

Logical Relations for Formally Verified

Authenticated Data Structures

Simon Oddershede Gregersen

joint work with Chaitanya Agarwal and Joseph Tassarotti

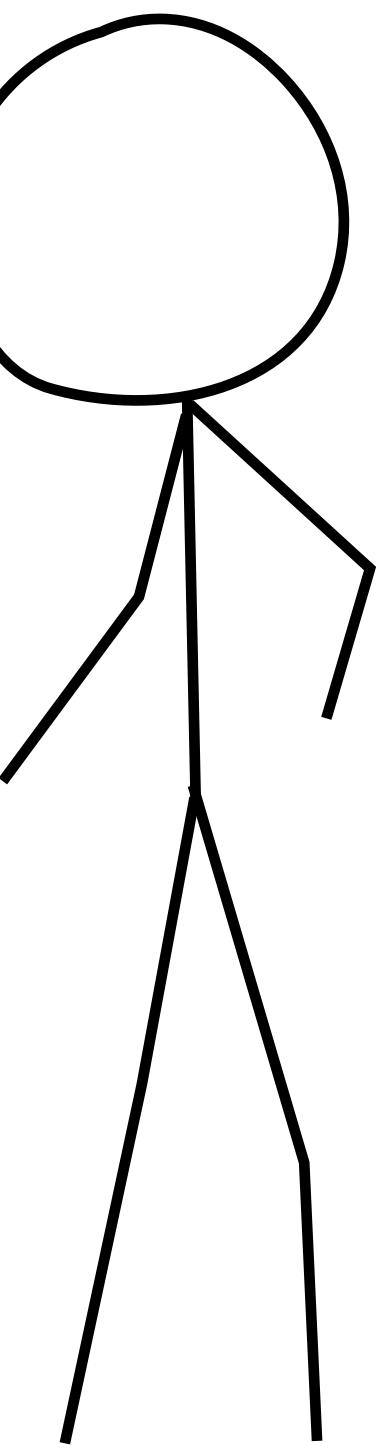
(to appear at CCS'25)



I have so much stuff to store!



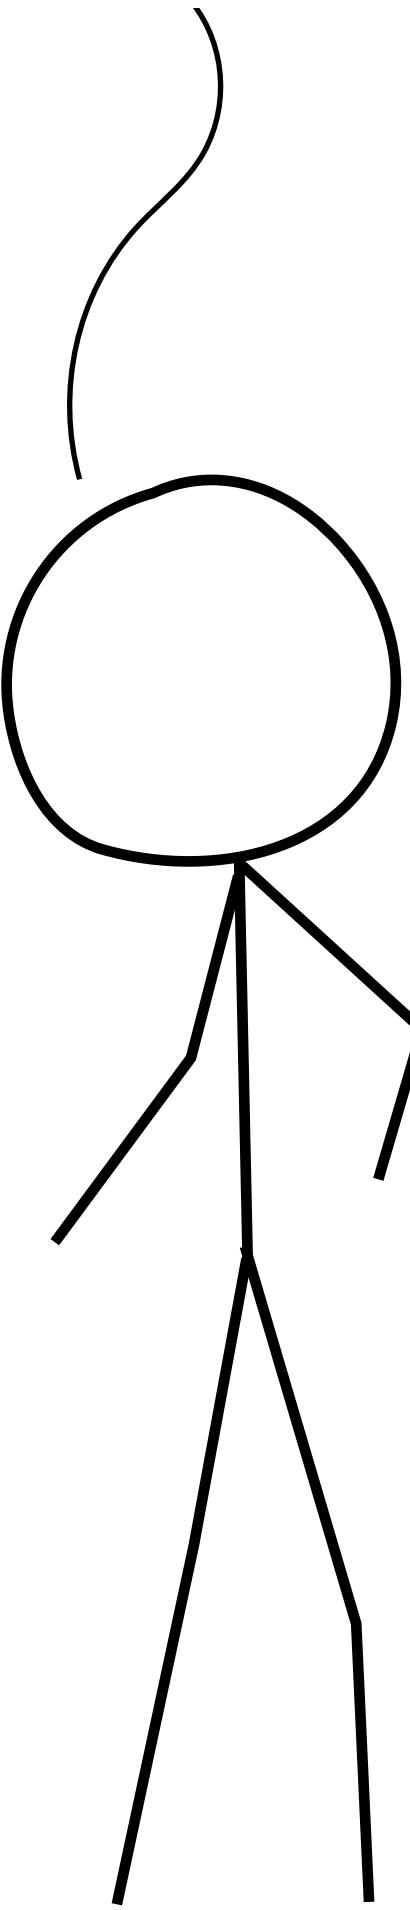
I have so much stuff to store!



I have so much stuff to store!



I can help!

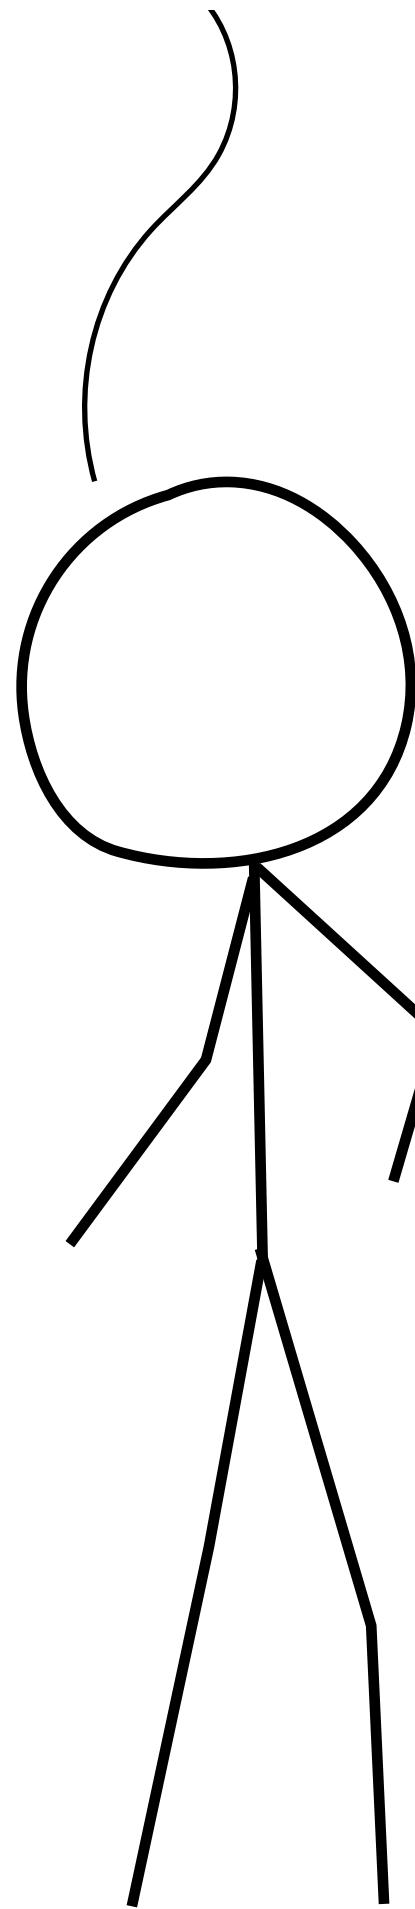


I have so much stuff to store!

Can I trust you to not mess it up?



I can help!



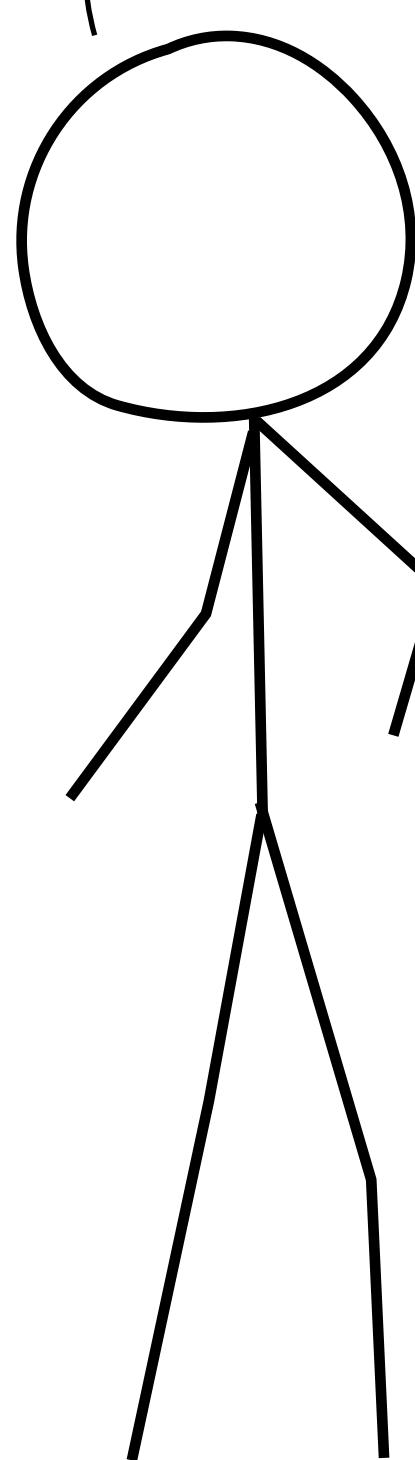
I have so much stuff to store!

Can I trust you to not mess it up?



I can help!

Of course!



How can Alice outsource data storage to Bob?

How can Alice outsource data storage to Bob?

If Alice can state her work as operations on an **authenticated data structure** then they can be outsourced to Bob, but later verified by Alice!

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

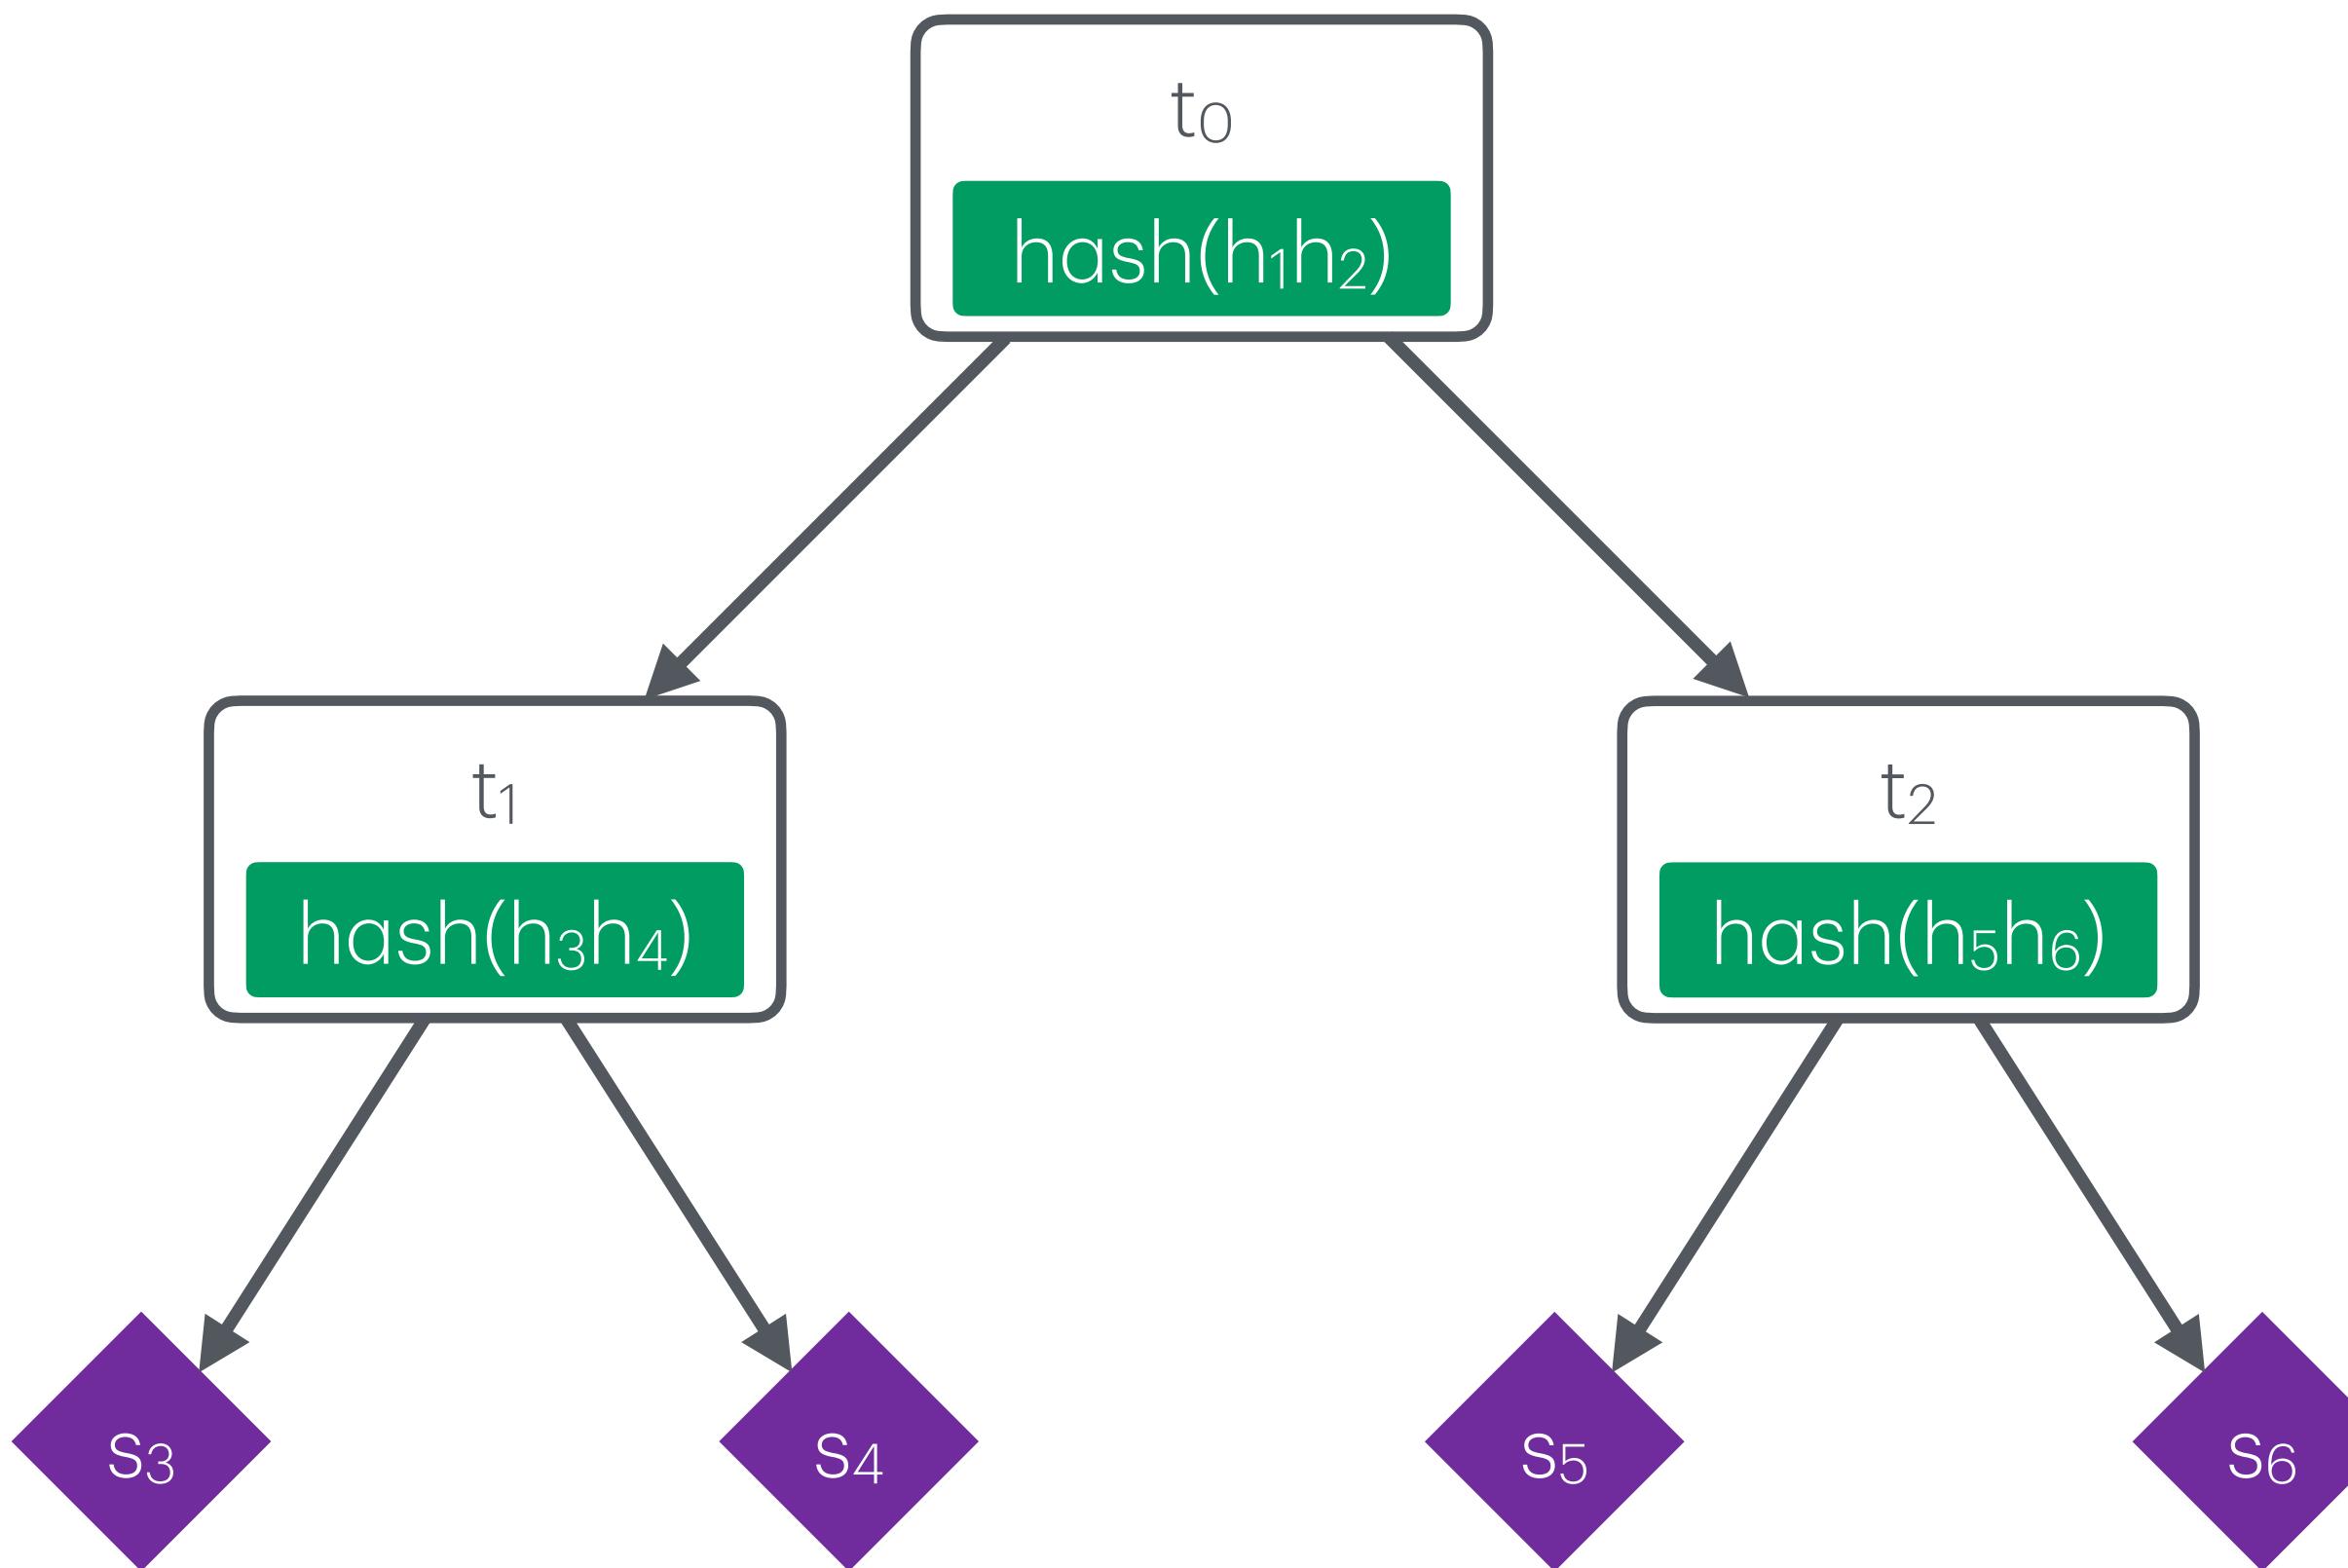
How can Alice outsource data storage to Bob?

If Alice can state her work as operations on an **authenticated data structure** then they can be outsourced to Bob, but later verified by Alice!

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

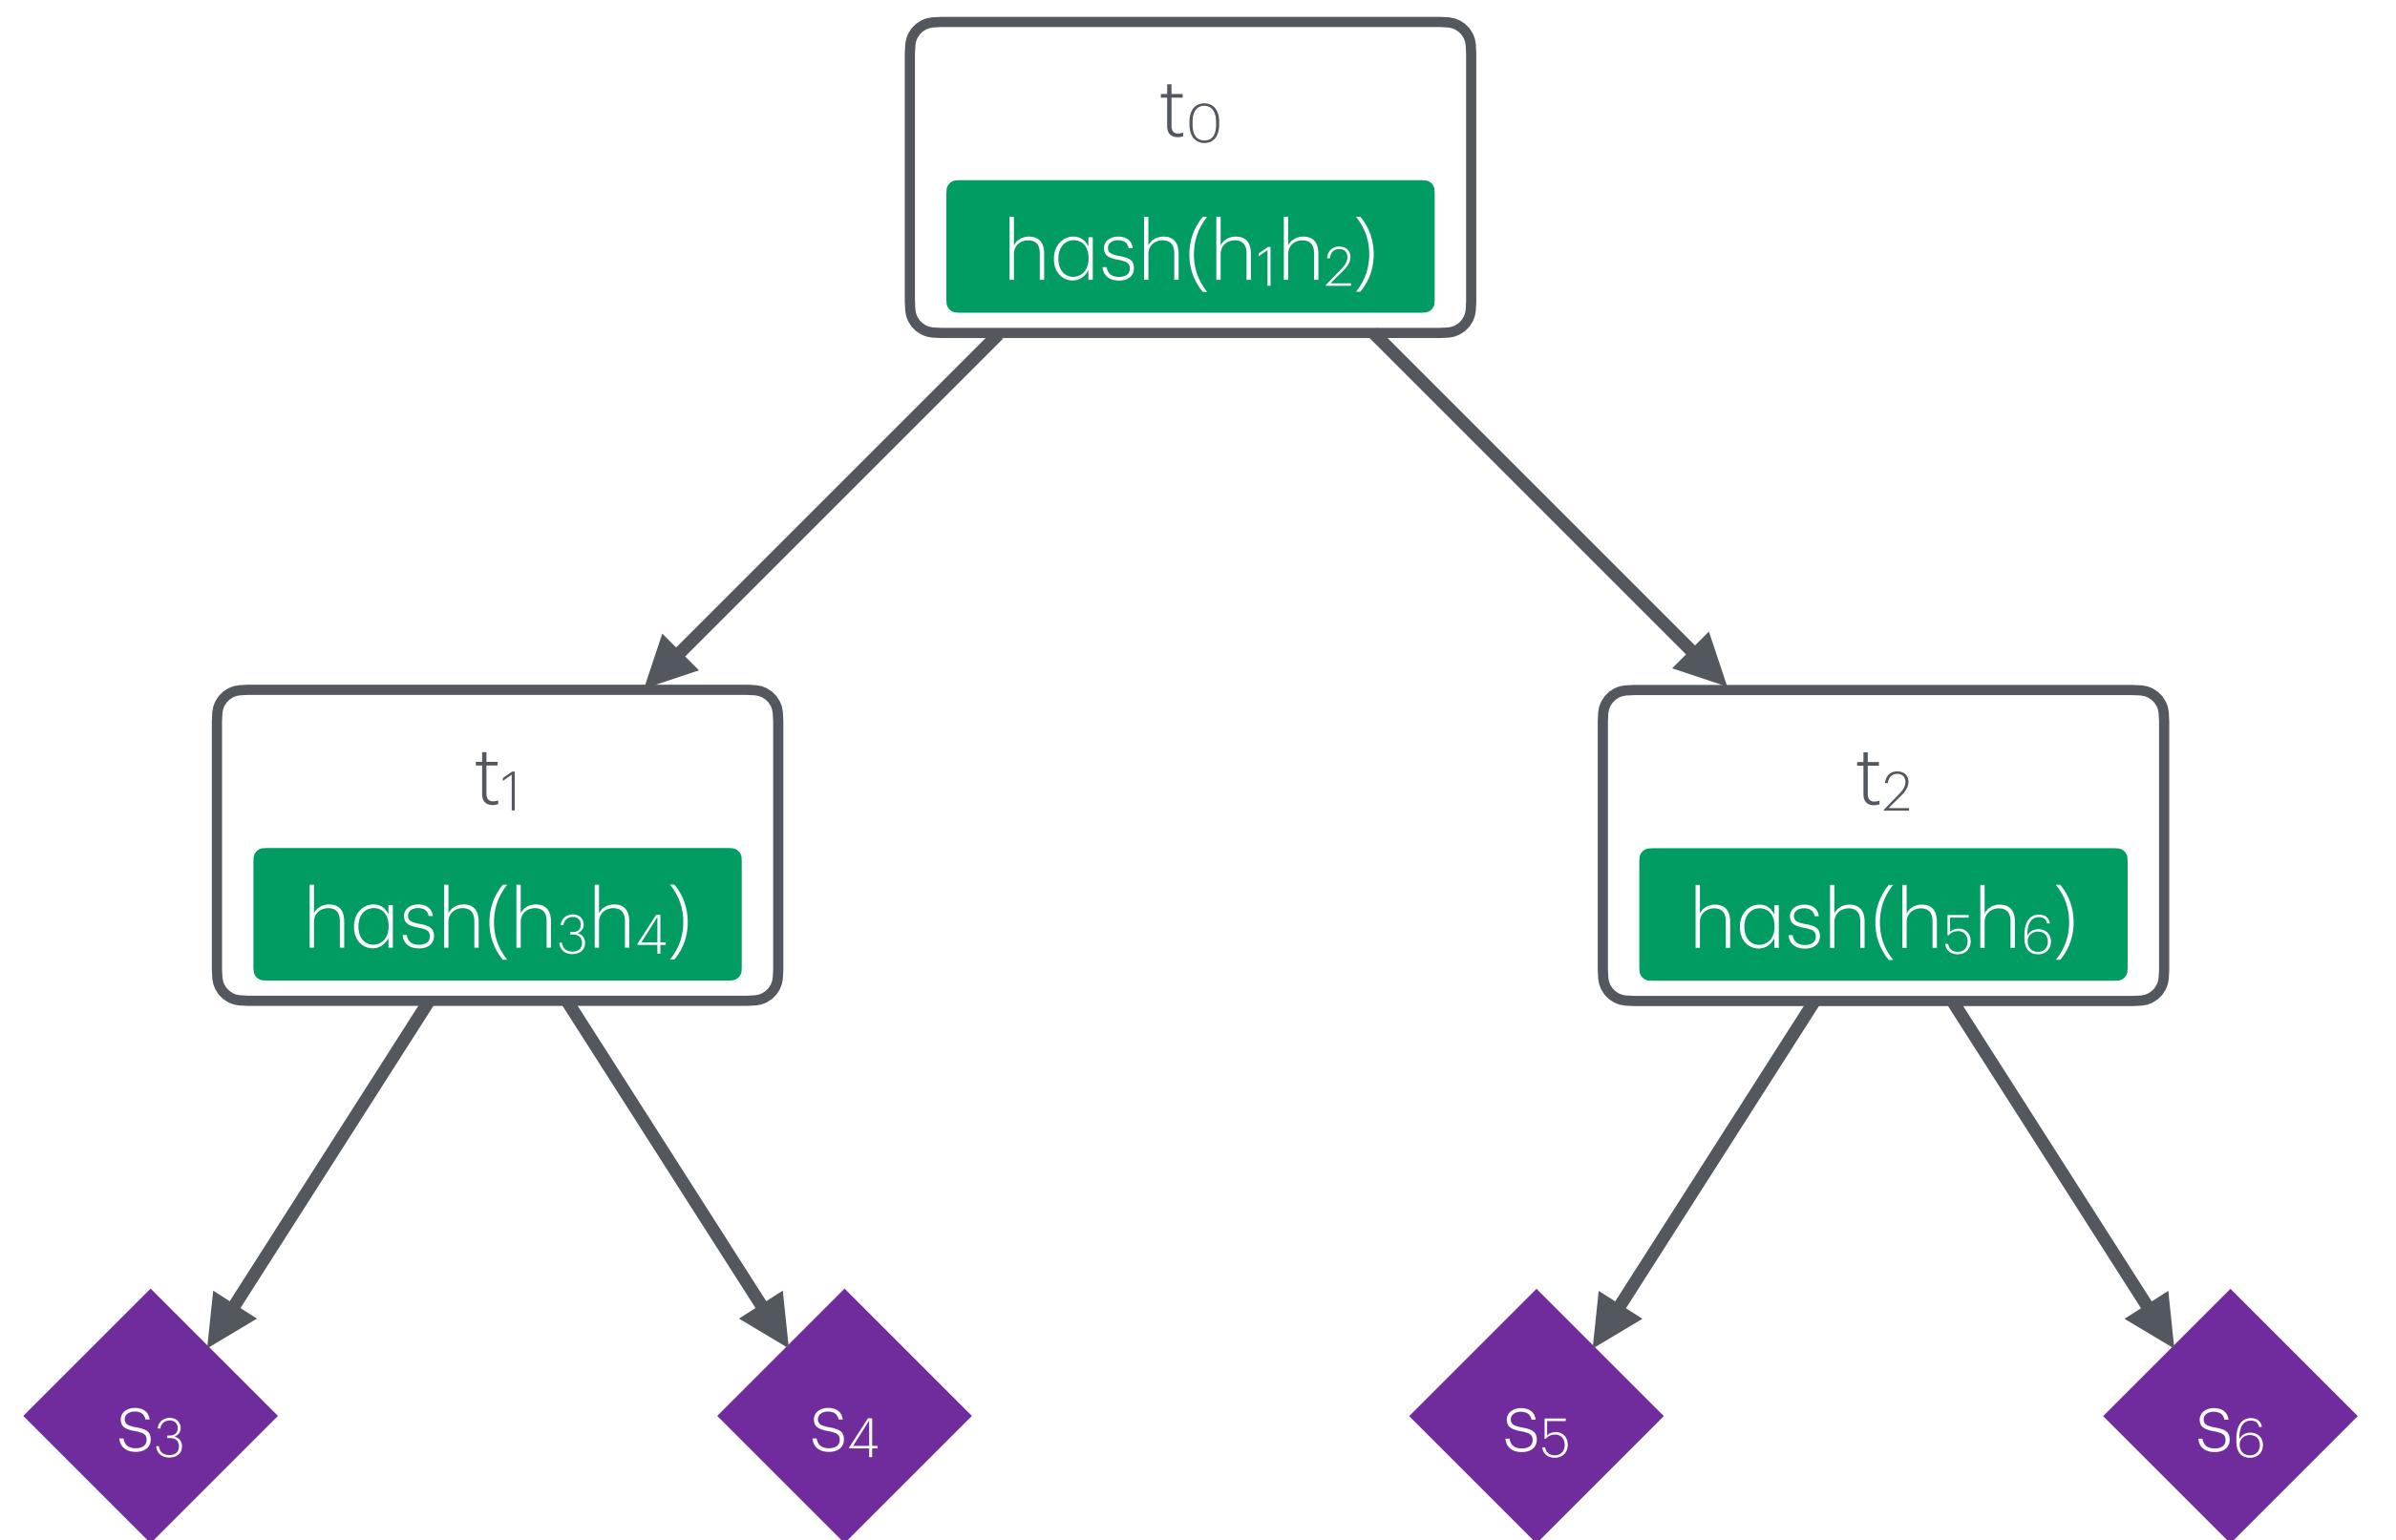
ADSs allow outsourcing data storage and processing tasks to untrusted servers without loss of integrity.

