

Logical Relations for Formally Verified

Authenticated Data Structures

Simon Oddershede Gregersen


joint work with Chaitanya Agarwal and Joseph Tassarotti



I have so much work to do!




I have so much work to do!



I have so much work to do!



I can help!




I have so much work to do!

Can I trust you to not mess it up?



I can help!




I have so much work to do!

Can I trust you to not mess it up?



I can help!

Of course!



How can Alice securely outsource work to Bob?

How can Alice securely outsource work to Bob?

The operations of an **authenticated data structure** can be carried out by Bob, but (efficiently) verified by, e.g., Alice!

This is done by having Bob produce a **compact proof** that Alice can check.


How can Alice securely outsource work to Bob?

The operations of an **authenticated data structure** can be carried out by Bob, but (efficiently) verified by, e.g., Alice!

This is done by having Bob produce a **compact proof** that Alice can check.


ADSs allow outsourcing data storage and processing tasks to untrusted servers without loss of integrity.

Example: Merkle Tree




where h_i denotes the hash of t_i/s_i


Example: Merkle Tree (Prover)




lookup([R, L], t_0) =




Example: Merkle Tree (Prover)




lookup([R, L], t_0) =




Example: Merkle Tree (Prover)




lookup([R, L], t_0) =



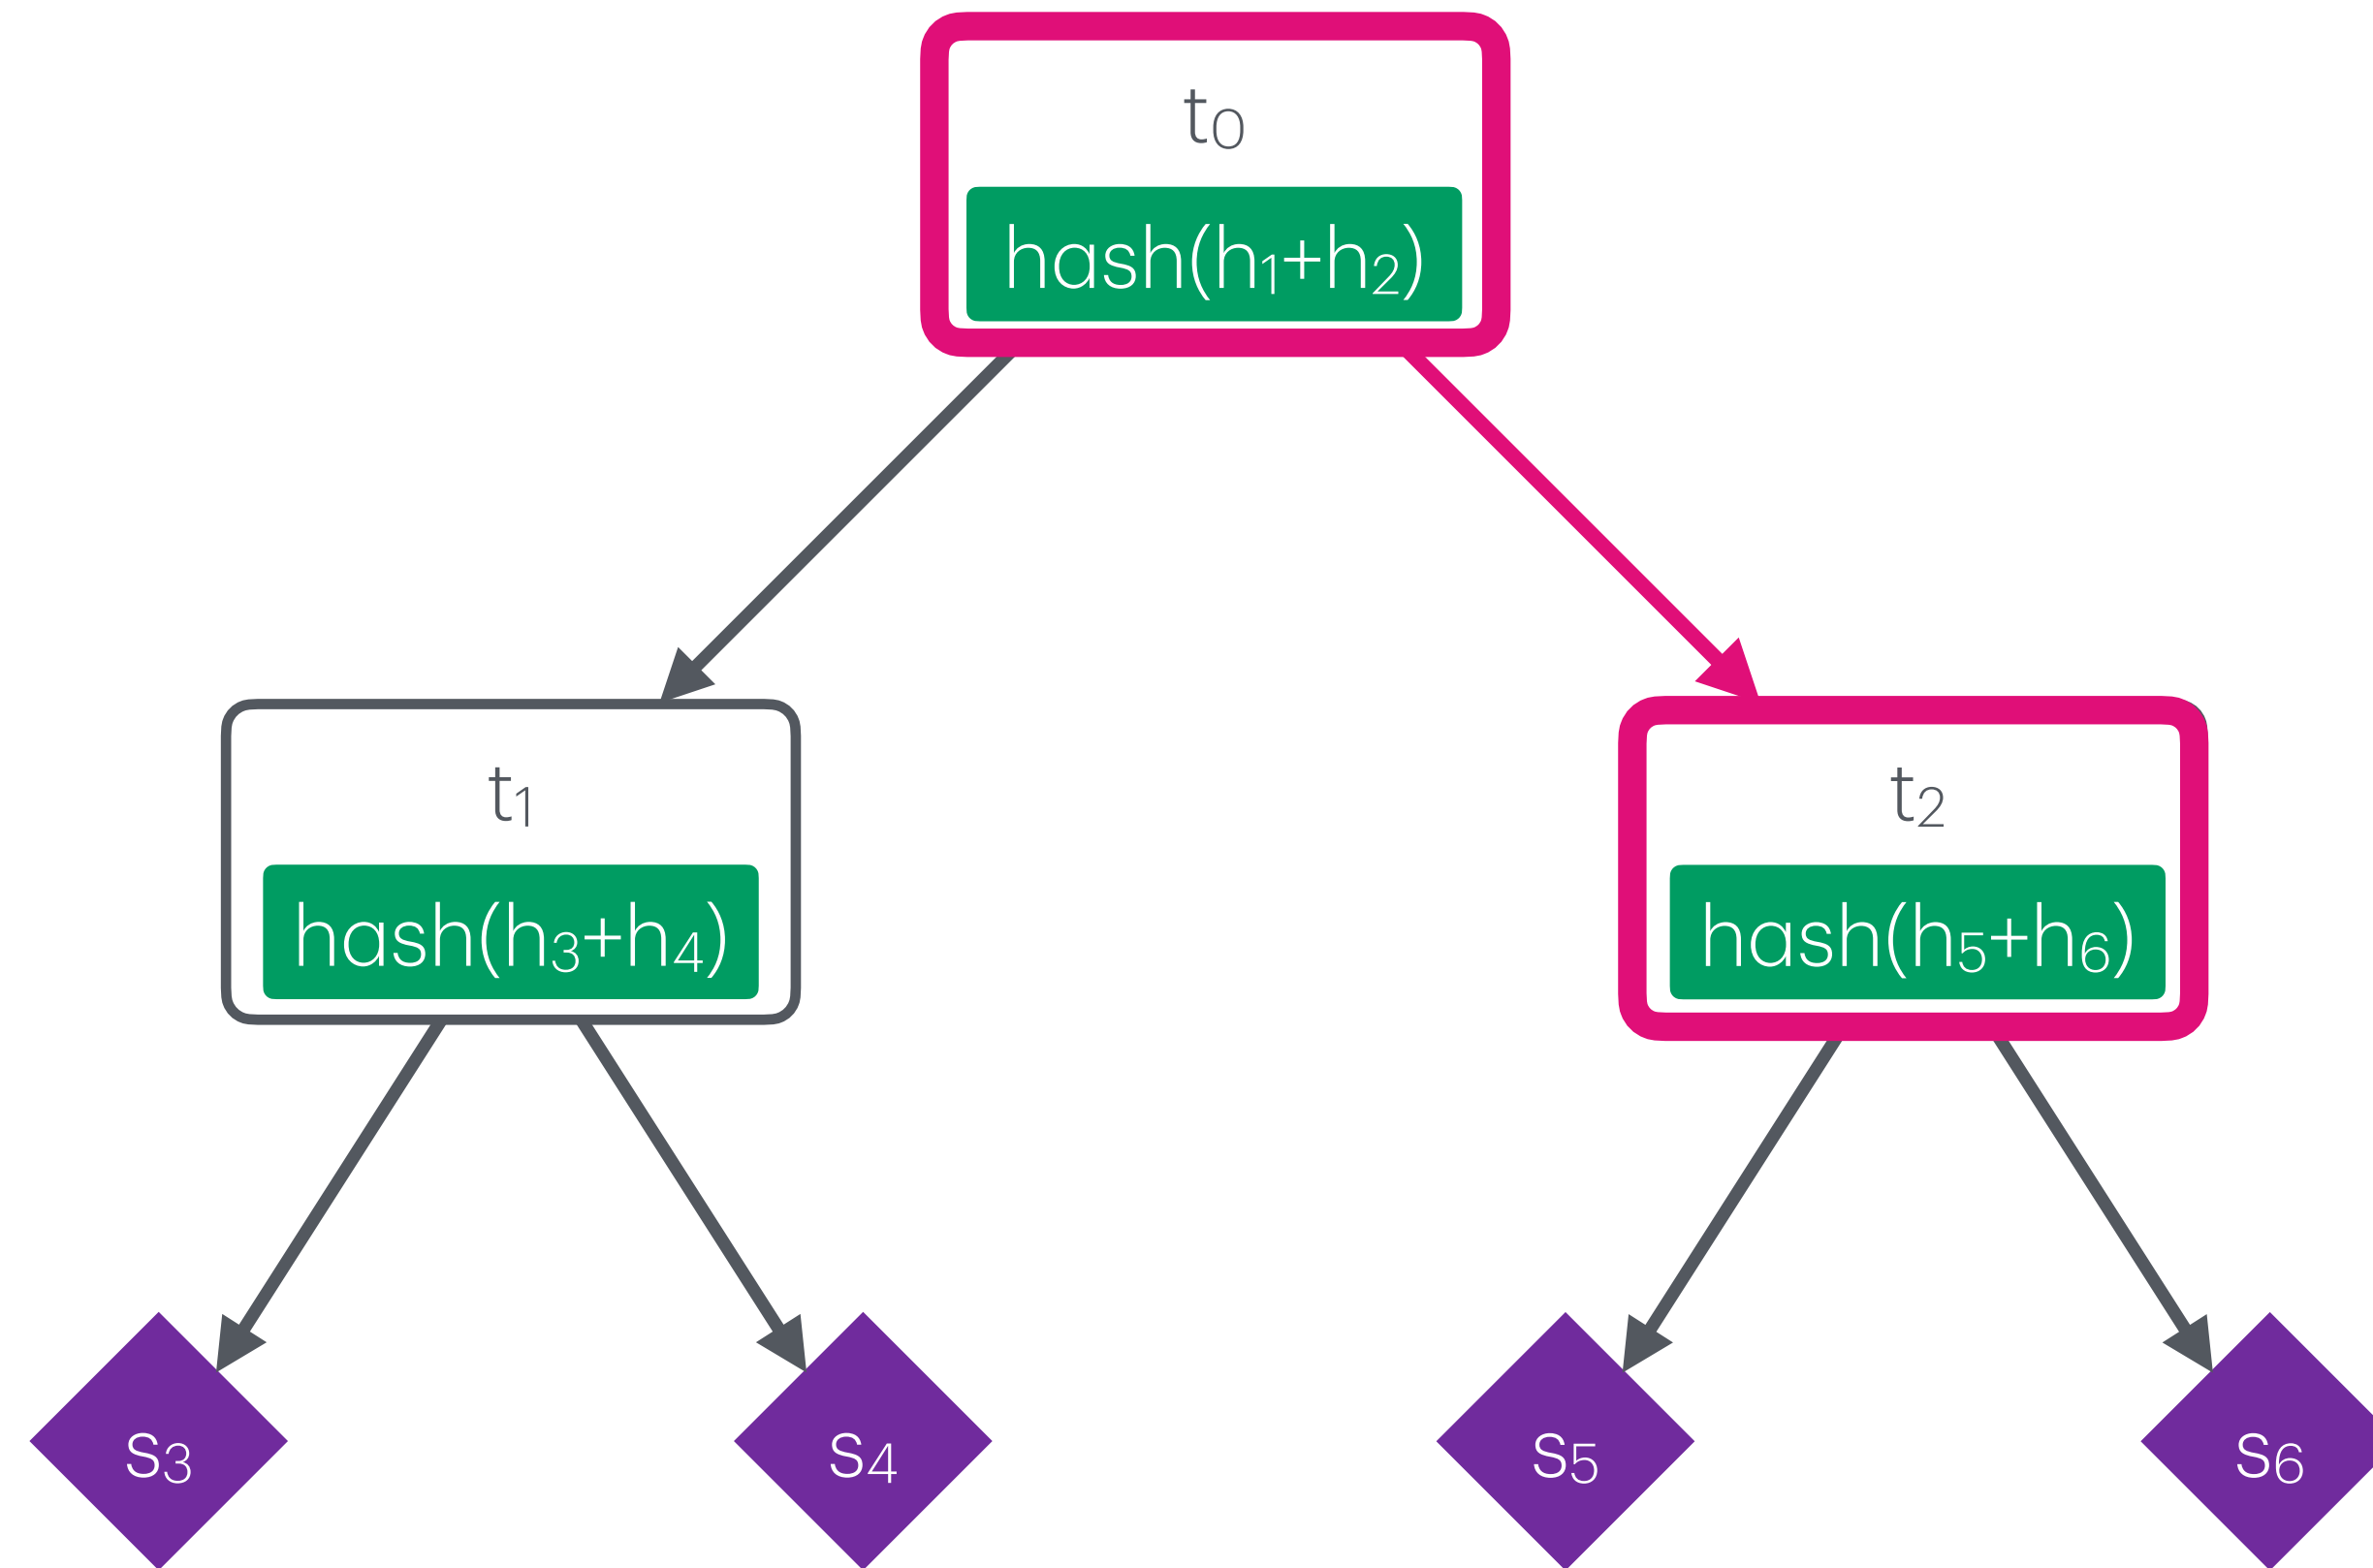
Example: Merkle Tree (Prover)




