary key and $q \neq \{\}$. The result is a set of tuples which build a contiguous sequence of leaf nodes in $MHT(r, A_i)$. In order to provide a \mathcal{VO} for q , the following approach is taken. First, identify $l := LCA(q \cup GLB(q) \cup LUB(q))$ in $MHT(r, A_i)$, and show a verifiable path from l to the root. Next, identify proximity subtrees showing that $GLB(q)$ ($LUB(q)$) occur consecutively to the smallest (largest) element of q . Now, the entire sub-tree from the elements of the set q to l can be constructed, using the hash values of the tuples in q . This verifies that the entire set occurs contiguously in the leaf nodes of the tree. To construct this subtree and to

verify the root hash on the $LCA(q)$ of this subtree, the length of the \mathcal{VO} is $O(|q| + \log_2(|r|))$. The proximity subtrees establish that no tuples are left out from the answer.

Case 4. If A_i is not a primary key and $q = \{\}$, we can apply the same approach as for case 2.

With these fundamental techniques, proving that the answers to a selection query over a relation r are inclusive and complete is simple. First, using normal query evaluation, the answer set q is determined. Since we only consider simple relational queries here, the answer set q is a contiguous subset of r based on $MHT(r, A_i)$. We also retrieve two additional tuples, $GLB(q)$ (respectively, $LUB(q)$) which is immediately smaller (larger) than the smallest (largest) tuple in r with regard to the answer set q (using the ordering based on A_i). It should be noted that if the answer set is empty, these two will occur consecutively in the Merkle tree. The \mathcal{VO} uses the values in q plus the paths to $GLB(q)$ and $LUB(q)$, so has $O(|q| + \log_2(|r|))$.

For $\Theta \equiv \neq$, we can make the following observation: The answer set to a query of the pattern $\sigma_{A_i \neq c}(r)$ determines at most two contiguous sets of leaf nodes in $MHT(r, A_i)$. For each of these sets, we have to follow an approach similar to case 3 shown above.

For $\Theta \in \{<, >\}$, the scenario is as follows. The answer set to a query of the pattern $\sigma_{A_i < c}(r)$ or $\sigma_{A_i > c}(r)$ determines at most one contiguous set of leaf nodes in $MHT(r, A_i)$. For this set, we have to follow an approach similar to case 3 discussed above. If q is empty, we just have to give the \mathcal{VO} for the tuple $t = \{t \mid t \in r \wedge t.A = \min\{s.A \mid s \in r\}\}$ (analogous for $\Theta \equiv >$). Cases for $\Theta \in \{\leq, \geq\}$ can be handled in a very similar fashion by just shifting the boundaries.

Lemma 6 If *publisher* cannot engineer collisions on the hash function or forge signatures on the root hash value, then if *client* computes the right authenticated root hash value using the \mathcal{VO} and the answer provided for a *selection* query, then the answer is indeed complete and inclusive. ■

The proof can be found in the appendix, along with a more formal discussion of the security assumptions.

4.2 Projections

For queries of the pattern $\pi_{\mathcal{A}}(R), \mathcal{A} \subset \text{schema}(R)$, the projection operator eliminates some attributes of the tuples in the relation r , and then eliminates duplicates from the set of shortened tuples, yielding the final answer q . There may be many different possible projections on a relation R . If the *client* wishes to choose among these dynamically, it may be best to let the *client* perform the projection. The *client* will then also have to eliminate duplicates because these are not automatically eliminated in SQL, unlike the relational algebra. So in this case, the *client* is provided with the whole relation r (or some subset thereof after intermediate selections etc) and the \mathcal{VO} for r before the projection; so the \mathcal{VO} will be linear in size $|r|$, rather than the smaller size $|q|$ of the final result. Note also that the projection may actually mask some attributes that the *client* is not allowed to see; if so, with just the hash of those attributes in each tuple, the *client* can compute the tuple hash, and the \mathcal{VO} for r will still work.

Consider the case where a particular projection $\pi_{\mathcal{A}}(r)$ (which is used often) projects onto attributes and where the values for the projected attributes are poorly distributed, i.e., many tuples have the same values for the attribute(s) \mathcal{A} . In this case, duplicate elimination will remove numerous tuples, leaving behind a small final answer q . Just given the pre-projection tuple set, the *client* would have to do all this work. Now, suppose we have a Merkle tree $MHT(r, \mathcal{A})$, i.e., we assume that the sets of retained attribute values can be mapped to single values (which corresponds to building

equivalence classes) with an applicable total order. In this case, we can provide a \mathcal{VO} for the projection step that is linear in the size of the projected result q .

Each tuple t in the result set q potentially results from a set of tuples $S(t) \subseteq r$. Each tuple in $S(t)$ has identical values for the projected attribute(s) \mathcal{A} . We need to establish that the set q is inclusive (i.e., each t is indeed in the projection) and complete (i.e., no tuple that should be in q is missing). This is accomplished as follows.

1. To show that $t \in q$, we find *any* witness tuple $y \in S(t) \subseteq r$, with the same attribute value for \mathcal{A} , and show the hash path from this tuple to the Merkle Root. This establishes that the tuple t belongs to the result set q . However, the witness tuple is preferably chosen as a “boundary” value, as described next.
2. We must show that there are no tuples missing, say between t and t' , ($t, t' \in q$). To do this, it is enough to show that the sets $S(t)$ and $S(t')$ are both in r and occur immediately next to each other in the sorted order. This is done by showing hash paths, which prove that two “boundary” tuples $y \in S(t)$ and $x \in S(t')$ occur next to each other in the Merkle tree $MHT(r, \mathcal{A})$.

