

CS251 Fall 2021
(cs251.stanford.edu)



Recursive SNARKs

Benedikt Bünz

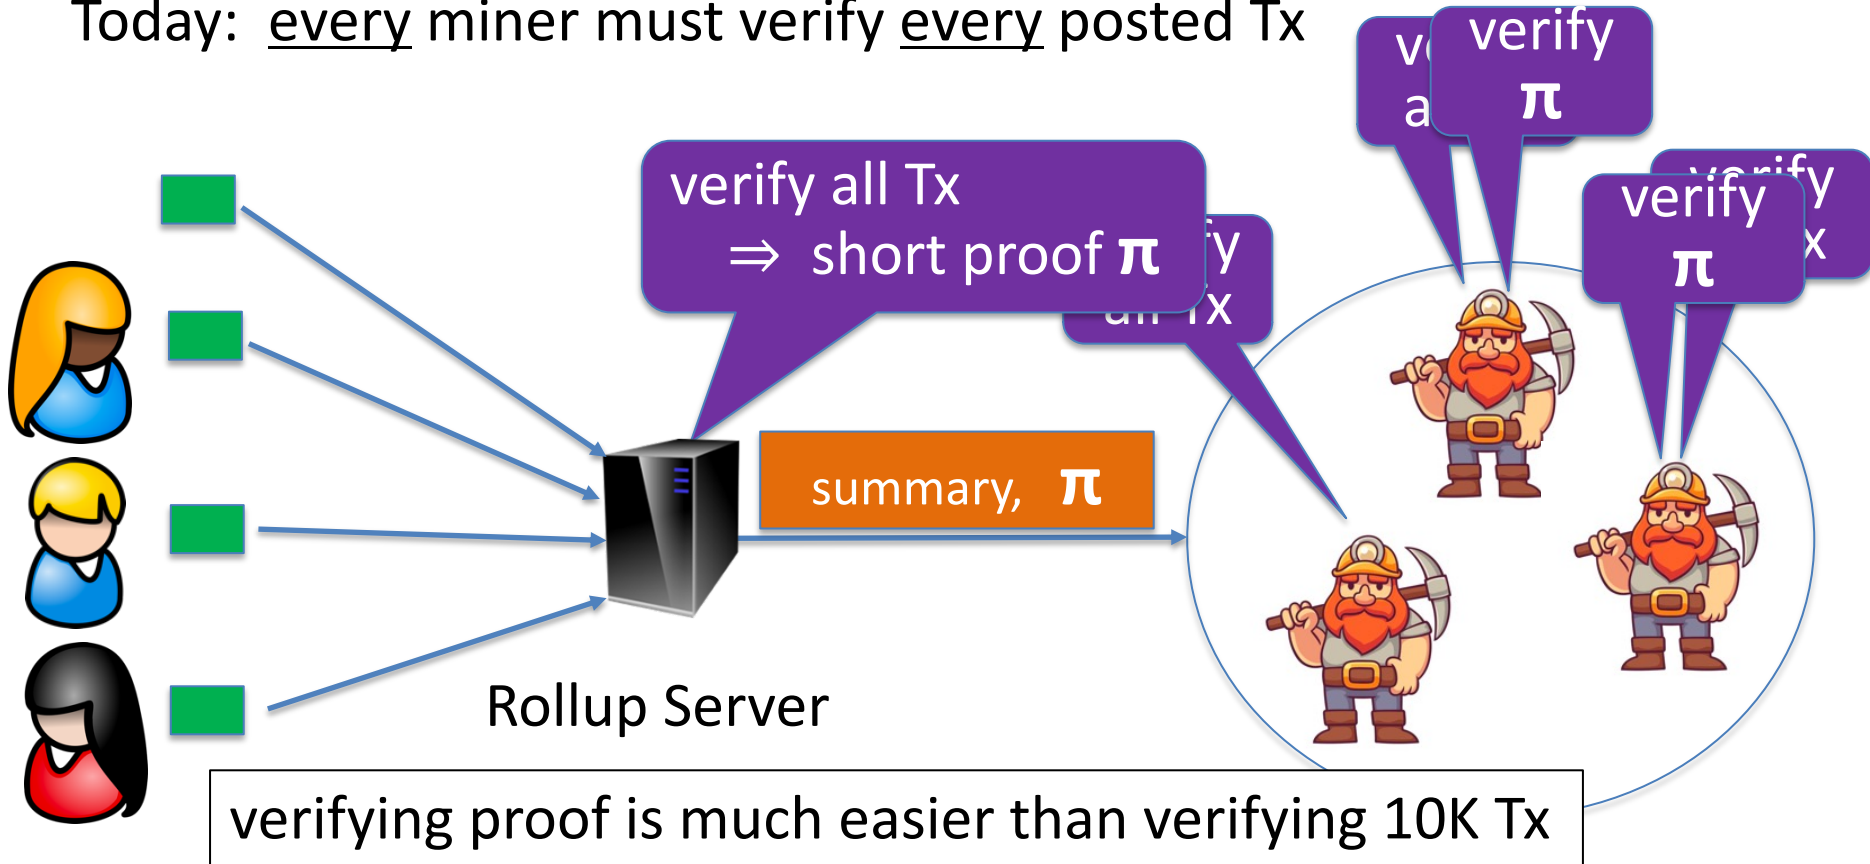
Recap: Non-interactive Proof Systems

A **non-interactive proof system** is a triple (S, P, V) :

- $S(C) \rightarrow$ public parameters (S_p, S_v) for prover and verifier
 (S_p, S_v) is called a *reference string*
- $P(S_p, \mathbf{x}, \mathbf{w}) \rightarrow$ proof π
- $V(S_v, \mathbf{x}, \pi) \rightarrow$ accept or reject

