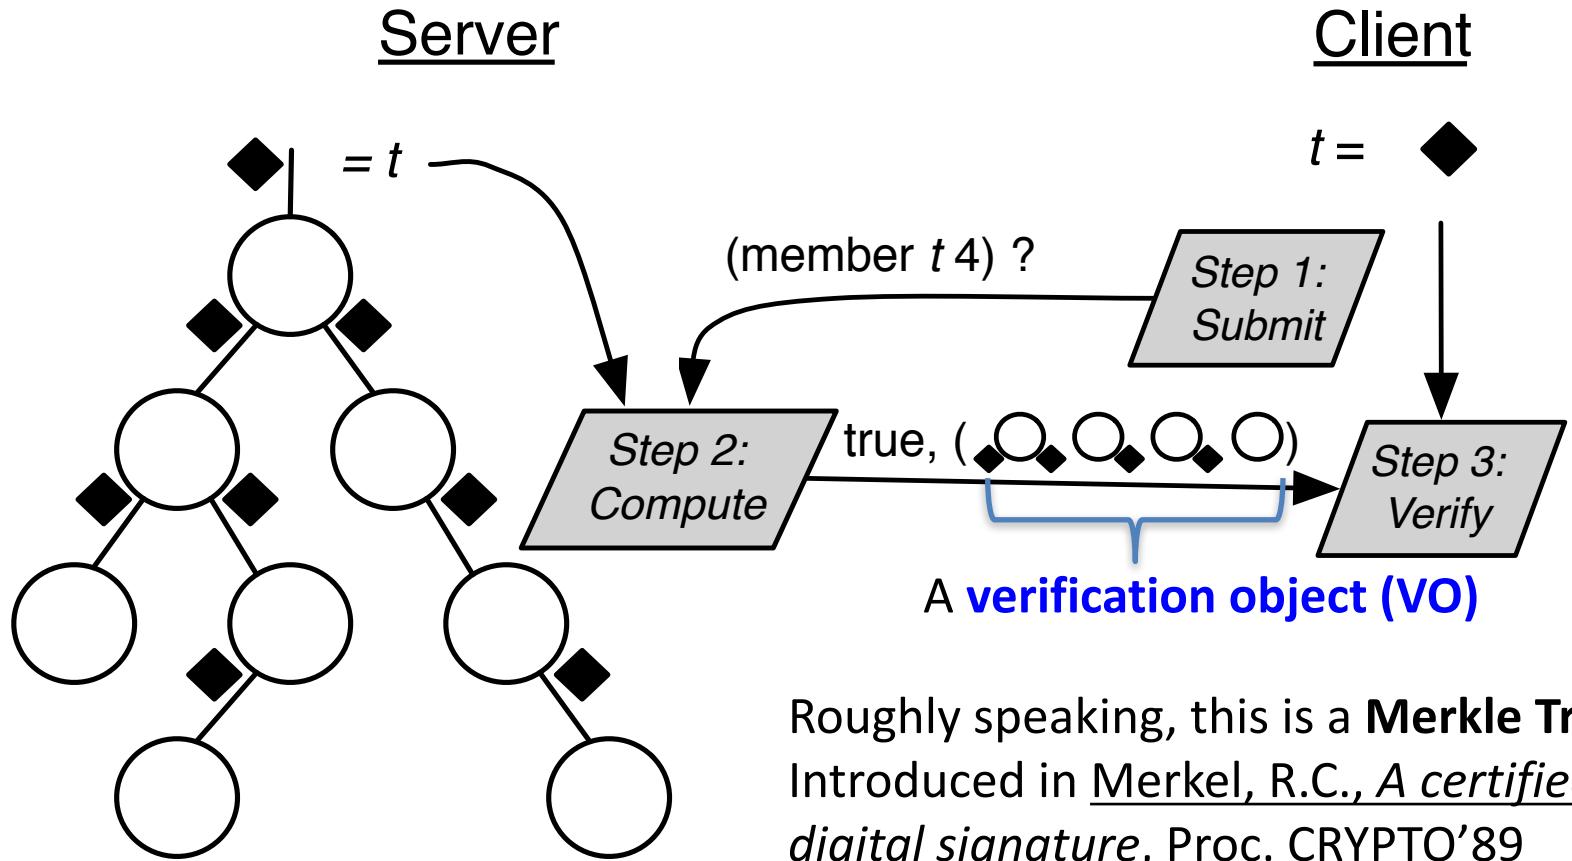


Authenticated Data Structures, Generically

Michael Hicks

with Andrew Miller, Elaine Shi, and Jonathan Katz
The University of Maryland, College Park

Authenticated Data Structures



Applications/benefits of ADSs

- Trustworthy mirroring/duplication
 - Duplicate data on many untrusted servers, but **ensure trustworthy results** (ensures integrity, not privacy)
 - **Low space requirement for client**
- Practical example: **Bitcoin block chain**
 - But implementation is suboptimal due to difficulty of combining the algorithm with the crypto
- Other applications possible: GPG key servers, Tor relay directories, Tahoe-LAFS mutable files, ...