We observe that both the above bits of evidence are provided by displaying at most $2|q|$ hash paths, each of length $\lceil \log_2 |r| \rceil$. This meets our constraint that the size of the authentication evidence be bounded by $O(|q| \log_2 |r|)$.

But how can we assume that a Merkle tree would exist on precisely the intermediate result just prior to the projection? This can be accomplished by multi-dimensional range trees, which are discussed in Section 5.

4.3 Joins

Joins between two or more relations are the most common type of operators in relational algebra used to formulate complex queries. This holds in particular for equi-joins where relations are combined based on primary/foreign key dependencies. There are many alternative realizations for Merkle Tree structures that can be used to provide \mathcal{VO} s for query results computed from joins. In this paper, we focus on pairwise joins of the pattern $R \bowtie_C S$ where C is a condition on join attributes of the pattern $A_R \Theta A_S$, $A_R \in \text{schema}(R)$, $A_S \in \text{schema}(S)$, $\Theta \in \{=, <, >\}$. We assume that the data types underlying A_S and A_R are compatible with respect to Θ . For Θ being the equality predicate, we obtain the so-called equi-join.

Given a query of the pattern $R \bowtie_C S$, a Merkle Hash Tree structure that supports the efficient computation of a \mathcal{VO} for the query result is based on the materialization (i.e., the physical storage) of the Cartesian Product $R \times S$. Note that such a structure has to be constructed by the *owner* before a *publisher* accepts any query from *clients*. The reason for choosing a materialization of the Cartesian Product is that this structure supports all three types of joins mentioned above. This is due to the fact that these joins can be formulated in terms of basic relational algebra operators, i.e., $R \bowtie_{A_R \Theta A_S} S := \sigma_{A_R \Theta A_S}(R \times S)$. The important issue in constructing the Merkle Hash Tree for queries of the pattern $R \bowtie_C S$ is that first the Cartesian Product is determined and then the resulting tuples are sorted on the *difference* between the values for A_R and A_S , assuming such an operation can be defined. We thus obtain three “groups” of leaf nodes in the Merkle Tree: (1) nodes for which the difference $t.A_R - s.A_S$ for two tuples $t \in R, s \in S$ is 0, thus supporting equi-joins,

(2) nodes where the difference is positive, thus supporting the predicate $\Theta \equiv >$, and (3) nodes where the difference is negative, thus supporting the predicate $\Theta \equiv <$.

For each of the three cases above, assume a query result set q . Our burden again is provide inclusiveness and completeness evidence for the query result.

1. For each tuple $t \in q$, we provide a hash path in the (sorted) Merkle tree for $R \times S$ showing that tuple is in the relation.
2. To show that no tuples are missing, we show two pairs of boundary hash paths (within and without) verifying the boundaries of the answer set.

If it is known that the queries will only result in equi-joins against the database, an optimized Merkle Tree structure can be used. We will only sketch the basic concept for using this structure here. Instead of a space-consuming materialization of the Cartesian Product $R \times S$, we materialize the *Full Outer Join* $R \bowtie_{\sqcup} S$ which pads tuples for which no matching tuples in the other relation exist with null values (see, e.g., [8, 25]). The result tuples obtained by the full outer-join operator again can be grouped into three classes: (1) those tuples $ts, t \in R, s \in S$, for which the join condition holds, (2) tuples from r for which no matching tuples in s exist, and (3) tuples from s for which no matching tuples in r exist. Constructing a \mathcal{VO} for the result of a query of the pattern $R \bowtie_{A_R \Theta A_S} S$ then can be done in the same fashion as outlined above.

4.4 Set Operations

All set operations involve two relations u and v . We may assume that u and v are intermediate results of a query evaluation, and are subsets of some relations r and s respectively, and that r and s are each sorted (possibly on different attributes) and have their own Merkle trees $MHT(r, \mathcal{A})$ and $MHT(s, \mathcal{A}')$, the root of which is signed as usual. We consider the set operations union and set intersection.

Union. In this case, the answer set is $q = u \cup v$. Evidence that $u \cup v$ is inclusive and complete is straightforward; it is sufficient to provide verification paths for each tuple in $u \cup v$ showing that it belongs to one of the sets. Additionally, a single pass over the union can show that no tuple of u or v is omitted. This can be done with a \mathcal{VO} of size $O(|q| \log_2(\max\{|r|, |s|\}))$. If u and v are presented as contiguous subsets of r and s , inclusiveness can be done even faster; all that is necessary to check the union is the order in which the tuples of $u \cup v$ occur in r and/or s , along with the authenticating information for the sets u and v within r and s , respectively. Given this order, inclusiveness and completeness can be evidenced by a \mathcal{VO} of linear size.