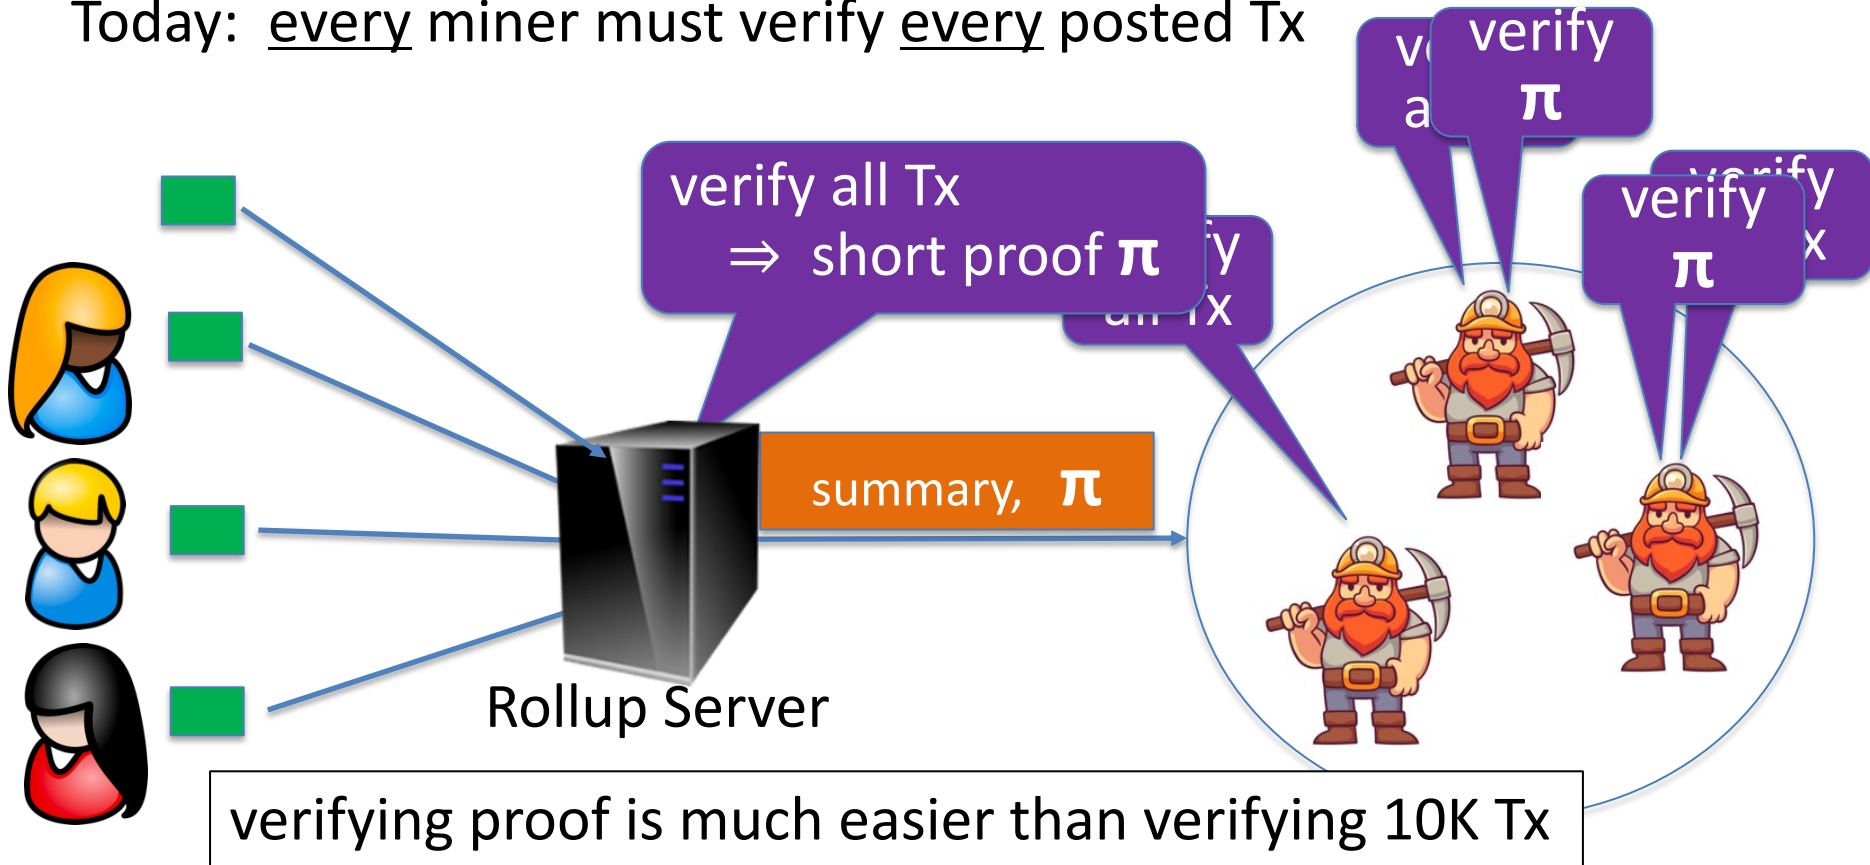
Recap: zkRollup

Today: every miner must verify every posted Tx

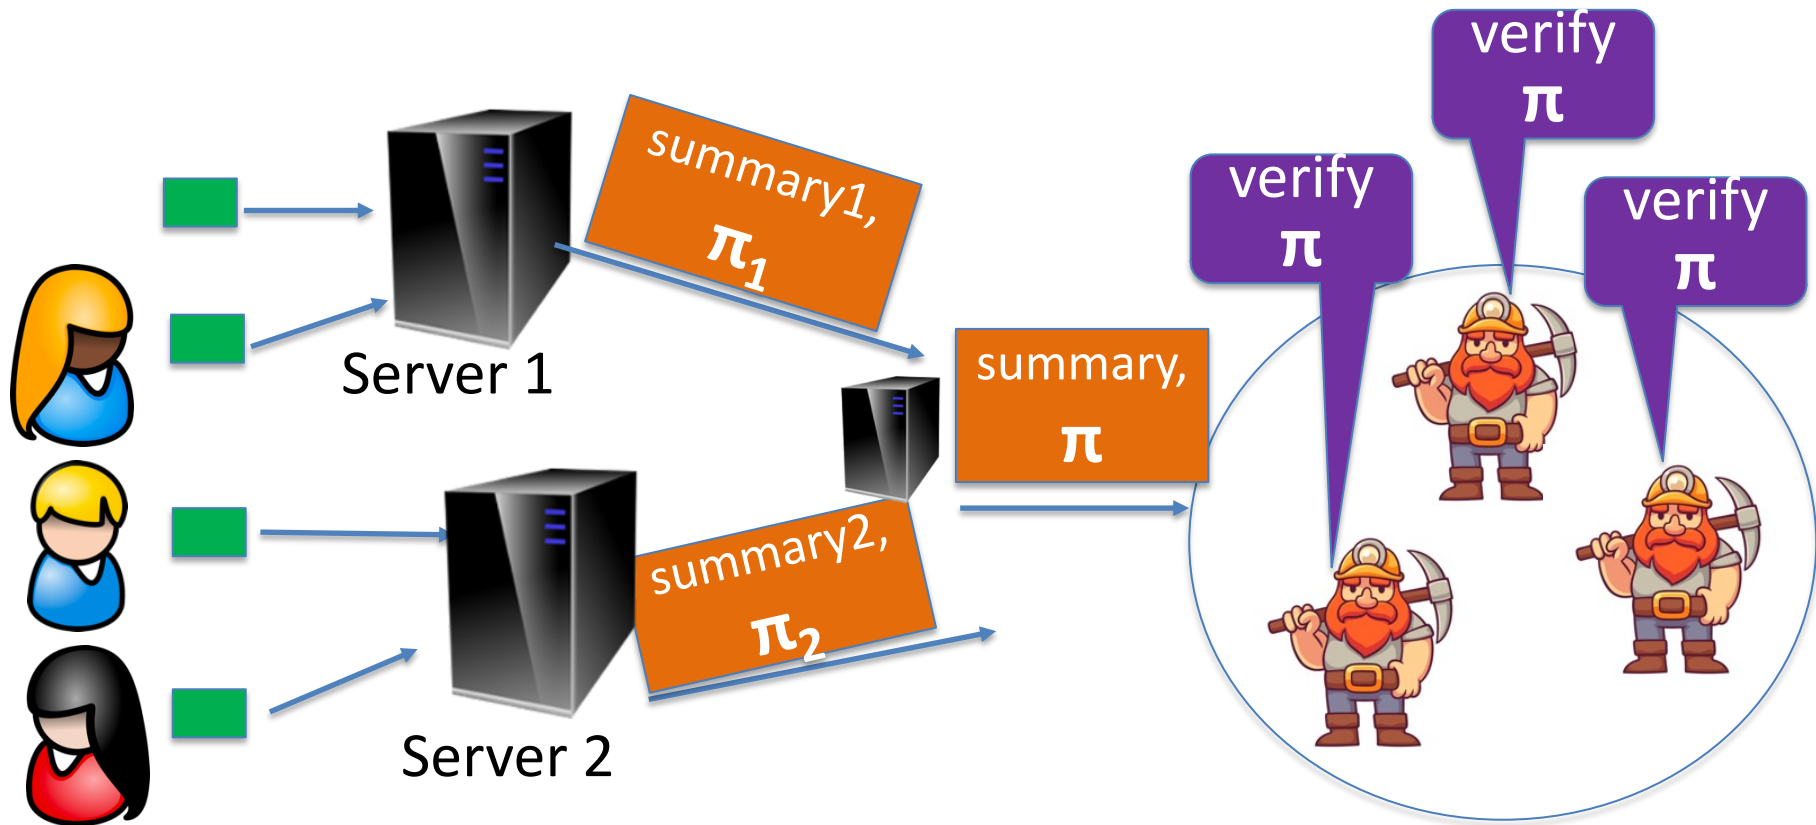


Recap: zkRollup

Today: every miner must verify every posted Tx



Rollup with many coordinators



Zk-zk-Rollup

- Multiple servers
- Each responsible for subset of users (no overlaps)
- Rollup aggregator (can be one of the servers)
- Rollup aggregator combines summaries (balance table) and creates one proof that
- How can we combine proofs?
- Trivial solution:
 - All servers forward all Tx
 - Rollup aggregator creates one SNARK
 - Does not save work

Recap: Non-interactive Proof Systems

A **non-interactive proof system** is a triple (S, P, V) :

- $S(C) \rightarrow$ public parameters (S_p, S_v) for prover and verifier
 (S_p, S_v) is called a *reference string*
- $P(S_p, \mathbf{x}, \mathbf{w}) \rightarrow$ proof π
- $V(S_v, \mathbf{x}, \pi) \rightarrow$ accept or reject

SNARK of a SNARK Proof

A **non-interactive proof system** is a triple (S, P, V) :

- $S(C) \rightarrow$ public parameters (S_p, S_v) for prover and verifier