lookup([R, L], t_0) =
([h_1




Example: Merkle Tree (Prover)




lookup([R, L], t_0) =
([h_1




Example: Merkle Tree (Prover)




lookup([R, L], t_0) =
([h_1




Example: Merkle Tree (Prover)




lookup([R, L], t_0) =
([h_1 , h_6




Example: Merkle Tree (Prover)




lookup([R, L], t_0) =
([h_1 , h_6




Example: Merkle Tree (Prover)



lookup([R, L], t_0) =
([h_1 , h_6 , s_5], s_5)



Example: Merkle Tree (Verifier)




$\text{lookup}([R, L], t_0) =$
 $([h_1, h_6, s_5], s_5)$



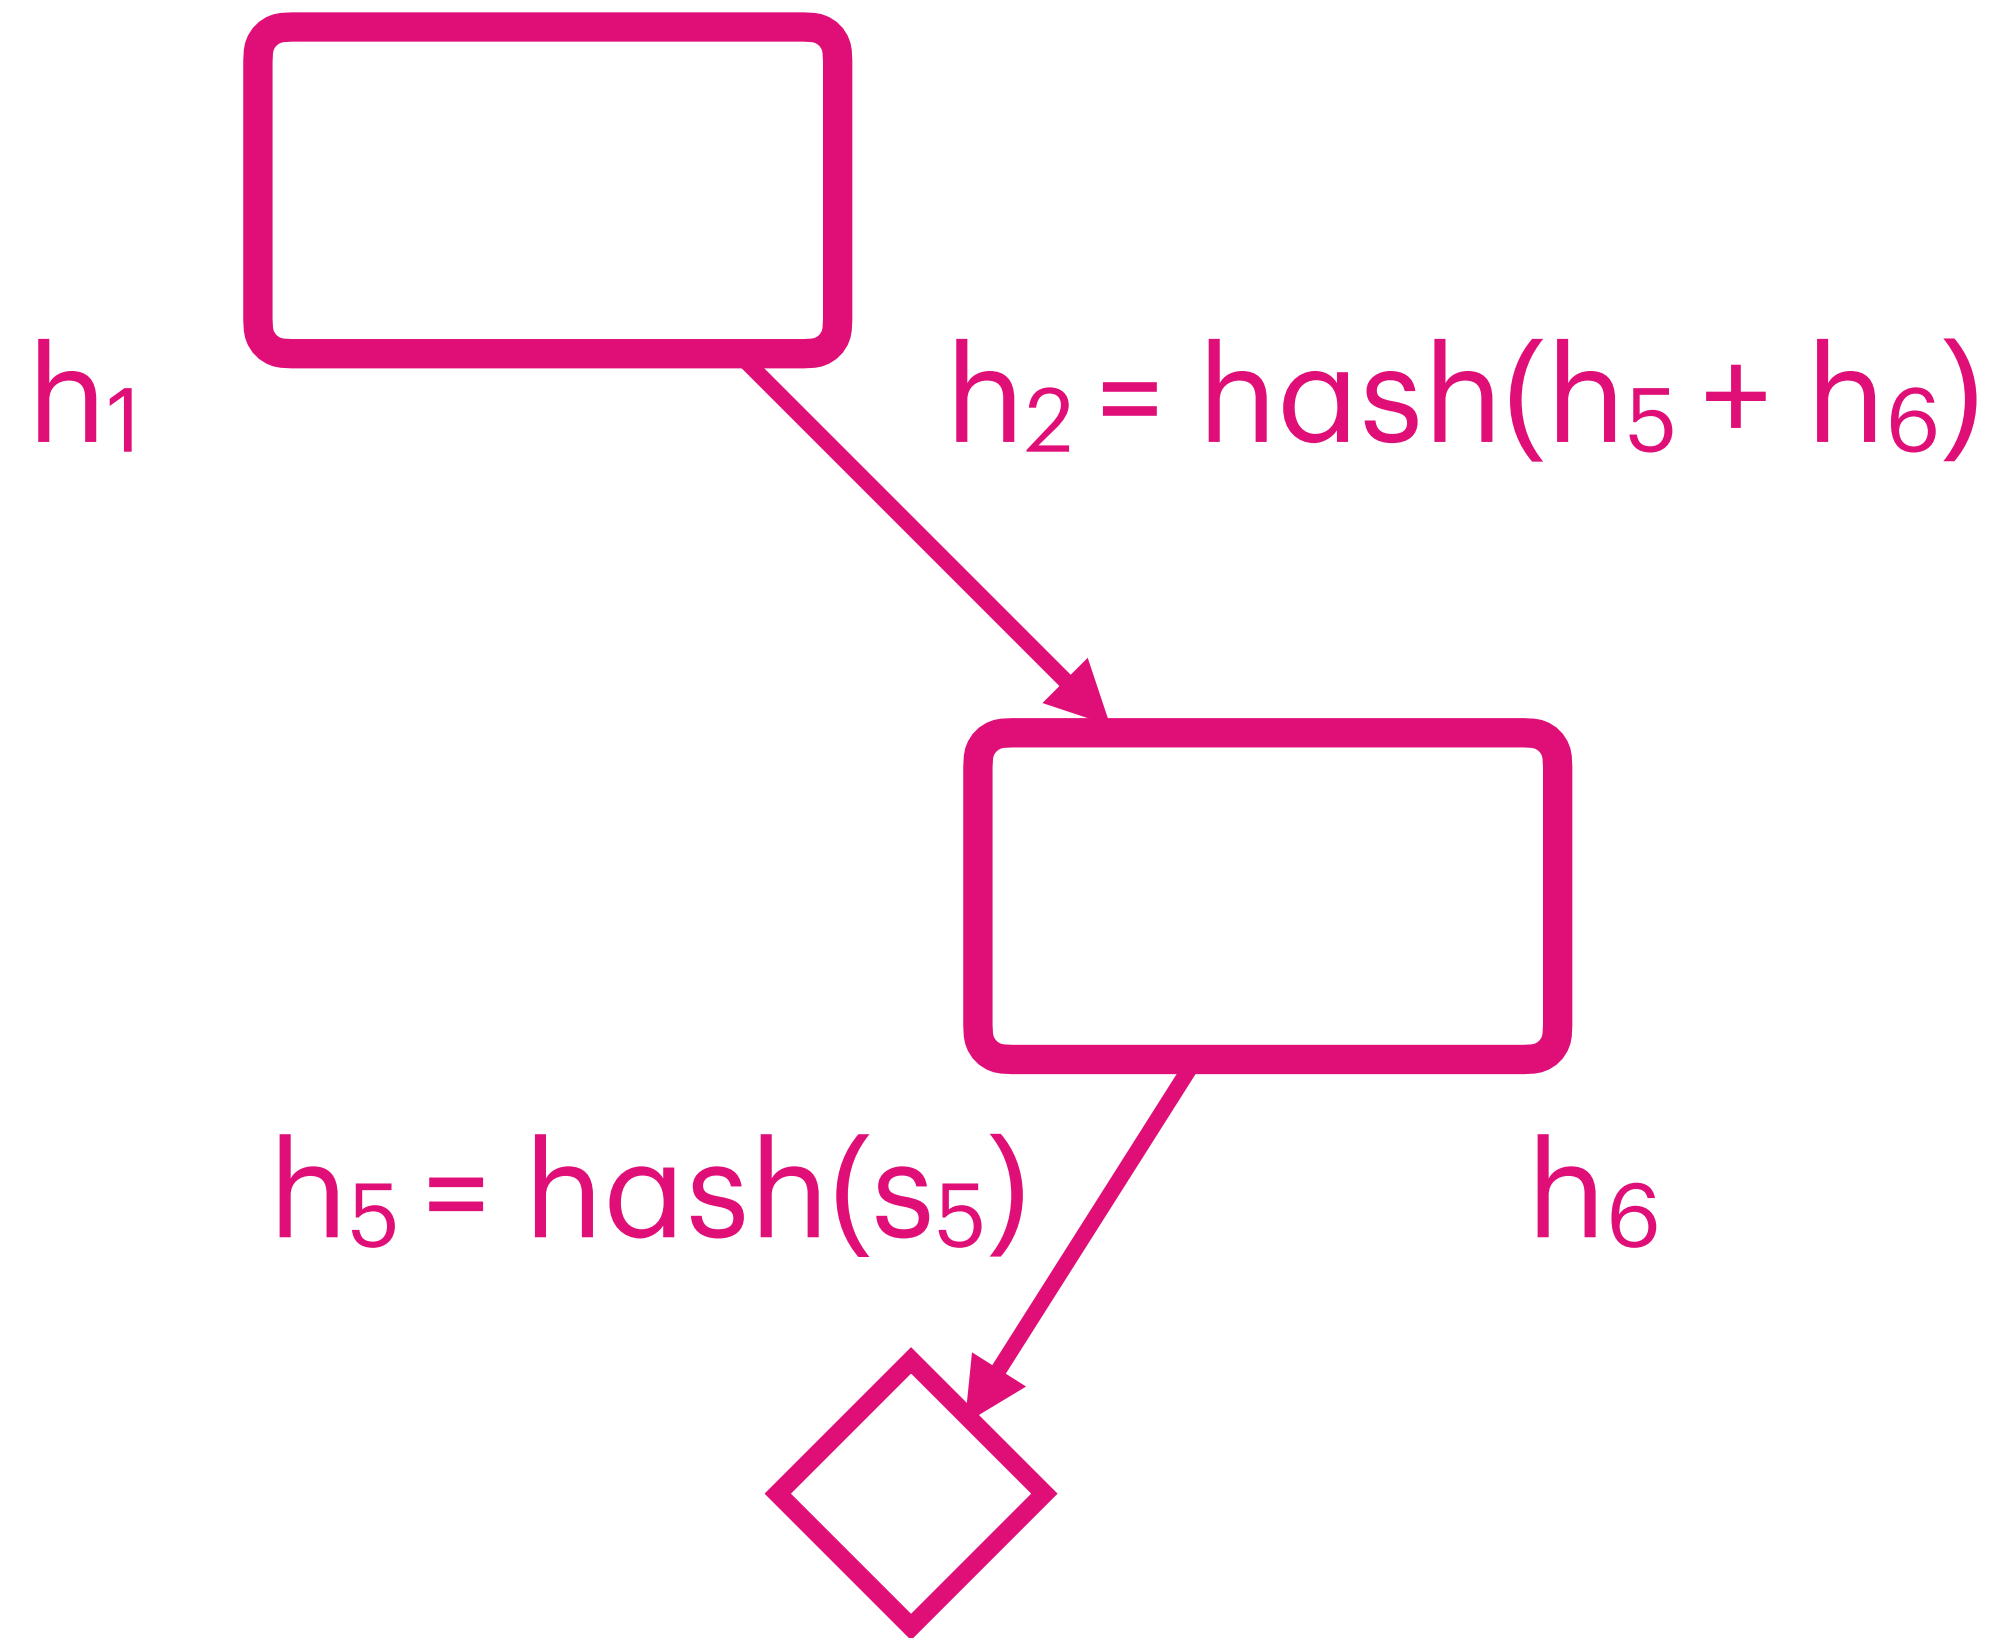
Example: Merkle Tree (Verifier)

lookup([R, L], t₀) =
([h₁, h₆, s₅], s₅)



Example: Merkle Tree (Verifier)

