

CS251 Fall 2021
(cs251.stanford.edu)



Using zk-SNARKs for Privacy on the Blockchain

Dan Boneh

Can we have private transactions on a public blockchain?

Naïve reasoning:

universal verifiability \Rightarrow transaction data must be public.
otherwise, how we can the public verify Tx ??

Goal for this lecture:

crypto magic \Rightarrow private Tx on a publicly verifiable blockchain

Crypto tools: **commitments** and **zero knowledge proofs**

The need for privacy in the financial system

Supply chain privacy:

A car company does not want to reveal how much it pays its supplier for tires, wipers, etc.



Payment privacy:

- A company that pays its employees in crypto needs to keep list of employees and their salaries private.
- Privacy for rent, donations, purchases

Business logic privacy: Can the code of a smart contract be private?

Last lecture





Neither Bitcoin nor Ethereum are private

etherscan.io:

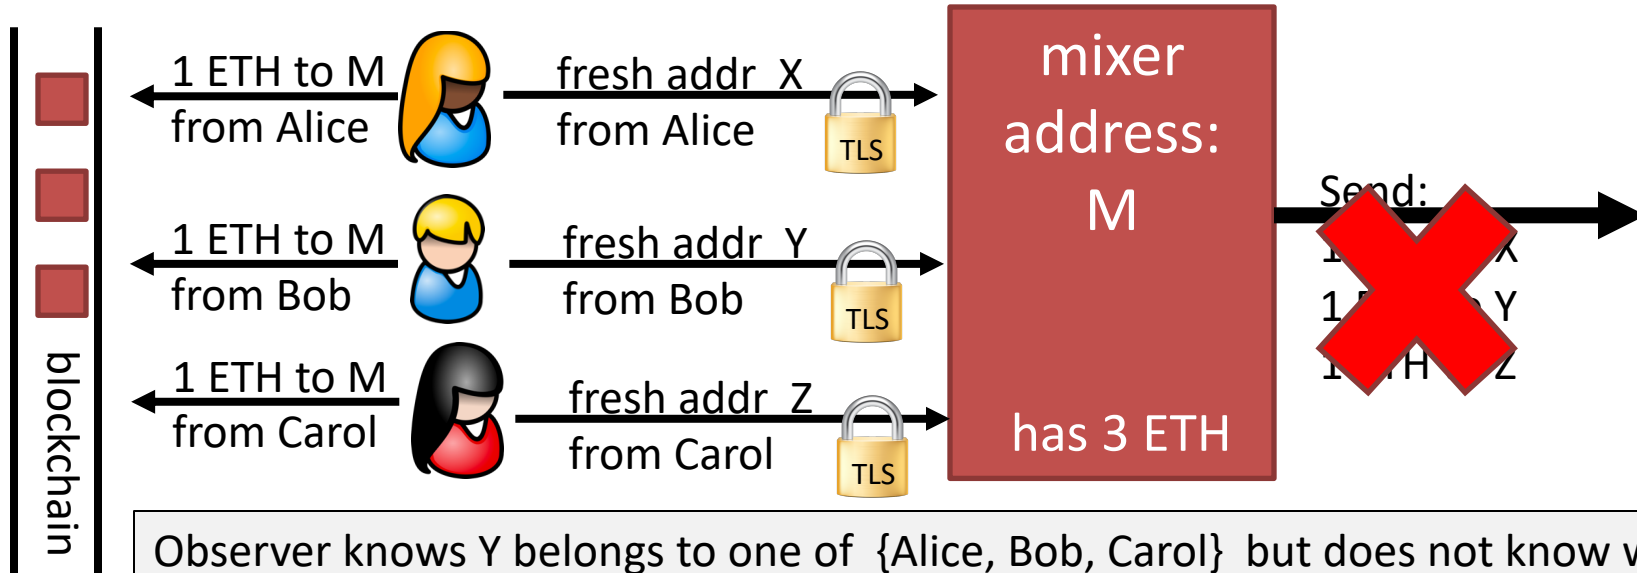
Address 0x1654b0c3f62902d7A86237...

Balance: 1.114479450024297906 Ether

Ether Value: \$4,286.34 (@ \$3,846.05/ETH)

	Txn Hash	Method ⓘ	Block
	0x0269eff8b4196558c07...	Set Approval For...	13426561
	0xa3dacb0e7c579a99cd...	Cancel Order_	13397993
	0x73785abcc7ccf030d6a...	Set Approval For...	13387834
	0x1463293c495069d61c...	Atomic Match_	13387703

Simple blockchain anonymity via mixing



Observer knows Y belongs to one of {Alice, Bob, Carol} but does not know which one
⇒ anonymity set of size 3.
⇒ Bob can mix again with different parties to increase anonymity set.

Problems: (i) mixer knows all, (ii) mixer can abscond with 3 ETH !!

Mixing without a mixer? on Bitcoin: **CoinJoin** (e.g., Wasabi), on Ethereum: **Tornado cash**

What is a zk-SNARK?