 (S_p, S_v) is called a *reference string*
- $P(S_p, \mathbf{x}, \mathbf{w}) \rightarrow$ proof π
- $V(S_v, \mathbf{x}, \pi) \rightarrow$ accept or reject

SNARK of SNARK

How can we aggregate proofs?

$$S(C) \rightarrow S_P, S_V$$
$$\pi \leftarrow P(S_P, x, w)$$

Now write a circuit C' that verifies π :

- Input x' is x
- Witness w' is π
- $C'(x', w') = 0$ iff $V(S_V, \pi, x) = \text{Accept}$

Finally:

$$S(C') \rightarrow S'_P, S'_V$$
$$\pi' \leftarrow P(S'_P, x', w')$$

SNARK of SNARKs

How can we aggregate proofs?

$$S(C) \rightarrow S_P, S_V$$

$$\pi_1 \leftarrow P(S_P, x_1, w_1) \quad \pi_2 \leftarrow P(S_P, x_2, w_2)$$

Now write a circuit C' that verifies π :


- Input x' is $x_1 || x_2$
- Witness w' is $\pi_1 || \pi_2$
- $C'(x', w') = 0$ iff $V(S_V, x_1, \pi_1) = \text{Accept}$ and $V(S_V, x_2, \pi_2) = \text{Accept}$

Finally:

$$S(C') \rightarrow S'_P, S'_V$$
$$\pi' \leftarrow P(S'_P, x', w')$$

SNARK of SNARKs

- Note that C' depends only on \mathbf{V} and S_v (not on C_1, C_2)
- We can express \mathbf{V} as a circuit:

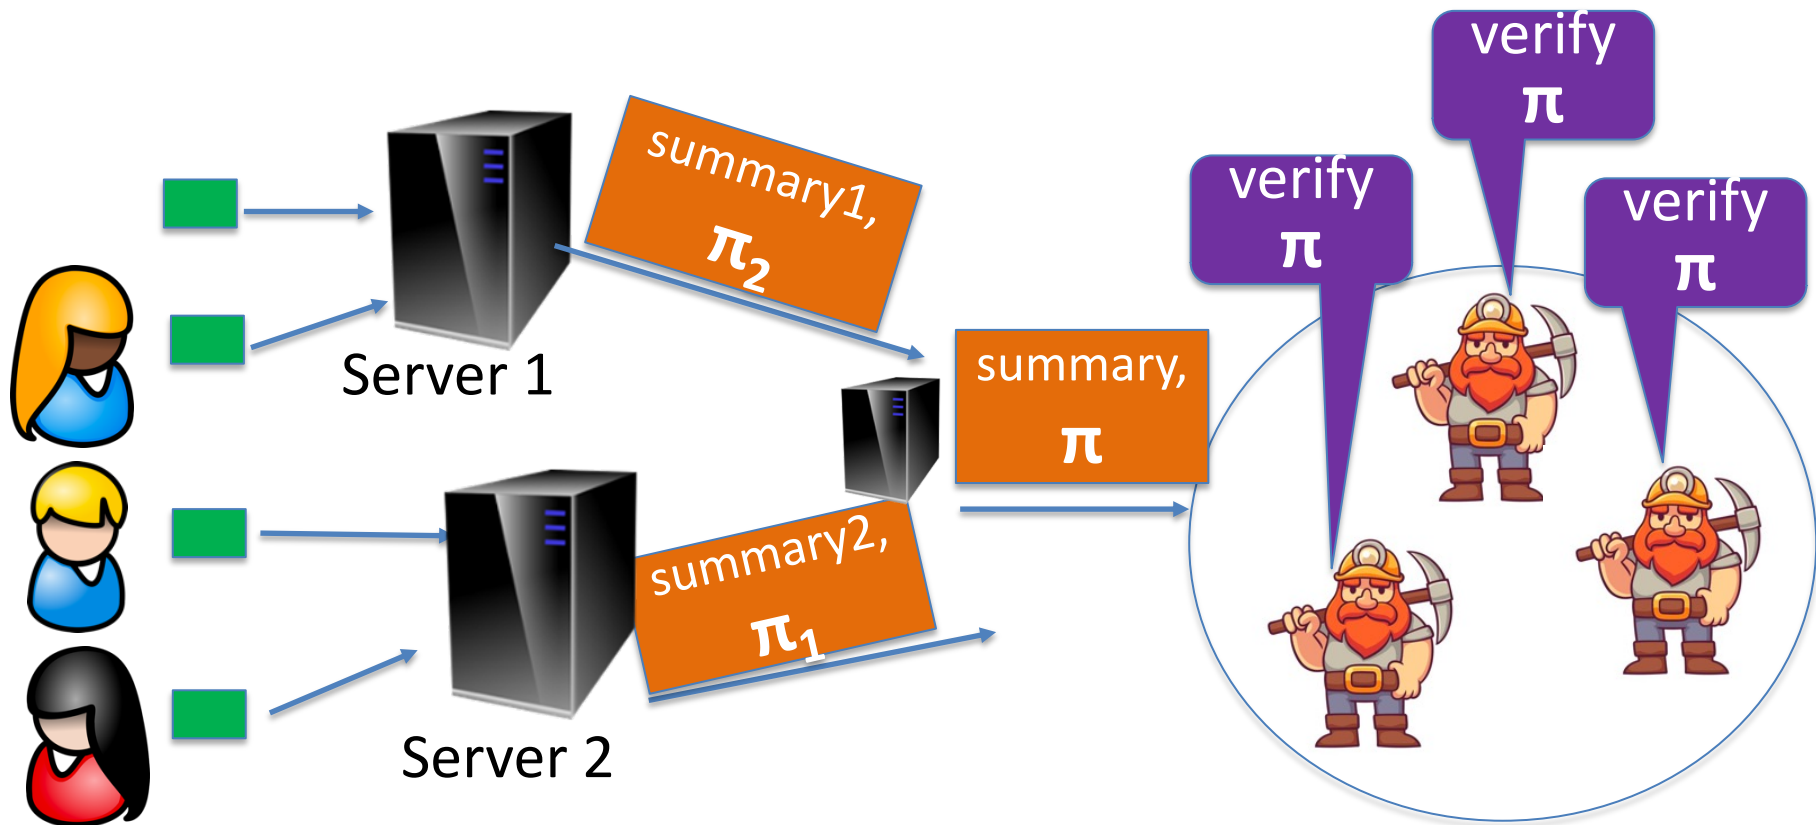


$w' = \pi'$ is independent of w_1, w_2
 $|C'| = 2 * |V| < |2 * C|$

Building SNARK of SNARKs

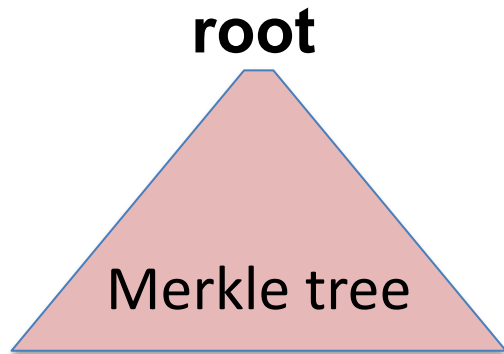
- How big is C' ?
- Comparison |SHA256 circuit| = 20k gates
- First SNARK of SNARK ~ 1 million gates with trusted setup (BCTV14)
- Today less than 50k gates (Halo, BCLMS20, Nova)
 - no trusted setup
- Independent of inner SNARK circuits!

Rollup with many coordinators



Zk-zk-Rollup

- Let \mathbf{root}_i be the Merkle Tree Root of summary i



Build one root from summaries

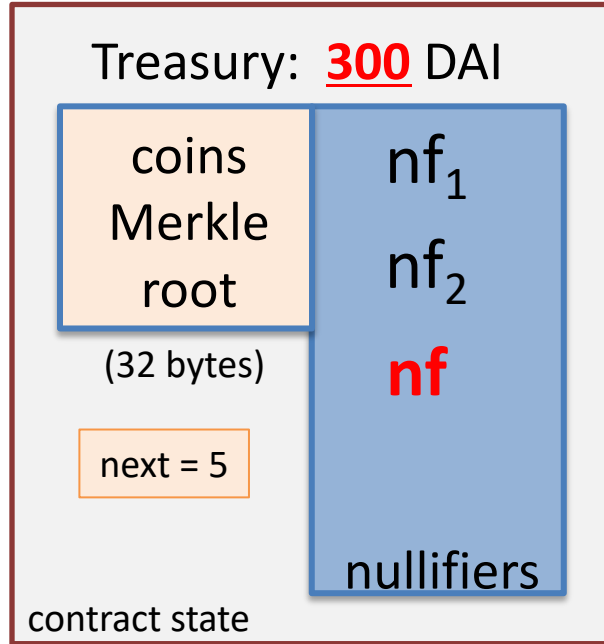
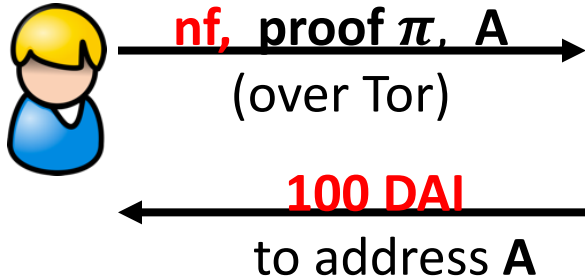
- $S_V, S_P \leftarrow \mathcal{S}(C_R) \quad // \quad C_R \text{ rollup circuit}$
 $root_1 \quad root_2 \quad root_3 \quad root_4$

$C_{zk^2}(x = S_V, \mathbf{root}; w = root_1, root_2 \dots, \pi_1, \pi_2, \dots):$
 $\mathbf{V}(S_V, x = root_i, \pi_i)$ for all i and $\mathbf{root} = \text{MT}(root_i)$