Intersection. The approach for union, however, does not produce compact \mathcal{VO} s for set intersection. Suppose $q = u \cap v$ where u and v are as before. Inclusiveness is easy, with the \mathcal{VO} providing $O(|q|)$ verification paths, as before, showing tuples of q belong to both u and v ; but completeness is harder. The *client* needs assurance that all tuples of u or v that belong in q are present. One can pick the smaller set (say u) and for each tuple t in $u - q$, construct a \mathcal{VO} to show that $t \notin v$. In general, if u and v are intermediate results not occurring contiguously in the same Merkle tree, such a \mathcal{VO} is linear in the size of the smaller set. Consider, for example, a set of tuples $\langle \text{name}, \text{age}, \text{salary} \rangle$, where one wishes to select tuples in a specific salary and age range. Assume then that u has been obtained by performing a selection based on **salary**, and v based on **age**. u

and v would be verified by \mathcal{VO} s resulting from different Merkle hash trees: one sorted by **salary**, and one sorted by **age**. Computing the intersection $u \cap v$ would result in a \mathcal{VO} with size linear in $|u|$. This \mathcal{VO} would provide inclusiveness evidence (in u and v) for each tuple of $u \cap v$, and shows completeness by showing that each remaining tuple in $(u - (u \cap v))$ is not in v . This again leaves us with the unsatisfactory situation of a \mathcal{VO} being linear in the size of a potentially much larger intermediate result (if $|u| \gg |u \cap v|$). A similar problem occurs with set differences $u - v$.

We have not solved the general problem of constructing \mathcal{VO} s linear in the size of the result for intersections and set differences. Indeed, the question remains as to whether (in general) linear-size \mathcal{VO} s can even be constructed for these objects. However, in the following section, we provide an approach to constructing linear-size \mathcal{VO} s for a specific type of intersection, so-called *range query*. This is accomplished using a data structure drawn from computational geometry called a *multi-dimensional range tree*. This approach also works for set differences over range queries on different attributes.

5 Multi-dimensional Verification Objects

In d -dimensional computational geometry, when one is dealing with sets of points in d -space, one could ask a d -space range query. Consider a spatial interval $(\langle x_1^1, x_2^1 \rangle \dots \langle x_1^d, x_2^d \rangle)$: this represents an axis-aligned rectilinear solid in d -space. A query could ask for the points that occur within this solid. Such problems are solved efficiently in computational geometry using so-called *Range Trees* (See [4], Chapter 5). We draw an analogy between this problem and a database query of the form

$$\sigma_{c_1^1 < A_1 < c_1^2}(r) \cap \dots \cap \sigma_{c_d^1 < A_d < c_d^2}(r)$$

where $\{A_1, \dots, A_d\} \subseteq \text{schema}(R)$ for a relation R . We use the multi-dimensional range tree (*mdrt*) data structure to provide compact verification objects.

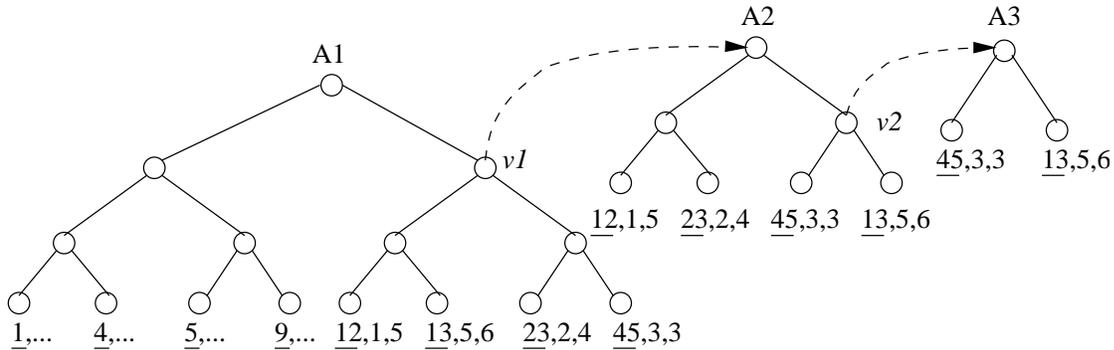


Figure 3: Excerpt of a 3-dimensional range tree, sorted by attributes A_1, A_2 and A_3

Consider the example *mdrt* shown in Figure 3. Let us assume a relation schema R , with 3 attributes A_1, A_2, A_3 on which we want to perform combined selections and provide a \mathcal{VO} for the final answer. Consider the first, 3-dimensional *mdrt* (labeled A_1). This is simply a tree which sorts the tuples according to attribute A_1 (primary key values are underlined). Each inner node in tree A_1 is the ancestor of a set of tuples. Consider such an inner node v_1 , which is the ancestor of tuples with primary key values 12, 13, 23, and 45. An *mdrt* now contains a link from v_1 to the root of an

associated tree A2, denoted as $T_{assoc}(v1, A_2)$. This 2-dimensional *mdrt* A2 contains the same set of tuples with primary keys 45, 11, 23, and 13; however, in this tree, they are sorted according to attribute A_2 . Likewise each inner node v_i in A1 is the ancestor to a set of tuples, and contains a pointer to an associated 2-dimensional *mdrt* $T_{assoc}(v_i, A_2)$ which sorts the tuples in the subtree below v_i by attribute A_2 . In general, each node v_i^j of a $\{d - j + 1\}$ -dimensional *mdrt* contains a pointer to a $\{d - j\}$ -dimensional *mdrt*. The nodes of the final 1-dimensional tree, corresponding to attribute A_d , do not have such pointers.

