Verifying Computations (with secrets) in the Blink of an Eye

matteo.campanelli@imdea.org

UCM, July 1st 2019

Matteo Campanelli **IMDEA** Software

"Just Stop Owning it" — The Rise of Cloud Computing

- Nowadays businesses usually outsource computation. • The bulk of it is rented, remote or both.
- Service Providers (SPs) can provide storage, computation, managed desktop, etc.

Advantages:

- Reduced infrastructure;
- Flexibility;
- Cost at scale.

Example:

A lone PhD student can analyze large genome data by renting 100 computers for 10 hours. And for less then \$200.

What Risks When Delegating?

• Complexity:

the computing capabilities I'm renting are complex systems and their execution may not always be correct.

• Incentives:

Service Providers do not necessarily have strong incentives to ensure correctness;

• **Opaqueness** (rented computers are "black boxes"): misconfigurations, HW problems, malicious operations, etc., these are all hard to detect.

How can we ever trust results computed by a third party then?

Some Solutions

Replicating

- Caveat: HW and SW configuration might be homogeneous
- •Auditing (checking the responses in a small sample) • Caveat: this doesn't help if the failures are too infrequent

Trusted HW

Caveat: assumes a chain of trust/additional assumptions

Proof-based verification:

What if the server returned its results along with a (probabilistic) proof that the results were computed correctly?

This Talk: Proofs and Computation

Part I: **Proving Correctness of Computation**

- The setting
- State of the art

Part II:

zkSNARKs and their applications

Beyond correctness – Proving Knowledge of Secrets

Part I: Proving Correctness of Computation

Proof-based Verification — Setting



(weak) Client



Verify(f, D, y, π) => accept/reject

Examples of functions *f*:

- video encoding • *ML training* image processing • data analysis



(powerful) Server

Proof-based Verification – Desiderata



(weak) Client



Desired Properties:

- inexpensive to verify
- little overhead for server

Is this possible?

Would (e.g.) an average phone be able to verify a herd of (untrusted) supercomputers?

• short proof (succinctness)



(powerful) Server

Inside Proof-Based Verification

Long list of works:

[ALMS98, GKR08, GGPR13, ..., Hyrax17, LegoSNARK19, ...]

Basic Approach:



What Can We Achieve Nowadays?

- How Efficiently?
 - Verification is super fast;
 - - matrix multiplication above: few minutes

Very short proofs for <u>any program (<1KB)</u> [Groth16]

• multiplication of two 1000x1000 matrices: few milliseconds

• Proving is meh (can give a 100-10000x overhead)

Main Challenges

- Main challenge: reducing resources for proving.
 - Example: Hashing a 13kB file (with SHA2) requires: >256GB RAM; few hours of work.

- Some recent works in this directions:
 - [WZCPS18]: Distributed generation of proofs (100x speedup)
 - [CFQ19]: LegoSNARK breaking proof generation in modular parts
 - in some cases 3 orders of magnitude speedup

Part II: Beyond Verifiable Computation – Proving Knowledge of Secrets

(Not So Much of) A Toy Example: Selling Secrets I know primes p,q such that: N = pqp,q Proof that it knows p,q, such that and N = pqp,q in **Seller/Prover** This "box" should a) **convince** Buyer (proof) b) not reveal anything on the secrets (zero-knowledge)

The start of a sketchy transaction:



Buyer/Verifier



Proofs with Secrets: Zero-Knowledge Proofs

So far: Verifiable Computation



"Please compute function **f** on data **D**"

"Here is result y = f(D)and a proof π "



Focus of this talk: zkSNARKs

(zk Succinct Non-Interactive Arguments of Knowledge)

Informally: leaks nothing except that the relation holds.

Now: Zero-Knowledge Proofs



"Given data **D**, show me you know secret **s** such that **R(D, s)** holds.

Relation



Example from before:

D was N and

s was (p,q) and ?

R was ~ "decryption and product"

Informally: we *could* efficiently "extract" — the secret from the prover.

zkSNARKs and Verifiable Computation

Techniques (and most of the efficiency) are the same:



NEXT: Applications.

Some General Applications of zkSNARKs

Authenticating Pictures





Q: Can we still authenticate pictures if we first need to (slightly) modify them?



