

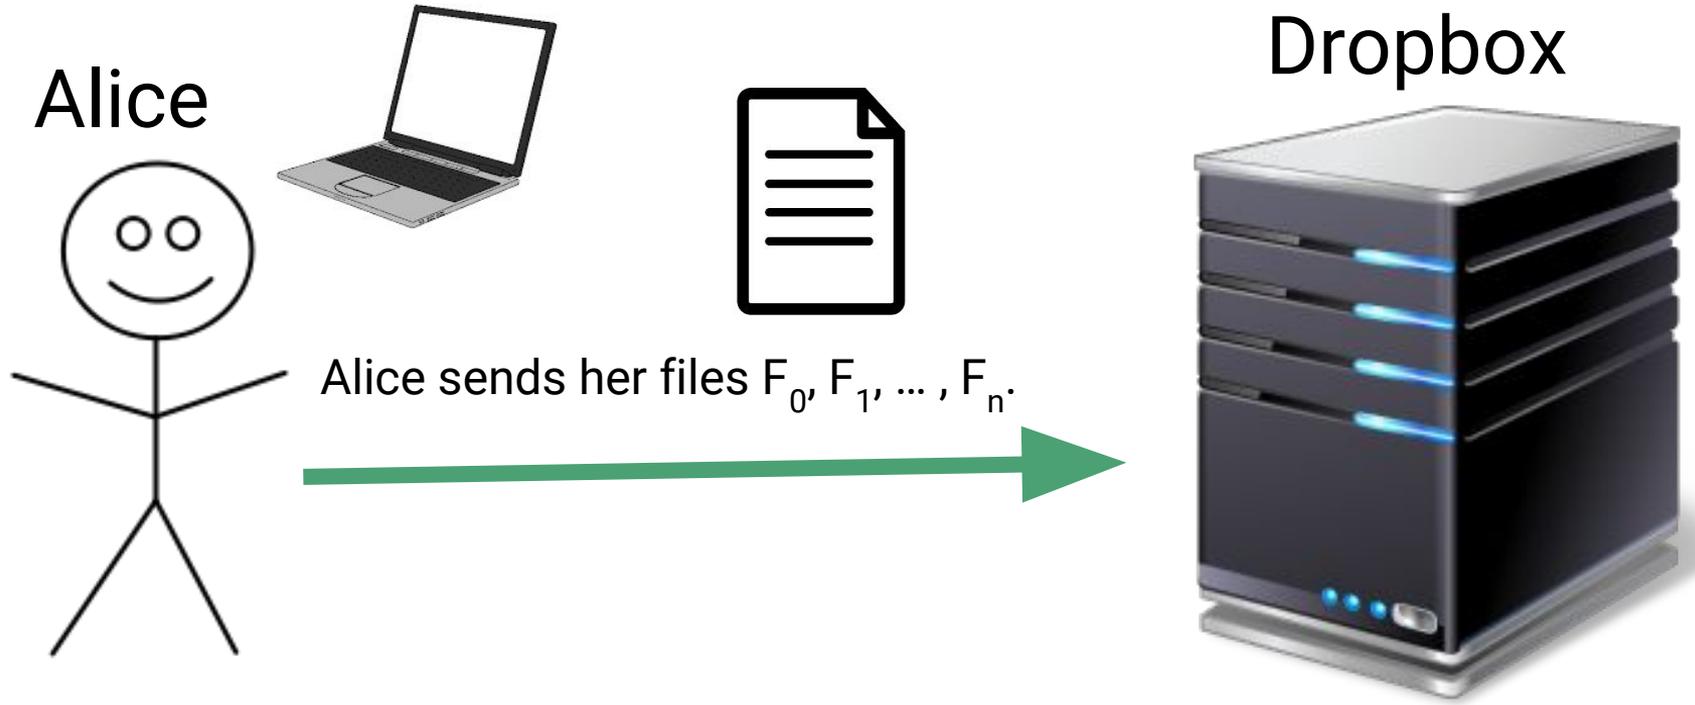
Verkle Trees: Ver(y Short Mer)kle Trees

John Kuszmaul

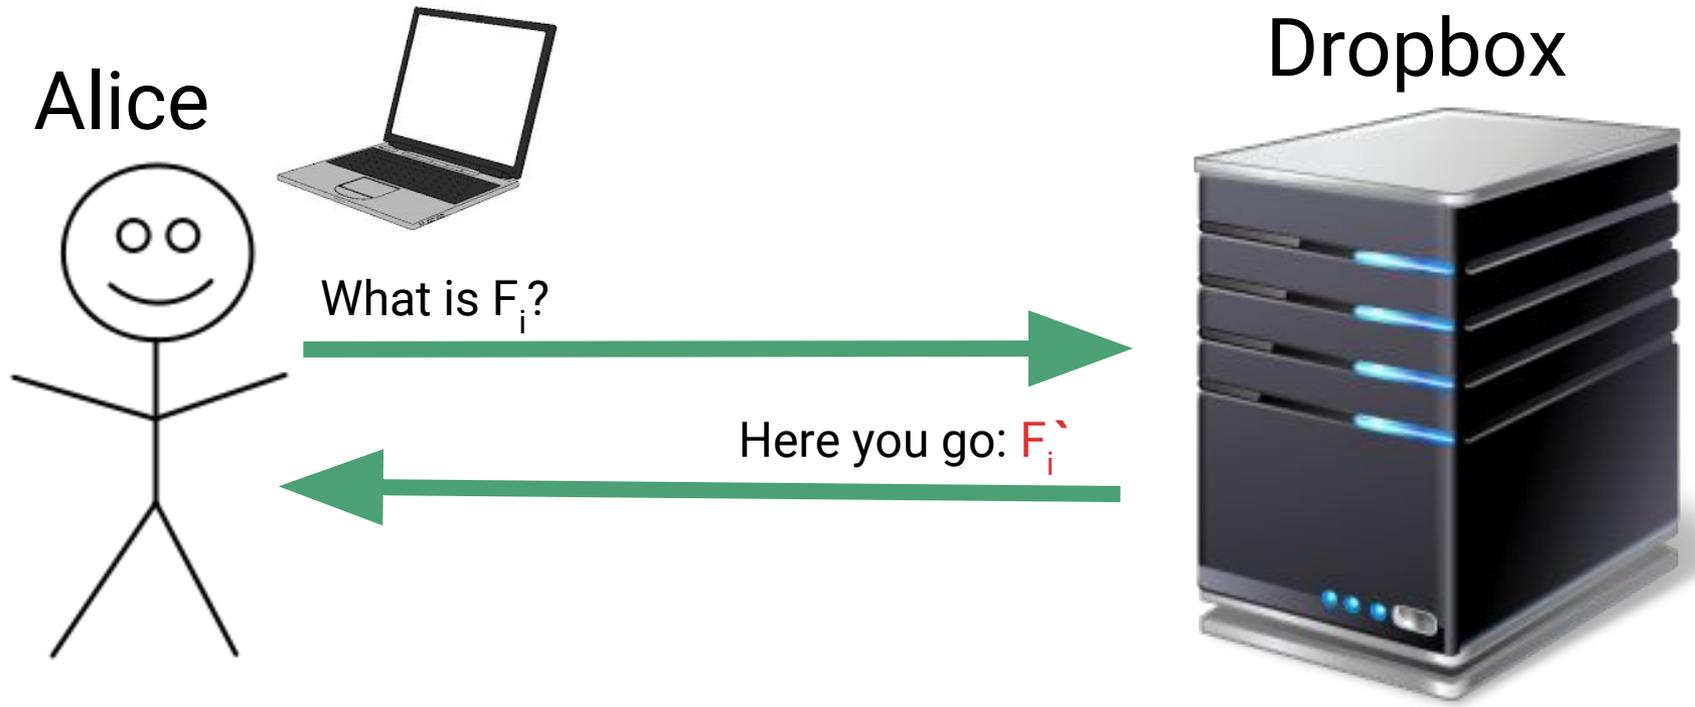
Mentored by Alin Tomescu

PRIMES Computer Science Conference - 10/13/18

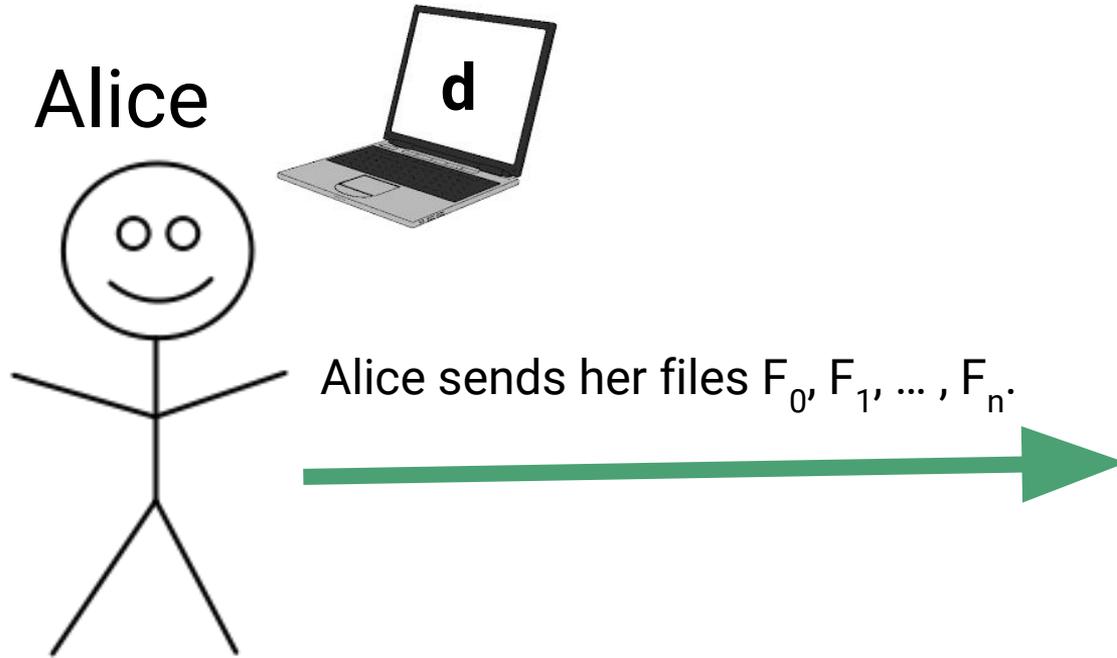
Storing Files Remotely



Storing Files Remotely



Proving/Verifying Integrity

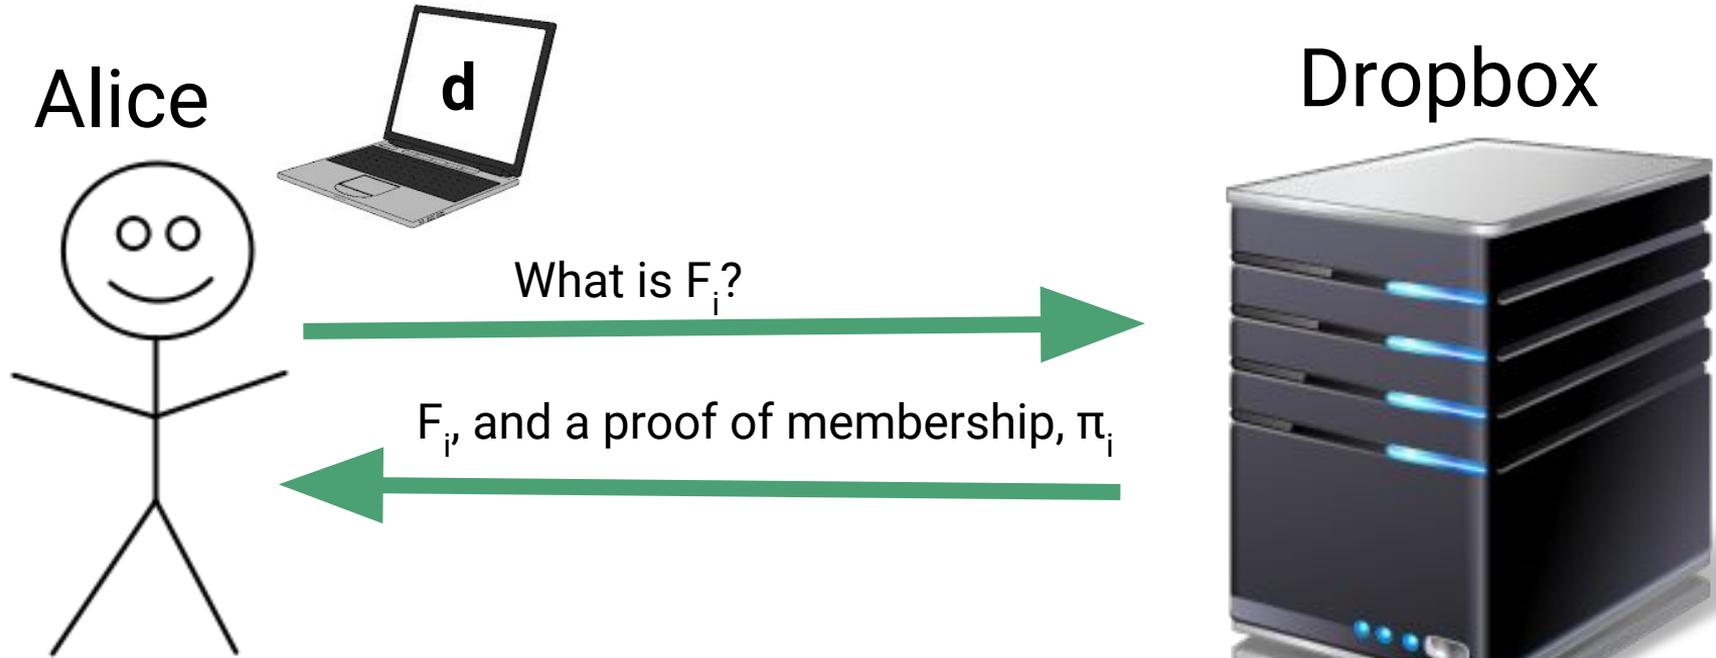


Dropbox



Alice generates a **digest** d of her files.

