

Authenticated Data Structures, Generically

Andrew Miller, Michael Hicks, Jonathan Katz, and Elaine Shi

University of Maryland, College Park, USA

Abstract

An authenticated data structure (ADS) is a data structure whose operations can be carried out by an untrusted *prover*, the results of which a *verifier* can efficiently check as authentic. This is done by having the prover produce a compact proof that the verifier can check along with each operation’s result. ADSs thus support outsourcing data maintenance and processing tasks to untrusted servers without loss of integrity. Past work on ADSs has focused on particular data structures (or limited classes of data structures), one at a time, often with support only for particular operations.

This paper presents a generic method, using a simple extension to a ML-like functional programming language we call $\lambda\bullet$ (lambda-auth), with which one can program authenticated operations over any data structure defined by standard type constructors, including recursive types, sums, and products. The programmer writes the data structure largely as usual and it is compiled to code to be run by the prover and verifier. Using a formalization of $\lambda\bullet$ we prove that all well-typed $\lambda\bullet$ programs result in code that is secure under the standard cryptographic assumption of collision-resistant hash functions. We have implemented $\lambda\bullet$ as an extension to the OCaml compiler, and have used it to produce authenticated versions of many interesting data structures including binary search trees, red-black+ trees, skip lists, and more. Performance experiments show that our approach is efficient, giving up little compared to the hand-optimized data structures developed previously.

Categories and Subject Descriptors D.3.3 [Programming Languages]: Language Constructs and Features—Data types and structures

General Terms Security, Programming Languages, Cryptography

1. Introduction

Suppose data provider would like to allow third parties to mirror its data, providing a query interface over it to clients. The data provider wants to assure clients that the mirrors will answer queries over the data truthfully, even if they (or another party that compromises a mirror) have an incentive to lie. As examples, the data provider might be providing stock market data, a certificate revocation list, the Tor relay list, or the state of the current Bitcoin ledger [22].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

POPL ’14, January 22–24, 2014, San Diego, CA, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2544-8/14/01...\$15.00.

<http://dx.doi.org/10.1145/2535838.2535851>

Such a scenario can be supported using *authenticated data structures* (ADS) [5, 24, 31]. ADS computations involve two roles, the *prover* and the *verifier*. The mirror plays the role of the prover, storing the data of interest and answering queries about it. The client plays the role of the verifier, posing queries to the prover and verifying that the returned results are authentic. At any point in time, the verifier holds only a short *digest* that can be viewed as summarizing the current contents of the data; an authentic copy of the digest is provided by the data owner. When the verifier sends the prover a query, the prover computes the result and returns it along with a *proof* that the returned result is correct; both the proof and the time to produce it are linear in the time to compute the query result. The verifier can attempt to verify the proof (in time linear in the size of the proof) using its current digest, and will accept the returned result only if the proof verifies. If the verifier is also the data provider, the verifier may also update its data stored at the prover; in this case, the result is an updated digest and the proof shows that this updated digest was computed correctly. ADS computations have two properties. *Correctness* implies that when both parties execute the protocol correctly, the proofs given by the prover verify correctly and the verifier always receives the correct result. *Security*¹ implies that a computationally bounded, malicious prover cannot fool the verifier into accepting an incorrect result.

Authenticated data structures can be traced back to Merkle [18]; the well-known *Merkle hash tree* can be viewed as providing an authenticated version of a bounded-length array. More recently, authenticated versions of data structures as diverse as sets [23, 27], dictionaries [1, 12], range trees [16], graphs [13], skip lists [11, 12], B-trees [21], hash trees [26], and more [15] have been proposed. In each of these cases, the design of the data structure, the supporting operations, and how they can be proved authentic have been reconsidered from scratch, involving a new, potentially tricky proof of security. Arguably, this state of affairs has hindered the advancement of new data-structure designs as previous ideas are not easily reused or reapplied. We believe that ADSs will make their way into systems more often if they become easier to build.

This paper presents $\lambda\bullet$ (pronounced “lambda auth”), a language for programming authenticated data structures. $\lambda\bullet$ represents the first *generic*, language-based approach to building dynamic authenticated data structures with provable guarantees. The key observation underlying $\lambda\bullet$ ’s design is that, whatever the data structure or operation, the computations performed by the prover and verifier can be made structurally the same: the prover constructs the proof at key points when executing a query, and the verifier checks a proof by using it to “replay” the query, checking at each key point that the computation is self-consistent.

$\lambda\bullet$ implements this idea using what we call *authenticated types*, written $\bullet\tau$, with coercions *auth* and *unauth* for introducing and eliminating values of an authenticated type. Using standard func-

¹ This property is sometimes called *soundness* but we eschew this term to avoid confusion with its standard usage in programming languages.

tional programming features, the programmer writes her ADS’s datatype definition and its corresponding operations (e.g., queries and updates) to use authenticated types. For example, as we show later in the paper, the programmer could write an efficient authenticated binary search tree using the (OCaml-style) type definition $\text{bst} = \text{Tip} \mid \text{Bin of } \bullet\text{bst} \times \text{int} \times \bullet\text{bst}$ along with essentially standard routines for querying and insertion. Then, given such a program, the $\lambda\bullet$ compiler produces code for both a prover and a verifier that will produce or confirm, respectively, a proof of the correct execution of the corresponding operation. Proofs consist of a stream of what we call *shallow projections* of the data the prover visits while running its routine: the prover’s code adds to this stream at each *unauth* call, while the verifier’s code draws from the stream at the corresponding call, checking for consistency. We give a more detailed overview of how this approach works, and how authenticated types are represented, in Section 2. Importantly, as we show in Sections 3 and 4, *any* well-typed program written in $\lambda\bullet$ compiles to a prover and verifier which are correct and secure, where security holds under the standard cryptographic assumption of *collision-resistant hash functions*.²

$\lambda\bullet$ provides two key benefits over prior work. First, it is extremely *flexible*. We can use $\lambda\bullet$ to implement any dynamic data structure, both queries and updates, expressible using ML-style data types (involving sums, products, and recursive types). Our theoretical development, though not our implementation, also supports authenticated functions. Previous work by Martel et al. [17] can also be used to build DAG-oriented ADSs, but it supports only queries and not (incremental) updates, requires the data structure have a single root, and does not support authenticated functions. $\lambda\bullet$ ’s flexibility does not compromise its performance. To the best of our knowledge the asymptotic performance of every prior ADS construction from the literature based on collision-resistant hashing can be matched by $\lambda\bullet$. We have implemented an optimizing $\lambda\bullet$ compiler as an extension to the OCaml compiler (described in Section 5), and using it we have implemented Merkle trees, authenticated binary search trees, red-black+ trees, skip lists, and planar separator trees, as well as improvements to standard Bitcoin data structures. Experiments described in Section 6 confirm the expected asymptotic performance of $\lambda\bullet$ ADSs, show the benefit of the two compiler optimizations we implemented (which exploit space/time tradeoffs), and demonstrate that the performance of $\lambda\bullet$ ADSs is competitive with hand-rolled versions.

$\lambda\bullet$ ’s second main benefit is *ease of use*. We find that it is relatively simple to construct an ADS using $\lambda\bullet$: just write the standard data structure in a purely functional style, and sprinkle in authenticated types; we give a flavor for this in the next section. Pleasantly, there is no need for the ADS designer to prove anything: Assuming the resulting program type checks, the programmer is assured that the produced prover and verifier code enjoy both correctness and security. By contrast, Martel et al. [17] provide no such support; programmers must write and check their implementations manually. $\lambda\bullet$ ADSs can be freely composed and customized just as one might expect with normal data structures, a fact which we hope will make them more readily deployable. All of this is in contrast to the state of practice with ADSs today, summarized in Section 7, which tends to favor hand-rolled versions that are hard to build, customize, and compose.

In summary, this paper makes the following contributions:

1. We present $\lambda\bullet$, a purely functional language in which one can write a rich array of authenticated data structures using a novel feature we call *authenticated types*.

² Informally, hash is collision-resistant if it is computationally infeasible to find distinct inputs x, x' such that $\text{hash}(x) = \text{hash}(x')$. We treat this more formally in Section 4.4.




Figure 1. A Merkle tree and proof π for `fetch(2)` describing the highlighted path (with bold edge and node outlines).

2. We formalize the semantics and type rules for $\lambda\bullet$ and prove that all well-typed $\lambda\bullet$ programs produce ADS protocols that are both correct and secure.
3. We have implemented $\lambda\bullet$ as an extension to the OCaml compiler³ and used it to program a variety of existing and new ADSs, showing good asymptotic performance that is competitive with hand-rolled implementations.

2. Overview

This section presents an overview of our approach. We begin by describing Merkle trees, the canonical example of an authenticated data structure. Then we give a flavor for $\lambda\bullet$ by showing how we can use it to implement Merkle trees. We conclude with a discussion of the flexibility and ease-of-use benefits of using $\lambda\bullet$ to write efficient authenticated data structures.

2.1 Background: Merkle trees

The canonical example of an ADS is a *Merkle tree* [19], which is the authenticated version of a full binary tree where data is stored at the leaves but not the interior nodes. A Merkle tree of height h can represent an array of $n = 2^{h-1}$ elements, x_0, \dots, x_{n-1} . Each leaf node is coupled with a *digest* that consists of the hash of the associated element, while each internal node contains a digest that is the hash of the concatenation of the digests of its two children. A depiction of a Merkle tree for $h = 3$ is given in Figure 1. Each leaf is associated with a string $\text{str}0, \text{str}1$, etc. Each node is numbered according to its position in the tree, with x, y indicating x as the row and y as the column.

The canonical Merkle-tree query fetches the value x_i at index $i \in [0, n - 1]$. When thus queried, the prover (call it P) returns the value x_i along with the set of digests π needed to compute the root digest. The verifier (call it V) keeps a copy of the root digest itself, and checks the proof by recomputing this digest from the proof to make sure the two match. Figure 1 shows the proof π for a fetch at position $i = 2$ (i.e., the leaf at position $N0, 2$). It consists of three elements in sequence, the string $\text{str}2$, the hash h_7 , and the hash h_2 ; these are labeled with the salient nodes of the tree that they relate to. Verification proceeds bottom up: V computes the hash of $\text{str}2$, which is h_6 , and concatenates that with the hash h_7 provided in π . It then concatenates these two and takes the hash to compute what should be the digest for node $N1, 1$, i.e., h_3 . Then it concatenates h_2 , the hash for $N1, 0$ provided in π , with its computed digest for $N1, 1$ and hashes the result. It then confirms that this computed digest equals h_1 , the digest it stores for the whole tree.

Performance analysis. Because the tree is perfectly balanced, the size of the proof is always $\log_2 n$; additionally the computational cost for each of P and V is $\log_2 n$. The overall size of the data structure stored by P is $O(n)$, whereas V at no point requires more than a constant amount of storage or memory. In particular, V only stores a constant-sized digest of the tree between fetch operations,

³ The full open-source code for our implementation is available at the following URL: <http://www.cs.umd.edu/~amiller/gpads/>

```

type tree = Tip of string | Bin of  $\bullet$ tree  $\times$   $\bullet$ tree
type bit = L | R
let rec fetch (idx:bit list) (t: $\bullet$ tree) : string =
  match idx, unauth t with
  | [], Tip a  $\rightarrow$  a
  | L :: idx, Bin(l,_)  $\rightarrow$  fetch idx l
  | R :: idx, Bin(_,r)  $\rightarrow$  fetch idx r

```

Figure 2. Merkle trees in $\lambda\bullet$. The tree is assumed to be complete, i.e., with a power-of-two number of leaves.

and because V processes each (constant-sized) hash in order, it can discard it immediately after it is read. For this reason, we often refer to π as a proof *stream*.