$$h_0' = \text{hash}(h_1 + h_2)$$

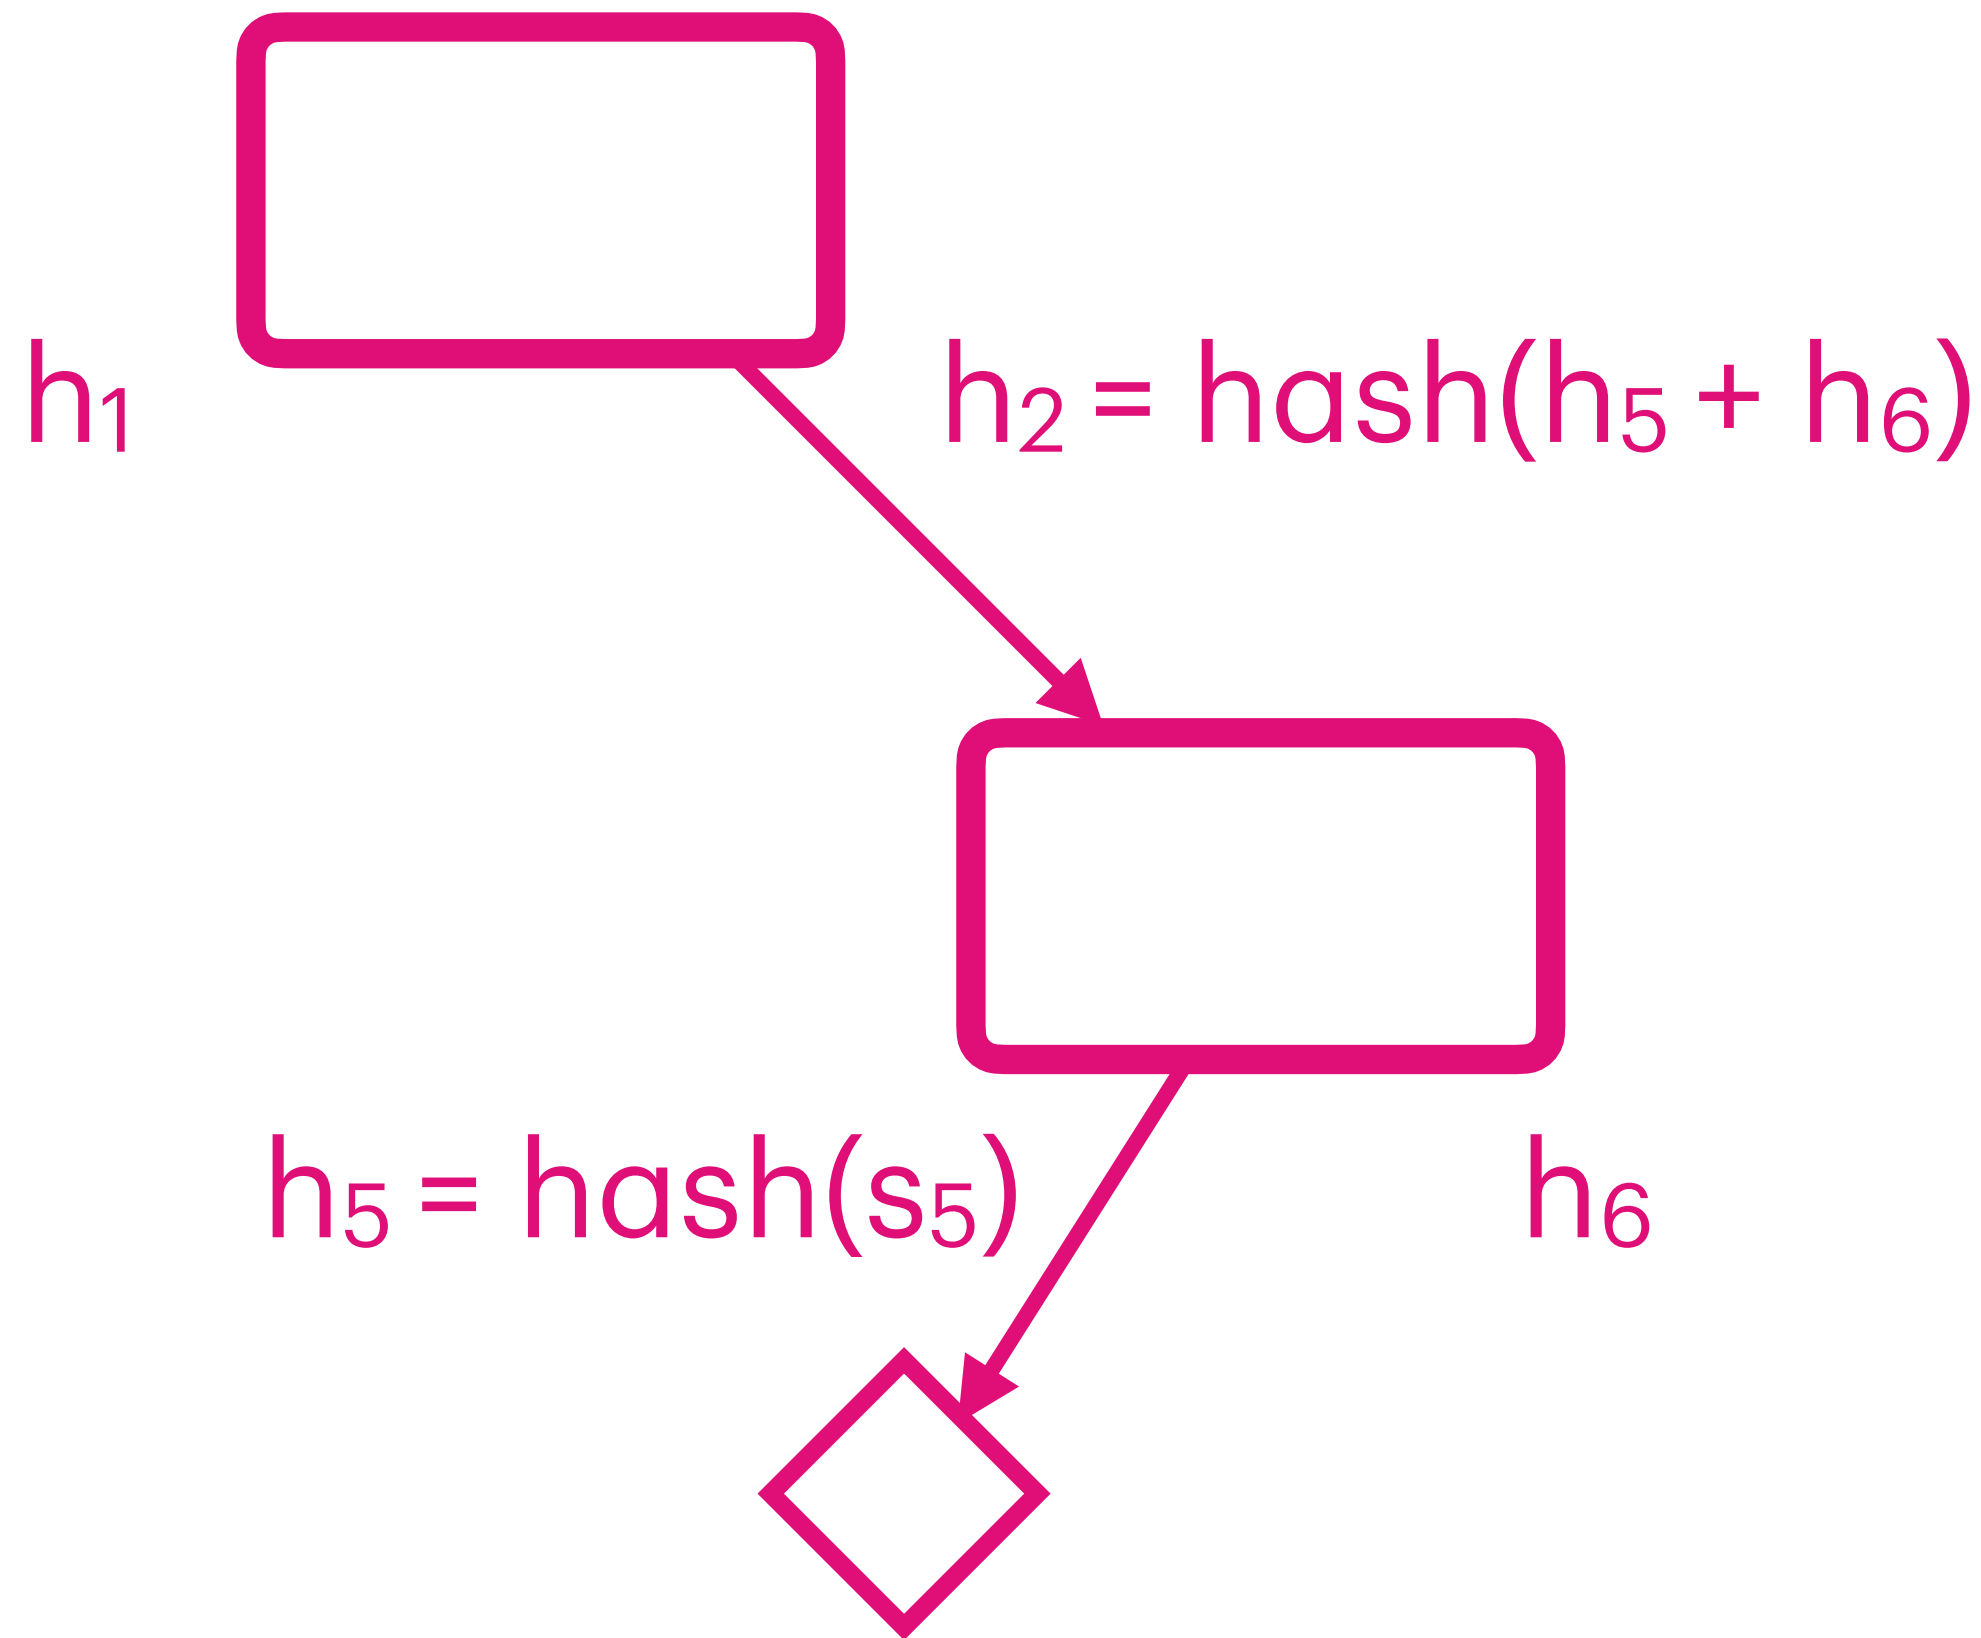


$$\text{lookup}([R, L], t_0) = ([h_1, h_6, s_5], s_5)$$





Example: Merkle Tree (Verifier)