Example: Merkle Tree

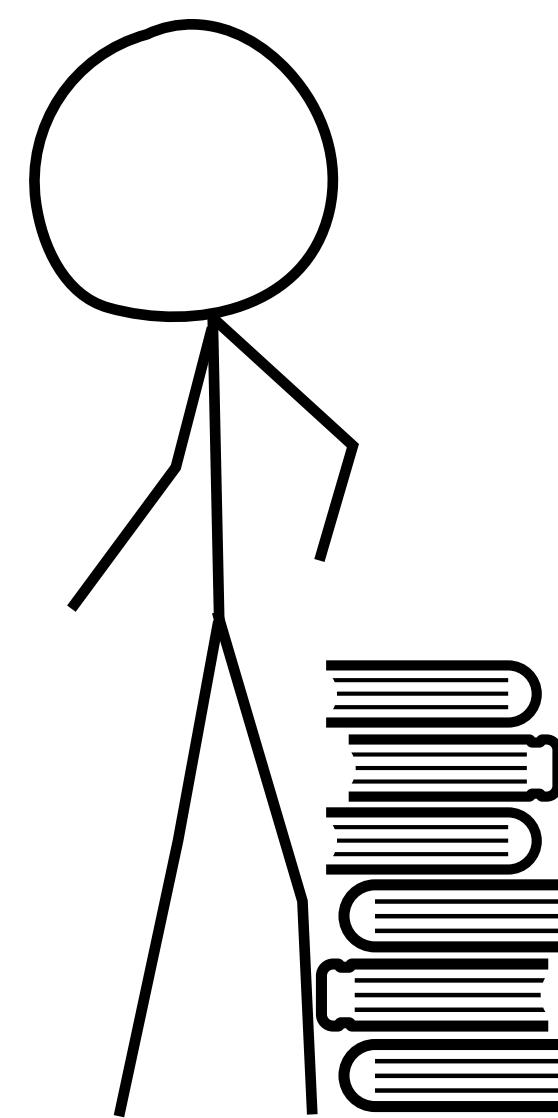


where h_i denotes the hash of t_i / s_i

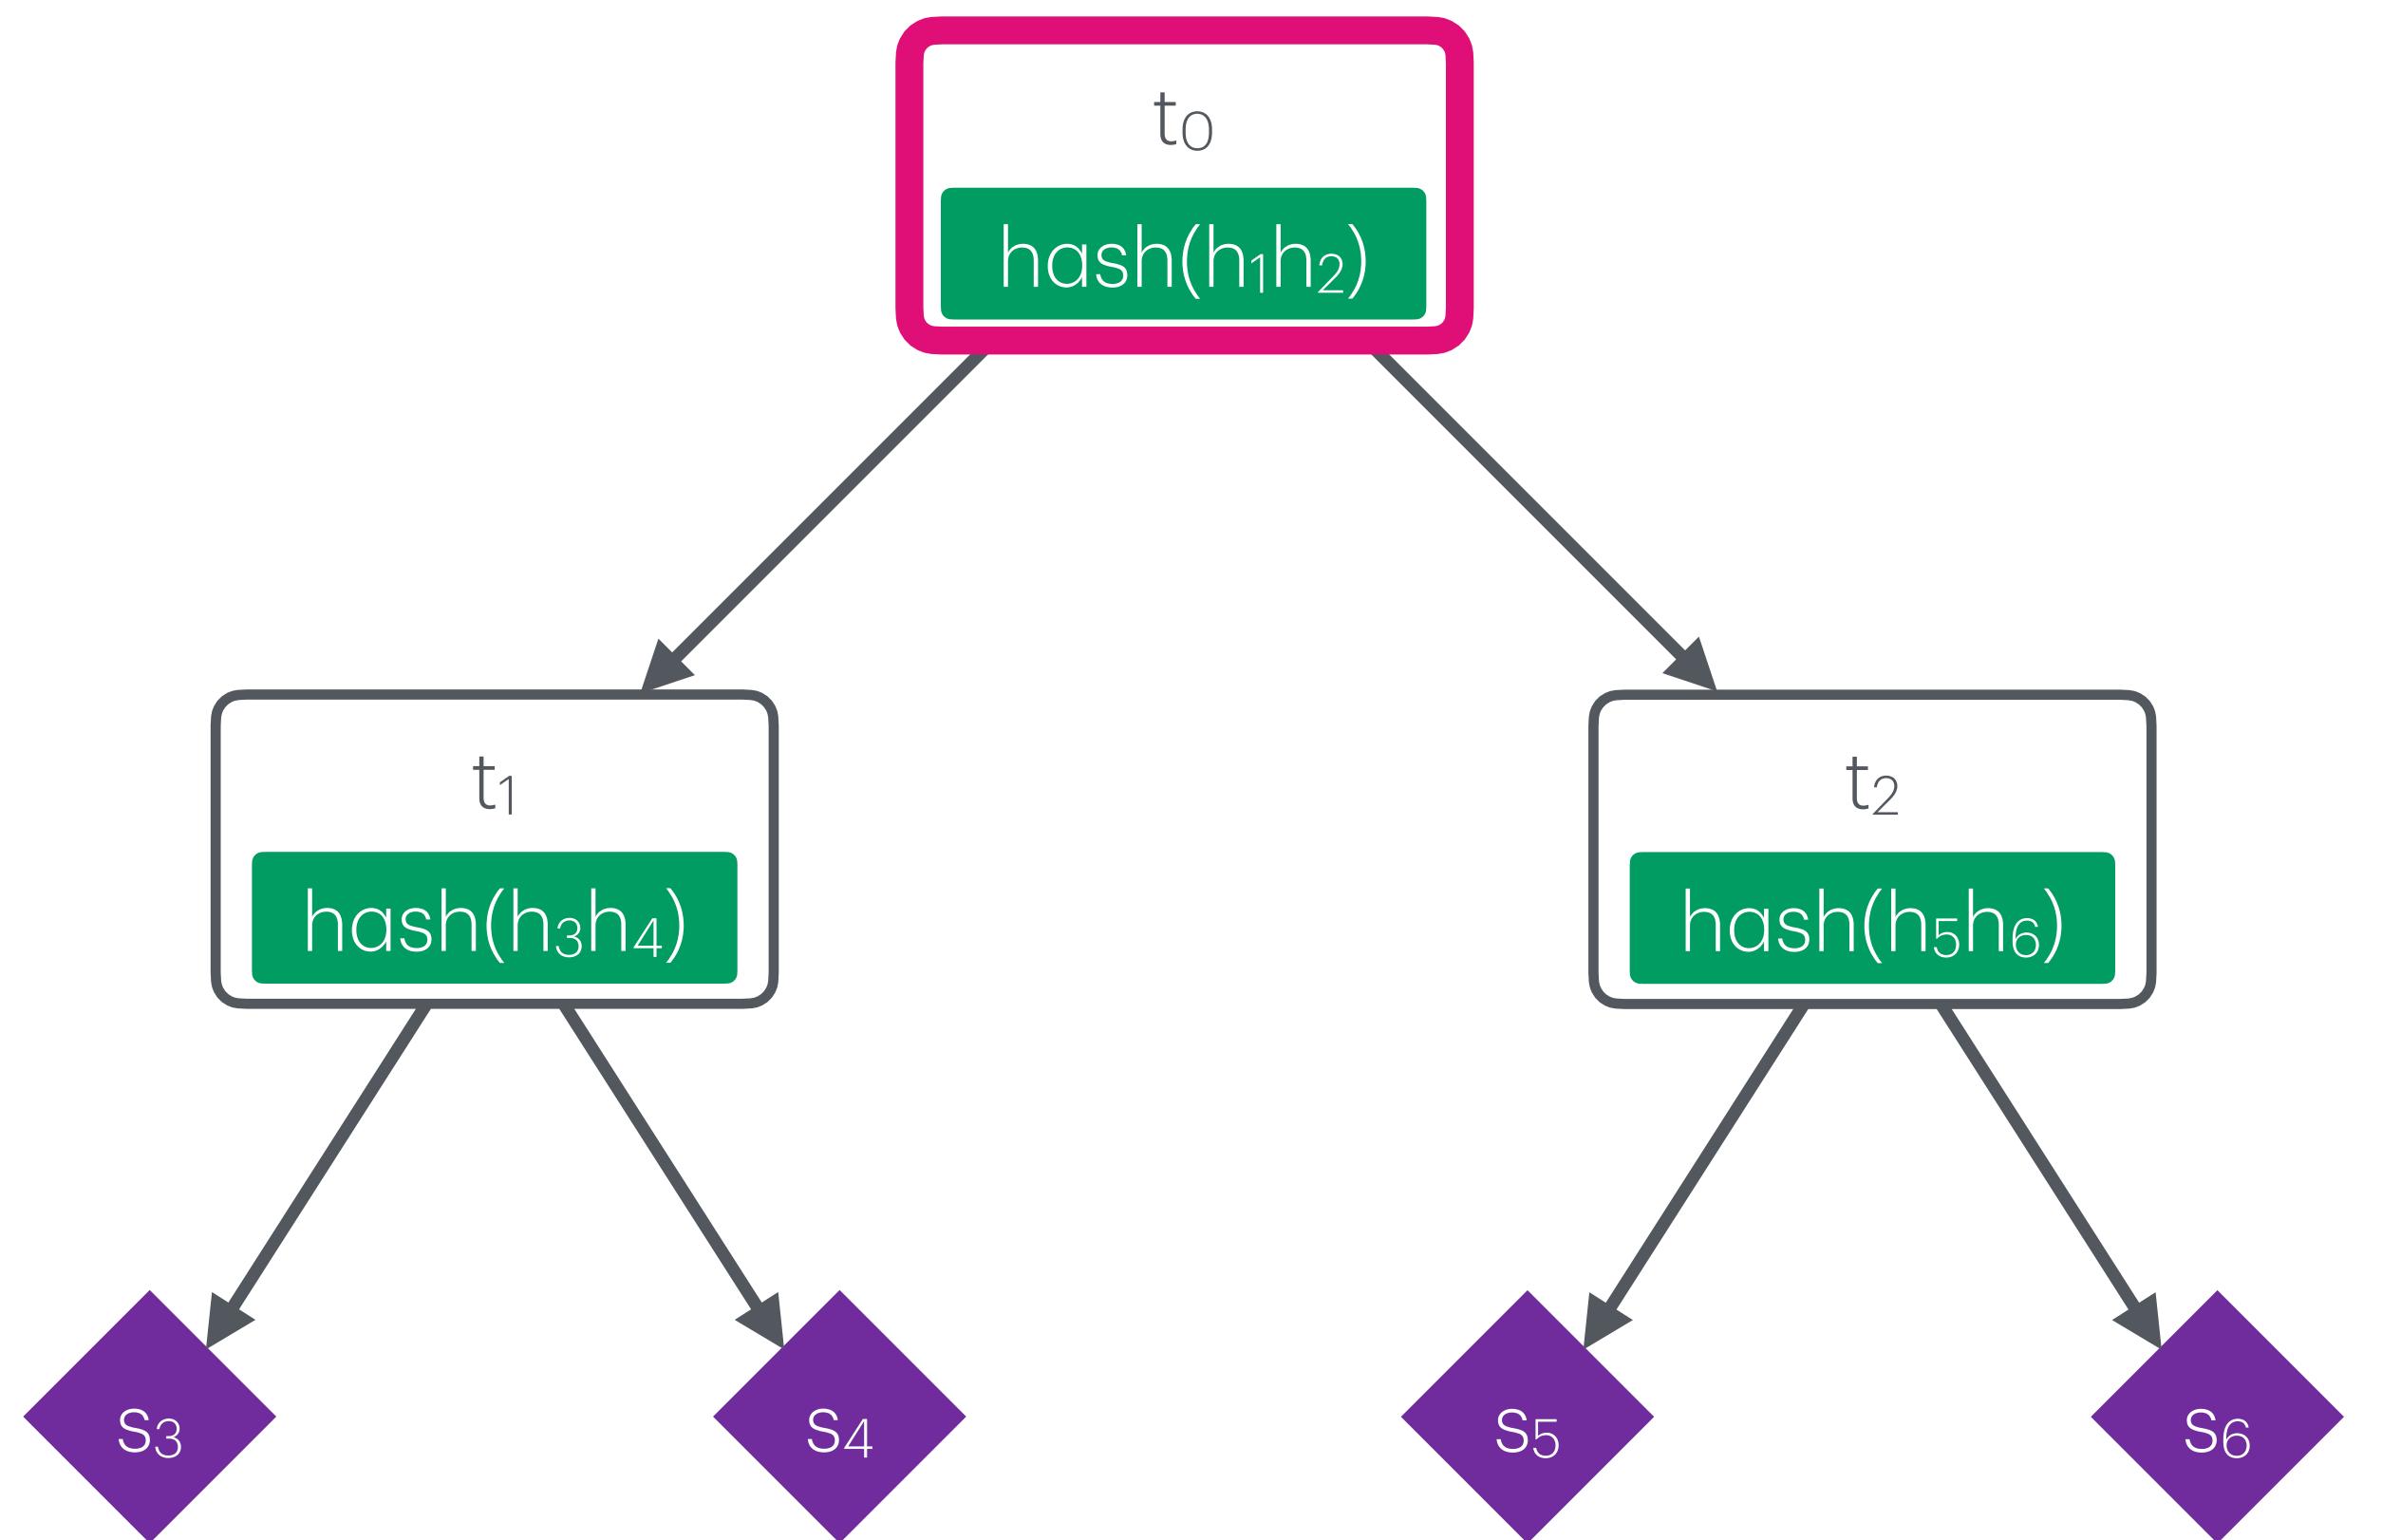
Example: Merkle Tree (Prover)



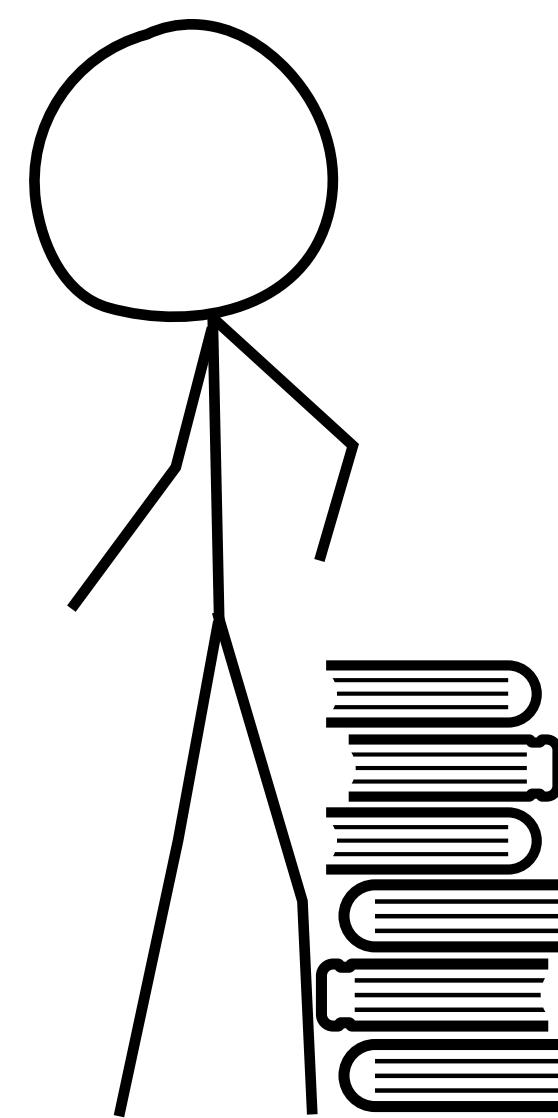
`fetch([R, L], t0) =`



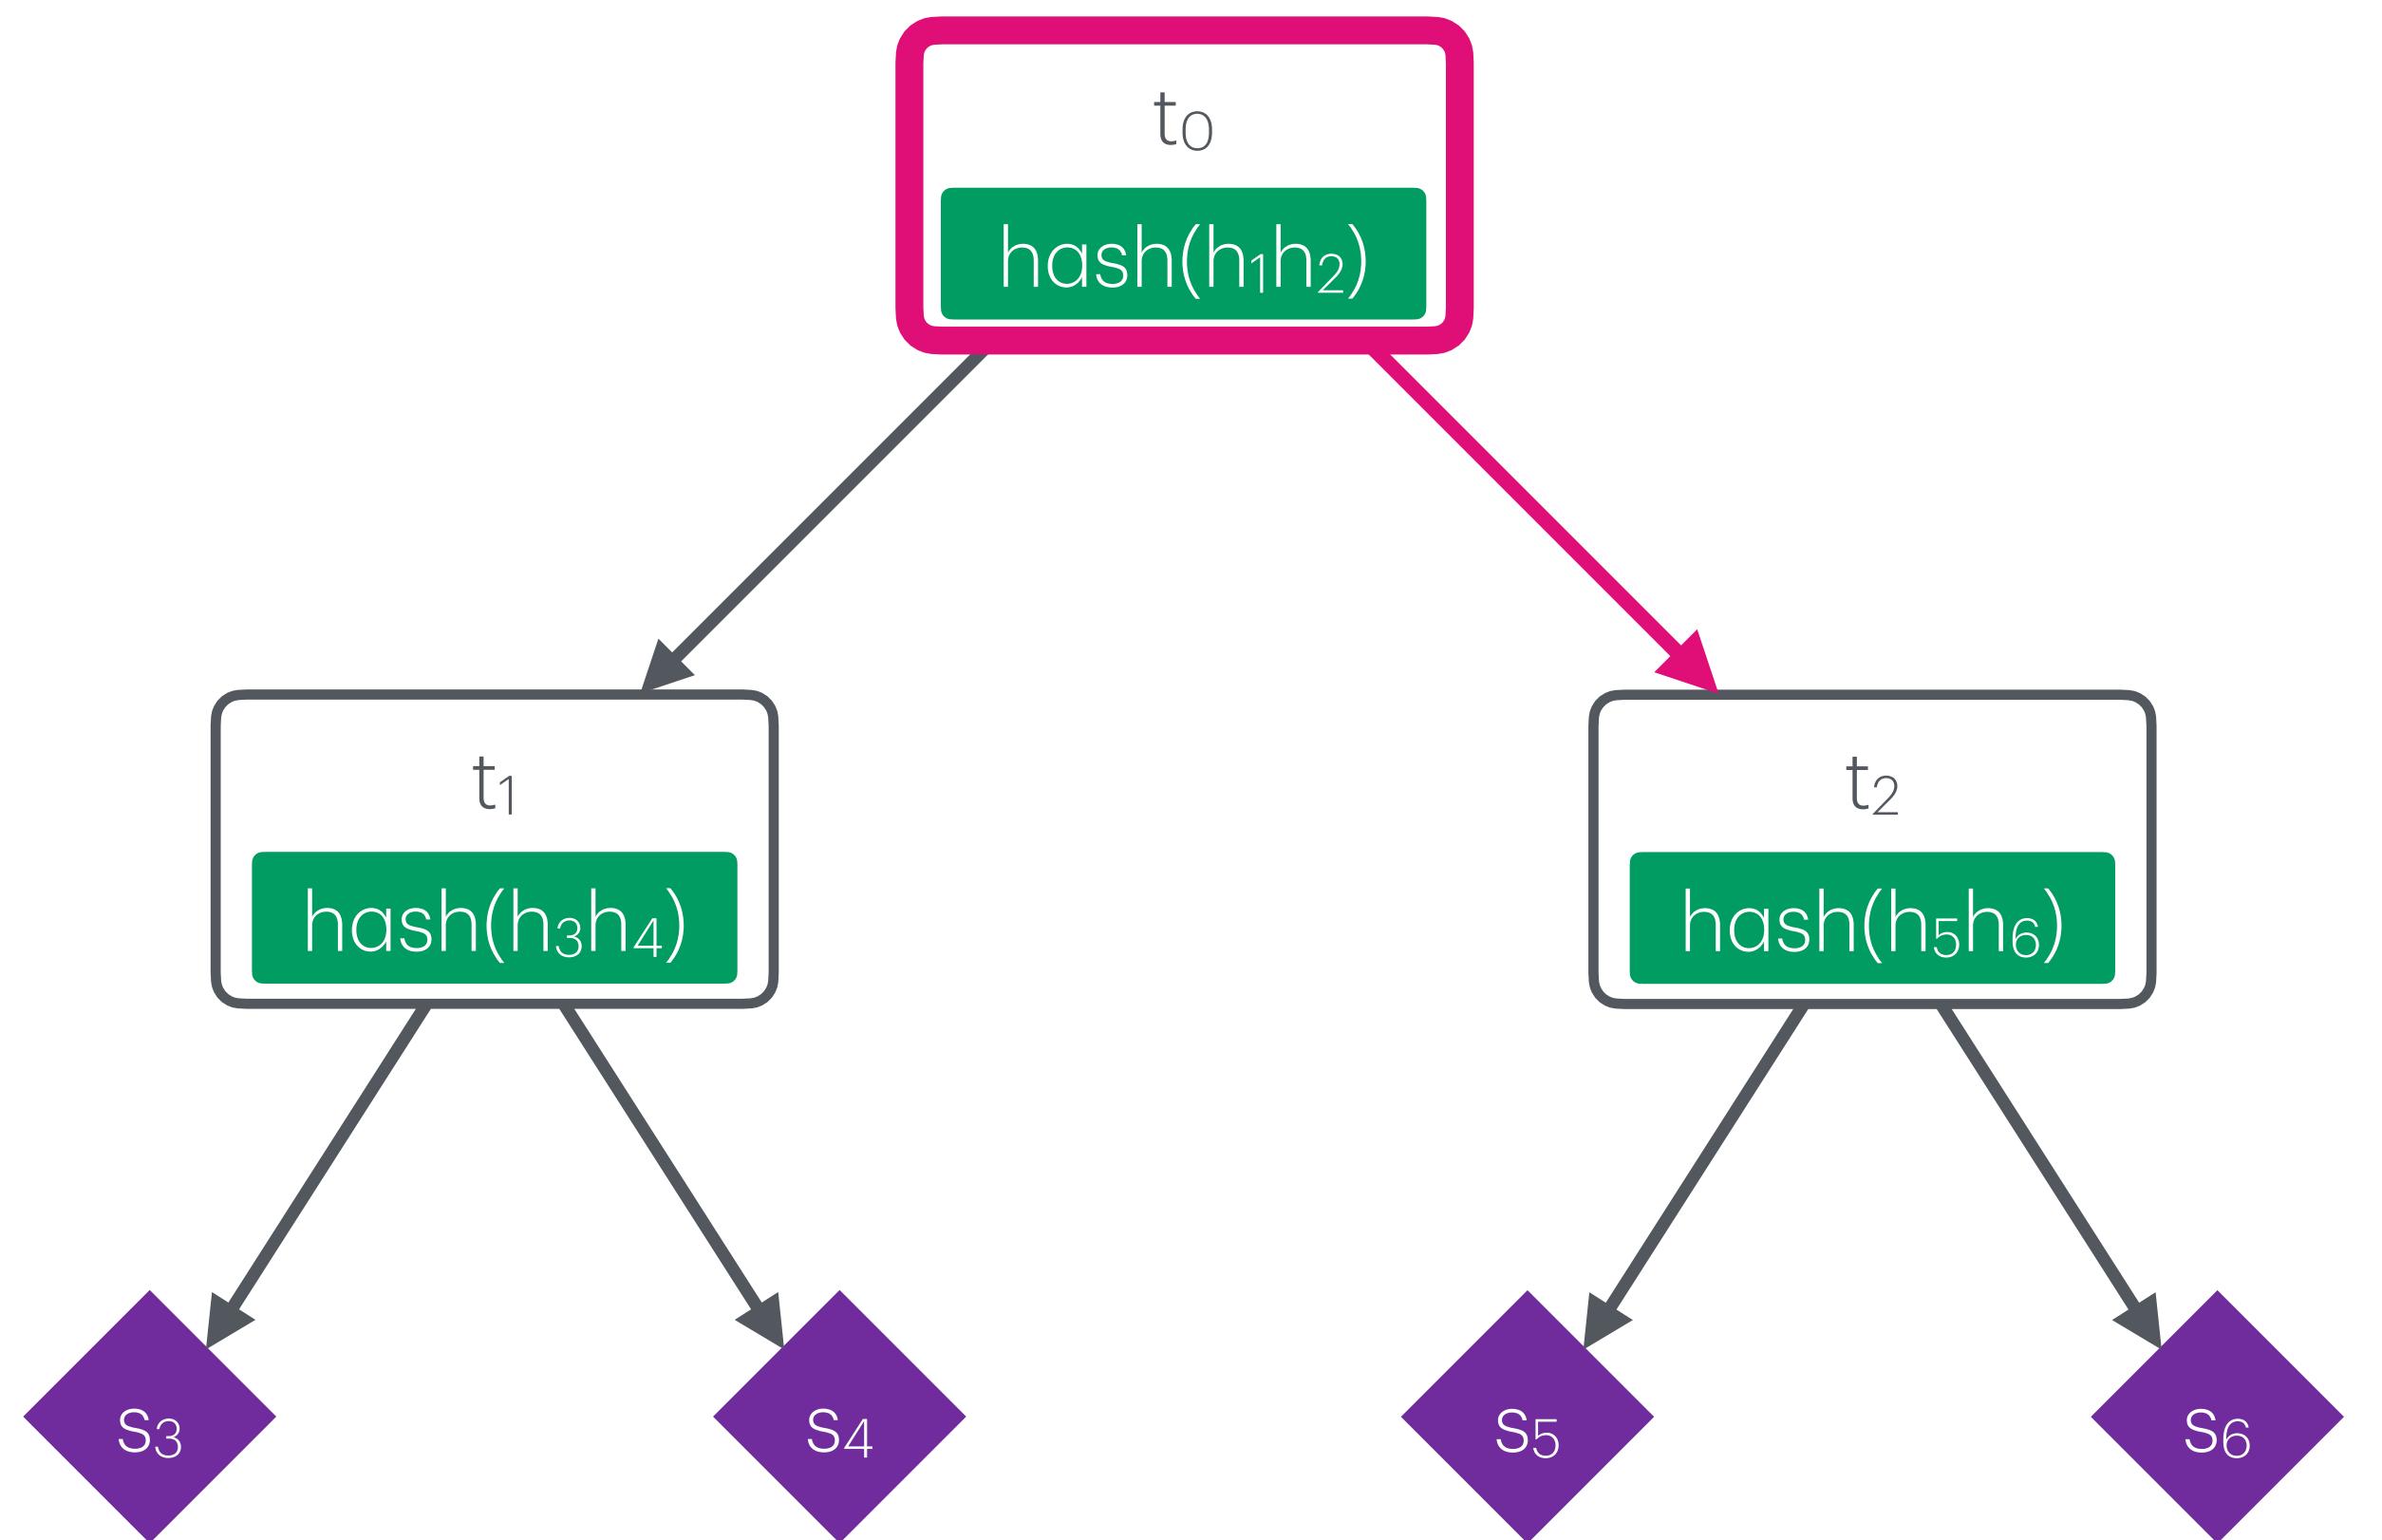
Example: Merkle Tree (Prover)