An central tool for
privacy on the blockchain

zk-SNARK: Blockchain Applications

Private Tx on a public blockchain:

- Confidential transactions
- Tornado cash, Zcash, IronFish
- Private Dapps: Aleo


Compliance:

- Proving solvency in zero-knowledge
- Zero-knowledge taxes

Scalability: privacy in zk-SNARK Rollup (next week)

(1) arithmetic circuits

- Fix a finite field $\mathbb{F} = \{0, \dots, p - 1\}$ for some prime $p > 2$.
- **Arithmetic circuit:** $C: \mathbb{F}^n \rightarrow \mathbb{F}$
 - directed acyclic graph (DAG) where internal nodes are labeled $+$, $-$, or \times
inputs are labeled $1, x_1, \dots, x_n$
 - defines an n -variate polynomial with an evaluation recipe
- $|C| = \# \text{ gates in } C$



Interesting arithmetic circuits

Examples:

- $C_{\text{hash}}(h, \mathbf{m})$: outputs 0 if $\text{SHA256}(\mathbf{m}) = h$, and $\neq 0$ otherwise
 $C_{\text{hash}}(h, \mathbf{m}) = (h - \text{SHA256}(\mathbf{m}))$, $|C_{\text{hash}}| \approx 20\text{K gates}$
- $C_{\text{sig}}(\text{pk}, m, \sigma)$: outputs 0 if σ is a valid ECDSA signature on m with respect to pk

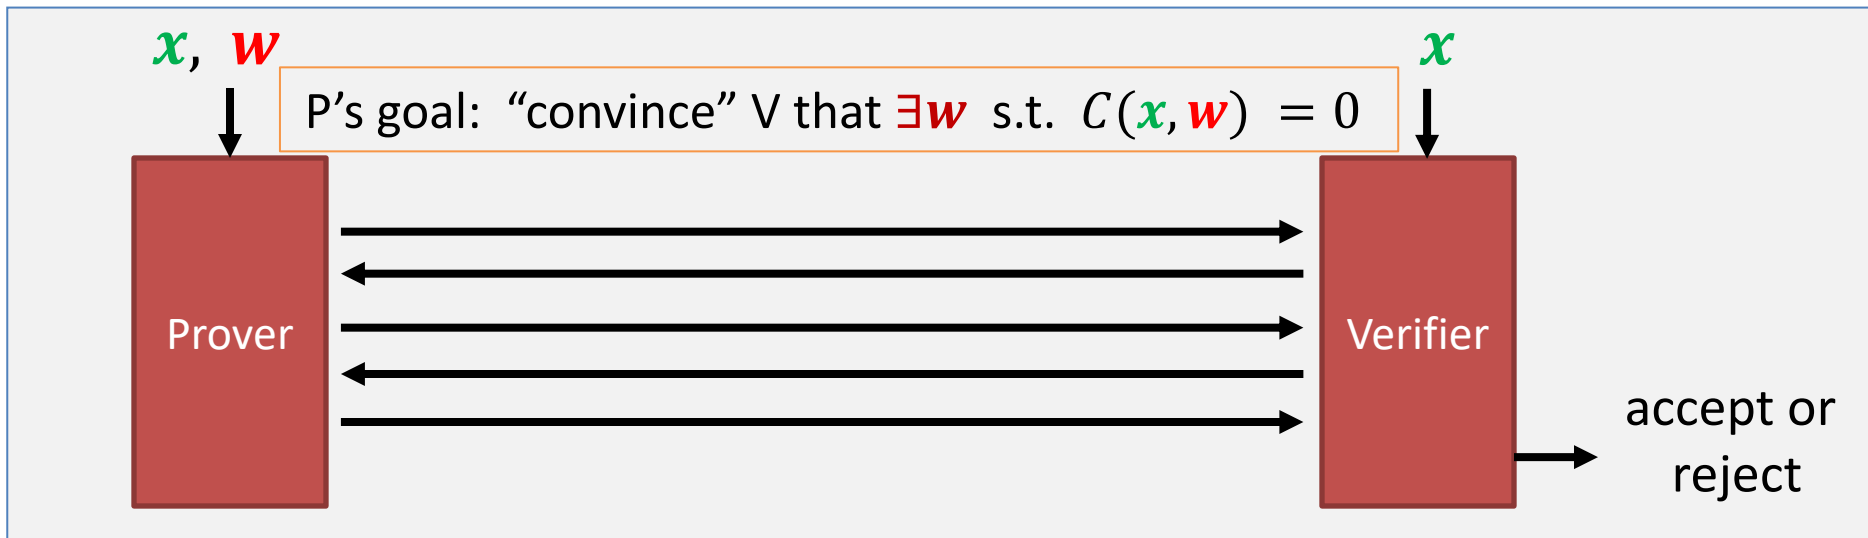
(2) Argument systems

(for NP)

Public arithmetic circuit: $C(x, w) \rightarrow \mathbb{F}$

public statement in \mathbb{F}^n

secret witness in \mathbb{F}^m



(non-interactive) Preprocessing argument systems

Public arithmetic circuit: $C(x, w) \rightarrow \mathbb{F}$

public statement in \mathbb{F}^n

secret witness in \mathbb{F}^m


Preprocessing (setup): $S(C) \rightarrow$ public parameters (S_p, S_v)

S_p, x, w



proof π

S_v, x



accept or reject

Preprocessing argument System


A preprocessing argument system is a triple (S, P, V) :

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

Argument system: requirements (informal)

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

Verifier $V(S_v, \mathbf{x}, \pi)$



Complete: $\forall x, w: C(\mathbf{x}, \mathbf{w}) = 0 \Rightarrow \Pr[V(S_v, x, P(S_p, \mathbf{x}, \mathbf{w})) = \text{accept}] = 1$

Knowledge sound: $V \text{ accepts} \Rightarrow P \text{ "knows" } \mathbf{w} \text{ s.t. } C(\mathbf{x}, \mathbf{w}) = 0$

P^* does not "know" $\mathbf{w} \Rightarrow \Pr[V(S_v, x, \pi) = \text{accept}] < \text{negligible}$

Optional: **Zero knowledge:** (S_v, \mathbf{x}, π) "reveals nothing" about \mathbf{w}

SNARK: a Succinct ARgument of Knowledge

A succinct preprocessing argument system is a triple (S, P, V) :

- $S(C) \rightarrow$ public parameters (S_p, S_v) for prover and verifier
- $P(S_p, \mathbf{x}, \mathbf{w}) \rightarrow$ short proof π ; $|\pi| = O(\log(|C|), \lambda)$
- $V(S_v, \mathbf{x}, \pi)$ fast to verify ; $\text{time}(V) = O(|x|, \log(|C|), \lambda)$

short “summary” of circuit

Why preprocess C ??