Security analysis. As described in the Introduction, we are interested in two properties of this scheme. *Correctness* says that when P executes a query f over a tree t correctly, then V gets the same result as it would have if it had just computed $f(t)$ normally. The second property, *security*, says that a computationally bounded, cheating prover P^* cannot cause V to accept an incorrect answer. The basis of this property is the use of collision-resistant hashes: we can show that if P^* can cause V to accept an incorrect answer then the proof returned by P^* will yield a collision. We state these properties precisely, in the context of $\lambda\bullet$, in Section 4.

2.2 Introducing $\lambda\bullet$, a language for programming ADSs

The Merkle tree verification procedure was carefully designed with the properties of the underlying data structure in mind. In particular, there can be but one *path* from the root to a given leaf, and from this path we can determine digests sufficient to recompute the root digest. The question is: how might we generalize this approach to arbitrary data structures t involving arbitrary computations f ? We designed $\lambda\bullet$ as a solution to this problem.

$\lambda\bullet$ is a completely standard, purely functional programming language extended with *authenticated types* $\bullet\tau$, along with coercions *auth* and *unauth*, which have type $\forall\alpha.\alpha \rightarrow \bullet\alpha$ (for introducing authenticated values) and type $\forall\alpha.\bullet\alpha \rightarrow \alpha$ (for eliminating authenticated values), respectively. A function f using authenticated types is compiled to variations f_P and f_V for the prover and verifier, respectively. Data of type $\bullet\tau$ stored at the prover is like a normal value of type τ but augmented with digests, while data of type $\bullet\tau$ stored at the verifier is simply a compact digest. The *auth/unauth* coercions at the prover facilitate proof generation, while at the verifier they check a provided proof. In short, $\lambda\bullet$'s design exploits the observation that proof generation and proof verification can be made structurally *identical* essentially by piggy-backing them on top of the ideal computation of $f(t)$.

Example. As an illustration, Figure 2 shows a version of Merkle trees written in our OCaml-based $\lambda\bullet$ implementation. The type tree is simply a binary tree with strings stored at the leaves. The fetch function takes an index expressed as a list of bits, which is interpreted as a path through the tree, with L bits directing the traversal to the left, and R bits directing it right. The function returns the string associated with the Tip that is eventually reached. Notice that since the argument t has type \bullet tree, the function must call *unauth* t to coerce it to a tree to be matched against (we give a use of *auth* in Section 2.3).

Interpretation of authenticated types. All standard constructs have the usual semantics in both f_P and f_V , but authenticated types are interpreted differently.

Prover For f_P , values of type $\bullet\tau$ consist of pairs $\langle h, v \rangle$ where v has type τ and h is its digest, i.e., a hash of the *shallow pro-*




Figure 3. A $\lambda\bullet$ Merkle tree t at the prover (of type \bullet tree from Figure 2) and proof stream π for query `fetch t [R; L]`. Hashes of relevant shallow projections are given in the lower right.

jection of v . The shallow projection of a value is just the value itself for all values of type τ that do not consist of any authenticated types $\bullet\tau_0$, while the shallow projection of an authenticated value $\langle h, v \rangle$ is just the digest h (the formal definition is in Figure 9). Looking at the definition of type tree in the figure, we can see that recursive references to the tree in the Bin case are authenticated. As such, the prover's representation is just as described in the previous subsection: each node of the tree has the form $\text{Bin}(\langle h_1, v_1 \rangle, \langle h_2, v_2 \rangle)$, which consists of the left subtree v_1 and its digest h_1 , and the right subtree v_2 and its digest h_2 . Each digest consists of the hash of the shallow projection of its respective tree. So, if v_1 was a leaf $\text{Tip}(s)$, the shallow projection is just $\text{Tip}(s)$ itself, and thus h_1 is the hash of $\text{Tip}(s)$. On the other hand, suppose v_1 is a node $\text{Bin}(\langle h_{11}, v_{11} \rangle, \langle h_{12}, v_{12} \rangle)$. Then $\text{Bin}(h_{11}, h_{12})$ is this node's shallow projection, and h_1 is its digest.

Verifier For f_V , values of type $\bullet\tau$ consist solely of the digest h of some value of type τ . As such, for our example, while the prover maintains the entire tree data structure, the verifier only keeps the digest of the root. In general, values in f_V are the shallow projections of their corresponding values in f_P ; we define this notion formally in the next section.

Turning to the coercions, for both the prover and verifier the *auth* v coercion computes the hash h of the shallow projection of v ; for the prover, this hash is paired with v while the verifier retains only h itself. The interesting part is the semantics of *unauth*. For the prover, *unauth* is called with $\langle h, v \rangle$ and it simply returns v . In addition, it computes the shallow projection of v and adds it to the proof π , which is just a list of such shallow projections. We often refer to π as a proof *stream* to emphasize the list structure. For the verifier, *unauth* takes a hash h and compares it to the hash of the element s at the head of the proof stream, which should be a shallow projection of type τ . If all is well, this element is the one the Prover put there and so the hashes will match and the coercion returns s . Otherwise there is a problem and verification fails.

Example Merkle-tree query. Figure 3 depicts the prover's version of an object of type \bullet tree corresponding to the Merkle tree from Figure 1. These trees are structurally similar but not identical; in particular, notice that a node's digest is stored with the pointer to that node, rather than at the node itself. Suppose the prover executes the query (`fetch [R; L]` t), which corresponds to the query from Section 2.1. The figure also depicts the proof stream π it produces, along with the hashes of shallow projections of relevant tree elements. The first thing the prover will execute is *unauth* t , which returns the pointer to the first node, and stores its shallow projection $\text{Bin}(h_2, h_3)$ in the proof stream—notice that this is the same as the pointed-to node but the sub-tree pointers have been dropped. Execution continues to the third case of the **match**, which recursively calls (`fetch [L]` r), where r is bound to the authenticated value $\langle h_3, v \rangle$ such that v is the right subtree. The prover then invokes *unauth* on this pair, returning v and adding $\text{Bin}(h_6, h_7)$ to the proof

stream. This time we take the second case of the **match**, recursing on $\langle h_6, \text{Tip}(str2) \rangle$, so the call to *unauth* returns $\text{Tip}(str2)$ and adds its shallow projection ($\text{Tip}(str2)$ itself) to the proof stream. Execution concludes with *str2* as the final result while the final proof stream π consists of three elements, representing the three nodes visited.

The verifier begins with the proof stream π and just the digest of t , which is h_1 . It then runs ($\text{fetch } [R; L] t$) using its version of the code. It first executes *unauth* h_1 , which compares h_1 to the hash of the first element s_0 of the proof stream, which is $\text{Bin}(h_2, h_3)$. The hashes match, as per the equations given in the lower right of the figure, and thus execution continues using s_0 . Execution proceeds to the third case of the **match**, recursively calling *fetch* with $[L]$ and h_3 . This time, calling *unauth* h_3 results in comparing h_3 to the hash of the second element in the proof stream, which is $\text{Bin}(h_6, h_7)$, and once again the hashes match and the proof stream element is returned. The second branch of the **match** fires, so the recursive call passes $[]$ and h_6 . Finally, *unauth* h_6 compares h_6 to the hash of the final element of the proof stream, $\text{Tip}(str2)$, which is returned as the hashes match. Thus execution concludes with the final result as *str2*. As all hash checks succeeded, the verifier has confirmed the prover's execution is correct.

Analysis. $\lambda\bullet$ Merkle trees are asymptotically as efficient as the originals, and as secure. As before, the verifier maintains only the constant-sized digest between queries, and the size of the fetch proof and the time to generate and verify it is $O(\log_2 n)$: the proof stream consists of one (constant-size) shallow projection for each recursive call to *fetch*. The argument for security once again rests on collision-resistant hashes, though $\lambda\bullet$ verification checks the root digest top-down rather than bottom-up. Our proof stream has some redundancy (it contains hashes h_2, h_3, h_6 , and h_7 , whereas in Figure 1 the proof contains only h_2 and h_7) but this is only a constant factor and can be optimized away (cf. Section 5.2).

2.3 Discussion: Benefits of $\lambda\bullet$

The primary benefit of writing ADSs in $\lambda\bullet$ over prior approaches is *flexibility* and *ease of use*. $\lambda\bullet$ can support essentially any computation over a DAG-oriented data structure that is expressed as a functional program. Moreover, as proved in Section 4, writing an ADS in $\lambda\bullet$ ensures it is both correct and secure; there is no need for a designer to make a new argument for each new data structure. As far as we are aware, $\lambda\bullet$ can be used to implement any previously proposed ADS based on collision-resistant hashing. As described in Section 6, so far we have successfully implemented Merkle trees and authenticated versions of binary search trees, red-black+ trees [24], skip lists [29], and variations of the Bitcoin block chain [22], all of which enjoy asymptotically identical, or better, performance than their specially-designed counterparts.

Support for updates. Martel et al. [17] also previously proposed a general-purpose scheme that supports ADSs based on DAGs. In principle, their scheme could also support the above-mentioned data structures, but only for query computations, not updates. By contrast, updates are completely natural in $\lambda\bullet$. For example, the function *update* in Figure 4 updates a Merkle tree. The verifier could submit a request to the prover to run ($\text{update } [R; L] t str4$). The prover will produce a proof stream π for the operation along with a new authenticated tree t' that contains the modification, and which shares much of the structure of the original tree t , as per standard functional programming style. The prover can then update its root to now be t' and then send the verifier the result of the execution, which is the digest portion of t' and the proof stream

```

let rec update (idx:bit list) (t:•tree) (newval:string) : •tree =
  match idx, unauth t with
  | [], Tip _ → auth(Tip newval)
  | L::idx', Bin(l,r) → auth(Bin(update idx' l newval, r))
  | R::idx', Bin(l,r) → auth(Bin(l, update idx' r newval))

let update_cps (idx:bit list) (t:•tree) (newval:string) : •tree =
  let rec _update (k : •(•tree → •tree)) idx t x : •tree =
    match idx, unauth t with
    | [], Tip _ → (unauth k) (auth(Tip x))
    | L :: idx', Bin(l,r) →
      _update (auth(fun t → auth(Bin(t,r)))) idx' l x
    | R :: idx', Bin(l,r) →
      _update (auth(fun t → auth(Bin(l,t)))) idx' r x
  in _update (auth(fun t → t)) idx t newval

type stack = E | SL of •stack × •tree | SR of •stack × •tree
let update_stk (idx:bit list) (t:•tree) (newval:string) : •tree =
  let rec build idx t (s:stack) : •tree × stack =
    match idx, unauth t with
    | [], Tip _ → auth(Tip newval), s
    | L::idx', Bin(l,r) → build idx l (SL(auth s, r))
    | R::idx', Bin(l,r) → build idx r (SL(auth s, l)) in
  let rec apply (child:•tree, s:stack) : •tree =
    match s with
    | E → child
    | SL(s, r) → apply (auth(Bin(child, r)), unauth s)
    | SR(s, l) → apply (auth(Bin(l, child)), unauth s) in
  apply (build idx t E)

```

Figure 4. Functions for updating a Merkle tree in $\lambda\bullet$.