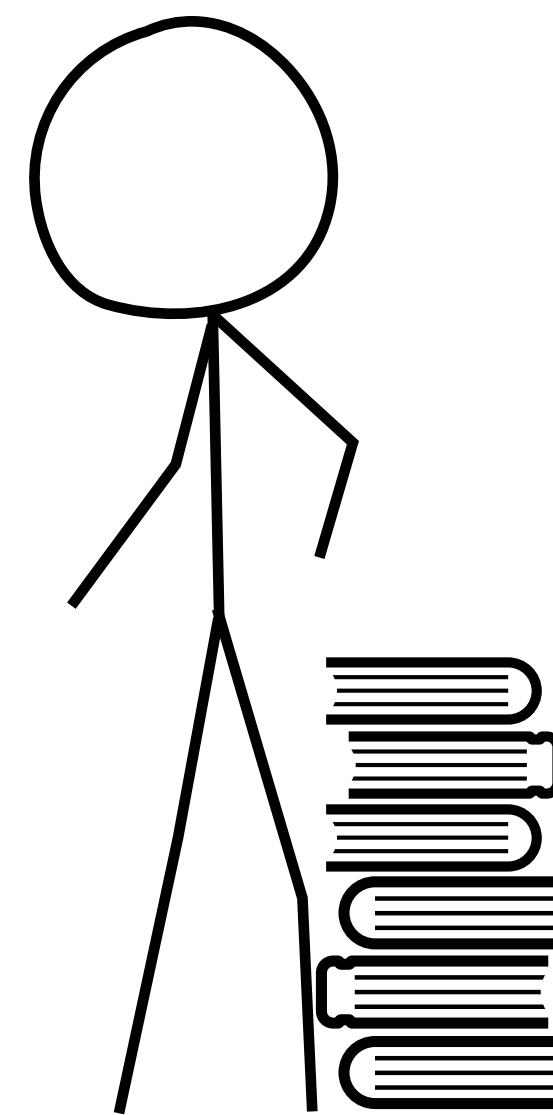
`fetch([R, L], t0) =`



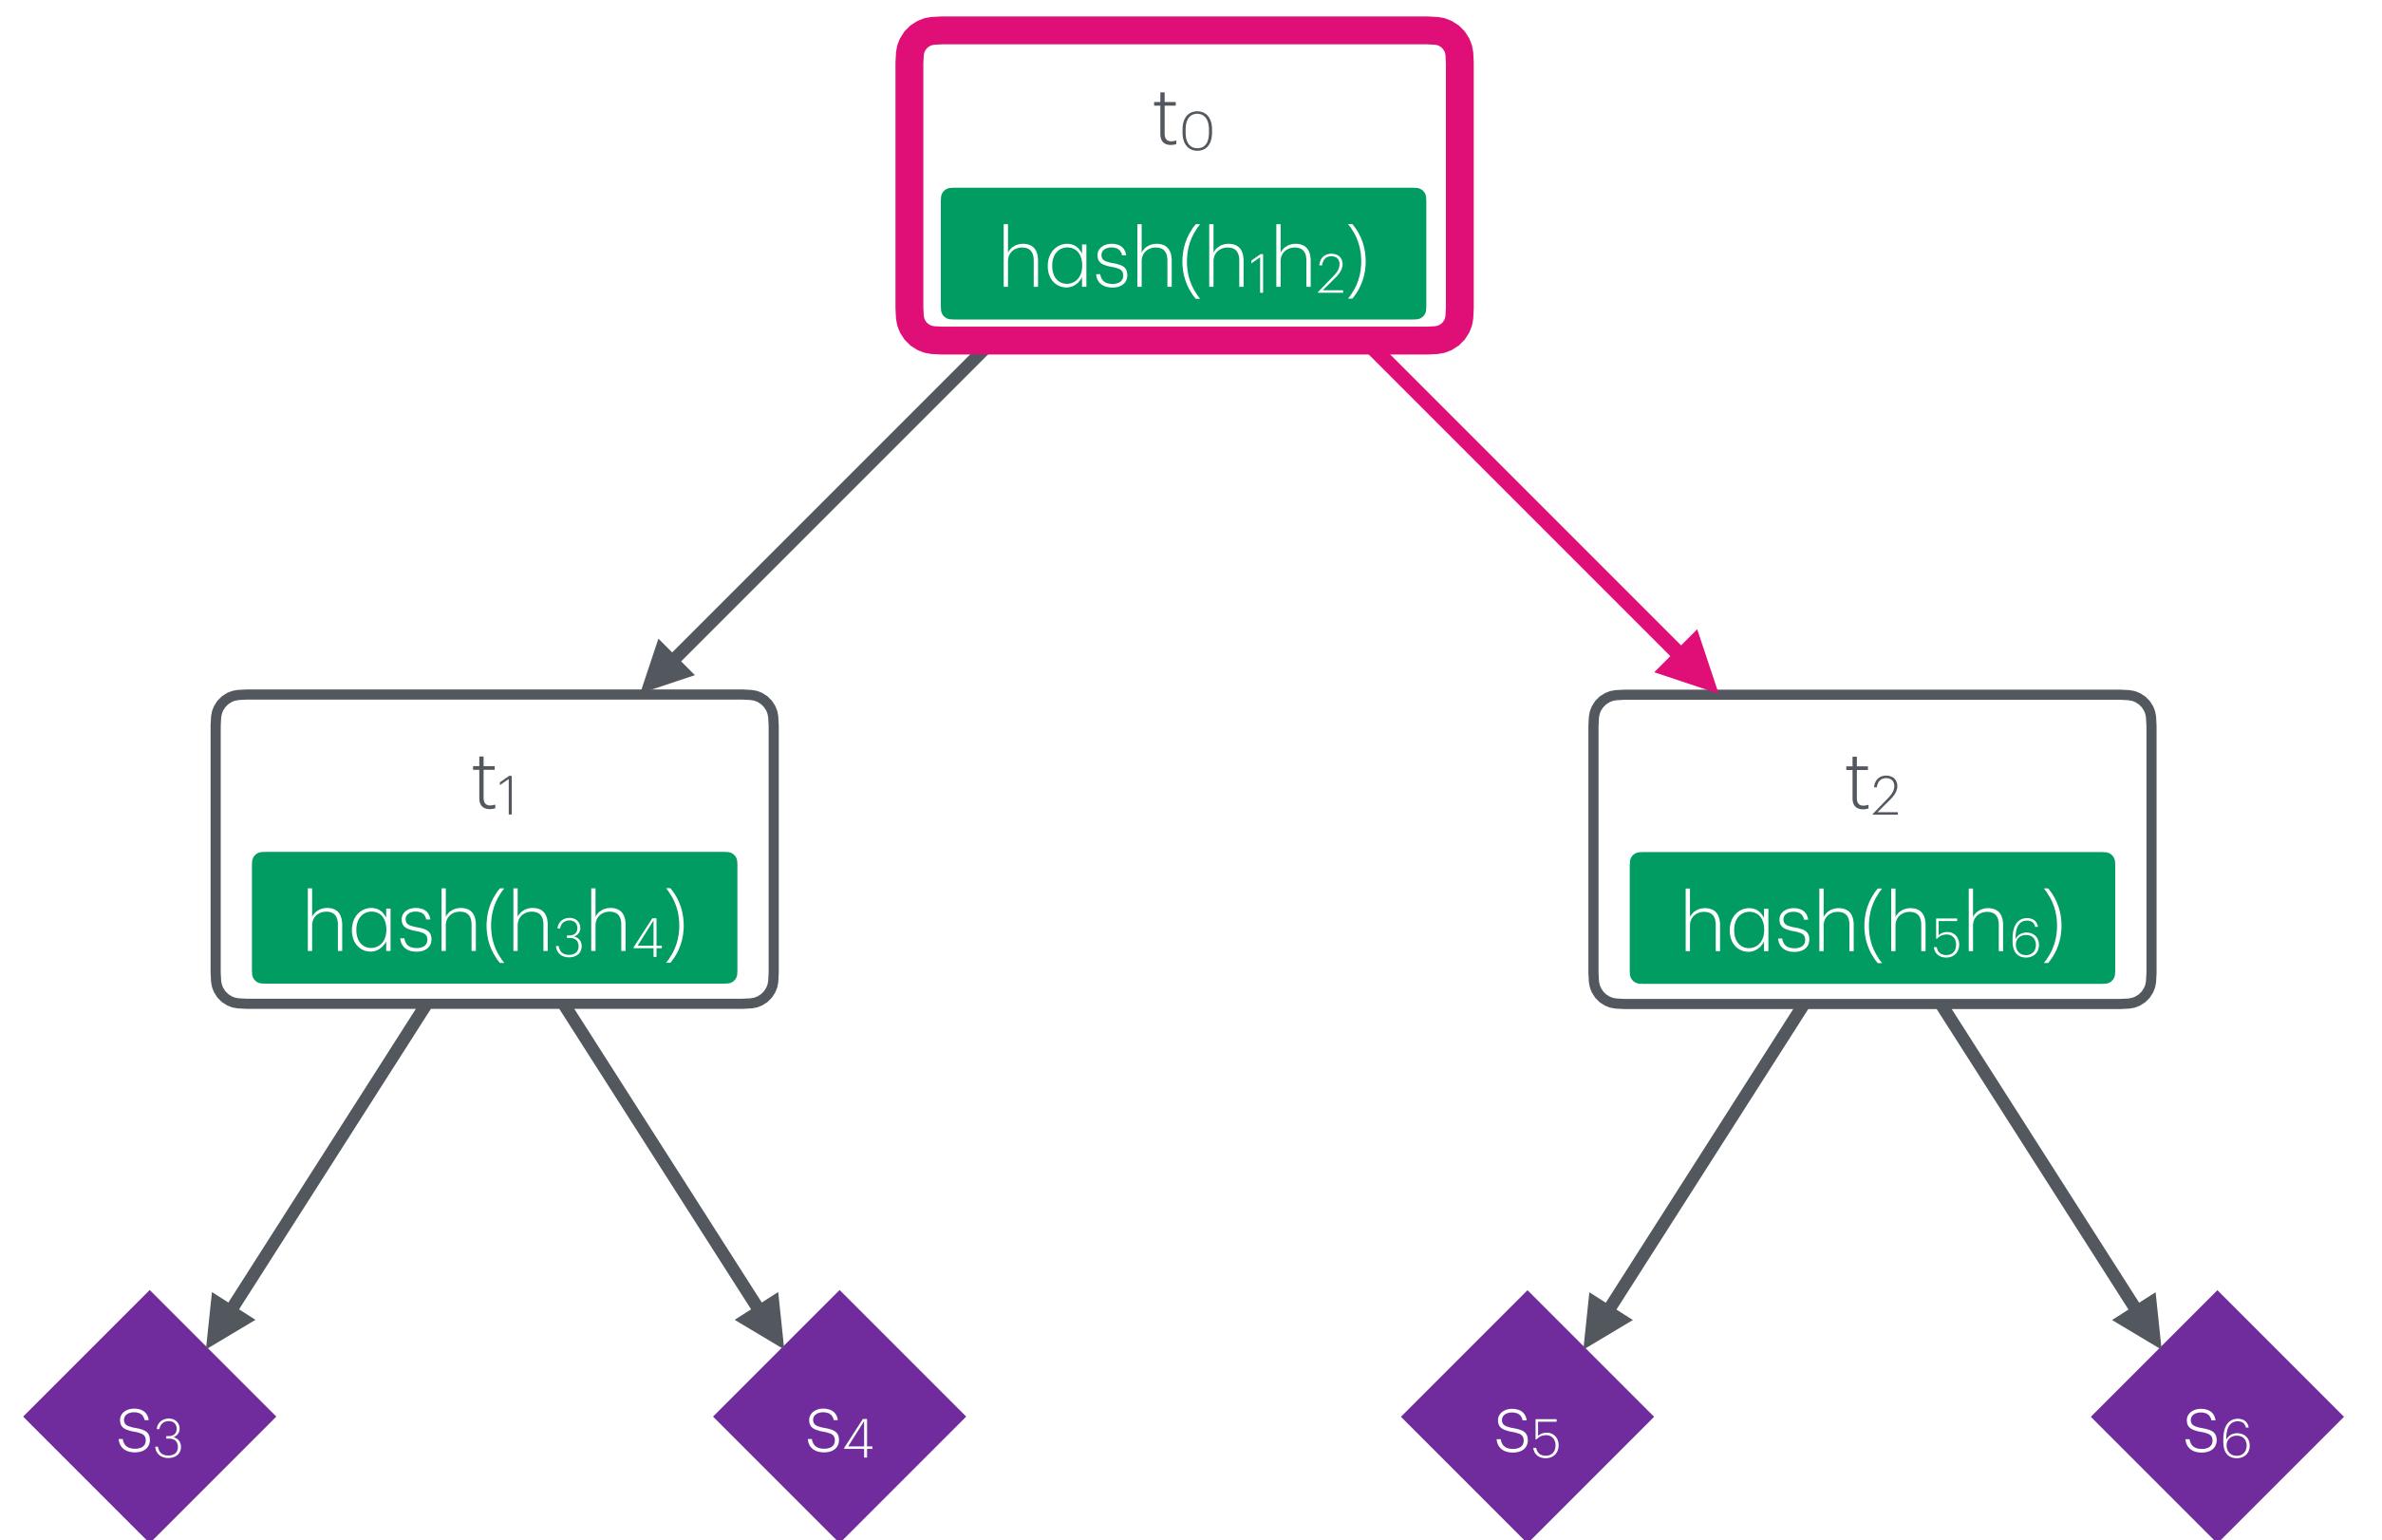
Example: Merkle Tree (Prover)



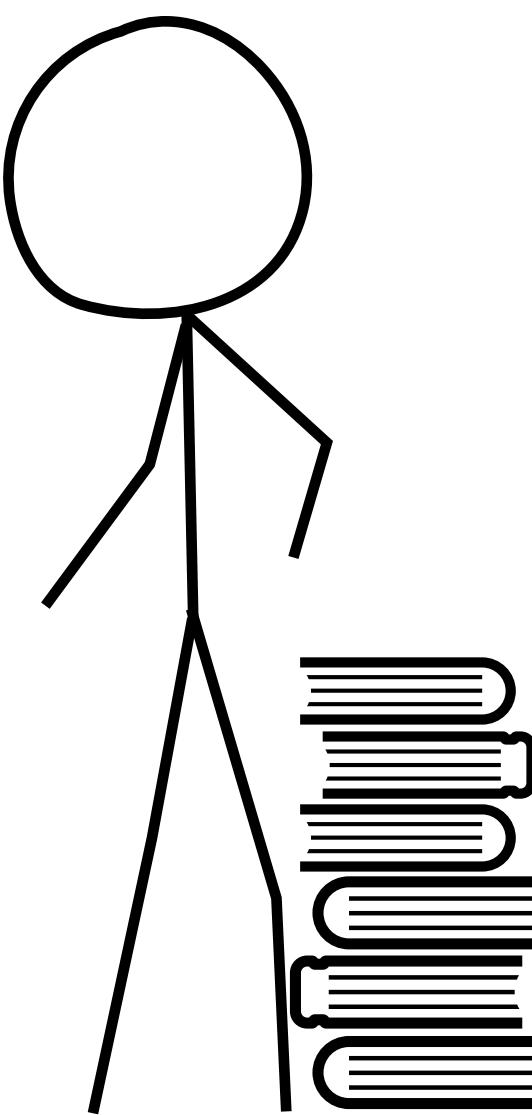
`fetch([R, L], t0) =`



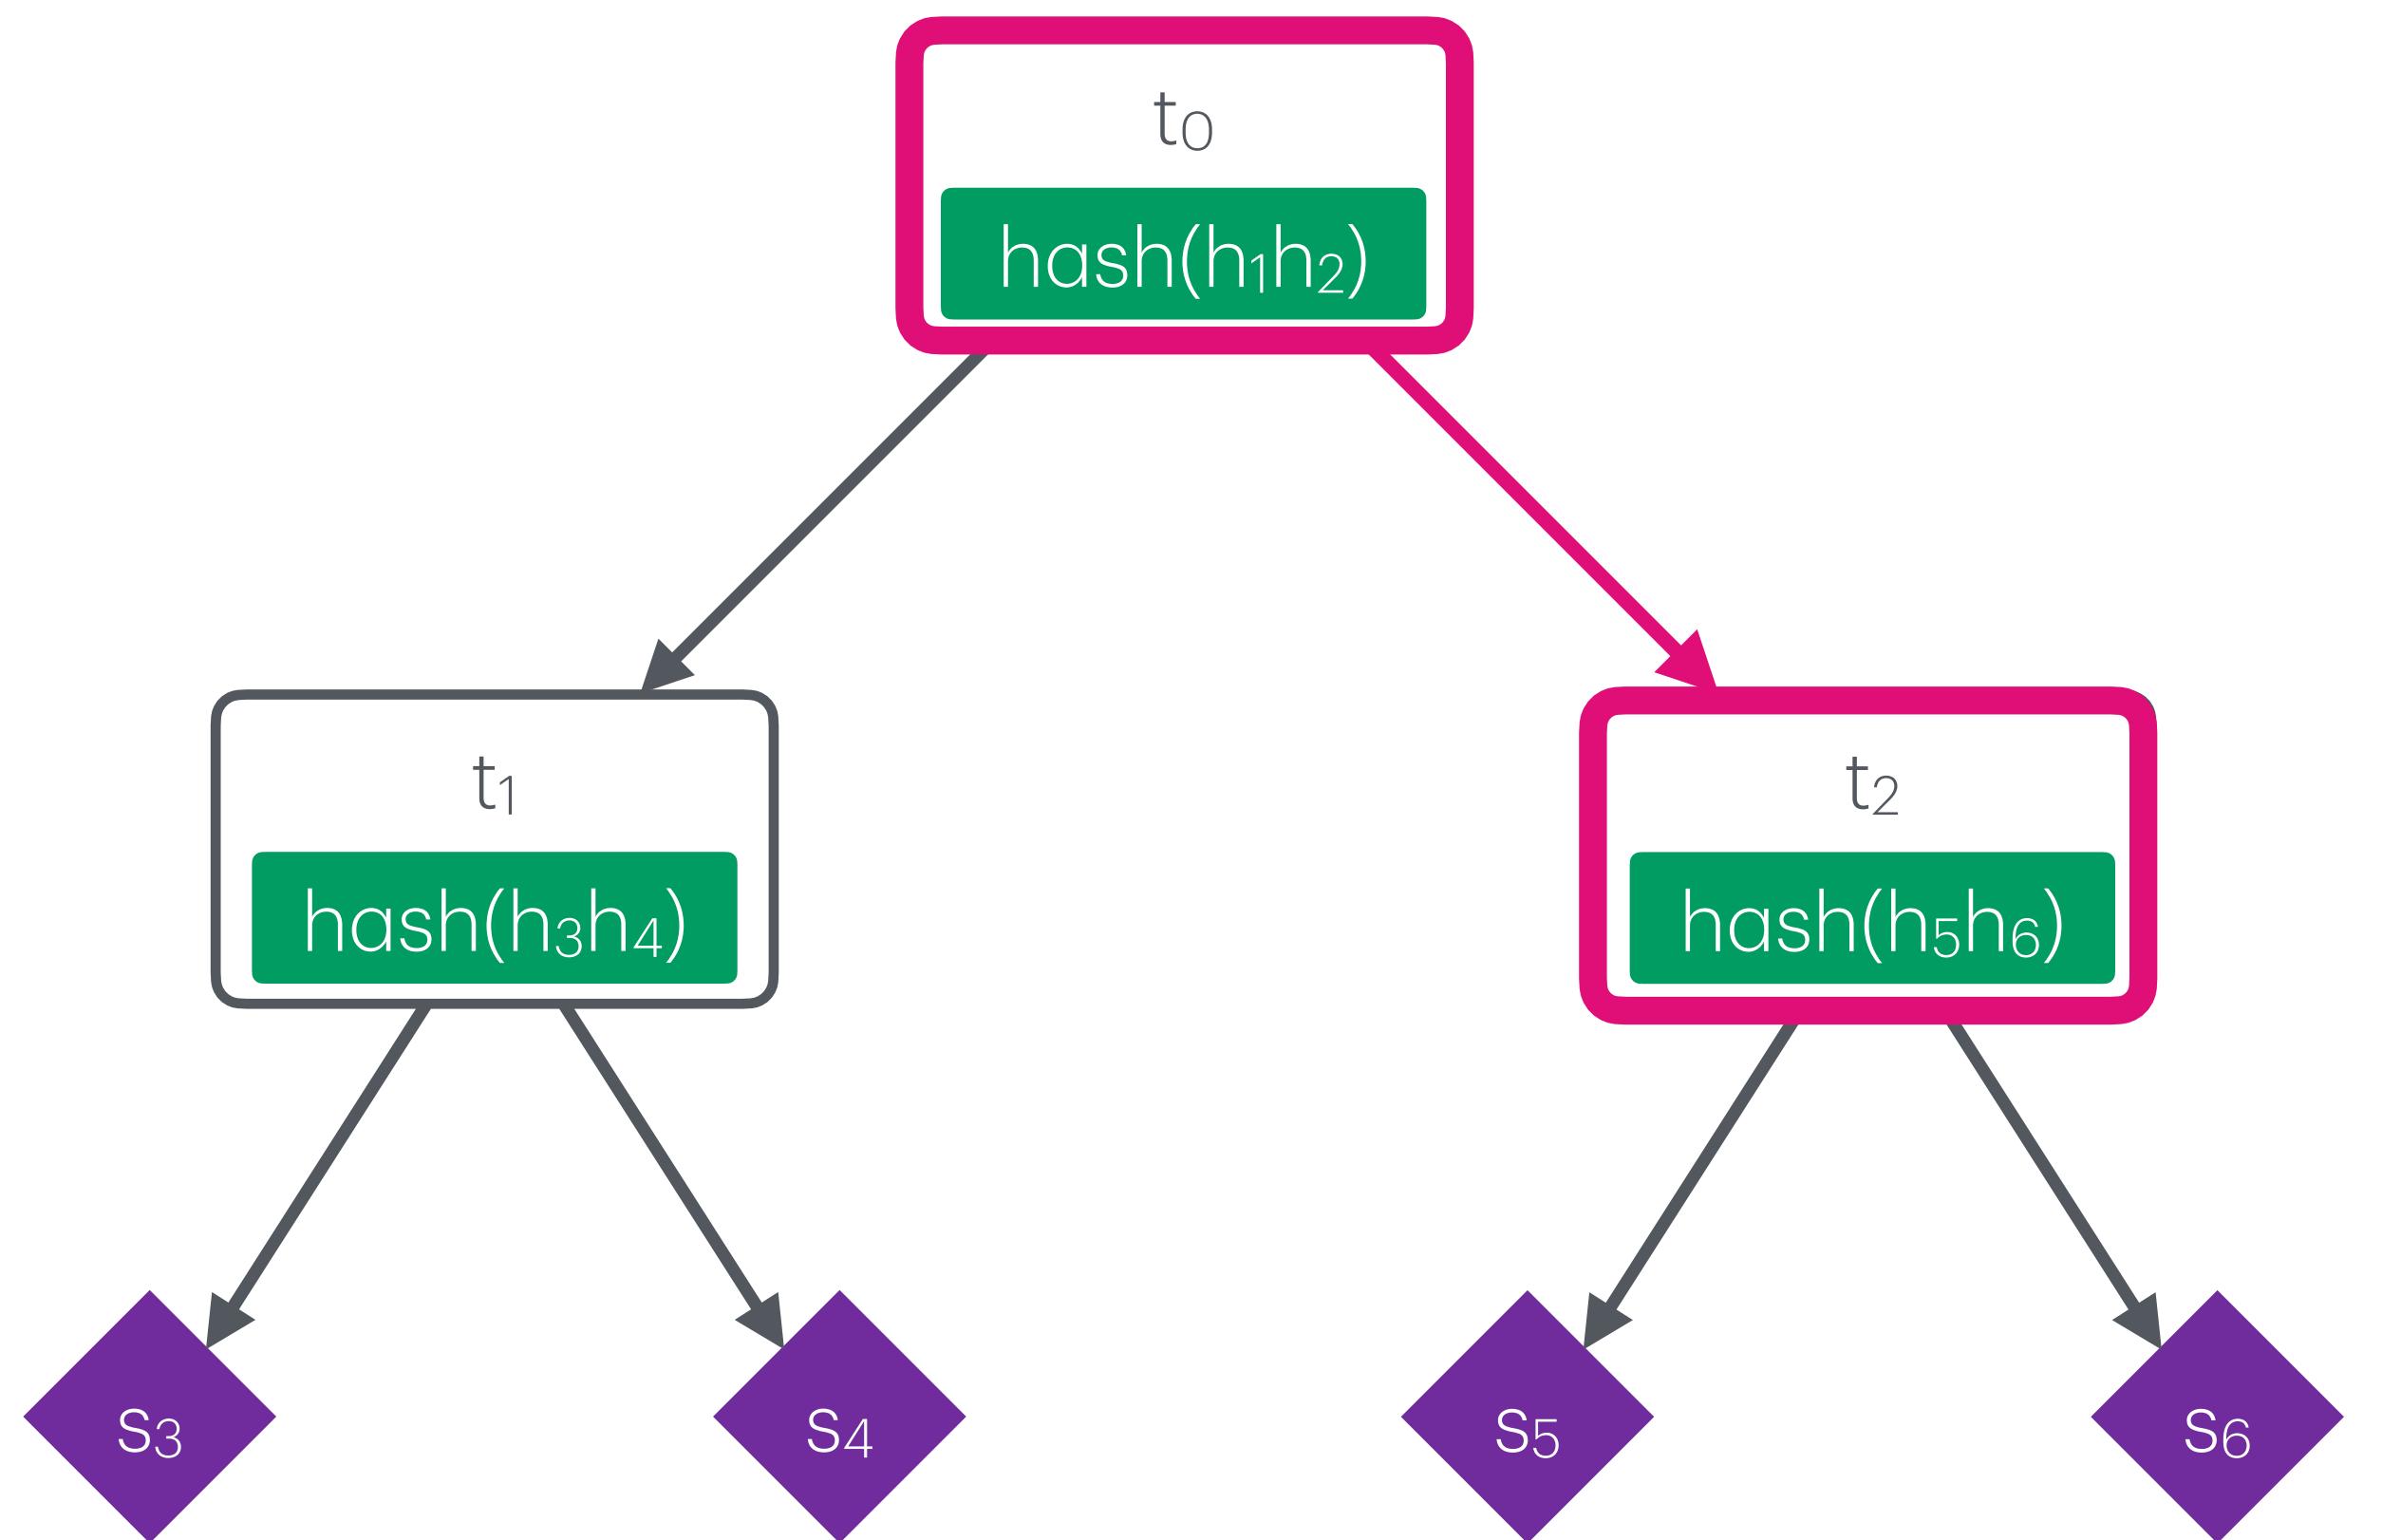
Example: Merkle Tree (Prover)



