From Zero Knowledge to Private Transactions

Alessandro Chiesa

UC Berkeley

StarkWare

Zcash

Plan for today:

1. in peer-to-peer systems there is a conflict between privacy & integrity

2. zero knowledge proofs enable achieving <u>both</u> privacy & integrity

3. frontiers of ZKP on blockchains

The Conflict Between Privacy & Integrity in Peer-to-Peer Systems

A Motivating Example: Bitcoin

Bitcoin has three main parts.

(1) consensus

a peer-to-peer protocol that achieves an append-only ledger

$$\bullet tx_{11} \bullet tx_{12} \bullet tx_{13} \bullet tx_{14} \bullet tx_{15}$$

(2) incentives to participate

nodes that participate get *donations/fees* from processed transactions



Every payment transaction reveals: sender, receiver, amount.

This should raise some worries!

Payment History Reveals A Lot

medical information (specialty of your doctors)



insurance companies could use it to increase premium or even deny coverage

current and past locations (your travel patterns)



gold mine for stalkers, burglars, assassins, ...

merchant cash flow (suppliers, daily sales, ...)



intelligence for competitors

One of the reasons for why banks are highly regulated.

E.g., in the US the Gramm-Leach-Bliley Act states that: The GLB Act requires financial institutions ... to explain their information-sharing practices to their customers and to **safeguard sensitive data**. [www.ftc.gov]

Transactions In Bitcoin (simplified)

But in Bitcoin payment transactions do not include names!

Instead they include *cryptographic addresses*:





The Transaction Graph

Your addressess are known by everyone you interact with. And literally anyone can analyze the ledger:



transaction graph + side-info \rightarrow addresses become names of people!

Obfuscation

Use new address for each payment.

Launder money with others.



"Seems" harder to analyze.

But these "obfuscation" approaches leave tracks.

Bitcoin history is publicly stored forever. Methods of analysis only get stronger.

Transaction Graph Analytics

Numerous de-anonymization studies by researchers for Bitcoin and altcoins: [RM11] [BBSU12] [RS12] [RS13] [MPJLMVS13] [RS14] ... [MMLN17] [KFTS17] ...

Quantitative Analysis of the Full Bitcoin Transaction Graph

A Fistful of Bitcoins: Characterizing Payments Among Men with No Names

Dorit Ron and Adi Shamir

Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Israel {dorit.ron,adi.shamir}@weizmann.ac.il Sarah Meiklejohn Marjori Pomarole Grant Jordan Kirill Levchenko Damon McCoy[†] Geoffrey M. Voelker Stefan Savage

University of California, San Diego George Mason University[†]

Several companies provide analytics services:



Deanonymization: Good or Bad?

The FBI famously used "blockchain forensics" to support investigations that shut down the darknet market *Silk Road* (which sold illegal drugs, data, contraband, ...).



In *United States v. Coinbase, Inc*: **coinbase** was ordered by the government to disclose information about **nearly all its customers**, and it did so. The IRS has stated that it is using this information for tax fraud investigations.

At the same time **every user** is at risk of deanonymization and exploitation:

- Bitcoin transactions are like "twitter for your bank account" [Ian Miers]
- Bitcoin transaction data can be exploited by anyone (domestic or foreign)

Protecting user privacy is a must.

And achieing privacy via p2p systems is a socially desirable goal. \rightarrow



Fungibility

a dollar is a dollar, regardless of its history

Recognized as crucial property of money 350+ years ago. (*Crawfurd v. The Royal Bank*, 1749)

Bitcoin is **NOT** fungible because a coin's history is public.

In particular, a coin's value is ill-defined:

- different people may value the same coin differently
- the same person may value different coins differently

This leads to the heuristic "new coins are more valuable than old ones". But how should people agree on the correct value? A central party?

If privacy is so important then why isn't Bitcoin private?



How does the world know that Bob has 1 Bitcoin to spend?

Check that Bob received it, and that Bob did not spend it.

What if users encrypted their payment transactions?