Tornado cash

100 DAI pool:
each coin = 100 DAI

Withdraw coin #3
to addr A:



$$H_1, H_2: \mathbb{R} \rightarrow \{0,1\}^{256}$$

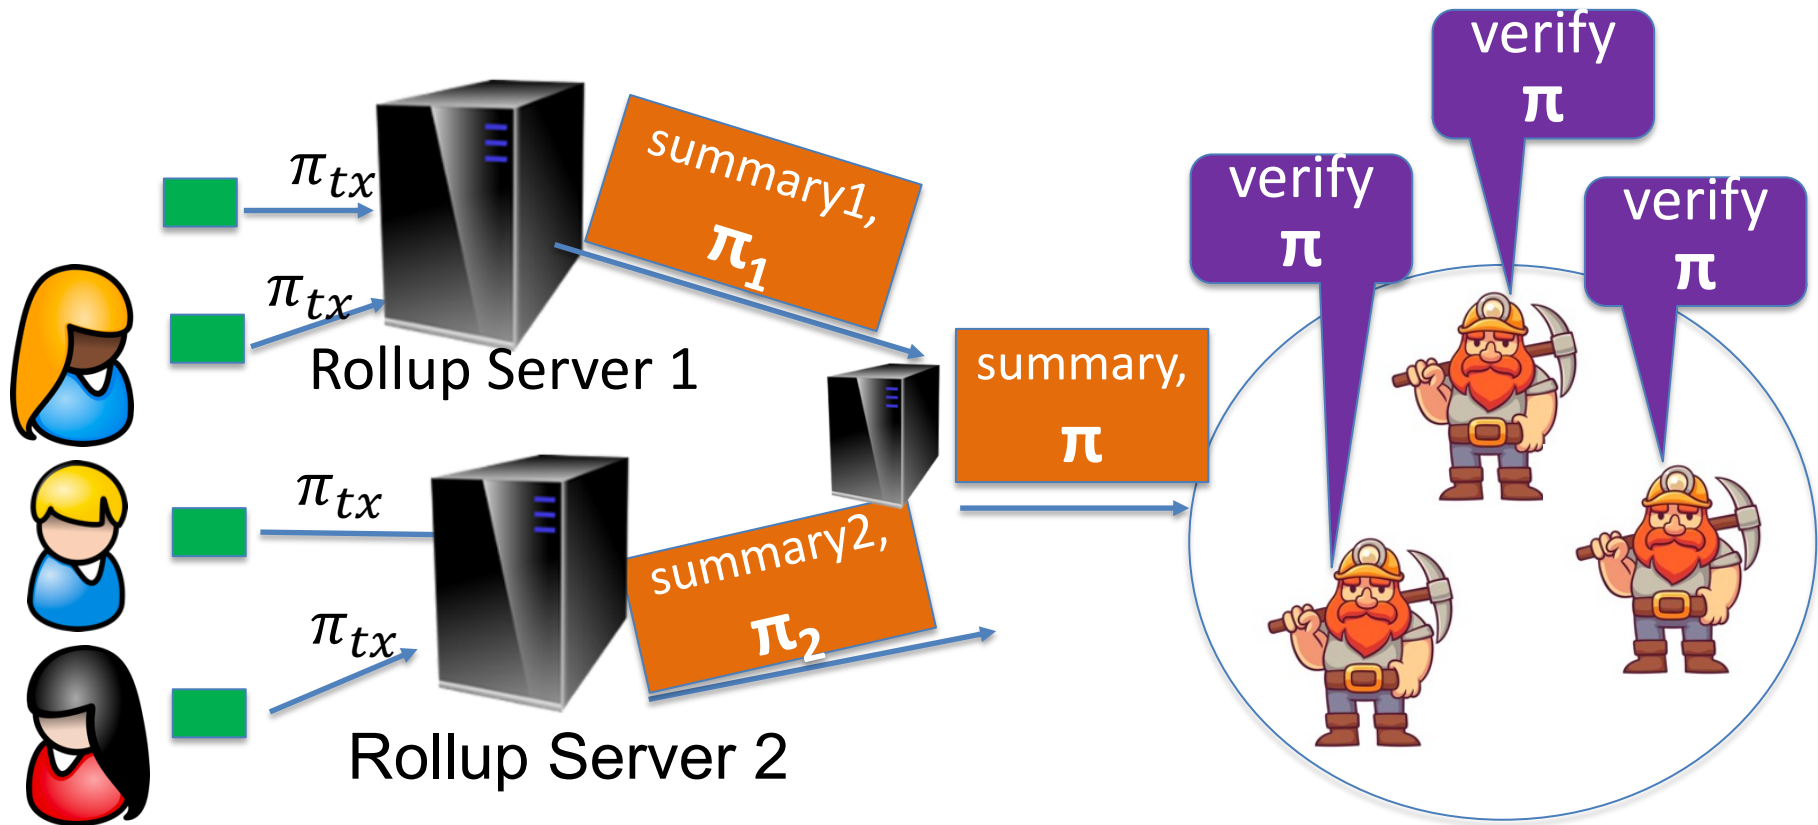


public list of coins
... but observer does not know which are spent

nf and π reveal nothing about which coin was spent.

But, coin #3 cannot be spent again, because $nf = H_2(k')$ is now nullified.

zk³-Rollup (tornado cash rollup)



zk³-Rollup

- Users create SNARK for TC Circuit C_{TC}
 - $S_V, S_P \leftarrow \mathcal{S}(C_{TC})$
 - $\pi_{TC} \leftarrow P(S_P, tx, w)$
- Rollups create SNARKs for $C_R = \forall_i V(S_V, tx_i, \pi_i) = \text{“accept”}$
 - $tx\ root = MT(tx_1, \dots, tx_n)$
 - $\pi' = \pi_{TC,1} || \dots || \pi_{TC,n}$
 - $S_V', S_P' \leftarrow \mathcal{S}(C_R)$
 - $\pi_R = P(S_P', tx\ root, \pi')$
- Rollup Aggregator creates SNARK for $C_A = \forall_i V(S_V', root_i, \pi_{R,i})$
 - $S_V'', S_P'' \leftarrow \mathcal{S}(C_A)$
 - $root = MT(root_1, \dots, root_k)$
 - $\pi_R' = \pi_{R,1} || \dots || \pi_{R,k}$
 - $\pi_A = P(S_P'', root, \pi_R')$

Enhancing transactions with SNARKs

- We've seen that private transactions require zero-knowledge proofs
- Add ZK-SNARKs to every transaction
- Level 1 coordinators verify transaction by verifying transaction ZK-SNARKs
- Additionally, we can have more complicated transactions (Smart Contracts)
 - Transaction verification is constant time regardless of proof complexity
- *Can we also hide the smart contract?*

ZEXE private execution

- ZEXE is a model of computation (like UTXOs/Scripts or Accounts/EVM)
- The basic unit is a record (similar to a UTXO)
- Every transaction consumes records and creates records
- Universal predicate: Prevents double spends
- Birth predicate: Says how a record can be created
- Death predicate: Says how a record can be consumed

ZEXE private execution

Record 1:
Birth predicate 1
Death predicate 1
Payload 1

Record 2:
Birth predicate 1
Death predicate 1
Payload 1




Record 3:
Birth predicate 3
Death predicate 3
Payload 3

TX checks that Record 1 and Record 2 have not been spent
 $\text{Birth3}(R1, R2, R3)$ and $\text{Death1}(R1, R2, R3)$ and $\text{Death2}(R1, R2, R3)$