π . The verifier can then use π in the usual way to verify that (the digest of) t' is indeed the right result and then update its local root.⁴

Controlling performance. The fact that $\lambda\bullet$ is a general-purpose programming language means that it affords substantial flexibility to the ADS designer in customizing an ADS design to her needs.

As one possible customization, the designer might refactor operations to better control space usage. Consider the *update* function once again. While the proof stream contributes only a constant space overhead, since the verifier can discard each element after it is read, we observe that executing *update* will require $O(\log n)$ stack space, since the function is not tail recursive. One way to eliminate this overhead is to rewrite *update* in continuation-passing style (CPS) such that the continuation itself is authenticated, as for the function *update_cps* given in the middle of Figure 4. As such, recursive uses of nested continuations will be replaced with a hash, effectively bounding the depth of the stack encoded in the continuation. To the best of our knowledge, no prior work has considered authenticated closures. Another way to achieve the same effect, but perhaps less elegantly, is to use an explicit authenticated stack as is done by *update_stk* given at the bottom of the figure.

The designer could also tune performance by adjusting the definition of the data structure itself. For example, we could have defined Merkle trees instead as follows:

```

type tree = Tip of string | Bin of •(tree × tree)

```

In this case, we are only hashing nodes, and will never hash tips. This definition makes more sense when the hash of the *Tip* is larger

⁴In general, the prover will return the shallow projection of the result of a computation back to the verifier; when the result is a normal value the prover will thus return the value itself (as with the result in our ($\text{fetch } [R; L] t$) example query).

Types $\tau ::= 1 \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 + \tau_2 \mid \tau_1 \times \tau_2 \mid \mu\alpha.\tau \mid \alpha \mid \bullet\tau$
 Values $v ::= () \mid x \mid \lambda x.e \mid \mathbf{rec} x.\lambda y.e$
 $\quad \mid \mathbf{inj}_1 v \mid \mathbf{inj}_2 v \mid (v_1, v_2) \mid \mathbf{roll} v$
 Exprs $e ::= v \mid \mathbf{let} x = e_1 \mathbf{in} e_2 \mid v_1 v_2 \mid \mathbf{case} v v_0 v_1$
 $\quad \mid \mathbf{prj}_1 v \mid \mathbf{prj}_2 v \mid \mathbf{unroll} v \mid \mathbf{auth} v \mid \mathbf{unauth} v$

Figure 5. Syntax for types and terms

than the representation of Tip itself, e.g., if the tree stored integers instead of strings. As another variation, we might imagine defining a tree that only optionally authenticates its children:

```

type tree =
  | Tip of string
  | Bin of tree × tree
  | AuthBin of •(tree × tree)

```

Then the tree might go several levels using Bin before using AuthBin. This design thus increases the constant factor on asymptotic space usage, but may reduce proving/verification time.

All of these customizations are possible, and easy to experiment with, thanks to the fact that $\lambda\bullet$ is a general-purpose programming language. However, this flexibility cuts both ways: there is nothing (at the moment) stopping the programmer from producing a suboptimal design. As an extreme example, the programmer could write `type tree = Tip of string | Bin of tree × tree`—i.e., without any use of authenticated types! This design will be secure and correct, as with every $\lambda\bullet$ program, but will effectively require the verifier to maintain the entire tree, not simply a digest. Fortunately, there is a simple rule of thumb that may have already become evident to the reader by this point: the data-structure type definition should authenticate recursive references, thus aiming for shallow projections to be constant-sized. We leave to interesting future work the task of automating the transformation of a \bullet -free type definition to an efficient authenticated one.

3. $\lambda\bullet$: A Language with Authenticated Types

This section formalizes $\lambda\bullet$, our language for writing computations over authenticated data structures. We present the syntax, typing rules, and operational semantics for $\lambda\bullet$ programs. The next section proves that $\lambda\bullet$ computations produce correct and secure results.

3.1 Syntax

Figure 5 presents the syntax for $\lambda\bullet$. Other than authenticated types $\bullet\tau$, the type language is entirely standard, consisting of the unit type 1, function types $\tau_1 \rightarrow \tau_2$, sum types $\tau_1 + \tau_2$, product types $\tau_1 \times \tau_2$, recursive types $\mu\alpha.\tau$, and variable types α arising from these. In this syntax, our authenticated tree type defined in Figure 2 would be written $\mu\alpha.\mathbf{string} + (\bullet\alpha \times \bullet\alpha)$, where `string` would itself be encoded, e.g., as a list of Peano-style integers. Our formal language does not include parametric polymorphism for simplicity, but adding it would present no difficulties. The language does not support references because mutations would risk invalidating hashes for $\bullet\tau$ values. In particular, given an authenticated value $\langle h, v \rangle$ where v is a reference, a mutation via v may invalidate h .

Terms (values v and expressions e) are in administrative normal form [7] to keep the semantics simple. In this form, the grammar forces us to write `let $x = e_1$ in let $y = e_2$ in $x y$` instead of the more familiar `$e_1 e_2$` , for example. In addition to variables x , the term language includes functions $\lambda x.e$ and function application $v_1 v_2$; sum-type values $\mathbf{inj}_1 v$ and $\mathbf{inj}_2 v$ which are eliminated by `case $v v_0 v_1$` , where v_0 and v_1 are expected to be functions;

$$\frac{\Gamma \vdash v : \tau_1}{\Gamma \vdash \mathbf{inj}_1 v : \tau_1 + \tau_2} \quad \frac{\Gamma \vdash v : \tau_2}{\Gamma \vdash \mathbf{inj}_2 v : \tau_1 + \tau_2}$$

$$\frac{\Gamma \vdash v : \tau_1 + \tau_2 \quad \Gamma \vdash v_1 : \tau_1 \rightarrow \tau \quad \Gamma \vdash v_2 : \tau_2 \rightarrow \tau}{\Gamma \vdash \mathbf{case} v v_1 v_2 : \tau}$$

$$\frac{\Gamma \vdash v : \tau}{\Gamma \vdash \mathbf{auth} v : \bullet\tau} \quad \frac{\Gamma \vdash v : \bullet\tau}{\Gamma \vdash \mathbf{unauth} v : \tau}$$

Figure 6. Selected typing rules

products (v_1, v_2) eliminated by expressions $\mathbf{prj}_1 v$ and $\mathbf{prj}_2 v$; values of recursive type introduced via `roll v` and eliminated by `unroll v` ; and finally fixpoints `rec $x.\lambda y.e$` for defining recursive functions (where inside of $\lambda y.e$ references to x refer to the function itself). Authenticated types $\bullet\tau$ are introduced by coercion `auth` and eliminated by `unauth`.

3.2 Typing

The typing judgment for $\lambda\bullet$ programs is the usual one, written $\Gamma \vdash e : \tau$. It states that expression e has type τ under environment Γ , where Γ is a map from variables x to types τ . Typing rules for most constructs are standard. Selected rules are given in Figure 6.

3.3 Operational semantics

In practice, our compiler takes a program like the one in Figure 2 and outputs versions to be run by the prover and the verifier. In our formalization, we define distinct semantics for the same program, as determined by an execution mode m , where $m = P$ for the prover’s execution, and $m = V$ for the verifier’s. We also define a mode I for the *Ideal* case, representing a computation that happens in the normal way, ignoring authenticated types; this is needed for stating the security and correctness properties.

We define a small-step operational semantics of the form $\ll \pi, e \gg \rightarrow_m \ll \pi', e' \gg$, where m is the mode and π is the proof stream, which is a list of shallow projections s . This can be read: *An expression e coupled with a proof stream π can evaluate one step in mode m to produce an expression e' and an updated proof stream π' .* We define $\ll \pi, e \gg \rightarrow_m^i \ll \pi', e' \gg$ to be the transitive multi-step application of the single-step relation; it states that *e evaluates, in i steps, to e' in mode m , starting with proof stream π and finishing with π' .* The proof stream is *produced* in mode P , so π is a prefix of π' in this mode. The stream is *consumed* in mode V , and thus π' is a suffix of π . The proof stream is ignored in mode I . We use the operator $\textcircled{\@}$ to denote the concatenation of two proof streams, treating $\textcircled{\@}$ as associative with the empty stream $[]$ as the identity. We write $[s]$ as the singleton stream containing element s .

The rules for standard language features are identical in all three modes, and are standard. They are defined in the top portion of Figure 7, and we discuss them briefly in order. The rule for function application, $(\lambda x.e) v$, substitutes v for x in e —this substitution is written $e[v \setminus x]$. Application of a recursive function is similar: when the function (`rec $x.\lambda y.e$`) is on the left-hand side of an application, we substitute x in the function body e with the recursive function itself. Let-binding is used to sequence computations, either evaluating the bound expression e_1 one step or else, if this expression is a value v_1 , substituting that value for x in the body e_2 . The semantics of `case` depends on whether it is given $\mathbf{inj}_1 v$ or $\mathbf{inj}_2 v$; in the former case we substitute v in the first lambda term (the “true branch”), else we substitute it in the second (“false”) one. Projection from a pair (v_1, v_2) produces v_1 for \mathbf{prj}_1 and v_2 for \mathbf{prj}_2 . Finally, the recursive type coercions `unroll` and `roll` nullify each other.

$$\begin{array}{l}
\ll \pi, (\lambda x.e) v \gg \quad \rightarrow_m \ll \pi, e[v \setminus x] \gg \\
\ll \pi, (\mathbf{rec} x.\lambda y.e) v \gg \quad \rightarrow_m \ll \pi, (\lambda y.e') v \gg \\
\text{where } e' = e[(\mathbf{rec} x.\lambda y.e) \setminus x] \\
\ll \pi, \mathbf{let} x = v_1 \mathbf{in} e_2 \gg \quad \rightarrow_m \ll \pi, e_2[v_1 \setminus x] \gg \\
\ll \pi, \mathbf{case} (\mathbf{inj}_1 v) (\lambda x.e_1) (\lambda x.e_2) \gg \quad \rightarrow_m \ll \pi, e_1[v \setminus x] \gg \\
\ll \pi, \mathbf{case} (\mathbf{inj}_2 v) (\lambda x.e_1) (\lambda x.e_2) \gg \quad \rightarrow_m \ll \pi, e_2[v \setminus x] \gg \\
\ll \pi, \mathbf{prj}_1 (v_1, v_2) \gg \quad \rightarrow_m \ll \pi, v_1 \gg \\
\ll \pi, \mathbf{prj}_2 (v_1, v_2) \gg \quad \rightarrow_m \ll \pi, v_2 \gg \\
\ll \pi, \mathbf{unroll} (\mathbf{roll} v) \gg \quad \rightarrow_m \ll \pi, v \gg \\
\hline
\ll \pi, e_1 \gg \rightarrow_m \ll \pi', e'_1 \gg \\
\ll \pi, \mathbf{let} x = e_1 \mathbf{in} e_2 \gg \rightarrow_m \ll \pi', \mathbf{let} x = e'_1 \mathbf{in} e_2 \gg \\
\hline
\ll \pi, e \gg \xrightarrow{i} \ll \pi', e' \gg \\
\ll \pi', e' \gg \rightarrow_m \ll \pi'', e'' \gg \\
\ll \pi, e \gg \xrightarrow{i+1} \ll \pi'', e'' \gg \quad \ll \pi, e \gg \xrightarrow{0} \ll \pi, e \gg
\end{array}$$