SNARK: a Succinct ARgument of Knowledge

A succinct preprocessing argument system is a triple (S, P, V) :

- $S(C) \rightarrow$ public parameters (S_p, S_v) for prover and verifier
- $P(S_p, \mathbf{x}, \mathbf{w}) \rightarrow$ short proof π ; $|\pi| = O(\log(|C|), \lambda)$
- $V(S_v, \mathbf{x}, \pi)$ fast to verify ; $\text{time}(V) = O(|x|, \log(|C|), \lambda)$

SNARK: (S, P, V) is **complete**, **knowledge sound**, and **succinct**

zk-SNARK: (S, P, V) is a SNARK and is **zero knowledge**

The trivial argument system

- (a) Prover sends w to verifier,
- (b) Verifier checks if $C(x, w) = 0$ and accepts if so.

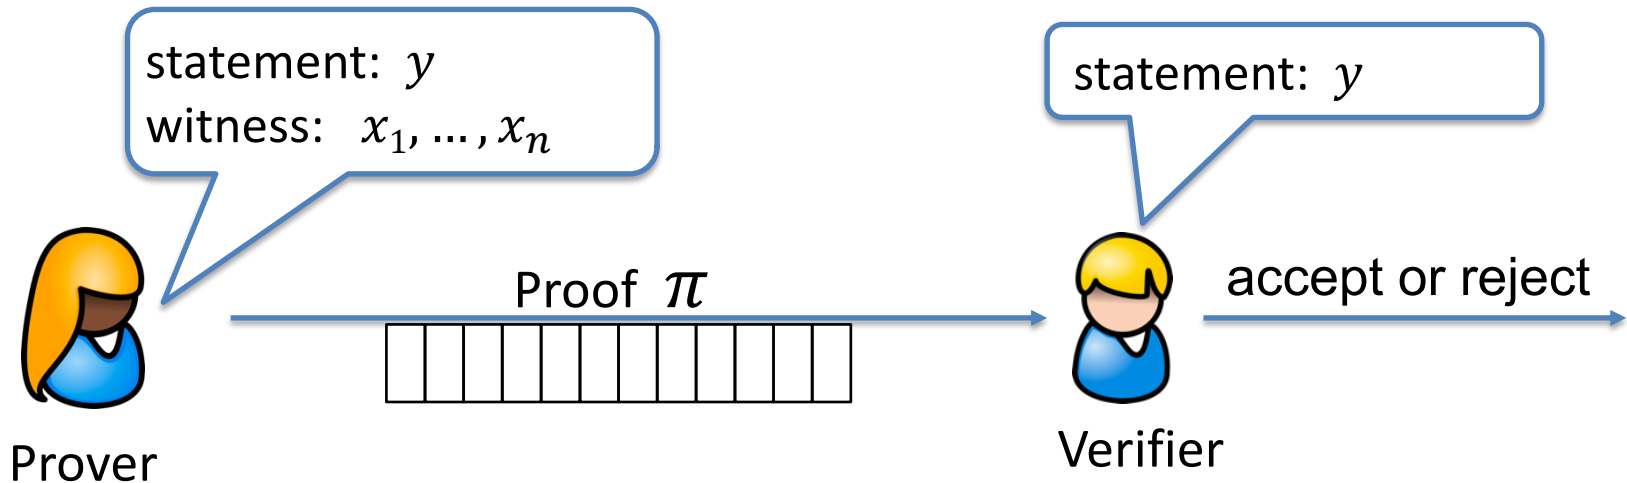
Problems with this:

- (1) w might be secret: prover does not want to reveal w to verifier
- (2) w might be long: we want a “short” proof
- (3) computing $C(x, w)$ may be hard: we want a “fast” verifier