Not clear how to check a payment's validity.

privacy and integrity are in conflict

Zero Knowledge Proofs Enable Achieving Privacy **and** Integrity

Today's Example: The Zerocash Protocol

A cryptographic protocol achieving a ledger-based currency that is

privacy-preserving

Anyone can publish a payment transaction to anyone else, while **provably hiding the payment's sender, receiver, amount**.

Bonus: the protocol is efficient

The privacy-preserving payment transactions:

- take few seconds to produce
- are less than 1KB in size
- take a few milliseconds to verify

Ideas introduced in this protocol are now endemic to privacy in blockchains.

The Basic Intuition



I am publishing three ciphertexts C1, C2, C3.

They contain the encryptions of a sender address, a receiver address, and a transfer amount respectively.

Moreover, the amount transfered has not been double spent.

I have generated a cryptographic proof π " that all of this is true.

Q1: what kind of cryptographic proof? Q2: what exactly is the statement being proved?

Beyond The Intuition



Q1: what kind of cryptographic proof?

zero knowledge succinct non-interactive proof of knowledge

(nothing revealed beyond truth of statement)

(proof is very short and cheap to verify)

(need to write it down!)

(true statements have proofs, false ones do not)

(technical... allows using crypto in statement)

zkSNARK

There has been tremendous theory+practice progress over the last 10 years for zkSNARKs. In particular, there are many popular open-source libraries for these tools: libsnark, bellman, arkworks, dalek, gnark, ethstark, openzkp, ...

A Little Beyond Intuition



Q2: what exactly is the statement being proved?

This is not trivial. Let's have some design fun.



Functionality: anyone can send a payment to anyone else.

Integrity: users cannot double spend their coins.

Privacy: all transactions "look the same".

The main tool is **ZKPs**: P(f,x,w) outputs *π* attesting that "*for public* f *and* x*, I know secret* w *s.t.* f(x,w)=true" and V(f,x,π) checks it.

Two additional basic tools:

- **commitments**: cm=COMM(m;r) *hides* m and *binds* to m
- pseudorandom functions:

y=PRF(m;sk) looks random to someone not knowing sk





Attempt #1: plain serial numbers



Transaction types

mint

sn

Consume 1 BTC to create a value-1 coin w/ serial number sn.

spend sn Consume

Consume the coin w/ serial number sn.



Good:

cannot double spend

Bad:

. . .

spend linkable to its mint anyone can spend!

Attempt #2: committed serial numbers



Transaction types

mint cm	Consume 1 BTC to create a value-1 coin w/ commitment cm.
spend	Consume the coin w/ serial number sn,
sn, r	by revealing the secret randomness r.



Good:

cannot double spend others can't spend my coins

Bad:

. . .

spend linkable to its mint

Attempt #3: ZKPoK of commitment



Transaction types

mint Consume 1 BTC to create a value-1 coin w/ commitment					
cm μ	Here is ZKP μ that for cm I know secret (sn,r) s.t. cm=COMM(sn;r).				
spend	Consume the coin w/ serial number sn.				
sn σ	Here is ZKP σ that for ${ m sn}$ I know secret randomness r s.t.				
exists \bullet cm \in "list of all prior commitments"					
well-formed	• $cm = COMM(sn;r)$				



Good:

cannot double spend

others can't spend my coins

spend and mint unlinkable

Bad:

. . .

fixed denomination

Attempt #4: variable denomination



Transaction types

mint	Consume v BTC to create a value-v coin w/ commitment cm.
cm,v µ	Here is ZKP μ that for (cm,v) I know secret (sn,r) s.t. cm=COMM(v,sn;r).
spend	Consume the value-v coin w/ serial number sn .
sn,v σ	Here is ZKP σ that for (sn,v) I know secret (cm,r) s.t.
existe	• $cm \in$ "list of all prior commitments"

well-formed • cm=COMM(v,sn;r)



Good:

cannot double spend others can't spend my coins spend and mint unlinkable variable denomination

Bad:

. . .