Publishing



Authenticating Sensitive Pictures



Q: How to prove the edited pic is a transformation of an authentic pic? (without leaking the original pic)

A: just use a zkSNARK stating knowledge of "a signed pic to which blurring has been applied."

(What's public data **D**? What's secret **s**?)

Publishing





Anonymous Credentials



(What's public data **D**? What's secret **s**?)

Verifiable ML





Then how can the public check 👹

(from [WZCPS18])

Goal: we want this ML model to be publicly auditable.





Some **Blockchain-specific** Applications of zkSNARKs

Brush-up on Bitcoin

Bitcoin: a protocol to agree on exchange of tokens (e.g. which tokens are around? Who controls a token?)

Α

Key Property: If A gives a token to B:1. A can't use it again (at least consensus-wise).2. B can use it.



Little Privacy in Bitcoin



Randomized addresses are not enough!

We can trace back transactions to the actual senders/receivers. [GKRDN17,...]

(Attempts of) Pseudo-anonymity in Bitcoin: randomized addresses.

Improving Privacy in Bitcoin

Privacy Issues of Bitcoin:

- All transactions are public (and traceable)
- Side-information (e.g. cookies...)

Possible Solutions:

- Tumblers (aka mixers or mixnets)
- Private crypto-currencies (e.g. ZCash) **Our focus**



Private Crypto-Currencies

Warm up: Transactions in Bitcoin

Existing tokens

(Once spent, they go to the right and "produce" a new coin)

A (OXAAA)

B (oxZZZ)

Scenario: A wants to pay B

Spent tokens (you can't use them again)



Warm up: Transactions in Bitcoin

Existing tokens

(Once spent, they go to the right and "produce" a new coin)

A (OXAAA)

"New coin" produced from 0xred





Scenario: A wants to pay B

NB: This is all public info!



Public transaction reads:

***A** makes (s) spent; **B** receives it as (s) *

ZCash: Anonymous Transactions

Traceable transaction in Bitcoin:

"A makes (\$) spent; B receives it as (\$) "

Roadmap (two steps):

- - valid transaction.

Our Goal:

- **Remove** receiver/sender as public;
- Keep validity of transaction.



1. Anonymize sets of spent/existing tokens; 2. use zkSNARKs to prove we are performing a

Before: completely public.

Existing tokens

Spent tokens

A (OXAAA) **Oxred**



Anonymized Spent Tokens

Now: somewhat private through hashing.



we publish the output of two hash functions instead of actual coins.



Transactions in ZCash



What A does to spend her coin to B:

- makes new Oxblue (\$) and send it (privately) to B
- broadcast (*newtok*, *oldtok*) where ullet

newtok = H_exist(0xblue)

oldtok = H_spent(0xred)

broadcast a zkSNARK proof π to validate transaction





zkSNARKs in ZCash: Summary

- Bitcoin is not anonymous
- We can make it anonymous by:
 - extending it with anonymized sets ●
 - receiver

using SNARKs to prove transactions are valid without leaking sender/

Fair Exchange of Digital Goods with Bitcoin

Recall the sketchy exchange from before:





Buyer/Verifier

Q: What's going to happen next? And how to ensure **fairness** with no trusted party?

Proof that it knows p,q, such that p,q in \bigcap and N = pq



Seller/Prover





The Seller



The Buyer





The Seller





The Seller



The Buyer





The Seller





The Buyer







The Seller





 5
 1
 5
 3
 7
 2

 8
 7
 4
 1
 9
 6

The Buyer







The Seller





The Seller





The Seller





The Seller





The Seller





The Seller





The Seller





The Seller





Wrapping up

- With verifiable computation we do not need trust a remote computer.
- We can give **proofs** of correctness that are **short** and **quickly verifiable**.
 - Proving often requires a large overhead.
- Beyond Verifiable Computation: we can prove knowledge of secrets (zkSNARKs).
- zkSNARKs have applications both to society at large and to blockchains.

Thanks! And questions?

For more info:

- visit <u>https://github.com/scipr-lab/libsnark</u>
- or drop me a line at matteo.campanelli@imdea.org