An example

Prover: I know $(x_1, \dots, x_n) \in X$ such that $H(x_1, \dots, x_n) = y$

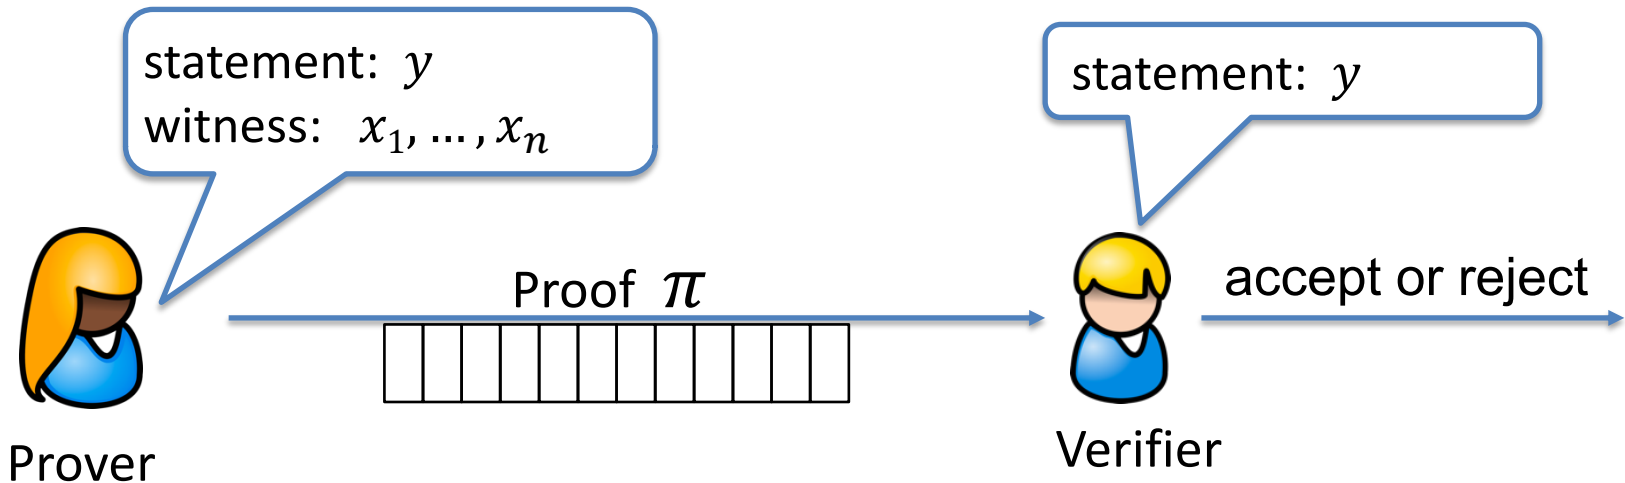
SNARK: $\text{size}(\pi)$ and $\text{VerifyTime}(\pi)$ is $O(\log n)$!!



An example

How is this possible ???

SNARK: $\text{size}(\pi)$ and $\text{VerifyTime}(\pi)$ is $O(\log n)$!!



Types of preprocessing Setup

Recall setup for circuit C : $S(C; r) \rightarrow$ public parameters (S_p, S_v)

random bits

Types of setup:

trusted setup per circuit: $S(C; r)$ random r must be kept secret from prover
prover learns $r \Rightarrow$ can prove false statements

trusted but universal (updatable) setup: secret r is independent of C

$S = (S_{init}, S_{index})$: $S_{init}(\lambda; r) \rightarrow pp$, $S_{index}(pp, C) \rightarrow (S_p, S_v)$
one-time no secret data from prover

transparent setup: $S(C)$ does not use secret data (no trusted setup)