Given a 2-dimensional range query $\sigma_{x_1^1 \leq A_1 \leq x_2^1}(r) \cap \sigma_{x_1^2 \leq A_2 \leq x_2^2}(r)$ the structure is used as follows. First, the tuple set $q \subseteq r$ with values for attribute A_1 in the range $\langle x_1^1, x_2^1 \rangle$ is identified using tree A1. For simplicity, let us assume (we relax this assumption later) that the tuples in q form the leaves of a balanced tree with root $LCA(q)$. With $|r|=n$, this range can be found in roughly time $O(\log_2 n)$, the time it takes to find the two end-points of the interval in the first tree. We now follow the link to the associated tree for attribute A_2 ; this tree sorts just the tuple set q according to attribute A_2 . So we can find the subset of q satisfying $x_1^2 \leq A_2 \leq x_2^2$ also in $O(\log_2 n)$ time. This gives the intuition behind the efficient processing of conjunctive range queries using *mdrts*. We now relax the assumption that the result of the first selection q includes just the leaves of a balanced tree.

Let us call the leaves of the subtree rooted at node v the *canonical subset* of v , denoted as $P(v)$. If v is a leaf, $P(v) = v$. In [4] (pp. 103-107) it is shown that *any* subset of leaves which lies in a range can be expressed as a disjoint union of $O(\log_2 n)$ canonical subsets for the given range query. The roots of these can be found in $O(\log_2 n)$ time in the process of finding the bounding paths for the interval. Given a range $\langle x, x' \rangle$, we search for them in the tree until we find node v_{split} where the paths split. Now we search x (x') in the left (right) subtree. At every point in the search for x where the path goes left, the right subtree belongs to the range; the search for x' goes just the opposite way. The result is a quick identification of roots of the canonical subtrees that precisely cover the leaves whose values are in the interval (see Figure 4).

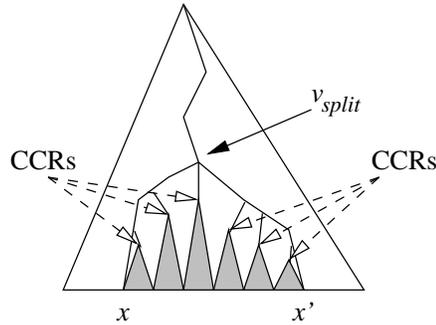


Figure 4: Finding the canonical covering roots (CCRs)

We call these the *covering canonical roots (CCRs)*. Consider, for example, a 2-dimensional range query over attributes A_1 and A_2 . First, the CCRs in tree A1 for the given range in attribute A_1 are found; there are $O(\log n)$ of these. Then for each of the CCRs, we go to the associated trees for attribute A_2 , and find the CCRs in *that* tree for the given range over A_2 . This results in $O(\log^2 n)$ CCRs in tree A2. The union of all the leaves under these CCRs in A2 constitutes the answer. In general, it is shown [4] that d -dimensional range queries can be computed in time $O(\log^{d-1} n)$. Range trees require $O(n \log^{d-1} n)$ storage space, and can be constructed in time $O(n \log^{d-1} n)$.

We now show how to produce a verification object \mathcal{VO} of size $O(\log^d n + |q|)$ for an intersection query with answer q of the form

$$\sigma_{c_1^1 < A_1 < c_1^2}(r) \cap \dots \cap \sigma_{c_d^1 < A_d < c_d^2}(r)$$

using *mdrts*. First, we construct a Merkle hash tree over a d -dimensional *mdrt*. Assume that the associated tree for a node i (for the next attribute) in a range tree is given by $A(i)$, and the hash value of a node i is given by $H(i)$.

Base Case For the base (1-dimensional) *mdrts*, we build the Merkle hash tree as before.

Induction Given the root hashes for the $\{l - 1\}$ -dimensional *mdrts*, we begin with the leaves (height $j = 0$) of the l -dimensional *mdrt*, and compute the tuple hashes, in the usual way. For a node i at height $j > 0$, we compute the hash value thus. First, we append the hashes of all the children of i (say i_1, i_2) together, and also the hash of the root of the associated $l - 1$ -dimensional range tree, and then hash the result.

$$H(i) = h(H(i_1) \parallel H(i_2) \parallel H(A(i)))$$

This construction can be completed in time $O(n \log^{d-1} n)$, and can be overlapped with the construction of the range tree itself.

Now consider the construction of a \mathcal{VO} for a d -way intersection range selection query. The \mathcal{VO} essentially follows the search algorithm. In any given tree, a range query corresponding to an interval results in a set of CCRs. The \mathcal{VO} for this group of $O(\log n)$ CCRs is presented by providing a verification path for the node v_{split} to the root, and verifiable paths (length $O(1)$) from each CCR to the path to v_{split} . In addition, it can readily be seen that by construction of the path from the CCRs to v_{split} , the canonical subtrees of the CCRs form a contiguous non-overlapping sequence over of the answer set. We also provide proximity trees for the LUB and GLB of the interval and the smallest and largest contained intervals; finally, we must show that each of the $k, k = |q|$ selected answer tuples, belongs under a CCR. This process needs to repeat, showing verification paths for all the $O(\log^{d-1} n)$ CCRs found in the process of evaluating the d -way intersecting range selection query. This contributes size $O(\log^d n)$ to the \mathcal{VO} . In addition, the $O(\log^{d-1} n)$ searches in the final tree will each use paths of size $O(\log n)$ and in addition will use the k tuple values returned to produce the root digests for each CCR. This gives us a total \mathcal{VO} size of $O(k + \log^d n)$, which would show that the k tuples belong to the answer q , and that the $O(\log^{d-1} n)$ CCRs together cover the intervals prescribed by the query. In situations where the results of each selection may be large and the final intersection is small, this approach gives us attractively small \mathcal{VO} s.