ZEXE private execution

- Universal predicate (similar to tornado cash)
 - Uses nullifiers
 - Checks that nullifier= $H(\text{sk}, \text{records})$ is properly created
 - Checks that nullifier only appears once
 - Prevents double spends



Implementing assets with ZEXE

- Record payload has a value v and an asset id
- Birth predicate
 - Defines the token
 - New record id needs to match consumed predicate ids
 - New record value is sum of inputs
- Death predicate
 - Defines the SCRIPT
 - E.g. spendable by signature
 - E.g. Spendable by multisignature + preimage of hash

Implementing smart contracts with ZEXE

- Record payload is state of smart contract, smart contract instance id
- Birth predicate
 - Either creates smart contract or
 - One of the inputs needs to be the old smart contract record
- Death predicate
 - Defines the smart contract logic

ZEXE game of Chess

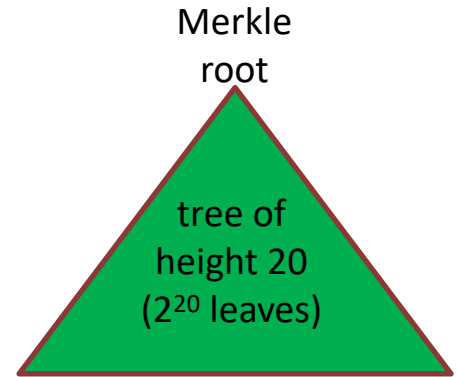
- Record payload is state of smart contract, smart contract instance id
- Birth predicate
 - Starts new game (and assigns pks to black/white) or
 - One of the inputs needs to be the old chess game
- Death predicate
 - If game finished then pay money to the winner
 - Otherwise input records must be game record + one move record
 - Move record must be signed by the right player
 - Move record must contain a valid move

Making ZEXE private

- $S_{P_U}, S_{V_U} \leftarrow S(C_U)$ (Universal predicate)
- $S_{P_B}, S_{V_B} \leftarrow S(C_B)$ (Birth predicate)
- $S_{P_D}, S_{V_D} \leftarrow S(C_D)$ (Death predicate)
- $S_{P_{TX}}, S_{V_{TX}} \leftarrow S(C_{TX})$ (TX circuit)
- $C_{TX} = V(S_{V_U}, \dots) = 0$ and $V(S_{V_B}, \dots) = 0$ and $V(S_{V_D}, \dots) = 0$

And Record = $H(\text{payload}, S_{V_B}, S_{V_D}, r)$ // r random

- TX: Input records || Output records
- Compute nullifiers nf_1, \dots, nf_n from input records
- To create a TX, create three ZK-SNARKS (now ZK is important)
 - $x = \text{TX}$, $w = \text{payloads}, S_{V_B}, S_{V_D}$
 - $\pi_U \leftarrow P(S_{P_U}, x \mid |nf_1, \dots, nf_n, w \mid | \text{MT proofs})$
 - $\pi_B \leftarrow P(S_{P_B}, x, w)$
 - $\pi_D \leftarrow P(S_{P_D}, x, w)$
- Create $\pi_{TX} \leftarrow P(S_{P_{TX}}, x, w \mid | \pi_U, \pi_B, \pi_D)$




$R_1 R_2 R_3 \dots 0 0 0$

MT of all records


Birth and death
predicate as well as
records are private!

Hitchhikers guide to the galaxy



What if we want to verify that computation?

Input




Output (42)

Long Computation

SNARKs for long computations

Issues:

- P** takes very long
- Starts after proving *after* computation finished
- Can't hand off computation
- S** also runs at least linear in $|C|$
(ok if many proofs)




C – Circuit for long computation

$$\mathbf{S}(C) \rightarrow (S_p, S_v)$$

$$\mathbf{x} = (\text{input}, \text{output})$$

$$\mathbf{w} = \text{transcript}$$

Input



Long Computation, Transcript

Output (42)

$$\mathbf{P}(S_p, \mathbf{x}, \mathbf{w}) \rightarrow \pi$$

$$\mathbf{V}(S_v, \mathbf{x}, \pi) \rightarrow \text{accept}$$

Handing off computation

C_I – Circuit for long intermediate computation

$\mathbf{S}(C_I) \rightarrow (\mathbf{S}_p, \mathbf{S}_v)$

$\mathbf{x}_1 = (\text{input}, \text{int}_1), \mathbf{w}_1 = \text{transcript}_1$

$\mathbf{x}_2 = (\text{int}_1, \text{int}_2), \mathbf{w}_2 = \text{transcript}_2$

$\mathbf{x}_3 = (\text{int}_2, \text{output}), \mathbf{w}_3 = \text{transcript}_3$