better



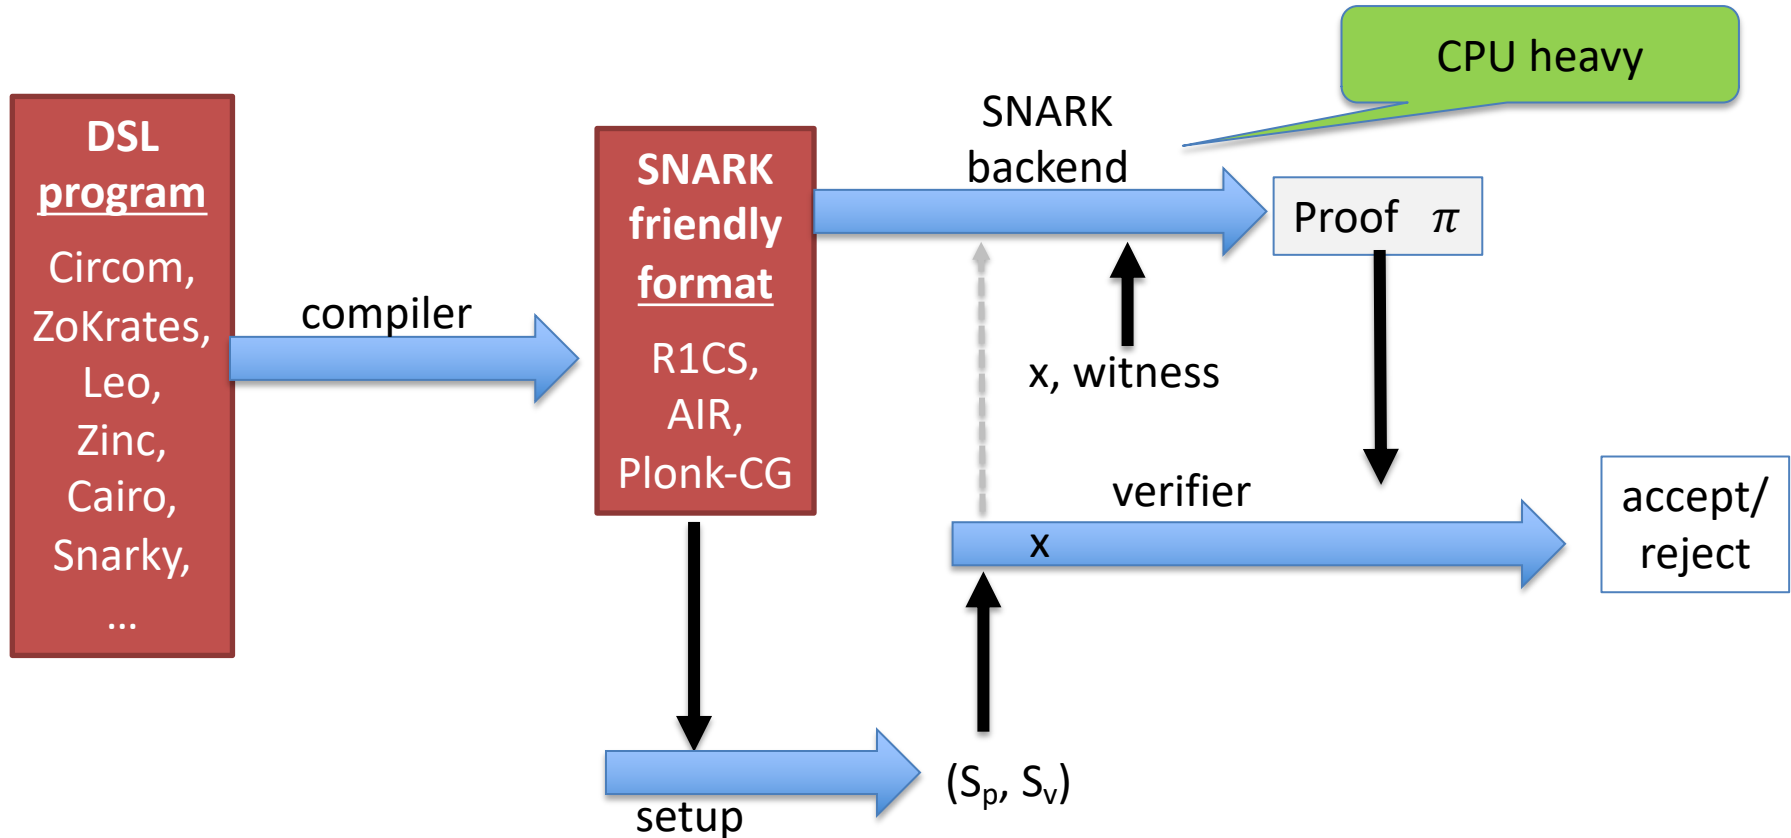
Significant progress in recent years

- **Kilian'92, Micali'94**: succinct transparent arguments from PCP
 - impractical prover time
- **GGPR'13, Groth'16, ...**: linear prover time, **constant size proof** $(O_\lambda(1))$
 - **trusted setup per circuit** (setup alg. uses secret randomness)
 - compromised setup \Rightarrow proofs of false statements
- **Sonic'19, Marlin'19, Plonk'19, ...** : universal trusted setup
- **DARK'19, Halo'19, STARK, ...** : no trusted setup (transparent)

Types of SNARKs (partial list)

	size of proof π	size of S_p (beyond C)	verifier time	trusted setup?
Groth'16	$O(1)$	$O(C)$	$O(1)$	yes/per circuit
Plonk/Marlin	$O(1)$	$O(C)$	$O(1)$	yes/universal
Bulletproofs	$O(\log C)$	$O(1)$	$O(C)$	no
STARK	$O(\log C)$	$O(1)$	$O(\log C)$	no
DARK	$O(\log C)$	$O(1)$	$O(\log C)$	no
⋮	⋮			⋮

A SNARK software system



How to define “knowledge soundness”
and “zero knowledge”?

Definitions: (1) knowledge sound

Goal: if V accepts then P “knows” w s.t. $C(x, w) = 0$

What does it mean to “know” w ??

informal def: P knows w , if w can be “extracted” from P



Definitions: (1) knowledge sound

Formally: (S, P, V) is **knowledge sound** for a circuit C if

for every poly. time adversary $A = (A_0, A_1)$ such that

$$S(C) \rightarrow (S_p, S_v), \quad (x, st) \leftarrow A_0(S_p), \quad \pi \leftarrow A_1(S_p, x, st):$$

$$\Pr[V(S_v, x, \pi) = \text{accept}] > 1/10^6 \quad (\text{non-negligible})$$

there is an efficient **extractor** E (that uses A_1 as a black box) s.t.

$$S(C) \rightarrow (S_p, S_v), \quad (x, st) \leftarrow A_0(S_p), \quad w \leftarrow E^{A_1(S_p, x, st)}(S_p, x):$$

$$\Pr[C(x, w) = 0] > 1/10^6 - \epsilon \quad (\text{for a negligible } \epsilon)$$

Definitions: (2) Zero knowledge

A story about the lady sipping tea



Definitions: (2) Zero knowledge (against an honest verifier)

(S, P, V) is **zero knowledge** if for every $x \in \mathbb{F}^n$
proof π “reveals nothing” about w , other than its existence

What does it mean to “reveal nothing” ??

Informal def: π “reveals nothing” about w if the verifier can
generate π **by itself** \implies it learned nothing new from π

(S, P, V) is **zero knowledge** if there is an efficient alg. **Sim**
s.t. $(S_p, S_v, \pi) \leftarrow \mathbf{Sim}(C, x)$ “look like” the real S_p, S_v and π .

Main point: **Sim** (C, x) simulates π without knowledge of w
(but also outputs S_p, S_v)

Definitions: (2) Zero knowledge (against an honest verifier)

Formally: (S, P, V) is (honest verifier) **zero knowledge** for a circuit C

if there is an efficient simulator ***Sim*** such that

for all $x \in \mathbb{F}^n$ s.t. $\exists w: C(x, w) = 0$ the distribution:

(S_p, S_v, x, π) : where $(S_p, S_v) \leftarrow S(C)$, $\pi \leftarrow P(S_p, x, w)$

is indistinguishable from the distribution:

(S_p, S_v, x, π) : where $(S_p, S_v, \pi) \leftarrow \mathbf{Sim}(C, x)$

How to build a zk-SNARK?