6 Pragmatic Issues and Related Work

We now examine some pragmatic considerations in using our approach and discuss related work.

6.1 Canonical Join-Select-Project Queries

A typical SQL “**select ... from ... where ...**” can be thought of as one or more joins, followed by a (combined) selection, followed by a projection. We describe how an *mdrt* can be used for both efficient evaluation and construction of compact \mathcal{VO} s for such queries. Specifically, consider a query that involves the join of two relations R and S , followed by a series of selections and a

final projection. Let us assume a Theta-join over attribute A_1 (A_1 occurring in both relations), followed by selections on attributes A_2 and A_3 , and a final projection on several attributes, jointly represented by A_4 (as discussed in Section 4.2).

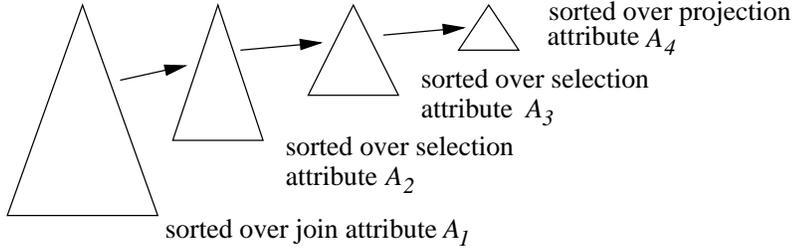


Figure 5: Excerpt of a 3-dimensional range tree, sorted by attributes A_1, A_2 and A_3

We begin this construction (see Figure 5) with the first range tree constructed over the join attribute A_1 , as explained in Section 4.3. Then follow the trees sorted by A_2 and A_3 , and finally on the projected attributes. Given a query, the evaluation plan (and the construction of the \mathcal{VO}) follows this set of range trees as described above, leading to both efficient evaluation of all the steps, and a \mathcal{VO} linear in the size of the final query result.

6.2 Query Flexibility

For efficient verification of query answering, we make use of different trees over sorted tuple sets. Without such trees, our approach cannot provide small verification objects. This points to a limitation of our approach—only queries for which Merkle trees have been pre-computed can be evaluated with compact verification objects. Our approach cannot support arbitrary interactive querying with compact verification objects. Arbitrary interactive querying, however, is quite rare in the presence of fixed applications at *client* sites.

In practice, however, data-intensive applications make use of a fixed set of queries. Indeed, via mechanisms such as embedded SQL (see, e.g., [25]) database queries are compiled into applications. These queries can still make use of parameters entered by a user and which are typically used in selection conditions. Our approach is compatible with such applications. Essentially, client applications commit *a priori* to queries they wish to execute; the owner and the publisher then pre-compute the required Merkle hash trees to produce short verification objects.

So while our approach cannot provide compact verification objects in the context of arbitrary interactive database querying, it is quite compatible with the widely-used practice of embedding pre-determined (and parameterizable) queries within data-intensive applications.

6.3 Conventions

It is important to note that all interested parties: the *owner*, *publisher* and *clients*, share a consistent schema for the databases being published. In addition there needs to be a secure binding between the schema, the queries and the query evaluation process over the constructed Merkle trees. A convention to include this information within the hash trees needs to be established. All parties also need to agree on the data structures used for the \mathcal{VO} . It is also important that the *publisher*

and the *clients* agree upon the format in which the \mathcal{VO} together with the query result is encoded and transmitted. Tagged data streams such as XML provide an attractive option.

6.4 Recent Query Evaluations

Clients must verify that query evaluation is due to an “adequately recent” snapshot of the database and not an old version. We assume the technique of recent-secure authentication for solving this problem. Risk takers (e.g., organizations relying on the accuracy of the data) specify freshness policies on how fresh the database must be. The digital signatures over the database include a timestamp of the last update as well as other versioning information. Based on assumptions concerning trusted synchronization paths and synchronization bounds, clients interpret the timestamps and verify the database is adequately recent with respect to their freshness policies.

6.5 Related Work

The use of Merkle hash trees for authentication of data is not new. This work is most closely related to the work of Naor & Nissim [20] for revocation. Haber and Stronetta [13] use similar techniques for timestamping. Similar schemes have also been used for micro-payments [23]. All these schemes (including ours) share a common theme of leveraging the trust provided by a few digital signatures from a trusted party over multiple hashes, hash paths or hash trees, with the goal of protecting the integrity of the content, with efficient verification, since hashes are more efficient than digital signatures. However, the use of such trees for authentic data publication is new.

There is quite bit of related work in the general area of database security, particularly on access control, statistical querying etc. [7, 16]. Anderson [3] discusses an approach to third-party publication of data in files, but without querying over the contents. Again, to our knowledge, this particular problem of authentic database publication has remained unexamined. Finally, our approach can be viewed as providing “proof-carrying” [21] answers to database queries.