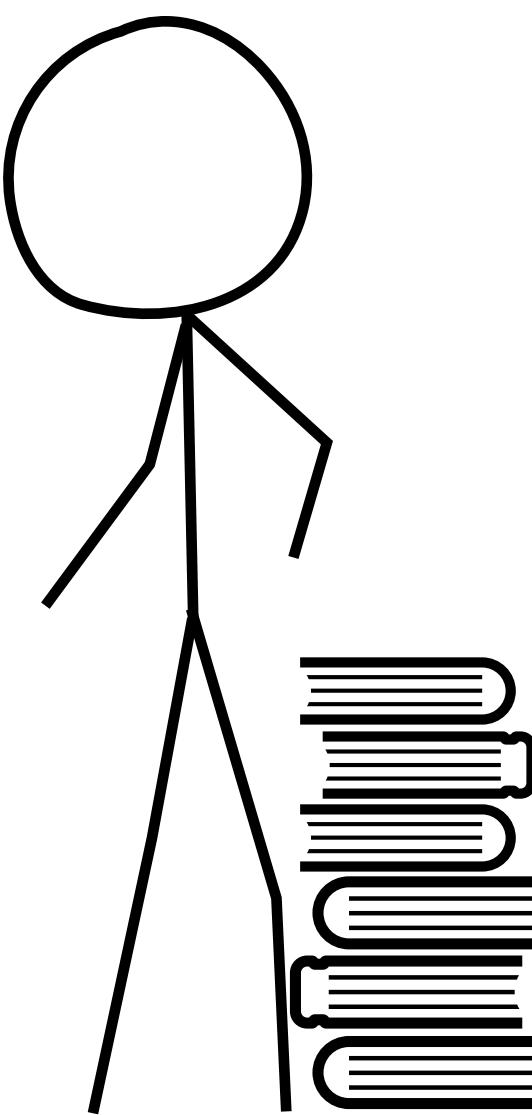
fetch([R, L], t₀) =
([h₁



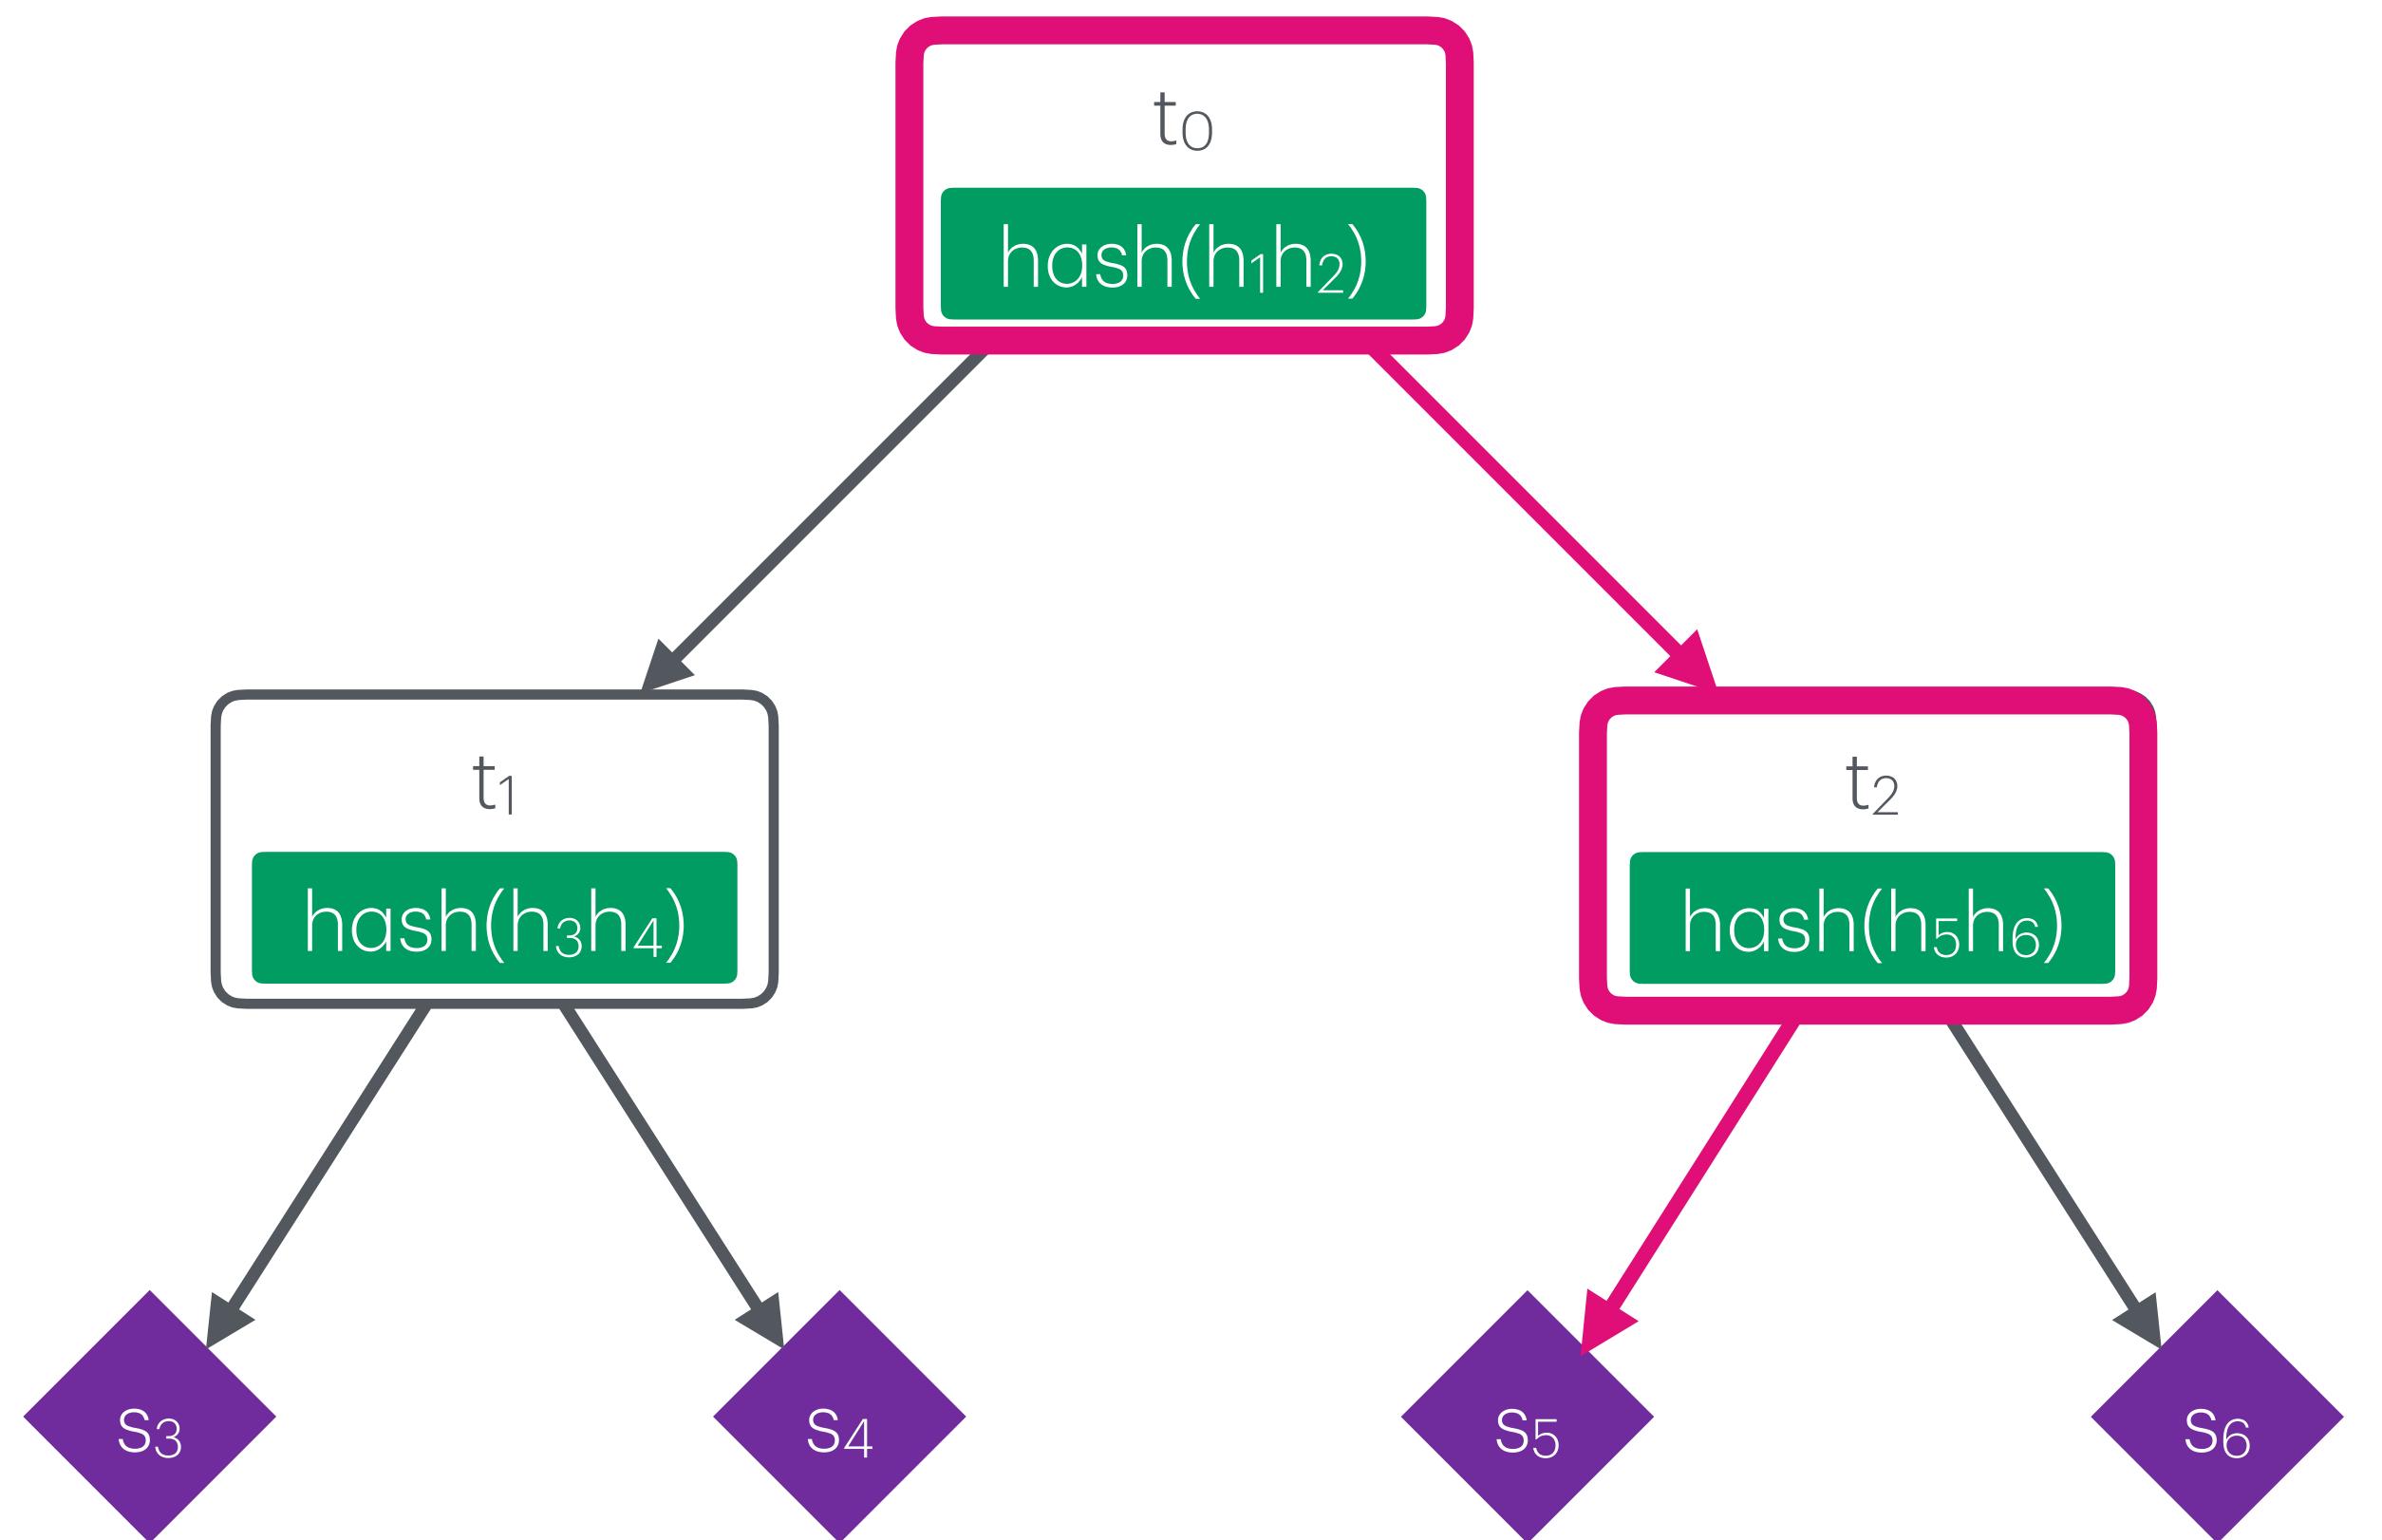
Example: Merkle Tree (Prover)



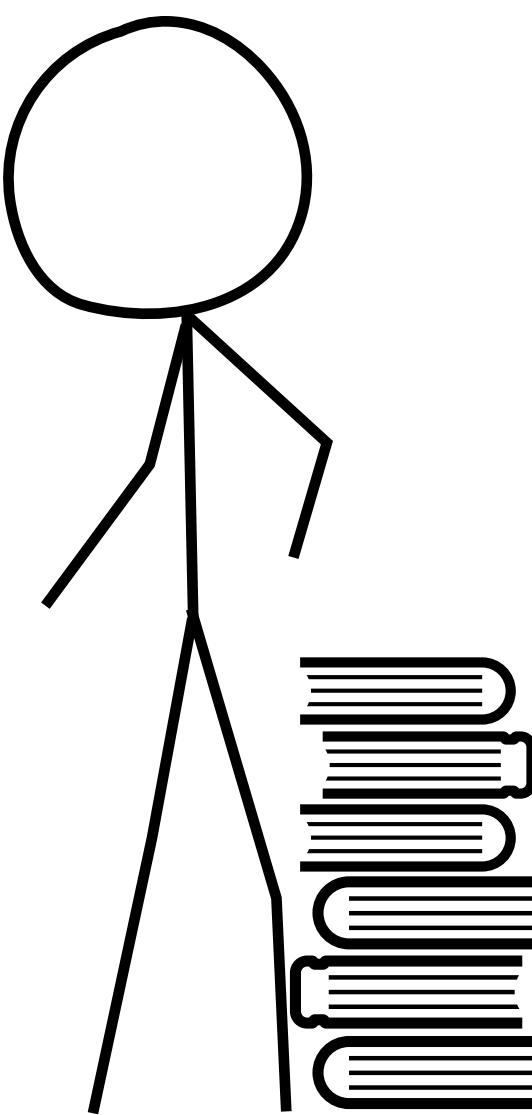
fetch([R, L], t₀) =
([h₁



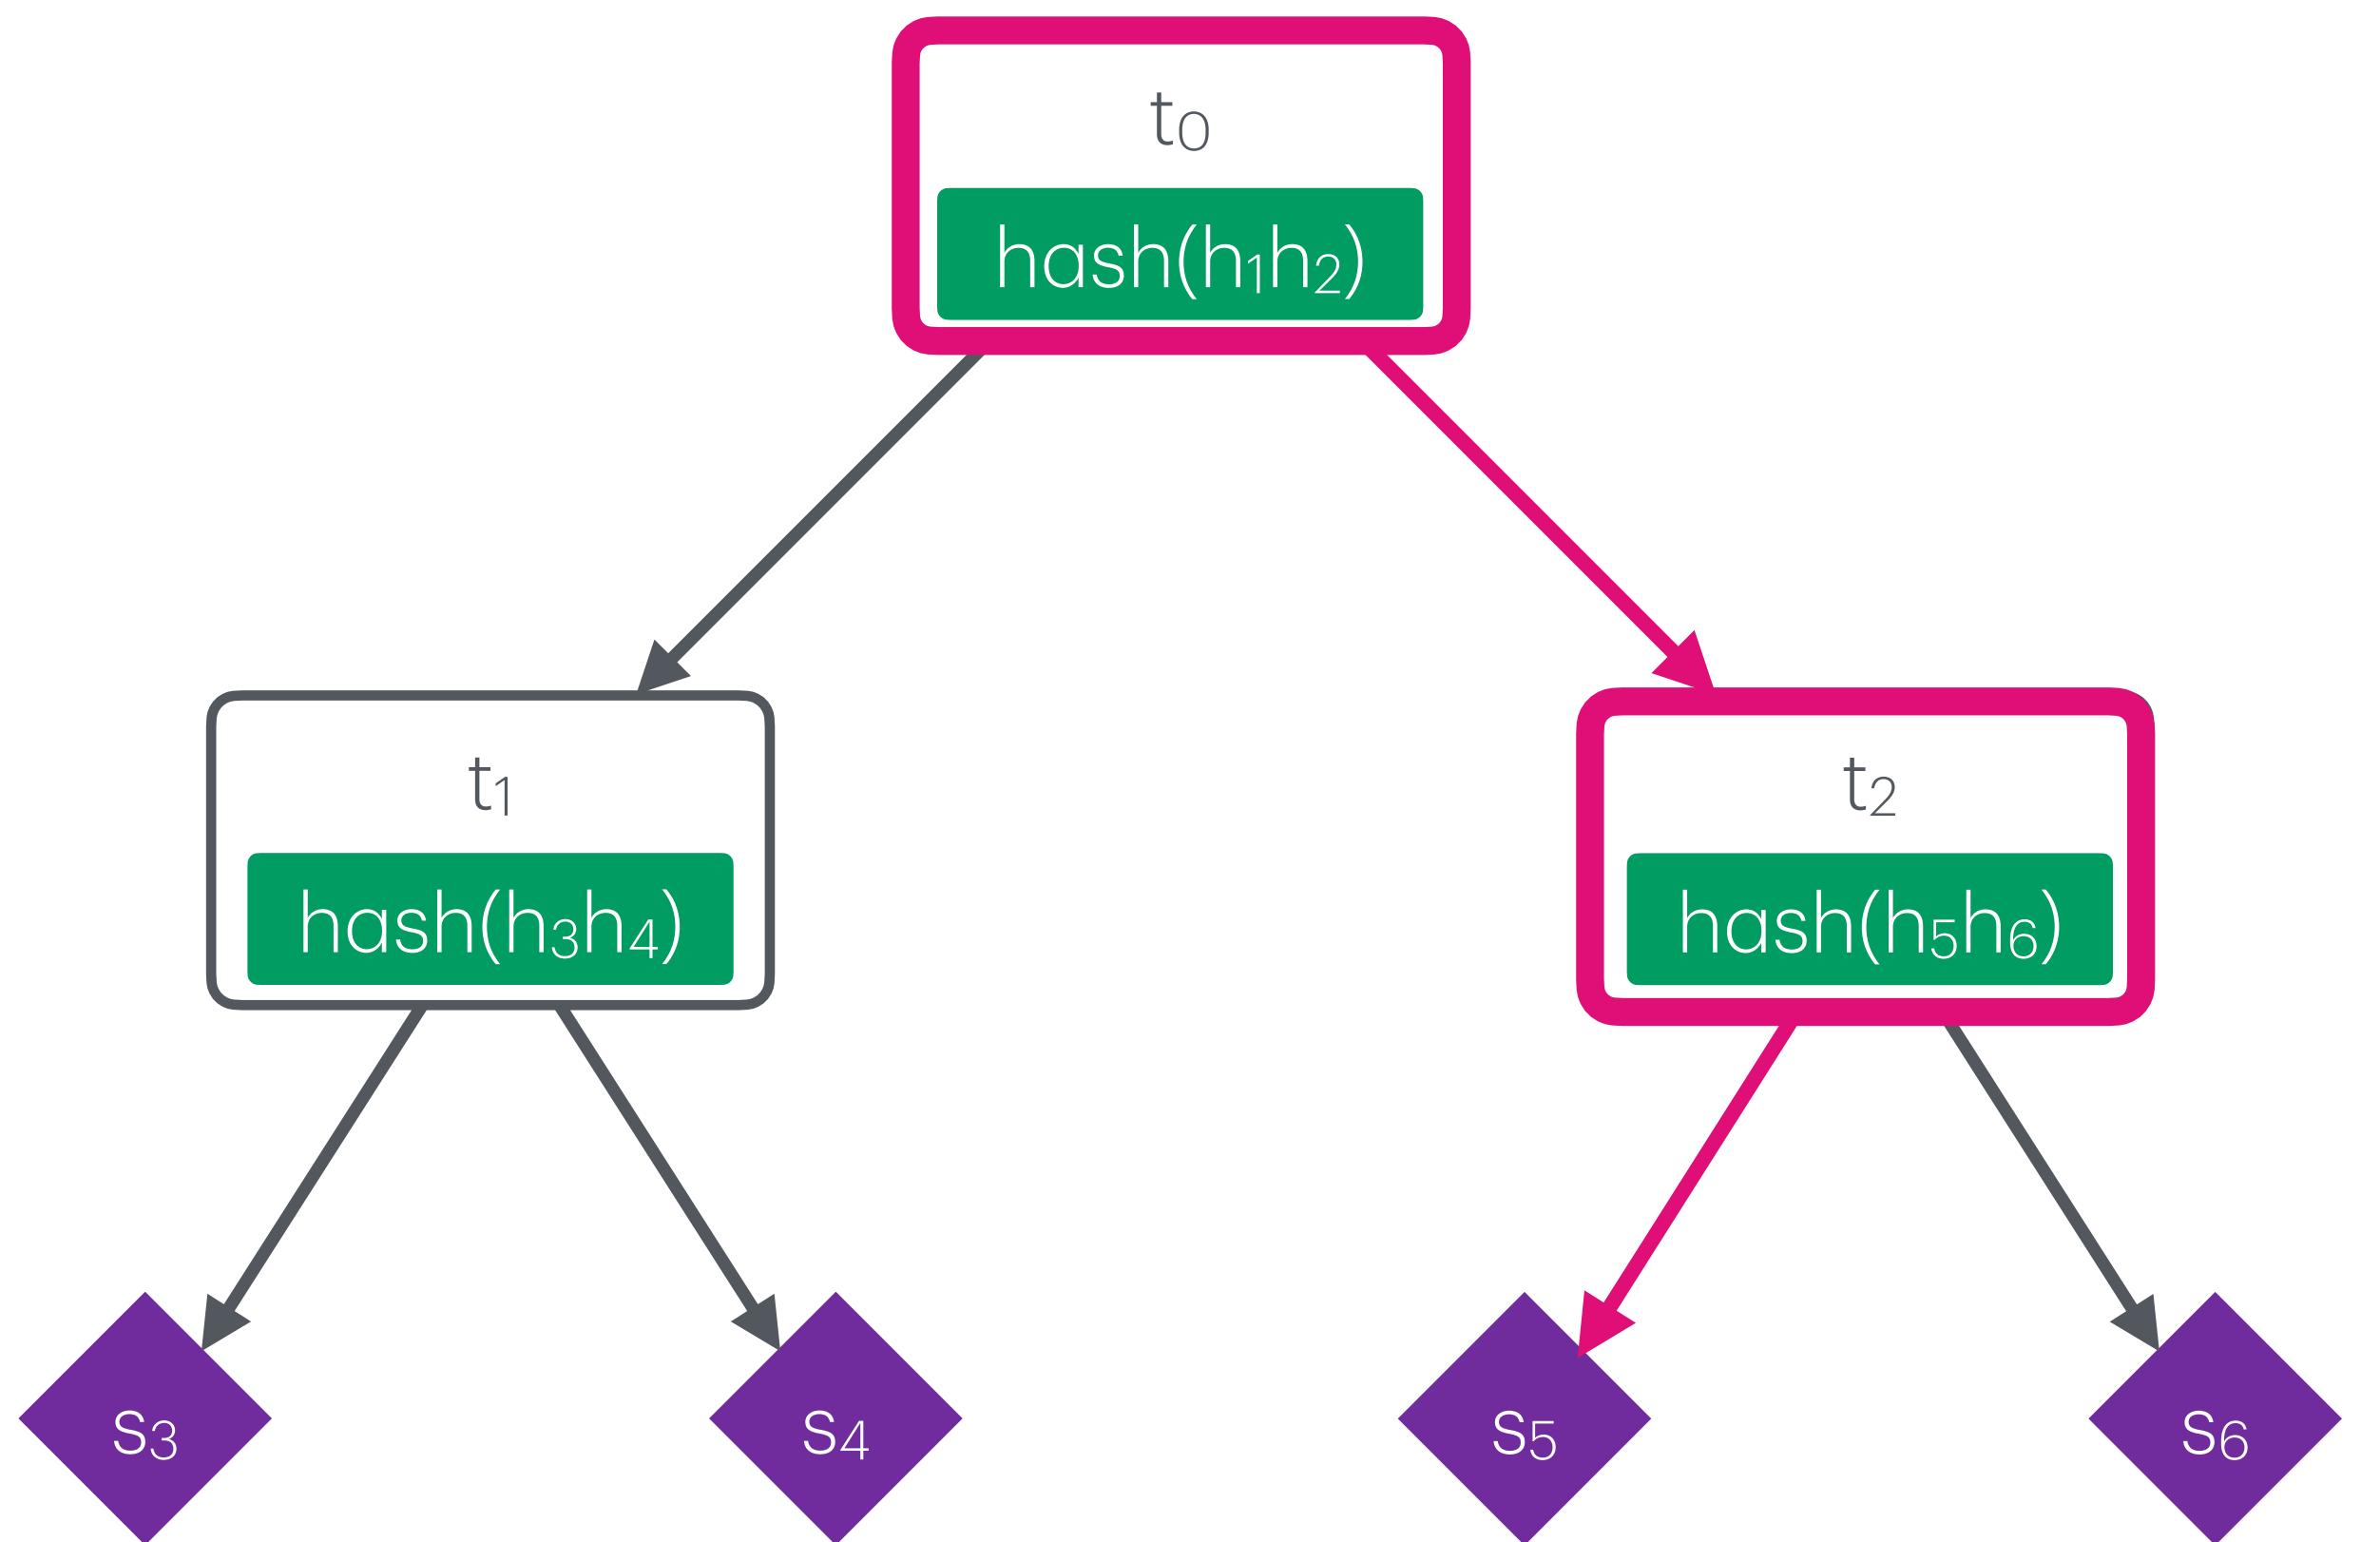
Example: Merkle Tree (Prover)