$\mathbf{P}(\mathbf{S}_p, \mathbf{x}_i, \mathbf{w}_i) \rightarrow \pi_i$

$\mathbf{V}(\mathbf{S}_v, \mathbf{x}_1, \pi_1)$

$\mathbf{V}(\mathbf{S}_v, \mathbf{x}_2, \pi_2)$

$\mathbf{V}(\mathbf{S}_v, \mathbf{x}_3, \pi_3)$

$|\pi| / \mathbf{V}$ linear in
#handoffs



Input

Int₁, π₁

Int₂, π₂

Output (42), π₃

transcript₁

transcript₂

transcript₃

Incremental Proofs

- We need updatable/incremental proofs

C_I – Circuit per computation step, t number of steps/handoffs

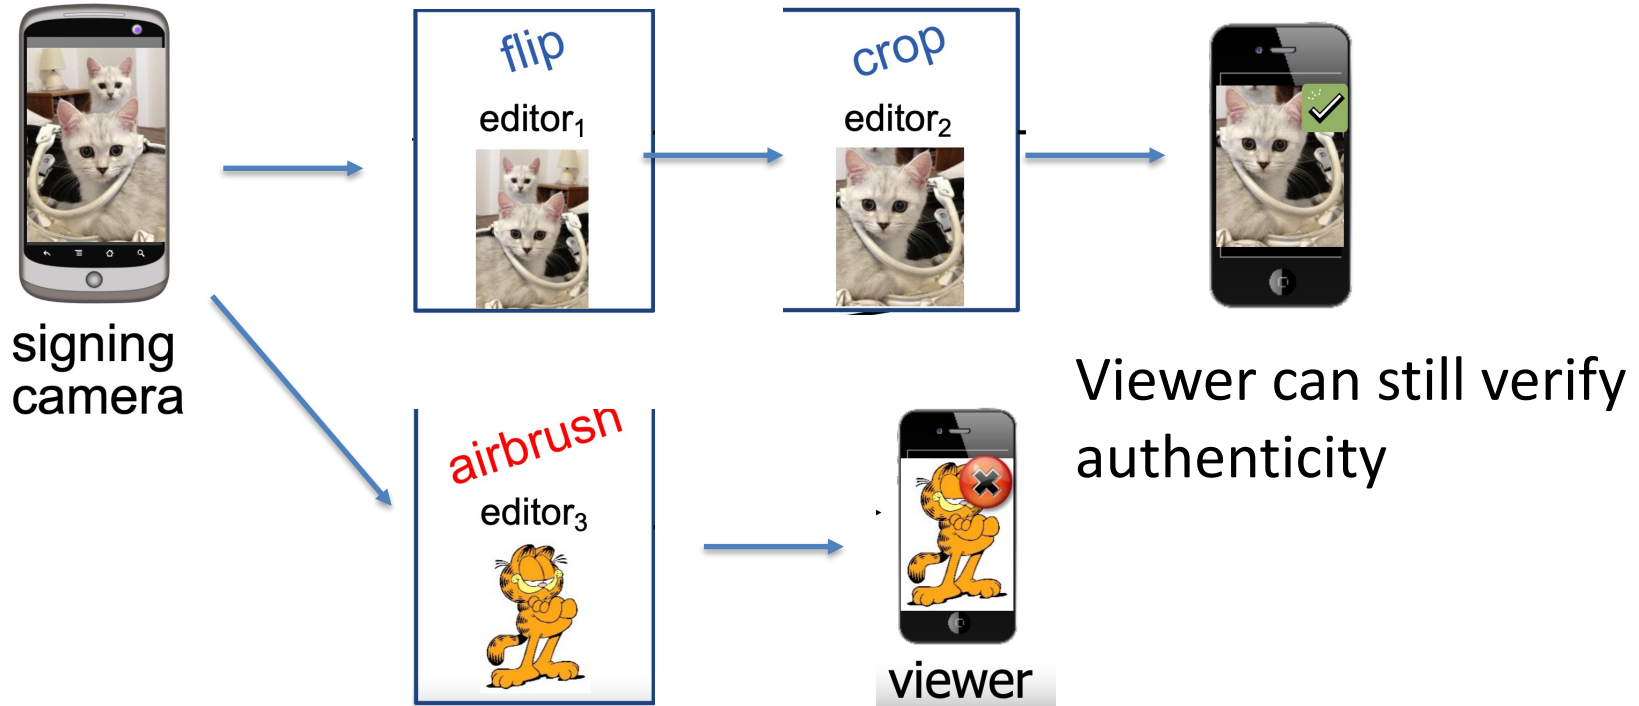
$\mathbf{S}(C_I) \rightarrow (\mathbf{S}_p, \mathbf{S}_v)$

$\mathbf{P}(\mathbf{S}_p, \mathbf{x}_i, \mathbf{w}_i, \pi_{i-1}) \rightarrow$ updated proof π_i // $\pi_0 = \perp$

$\mathbf{V}(\mathbf{S}_v, \mathbf{x}_0, \mathbf{x}_t, \pi_t, t) \rightarrow$ accept/reject