Proving/Verifying Integrity



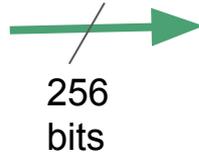
Alice verifies the proof π_i against d to make sure F_i has not been modified.

Secure Hash Functions

Original File F_i



Hash
Function

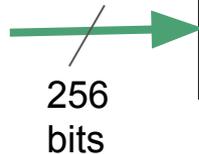


$H(F_i) =$
011010...110

Corrupted File F_i'



Hash
Function



$H(F_i') =$
100111...101

Ideally,

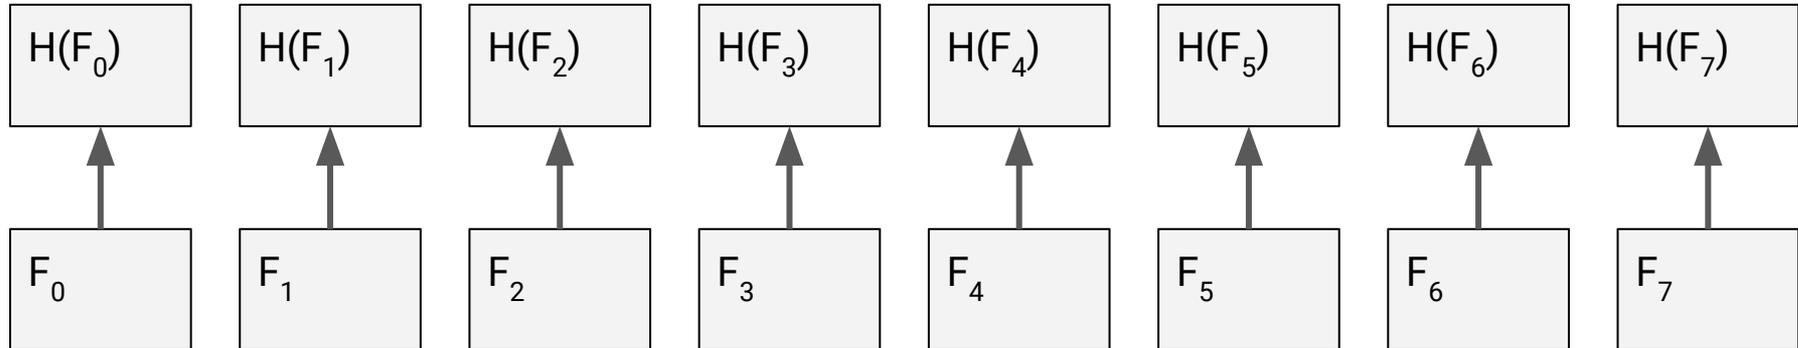
finding any two distinct
files, F_1, F_2 , s.t.

$$H(F_1) = H(F_2)$$

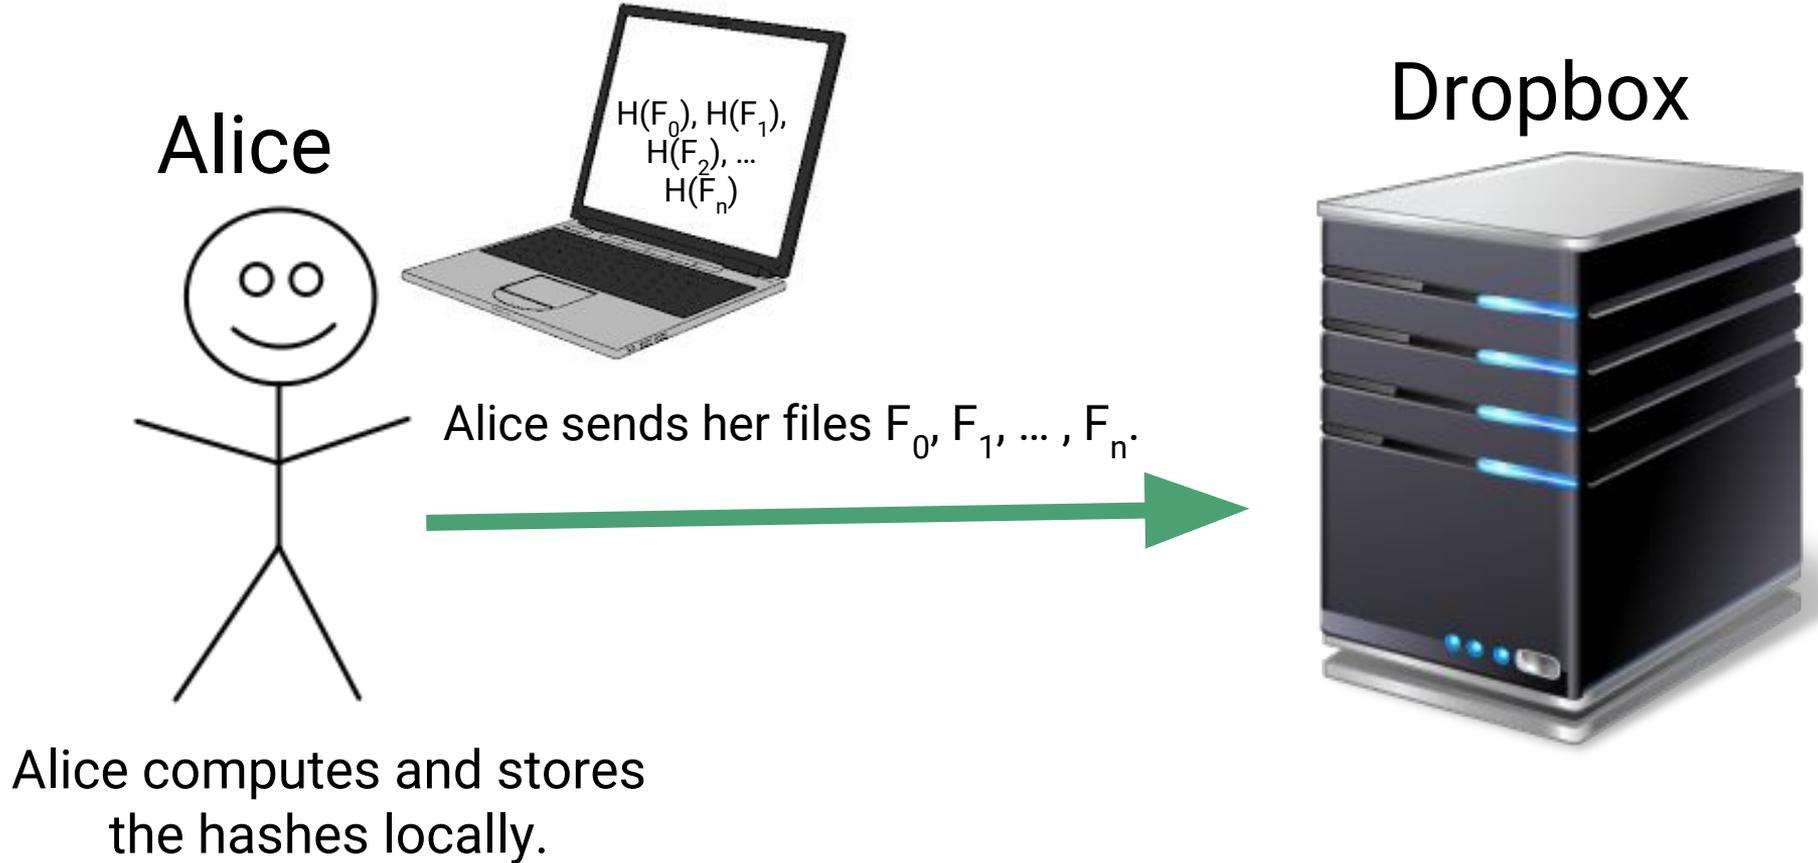
takes 2^{128} attempts.

A Simple Scheme for Verifying File Integrity

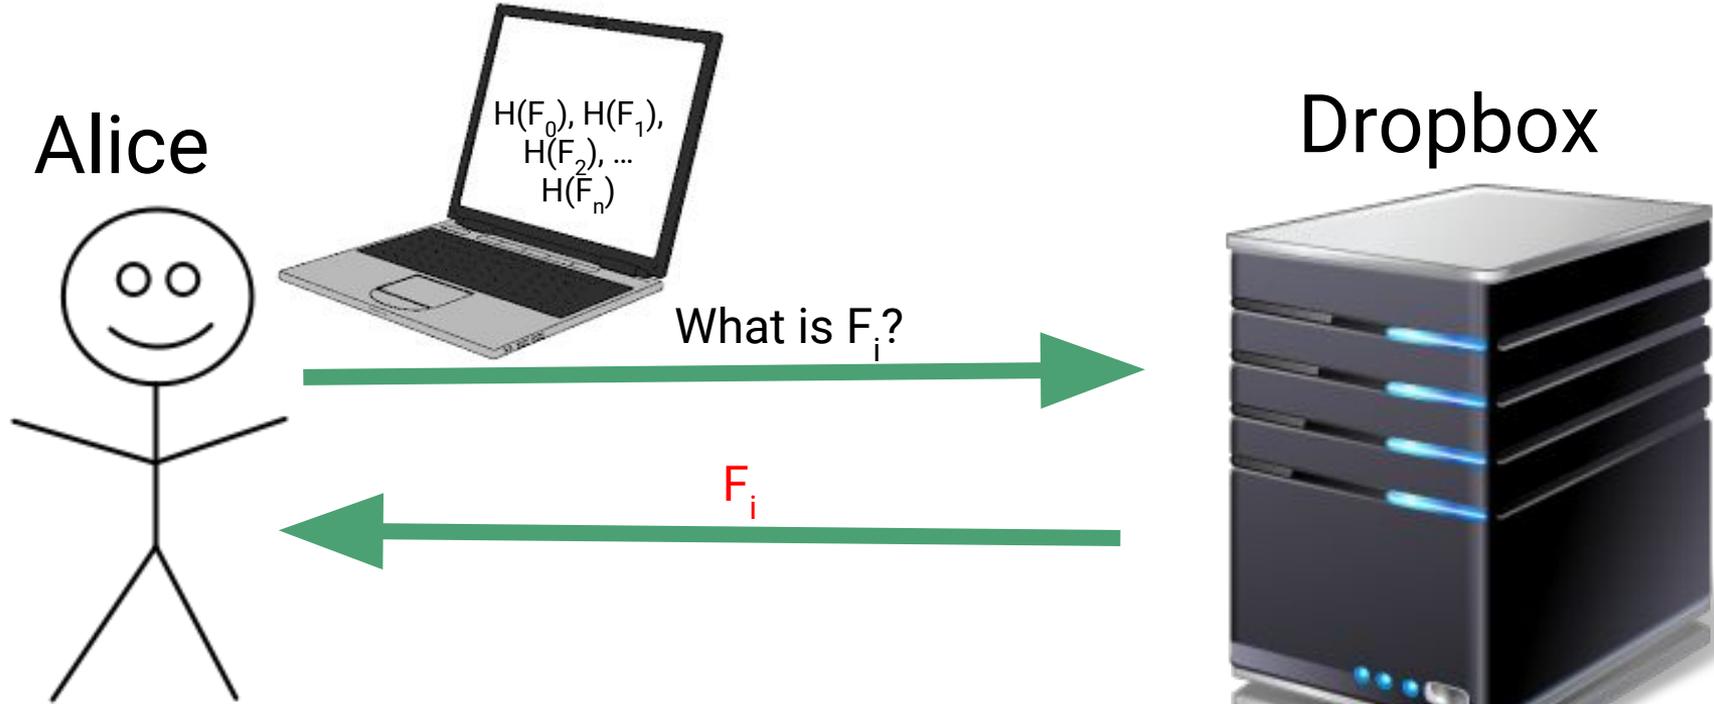
Alice hashes each of her files:



Proving/Verifying Integrity: Simple Scheme

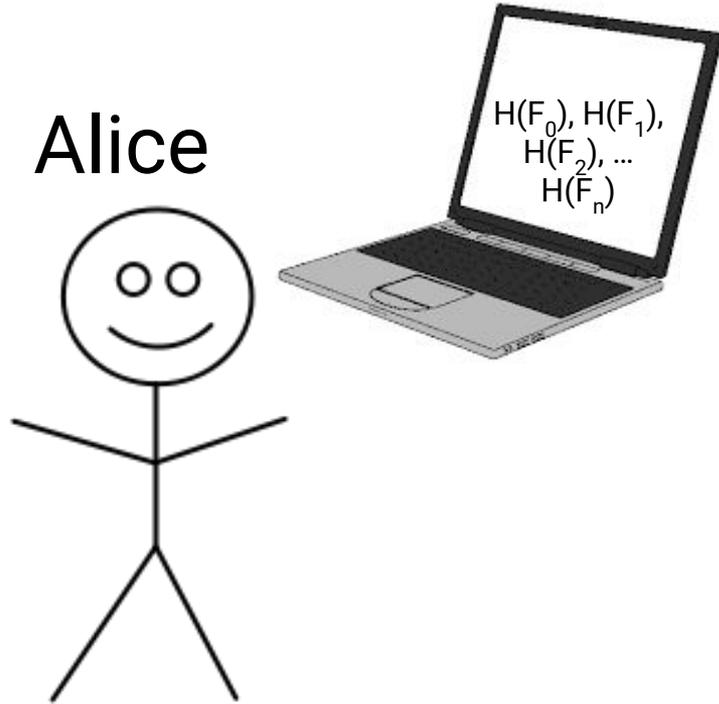


Proving/Verifying Integrity: Simple Scheme



Alice computes $H(F_i)$ and checks that it equals stored $H(F_i)$.

Problem: Alice has to store n hashes.



Alice's digest must be constant-sized.