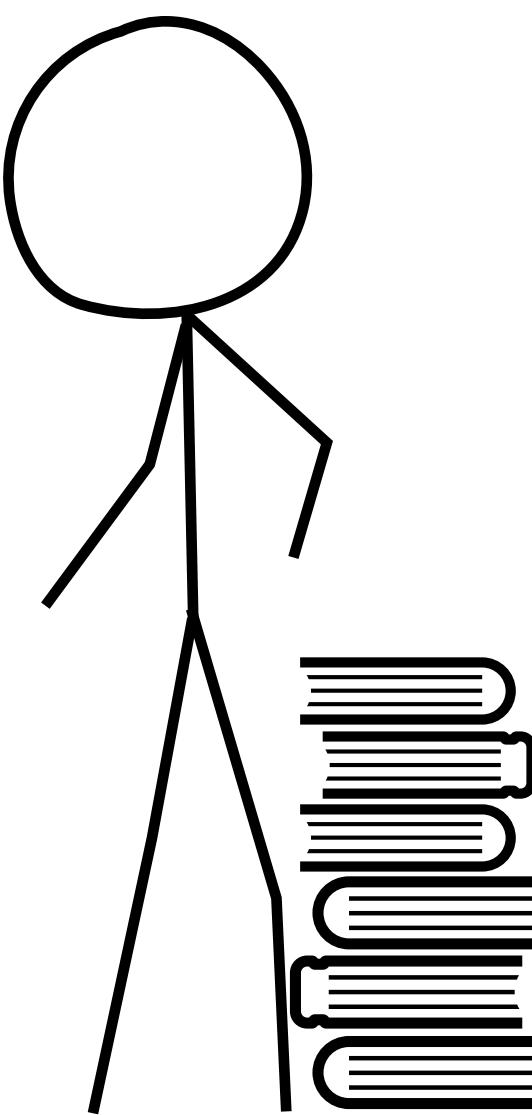
fetch([R, L], t₀) =
([h₁



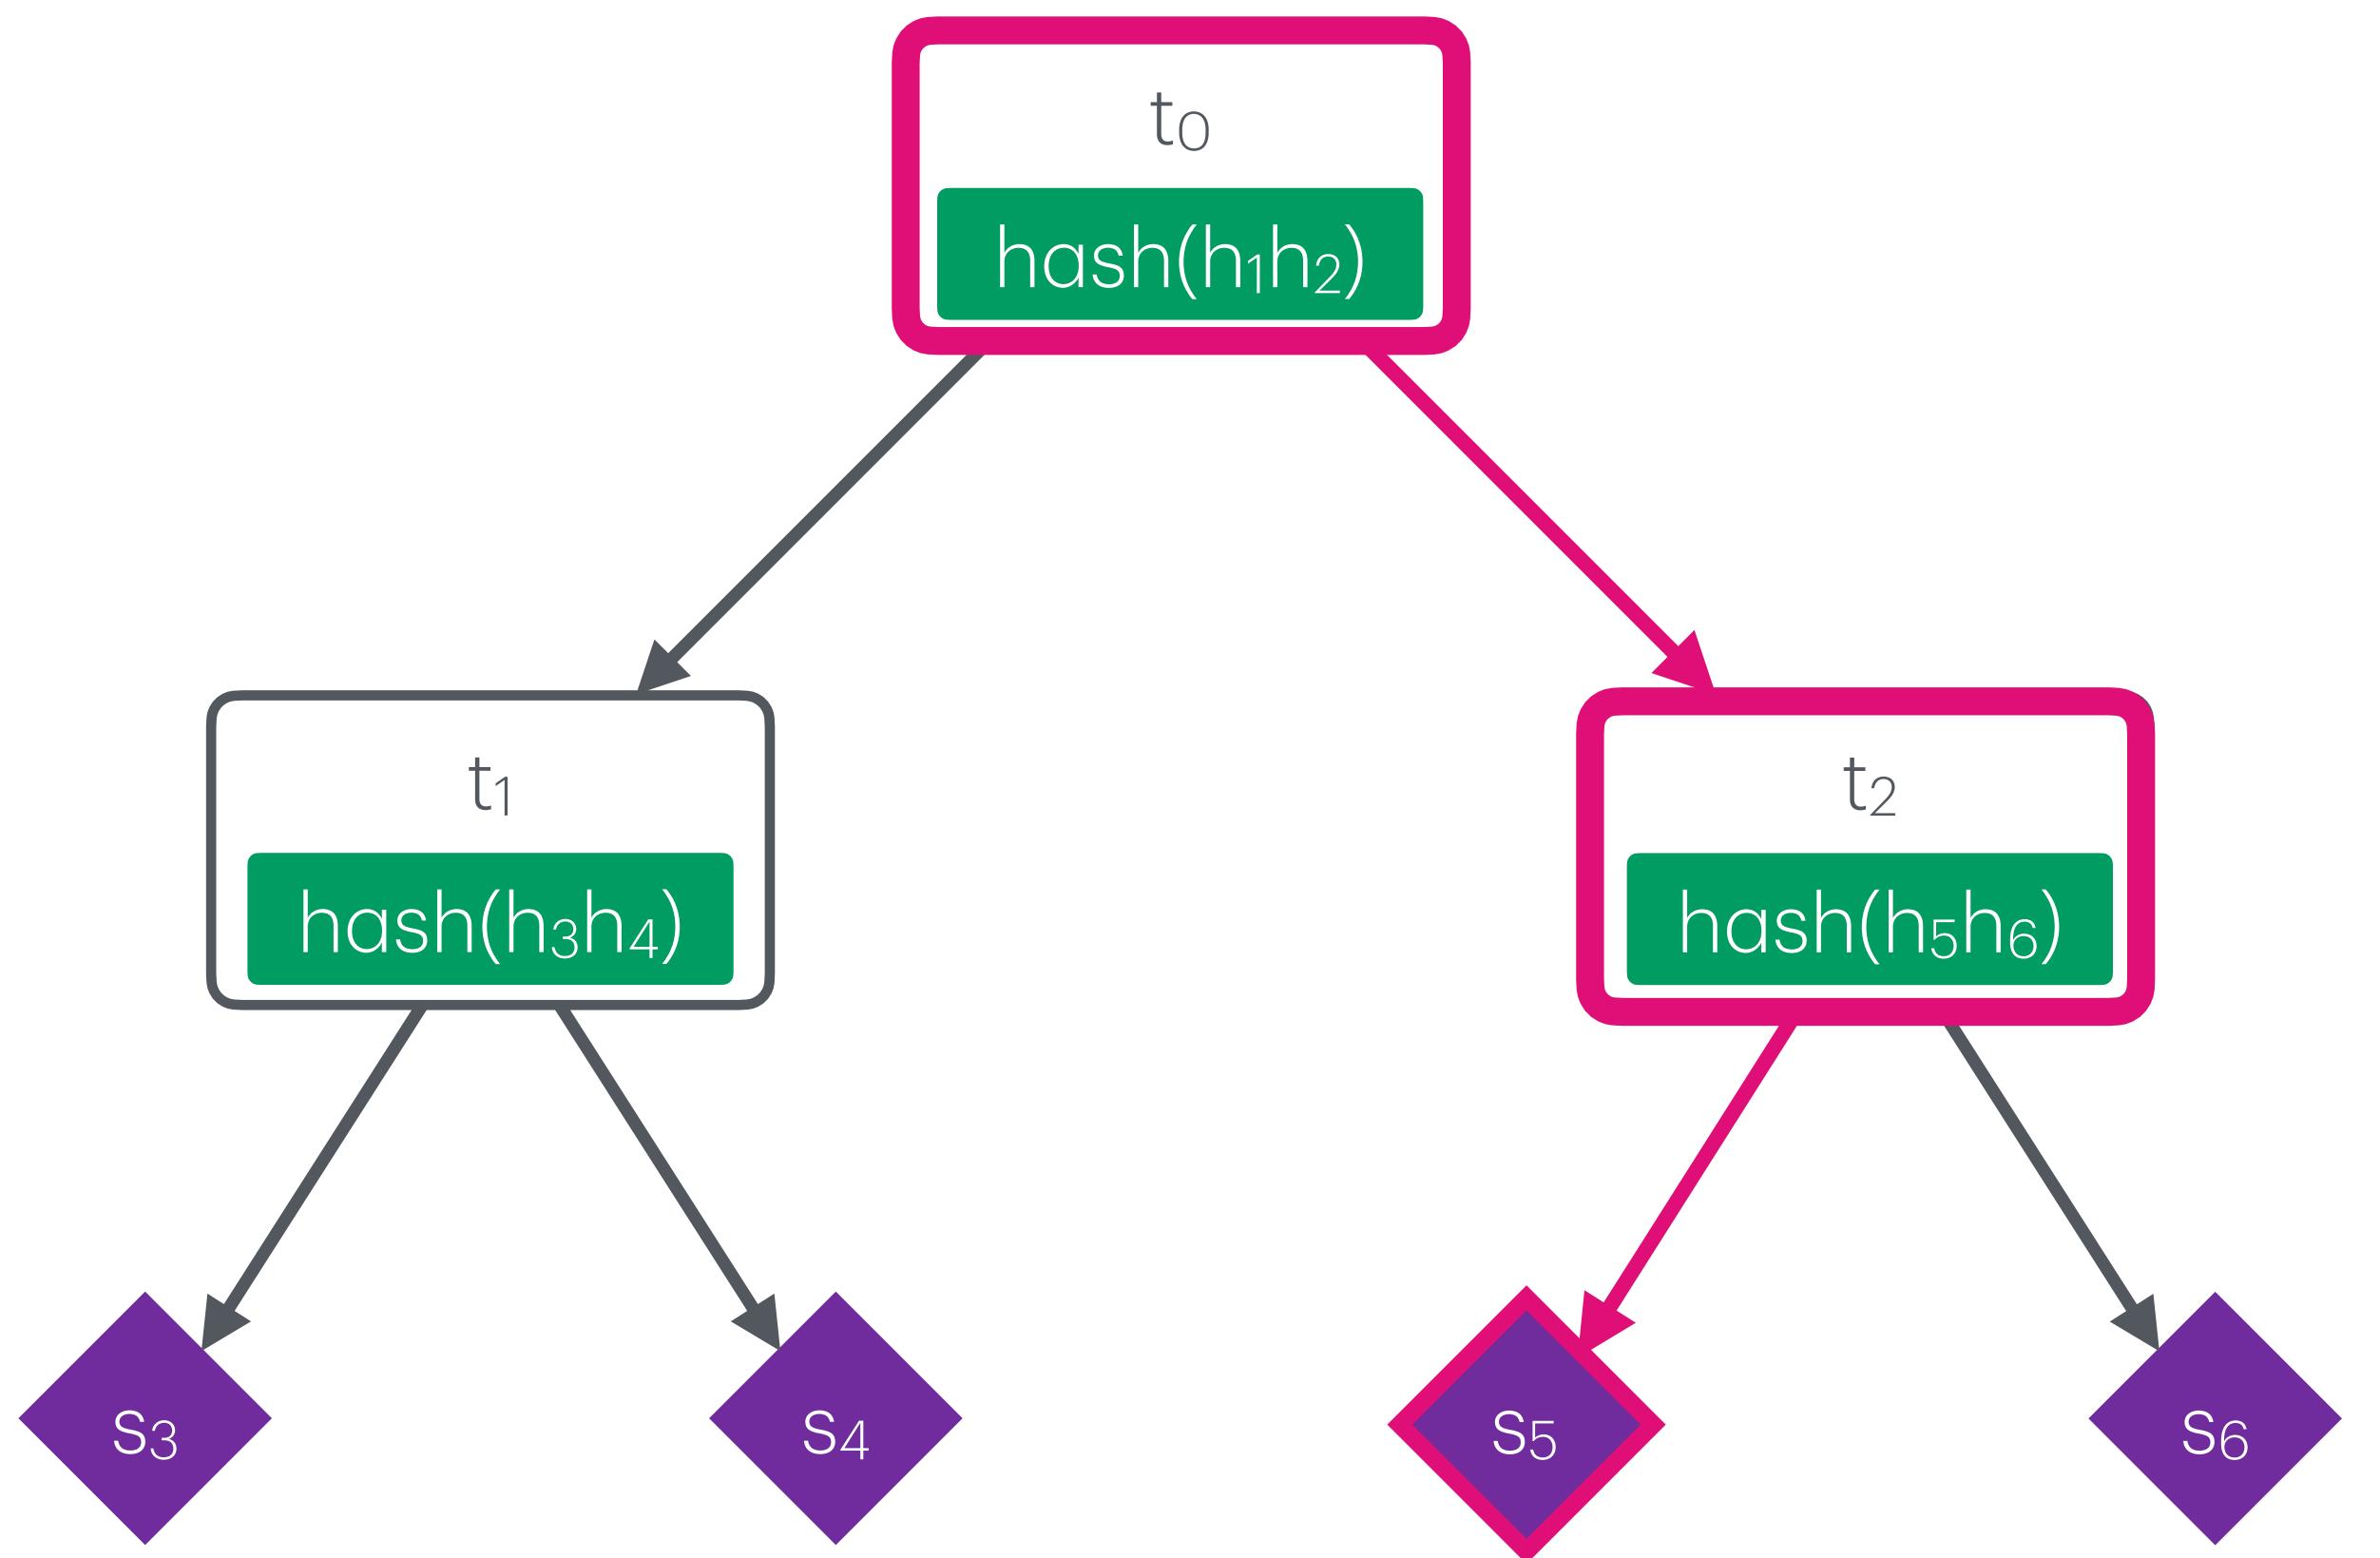
Example: Merkle Tree (Prover)



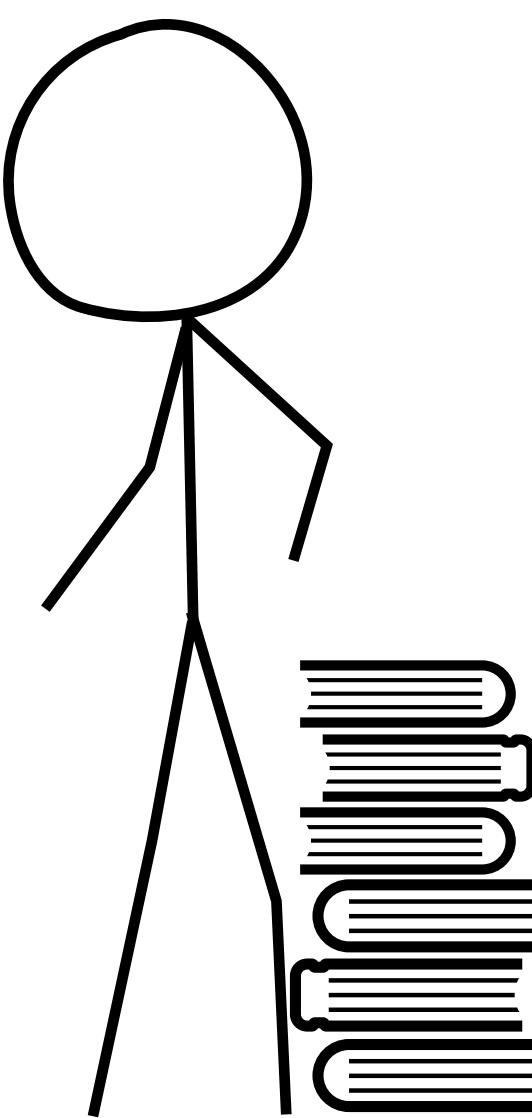
`fetch([R, L], t0) =`
`([h1, h6`



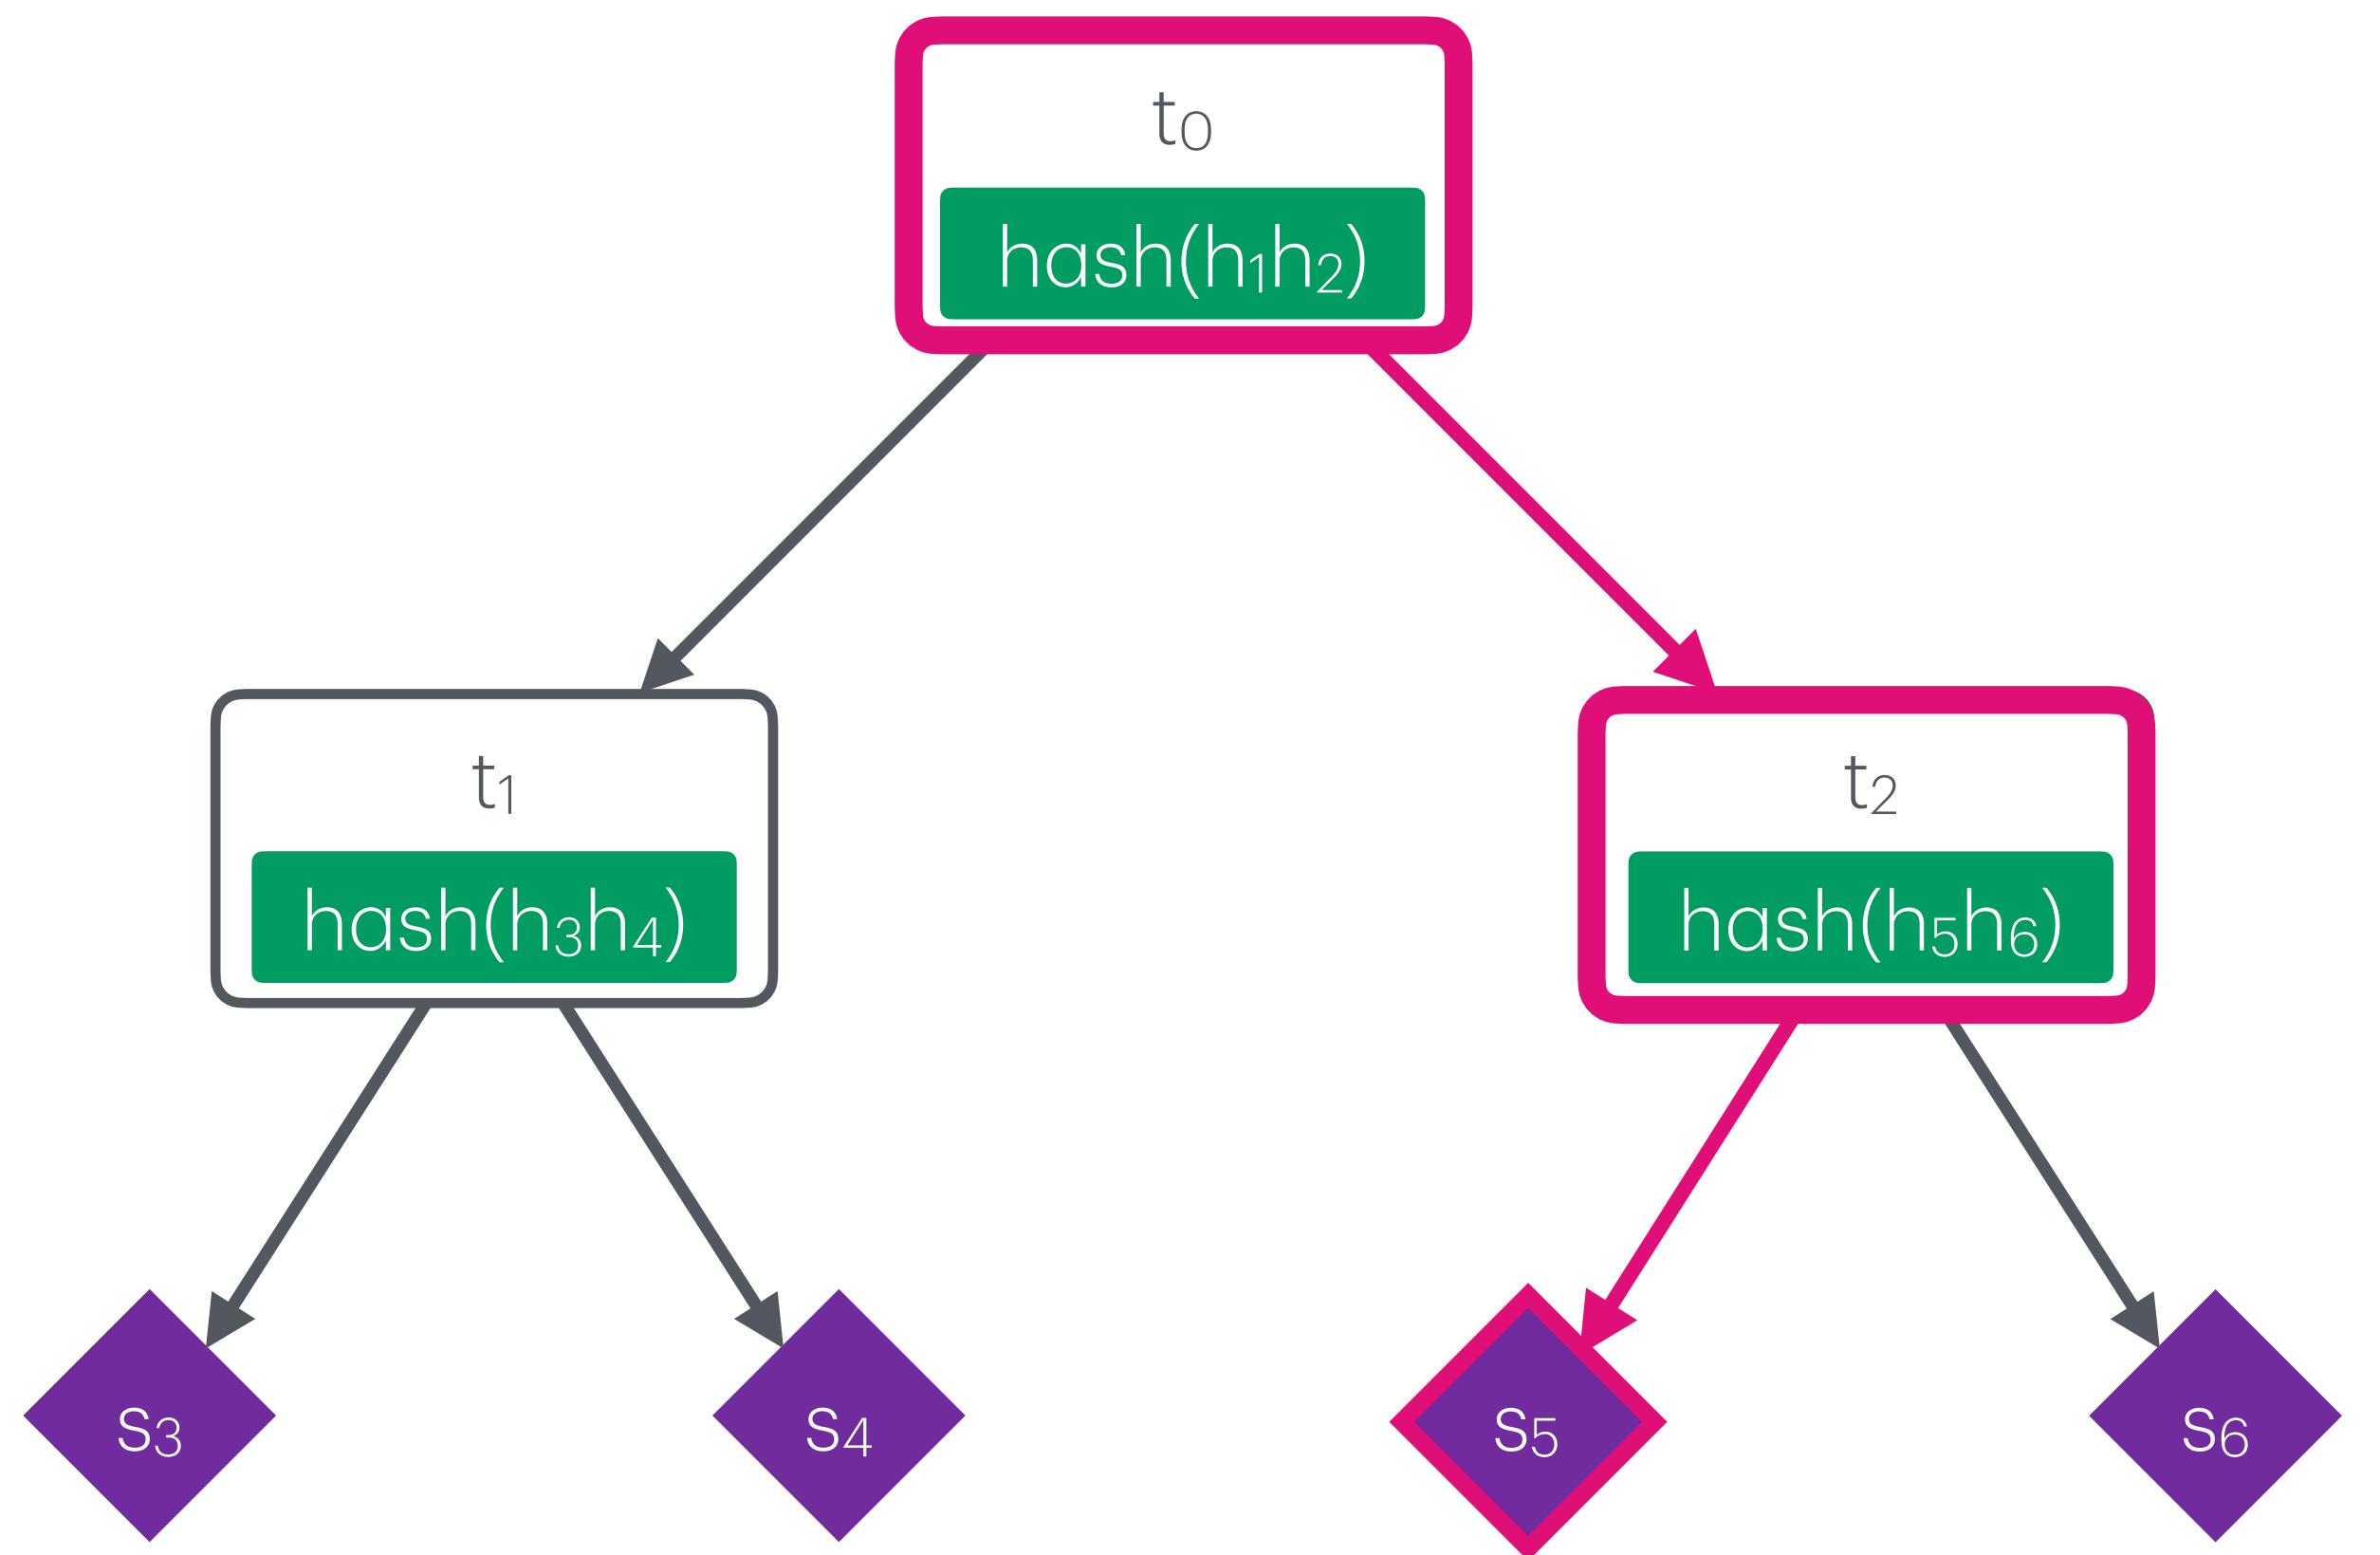
Example: Merkle Tree (Prover)