Generic method for building ADSs?

- State of the art: design different ADSs in an ad hoc (and heroic) fashion
 - Litany of papers on improvements to existing data structures, or variations
- Instead: Can we add something to a programming language to make it easy to build new ADSs?
 - (Yes!)

Approach

- Start with a pure functional programming language with standard features
 - E.g., Ocaml with datatypes, (recursive) functions, base types, etc. but no refs
- Add new type $\bullet\tau$ and two coercions
 - *auth*: $\tau \longrightarrow \bullet\tau$
 - *unauth*: $\bullet\tau \longrightarrow \tau$
- Develop two **computation modes**, one for the server, one for the client
 - Compiler produces code for each mode

Example: BST with auth types

type tree =

| Tip
| Bin of tree × int × tree

let rec member (t:tree) (x:int) : bool =

match unauth t **with**

Tip → false

| Bin (l,y,r) →

if y = x **then** true

else if x < y **then** member l x

else member r x

let rec insert (t:tree) (x:int) : tree =

match unauth t **with**

Tip → auth (Bin(auth Tip,x, auth Tip))

| Bin (l,y,r) →

if y = x **then** t

else if x < y **then** auth (Bin(insert l x,y,r))

else auth (Bin(l,y,insert r x))

Execution modes

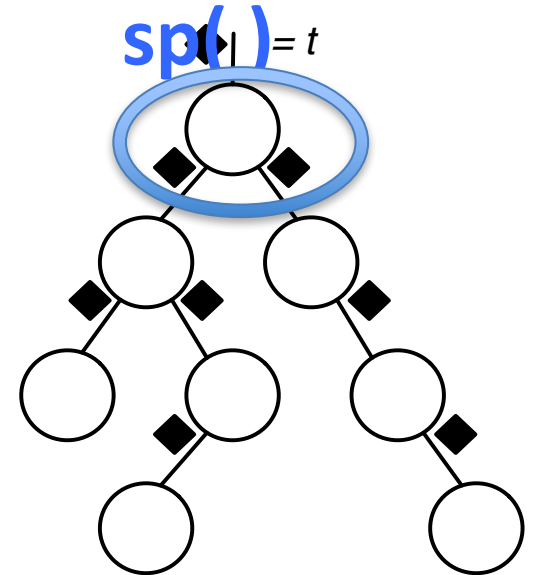
Server (mode **P**rover) produces VO.


A value of type $\bullet\tau$ is

- a value v of type τ , coupled with
- a cryptographic hash d of its **shallow projection** $sp(v)$
 - Informally: serialize the data up to, but not past, nested authenticated values, and hash that

Client (mode **V**erifier) consumes VO.

A value of type $\bullet\tau$ is a hash d



$t =$ 

Proof and Verification

- The VO is a list of shallow projections of authenticated values
- Server (mode P): `unauth <d,v>` enqueues `sp(v)` on the VO
 - Returns `v`
- Client (mode V): `unauth d` checks that `hash(hd(VO)) = d`
 - Dequeues and returns the head of the VO

Formalization

- Extension to standard programming language
 - auth and unauth, simple typing rules
- Small-step operational semantics
 - Three variants, indexed by modes I, P, V
 - Differ on auth, unauth, and whether VO consumed/produced
- **Theorem:** Starting with a well-typed program
 - The client/server together compute the right answer under normal operation
 - The client will only accept the wrong answer if the server can manufacture a hash collision (very hard!)

Security Proofs

- Proofs of key theorems use standard, syntactic methods common in programming languages
 - Like those we saw in earlier lectures
 - Taught in first graduate class on semantics
- Methods designed to speak about whole classes of computation, not single programs
 - Proof done once, for the language, not per ADS

Syntax and Types

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 $v ::= () \mid x \mid \lambda x.e \mid \mathbf{rec} x.\lambda y.e$
 $\mid \mathbf{inj}_1 v \mid \mathbf{inj}_2 v \mid (v_1, v_2) \mid \mathbf{roll} v$

Exprs $e ::= v \mid \mathbf{let} x = e_1 \mathbf{in} e_2 \mid v_1 v_2 \mid \mathbf{case} v v_0 v_1$
 $\mid \mathbf{prj}_1 v \mid \mathbf{prj}_2 v \mid \mathbf{unroll} v \mid \mathbf{auth} v \mid \mathbf{unauth} v$


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

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

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

...