$$h_0' = \text{hash}(h_1 + h_2)$$



$$\text{lookup}([R, L], t_0) = ([h_1, h_6, s_5], s_5)$$



Use cases

- **Certificate transparency**

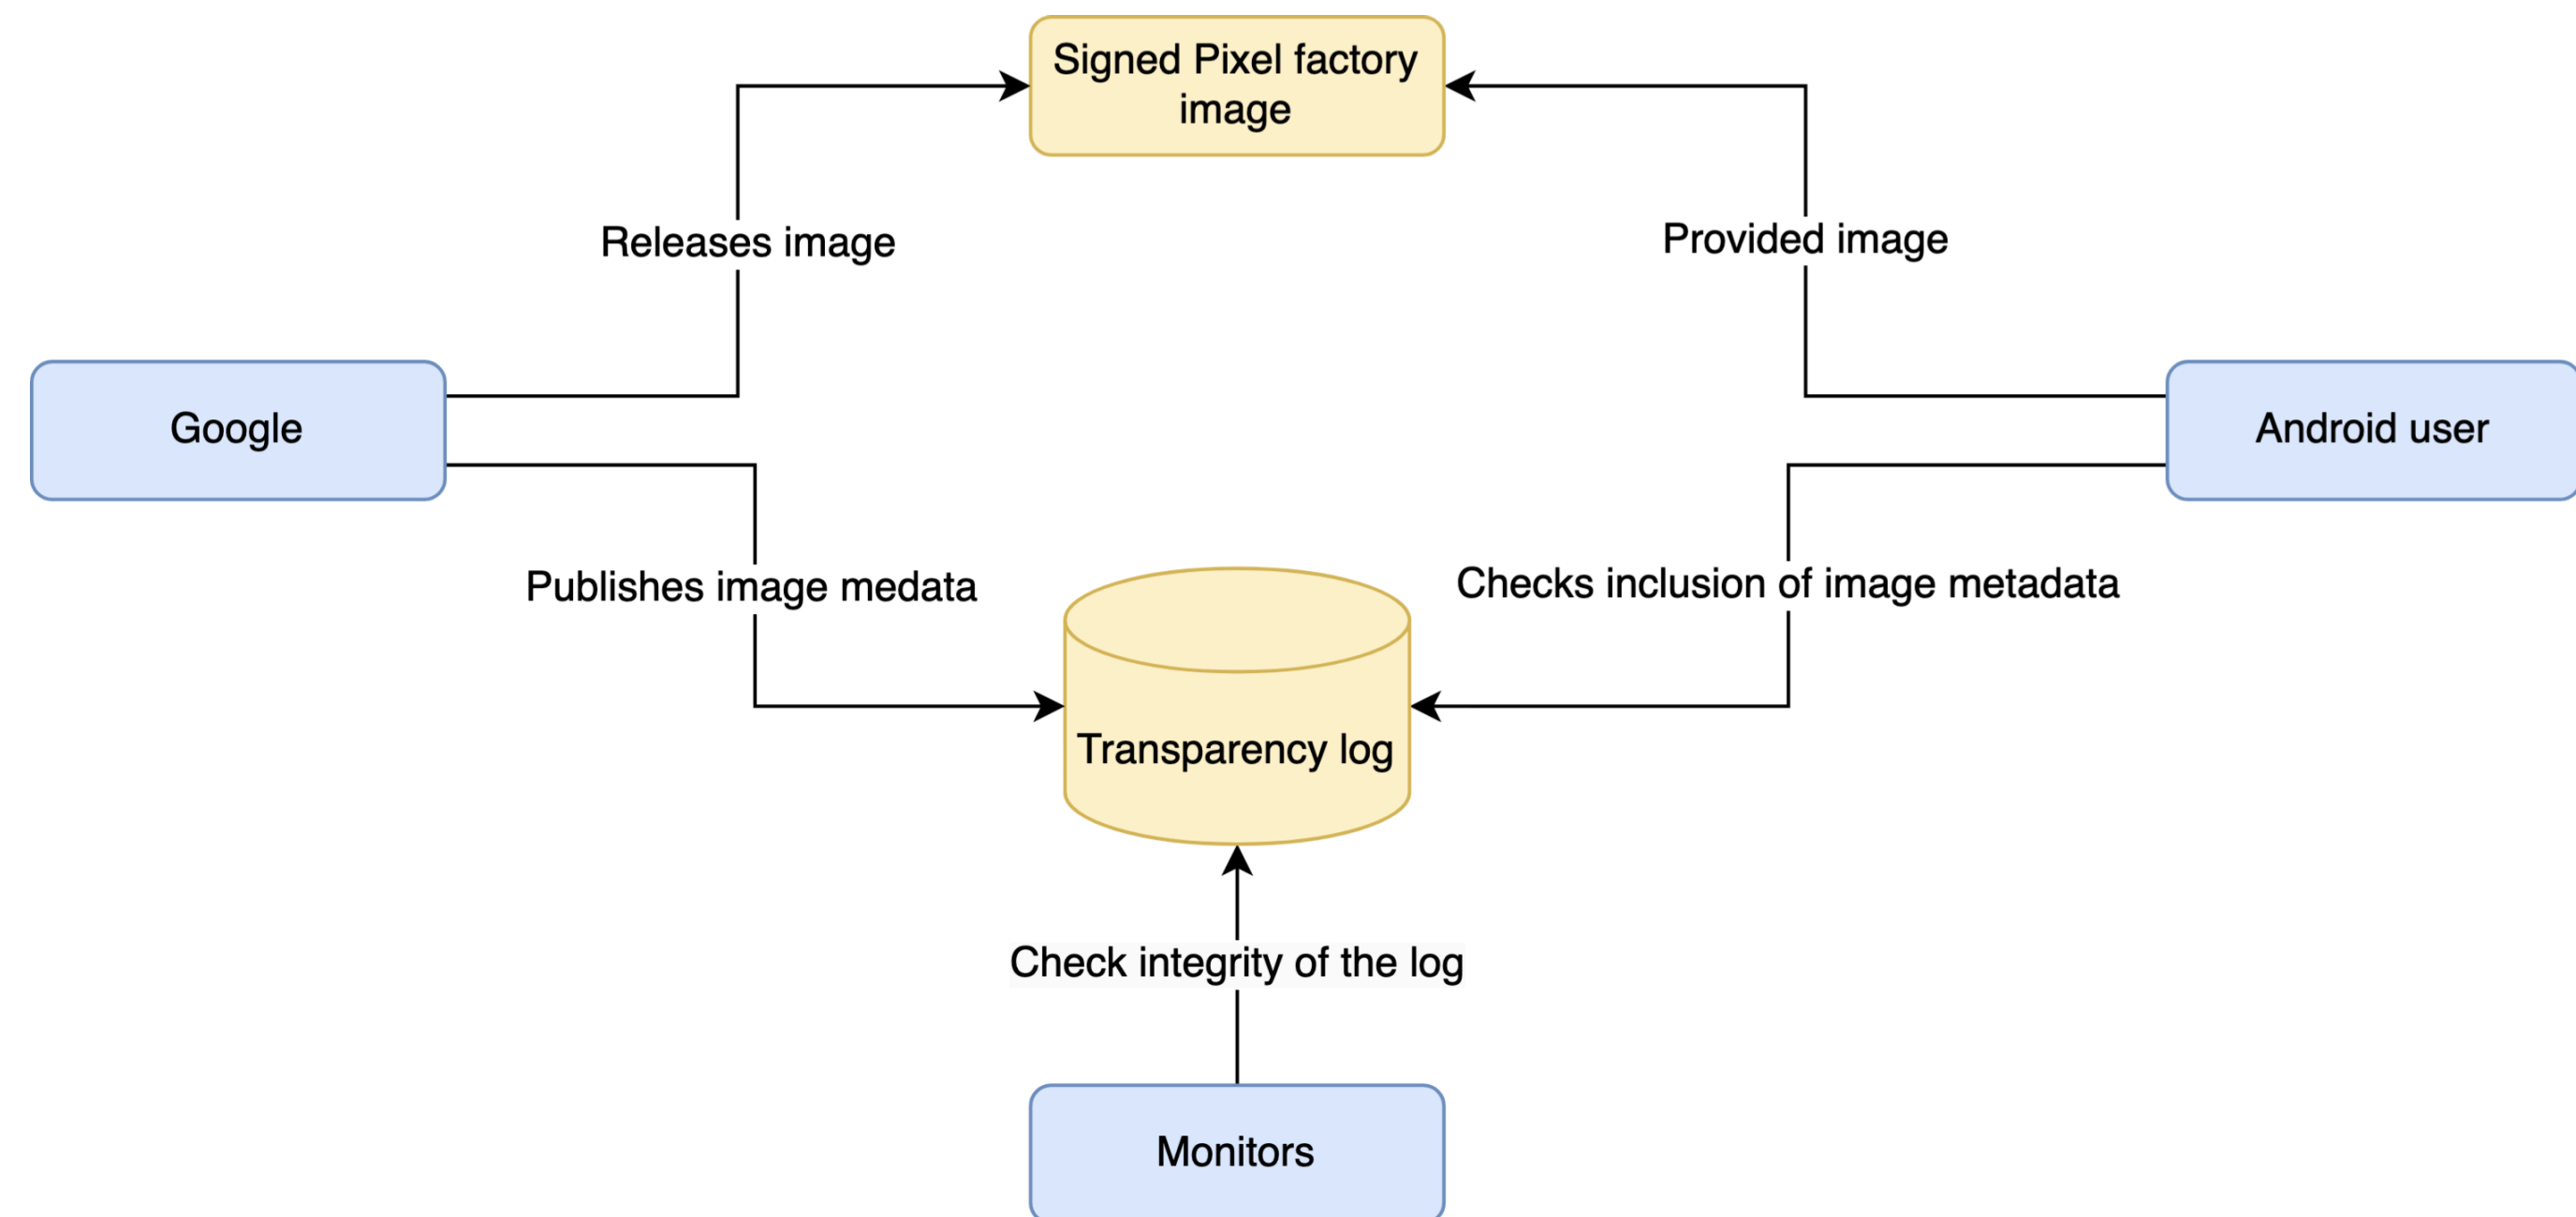
Google Chrome (2015), Cloudflare (2018), Let's Encrypt (2019), Firefox (2025)

- **Key transparency**


WhatsApp (2023), Signal (???)