Solution: Merkle Trees

$h_{14} = H(h_{12}, h_{13})$ ← **The root is the digest.**

$h_{12} = H(h_8, h_9)$

$h_{13} = H(h_{10}, h_{11})$

$h_8 = H(h_0, h_1)$

$h_9 = H(h_2, h_3)$

$h_{10} = H(h_4, h_5)$

$h_{11} = H(h_6, h_7)$

$h_0 = H(F_0)$ $h_1 = H(F_1)$

$h_2 = H(F_2)$ $h_3 = H(F_3)$

$h_4 = H(F_4)$ $h_5 = H(F_5)$

$h_6 = H(F_6)$ $h_7 = H(F_7)$

F_0

F_1

F_2

F_3

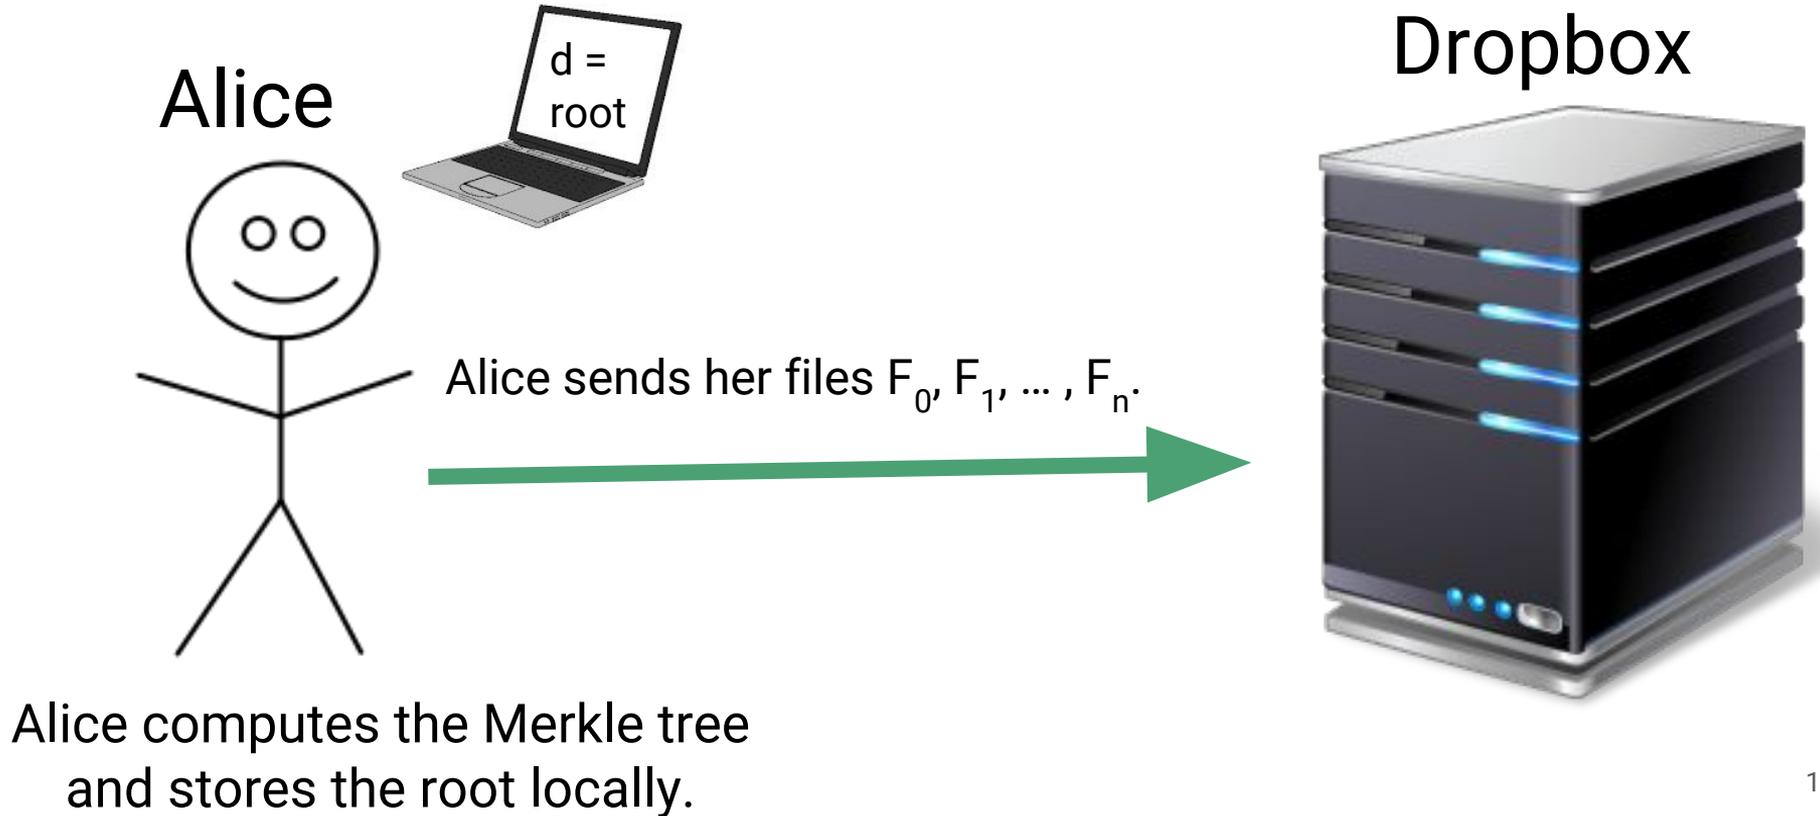
F_4

F_5

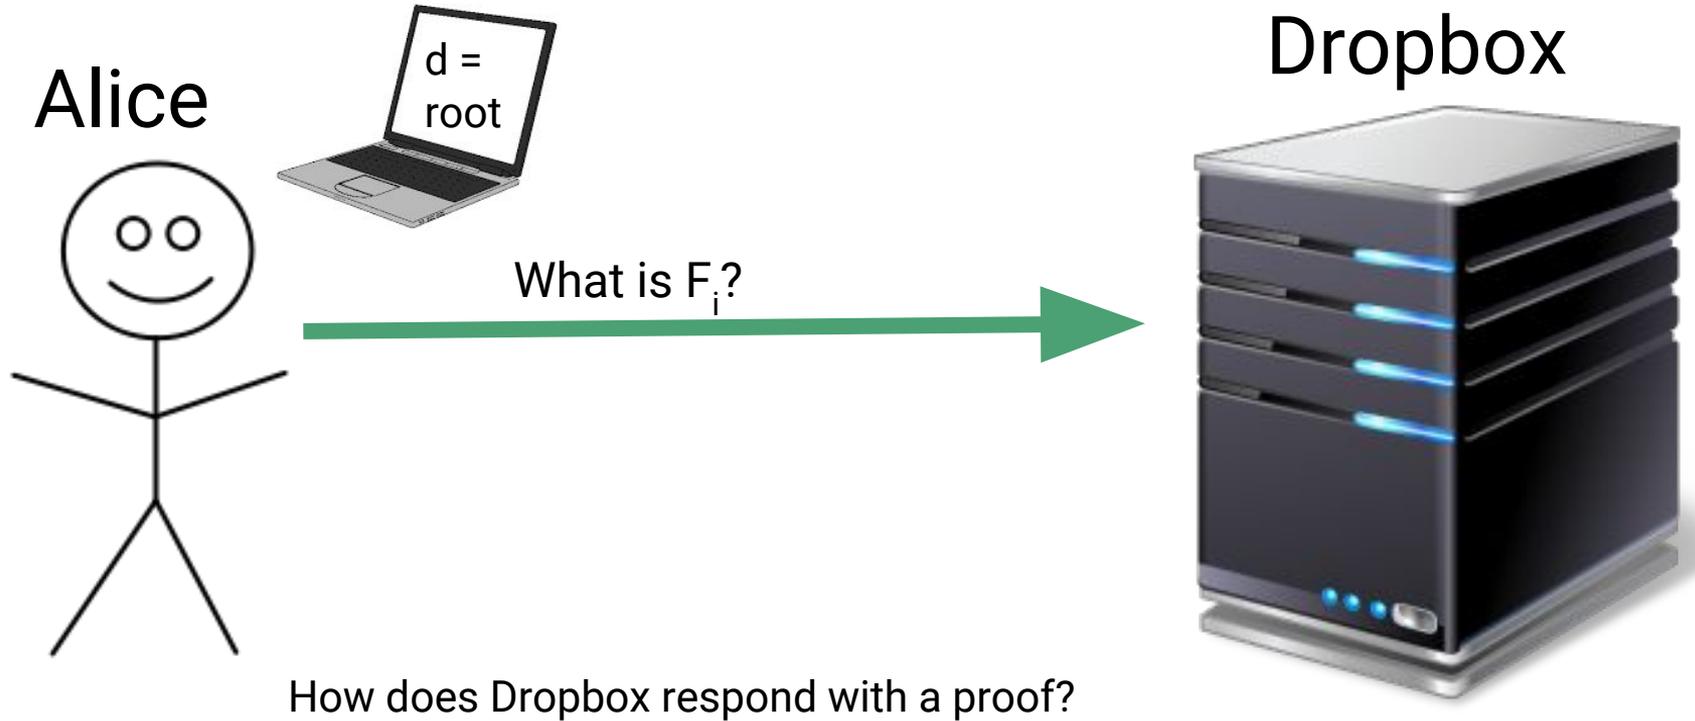
F_6

F_7

Proving/Verifying Integrity: Merkle Tree

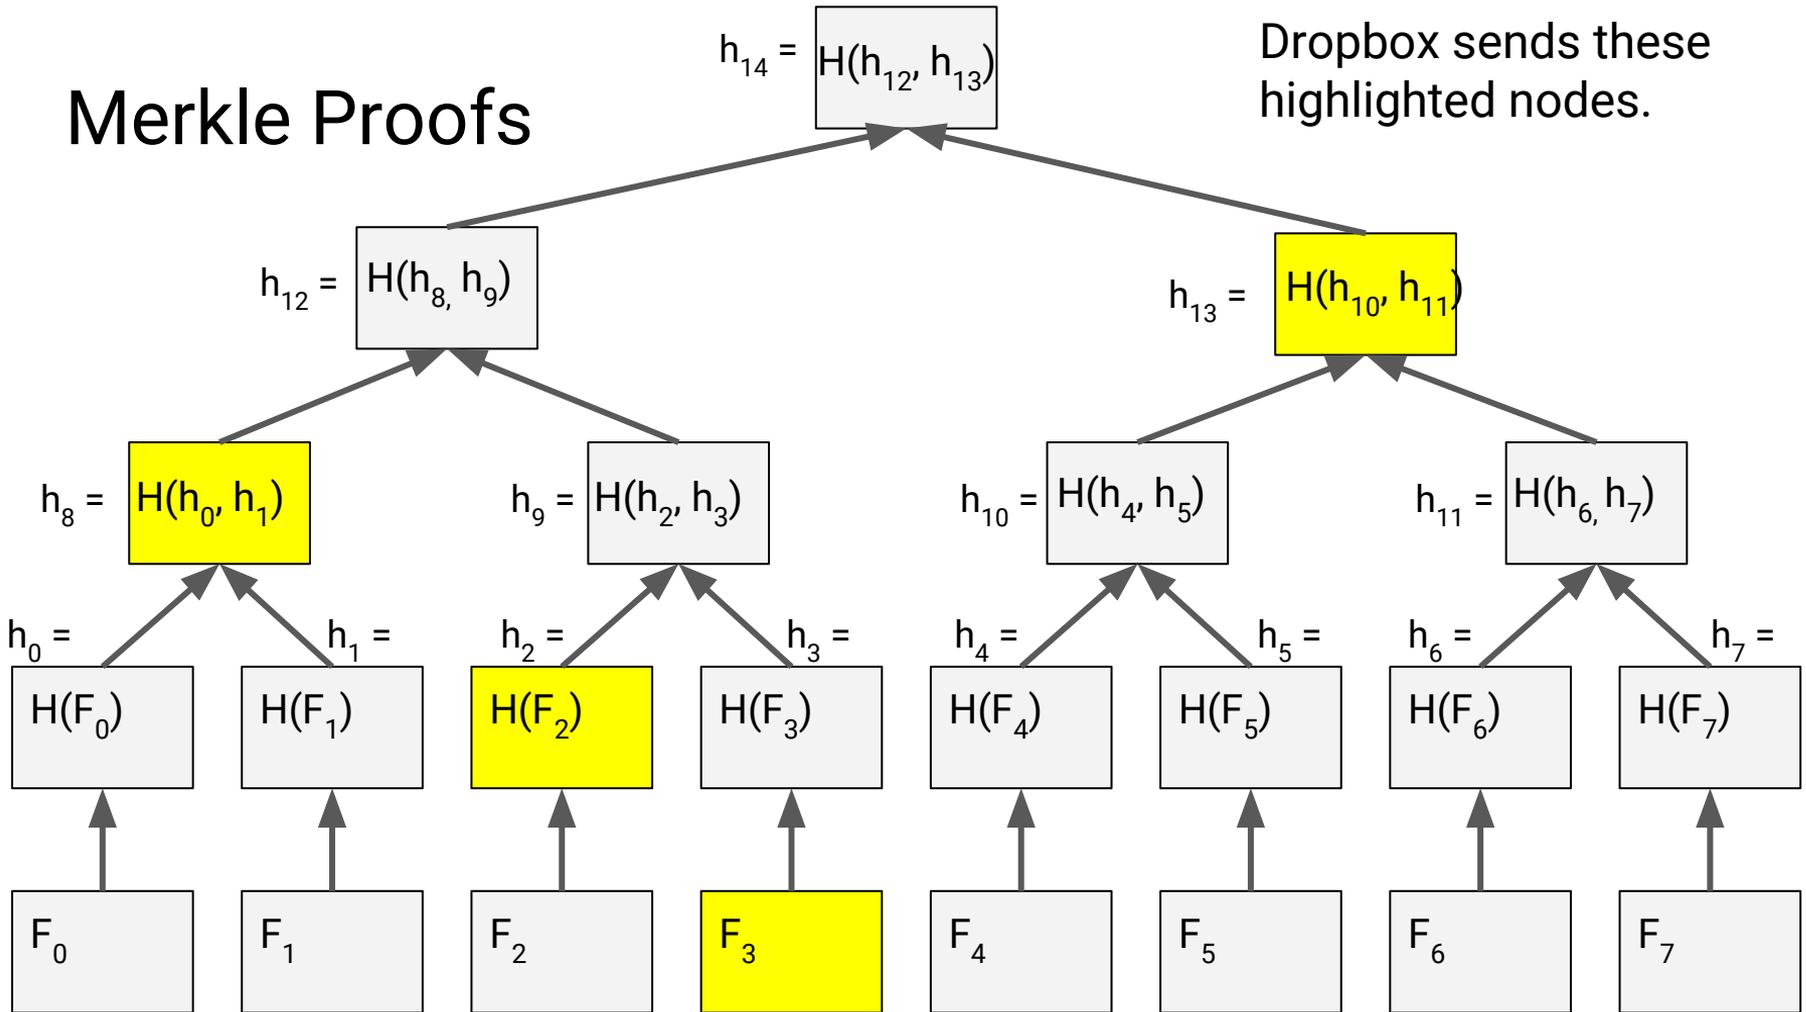


Proving/Verifying Integrity: Merkle Tree

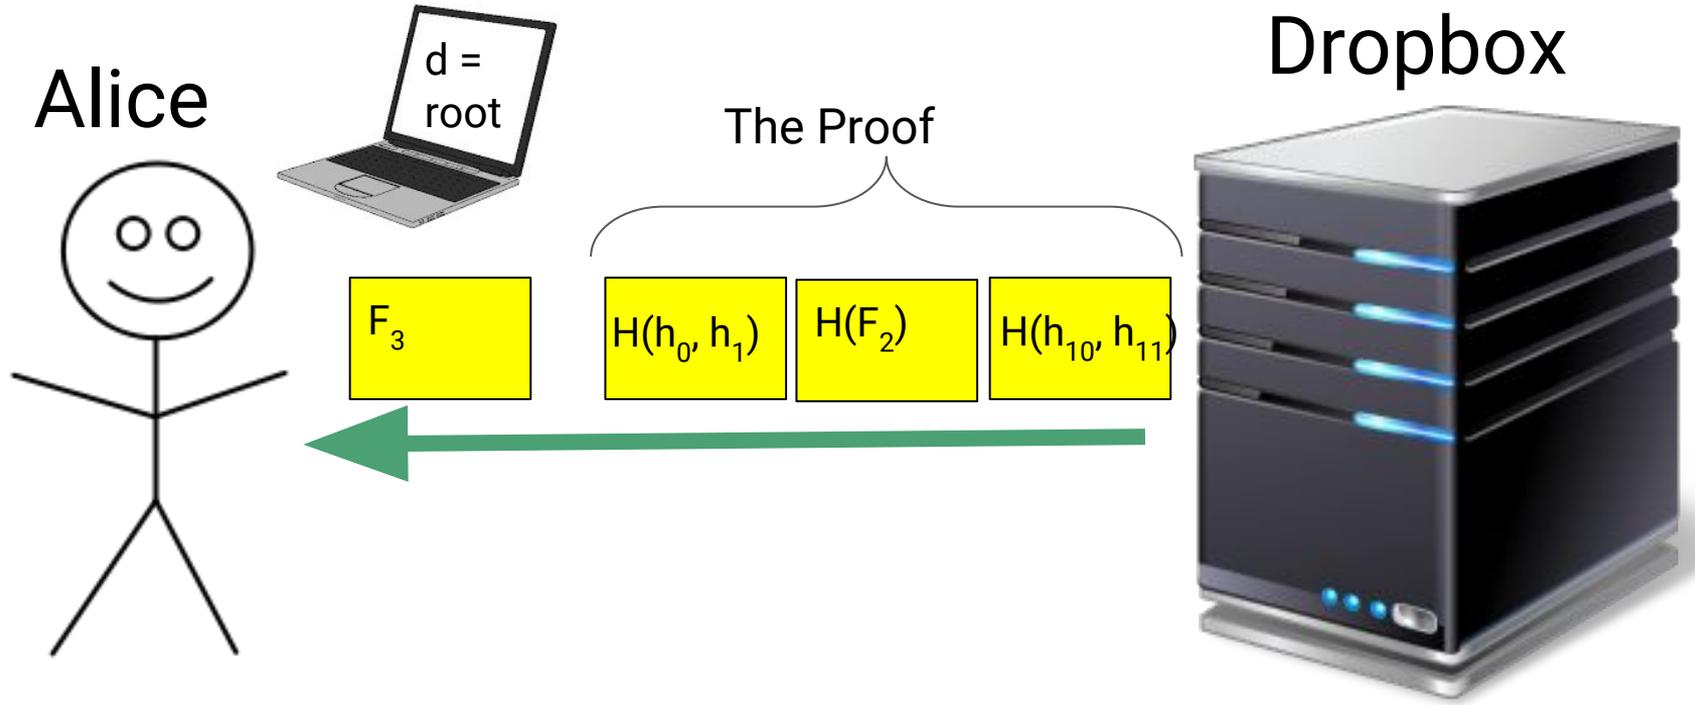


Merkle Proofs

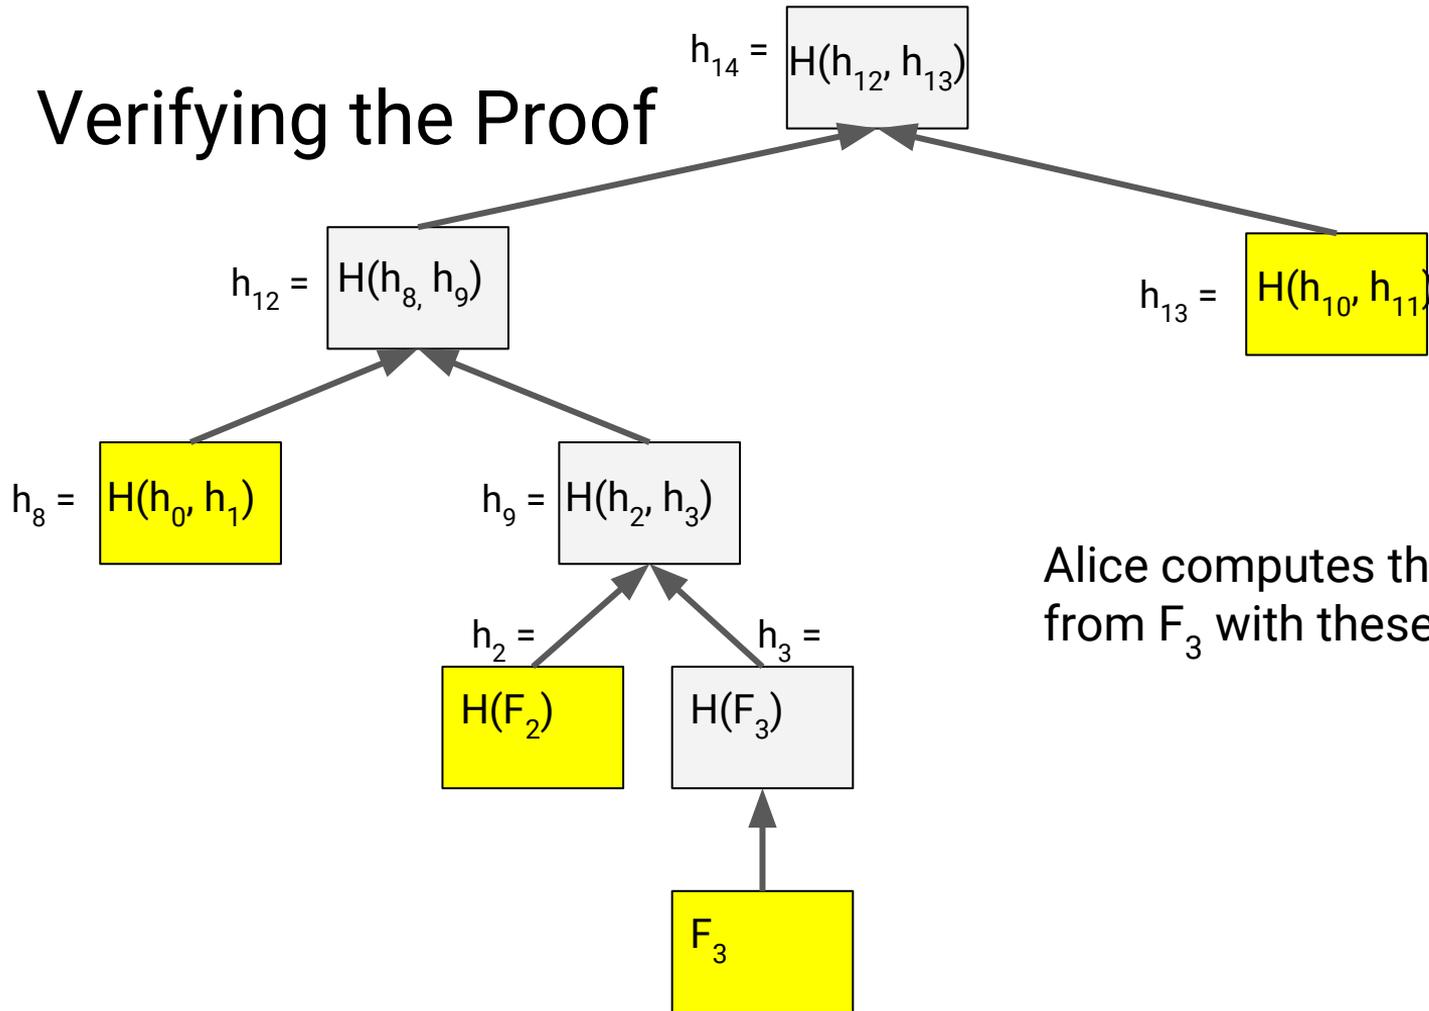
Dropbox sends these highlighted nodes.