7 Conclusions and Future Work

We have explored the problem of authentic third party data publication. In particular, we have developed several techniques that allow untrusted third parties to provide evidence of *inclusive* and *complete* query evaluation to clients without using public-key signatures. In addition, the evidence provided is linear in the size of the query answers, and can be checked in linear time. Our techniques may involve the construction of complex data structures, but the cost of this construction is amortized over more efficient query evaluation, as well as the production of compact verification objects. Such pre-computation of views and indexing structures are not uncommon in, e.g., data warehousing applications [22].

However, our techniques are restricted currently to the relational model. Our techniques do not allow interactive querying, but can be used with embedded queries in applications. We cannot currently construct linear-size \mathcal{VO} s for general SQL queries, such as ones including arbitrary intersections; we also leave open the (lower-bound) question as to whether such \mathcal{VO} s are possible. We believe, however, that our techniques are a start on an important problem area. Several recent extensions of this work show that our techniques can be used in broader contexts.

Recent results have shown many new settings in which data structures can be efficiently digested

and answers verified: skip lists [11] and persistent dictionaries [1]. In addition, a broad class of data structures was shown to be suitable for efficient use in authentic publication [17]. These results show that our approach to authentic publication may be useful in an increasingly wide range of settings.

We are currently investigating an authentic data publication framework in the context of the eXtended Markup Language (XML [6]), since XML is becoming a major factor in the representation and exchange of data over the Internet. A complicating factor in providing verification objects to queries over such hierarchically structured data is to come up with concise and compact structures that allow for an efficient computation of \mathcal{VO} s. In [10], we have presented an authentic publication framework for XML data using document type definitions (DTD). A similar approach, which does not assume the presence of a DTD, has been presented in [15]. We are currently implementing prototypes of the proposed frameworks for both relational and XML data. A major focus of our research currently is the extension of the queries supported by our approach.

References

- [1] A. Anagnostopoulos, M. T. Goodrich, and R. Tamassia. Persistent Authenticated Dictionaries and Their Applications. In *Proc. Information Security Conference (ISC 2001)*, Lecture Notes in Computer Science, vol. 2200, 379-393, 2001
- [2] N.M. Amato and M.C. Loui. Checking Linked Data Structures. In *Proceedings of the 24th Annual International Symposium on Fault-Tolerant Computing (FTPS)*, 164-173, 1994.
- [3] R. J. Anderson. The Eternity Service. In *Proceedings of Pragocrypt*, 1996.
- [4] M. D. Berg , M. V. Kreveld, M. Overmars and O. Schwarzkopf. *Computational Geometry*. Springer-Verlag, New York.
- [5] M. Blum, W. Evans, P. Gemmell, S. Kannan, and M. Noar. Checking the Inclusiveness of Memories. *Algorithmica*, 12(2/3):225-244, 1994.
- [6] T. Bray, J. Paoli, C. Sperberg-McQueen. Extensible Markup Language (XML) 1.0. W3C Recommendation, February 1998.
- [7] S. Castano, M. Fugini, G. Martella, P. Samarati. Database Security. Addison-Wesley, 1995
- [8] C.J. Date. An Introduction to Database Systems (7th Ed), Addison-Wesley, 1999.
- [9] C.J. Date and H. Darwen. A Guide to the SQL Standard (4th Ed), Addison-Wesley, 1997.
- [10] P. Devanbu, M. Gertz, A. Kwong, C. Martel, S. Stubblebine. Flexible Authentication of XML Documents. In *Proceedings of the 8th ACM Conference on Computer and Communications Security (CCS-8)*, ACM, 136-145, 2001.
- [11] M. Goodrich, R. Tamassia, A. Schwerin. Implementation of an Authenticated Dictionary with Skip Lists and Commutative Hashing. In *Proceedings of the 2nd DARPA Information Survivability Conference and Exposition (DISCEX II)*, 2001.
- [12] A. Gupta, I.S. Mumick. Maintenance of Materialized Views: Problems, Techniques, and Applications. *Data Engineering Bulletin* 18(2): 3-18, 1995,