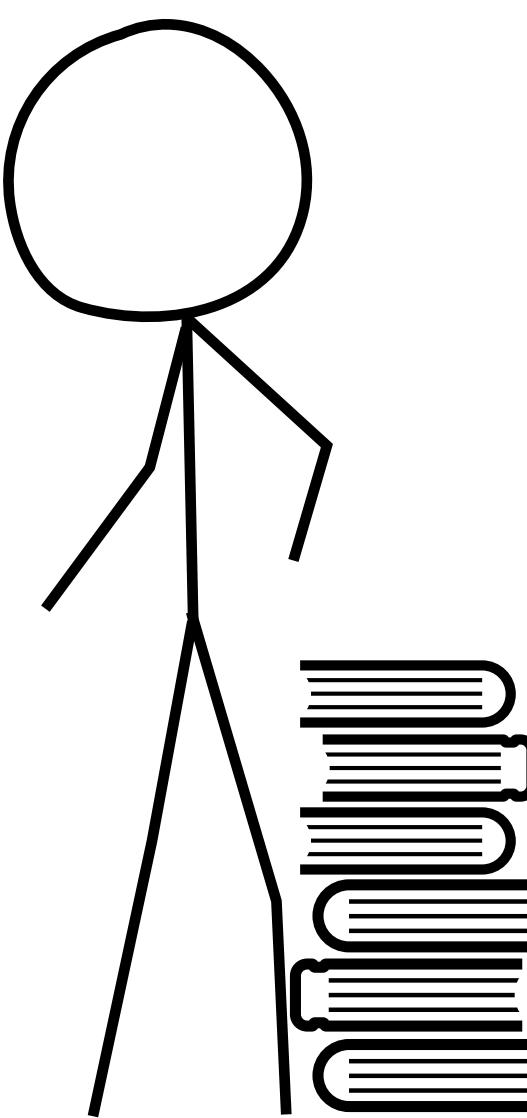
fetch([R, L], t₀) =
([h₁, h₆



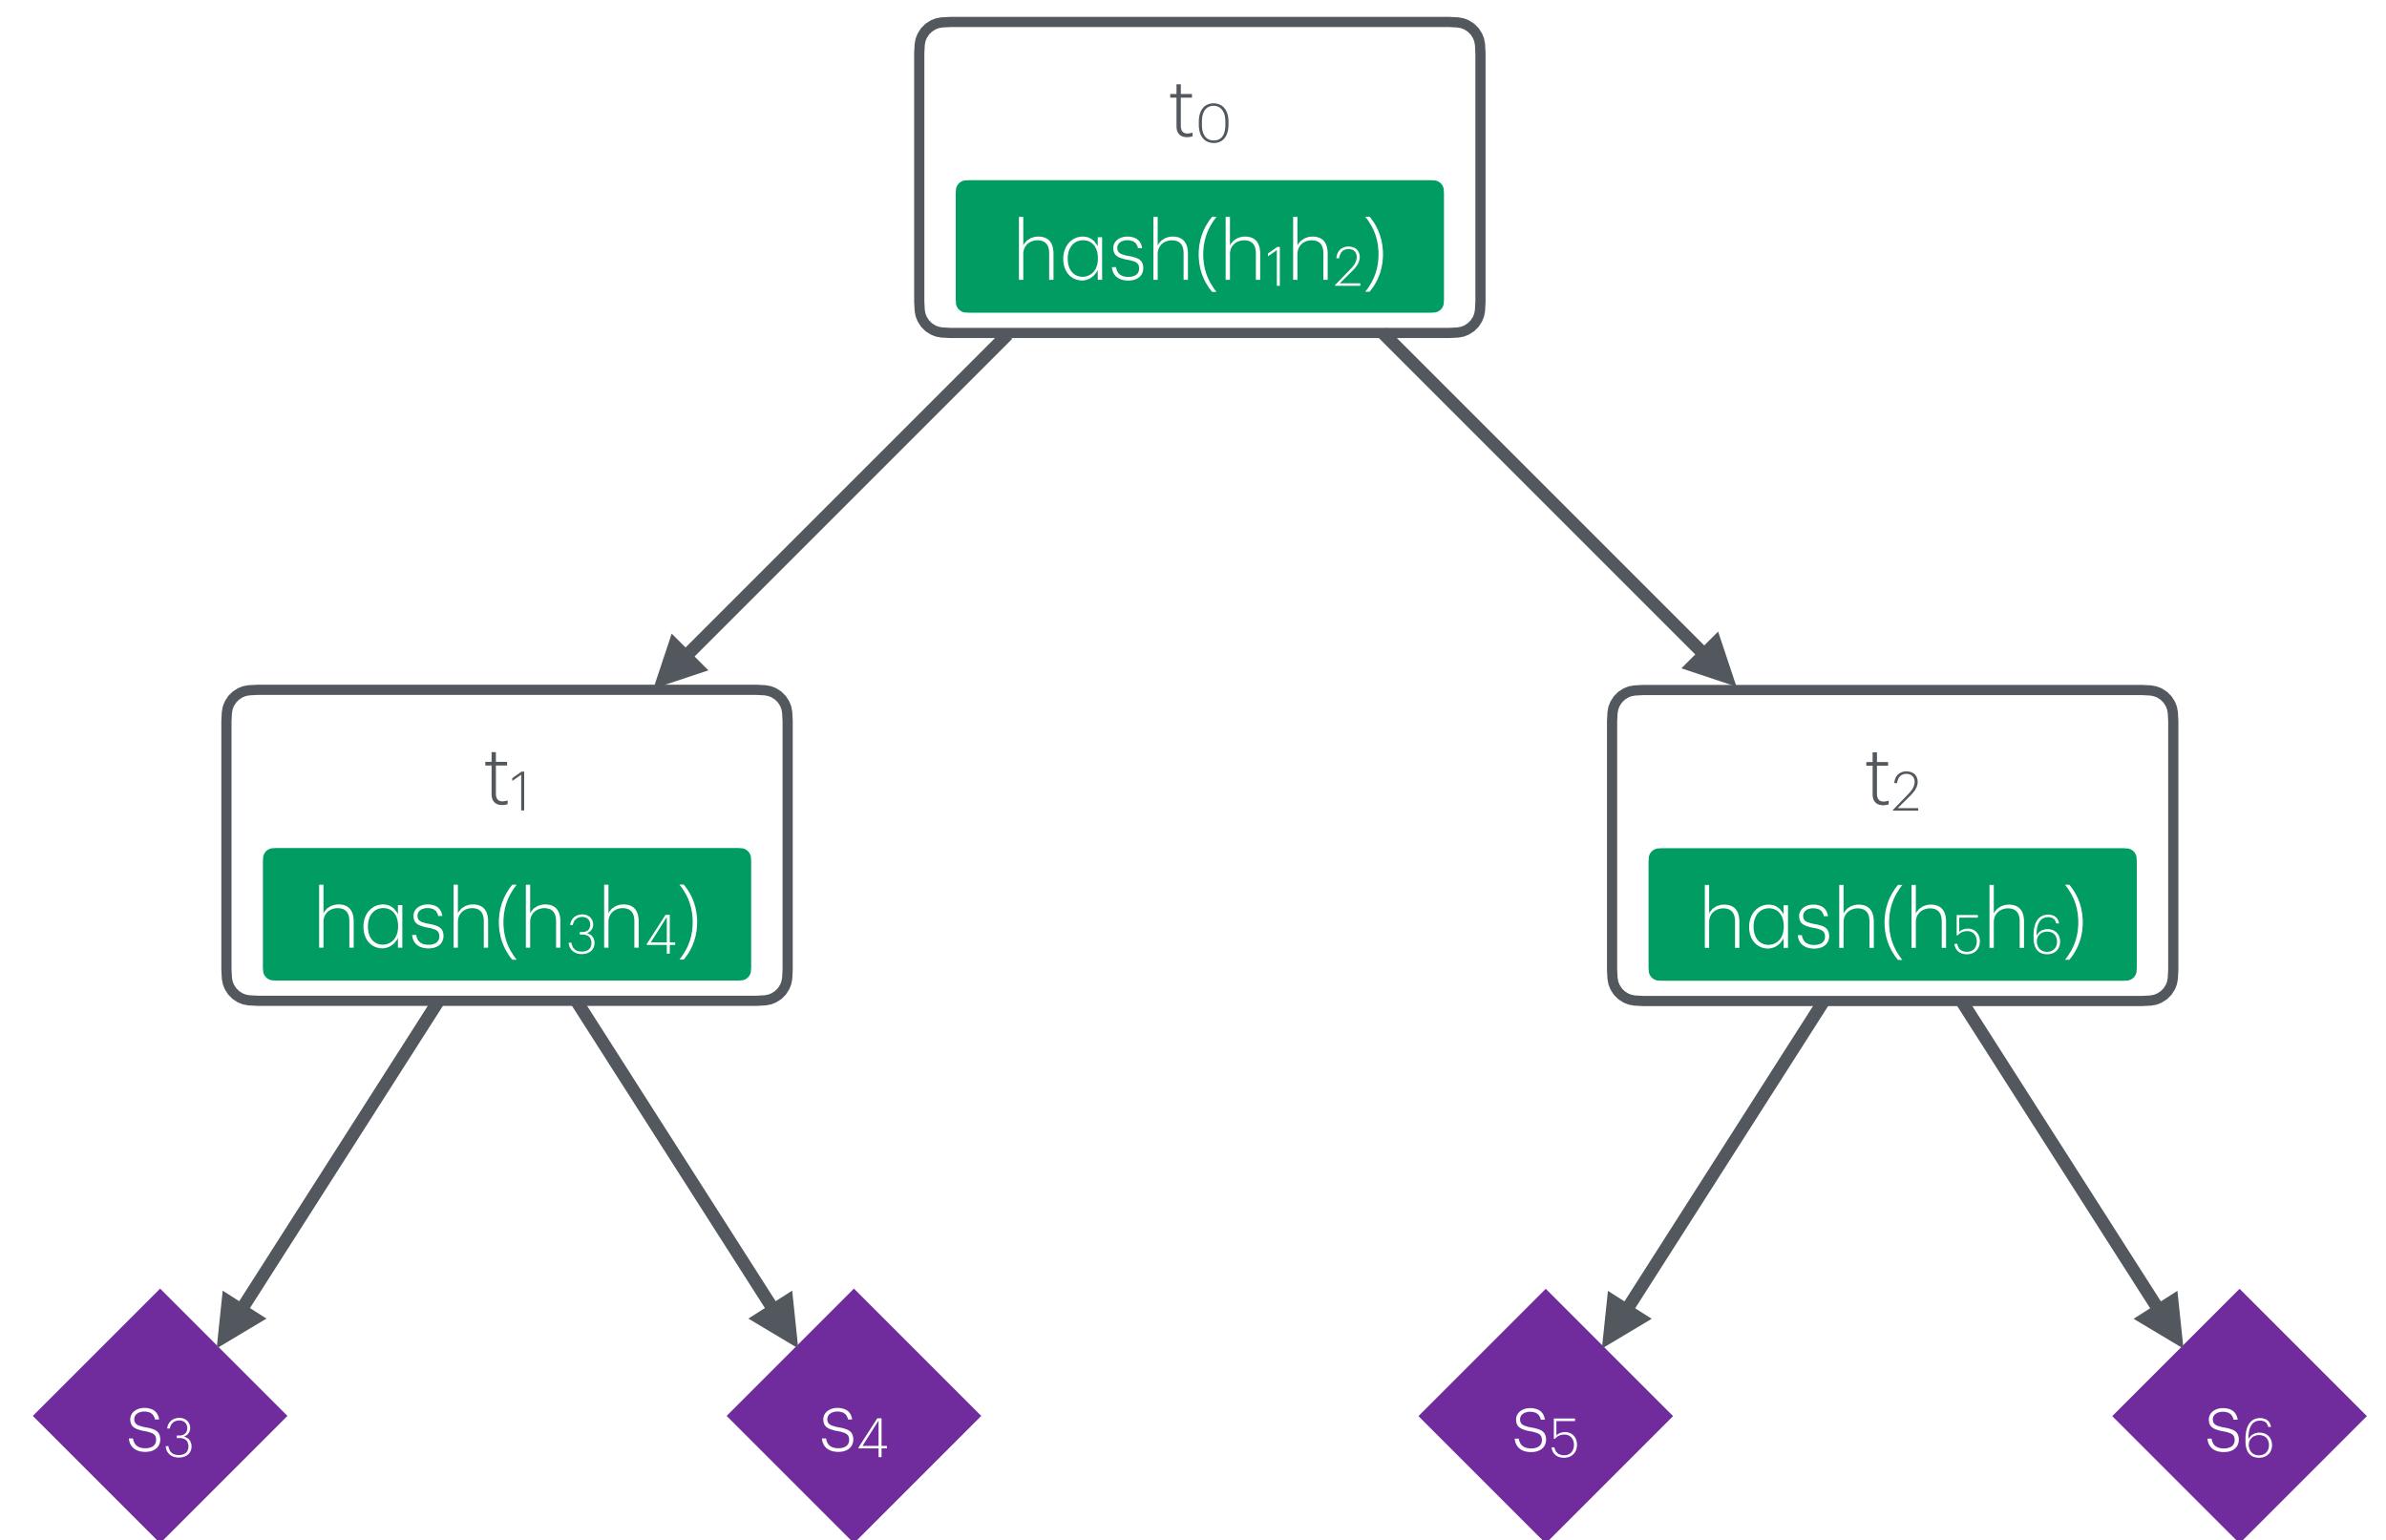
Example: Merkle Tree (Prover)



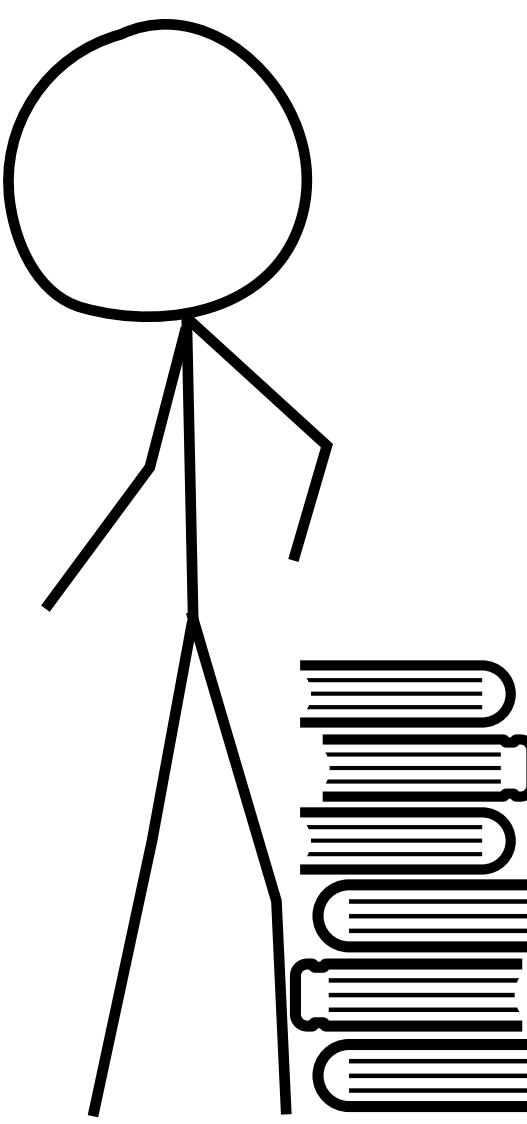
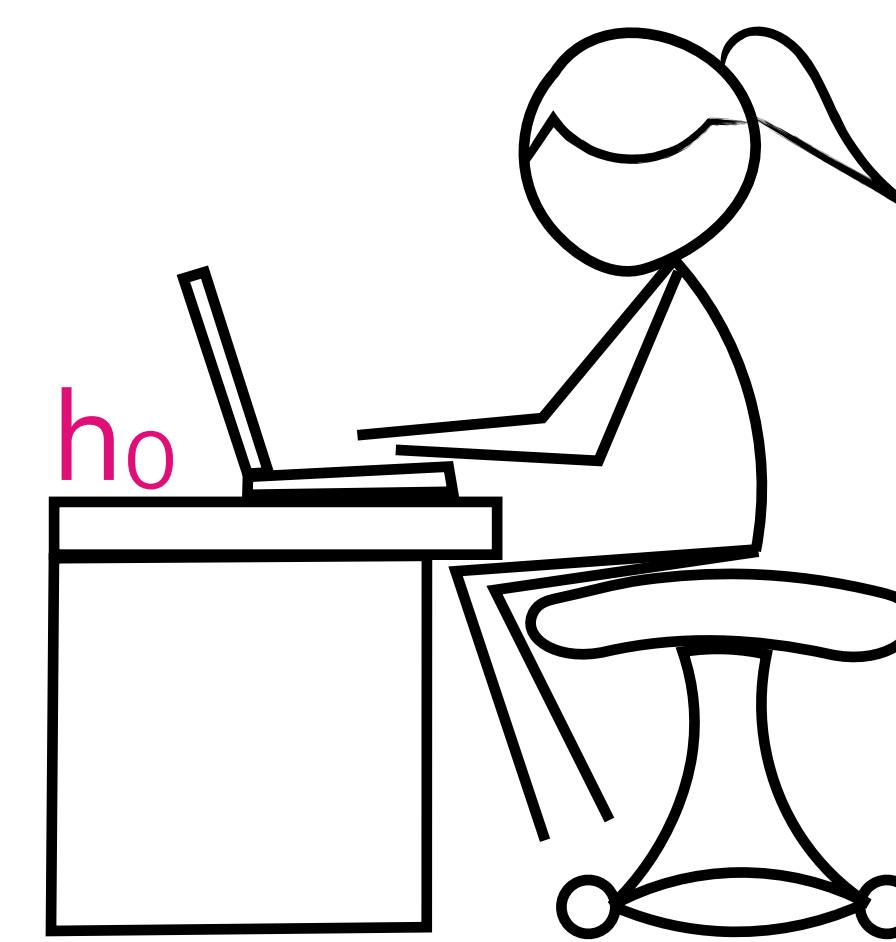
fetch([R, L], t₀) =
([h₁, h₆, s₅], s₅)



Example: Merkle Tree (Verifier)

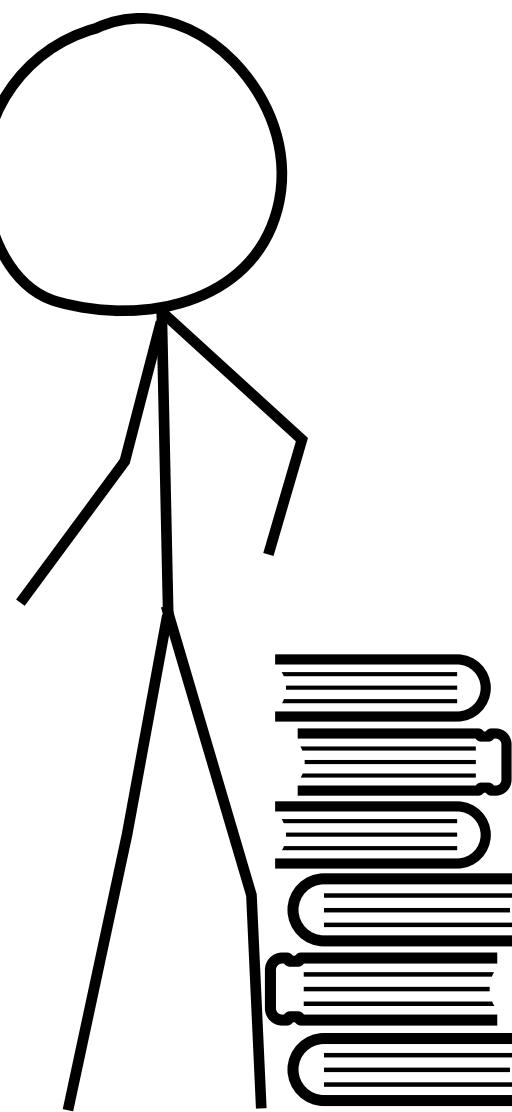


`fetch([R, L], t0) =
([h1, h6, s5], s5)`



Example: Merkle Tree (Verifier)

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



Example: Merkle Tree (Verifier)

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

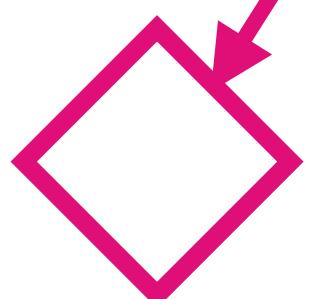
h_1



$h_2 = \text{hash}(h_5 h_6)$

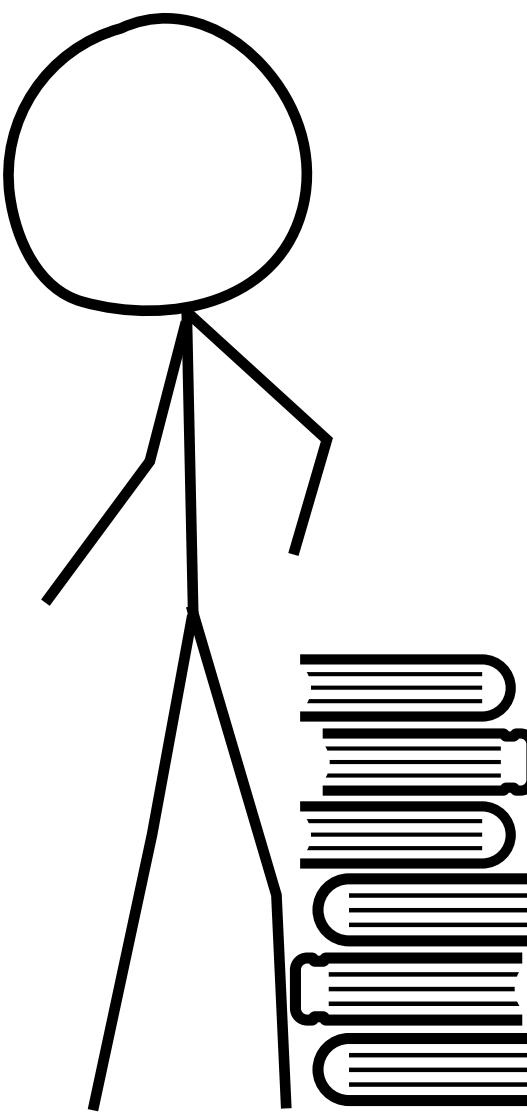


$h_5 = \text{hash}(s_5)$



h_6

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



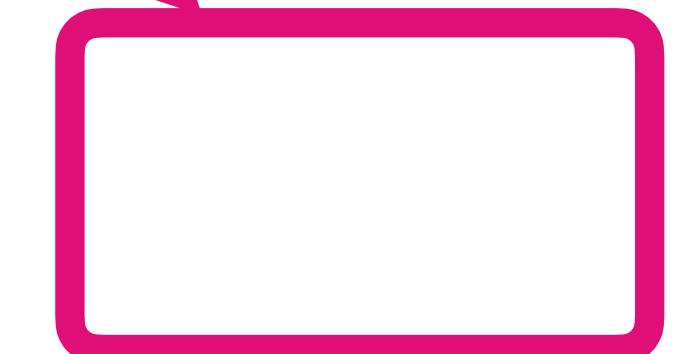
Example: Merkle Tree (Verifier)