Proving/Verifying Integrity: Merkle Tree



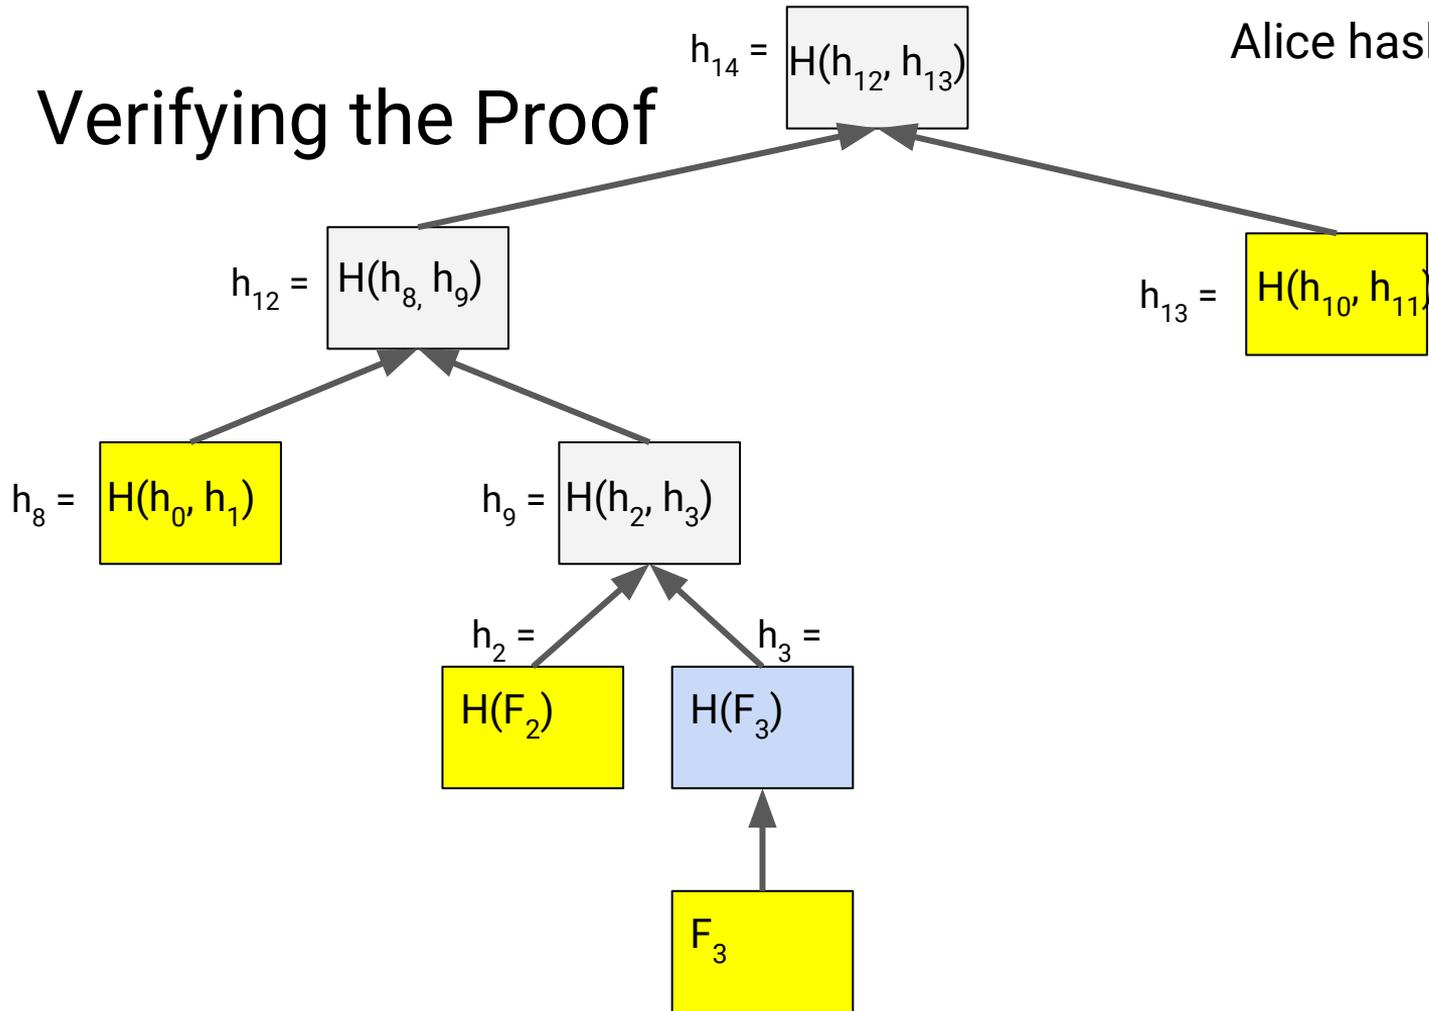
Verifying the Proof



Alice computes the root starting from F_3 with these highlighted proof.

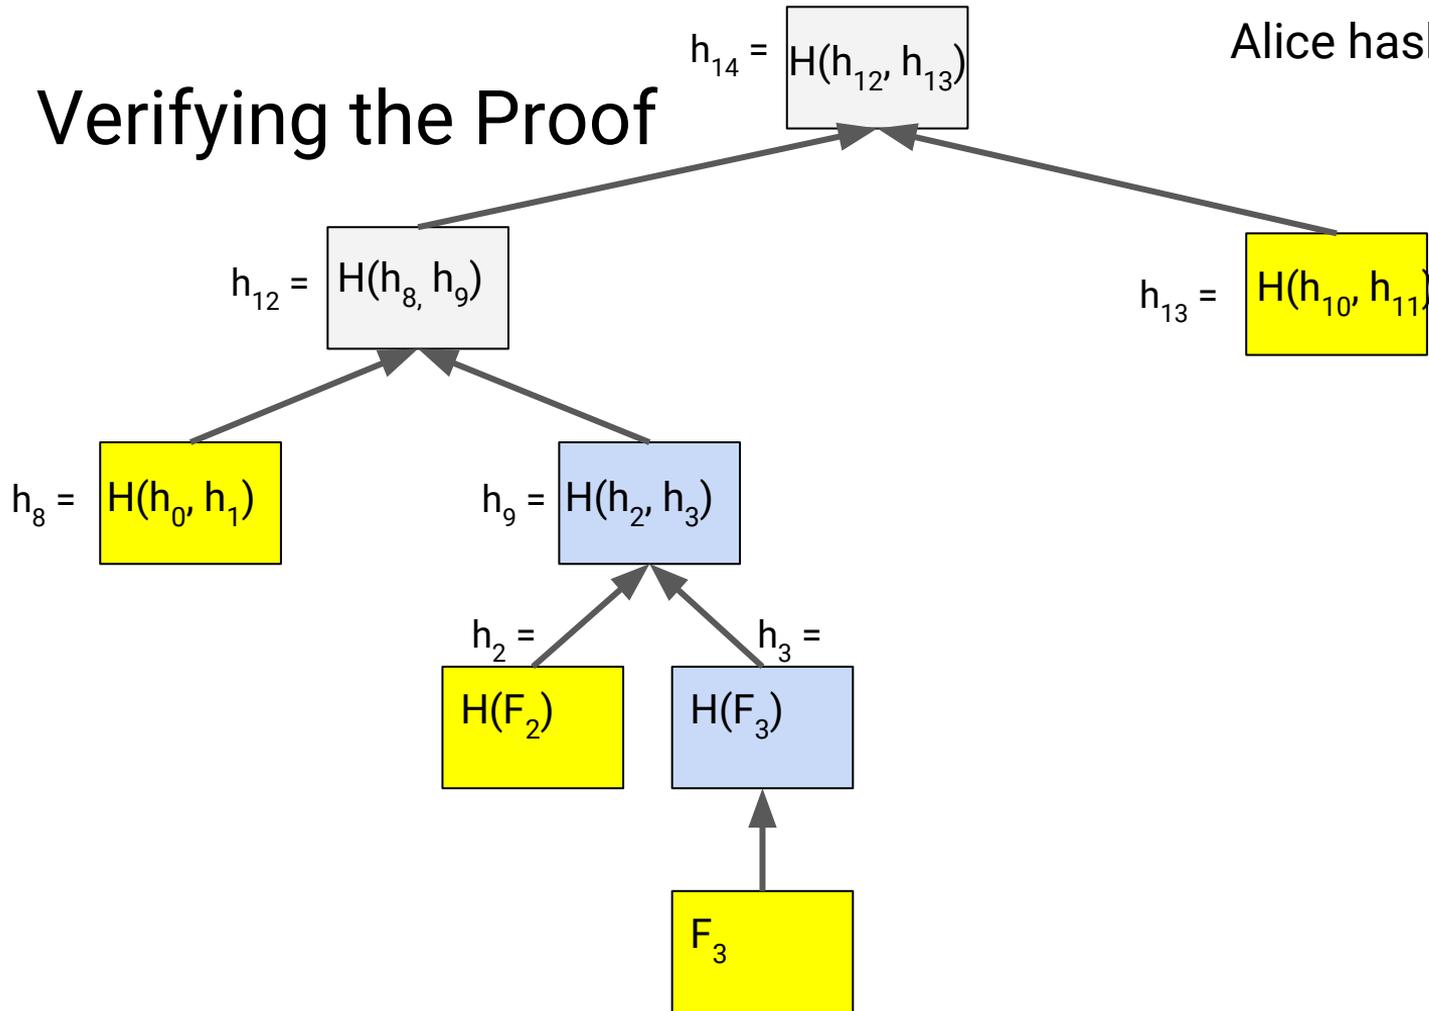
Verifying the Proof

Alice hashes up the tree.