Figure 7. Standard single-step and multi-step operational rules

$$\begin{array}{l}
\ll \pi, \mathbf{auth} v \gg \quad \rightarrow_I \ll \pi, v \gg \\
\ll \pi, \mathbf{unauth} v \gg \quad \rightarrow_I \ll \pi, v \gg \\
\ll \pi, \mathbf{auth} v \gg \quad \rightarrow_P \ll \pi, \langle \mathbf{hash} ([v]), v \rangle \gg \\
\ll \pi, \mathbf{unauth} \langle h, v \rangle \gg \quad \rightarrow_P \ll \pi @ [([v])], v \gg \\
\ll \pi, \mathbf{auth} v \gg \quad \rightarrow_V \ll \pi, \mathbf{hash} v \gg \\
\hline
\text{hash } s_0 = h \\
\ll [s_0] @ \pi, \mathbf{unauth} h \gg \rightarrow_V \ll \pi, s_0 \gg \\
\hline
\text{where } v ::= \dots | h | \langle h, v \rangle
\end{array}$$

Figure 8. Operational rules for authenticated values

$$\begin{array}{ll}
\langle () \rangle & = () & \langle [x] \rangle & = x \\
\langle \langle h, v \rangle \rangle & = h & \langle [\lambda x.e] \rangle & = \lambda x. \langle [e] \rangle \\
\langle [\mathbf{auth} v] \rangle & = \mathbf{auth} ([v]) & \langle [\mathbf{unauth} v] \rangle & = \mathbf{unauth} ([v]) \\
\langle \langle [v_1], [v_2] \rangle \rangle & = \langle ([v_1]), ([v_2]) \rangle & \langle [\mathbf{prj}_i v] \rangle & = \mathbf{prj}_i ([v]) \\
\langle [\mathbf{roll} v] \rangle & = \mathbf{roll} ([v]) & \langle [\mathbf{unroll} v] \rangle & = \mathbf{unroll} ([v]) \\
\langle [\mathbf{rec} x.\lambda y.e] \rangle & = \mathbf{rec} x. \langle [\lambda y.e] \rangle & \langle [\mathbf{inj}_i v] \rangle & = \mathbf{inj}_i ([v]) \\
\langle [\mathbf{case} v v_0 v_1] \rangle & = \mathbf{case} ([v]) ([v_0]) ([v_1]) \\
\langle [\mathbf{let} x = e_1 \mathbf{in} e_2] \rangle & = \mathbf{let} x = \langle [e_1] \rangle \mathbf{in} \langle [e_2] \rangle
\end{array}$$

Figure 9. Shallow projection of an expression e , written $\langle [e] \rangle$

The rules for the multi-step relation are given at the bottom of Figure 7, and are also standard.

The operational rules for authenticated values are given in Figure 8. For mode I , authenticated values of type $\bullet\tau$ are merely values of type τ and the $\mathbf{auth}/\mathbf{unauth}$ operations are no-ops. For mode P , values of type $\bullet\tau$ are implemented as a pair $\langle h, v \rangle$ of a hash h and a value v (of type τ). As shown in the \mathbf{auth} rule, the hash is computed by applying a hash function \mathbf{hash} over the shallow projection of v , written $\langle [v] \rangle$. We do not formalize the semantics of \mathbf{hash} explicitly; in practice it can be implemented by serializing the value it is given and hashing that using a collision-resistant hash function.⁵ The shallow projection operation is defined in Fig-

⁵For functions $\lambda x.e$, serialization involves pretty-printing the function’s code such that alpha-equivalent functions will have the same hash value. We discuss implementing authenticated functions more in Section 5.3.

ure 9. It is essentially a fold over the structure of the term, preserving that structure in every case but that of values $\langle h, v \rangle$: here we simply drop the value v and retain the hash h . Another interesting case is functions $\lambda x.e$: we recursively descend into e to translate any $\langle h, v \rangle$ values that appear there. Such values will not appear in source programs, but they can arise via substitution under lambdas. Returning to Figure 8, the mode- P semantics of $\mathbf{unauth} \langle h, v \rangle$ is to strip off the hash, returning v , while adding the shallow projection of v to the end of the proof stream. Finally, for mode V the representation of $\bullet\tau$ is the hash h of a value of type τ . The \mathbf{auth} rule constructs this representation, while the \mathbf{unauth} rule checks that the hash value matches the shallow projection at the head of the proof stream.

4. Metatheory

We want to show that well-typed $\lambda\bullet$ programs will (a) produce correct results—that is, results that all three modes agree on—or else (b) a malicious prover has been able to find a hash collision, which by assumption is computationally difficult. We call property (a) *correctness* and property (b) *security*. In this section we state and prove these two properties.

4.1 Type Soundness

To begin, we want to prove that $\lambda\bullet$ ’s design is sensible in that the ideal semantics is sound (and entirely ordinary). We can prove the standard type soundness lemma about the ideal mode’s semantics.

Lemma 1 (Type Soundness). *If $\Gamma \vdash e : \tau$, then either e is a value, or there exists e' and $i > 0$ such that $\ll [], e \gg \xrightarrow{i} \ll [], e' \gg$ and $\Gamma \vdash e' : \tau$.*

Proof. The proof is completely standard, using progress and preservation lemmas, and inducting on $\ll [], e \gg \xrightarrow{i} \ll [], e' \gg$. \square

4.2 Agreement

Now we must define what we mean when we say that the different execution modes “agree” on their results—it cannot be that these results are syntactically equal because each mode interprets authenticated values differently. For example, consider the update function from Figure 4. In the ideal setting, this function will return a normal tree v —because $\bullet\tau$ values are the same as those of type τ , the tree v will contain no digests. On the other hand, the prover will return a value $\langle h, v_P \rangle$, where h is the digest of v_P . For the same insertion on the same tree, the results v and $\langle h, v_P \rangle$ in I and P modes, respectively, should “agree” without being equal: v will just be a normal tree, while v_P will contain digests—but the elements and sub-trees, excepting the digests, should be the same. And, running the insertion at the verifier will return a digest h , which should match the digest in the prover’s returned value $\langle h, v_P \rangle$.

We formalize this connection as a three-way type-indexed relation $\Gamma \vdash e e_P e_V : \tau$, given in Figure 10, which states: *in environment Γ , ideal expression e , prover expression e_P , and verifier expression e_V all agree at type τ* . In every case but that of authenticated values (the last rule), agreement follows syntactic structure of the terms, and the shape of each rule matches that of the standard type rules. The rule for authenticated values formalizes the intuition given above; it states that $\Gamma \vdash v \langle h, v_P \rangle h : \bullet\tau$ holds when (a) the digest h of both the prover and verifier is the same; (b) this digest is the hash of the shallow projection of the prover’s value v_P ; (c) the prover’s value v_P agrees with the ideal value v .

Now we prove some useful facts about terms in agreement.

Lemma 2 (Agreement). *Suppose $\Gamma \vdash e e_P e_V : \tau$. Then*

1. $\langle [e_P] \rangle = e_V$.
2. $\Gamma \vdash e e'_P e'_V : \tau$ implies that $e'_P = e_P$ and $e_V = e'_V$.

$$\begin{array}{c}
\Gamma \vdash () () () : 1 \\
\frac{\Gamma(x) = \tau}{\Gamma \vdash x x x : \tau} \quad \frac{\Gamma, x:\tau_1 \vdash e \ e_P \ e_V : \tau_2}{\Gamma \vdash (\lambda x.e) (\lambda x.e_P) (\lambda x.e_V) : \tau_1 \rightarrow \tau_2} \\
\frac{\Gamma \vdash v_1 \ v_{1P} \ v_{1V} : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash v_2 \ v_{2P} \ v_{2V} : \tau_1}{\Gamma \vdash (v_1 \ v_2) (v_{1P} \ v_{2P}) (v_{1V} \ v_{2V}) : \tau_2} \\
\frac{\Gamma \vdash e_1 \ e_{1P} \ e_{1V} : \tau_1 \quad \Gamma, x:\tau_1 \vdash e_2 \ e_{2P} \ e_{2V} : \tau_2}{\Gamma \vdash (\mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2) (\mathbf{let} \ x = e_{1P} \ \mathbf{in} \ e_{2P}) (\mathbf{let} \ x = e_{1V} \ \mathbf{in} \ e_{2V}) : \tau_2} \quad \frac{\Gamma, x:\tau_1 \rightarrow \tau_2 \vdash (\lambda y.e) (\lambda y.e_P) (\lambda y.e_V) : \tau_1 \rightarrow \tau_2}{\Gamma \vdash (\mathbf{rec} \ x.\lambda y.e) (\mathbf{rec} \ x.\lambda y.e_P) (\mathbf{rec} \ x.\lambda y.e_V) : \tau_1 \rightarrow \tau_2} \\
\frac{\Gamma \vdash v \ v_P \ v_V : \tau_1}{\Gamma \vdash (\mathbf{inj}_1 \ v) (\mathbf{inj}_1 \ v_P) (\mathbf{inj}_1 \ v_V) : \tau_1 + \tau_2} \quad \frac{\Gamma \vdash v \ v_P \ v_V : \tau_2}{\Gamma \vdash (\mathbf{inj}_2 \ v) (\mathbf{inj}_2 \ v_P) (\mathbf{inj}_2 \ v_V) : \tau_1 + \tau_2} \\
\frac{\Gamma \vdash v \ v_P \ v_V : \tau_1 + \tau_2 \quad \Gamma \vdash v_P \ v_{1P} \ v_{1V} : \tau_1 \rightarrow \tau \quad \Gamma \vdash v_V \ v_{1V} \ v_{2V} : \tau_2 \rightarrow \tau}{\Gamma \vdash (\mathbf{case} \ v \ v_1 \ v_2) (\mathbf{case} \ v_P \ v_{1P} \ v_{2P}) (\mathbf{case} \ v_V \ v_{1V} \ v_{2V}) : \tau} \quad \frac{\Gamma \vdash v_1 \ v_{1P} \ v_{1V} : \tau_1 \quad \Gamma \vdash v_2 \ v_{2P} \ v_{2V} : \tau_2}{\Gamma \vdash (v_1, v_2) (v_{1P}, v_{2P}) (v_{1V}, v_{2V}) : \tau_1 \times \tau_2} \\
\frac{\Gamma \vdash v \ v_P \ v_V : \tau_1 \times \tau_2}{\Gamma \vdash (\mathbf{prj}_1 \ v) (\mathbf{prj}_1 \ v_P) (\mathbf{prj}_1 \ v_V) : \tau_1} \quad \frac{\Gamma \vdash v \ v_P \ v_V : \tau_1 \times \tau_2}{\Gamma \vdash (\mathbf{prj}_2 \ v) (\mathbf{prj}_2 \ v_P) (\mathbf{prj}_2 \ v_V) : \tau_2} \quad \frac{\Gamma \vdash v \ v_P \ v_V : \tau[\mu\alpha.\tau \setminus \alpha]}{\Gamma \vdash (\mathbf{roll} \ v) (\mathbf{roll} \ v_P) (\mathbf{roll} \ v_V) : \mu\alpha.\tau} \\
\frac{\Gamma \vdash v \ v_P \ v_V : \mu\alpha.\tau}{\Gamma \vdash (\mathbf{unroll} \ v) (\mathbf{unroll} \ v_P) (\mathbf{unroll} \ v_V) : \tau[\mu\alpha.\tau \setminus \alpha]} \quad \frac{\Gamma \vdash v \ v_P \ v_V : \tau}{\Gamma \vdash (\mathbf{auth} \ v) (\mathbf{auth} \ v_P) (\mathbf{auth} \ v_V) : \bullet\tau} \\
\frac{\Gamma \vdash v \ v_P \ v_V : \bullet\tau}{\Gamma \vdash (\mathbf{unauth} \ v) (\mathbf{unauth} \ v_P) (\mathbf{unauth} \ v_V) : \tau} \quad \frac{\vdash v \ v_P \ ([v_P]) : \tau \quad \mathbf{hash}([v_P]) = h}{\Gamma \vdash v \ \langle h, v_P \rangle \ h : \bullet\tau}
\end{array}$$

Figure 10. Agreement relation: defines those expressions that *agree* (i.e., that are morally, if not syntactically, the same) in the Ideal, Prover, and Verifier modes. The most interesting rule is the last one, while the rest are three-way versions of the standard type rules.