- [13] S. Haber and W. S. Stornetta. How to Timestamp a Digital Document. *Journal of Cryptology*, 3(2), 1991.
- [14] S. Jajodia, P. Samarati, V. S. Subramanian, E. Bertino. A Unified Framework for Enforcing Multiple Access Control Policies. In *Proceedings ACM SIGMOD International Conference on Management of Data*, 474-485, 1997.
- [15] A. Kwong and M. Gertz: Authentic Publication of XML Document Data. In *2nd International Conference on Web Information Systems Engineering (WISE 01)*, 331-340, IEEE Computer Society, 2001.
- [16] T. Lunt, (Ed.) *Research Directions in Database Security*. Springer-Verlag, 1992
- [17] C. Martel, G Nuckolls, P. Devanbu, M. Gertz, A. Kwong, S. Stubblebine. A General Model for Authenticated Data Structures. Technical Report CSE-2001-9, December 2001, Department of Computer Science, University of California, Davis, 2001.
- [18] R.C. Merkle. A Certified Digital Signature. In *Advances in Cryptology – CRYPTO '89, 9th Annual International Cryptology Conference*, Lecture Notes in Computer Science, vol. 435, 218-238, 1989.
- [19] J. Melton, A.R. Simon. *SQL: 1999 - Understanding Relational Language Components*. Morgan Kaufmann, 2001.
- [20] M. Naor and K. Nissim. Certificate Revocation and Certificate Update. In *Proceedings of the 7th USENIX Security Symposium*, 1998.
- [21] G. Neula. Proof-carrying code. In *24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ACM Press, 106-119, 1997.
- [22] W.H. Inmon. *Building the Data Warehouse* John Wiley & Sons, 1996.
- [23] S. Charanjit and M. Yung. Paytree: Amortized Signature for Flexible Micropayments. In *Proceedings of the Second Usenix Workshop on Electronic Commerce*, 1996.
- [24] N. Roussopoulos. Materialized Views and Data Warehouses. *SIGMOD Record* 27(1): 21-26, 1998.
- [25] A. Silberschatz, H. Korth, S. Sudarshan. *Database System Concepts (4th Edition)*, McGraw-Hill, 2002.
- [26] J. D. Tygar. Open Problems In Electronic Commerce. In *Proceedings of the Eighteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, 1999.
- [27] B. Yee and D. Tygar. Secure Coprocessors in Electronic Commerce Applications In *Proceedings of The First USENIX Workshop on Electronic Commerce*, New York, 1995.

Appendix: Proof of correctness for selection verification

For now, we include here only the proof that simple selections are hard to forge. This proof is an induction on the height of the Merkle tree corresponding to that selection. The full proof shows that checks for all other operators are also hard to forge; then we need another induction on the nesting depth of the operators in the query. This proof is also only for binary trees, but is easily extended to multiway trees such as B-trees.

We assume that the client knows H_r , the correct root hash value for the Merkle tree associated with our relation/attribute. We also assume that the client has a program (including the hash function h) which will correctly check the \mathcal{VO} and the answer. We now show that as long as an adversary cannot forge values which cause collisions in h (more precision on this below), the client will never verify an incorrect answer.

Before giving the proof we need some preliminaries that describe the process and syntax in more detail than we had in Section 4.1. Let the correct answer S to a selection query on relation r be the tuples $t_j \leq t_{j+1} \dots \leq t_k$. If t_j isn't the smallest tuple in r we also return t_{j-1} , and if t_k isn't the largest tuple in r we return t_{k+1} , which are the largest and smallest tuples not in S . We call this larger set S_b (S plus the boundary tuples). We assume that the tree also may have two special tuples which are smaller than anything in r and larger than anything in r (in case the smallest tuple is a right child or the largest a left child). Note that this assures that the root always has two children. In addition to S_b , the client is provided with the hash values of some internal nodes in the Merkle tree. Specifically, the answer (which is both data and the \mathcal{VO}) can be computed by the *publisher* as follows by initially calling the procedure **Answer** with $u = R$ the root of the tree. Note that we assume a syntax for the \mathcal{VO} of the form: $[A_l], [A_r]$ where A_l is the \mathcal{VO} for the left subtree and A_r is the \mathcal{VO} for the right subtree.

Answer(u)

begin

If all leafs in the subtree rooted at u are in S_b , then return the single tuple if u is a leaf, otherwise return the values of all the leaves in the form: [leaves in left subtree], [leaves in right subtree];

If no leaf in the subtree rooted at u is in S_b , then return the hash value of u in the Merkle tree;

If some but not all leaves in the subtree rooted at u are in S_b , then return [*Answer*(left-child(u))], [*Answer*(right-child(u))];

end

Note that this construction algorithm for the answer implies a simple checking mechanism: compute the hash value v_l for the left child u_l of the root using the string in the left group of [], then compute the right child value v_r , finally compute H_r using these two values. To check validity the *client* will also compute a status variable for the current subtree rooted at u . A node's status can be

empty: no tuples under u were returned;

right-terminated (RT): a tuple too large to be in S is under u ;

right-full (RF): valid, but not right-terminated;

left-full (LF) and left-terminated (LT) are defined similarly;

<i>left</i>	<i>right</i>	<i>parent</i>
empty	empty	invalid
empty	LF,R	invalid
empty	LT,R	LT, R
non-empty	LT,R	invalid
L,RF	LF,R	L,R
L,RT	empty	L,RT
L,RF	empty	invalid
L,RT	non-empty	invalid.

invalid: the tuples under v are known to be an invalid answer.

The status makes sure that an answer is put together properly, e.g., if the left subtree of a node u is *right-terminated* then no tuples from the right subtree of u can be in the answer, so the right subtree status must be empty.

We compute the status as follows: a leaf is both left-full and right-full if it is in S ; left-terminated and right-full if too small; and left-full and right-terminated if too large.

In general, for a node u , if either child is invalid, u will also be invalid. If both children are valid the following table shows how to get the parent's status. L stands for LT or LF and R for RT or RF. If an L or R appears as a child and parent it has the same status both times. Note that each valid, non-empty node has two status values (for its left and right subtrees).

Attack Model We assume that the *owner* has constructed a Merkle hash tree of a relation r over the selection attribute A , and the client knows the correct root hash value H_r .