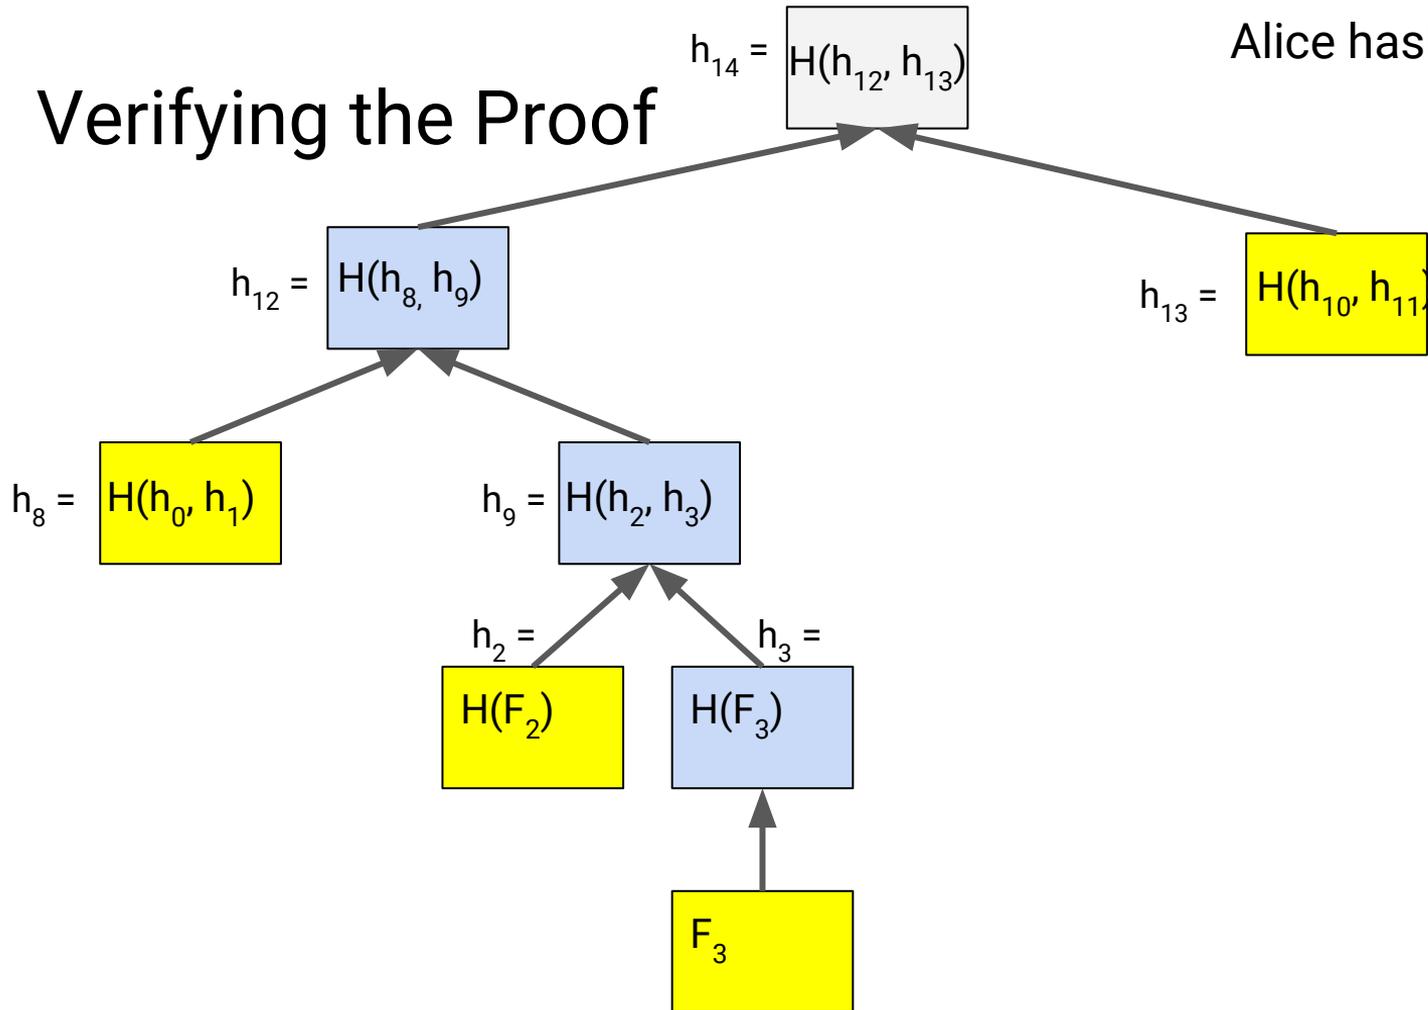
Verifying the Proof

Alice hashes up the tree.

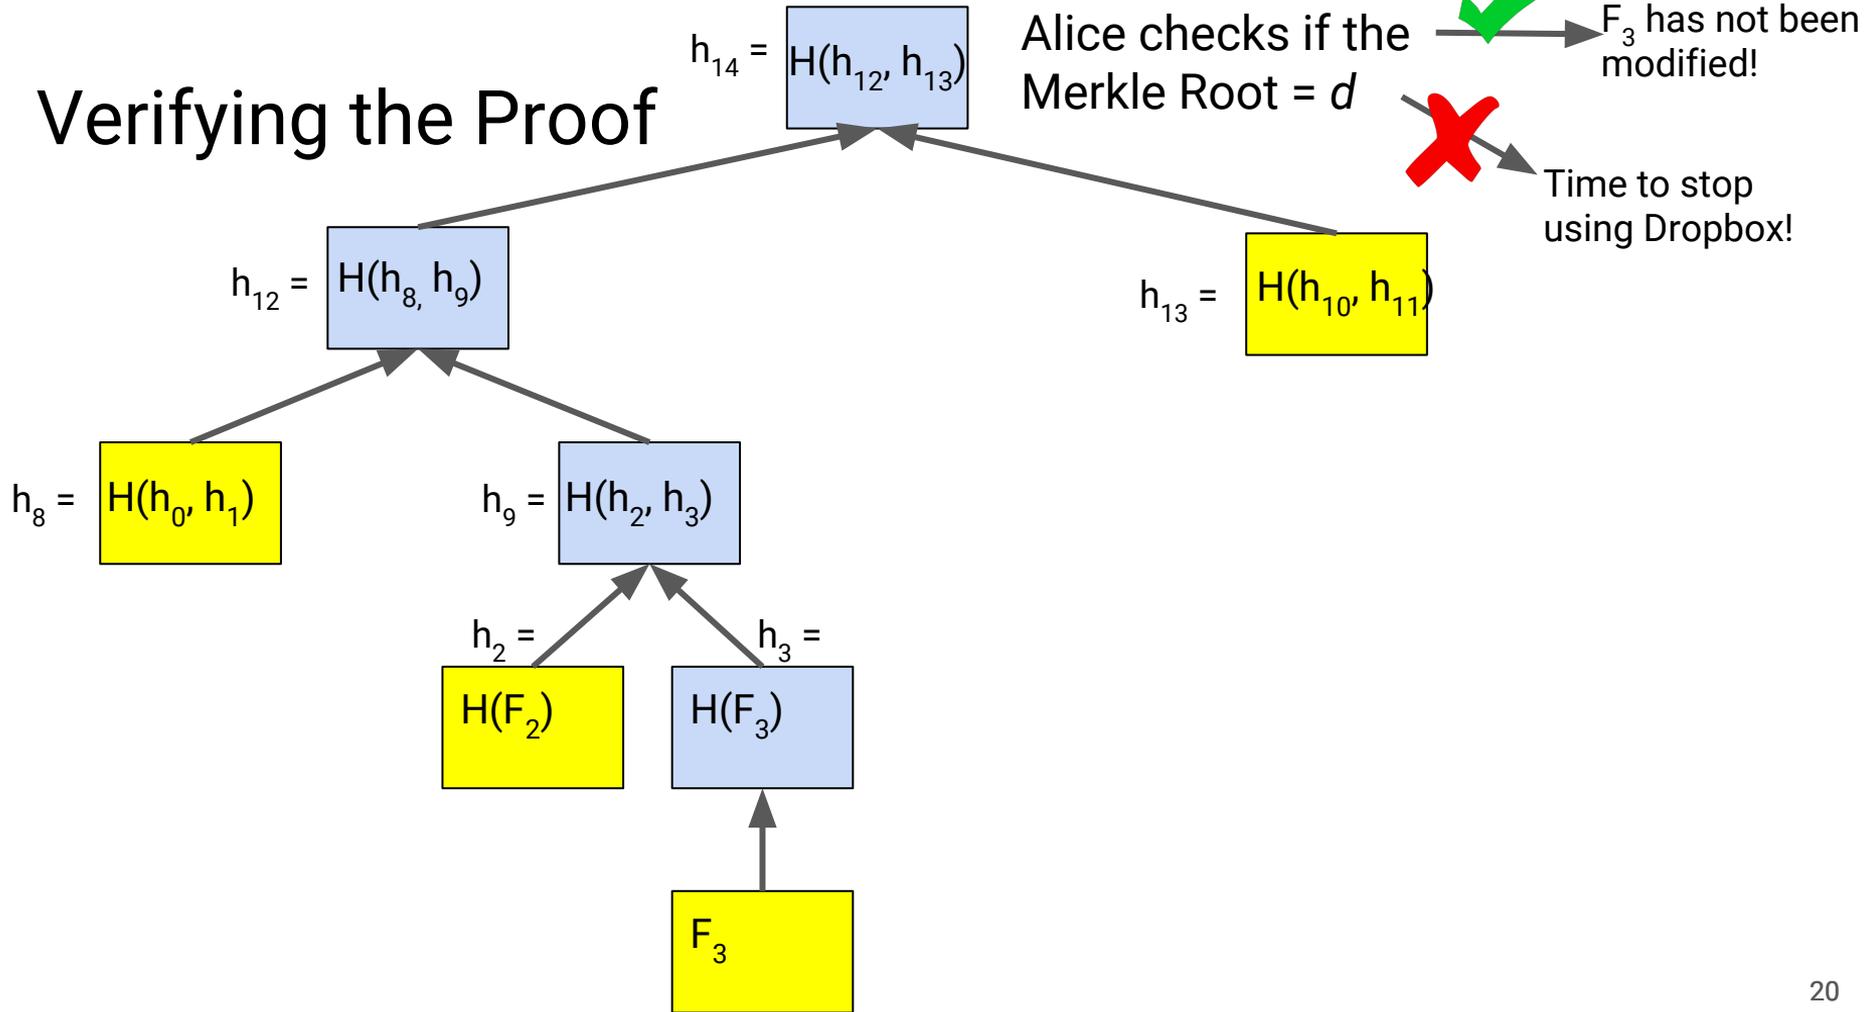


Verifying the Proof

Alice hashes up the tree.

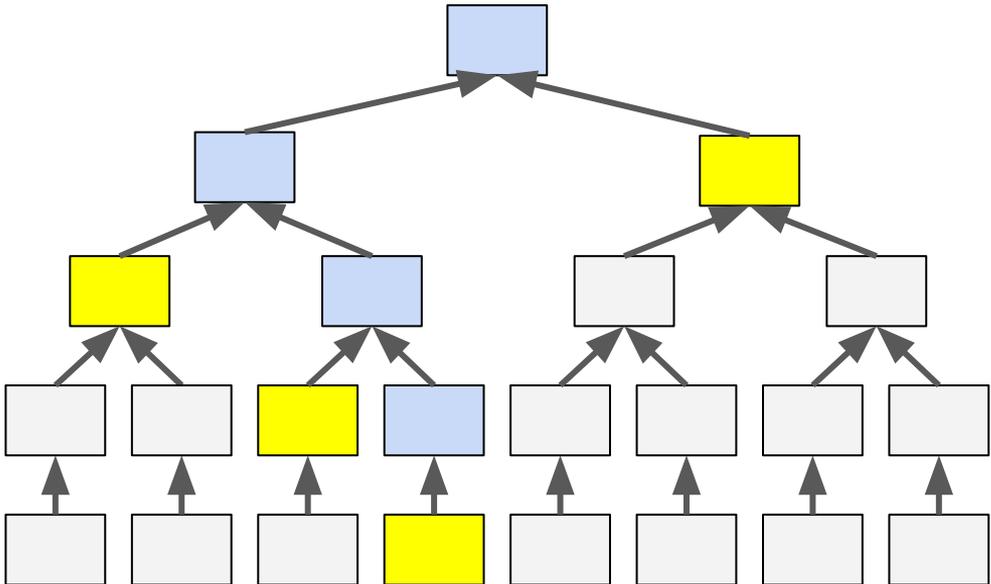


Verifying the Proof



Everyone loves Merkle Trees!

- They're beautiful.
- They're efficient.



n = number of leaves (files)