3. $\Gamma \vdash e : \tau$

4. Either e , e_P , and e_V are all values, or none of them are.

Proof. By induction on $\Gamma \vdash e \ e_P \ e_V : \tau$. \square

The first part shows the agreement relation is intimately connected to the shallow projection operator—a prover’s term only ever agrees with a verifier’s term when the latter is the shallow projection of the former. Moreover, we prove for any given ideal term e , there is at most one pair of terms e_P and e_V that agree with it under a given environment Γ and type τ , that agreement implies e is well-typed, and that values only agree with other values.

Client and server agree. In a client/server application scenario, a query/update sent by the client will reference the data structure stored at the server using a free variable, e.g., the t in the query `fetch t`. To run this query on the server, we substitute the prover’s representation for t , while to verify the result at the client, we substitute t ’s digest. These representations should agree. The following lemma states that, given an expression e containing free variables with authenticated types, substituting authenticated values that agree for the free variables of e produces versions e_I , e_P , and e_V that also agree.

Lemma 3. *Given the following:*

1. $\Gamma \vdash e : \tau$ where e contains no values of type $\bullet\tau$
2. For all $x_i \in \text{domain}(\Gamma)$,
 - (a) $\Gamma(x_i) = \bullet\tau_i$ for some τ_i
 - (b) $\vdash v_i \ \langle h_i, v_{P_i} \rangle \ h_i : \bullet\tau_i$ for some $\langle h_i, v_i \rangle$ and v_{P_i}
3. $e_P = e[\langle h_1, v_{P_1} \rangle \setminus x_1] \dots [\langle h_n, v_{P_n} \rangle \setminus x_n]$
 $e_V = e[h_1 \setminus x_1] \dots [h_n \setminus x_n]$
 $e_I = e[v_1 \setminus x_1] \dots [v_n \setminus x_n]$

Then $\vdash e_I \ e_P \ e_V : \tau$.

The proof of this lemma follows from straightforward application of the following substitution lemma:

Lemma 4 (Substitution). *If $\Gamma, x:\tau' \vdash e \ e_P \ e_V : \tau$ and $\vdash v \ v_P \ v_V : \tau'$, then $\Gamma \vdash (e[v \setminus x]) (e_P[v_P \setminus x]) (e_V[v_V \setminus x]) : \tau$.*

Proof. The proof is by induction on $\Gamma, x:\tau' \vdash e \ e_P \ e_V : \tau$. The only interesting case is when $\Gamma, x:\tau' \vdash v' \ \langle h, v'_P \rangle \ h : \bullet\tau$. The empty environment in the premise $\vdash v' \ v'_P \ ([v'_P]) : \tau$ ensures that v' and v'_P contain no variables, so the substitution will be the identity and the result follows by assumption. \square

4.3 Correctness and Security

Now we can state and prove our main theorem, Theorem 1, which encapsulates the two properties of interest, correctness and security. For both properties, we start with the assumption that terms e , e_P , and e_V agree (which will be the case at the start of evaluating a query/update as per Lemma 3). The ideal-mode evaluation represents the specification of correctness: if e can evaluate to e' in ideal mode in i steps, then the verifier, when consuming the proof stream π produced by the prover’s evaluation of e_P , should evaluate to some e'_V which (along with the prover’s resulting term e'_P) agrees with e' . On the other hand, if the verifier does not consume π but rather some other, adversarially chosen stream π_A that does not contain π as a prefix (e.g., because the server is behaving maliciously or incorrectly), then the only way the verifier can accept an *incorrect* result is if the adversary has found a *hash collision*. That is, the consumed stream π_A contains an element s^\dagger that corresponds to an element s in π such that $s \neq s^\dagger$ but $\text{hash } s = \text{hash } s^\dagger$.

As discussed further in Section 4.4, this implies the standard cryptographic notion of security for this setting if hash is collision-resistant. Here is the theorem, stated formally:

Theorem 1. *Suppose that $\vdash e \ e_P \ e_V : \tau$.*

Correctness: *If $\ll \square, e \gg \rightarrow_I^i \ll \square, e' \gg$ then there exist e'_P, e'_V, π such that*

- $\ll \square, e_P \gg \rightarrow_P^i \ll \pi, e'_P \gg$
- $\ll \pi, e_V \gg \rightarrow_V^i \ll \square, e'_V \gg$
- $\vdash e' \ e'_P \ e'_V : \tau$

Security: *If $\ll \pi_A, e_V \gg \rightarrow_V^i \ll \pi', e'_V \gg$ then*

1. *there exist e', e'_P, π , such that*
 - $\ll \square, e \gg \rightarrow_I^i \ll \square, e' \gg$
 - $\ll \square, e_P \gg \rightarrow_P^i \ll \square @ \pi, e'_P \gg$
 - $\pi_A = \pi @ \pi'$
 - $\vdash e' \ e'_P \ e'_V : \tau$
2. *or else there exist $j \leq i, e'_P, \pi_0, s$ and s^\dagger such that*
 - $\ll \square, e_P \gg \rightarrow_P^j \ll \square @ \pi_0 @ [s], e'_P \gg$
 - $\pi_A = \pi_0 @ [s^\dagger] @ \pi'$
 - $s \neq s^\dagger$ but $\text{hash } s = \text{hash } s^\dagger$.

The proof is by induction on the length i of the multi-step derivations, relying on two lemmas about the Correctness and Security of single-step evaluation, which we present next.

Lemma 5 (Correctness). *If $\vdash e \ e_P \ e_V : \tau$ and $\ll \square, e \gg \rightarrow_I \ll \square, e' \gg$ then there exist e'_P, e'_V , and π such that for all π', π_p*

1. $\vdash e' \ e'_P \ e'_V : \tau$
2. $\ll \pi_p, e_P \gg \rightarrow_P \ll \pi_p @ \pi, e'_P \gg$
3. $\ll \pi @ \pi', e_V \gg \rightarrow_V \ll \pi', e'_V \gg$.

Proof. By induction on $\vdash e \ e_P \ e_V : \tau$. Most cases are straightforward, and follow by application of the Substitution lemma. The two interesting cases deal with authenticated computations:

- Suppose e, e_P , and e_V are *auth* v , *auth* v_P and *auth* v_V , respectively. Each can take a step in its respective mode, producing $v, \langle \text{hash}([v_P]), v_P \rangle$, and $\text{hash } v_V$, respectively, leaving the proof stream unchanged (i.e., $\pi = \square$). Now we must prove $\vdash v \ \langle \text{hash}([v_P]), v_P \rangle \ \text{hash } v_V : \bullet\tau$, which in turn requires proving $\vdash v \ v_P \ ([v_P]) : \tau$ and $\text{hash}([v_P]) = \text{hash } v_V$. Both are the consequence of Lemma 2.1 and $\vdash e \ e_P \ e_V : \tau$.
- Suppose e, e_P , and e_V are *unauth* v , *unauth* v_P and *unauth* v_V , respectively. By inversion on $\vdash e \ e_P \ e_V : \tau$ we know that $\vdash v \ \langle h, v_P \rangle \ h : \bullet\tau$ and by inversion on this we know $\vdash v \ v_P \ ([v_P]) : \tau$ and $h = \text{hash}([v_P])$. We can set $\pi = [[v_P]]$, and then each term can take a step in its respective mode to v, v_P , and $[v_P]$, which agree by Lemma 2.1. \square

Finally, we demonstrate that a verifier term that begins in agreement and takes a step remains in agreement unless an adversary has managed to find a hash collision.

Lemma 6 (Security). *Given the following:*

- $\vdash e \ e_P \ e_V : \tau$
- $\ll \pi_A, e_V \gg \rightarrow_V \ll \pi', e'_V \gg$

then there exist e', e'_P , and π such that for all π_p ,

- $\ll \square, e \gg \rightarrow_I \ll \square, e' \gg$
- $\ll \pi_p, e_P \gg \rightarrow_P \ll \pi_p @ \pi, e'_P \gg$

and either

1. $\vdash e' \ e'_P \ e'_V : \tau$ and $\pi_A = \pi @ \pi'$, or else

2. *there exists a pair s and s^\dagger such that $\pi = [s]$ and $\pi_A = [s^\dagger] @ \pi'$ with $s \neq s^\dagger$ but $\text{hash } s = \text{hash } s^\dagger$.*

Proof. By induction on $\vdash e \ e_P \ e_V : \tau$. Since e_V is not a value, we know from Lemma 2.4 that neither are e or e_P , so we can always introduce e' and e'_P . Most cases are straightforward because evaluation yields $\pi = \square$ and $\pi_A = \pi'$, and $\vdash e' \ e'_P \ e'_V : \tau$ can be obtained directly from inversion on $\vdash e \ e_P \ e_V : \tau$. The remaining cases are for let binding and *unauth*. The former follows by induction. For the latter we have e, e_P , and e_V are *unauth* $v, \text{unauth } v_P$ and *unauth* v_V , respectively. Since these terms agree by assumption, we know that $v_P = \langle h, v'_P \rangle$ and $v_V = h = \text{hash}([v'_P])$ for some v'_P . From the stepping rule for \rightarrow_P , we know $\pi = [[v'_P]]$. Then there are three possible outcomes depending on π_A :

1. $\pi_A = \square$, or $\pi_A = [s^\dagger] @ \pi'$ and $\text{hash } s^\dagger \neq h$, in which case $\ll \pi_A, e_V \gg$ is stuck and we have a contradiction
2. $\pi_A = [[v'_P]] @ \pi'$ and $e'_V = ([v'_P])$, from which $\vdash e' \ e'_P \ e'_V : \tau$ follows directly
3. $\pi_A = [s^\dagger] @ \pi'$ and $([v'_P]) \neq s^\dagger$, but $\text{hash}([v'_P]) = \text{hash } s^\dagger$. \square

Although the Security property guarantees that the ideal computation can always take as many steps as the verifier (in particular, the verifier cannot run forever if the ideal computation terminates), we would also like to show that ideal computation can take as many steps as the prover.