Recall: A zero knowledge preprocessing argument system.

Prover generates a short proof that is fast to verify

How to build a zk-SNARK ??

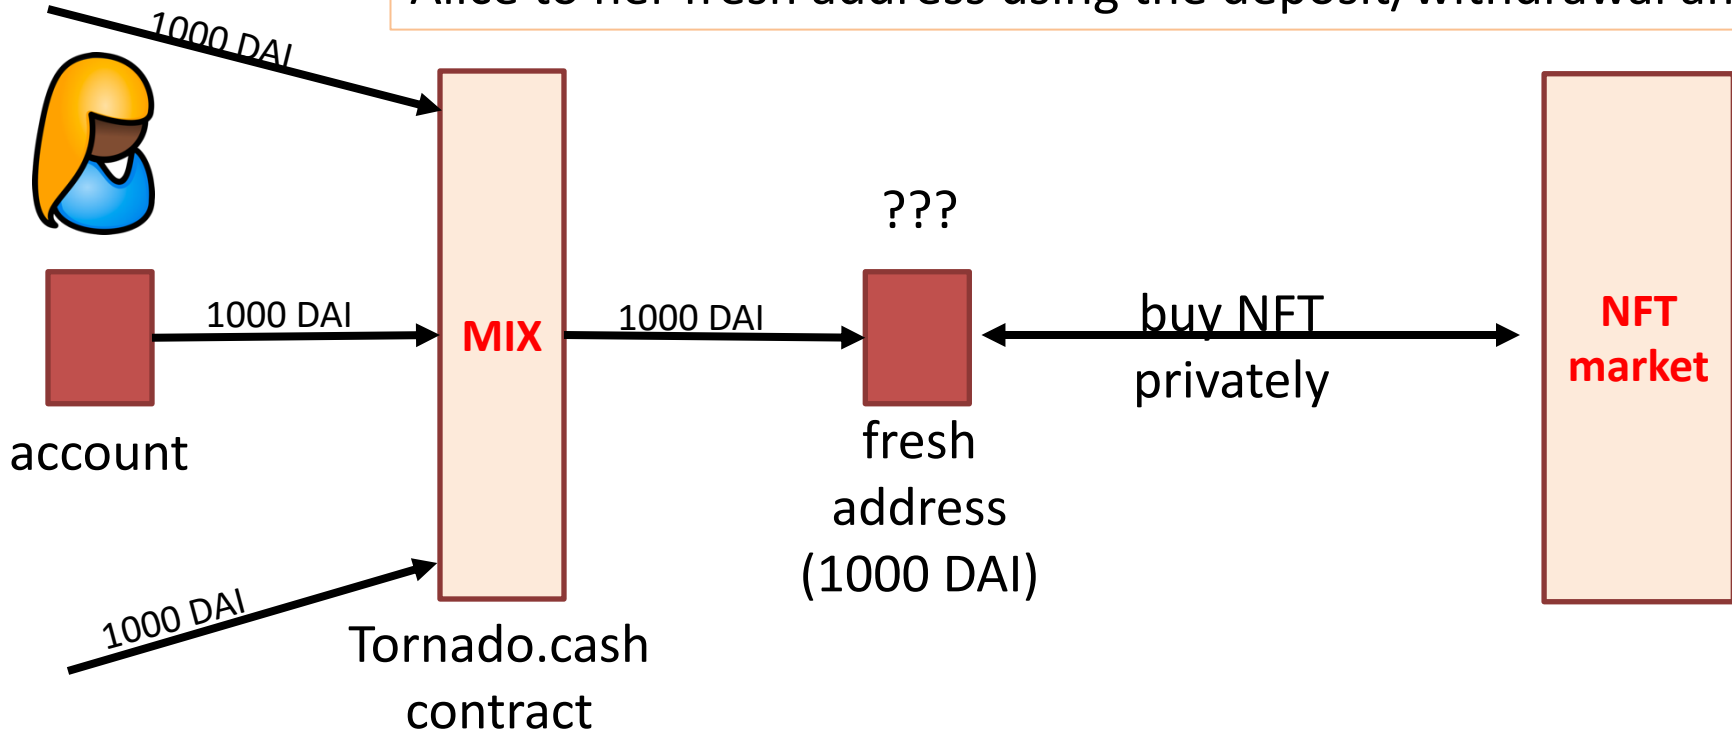
Next lecture

Tornado cash: a zk-based mixer

Launched on the Ethereum blockchain on May 2020 (v2)

Tornado Cash: a ZK-mixer

A common denomination (1000 DAI) is needed to prevent linking Alice to her fresh address using the deposit/withdrawal amount

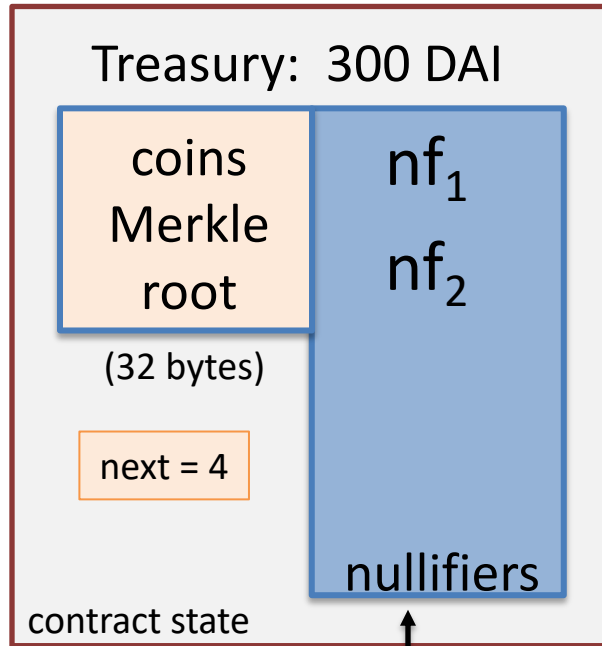


The tornado cash contract (simplified)