	Merkle Tree
Construct Tree	$O(n)$
Proof size	$O(\log n)$
Update File	$O(\log n)$

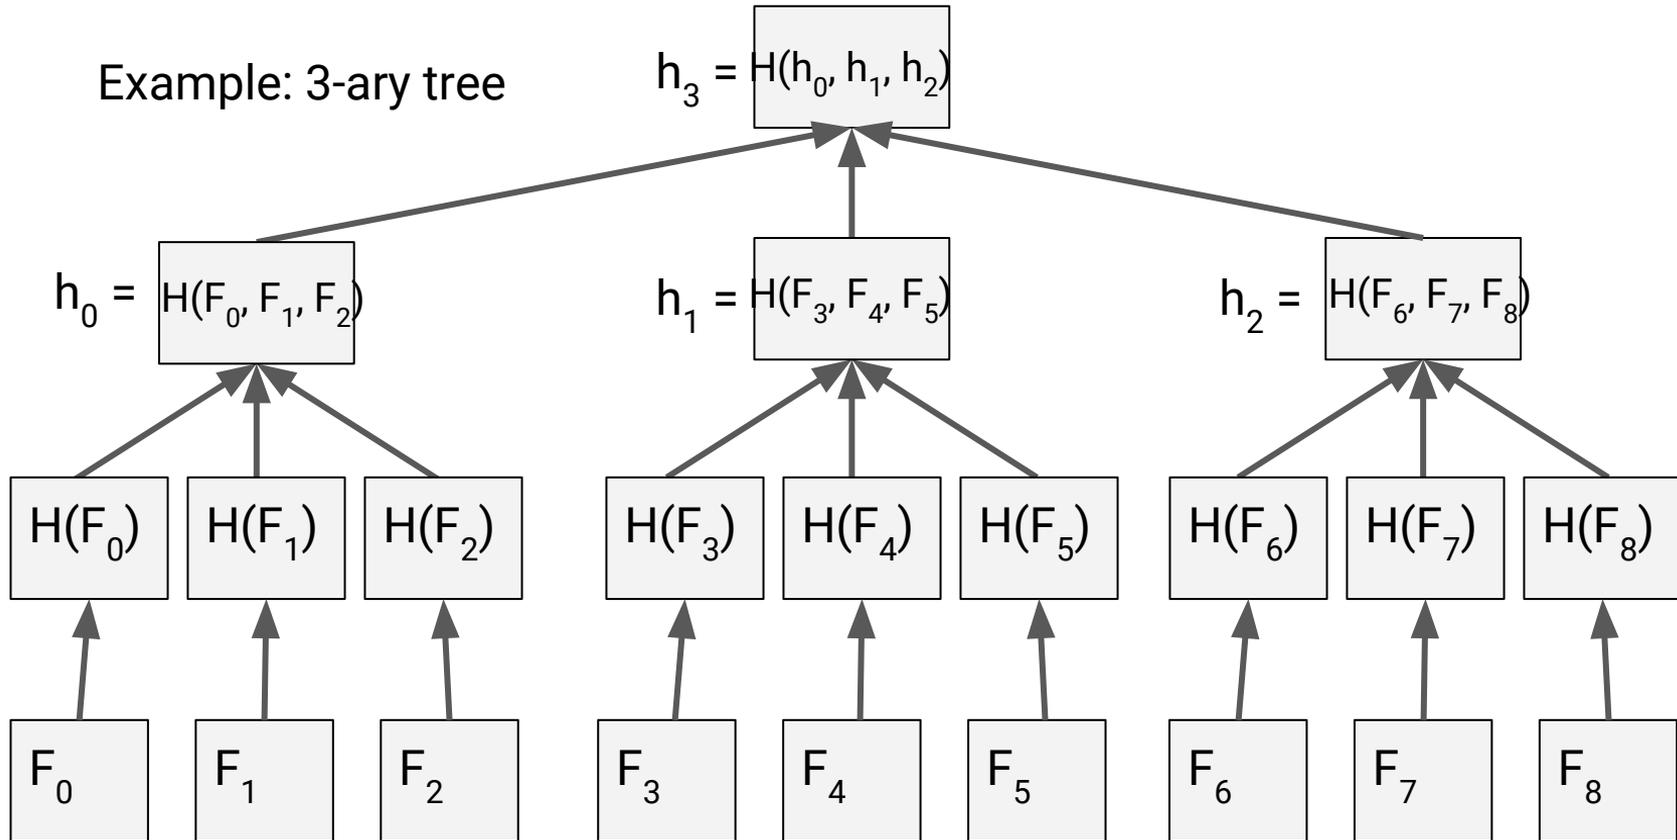
Problem: Many small files \Rightarrow Merkle proofs too large.

Problem: Many small files \Rightarrow Merkle proofs too large.

- Suppose Alice has one billion $\approx 2^{30}$ files.

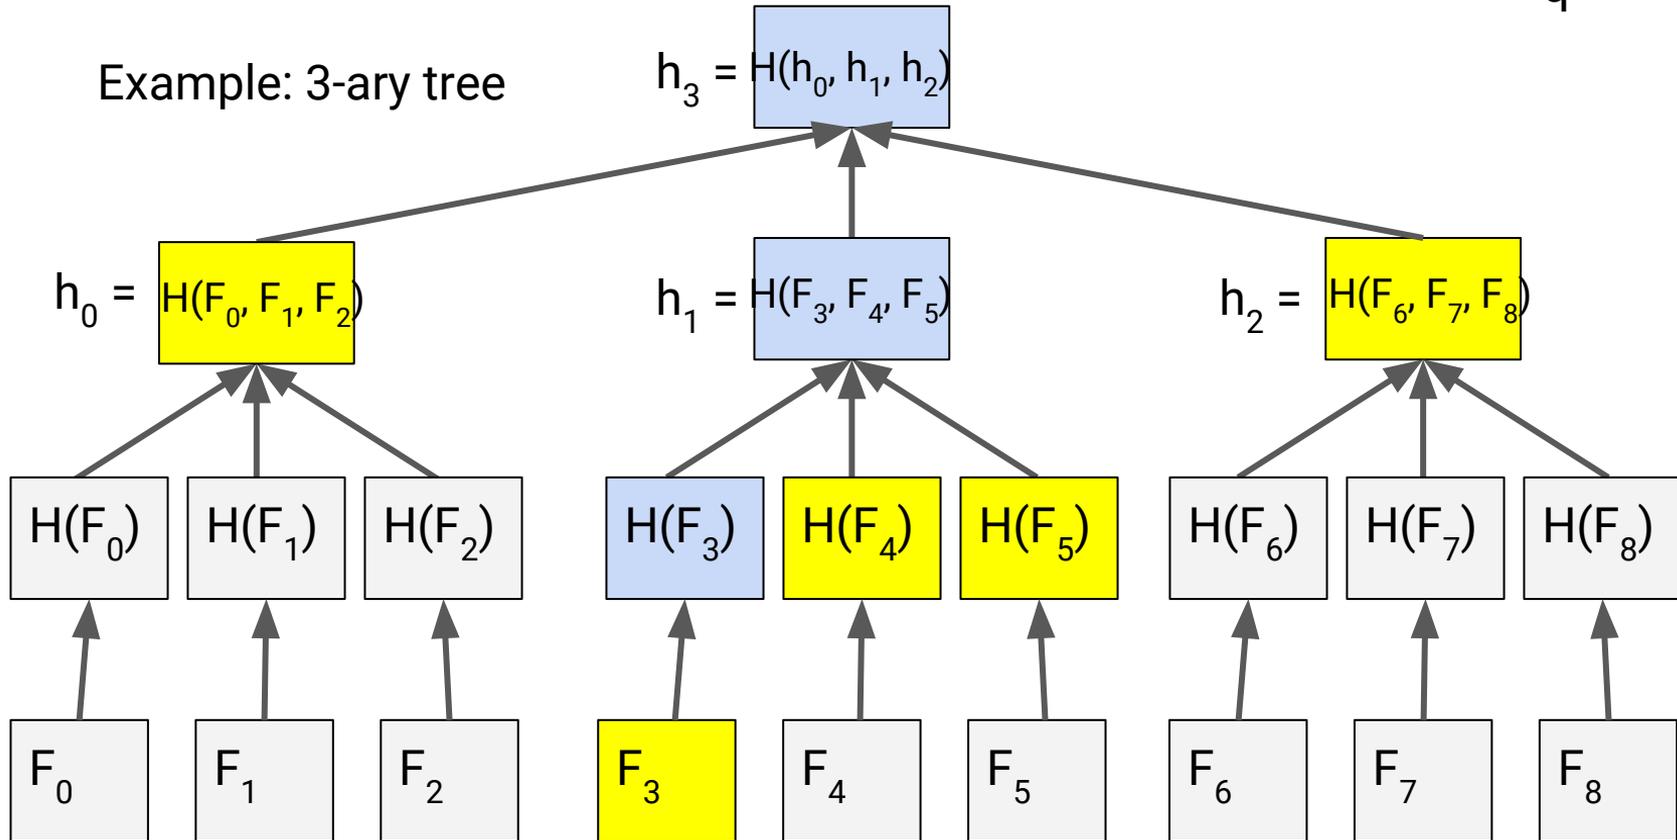
Possible Solution: q-ary Merkle Tree

Example: 3-ary tree



Problem: The Proof Becomes Bigger, $O(q \log_q n)$

Example: 3-ary tree



Our Work: Verkle Trees reduce the proof size

- We pick a q .
- We reduce the proof size from $\log_2 n$ to $\log_q n = \log_2 n / \log_2 q$.
- Factor of **$\log_2 q$ less bandwidth!**
- At the cost of **q times more computation**
- (e.g., $q = 1024 \Rightarrow \log_2(q) = 10x$ less bandwidth)

Wow, that's big!

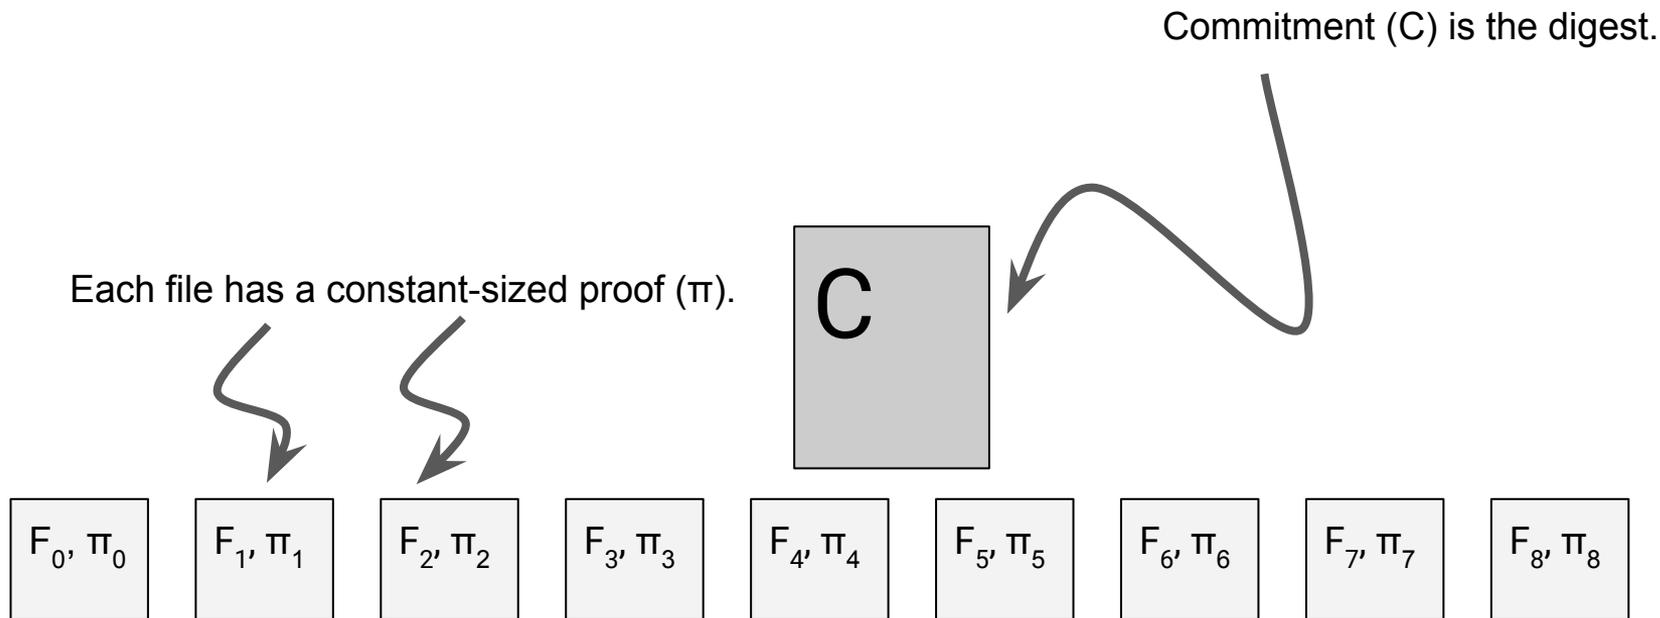


Does this matter? (Hint: Yes)

- Merkle hash trees are everywhere in cryptography:
 - Consensus Protocols
 - Public-Key Directories
 - Cryptocurrencies
 - Encrypted Web Applications
 - Secure File Systems



Vector Commitment (VC) Schemes by Catalano and Fiore (2013)