- **Binary transparency**

Pixel Binaries, Go modules



Miller et al. realized that the prover and verifier can be **compiled** from a single implementation.



Miller et. al's approach

OCaml is extended with three new primitives:

- authenticated types • τ
- $\text{auth} : \forall \alpha . \alpha \rightarrow \bullet \alpha$
- $\text{unauth} : \forall \alpha . \bullet \alpha \rightarrow \alpha$



Miller et. al's approach

OCaml is extended with three new primitives:

- authenticated types • τ
- $\text{auth} : \forall \alpha . \alpha \rightarrow \bullet \alpha$
- $\text{unauth} : \forall \alpha . \bullet \alpha \rightarrow \alpha$

```
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
```



To justify the correctness of their approach, they define a core calculus and show **security** and **correctness**:

To justify the correctness of their approach, they define a core calculus and show **security** and **correctness**:

Security: If the **verifier** accepts a proof p and returns v then

- the **ideal** execution returns v or
- a hash collision occurred.

To justify the correctness of their approach, they define a core calculus and show **security** and **correctness**:

Security: If the **verifier** accepts a proof p and returns v then

- the **ideal** execution returns v or
- a hash collision occurred.

Correctness: If the **prover** generates a proof p and a result v then

- the **ideal** execution returns v and
- the **verifier** accepts p and returns v as well.

Limitations

1. Maintaining a custom compiler frontend imposes development burden.
2. To construct compact proofs, the compiler implements several optimizations that are not covered by the security and correctness theorems.
3. Even with optimizations, the generated data structures are not always producing proofs as compact as hand-written implementations.

[About](#)[Blog](#)[Publications](#)

BOB ATKEY

Authenticated Data Structures, as a Library, for Free!

Let's assume that you're querying to some database stored in the cloud (i.e., on someone else's computer). Being of a sceptical mind, you worry whether or not the answers you get back are from the database you expect. Or is the cloud lying to you?

Published: Tuesday 12th April
2016

Authenticated Data Structures (ADSs) are a proposed solution to this problem. When the server sends back its answers, it also sends back a "proof" that the answer came from the database it claims. You, the client, verify this proof. If the proof doesn't verify, then you've got evidence that the server was lying. If the

[About](#)[Blog](#)[Publications](#)

BOB ATKEY

Authenticated Data Structures, as a Library, for Free!

Let's assume that you're querying to some database
stored in the cloud (i.e., on someone else's computer).

Published: Tuesday 12th April
2016

```
module type MERKLE = functor (A : AUTHENTIKIT) -> sig
  open A

  (* ... *)

  val fetch : path -> tree auth -> string option auth_computation
end
```

This work

- Two **logical-relations models** and a proof of security and correctness of the typed module construction in a general-purpose programming language.
- We address the remaining two limitations:
 - We verify several **optimizations** (as supported by the compiler).
 - We show how to prove that manually verified code can be **safely linked** with automatically generated code.
- Full mechanization in the Rocq theorem prover.

```
module type AUTHENTIKIT = sig
  type 'a auth

  (* ... *)

  (* ... *)

  val auth      : 'a Serializable.evidence -> 'a -> 'a auth
  val unauth    : 'a Serializable.evidence -> 'a auth -> 'a auth_computation
end
```

```
module type AUTHENTIKIT = sig
  type 'a auth
  type 'a auth_computation

  val return : 'a -> 'a auth_computation
  val bind   : 'a auth_computation -> ('a -> 'b auth_computation) -> 'b auth_computation

  (* ... *)

  val auth      : 'a Serializable.evidence -> 'a -> 'a auth
  val unauth    : 'a Serializable.evidence -> 'a auth -> 'a auth_computation
end
```

```

module type AUTHENTIKIT = sig
  type 'a auth
  type 'a auth_computation

  val return : 'a -> 'a auth_computation
  val bind   : 'a auth_computation -> ('a -> 'b auth_computation) -> 'b auth_computation

module Serializable : sig
  type 'a evidence
  val auth      : 'a auth evidence
  val pair      : 'a evidence -> 'b evidence -> ('a * 'b) evidence
  val sum       : 'a evidence -> 'b evidence -> ['left of 'a | `right of 'b] evidence
  val string    : string evidence
  val int       : int evidence
end

  val auth      : 'a Serializable.evidence -> 'a -> 'a auth
  val unauth    : 'a Serializable.evidence -> 'a auth -> 'a auth_computation
end

```

```
module Merkle : MERKLE = functor (A : AUTHENTIKIT) -> struct
  open A

  type path = [`L | `R] list
  type tree = [`leaf of string | `node of tree auth * tree auth]

  (* ... *)

  (* ... *)

end
```

```

module Merkle : MERKLE = functor (A : AUTHENTIKIT) -> struct
  open A

  type path = [`L | `R] list
  type tree = [`leaf of string | `node of tree auth * tree auth]

  let tree_evi : tree Serializable.evi = Serializable.(sum string (pair auth auth))

  let make_leaf (s : string) : tree auth = auth tree_evi (`leaf s)
  let make_branch (l r : tree auth) : tree auth = auth tree_evi (`node (l, r))

  (* ... *)

end

```



```

module Merkle : MERKLE = functor (A : AUTHENTIKIT) -> struct
  open A


  type path = [`L | `R] list
  type tree = [`leaf of string | `node of tree auth * tree auth]

  let tree_evi : tree Serializable.evi = Serializable.(sum string (pair auth auth))

  let make_leaf (s : string) : tree auth = auth tree_evi (`leaf s)
  let make_branch (l r : tree auth) : tree auth = auth tree_evi (`node (l, r))

  let rec fetch (p : path) (t : tree auth) : string option auth_computation =
    bind (unauth tree_evi t) (fun t ->
      match p, t with
      | [], `leaf s -> return (Some s)
      | `L :: p, `node (l, _) -> fetch p l
      | `R :: p, `node (_, r) -> fetch p r
      | _, _ -> return None)
end

```




```
type proof = string list

module Prover : AUTHENTIKIT =
  type 'a auth = 'a * string
  type 'a auth_computation = () -> proof * 'a

  (* ... *)

  (* ... *)

end
```



```
type proof = string list

module Prover : AUTHENTIKIT =
  type 'a auth = 'a * string
  type 'a auth_computation = () -> proof * 'a


  let return a () = ([], a)
  let bind c f =
    let (prf, a) = c () in
    let (prf', b) = f a () in
    (prf @ prf', b)

  module Serializable = struct
    type 'a evidence = 'a -> string

    (* ... *)
  end

  (* ... *)

end
```



```
type proof = string list

module Prover : AUTHENTIKIT =
  type 'a auth = 'a * string
  type 'a auth_computation = () -> proof * 'a

  let return a () = ([], a)
  let bind c f =
    let (prf, a) = c () in
    let (prf', b) = f a () in
    (prf @ prf', b)

  module Serializable = struct
    type 'a evidence = 'a -> string

    (* ... *)
  end


  let auth evi a = (a, hash (evi a))
  let unauth evi (a, _) () = ([evi a], a)
end
```

```
module Verifier : AUTHENTIKit =  
  type 'a auth = string  
  type 'a auth_computation =  
    proof -> [`Ok of proof * 'a | `ProofFailure]
```

```
(* ... *)
```

```
(* ... *)
```

```
end
```



```

module Verifier : AUTHENTIKIT =
  type 'a auth = string
  type 'a auth_computation =
    proof -> [`Ok of proof * 'a | `ProofFailure]

  let return a prf = `Ok (prf, a)
  let bind c f prf =
    match c prf with
    | `ProofFailure -> `ProofFailure
    | `Ok (prf', a) -> f a prf'


  module Serializable = struct
    type 'a evidence =
      { serialize : 'a -> string; deserialize : string -> 'a option }

    (* ... *)
  end

  (* ... *)

end

```



```

module Verifier : AUTHENTIKIT =
  type 'a auth = string
  type 'a auth_computation =
    proof -> [`Ok of proof * 'a | `ProofFailure]


  let return a prf = `Ok (prf, a)
  let bind c f prf =
    match c prf with
    | `ProofFailure -> `ProofFailure
    | `Ok (prf', a) -> f a prf'


  module Serializable = struct
    type 'a evidence =
      { serialize : 'a -> string; deserialize : string -> 'a option }

    (* ... *)
  end

  let auth evi a = hash (evi.serialize a)
  let unauth evi h prf =
    match prf with
    | p :: ps when hash p = h ->
      match evi.deserialize p with
      | None -> 'ProofFailure
      | Some a -> `Ok (ps, a)
    | _ -> 'ProofFailure
  end

```





```
module Ideal : AUTHENTIKIT = struct
  type 'a auth = 'a
  type 'a auth_computation = () -> 'a

  let return a () = a
  let bind a f () = f (a ()) ()

  (* ... *)

  let auth _ a = a
  let unauth _ a () = a
end
```


Takeaway

- In the end, it is not so difficult to prove that **one particular client** has the security and correctness property.
- The challenge is to prove that **any well-typed client** has these properties!
- Authentikit relies on a **parametricity** property of OCaml's module system.