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

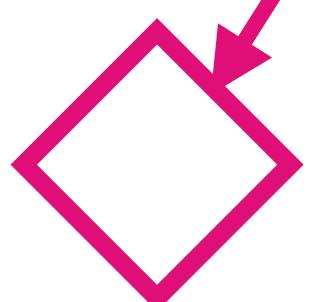
h_1



$h_2 = \text{hash}(h_5 h_6)$

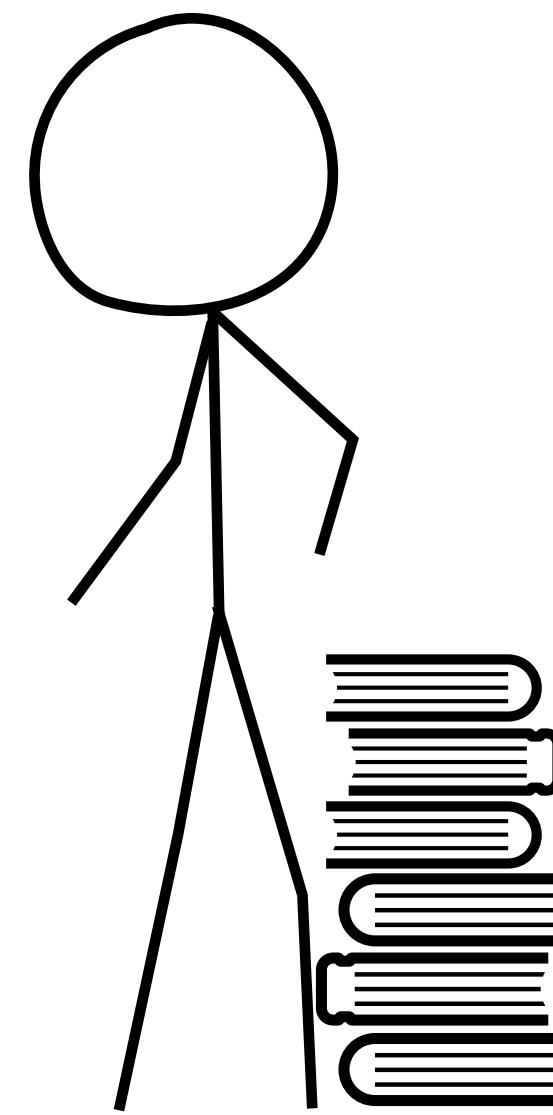
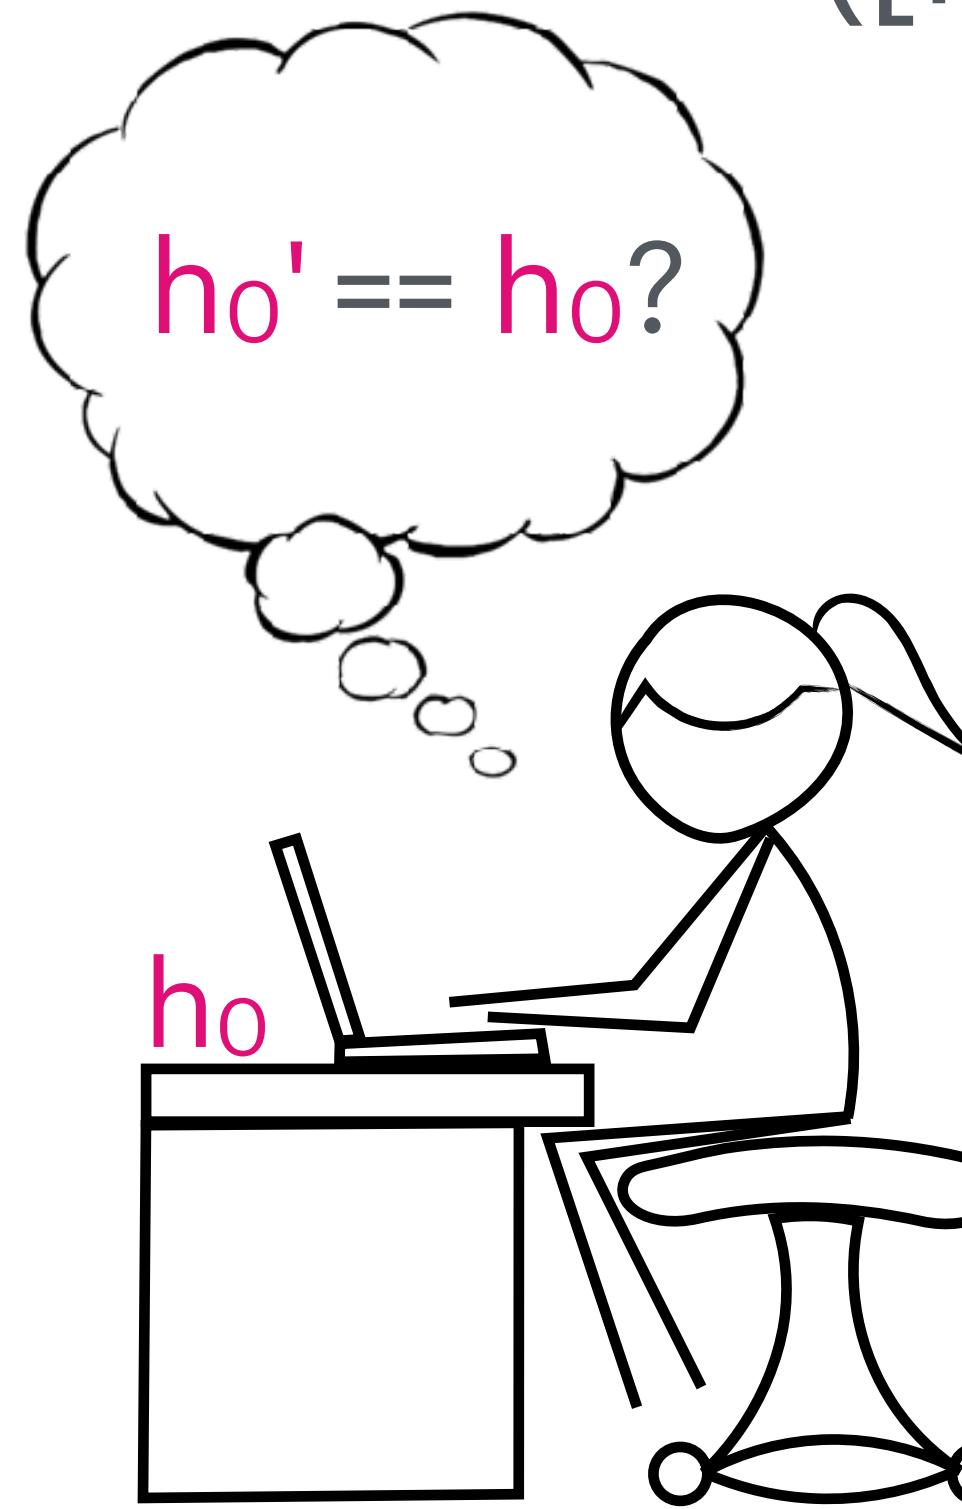


$h_5 = \text{hash}(s_5)$



h_6

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



Use cases

- **Certificate transparency**

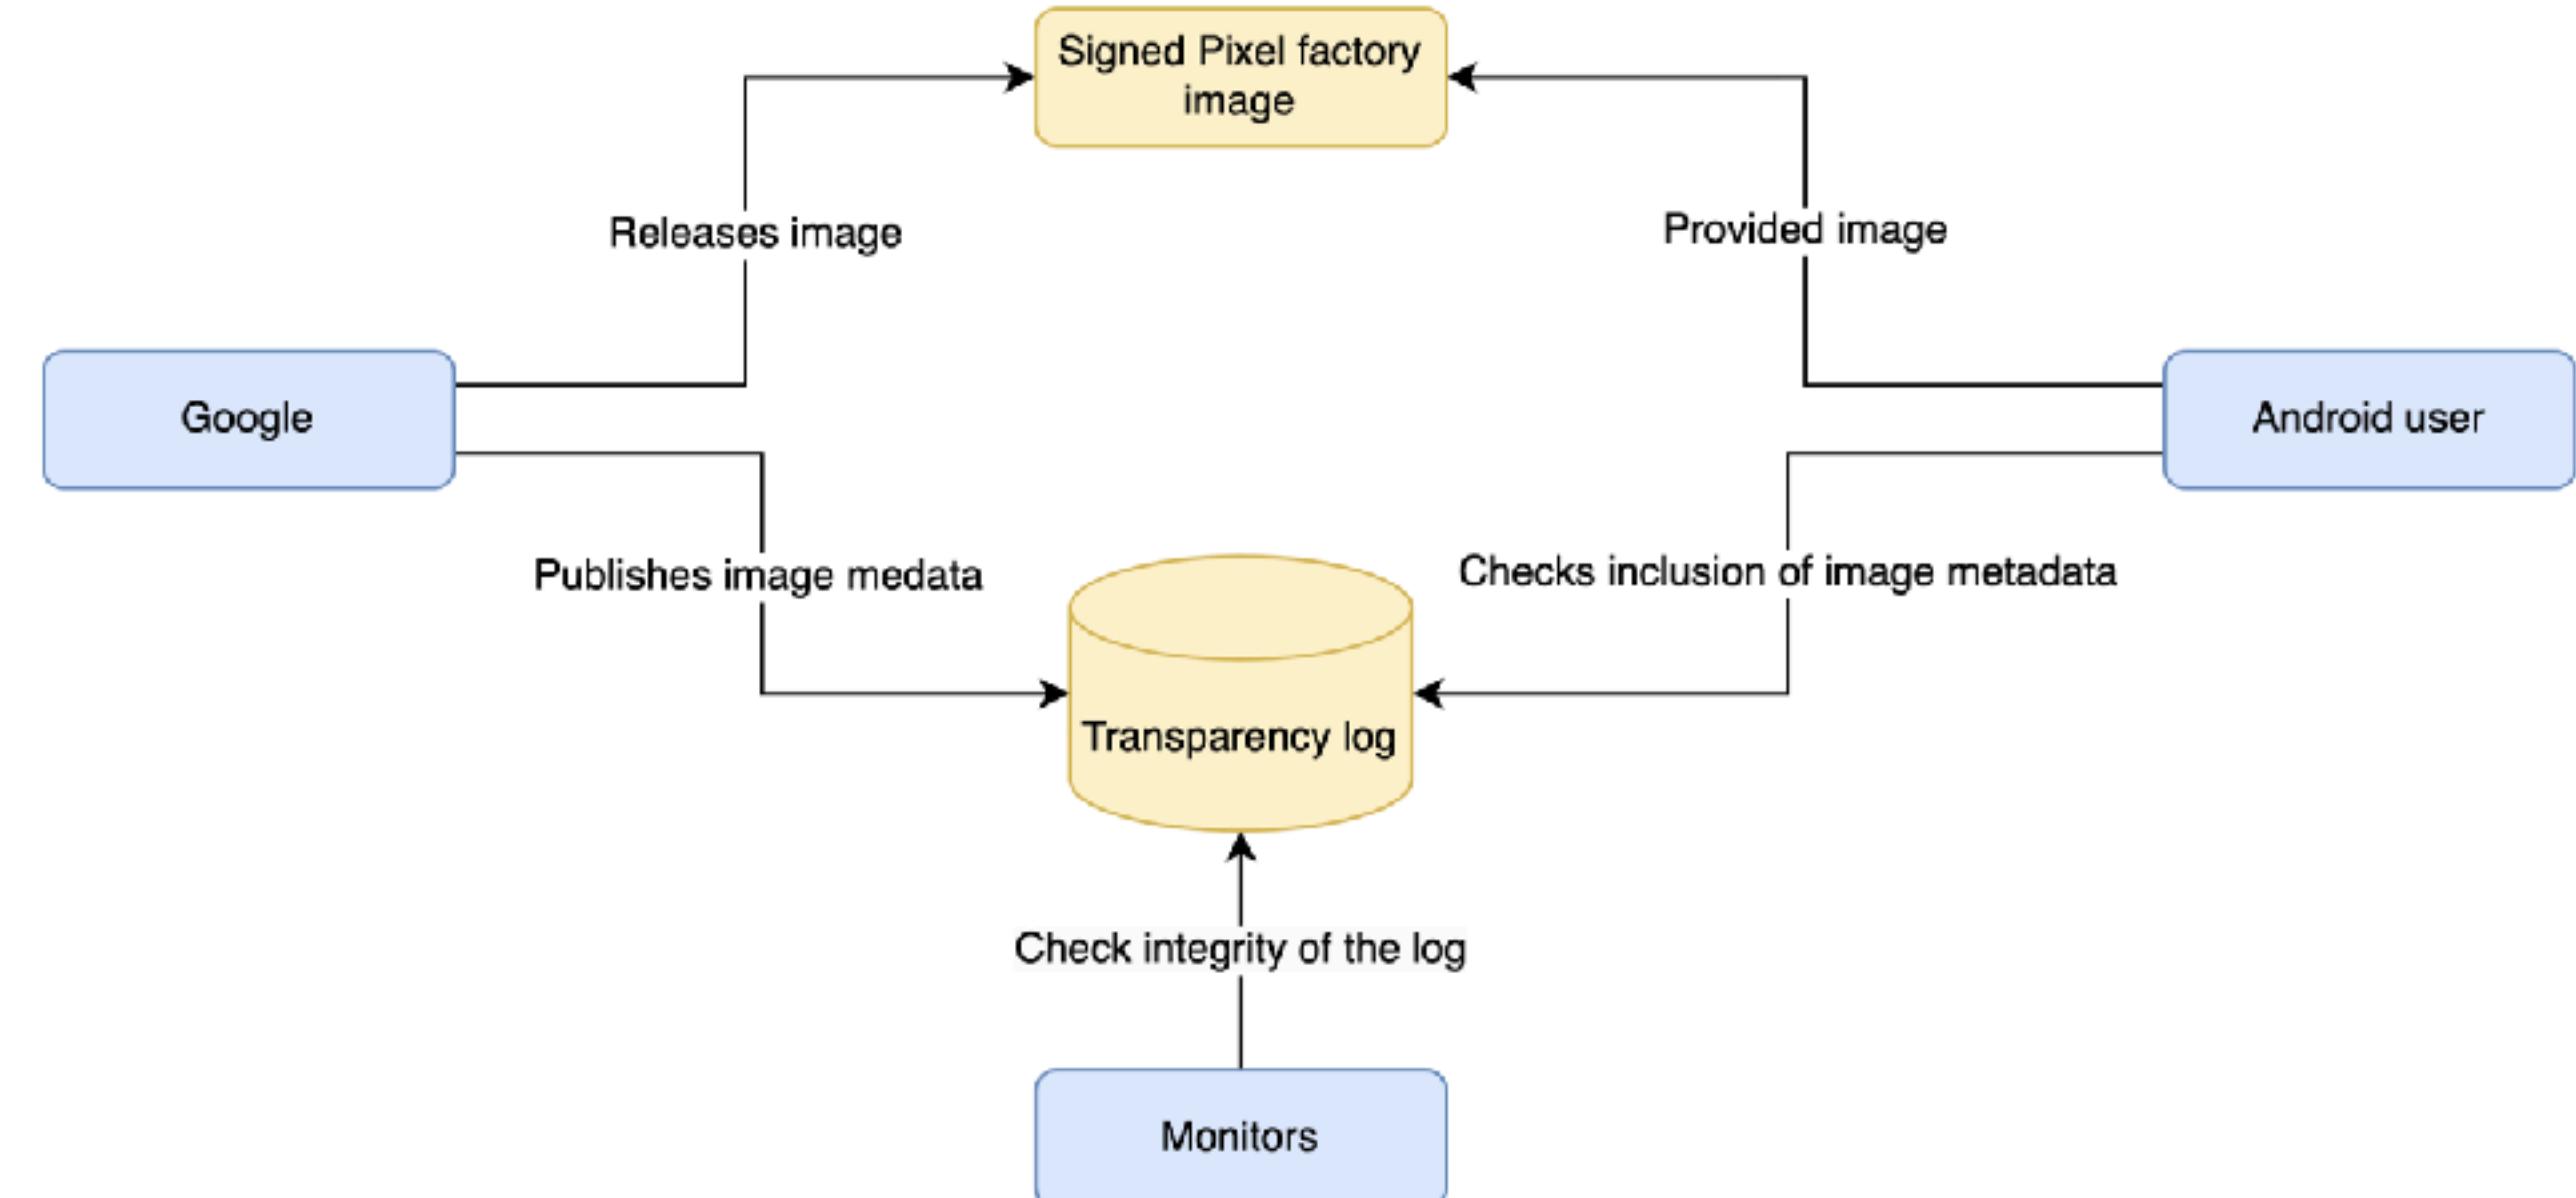
Google Chrome (2015), Cloudflare (2018), Let's Encrypt (2019), Firefox (2025)

- **Key transparency**

WhatsApp (2023), Signal

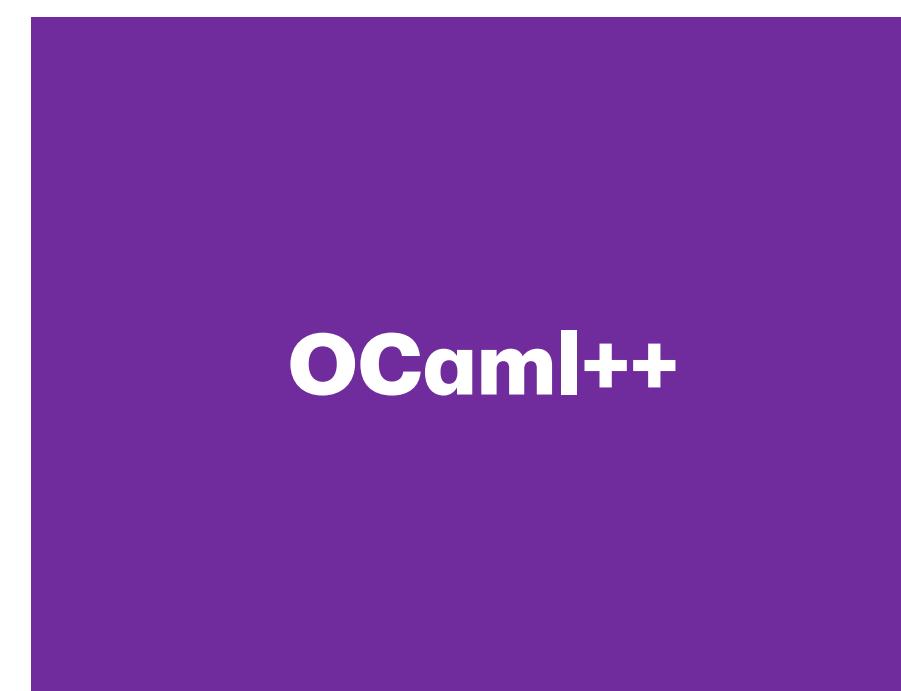
- **Binary transparency**

Pixel Binaries, Go modules



Miller et al. realized that the prover and verifier can be **compiled** from a single implementation of the “non-authenticated” data structure.

data
structure



→ prover

→ verifier

- - - → “ideal”



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 at distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyright 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 fee. Request permission from permissions@acm.org.
POPL ’14, January 22–24, 2014, San Diego, CA, USA
Copyright is held by the owner(s). Publication rights licensed to ACM.
ACM 978-1-4503-2544-8/14/01...\$15.00
<http://doi.acm.org/10.1145/2515538.2515851>

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

Miller et al.'s approach

OCaml is extended with three new primitives:

- authenticated types $\bullet \tau$
- auth : $'a \rightarrow \bullet 'a$
- unauth : $\bullet 'a \rightarrow 'a$



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. Copyright 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 fee. Request permission from permissions@acm.org.

POPL '14, January 22–24, 2014, San Diego, CA, USA
Copyright is held by the owner(s). Publication rights licensed to ACM.
ACM 978-1-4503-2544-8/14/01...\$15.00
<http://doi.acm.org/10.1145/2515538.2515851>

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

Miller et al.'s approach

OCaml is extended with three new primitives:

- authenticated types $\bullet \tau$
- auth : $'a \rightarrow \bullet 'a$
- unauth : $\bullet 'a \rightarrow 'a$

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.

Caterpillar and Cubicle: Practical Data Structures for the Cloud

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 ν then

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

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

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

Limitations

Limitations

1. A custom compiler frontend imposes development burden.

Limitations

1. A custom compiler frontend imposes development burden.
2. The compiler implements several optimizations that are not covered by the security and correctness theorems.

Limitations

1. A custom compiler frontend imposes development burden.
2. The compiler implements several optimizations that are not covered by the security and correctness theorems.
3. The generated data structures are not always as efficient or produce proofs as compact as hand-written implementations.

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



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

  (* ... *)

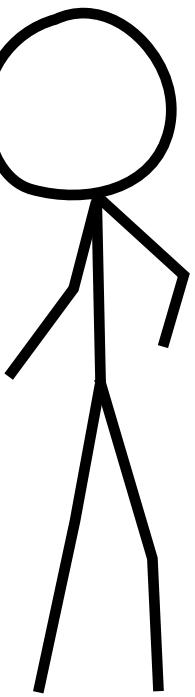
  val fetch : path -> tree auth -> string option auth_computation = (* ... *)
end
```

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

  (* ... *)

  val fetch : path -> tree auth -> string option auth_computation = (* ... *)
end
```

module Prover : AUTHENTIKIT



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

  (* ... *)

  val fetch : path -> tree auth -> string option auth_computation = (* ... *)
end
```

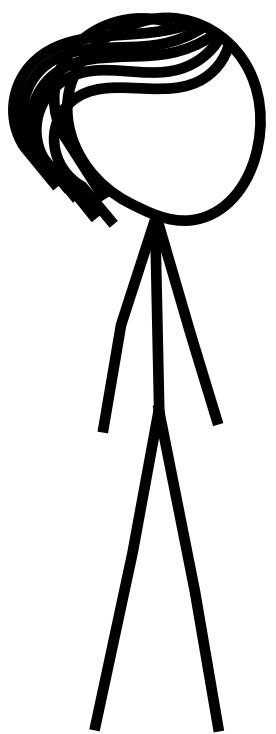
```
module Merkle = functor (A : AUTHENTIKIT) -> struct
  open A
  (* ... *)
  val fetch : path -> tree auth -> string option auth_computation = (* ... *)
end
```

module Prover : AUTHENTIKIT



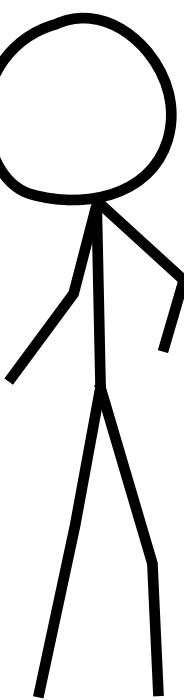
module Verifier : AUTHENTIKIT





module Ideal : AUTHENTIKIT

module Prover : AUTHENTIKIT



```
module Merkle = functor (A : AUTHENTIKIT) -> struct
  open A
  (* ... *)
  val fetch : path -> tree auth -> string option auth_computation = (* ... *)
end
```

module Verifier : AUTHENTIKIT



This work

- Two **logical relations** and a proof of security and correctness of the Authentikit module functor construction in OCaml.
- We address the remaining two limitations:
 - We verify several **optimizations** (as supported by the compiler).
 - We show how to **safely link** manually verified code with code automatically generated by Authentikit.
- Full mechanization in the Rocq theorem prover.

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

  (* ... *)

  module Serializable : sig
    type 'a evidence

    (* ... *)

  end

  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

    (* ... *)

  end

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

```
module 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 = 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.evidence = (* ... *)

  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 = 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.evidence = (* ... *)

  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

```

Takeaway

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.
In fact, we prove security and correctness as “free” theorems.
- To do this, we define two logical relations.

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

    (* ... *)

  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

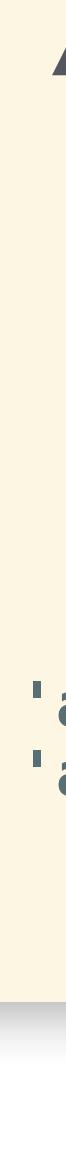
  val return : 'a -> 'a auth_computation
  val bind   : 'a auth_computation -> ('a -> 'b auth_computation) -> 'b auth_computation

  module Serializable : sig
    type 'a evidence

    (* ... *)

  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

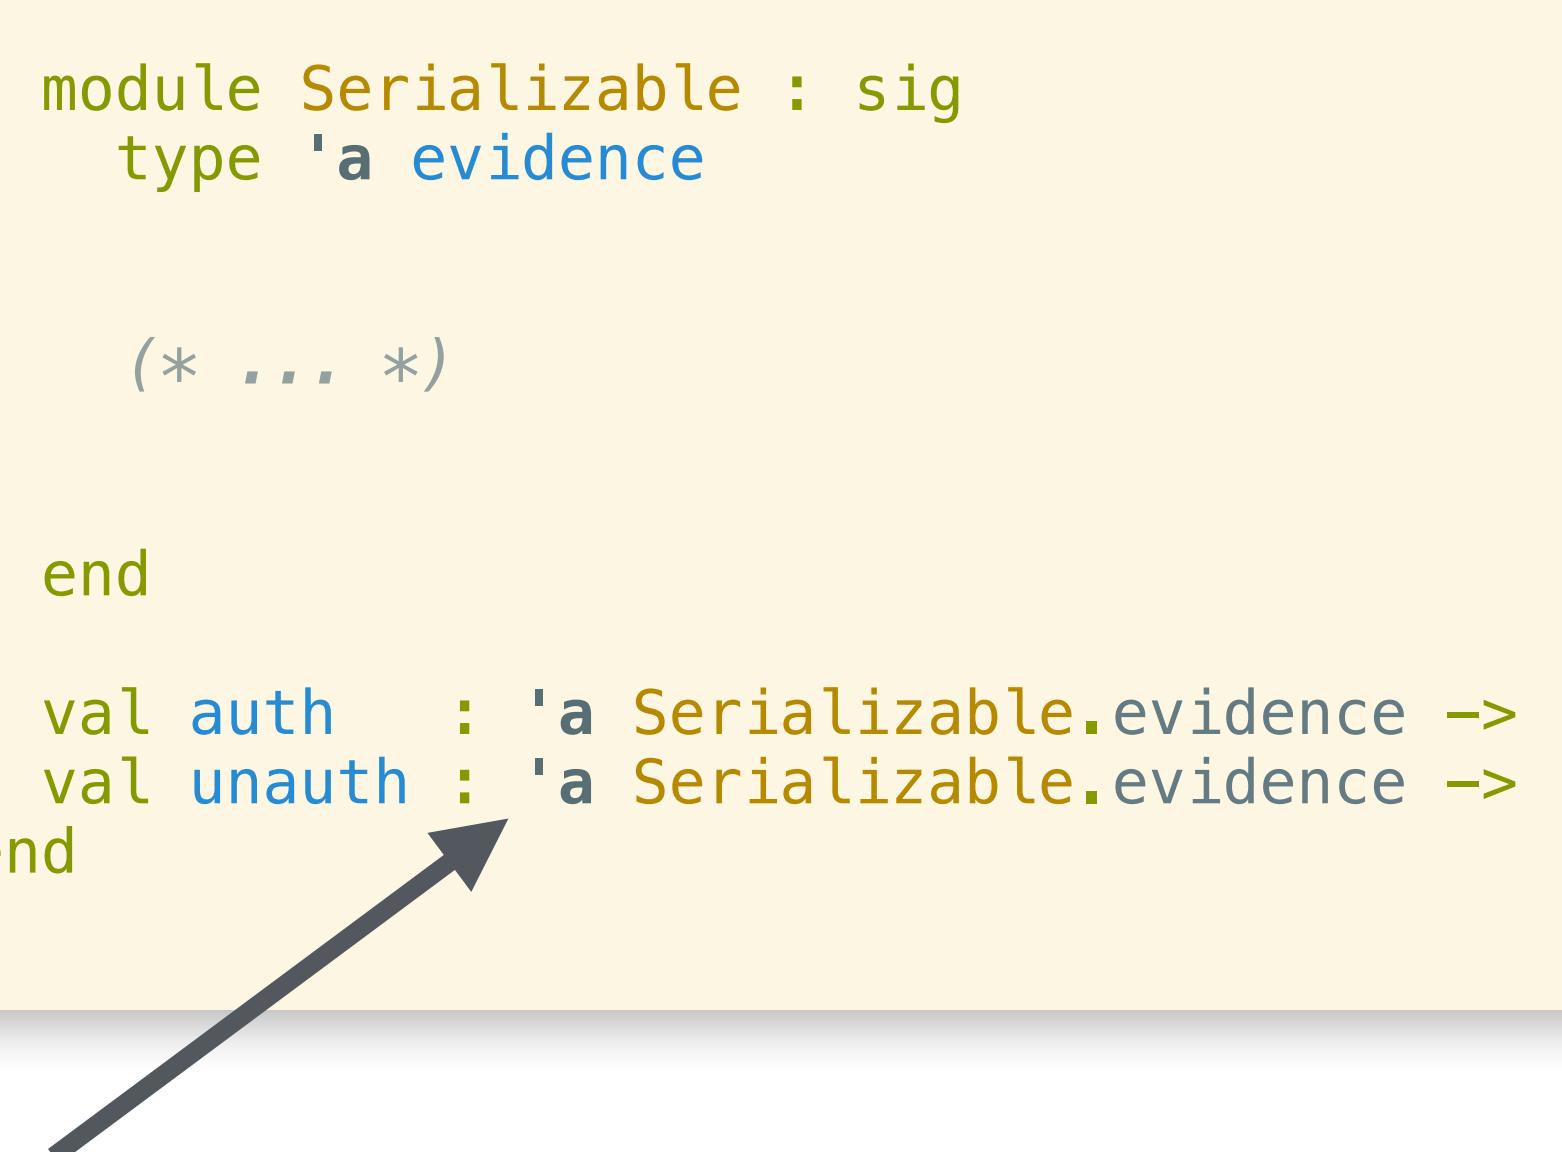
    (* ... *)

  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

abstract 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

    (* ... *)

  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

abstract 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

    (* ... *)

  end

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

polymorphism

```
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
```

recursive types

(higher-order) functions

Requirements

(abstract) type constructors

abstract 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
  end
(* ... *)
```

polymorphism

```
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

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

recursive types

(higher-order) functions

Requirements

(abstract) type constructors

abstract 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
  end
  (* ... *)
end
```

polymorphism

```
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
```

recursive types

state

(higher-order) functions

```
module Prover : AUTHENTIKIT
```

Our approach: “logical” logical relations

Our approach: “logical” logical relations

To show security and correctness we

Our approach: “logical” logical relations

To show security and correctness we

1. Define **Collision-Free Separation Logic** (in Iris).

Our approach: “logical” logical relations

To show security and correctness we

1. Define **Collision-Free Separation Logic** (in Iris).
2. Define **binary** and **ternary logical relations** for security and correctness.