Remark 1. *Suppose $\vdash e \ e_P \ e_V : \tau$ and*

$\ll \pi_p, e_P \gg \rightarrow_P^i \ll \pi_p @ \pi, e'_P \gg$. Then there exists e', e'_V such that $\vdash e' \ e'_P \ e'_V : \tau$, $\ll \square, e \gg \rightarrow_I^i \ll \square, e' \gg$, and $\ll \pi, e_V \gg \rightarrow_V^i \ll \square, e'_V \gg$.

Proof. This follows from straightforward induction on derivation length i and $\vdash e \ e_P \ e_V : \tau$, applying Lemmas 2 and 5.

4.4 Cryptographic Security

In the cryptographic security definition for a fixed ADS protocol, (e.g., as per Papamanthou et al. [25]), somewhat informally, there is an attacker who is assumed able to control the interaction between an honest prover and verifier. The attacker may specify a sequence of operations (“queries”) q_1, \dots, q_t that the verifier poses to the prover. For each such query, the prover generates and sends a proof to the verifier; both parties update their local state as appropriate. Finally, the attacker specifies a query q_{t+1} along with an adversarially generated proof string π_A ; the attacker *succeeds* if π_A causes the verifier to output an incorrect result for the given query. The ADS is parameterized by a *security parameter* k which we may identify with the output length of the hash function being used. The ADS is *secure* if no attacker running in polynomial time (in k) succeeds with non-negligible probability. Security is proven by contradiction with the collision-resistance of an appropriately chosen hash function.

To translate the standard cryptographic notion to our setting, we must address three technicalities. First, we must provide a notion of programs and inputs, such that the ADS is specified by an arbitrary $\lambda\bullet$ program, and the adversary is allowed to choose the inputs. Second, in order to show contradiction with collision-resistance, we must define a specific procedure by which the hash function is instantiated *after* the ADS program and the adversary are fixed. Finally, we must relate the number of reduction steps taken by a $\lambda\bullet$ program to a number of steps taken by a Turing machine. We then claim that the reduction argument in Theorem 1 implies that this security definition holds for every $\lambda\bullet$ program.

Inputs. We can treat the free variables in an open $\lambda\bullet$ expression as program inputs. In our running example, for instance, the expression `fetch idx t` has `idx` and `t` as free variables. The adversary chooses the inputs by computing well-typed and in-agreement values to substitute for each of the free variables.

Choosing the hash function. Our language is parameterized by an arbitrary hash function, which we have assumed is collision-resistant. However, to be precise, collision-resistance is only defined formally for a *family* of hash functions rather than some fixed hash function (see Katz and Lindell [14]). That is, collision-resistance is defined according to the following game: First, take an arbitrary adversary that runs in polynomial-time given a security parameter k . Next, choose hash randomly from a family of functions, and give the adversary a description of hash as input. The family of hash functions is *collision-resistant* if the the adversary outputs a collision in hash with negligible probability in k .

The following game captures this notion in the context of $\lambda\bullet$:

1. Let e be an arbitrary well-typed $\lambda\bullet$ program specifying the data structure being supported, and let \mathcal{A} be an arbitrary adversary that runs in polynomial-time (in k).
2. Choose hash at random from a collision-resistant family of k -bit hash functions, and then compile e_P and e_V .
3. \mathcal{A} *succeeds* if, given k and hash as input, it outputs in-agreement values to substitute for the free variables of e , e_P , and e_V , and a proof stream $\pi_{\mathcal{A}}$, such that after some polynomial of steps i , the verifier outputs an incorrect answer; i.e., a value e'_V such that $\ll \pi_{\mathcal{A}}, e_V \gg \rightarrow_V^i \ll \pi'_{\mathcal{A}}, e'_V \gg$ for some $\pi'_{\mathcal{A}}$, but there is no e' , e'_P such that $\ll [], e \gg \rightarrow_I^i \ll [], e' \gg$ and $\vdash e' e'_P e'_V : \tau$.

Note that the game begins with an arbitrary well-typed $\lambda\bullet$ program e , before the hash function has been chosen. Although $\lambda\bullet$ is parameterized by hash, the ideal terms are invariant to the choice of hash function, so we can fix e and $\Gamma \vdash e : \tau$ before choosing hash. However, the prover and verifier terms e_P and e_V may actually contain digests, so we must instantiate hash before compiling e_P and e_V .

If \mathcal{A} succeeds, then by Theorem 1 we can evaluate e_P and e_V (on $\pi_{\mathcal{A}}$) for $j \leq i$ steps, and extract the collision s and s^\dagger . Thus, it only remains for us to show that evaluation of e_V and e_P also takes polynomial time, and then we can obtain from \mathcal{A} a polynomial-time collision-finding algorithm with the same success probability.

Polynomial-time execution. It is well known that standard lambda-calculus evaluation and Turing-machine execution are polynomially equivalent [4]; that is, a lambda-calculus interpreter implemented as a Turing machine can simulate i lambda-calculus evaluation steps in $O(\text{poly}(i))$ Turing-machine steps. The only nonstandard terms in $\lambda\bullet$ are *auth* and *unauth*, and each involves at most one hash computation during evaluation. The time to compute a hash is proportional to the size of the serialized shallow projection, which depends on the hash output length k . Therefore, if an ideal program e takes i steps to reach a value, a Turing machine can simulate the execution of e in $O(\text{poly}(i))$ steps and the corresponding e_P and e_V in $O(\text{poly}(ik))$ steps. Since we assume i is bounded by $O(\text{poly}(k))$, the entire evaluation is $O(\text{poly}(k))$.

5. Implementation

This section describes our prototype extension to the OCaml compiler for supporting authenticated types. We discuss basic compilation,⁶ two optimizations we implement, and current limitations.

⁶Our technique for extending the OCaml compiler is based on a 2012 blog post by Jun Furuse: <https://bitbucket.org/camlspotter/compiler-libs-hack>

5.1 Compilation

The compilation process works as follows. The programmer writes an OCaml program p like that of Figure 2 that contains uses of authenticated types. The code will link against the ADS module, whose signature declares $\bullet\alpha$ as abstract⁷ and declares the (polymorphic) types of the *auth* and *unauth* coercions. Program p is then passed to our extended compiler which, depending on a command-line *mode* flag, replaces each application of *auth* and *unauth* it finds with a call to a prover- or verifier-specific implementation; the resulting code is linked with the ADS module.

This module, given in Figure 11, defines type $\bullet\alpha$ as either a digest (just the hash, represented as a string), or as a pair of the hash and a value of type α . The next four functions define the prover’s and verifier’s versions of *auth* and *unauth*, respectively; the calls to *auth* and *unauth* will be replaced by calls to these functions instead. We can see that their code largely matches the operational rules given in Figure 8, where the proof stream from the rules is implemented as OCaml channels, `prf_output` and `prf_input`. The one departure is that `auth_prover` and `unauth_prover` additionally take a function `shallow` that is invoked to perform the shallow projection operation. This operation is needed because OCaml does not provide a generic method for folding/mapping over the structure of a term. As such, our compiler generates type-specific shallow projection functions where needed, and includes them in the replaced calls to *auth* and *unauth*. The type of the shallow projection operator is determined by the concrete type inferred at each *auth/unauth* call. For example, the *unauth* in `let x : int = unauth y` is inferred to have type $\bullet\text{int} \rightarrow \text{int}$. Therefore, we need a shallow projection operation of type $\text{int} \rightarrow \text{int}$ (which is just the identity). The generated code will refer to the ADS module’s `shallow_•` function, shown at the bottom of the figure, for handling (nested) authenticated values.

Library functions are provided to enable the programmer to manipulate the proof and verification streams, for example by writing/reading to a file or a socket, or performing multiple authenticated operations with separate proofs in a single execution.

Figure 12 shows the result of compiling a variant of authenticated binary search trees. The top of the figure is the code provided by the programmer. Compilation will replace the call to *auth* with a call to `auth_bst1` and the call to *unauth* with a call to `unauth_bst`. These functions are defined at the bottom of the figure, and employ the needed, type-specific shallow projection operations.

The hash function referenced in Figure 11 is polymorphic, having type $\forall\alpha.\alpha \rightarrow \text{string}$. It is implemented by first serializing the argument and then hashing it using SHA1 (which is widely used as a collision-resistant hash function).

For serialization, we use OCaml’s default serializer implemented in the `Marshal` module. This choice has an implication for security: the worst-case cost to compute the hash of a malicious string is bounded only by the representation of an integer in OCaml, either 32 or 64 bits, depending on the OS. We used the *no-sharing* option for the `Marshal` module to guarantee that any two equal objects have equal serializations.⁸

⁷Type $\bullet\alpha$ is written in ASCII as `α authtype`, but we continue writing $\bullet\alpha$ for consistency.

⁸It may be the case that two terms of different types have equal serializations using `Marshal`; however, since we only interpret serialized values according to known static types (no polymorphism), it would take a collision *at the same type* to harm security. In our formalism, which is essentially dynamically typed, we assume serialization (encased in the abstract hash function) is bijective.

```

type  $\bullet\alpha = \mid$  Digest of string (* the digest *)
            $\mid$  Prover of string  $\times \alpha$ 

let auth_prover (shallow:  $\alpha \rightarrow \alpha$ ) (v: $\alpha$ ) :  $\bullet\alpha =$ 
  Prover(hash (shallow v), v)

let unauth_prover (shallow:  $\alpha \rightarrow \alpha$ ) (v: $\bullet\alpha$ ) :  $\alpha =$ 
  let Prover( $_,x$ ) = v in
  to_channel !prf_output (shallow x);
   $\times$ 

let auth_verifier (v: $\alpha$ ) :  $\bullet\alpha =$  Digest(hash v)

let unauth_verifier (v: $\bullet\alpha$ ) :  $\alpha =$ 
  let Digest(h) = v in
  let y = from_channel !prf_input in
  assert h = hash y;
  y

let shallow_ $\bullet$  (Prover(h, $_$ ) :  $\bullet\alpha$ ) :  $\bullet\alpha =$  Digest(h)

```

Figure 11. The implementation of type constructor \bullet and the Prover and Verifier’s *unauth* and *auth* coercions.

```

(* User-provided code *)
type bst = Tip
            $\mid$  Bin of  $\bullet$ bst  $\times$  int  $\times$   $\bullet$ bst
            $\mid$  AuthBin of  $\bullet$ (bst  $\times$  int  $\times$  bst)
let is_empty (t: $\bullet$ bst) : bool = (unauth t = Tip)
let mk_leaf (x:int) :  $\bullet$ bst = AuthBin(auth(Tip, x, Tip))

(* Generated Prover code *)
let rec shallow_bst : bst  $\rightarrow$  bst = function
   $\mid$  Tip  $\rightarrow$  Tip
   $\mid$  Bin (x, y, z)  $\rightarrow$  Bin(shallow_ $\bullet$  x, y, shallow_ $\bullet$  z)
   $\mid$  AuthBin (x)  $\rightarrow$  AuthBin (shallow_bst1 x)
and shallow_bst1 : bst  $\times$  int  $\times$  bst  $\rightarrow$  bst  $\times$  int  $\times$  bst
= function (x, y, z)  $\rightarrow$  (shallow_bst x, y, shallow_bst z)

let unauth_bst = unauth_prover shallow_bst
let auth_bst1 = auth_prover shallow_bst1

```

Figure 12. Example types and generated code (for prover)

5.2 Optimizations

Our compiler implements two optimizations that reduce the size of the proof stream, *reuse buffering*, and *suspended disbelief*.