Plan

1. Define a **type system** that can capture the module-based construction.
2. Define a **semantic model** that captures the type system.
3. Show that the inhabitants of the semantic model have the property of interest.
4. Show that the three Authentikit implementations inhabit the model.

Requirements

```
module type AUTHENTIKIT = sig
  type 'a auth
  type 'a auth_computation

  val return : 'a -> 'a auth_computation
  val bind   : 'a auth_computation -> ('a -> 'b auth_computation) -> 'b auth_computation

  module Serializable : sig
    type 'a evidence
    val auth      : 'a auth evidence
    val pair      : 'a evidence -> 'b evidence -> ('a * 'b) evidence
    val sum       : 'a evidence -> 'b evidence -> ['left of 'a | `right of 'b] evidence
    val string    : string evidence
    val int       : int evidence
  end

  val auth      : 'a Serializable.evidence -> 'a -> 'a auth
  val unauth    : 'a Serializable.evidence -> 'a auth -> 'a auth_computation
end
```

Requirements

```
module type AUTHENTIKIT = sig
  type 'a auth
  type 'a auth_computation

  val return : 'a -> 'a auth_computation
  val bind   : 'a auth_computation -> ('a -> 'b auth_computation) -> 'b auth_computation

  module Serializable : sig
    type 'a evidence
    val auth   : 'a auth evidence
    val pair   : 'a evidence -> 'b evidence -> ('a * 'b) evidence
    val sum    : 'a evidence -> 'b evidence -> ['left of 'a | `right of 'b] evidence
    val string : string evidence
    val int    : int evidence
  end

  val auth   : 'a Serializable.evidence -> 'a -> 'a auth
  val unauth : 'a Serializable.evidence -> 'a auth -> 'a auth_computation
end
```

(higher-order) functions



Requirements

```
module type AUTHENTIKIT = sig
  type 'a auth
  type 'a auth_computation

  val return : 'a -> 'a auth_computation
  val bind   : 'a auth_computation -> ('a -> 'b auth_computation) -> 'b auth_computation

  module Serializable : sig
    type 'a evidence
    val auth      : 'a auth evidence
    val pair      : 'a evidence -> 'b evidence -> ('a * 'b) evidence
    val sum       : 'a evidence -> 'b evidence -> ['left of 'a | `right of 'b] evidence
    val string    : string evidence
    val int       : int evidence
  end

  val auth      : 'a Serializable.evidence -> 'a -> 'a auth
  val unauth    : 'a Serializable.evidence -> 'a auth -> 'a auth_computation
end
```

polymorphism

(higher-order) functions

Requirements

```
module type AUTHENTIKIT = sig
  type 'a auth
  type 'a auth_computation

  val return : 'a -> 'a auth_computation
  val bind   : 'a auth_computation -> ('a -> 'b auth_computation) -> 'b auth_computation

module Serializable : sig
  type 'a evidence
  val auth   : 'a auth evidence
  val pair   : 'a evidence -> 'b evidence -> ('a * 'b) evidence
  val sum    : 'a evidence -> 'b evidence -> ['left of 'a | `right of 'b] evidence
  val string : string evidence
  val int    : int evidence
end

val auth   : 'a Serializable.evidence -> 'a -> 'a auth
val unauth : 'a Serializable.evidence -> 'a auth -> 'a auth_computation
end
```

abstract types



polymorphism



(higher-order) functions



Requirements

```
module Merkle : MERKLE = functor (A : AUTHENTIKIT) -> struct
  open A

  type path = [`L | `R] list
  type tree = [`leaf of string | `node of tree auth * tree auth]

  (* ... *)
end
```

```
module type AUTHENTIKIT = sig
  type 'a auth
  type 'a auth_computation

  val return : 'a -> 'a auth_computation
  val bind   : 'a auth_computation -> ('a -> 'b auth_computation) -> 'b auth_computation

  module Serializable : sig
    type 'a evidence
    val auth   : 'a auth evidence
    val pair   : 'a evidence -> 'b evidence -> ('a * 'b) evidence
    val sum    : 'a evidence -> 'b evidence -> [`left of 'a | `right of 'b] evidence
    val string : string evidence
    val int    : int evidence
  end

  val auth   : 'a Serializable.evidence -> 'a -> 'a auth
  val unauth : 'a Serializable.evidence -> 'a auth -> 'a auth_computation
end
```

abstract types

recursive types

polymorphism

(higher-order) functions

Requirements

(abstract) type constructors

abstract types

recursive types

```
module type AUTHENTIKIT = sig
  type 'a auth
  type 'a auth_computation

  val return : 'a -> 'a auth_computation
  val bind   : 'a auth_computation -> ('a -> 'b auth_computation) -> 'b auth_computation

  module Serializable : sig
    type 'a evidence
    val auth   : 'a auth evidence
    val pair   : 'a evidence -> 'b evidence -> ('a * 'b) evidence
    val sum    : 'a evidence -> 'b evidence -> ['left of 'a | `right of 'b] evidence
    val string : string evidence
    val int    : int evidence
  end

  val auth   : 'a Serializable.evidence -> 'a -> 'a auth
  val unauth : 'a Serializable.evidence -> 'a auth -> 'a auth_computation
end
```

```
module Merkle : MERKLE = functor (A : AUTHENTIKIT) -> struct
  open A

  type path = [`L | `R] list
  type tree = [`leaf of string | `node of tree auth * tree auth]

  (* ... *)
end
```

polymorphism

(higher-order) functions

Reminder

STLC: terms can depend on terms,

$$\frac{\Gamma, x : \sigma \vdash e : \tau}{\Gamma \vdash \lambda x. e : \sigma \rightarrow \tau}$$


System F: terms can depend on types,

$$\frac{\Theta, \alpha \mid \Gamma \vdash e : \tau}{\Theta \mid \Gamma \vdash \Lambda \alpha. e : \forall \alpha. \tau}$$

System F_ω: types can depend on types,

$$\frac{\Theta \vdash \tau \equiv \sigma \quad \Theta \mid \Gamma \vdash e : \sigma}{\Theta \mid \Gamma \vdash e : \tau}$$

$$\frac{}{\Theta \vdash (\lambda \alpha. \tau)\sigma \equiv \tau[\sigma/\alpha]}$$



The $F_{\omega, \mu}^{\text{ref}}$ language

$\kappa ::= \star \mid \kappa \Rightarrow \kappa$

(kinds)

$\tau ::= \alpha \mid \lambda \alpha : \kappa . \tau \mid \tau \tau \mid c$

(types)

$c ::= \dots \mid \times \mid + \mid \rightarrow \mid \text{ref} \mid \forall_{\kappa} \mid \exists_{\kappa} \mid \mu_{\kappa}$

(constructors)

The $F_{\omega, \mu}^{\text{ref}}$ language

$\kappa ::= \star \mid \kappa \Rightarrow \kappa$

(kinds)

$\tau ::= \alpha \mid \lambda \alpha : \kappa . \tau \mid \tau \tau \mid c$

(types)

$c ::= \dots \mid \times \mid + \mid \rightarrow \mid \text{ref} \mid \forall_{\kappa} \mid \exists_{\kappa} \mid \mu_{\kappa}$

(constructors)

$v ::= \dots \mid \text{rec } f x = e \mid \Lambda e \mid \text{pack } v$

(values)

$e ::= \dots \mid \text{hash } e$

(expressions)

The $F_{\omega, \mu}^{\text{ref}}$ language

$\kappa ::= \star \mid \kappa \Rightarrow \kappa$ (kinds)