100 DAI pool:
each coin = 100 DAI


Currently:

- three coins in pool
- contract has 300 DAI
- two nullifiers stored



explicit list:
one entry per **spent coin**

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




Tornado cash: deposit

(simplified)

100 DAI pool:

each coin = 100 DAI

Alice deposits 100 DAI:



100 DAI


C_4 , MerkleProof(4)

Build Merkle proof for leaf #4:

MerkleProof(4) (leaf=0)

choose random k, r in R


set $C_4 = H_1(k, r)$



explicit list:

one entry per **spent coin**

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




Tornado cash: deposit

(simplified)



100 DAI

C_4 , MerkleProof(4)



Tornado contract

Tornado contract does:

- (1) verify MerkleProof(4) with respect to current stored root
- (2) use C_4 and MerkleProof(4) to compute updated Merkle root
- (3) update state

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

Coins
Merkle
root


tree of
height 20
(2^{20} leaves)

C_1 C_2 C_3 0 0 ... 0

public list of coins


Tornado cash: deposit

(simplified)



100 DAI

C_4 , MerkleProof(4)



Tornado contract

Tornado contract does:

- (1) verify MerkleProof(4) with respect to current stored root
- (2) use C_4 and MerkleProof(4) to compute updated Merkle root
- (3) update state

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

updated
Merkle
root

tree of
height 20
(2^{20} leaves)

C_1 C_2 C_3 C_4 0 ... 0


public list of coins

Tornado cash: deposit

(simplified)

100 DAI pool:
each coin = 100 DAI

Alice deposits 100 DAI:

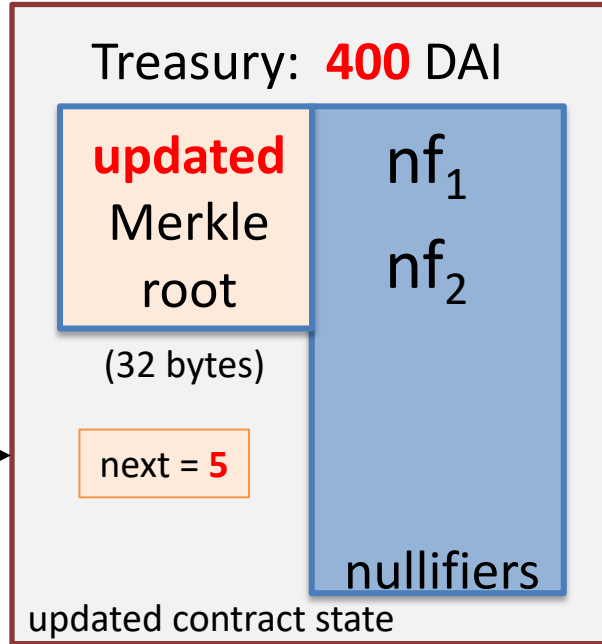


100 DAI


C_4 , MerkleProof(4)

note: (k, r)

Alice keeps secret
(one note per coin)



Every deposit: new Coin
added sequentially to tree



public list of coins

an observer sees who
owns which coins

Tornado cash: withdrawal

(simplified)


100 DAI pool:
each coin = 100 DAI

Withdraw coin #3
to addr A:




has note = (k', r')

set $nf = H_2(k')$



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



Bob proves “I have a note for some leaf in the coins tree, and its nullifier is **nf**”
(without revealing which coin)

Tornado cash: withdrawal

(simplified)

Withdraw coin #3 to addr A:




has note = (k', r') set **nf** = $H_2(k')$

Bob builds zk-SNARK proof π for
public statement $x = (\text{root}, \text{nf}, A)$
secret witness $w = (k', r', C_3, \text{MerkleProof}(C_3))$

where $\text{Circuit}(x, w) = 0$ iff:

- (i) $C_3 = (\text{leaf \#3 of root})$, i.e. $\text{MerkleProof}(C_3)$ is valid,
- (ii) $C_3 = H_1(k', r')$, and
- (iii) **nf** = $H_2(k')$.

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



(address A not used in Circuit)

Tornado cash: withdrawal

(simplified)

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

Withdrawal



The address A is part of the statement to ensure that a miner cannot change A to its own address and steal funds

Assumes the SNARK is **non-malleable**:

adversary cannot use proof π for x to build a proof π' for some “related” x' (e.g., where in x' the address A is replaced by some A')

C_1 C_2 C_3 C_4 0 ... 0

Bob builds zk-SNARK proof π for
public statement $x = (\text{root}, \text{nf}, A)$
secret witness $w = (k', r', C_3, \text{MerkleProof}(C_3))$

Tornado cash: withdrawal


(simplified)