Reuse buffering. We can reduce the size of the proof stream, and speed up proving/verification, when we anticipate that the same elements may appear in the proof stream more than once. For example, a client may submit a batch of operations to the server that end up re-traversing many of the nodes of the ADS. The client could cache the shallow projections of these elements locally instead of reading them from the proof stream multiple times.

This optimization requires modifying the *unauth* and *auth* functionality; the modification is similar for both prover and verifier. A counter is incremented each time *auth* or *unauth* is called. Each party maintains two data structures: *LRU-Map*, a mapping from *auth/unauth* counter values to shallow projections, indicating the least-recently-used ordering, and *Digest-Map*, a mapping from digests to *auth/unauth* counter values; collectively *Digest-Map* and *LRU-Map* are referred to as “the cache,” as they contain

corresponding elements. When *unauth* is called, if the digest exists in *Digest-Map*, the prover omits (resp., verifier fetches from the cache) the corresponding shallow projection, then updates the counter value associated with it in the cache; otherwise, the prover appends it to (resp., verifier reads it from) the proof stream and adds it to the cache. If the size of the cache exceeds the parameter, then the least recently used element (the smallest key in *LRU-Map*) is removed. The proof generation/checking/size benefits come at the cost of having to store the cache.

Suspending disbelief. This optimization eliminates redundancy in shallow projections that contain hashes computable from nested shallow projections appearing subsequently in the proof stream. This idea is implicit in the verification procedure for Merkle trees (Section 2.1). We can approximate the optimization for arbitrary data structures in $\lambda\bullet$ by modifying *unauth* (for both the prover and verifier) to use a buffer to “suspend disbelief.”

For the prover, the goal is to decide which digests in a shallow projection can safely be omitted, which are those that correspond to nodes that will be visited during subsequent calls to *unauth*. To achieve this, we extend the representation for $\bullet\alpha$ values:

```

type  $\bullet\alpha = \dots \mid$  Susp of string  $\times \alpha \times$  bool ref
            $\mid$  Sentinel

```

Each time *unauth* is called, the shallow projection is computed differently: each immediate child, *Prover*(d,a), would ordinarily be replaced with *Digest*(d) but is instead replaced with *Susp*(d,a,flag), where *flag* is a fresh mutable flag (initially **false**) that indicates whether the digest can be omitted. When one of these children is accessed by *unauth* the corresponding flag is set. Rather than writing the shallow projection immediately to the proof stream, it is appended to a queue. The queue is flushed when execution reaches a programmer-inserted *insist* annotation. This is a hint that belief need no longer be suspended; its placement has no effect on security but it is best used at the end of a distinct operation. Before a queue element is added to the proof stream, occurrences of *Susp*(d,a,flag) are (functionally) replaced with *Sentinel* when !flag is **true**, and replaced with *Digest*(d) when !flag is **false**.

When the verifier encounters a shallow projection object in the proof stream containing *Sentinel* values, its digest cannot be computed immediately so the object is stored in a set referred to as the *suspended-disbelief* buffer. If *unauth* is subsequently called on an immediate child of a node in this buffer, even if the shallow projection is available in the proof stream, the root hash is unknown so it cannot be immediately validated either. Thus we extend the $\bullet\alpha$ representation again with a new tag, *Suspension*.

```

type  $\bullet\alpha = \dots \mid$  Suspension of string ref  $\times$  (unit  $\rightarrow$  unit)

```

For every object placed in the buffer, each immediate *Sentinel* is replaced with a *Suspension* containing a callback closure and a mutable reference (initially empty) optionally containing a hash. When a leaf node is accessed, the *Suspension* reference is updated with the actual digest of the leaf, and the callback is invoked. The callback encapsulates a pointer to the parent node, and checks if all immediate *Suspension* children have been populated with digests. If so, the callback removes the node from the buffer, computes the node’s digest, and takes one of two actions: if the reference already contains a digest, then the two are compared for equality; if not, then the reference is updated with the computed digest, and the callback closure is invoked. Thus validation propagates upwards from the leaves to the root.

A caveat of this optimization is that it exposes the verifier to potential resource exhaustion attacks, as (potentially infinite) computation is performed on untrusted data before it is validated. A solution would be to bound the number of steps the program

should take before the buffer is cleared. Another caveat is that this optimization requires additional storage on the verifier in contrast to the original Merkle tree optimization.

5.3 Supporting full OCaml

Our compiler prototype supports authenticating the OCaml equivalent of the type language given in Figure 5 with the exception of function types. This support has been sufficient to program a variety of interesting data structures, as described in the next section.

To implement authenticated functions requires that we be able to perform the shallow projection of a lambda term. Our formalism does this by folding over the syntax of the lambda term’s body to find authenticated values $\langle h, v \rangle$ and replace them with values h . In an implementation this operation is tantamount to transforming a closure’s environment. We could do this quite naturally using Siskind and Pearlmutter’s map-closure operator [30], but unfortunately (as they point out) it is not clear how to implement this operator in a statically typed language, since the compiler cannot, in general, know the types of a given closure’s environment variables. We could imagine storing type information with the closure (e.g., Crary et al’s [3] term representations for types) in support of a generic, run-time shallow projection operation as per the formalism. In the meantime, the most natural use of authenticated closures we have found is to support shallow CPS transformations; we can use an explicit stack to the same effect, as shown in Figure 4.

Among other OCaml features not yet supported, the most desirable is authenticated polymorphic types. Similarly to closure environments, the compiler cannot know types needed to perform shallow projection—if a value given to *auth* and *unauth* is polymorphic, then its type is determined by how type variables are instantiated at a particular call-site. Once again, a generic shallow projection operator as per our formalism (and easily implemented in a dynamically typed language) would fit the bill. We could imagine requiring that *auth* and *unauth* each take an additional type parameter; in most cases it could be statically determined, but to support polymorphism it could be passed as an argument, e.g., to the polymorphic function containing the *auth/unauth* call.

6. Evaluation

To demonstrate the effectiveness and generality of our language and compiler, we have implemented a variety of ADSs. We analyze their performance and confirm it empirically with benchmarks for selected algorithms. Our benchmarks were conducted using an Amazon EC2 “m1.xlarge” instance (an Intel E5645 2.4GHz processor, with 16GB of RAM). All data structures were stored in RAM. Complete code is given in our extended technical report [20].

Merkle trees. Our version of Merkle trees was given in Figure 2. As Merkle trees are the most common authenticated data structure, and are readily implemented, we compared the running time of our compiled verifier routine to hand-written implementations in OCaml and in C (see Figure 13(c)). The benchmark consists of 10,000 random accesses to trees of height h for $h \in [5, 19]$. Each array element is a 1024 byte string; thus the largest tree contains approximately 250,000 elements and stores approximately 250 MB of data in total. Compared to the hand-optimized version in C, the program generated by our compiler is slower by only a factor of two; the hand-written OCaml code is about 30% slower than the C program. Profiling with *gprof* reveals that substantial overhead is due to the Marshal serialization routine used by our compiler; the hand-written OCaml code avoids this serialization by concatenating the child digests directly.

Red-black+ Trees. A red-black+ tree is a self-balancing binary tree—the + indicates that internal nodes only store keys, and the

values are stored only in the leaves. This data structure is appropriate for a (updatable) dictionary. We consider authenticated search trees to be the second oldest authenticated data structure, proposed by Naor and Nissim [24] to implement certificate revocation lists. Our results are asymptotically equivalent: the storage cost to the prover for the entire data structure is $O(n)$, while the computation cost per operation for both prover and verifier is $O(\log n)$; the size of the proof stream is also $O(\log n)$. Note that we also implement an authenticated version of normal binary search trees, too, and these also have the expected performance.

In Figure 13(a), we show the empirical runtime performance of the authenticated red-black+ tree for both the prover (P) and verifier (V) modes, as well the ideal mode with the \bullet annotations and *unauth/auth* commands erased, to show the overall computational overhead of authentication. The benchmark consists of 100,000 random insertions into a random tree containing 2^k elements, for each $k \in [4, 21]$. P runs approximately 25% faster than V ; this is because V computes a hash during both *unauth* and *auth* instructions whereas P only computes a hash during *auth*. According to profiler analysis (using *gprof*), V spends 55% of its time in the SHA1 hash routine and 30% in (de)serialization routines; P spends 28% of its time computing SHA1, 30% performing serialization and 22% in garbage collection. The overhead of P is approximately a factor of 100 compared to the ordinary data structure. We omit similar benchmarks for our other algorithms as the results are similar.

We measured the largest amount of memory allocated by OCaml during this benchmark as shown in Figure 13(b). This illustrates the key advantage of an authenticated data structure scheme: while P incurs space overhead by a factor of 3 vs the Ideal mode, the space requirement of the V is effectively constant.

Finally, we used this benchmark to evaluate the effectiveness of our two compiler optimizations on proof size (see Figure 13(d)). For a binary tree, the suspended-disbelief buffer results in a proof-size reduction of almost 50%, since only one of the left or right child digests must be transmitted for each node. The reuse cache is most effective when the tree is small and mostly fits in the cache; thus since the cache-size parameter in our benchmark is 1000, the proof is nearly empty up to trees of height 9. However, because the benchmark consists of random queries, the leaves and nodes toward the bottom are accessed almost uniformly at random, so only a constant number of nodes near the root are read from the cache each query. The cache would be more beneficial in applications where some nodes were accessed much more frequently. We only implemented the two optimizations separately; combining the two is left as future work.

Skip Lists. Skip lists [29] are randomized data structures providing similar performance (in expectation) to binary search trees.⁹ Our results are asymptotically equivalent to previous work on authenticated skip lists [12]: the storage cost to P for the entire data structure is $O(\log n)$, where n is the number of elements inserted; the expected computational cost to P and V as well as the size of the proof stream is $O(\log n)$, but $O(n)$ in the worst case.

Planar Separator Trees. Planar separator trees (PSTs) are data structures that can be used to efficiently query the distance of the shortest path between two vertices in a planar graph (e.g., many road maps) [6]. A separator of a graph is a collection of vertices inducing a binary partition on the graph, such that any path from a vertex in one partition to a vertex in the other partition must pass through a vertex in the separator. A consequence of the planar separator theorem is that every planar graph has a separator that is *small*

⁹Random algorithms can be used in $\lambda\bullet$ as long as P and V both use the same pseudorandom function and seed.

($O(\sqrt{n})$ where n is the number of vertices) yet induces a balanced partition (both partitions have at least $\frac{n}{3}$ elements); a search structure can be built over the graph by recursively constructing separators for each partition. While the naïve solution of storing the shortest distance between every pair of vertices requires $O(n^2)$ storage, the planar separator tree requires only $O(n^{\frac{3}{2}})$ space. The shortest distance between any two points can be computed in $O(\sqrt{n} \log n)$ time using this data structure. A potential application of an authenticated planar separator tree is for a mobile device user to query a service provider for travel directions, along with a proof that the response is actually the shortest path (rather than, e.g., one that routes the user out of their way past billboards for which the service provider might receive a profit). Using our authenticated compiler, the proof size and computation cost for P and V is also $O(\sqrt{n} \log n)$. We include this primarily as an example of a data structure that has not been treated as an ADS in prior work, but is straightforward to authenticate using our framework.