only hides sender

Attempt #5: payment addresses



Transaction types

mint	Consume v BTC to create a value-v coin w/ commitment cm.				
cm,v µ	Here is ZKP μ that for (cm,v) I know secret (sn,r) s.t. cm=COMM(v,sn;r).				
spend	Consume the value-v coin w/ serial number sn.				
sn,v o	Here is ZKP <i>o</i> that for (sn,v) I know secret (cm,r, <mark>p,pk,sk</mark>) s.t.				
exists well-formed mine	 cm ∈ "list of all prior commitments" cm=COMM(v,pk,p;r) sn=PRF(p;sk) & pk=PRF(0;sk) 				



Good:

cannot double spend others can't spend my coins spend and mint unlinkable variable denomination

Bad:

. . .

still only hides sender

Attempt #6: direct payments



Transaction types

COMM

value v

←r rand

pk public key

mint cm,v μ	Consume v BTC to create a value-v coin w/ commitment cm. Here is ZKP μ that for (cm,v) I know secret (sn,r) s.t. cm=COMM(v,sn;r).					
spend sn ^A ,cm ^B σ	Consume coin w/ serial number sn ^A & create coin w/ commitment cm ^B . Here is ZKP σ that for (sn ^A , cm ^B) I know secret (cm ^A , v ^A , r ^A , ρ^A , pk ^A , sk ^A) s.t.					
exists well-formed mine	 cm^A ∈ "list of all prior c cm^A=COMM(v^A,pk^A,ρ^A; sn^A=PRF(ρ^A;sk^A) & pk^A; 	ommitments" ;r ^A) =PRF(0;sk ^A)	(cm ^B ,v ^B ,r ^B ,ρ ^B ,pk ^B) send out-of-band			
well-formed same value	 cm^B=COMM(v^B,pk^B,p^B;r^B) v^A=v^B Good: 					
coin cm c	commitment serial	address	cannot c	double spend		

pk

PRF

┫

0

serial

number

sk

secret

key

sn

PRF

┫

 ρ seed

others can't spend my coins spend and mint unlinkable variable denomination hides sender, receiver, amt **Bad:** join and split coins?

Sketch of Final Design

pour pour view of mint mint mint sn₁ cm₃ sn₃ cm₆ cm_2, v_2 cm_5, V_5 blockchain cm_1, v_1 μ_2 μ_5 μ_1 sn₂ cm₄ sn₅ cm₇ Transaction types σ Consume v BTC to create a value-v coin w/ commitment cm. mint cm,v Here is ZKP μ that for (cm,v) I know secret (sn,r) s.t. cm=COMM(v,sn;r). μ Consume (my) **input** coins w/ serial numbers sn^A and sn^B in order to pour sn^A cm^C create two **output** coins (maybe not mine) w/ commitments cm^C and cm^D. Here is ZKP σ that I know secrets that demonstrate that sn^B cm^D σ • the input coins were minted at some point in the past, • the output coins are well-formed, • balance is preserved. Single tx type for:



- ✓ simple payments
- \checkmark join coins
- ✓ split coins
- ✓ making change
- \checkmark pay transaction fees

Research → Real World

2013.11: design of Zerocash protocol

2014.05: proof-of-concept implementation of Zerocash protocol

2015.02: commercial venture (Electric Coin Company)

2016.10: launch of **С**асн

The first large-scale deployment of ZKP technology.



Frontiers

Beyond Simple Payments

sn^A cm^C sn^B cm^D π I'm consuming my **unspent** coins in order to create new coins in a way that **value is preserved**. I'm not revealing the value, sender, or receiver.

& the receiver was a 501(c) organization but I am not revealing which one

& the value transfered is less than 10K

Exciting research direction:

Which policies are desirable (and feasible!) to balance privacy/fungibility and oversight/integrity?

Recent Work: Zero knowledge EXEcution (ZEXE)



Each transaction contains:

- serial numbers of **consumed** (input) **records**

- commitments of created (output) records

A zkSNARK attesting that the output records were created according to a **private offline computation**.

The transaction hides the record's contents, owners, and thei relation. (Data and function privacy.) Think "privacy-preserving Bitcoin scripts".



Thanks!