Operational Semantics in 3 Modes



$$\begin{aligned} \ll \pi, \mathbf{case} (\mathbf{inj}_1 v)(\lambda x.e_1)(\lambda x.e_2) \gg &\rightarrow_m \ll \pi, e_1[v \setminus x] \gg \\ \ll \pi, \mathbf{case} (\mathbf{inj}_2 v)(\lambda x.e_1)(\lambda x.e_2) \gg &\rightarrow_m \ll \pi, e_2[v \setminus x] \gg \\ \ll \pi, \mathbf{prj}_1 (v_1, v_2) \gg &\rightarrow_m \ll \pi, v_1 \gg \\ \ll \pi, \mathbf{prj}_2 (v_1, v_2) \gg &\rightarrow_m \ll \pi, v_2 \gg \end{aligned}$$

- m is either I, P, or V
- carry around the VO π
- most transitions leave it unchanged

$$\begin{aligned} \ll \pi, \mathbf{auth} v \gg &\rightarrow_I \ll \pi, v \gg \\ \ll \pi, \mathbf{unauth} v \gg &\rightarrow_I \ll \pi, v \gg \end{aligned}$$

auth/unauth are no-ops in ideal mode

Operational Semantics in 3 Modes

$$\begin{aligned}
 \langle\langle () \rangle\rangle &= () \\
 \langle\langle \langle h, v \rangle \rangle\rangle &= h \\
 \langle\langle \mathit{auth} \ v \rangle\rangle &= \mathit{auth} \ \langle\langle v \rangle\rangle \\
 \langle\langle \langle v_1, v_2 \rangle \rangle\rangle &= (\langle\langle v_1 \rangle\rangle, \langle\langle v_2 \rangle\rangle) \\
 \langle\langle \mathbf{roll} \ v \rangle\rangle &= \mathbf{roll} \ \langle\langle v \rangle\rangle \\
 \langle\langle \mathbf{rec} \ x. \lambda y. e \rangle\rangle &= \mathbf{rec} \ x. (\lambda y. e)
 \end{aligned}$$

Shallow projection function, written $(|e|)$

...

$$\begin{aligned}
 \langle\langle \pi, \mathit{auth} \ v \rangle\rangle &\rightarrow_P \langle\langle \pi, \langle \mathbf{hash} \ \langle\langle v \rangle\rangle, v \rangle \rangle \\
 \langle\langle \pi, \mathit{auth} \ v \rangle\rangle &\rightarrow_V \langle\langle \pi, \mathbf{hash} \ v \rangle \rangle
 \end{aligned}$$

auth in Prover/Verifier builds new digest

$$\langle\langle \pi, \mathit{unauth} \ \langle h, v \rangle \rangle \rangle \rightarrow_P \langle\langle \pi @ [\langle\langle v \rangle\rangle], v \rangle \rangle$$

unauth in Prover adds to the stream

$$\frac{\mathbf{hash} \ s_0 = h}{\langle\langle [s_0] @ \pi, \mathit{unauth} \ h \rangle \rangle \rightarrow_V \langle\langle \pi, s_0 \rangle \rangle}$$

unauth in Verifier consumes from the stream

Implementation

- We have extended the Ocaml compiler to support authenticated types
 - produces versions of data structure for the *prover* (server) and *verifier* (client)
- Implemented several ADSs
 - BST, Red-black trees, skip lists, planar separator DS for shortest paths, efficient blockchain
- Confirmed expected space/time performance

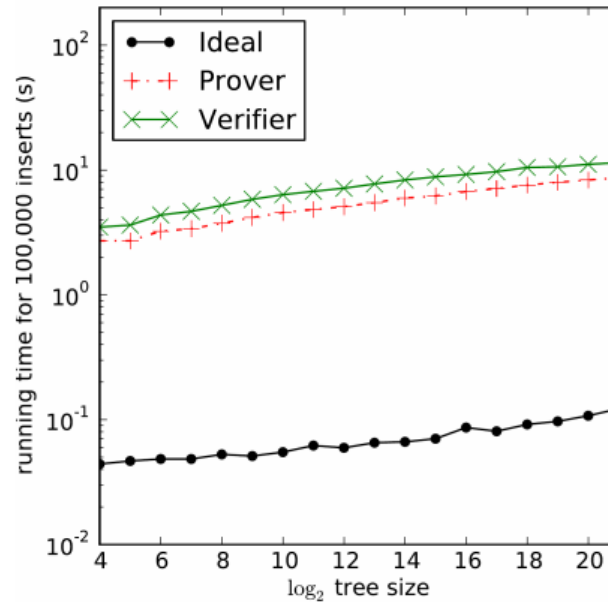
Red+black Tree Performance

Verifier:


- 55% in SHA1
 - used in both *auth* and *unauth*
- 30% in Marshal (used in hash(v))

Prover:

- 28% in SHA1
 - used in *auth* only
- 30% in Marshal
- 22% in GC




(a) Running time



(b) Memory usage

Merkle Tree Performance

Compared against
hand-rolled
C and OCaml



Summary

- Generic implementation of hash-based ADS
- Write program once in ordinary functional language, automatically derive Client/Server modes
- Once-and-for-all Security Theorem applies to every program
- Performance comparable to hand-optimized