$\tau ::= \alpha \mid \lambda \alpha : \kappa . \tau \mid \tau \tau \mid c$ (types)

$c ::= \dots \mid \times \mid + \mid \rightarrow \mid \text{ref} \mid \forall_{\kappa} \mid \exists_{\kappa} \mid \mu_{\kappa}$ (constructors)

$v ::= \dots \mid \text{rec } f x = e \mid \Lambda e \mid \text{pack } v$ (values)

$e ::= \dots \mid \text{hash } e$ (expressions)

We write, e.g., $\forall \alpha : \kappa . \tau$ to mean $\forall_{\kappa} (\lambda \alpha : \kappa . \tau)$ and $\tau_1 \times \tau_2$ for $\times \tau_1 \tau_2$

Authentikit in $F_{\omega, \mu}^{\text{ref}}$

```
module type AUTHENTIKIT = sig
  type 'a auth
  type 'a auth_computation

  val return : 'a -> 'a auth_computation
  val bind   : 'a auth_computation ->
    ('a -> 'b auth_computation) ->
    'b auth_computation

  module Serializable : sig
    type 'a evidence
    val auth   : 'a auth evidence
    val pair   : 'a evidence -> 'b evidence -> ('a * 'b) evidence
    val sum    : 'a evidence -> 'b evidence ->
      ['`left of 'a | `right of 'b] evidence
    val string : string evidence
    val int    : int evidence
  end

  val auth   : 'a Serializable.evidence -> 'a -> 'a auth
  val unauth : 'a Serializable.evidence ->
    'a auth -> 'a auth_computation
end
```

$AUTHENTIKIT \triangleq \exists \text{auth}, m : \star \implies \star . \text{Authentikit auth } m$

$\text{Authentikit} \triangleq \lambda \text{auth}, m : \star \implies \star .$

$(\forall \alpha : \star . \alpha \rightarrow m \alpha) \times$

$(\forall \alpha, \beta : \star . m \alpha \rightarrow (\alpha \rightarrow m \beta) \rightarrow m \beta) \times$

\vdots

$(\forall \alpha : \star . \text{evi } \alpha \rightarrow \alpha \rightarrow \text{auth } \alpha) \times$

$(\forall \alpha : \star . \text{evi } \alpha \rightarrow \text{auth } \alpha \rightarrow m \alpha)$

Our approach

To show security and correctness we

1. Define a **program logic** that is expressive enough for proving that programs have the property in question, e.g., a variant of Hoare logic.
2. Define a **semantic model** of the type system, in which types are given meaning through Hoare triples of the program logic.

Using the rules of the logic, we then show that the model is sound and that well-typed terms inhabit the model.

Collision-free reasoning

We first define a relational **Collision-Free Separation Logic (CF-SL)** on top of Iris.

$$\{P\} e_1 \sim e_2 \{Q\}$$

CF-SL statements hold **“up to”** hash collision: given P holds for the initial state,

if e_1 evaluates to v_1 and e_2 evaluates to v_2 then $Q(v_1, v_2)$ holds

or a hash collision occurred.

Collision-f

Security: If the **verifier** accepts a proof p and returns v then

- the **ideal** execution returns v or
- a hash collision occurred.

We first define a relational **Collision-Free Separation Logic (CF-SL)** on top of Iris.

$$\{P\} e_1 \sim e_2 \{Q\}$$

CF-SL statements hold “**up to**” hash collision: given P holds for the initial state,

if e_1 evaluates to v_1 and e_2 evaluates to v_2 then $Q(v_1, v_2)$ holds

or a hash collision occurred.

CF-SL

CF-SL satisfies all the standard program-logic rules, e.g.,

$$\frac{\{P\} e_1 \sim e'_2 \{Q\} \quad e_2 \rightsquigarrow e'_2}{\{P\} e_1 \sim e_2 \{Q\}}$$

$$\frac{\{\ell \mapsto w\} () \sim e_2 \{Q\}}{\{\ell \mapsto v\} \ell := w \sim e_2 \{Q\}}$$

CF-SL

CF-SL satisfies all the standard program-logic rules, e.g.,

$$\frac{\{P\} e_1 \sim e'_2 \{Q\} \quad e_2 \rightsquigarrow e'_2}{\{P\} e_1 \sim e_2 \{Q\}}$$

$$\frac{\{\ell \mapsto w\} () \sim e_2 \{Q\}}{\{\ell \mapsto v\} \ell := w \sim e_2 \{Q\}}$$

but introduces a new proposition **hashed(s)** satisfying

$$\frac{\{P \star \text{hashed}(s)\} \text{hash}(s) \sim e_2 \{Q\}}{\{P\} \text{hash } s \sim e_2 \{Q\}}$$

$$\frac{\text{collision}(s_1, s_2)}{\text{hashed}(s_1) \star \text{hashed}(s_2) \vdash \text{False}}$$

Security

To show security of Authentikit, we use CF-SL to define a **logical relation**

$$\Theta \mid \Gamma \vDash e_1 \sim e_2 : \tau$$

and show

1. If $\Theta \mid \Gamma \vdash e : \tau$ then $\Theta \mid \Gamma \vDash e \sim e : \tau$
2. If $\Theta \mid \Gamma \vDash e_1 \sim e_2 : \tau$ then e_1 and e_2 are secure (as verifier and ideal)
3. $\emptyset \mid \emptyset \vDash \text{Authentikit}_V \sim \text{Authentikit}_I : \text{AUTHENTIKIT}$

Logical relation, sketch

Intuitively, the judgment $\emptyset \mid \emptyset \vDash e_1 \sim e_2 : \tau$ means

$$\{\text{True}\} e_1 \sim e_2 \{ \llbracket \tau \rrbracket \}$$

where $\llbracket \tau \rrbracket : \text{Val} \times \text{Val} \rightarrow \text{iProp}$ is an **interpretation of types**. E.g.

$$\llbracket \mathbb{N} \rrbracket(v_1, v_2) \triangleq \exists n \in \mathbb{N}. v_1 = v_2 = n$$

$$\llbracket \tau_1 \rightarrow \tau_2 \rrbracket(v_1, v_2) \triangleq \forall w_1, w_2. \{ \llbracket \tau_1 \rrbracket(w_1, w_2) \} v_1 w_1 \sim v_2 w_2 \{ \llbracket \tau_2 \rrbracket \}$$

Theorem (Security)

If e is a program parameterized by an Authentikit implementation, i.e.,

$$\emptyset \mid \emptyset \vdash e : \forall \text{auth}, m . \text{Authentikit } \text{auth } m \rightarrow m \tau$$

then for all proofs p , if

$$e \text{ Authentikit}_V p \rightarrow_{\text{cf}}^* \text{Some}(v)$$

then

$$e \text{ Authentikit}_I \rightarrow^* v$$

Theorem (Correctness)

If e is a program parameterized by an Authentikit implementation, i.e.,

$$\emptyset \mid \emptyset \vdash e : \forall \text{auth}, m . \text{Authentikit } \text{auth } m \rightarrow m \tau$$

then if

$$e \text{ Authentikit}_P \rightarrow_{\text{cf}}^* (p, v)$$

then

$$e \text{ Authentikit}_V p \rightarrow^* \text{Some}(v) \quad \text{and} \quad e \text{ Authentikit}_I \rightarrow^* v$$

Optimizations of Authentikit

- Proof accumulator
- Proof-reuse buffering
- Heterogeneous buffering
- Stateful buffering

```
module Verifier : AUTHENTIKIT =
  type 'a auth_computation =
    pfstate -> [`Ok of pfstate * 'a | `ProofFailure]

  (* ... *)

  let unauth evi h pf =
    match Map.find_opt h pf.cache with
    | None ->
      match pf.pf_stream with
      | [] -> `ProofFailure
      | p :: ps when hash p = h ->
        match evi.deserialize p with
        | None -> `ProofFailure
        | Some a ->
          `Ok ({pf_stream = ps;
              cache = Map.add h p pf.cache}, a)
      | _ -> `ProofFailure
    | Some p ->
      match evi.deserialize p with
      | None -> `ProofFailure
      | Some a -> `Ok (pf, a)

  end
```


Manual proofs

The naïve implementation of Authentikit does not emit the minimal proofs, e.g.,

$$\text{lookup}([R, L], t_0) = ([(h_1, h_2), (h_5, h_6), s_5], s_5)$$



Instead, we can manually implement and “semantically type” the optimal strategy:

$$\llbracket \text{path} \rightarrow \text{auth tree} \rightarrow m \text{ (option string)} \rrbracket (\text{fetch}_V, \text{fetch}_I)$$



Summary

- **Authentikit** is a library for implementing ADSs generically.
- Two **logical-relations models** and a proof of security and correctness of the typed module construction in a general-purpose programming language.
 - We verify several **optimizations**.
 - We show how to prove that manually verified code can be **safely linked** with automatically generated code.
- Full mechanization in the Rocq theorem prover.



That's it, folks !