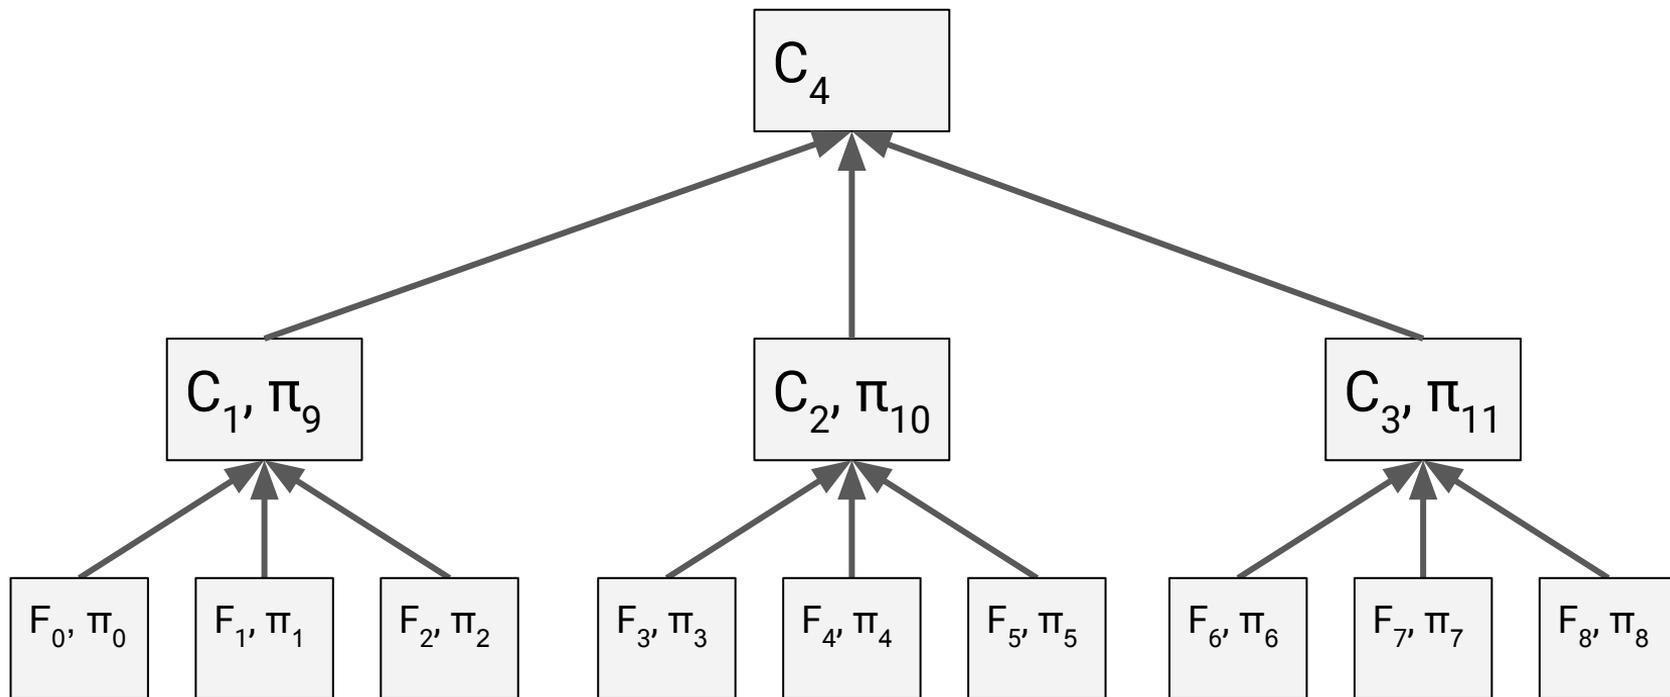


VC Schemes are Computationally Impractical

Scheme/op	Construct	Proof size
Merkle	$O(n)$	$O(\log_2 n)$
VC scheme	$O(n^2)$	$O(1)$

Our Solution: Replace Hash Functions with VC Schemes

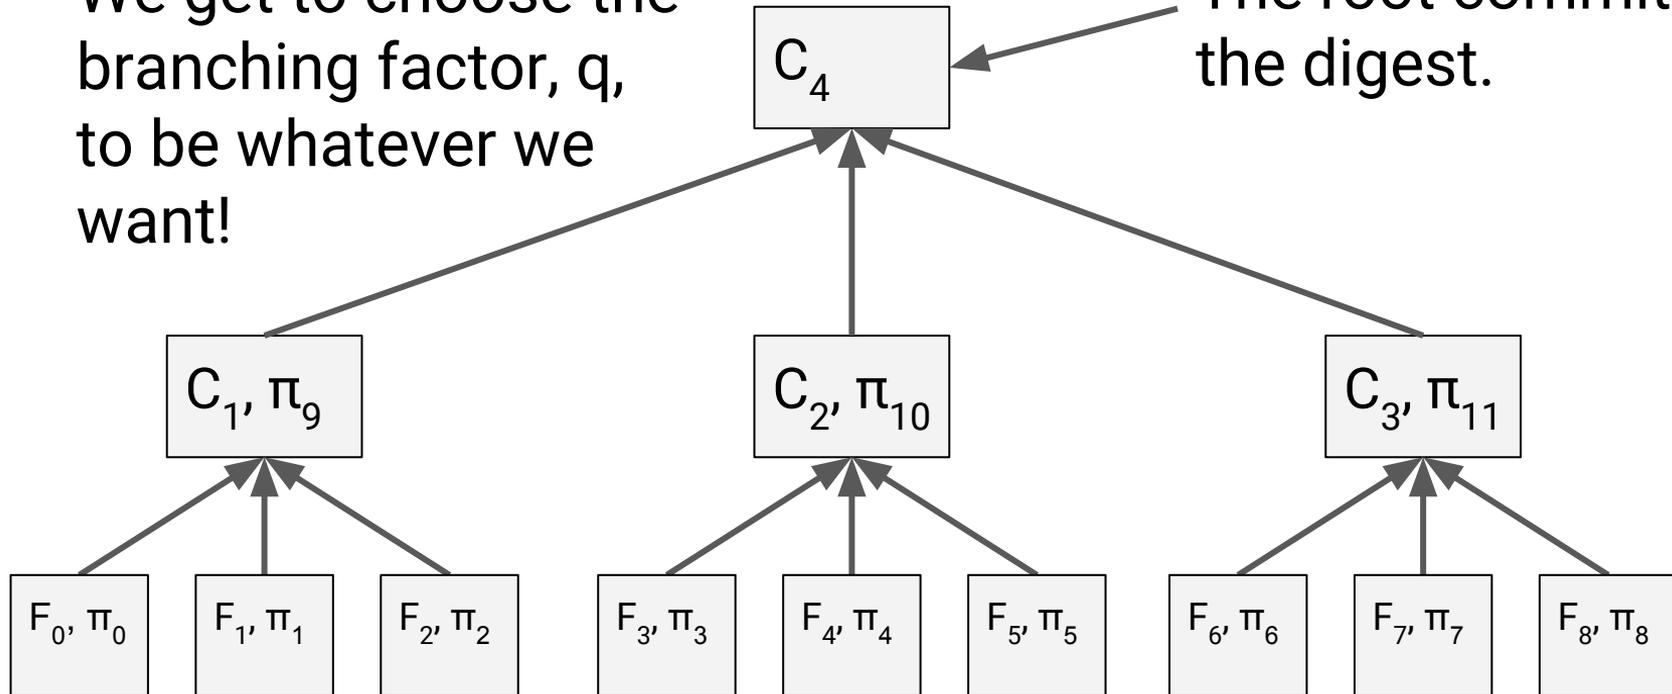
This is the Verkle Tree.



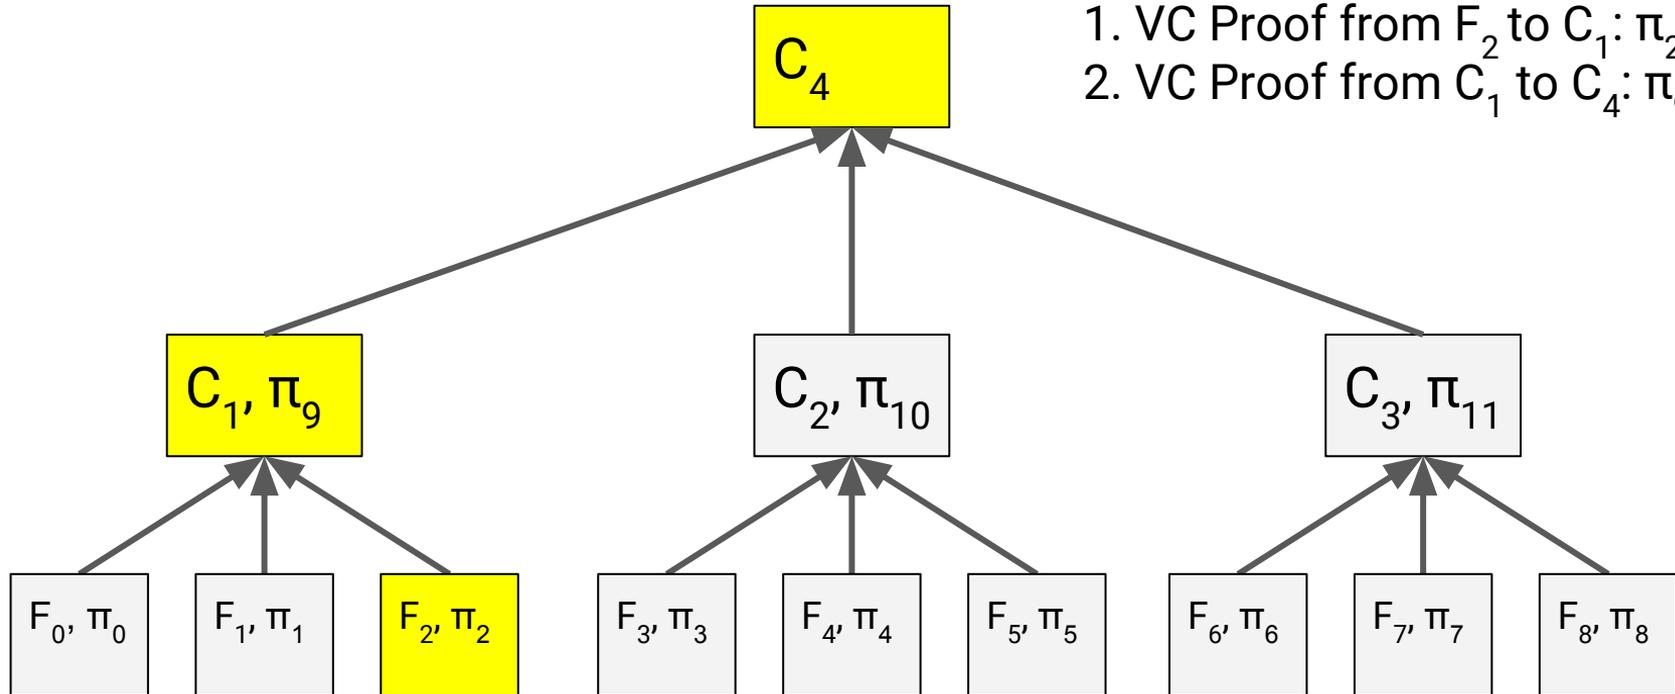
We now have a Verkle Tree!

We get to choose the branching factor, q , to be whatever we want!

The root commitment is the digest.



Alice Receives $\log_q n$ Constant-Sized π 's.



Alice verifies:

1. VC Proof from F_2 to C_1 : π_2
2. VC Proof from C_1 to C_4 : π_9

Comparison

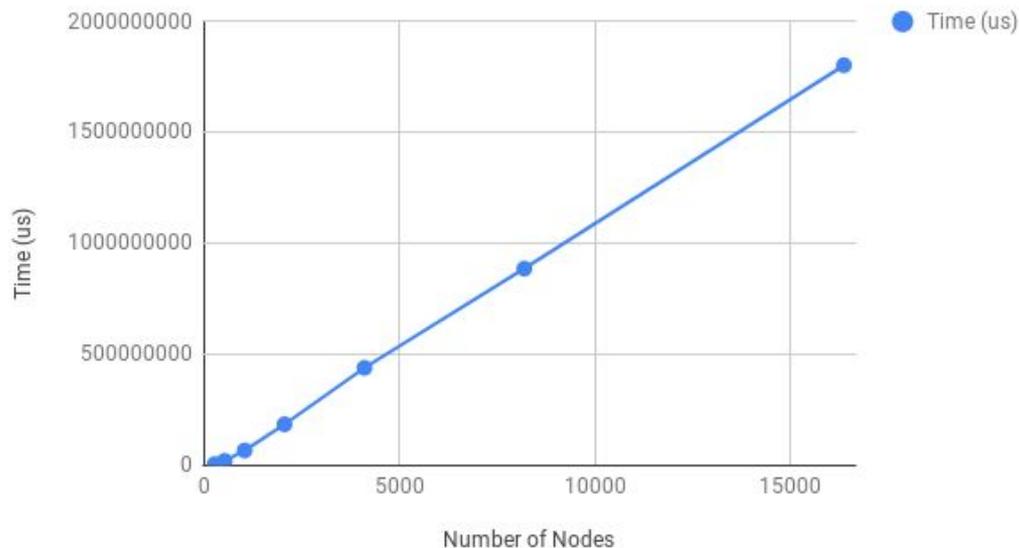
Scheme/op	Construct	Update file	Proof size
Merkle	$O(n)$	$O(\log_2 n)$	$O(\log_2 n)$
q-ary Merkle	$O(n)$	$O(q \log_q n)$	$O(q \log_q n)$
VC scheme	$O(n^2)$	$O(n)$	$O(1)$
q-ary Verkle	$O(qn)$	$O(q \log_q n)$	$O(\log_q n)$

Verkle Trees let us trade off proof-size vs. construction time.

My Contribution

- I proved complexity bounds for Verkle Trees.
- I implemented Verkle Trees in C++.
- I am measuring performance.

Verkle Tree Construction: $q = 1024$



Acknowledgements

- Thank you Alin!
- Thank you PRIMES!
- Thank you Mom and Dad!