Bitcoin. Bitcoin [22] is a peer-to-peer network implementing a virtual currency; it features an authenticated data structure called the *blockchain*, which represents a history of transactions. A global ledger can be computed from the blockchain, which, abstractly speaking, contains a set of currently valid coins. A valid transaction removes a past coin from the ledger and adds a new one (assigned to the recipient of the coin).

The current Bitcoin data structure suffers from two drawbacks: 1) a miner (verifier) needs to maintain an entire copy of the global ledger to efficiently perform validation of transactions, which may become unscalable as the number of coins increases; and 2) at install time, a new client needs to download the entire transaction history and perform a linear computation (in the size of the block chain) to construct the ledger even when it trusts the root digest.

These drawbacks can be resolved if we build the ledger (as an authenticated set) into the block chain data structure. We have done this in our language by composing the blockchain data structure with a ledger consisting of our authenticated red-black+ tree, thereby reducing client storage to $O(\log M)$, where M bounds the number of coins in the ledger. Full details are given our supplemental technical report [20].

7. Related Work

As mentioned earlier, authenticated data structure research has had a flurry of results [1, 12, 15, 16, 19, 23] from the cryptography community, proposing authenticated versions of set (non)-membership, dictionaries, range queries, certain graph queries, B-trees, etc.

To the best of our knowledge, no one has offered a general authenticated data structure implementation for generic programs. Those closest work is by Martel *et al.* [16], which proposes a method for designing authenticated data structures for a class of data structures referred to as “search DAGs.” Their model is limited to *static* data structures, i.e., it does not support updates, which are supported by $\lambda\bullet$. Our approach is also easier to use, since the programmer writes largely standard (purely functional) implementations of data structures and their operations, and the compiler generates the prover/verifier code; in their approach, the task of designing and implementing ADS still must be done by hand.

Our programmatically generated ADS implementations have performance competitive with known customized ADS constructions based on collision-resistant hashes. We note, however, that while most are, not all known ADS constructions are based on collision resistant hash functions. For example, bilinear-group based ADS constructions exist for set operations [27]. Using alternative algebraic primitives other than collision resistant hashes can sometimes yield asymptotically better ADS constructions.




Figure 13. Empirical performance evaluation results. For 100,000 insertions into a red-black+ tree, (a) running time for Prover, Verifier, and Ideal; (b) memory usage (heap size). (c) Running time for fetch operation in a balanced merkle tree, for the program generated by our compiler, hand-written OCaml code, and hand-written C code; For 100,000 random insertions in a binary search tree, (d) the effect of two optional optimizations on proof stream size, compared with two optional optimizations: a reuse-cache (of size 1000) and the belief-suspension buffer.

Beyond the optimizations we have implemented (see Section 5.2), there are other known optimizations that we have not included. An example is the technique of “commutative hashing” due to Goodrich *et al.* [12] which reduces the proof size when it is irrelevant which (of possibly several) child nodes are traversed. We believe it is likely that optimizations for specific data structures could be incorporated generally into our compiler. Other optimizations lie further outside our model; an example is a line of work beginning with Merkle’s original paper [2, 18] in which the stored data is not arbitrary, but is instead generated algorithmically (e.g., using a pseudo-random number generator). In this case, the prover’s storage costs can be significantly reduced by recomputing data on the fly rather than storing it. In our performance evaluation we consider memory usage and running time; however some work in authenticated data structures [16, 21] considers I/O efficiency, another practical characteristic.

Verified computation [9] and succinct non-interactive arguments of knowledge (SNARKs) [10, 28] can also yield asymptotically better protocols for ensuring integrity of outsourced computation (e.g., with $O(1)$ amount of client computation other than reading the input and output). However, while theoretically attractive, known verified computation schemes and SNARKs are orders of magnitude more expensive than constructions based on collision-resistant hashes in practice [28]—partly due to the use of heavyweight cryptographic primitives such as fully homomorphic encryption.

ZQL [8] is a language for writing verified computations over hidden inputs; this is achieved by compiling programs to custom zero-knowledge protocols. Thus, unlike $\lambda\bullet$, ZQL provides correctness, security, and privacy. However, this comes at the price of a more limited language, e.g., ZQL does not support branching or higher-order functions generally. In addition, ZQL's absolute performance is much worse due to the use of more heavyweight cryptography.

8. Conclusions

We have presented $\lambda\bullet$, the first programming language for authenticated data structures (ADS). We have formally proven that every well-typed $\lambda\bullet$ program compiles to a secure protocol, and we have implemented $\lambda\bullet$ as a simple compiler extension to OCaml so that a programmer can easily derive an authenticated data structure from any ordinary one. The protocols generated by our compiler are competitive with the state-of-the-art in hand-rolled ADSs.

We believe this work is long past-due. ADS are a 30-year-old technique in cryptography, and yet researchers have overlooked the simple connection between ADSs encapsulated in our notion of authenticated types. We plan to extend our language to be more expressive, and to include more efficient techniques based on advanced cryptographic primitives. We hope our work encourages ADS adoption in future secure computing infrastructure.

Acknowledgments We thank Zooko and the Tahoe-LAFS team for several conversations that inspired our approach, and Nikhil Swamy and the anonymous reviewers for invaluable comments on drafts of this paper. This research was sponsored in part by NSF grant CNS-1111599, and by the US Army Research Laboratory and the UK Ministry of Defence under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defense, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

- [1] A. Anagnostopoulos, M. T. Goodrich, and R. Tamassia. Persistent authenticated dictionaries and their applications. In *Proc. ISC*, pages 379–393, London, UK, UK, 2001. Springer-Verlag.
- [2] P. Berman, M. Karpinski, and Y. Nekrich. Optimal trade-off for Merkle tree traversal. *Theor. Comput. Sci.*, 372(1):26–36, Mar. 2007.
- [3] K. Cray, S. Weirich, and G. Morrisett. Intensional polymorphism in type-erasure semantics. *Journal of Functional Programming*, 12(6), 2002.
- [4] U. Dal Lago and S. Martini. An invariant cost model for the lambda calculus. In *Logical Approaches to Computational Barriers*, pages 105–114. Springer, 2006.
- [5] P. Devanbu, M. Gertz, C. Martel, and S. G. Stubblebine. Authentic third-party data publication. In *Data and Application Security*, pages 101–112. Springer, 2002.
- [6] R. Duan. Planar separator theorem and its applications. lecture slides, Advanced Graph Algorithms, Summer 2012. max planck institut informatik. <http://www.mpi-inf.mpg.de/departments/d1/teaching/ss12/AdvancedGraphAlgorithms/Slides10.pdf>.
- [7] C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. In *Proc. PLDI*, 1993.
- [8] C. Fournet, M. Kohlweiss, G. Danezis, and Z. Luo. ZQL: A compiler for privacy-preserving data processing. In *USENIX Security*, 2013.
- [9] R. Gennaro, C. Gentry, and B. Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *CRYPTO 2010*, pages 465–482. Springer, 2010.
- [10] R. Gennaro, C. Gentry, B. Parno, and M. Raykova. Quadratic span programs and succinct NIZKs without PCPs. Cryptology ePrint Archive, Report 2012/215, 2012. <http://eprint.iacr.org/>.
- [11] M. T. Goodrich, C. Papamanthou, and R. Tamassia. On the cost of persistence and authentication in skip lists. In *Proc. Intl. Workshop on Experimental Algorithms*, volume 4525 of *LNCS*, pages 94–107. Springer, 2007.
- [12] M. T. Goodrich, R. Tamassia, and A. Schwerin. Implementation of an authenticated dictionary with skip lists and commutative hashing. In *Proc. DARPA Information Survivability Conference and Exposition II (DISCEX II)*, pages 68–82, 2001.
- [13] M. T. Goodrich, R. Tamassia, and N. Triandopoulos. Efficient authenticated data structures for graph connectivity and geometric search problems. *Algorithmica*, 60(3):505–552, 2011.
- [14] J. Katz and Y. Lindell. *Introduction to Modern Cryptography*. Chapman & Hall/CRC Press, 2007.
- [15] F. Li, K. Yi, M. Hadjieleftheriou, and G. Kollios. Proof-infused streams: Enabling authentication of sliding window queries on streams. In *VLDB*, pages 147–158, 2007.
- [16] C. Martel, G. Nuckolls, P. Devanbu, M. Gertz, A. Kwong, and S. Stubblebine. A general model for authentic data publication, 2001. Available from <http://www.cs.ucdavis.edu/~devanbu/files/model1-paper.pdf>.
- [17] C. Martel, G. Nuckolls, P. Devanbu, M. Gertz, A. Kwong, and S. G. Stubblebine. A General Model for Authenticated Data Structures. *Algorithmica*, 39(1):21–41, Jan. 2004.
- [18] R. C. Merkle. Secure communications over insecure channels. *Communications of the ACM*, 21(4):294–299, Apr. 1978.
- [19] R. C. Merkle. A certified digital signature. In G. Brassard, editor, *Proc. CRYPTO '89*, volume 435 of *LNCS*, pages 218–238. Springer-Verlag, 1989.
- [20] A. Miller, M. Hicks, J. Katz, and E. Shi. Full version: Authenticated data structures, generically. <http://www.cs.umd.edu/~amiller/gpads-full.pdf>, Jan. 2014.
- [21] E. Mykletun, M. Narasimha, and G. Tsudik. Authentication and integrity in outsourced databases. *Trans. Storage*, 2(2):107–138, 2006.
- [22] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Technical report, unpublished, 2009.
- [23] M. Naor and K. Nissim. Certificate revocation and certificate update. In *Proc. USENIX*, pages 217–228, Berkeley, 1998.
- [24] M. Naor and K. Nissim. Certificate revocation and certificate update. *IEEE J. on Sel. Areas Commun.*, 18(4):561–570, 2000.
- [25] C. Papamanthou and R. Tamassia. Time and space efficient algorithms for two-party authenticated data structures. In *Information and Communications Security*, pages 1–15. Springer, 2007.
- [26] C. Papamanthou, R. Tamassia, and N. Triandopoulos. Authenticated hash tables. In *Proc. ACM Conference on Computer and Communications Security (CCS)*, pages 437–448. ACM, October 2008.
- [27] C. Papamanthou, R. Tamassia, and N. Triandopoulos. Optimal verification of operations on dynamic sets. In *CRYPTO*, pages 91–110, 2011.
- [28] B. Parno, C. Gentry, J. Howell, and M. Raykova. Pinocchio: Nearly practical verifiable computation. In *Proc. IEEE SSP*, 2013.
- [29] W. Pugh. Skip lists: a probabilistic alternative to balanced trees. *Communications of the ACM*, 33(6):668–676, 1990.
- [30] J. M. Siskind and B. A. Pearlmutter. First-class nonstandard interpretations by opening closures. In *POPL*, 2007.
- [31] R. Tamassia. Authenticated data structures. In *11th Annual European Symposium on Algorithms*, Sept. 2003.