We assume that *publisher* is the adversary, has access to the relation r , and can construct the same Merkle hash tree; we also assume that the hash function chosen by *owner* cannot be feasibly forged. Specifically, we make the following assumption about the adversary. For any selection query there is a correct answer which the client is supposed to receive, and a correct verification computation. In the course of the correct verification computation the *client* first hashes each tuple given, and then creates a sequence of triples x_j, y_j, z_j where $z_j = h(x_j, y_j)$ is the result of the j th application of the hash function. We assume that the adversary is prohibited from

1. providing a tuple t such that t is not in r and $h(t) = h(t_j)$ for t_j in r ,
2. providing an answer such that in the course of the protocol *client* would compute $z_j = h(a_j, b_j)$ and $a_j \neq x_j$ or $b_j \neq y_j$ (intuitively this would mean the adversary has found an alternate pair of values which hash to the same answer), and
3. providing an answer such that in the course of the protocol *client* would compute $h(a_j, b_j) = h(t_j)$ for t_j in r (we don't need this if the *client* knows the height of the tree).

Lemma 7 Under the above cryptographic assumptions, when *client* executes the verification protocol on an *answer* ADV provided by *publisher*, if *client* computes the correct root hash value H_r and the root status is valid, ADV is in fact the correct answer and the status of the root is correct. ■

Basis We show this by induction. Consider a Merkle hash tree of height 1; this must be a root with just two leaves t_1, t_2 . Since the root hash H_r is known, by definition, the final step of an accepting

protocol must be to compute $h(v_l, v_r) = H_r$ where $v_l = h(t_1)$ and $v_r = h(t_2)$ are the correct hash values of the left and right subtrees. By our security assumption the only way the protocol could get v_l, v_r is either by being given them directly or by being given a correct value of t_1 or t_2 and hashing them. The protocol cannot be given both v_l and v_r directly since this would mean S_b is empty which the *client* knows is impossible. Thus the only way ADV might be accepted is if it provides two tuples or one tuple and a value.

Case 1: The answer ADV is $[t_1], [t_2]$ where the two tuples t_1 and t_2 are children of the root (if r has only one tuple the other tuple will be the dummy tuple for maximum/minimum attribute value, if r is empty both tuples will be dummy). Thus t_1 and t_2 are both claimed to be in S_b . The *client* will always be able to look at the tuples to see which are actually part of the correct answer and which are boundary tuples. The user then computes $v_l = h(t_1)$, $v_r = h(t_2)$, and finally $V = h(v_l, v_r)$. By our assumption this is the only hash computation that can occur while processing ADV which can have R_h as its answer. Furthermore, no hash computation other than those listed above can have v_l or v_r as its results. Thus $V = H_r$ if and only if t_1 and t_2 are the correct tuple values and no other values are provided. The protocol also has a completeness proof. For example, if t_1 is in S and t_2 is too large to be in r , since the *client* knows that t_2 is t_1 's right sibling, t_1 must be the largest tuple in S . In this case the root status would be left-full, right-terminated.

Case 2: The adversary might also get the correct root value computed by providing $\text{ADV} = [t_1], [v_r]$. This is one tuple value (t_1), and the hash value of the other tuple (in essence claiming that $S_b = \{t_1\}$). However, this means the protocol would see the right subtree as empty, so the root is only valid if the left subtree is right-terminated. This would only be the case if t_1 is too large to be in S , and since the algorithm has the correct value of t_1 it will know if this is the case.

The setting where ADV is $[v_l], [t_2]$ is analogous.

Induction Step We now assume that the lemma holds for any tree of height less than i , namely that when we apply our protocol to a non-empty tree of height less than i , we get the correct root value and valid status if and only if we are given a correct answer (tuples and hash values) for that tree.

Now consider a Merkle hash tree of height i rooted at R of height $i \geq 2$, with two immediate subtrees rooted at u_l and u_r with hash values v_l and v_r . ADV must be of the form $[A_l], [A_r]$ where A_l and A_r are answer strings which our Answer protocol can parse. If the *answer* does not have this high level form it will be immediately rejected since a correct answer cannot be empty (it always has the boundary tuples).

As the final step of an accepting protocol we must compute $H_r = h(v_l, v_r)$. By our assumption, the only way the protocol when run on ADV can yield H_r as its final value is if the protocol evaluates A_l to v_l and A_r to v_r and ends with a valid status for both nodes.

The trees rooted at u_l and u_r are of height less than i . Furthermore, the protocol treats non-empty subtrees exactly as it does the entire tree. Thus by the induction hypothesis, if A_l is non-empty, and produces the correct hash value and a valid status, then it must include the correct tuples and end with the correct status for u_l (and similarly for u_r).

If A_l, A_r are both non-empty, evaluate to the correct hash values, and have a valid status, they are the correct trees and thus must fit together in a valid way (e.g., we will never have them both be right-terminated). However, if one of the trees is empty we cannot apply our induction hypothesis to it (since this only works for non-empty trees). So suppose that we have simply been given $A_l = v_l$ as part of ADV (this is easy for the adversary since v_l is known). In order for the root to get a valid status, the status of u_r must be left-terminated (we don't care if its right-terminated or

full, either is OK). However, if u_r is left-terminated and correct, we know that u_l should be empty, and we are again correct to conclude that the root status is valid. The case where u_r is empty is analogous.

Thus in all cases we can extend the correctness of our protocol to a tree of height i and the lemma follows.