$|\pi_i| = |\pi_{i-1}|$ // proofs don't grow

PhotoProof



Allow valid updates of photo and provide proof

PhotoProof

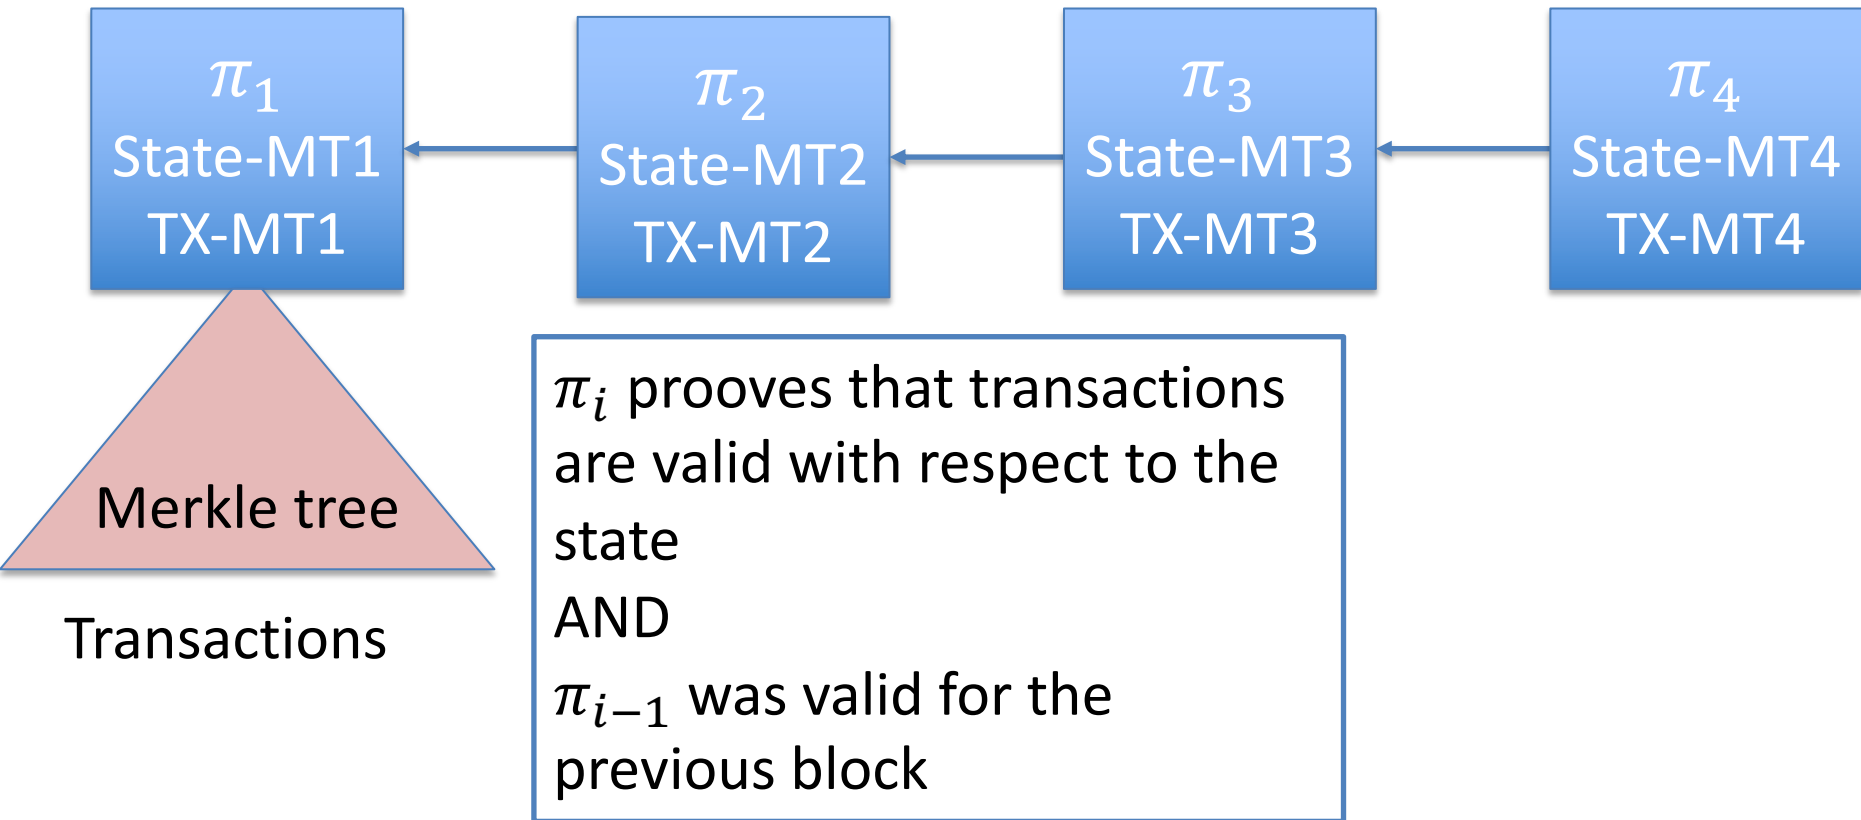


Proof allows valid edits only, Incrementally updated

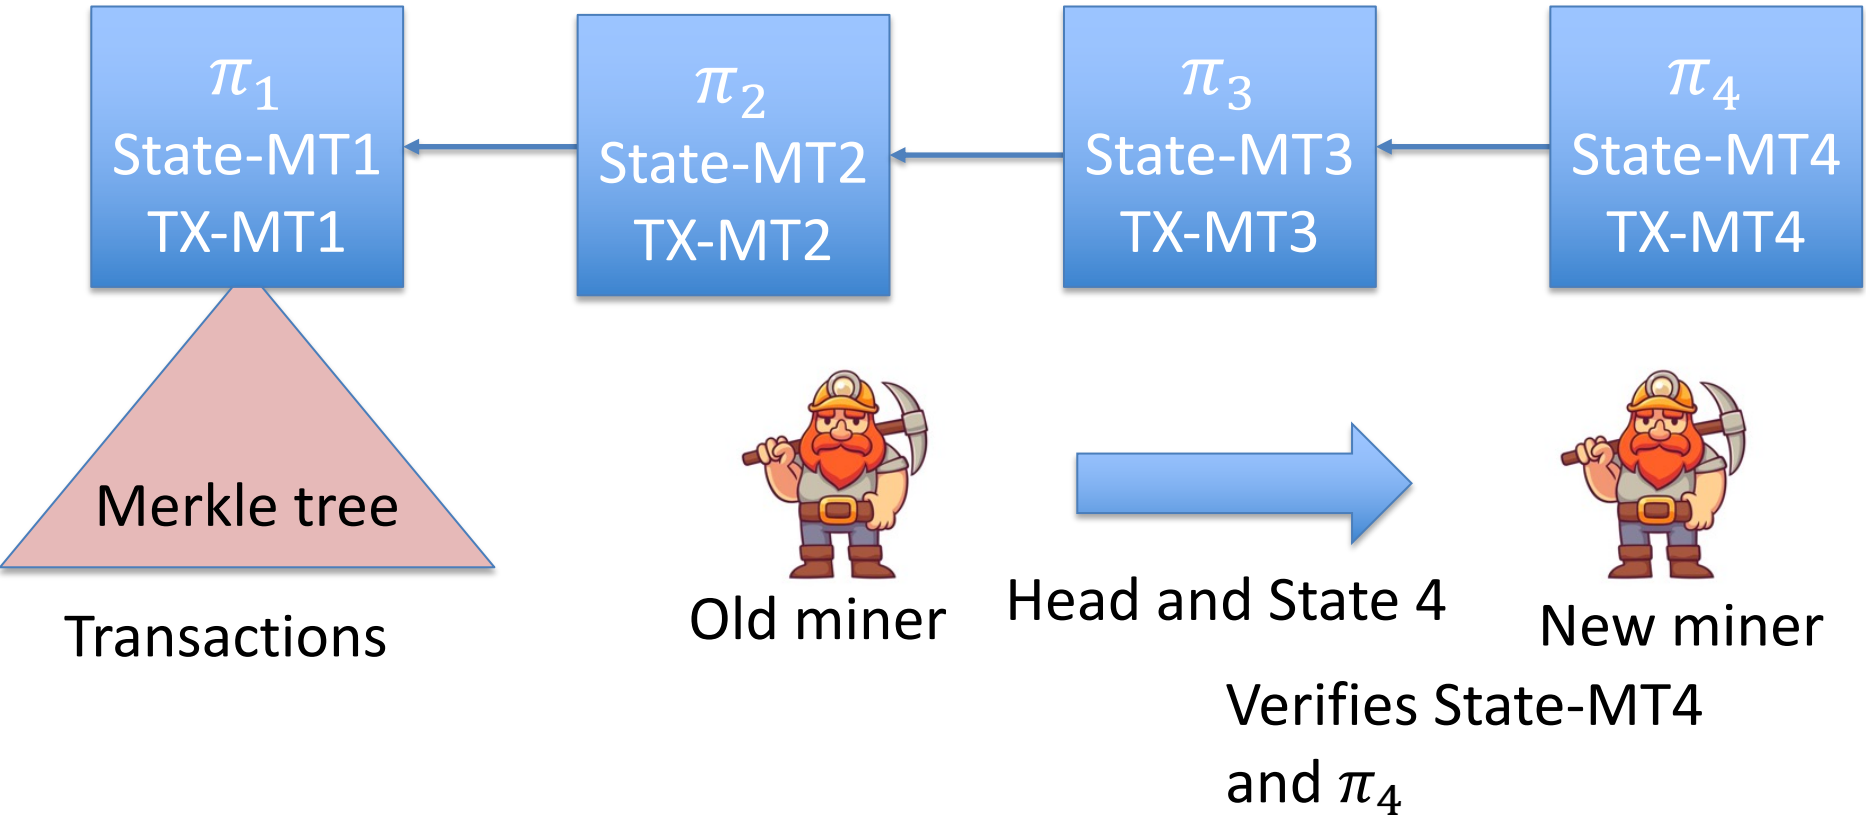
Constant size blockchains

- Rollup reduces the verification cost
- Still linear in the number of state updates
- When a node joins the network they need to verify one rollup proof per block!
- In general starting a full node requires verification of all blocks
 - Can take days!

Constant size Blockchain



Constant size Blockchain



Constant size Blockchain

- Light clients can verify every block!
 - Low memory, low computation
 - Independent of length of chain or #transactions
- Relies on data serving nodes for synching
- Practical today!

END OF LECTURE

Next lecture: Crypto tricks and open discussion
Please attend last two lectures if you can