100 DAI pool:


each coin = 100 DAI

Withdraw coin #3

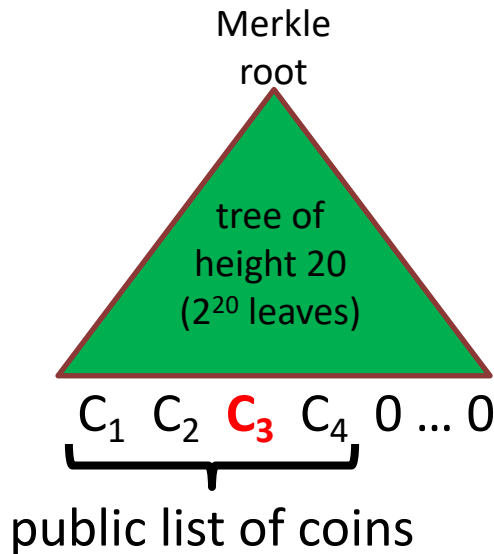
to addr A:



Bob's ID and coin C_3
are not revealed



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



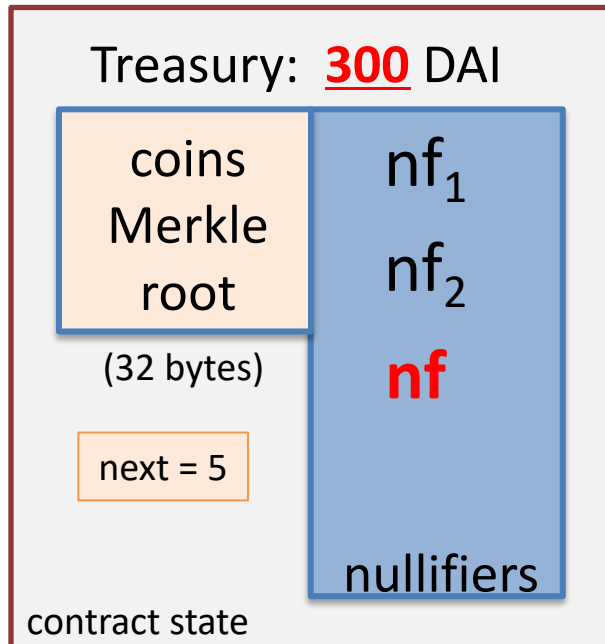
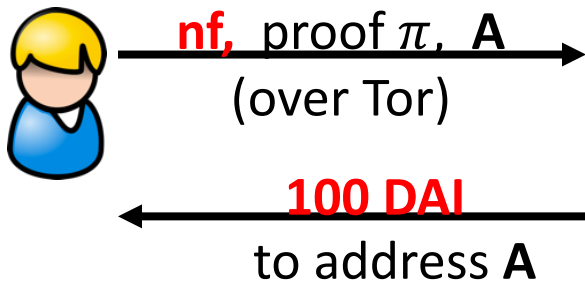
Contract checks (i) proof π is valid for (root, **nf**, **A**), and
(ii) **nf** is not in the list of nullifiers

Tornado cash: withdrawal


(simplified)

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.

Who pays the withdrawal gas fee?

Problem: how does Bob pay for gas for the withdrawal Tx?

- If paid from Bob's address, then fresh address is linkable to Bob

Tornado's solution: **Bob uses a relay**



Note: relay and Tornado also charge a fee

Tornado Cash: the UI

Deposit Withdraw

Token

DAI

Amount ⓘ

100 DAI 1K DAI 10K DAI 100K DAI

Deposit Withdraw

Note ⓘ

Please enter note here

Recipient Address [Donate](#)

Please enter address here

After deposit: get a note

Later, use note to withdraw

(wait before withdrawing)

Anonymity set

88,036
Total deposits

\$3,798,916,834
Total USD deposited

leaves occupied
over all pools

1 ETH pool

30141 equal user deposits

Latest deposits

30141. 4 minutes ago

30136. 3 hours ago

30140. 9 minutes ago

30135. 4 hours ago

30139. 2 hours ago

30134. 5 hours ago

30138. 3 hours ago

30133. 5 hours ago

30137. 3 hours ago

30132. 6 hours ago

Oct. 2021

Compliance tool

Tornado.cash compliance tool

Maintaining financial privacy is essential to preserving our freedoms.

However, it should not come at the cost of non-compliance. With Tornado.cash, you can always provide cryptographically verified proof of transactional history using the Ethereum address you used to deposit or withdraw funds. This might be necessary to show the origin of assets held in your withdrawal address.

To generate a compliance report, please enter your Tornado.Cash Note below.

Note

Please enter note here

Compliance tool

Note


Deposit 1 ETH Verified

Withdrawal 0.942 ETH Verified
Relayer fee 0.058 ETH

Date Transaction From Commitment

Date Transaction To Nullifier Hash

Generate PDF report



Reveals source address and destination address of funds

ZCASH / IRONFISH

Two L1 blockchains that extend Bitcoin.

Sapling (Zcash v2) launched in Aug. 2018.

Similar use of Nullifiers, support for any value Tx.

Quick review

A zk-SNARK for a circuit C :

- For a public statement x , prover outputs a proof that “convinces” verifier that prover knows w s.t. $C(x, w) = 0$.
- Proof is short and fast to verify

What is it good for?

- Private payments and private Dapp business logic (Aleo)
- Private compliance and L2 scalability with privacy

More to think about:

- private DAO participation? private governance?

Further topics

Privately communicating with the blockchain: Nym

- How to privately compensate proxies for relaying traffic

Next lecture: how to build a SNARK

END OF LECTURE

Two types of argument systems: interactive vs. non-interactive

Interactive: proof takes multiple $P \leftrightarrow V$ rounds of interaction

- Useful when there is a single verifier, e.g. a compliance auditor
- Example: zero-knowledge proof of taxes to tax authority

Non-interactive: prover sends a single message (proof) to verifier

- Used when many verifiers need to verify proof
- SNARK: short proof that is fast to verify