Our approach: “logical” logical relations

To show security and correctness we

1. Define **Collision-Free Separation Logic** (in Iris).
2. Define **binary** and **ternary logical relations** for security and correctness.
3. Show implementations of the **Prover**, **Verifier**, and **Ideal** inhabit the model.

Theorem (Security)

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

Theorem (Security)

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

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

Theorem (Security)

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

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

then for all proofs p , if

Theorem (Security)

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

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

then for all proofs p , if

e instantiated with **Verifier** accepts p and returns v

Theorem (Security)

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

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

then for all proofs p , if

e instantiated with **Verifier** accepts p and returns v

then

Theorem (Security)

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

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

then for all proofs p , if

e instantiated with **Verifier** accepts p and returns ν

then

- e instantiated with **Ideal** returns ν or

Theorem (Security)

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

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

then for all proofs p , if

e instantiated with **Verifier** accepts p and returns ν

then

- e instantiated with **Ideal** returns ν or
- a **hash collision** occurred

Theorem (Correctness)

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

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

then if

e instantiated with **Prover** produces a proof p and returns ν

then

- e instantiated with **Verifier** accepts p and returns ν and
- e instantiated with **Ideal** returns ν as well.

Theorem (Correctness)

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

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

then if

e instantiated with **Prover** produces a proof p and returns ν

then

- e instantiated with **Verifier** accepts p and returns ν and
- e instanti

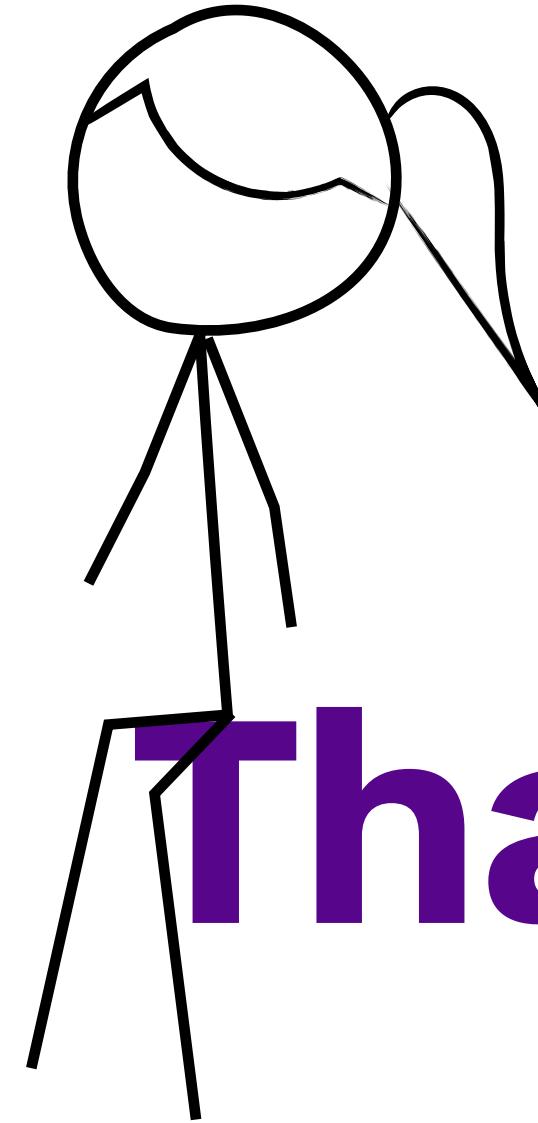
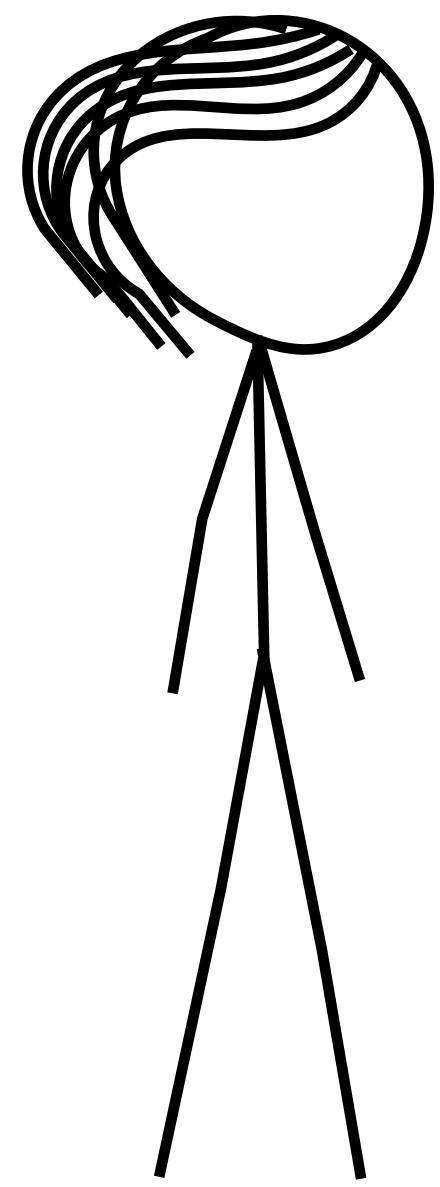
This proof requires prophecy variables!

Come talk to me later if you want to know more.

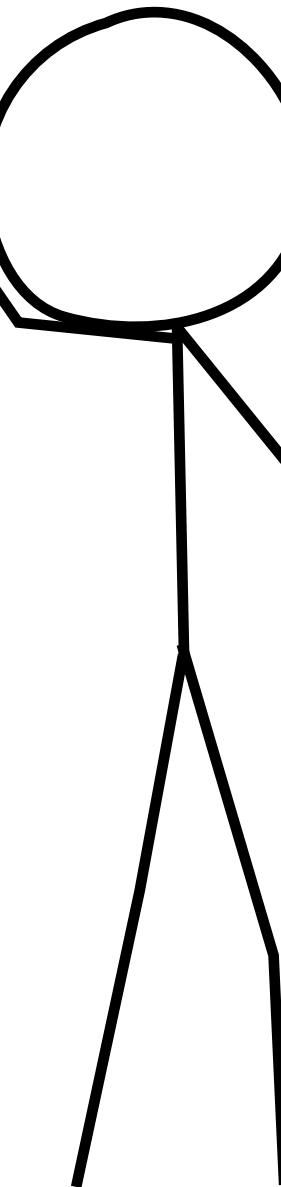
Summary

- **Authentikit** is a library for implementing ADSs generically.
- Two **logical-relations models** and a proof of security and correctness of the Authentikit module functor construction in OCaml.
 - We verify several **optimizations**.
 - We show how to **safely link** manually verified code with code automatically generated using Authentikit.
- Full mechanization in the Rocq theorem prover.

<https://arxiv.org/abs/2501.10802>



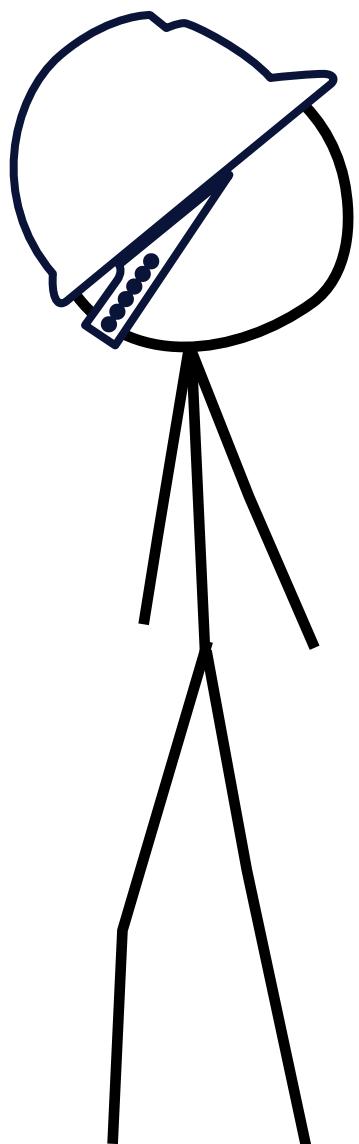
That's it, folks!



```
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

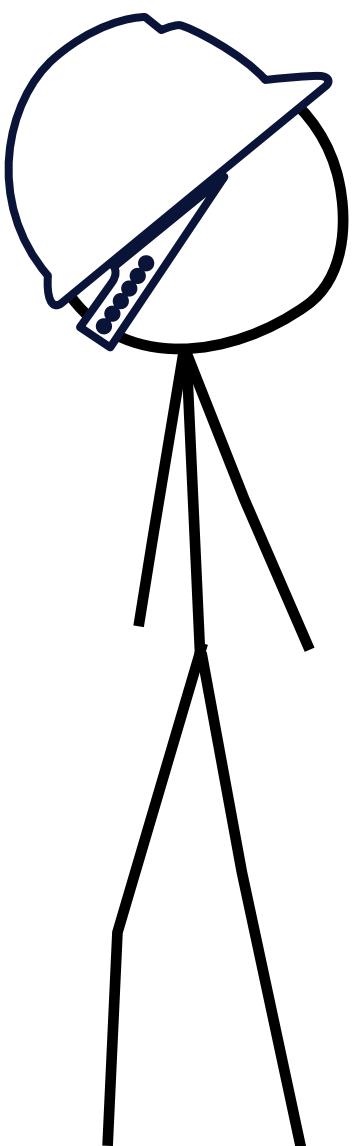
  let return a () = ([], a)
  let bind c f =
    let (prf, a) = c () in
    let (prf', b) = f a () in
    (prf @ prf', b)

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

  (* ... *)
end

(* ... *)

end
```



```
type proof = string list

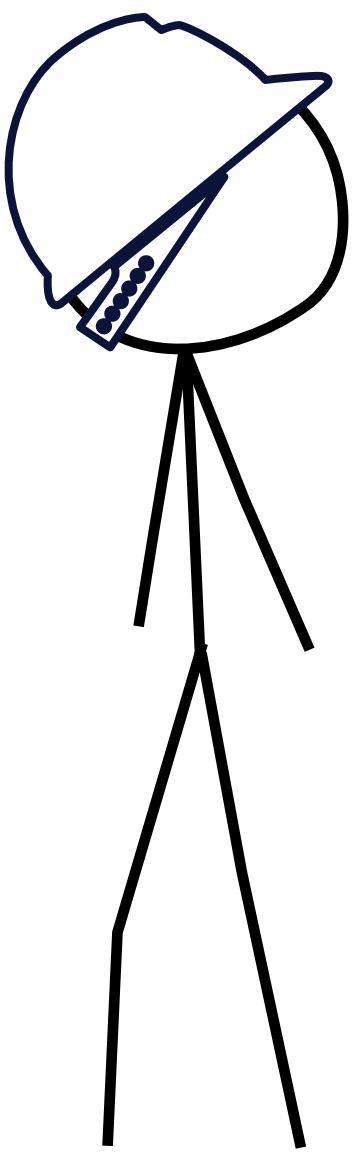
module Prover : AUTHENTIKIT =
  type 'a auth = 'a * string
  type 'a auth_computation = () -> proof * 'a

  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
```

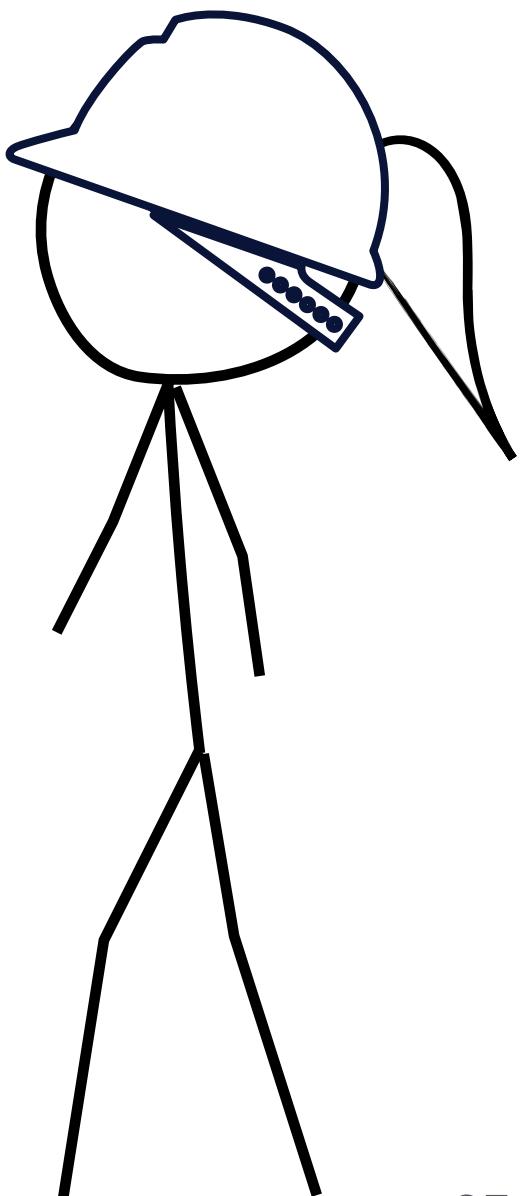


```
module Verifier : AUTHENTIKIT =
  type 'a auth = string
  type 'a auth_computation =
    proof -> [ `Ok of proof * 'a | `ProofFailure]
```

(* ... *)

(* ... *)

end



```

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

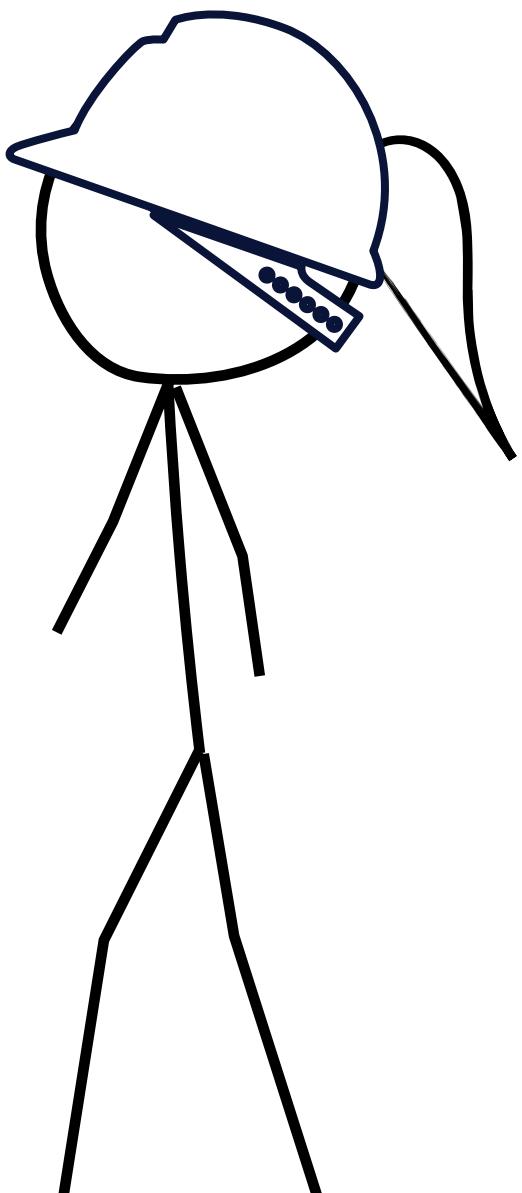
  let return a prf = `Ok (prf, a)
  let bind c f prf =
    match c prf with
    | `ProofFailure -> `ProofFailure
    | `Ok (prf', a) -> f a prf'

  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]

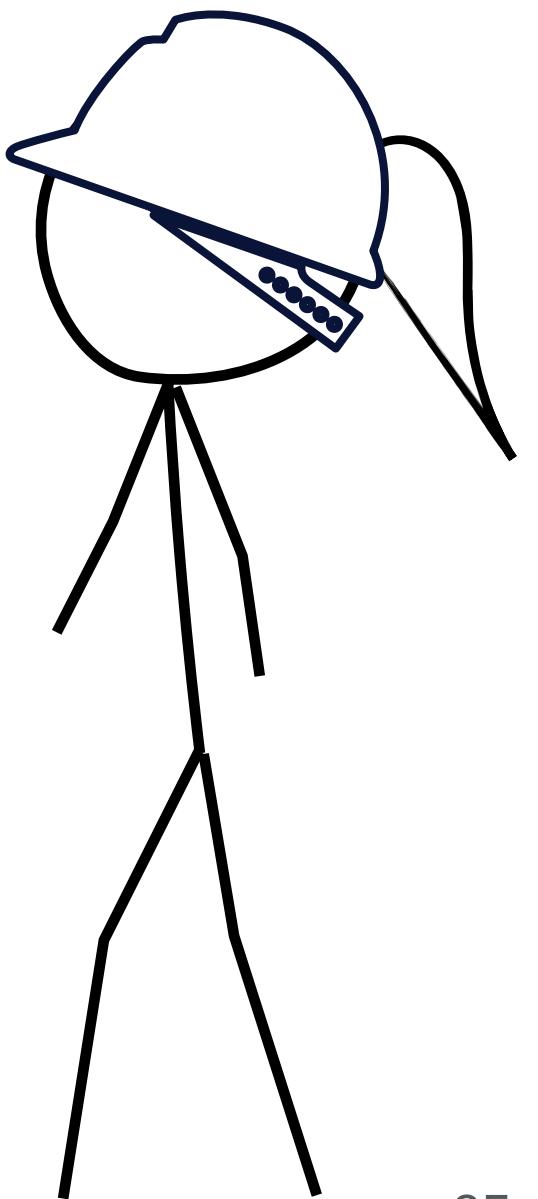
  let return a prf = `Ok (prf, a)
  let bind c f prf =
    match c prf with
    | `ProofFailure -> `ProofFailure
    | `Ok (prf', a) -> f a prf'

  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

```

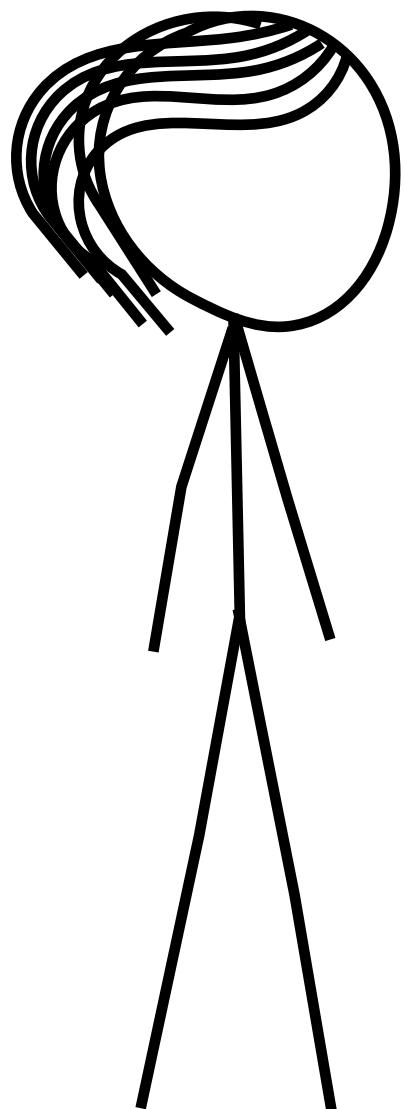


```
module Ideal : AUTHENTIKIT = struct
  type 'a auth = 'a
  type 'a auth_computation = () -> 'a

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

  (* ... *)

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



Reminder

STLC: terms can depend on terms,

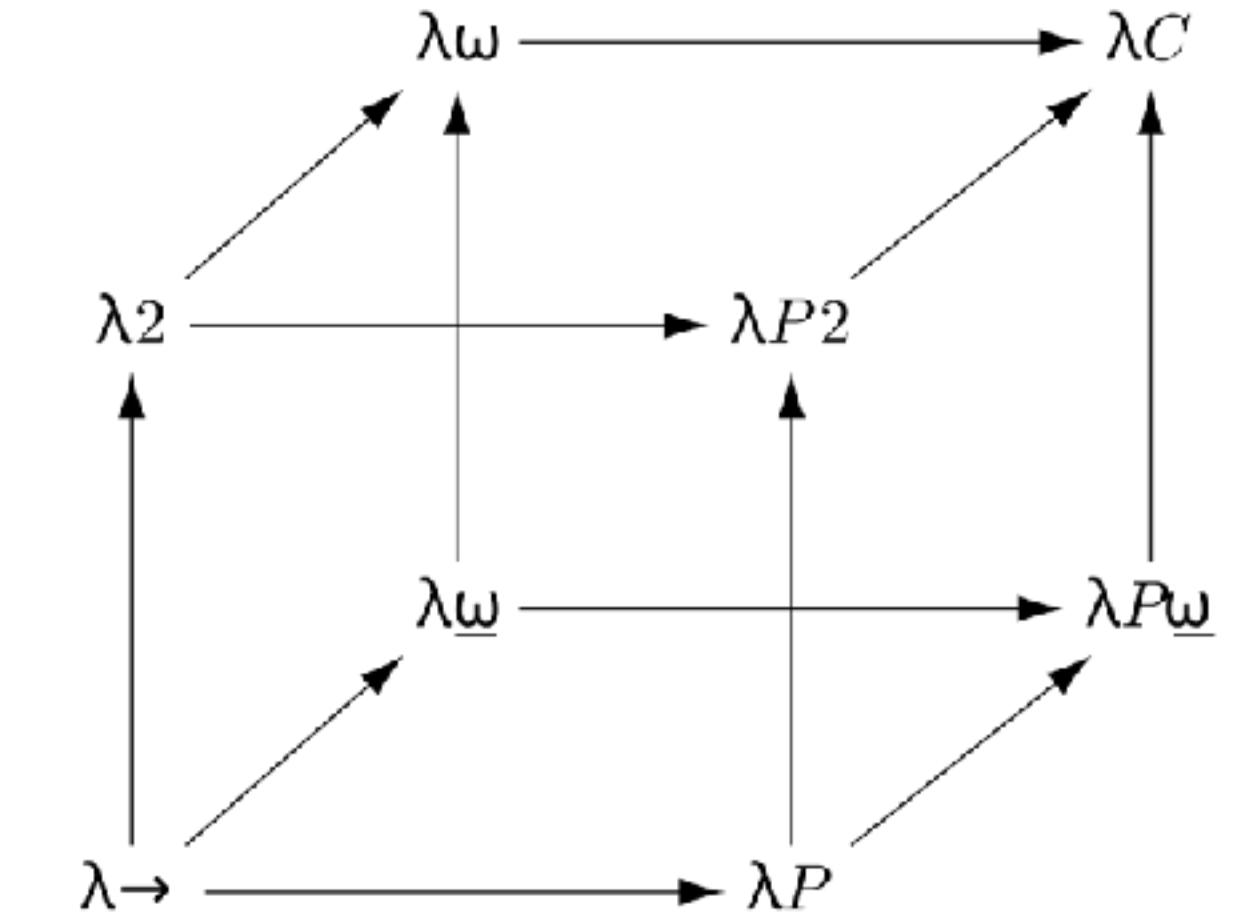
$$\frac{\Gamma, x : \sigma \vdash e : \tau}{\Gamma \vdash \lambda x. e : \sigma \rightarrow \tau}$$

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,

$$\frac{\Theta \vdash \tau \equiv \sigma \quad \Theta \mid \Gamma \vdash e : \sigma}{\Theta \mid \Gamma \vdash e : \tau} \qquad \frac{}{\Theta \vdash (\lambda \alpha. \tau) \sigma \equiv \tau[\sigma/\alpha]}$$



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

```

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

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

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

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

⋮

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

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

Collision-free reasoning

We define 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-free

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

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

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

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

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

Security

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

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

and show

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

Logical relation, sketch

Intuitively, the judgment $\emptyset \mid \emptyset \models e_1 \sim e_2 : \tau$ 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.

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

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

Security proof

The main work is to show

$$\llbracket \text{Authentikit auth m} \rrbracket(\text{Authentikit}_V, \text{Authentikit}_I)$$

The challenging part is finding the right interpretation of the type variables.

$$\llbracket \text{auth} \rrbracket(A)(v_1, v_2) \triangleq \exists a, t. v_1 = \text{hash}(\text{serialize}_t(a)) * A(a, v_2) * \text{hashed}(\text{serialize}_t(a))$$

$$\llbracket \text{m} \rrbracket(A)(v_1, v_2) \triangleq \forall p. \{\text{isProof}(p)\} v_1 p \sim v_2 () \{Q_{\text{post}}\}$$

$$Q_{\text{post}}(u_1, u_2) \triangleq u_1 = \text{None} \vee (\exists a_1, p'. u_1 = \text{Some}(p', a_1) * \text{isProof}(p') * A(a_1, u_2))$$

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 client proofs

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

$$\text{lookup}([R, L], t_0) =([(h_1, \textcolor{red}{h}_2), (\textcolor{red}{h}_5, h_6), s_5], s_5)$$

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

$$[\![\text{path} \rightarrow \text{auth tree} \rightarrow m \text{ (option string)}]\!](\text{fetch}_V, \text{fetch}_I)$$

