



# Blockchains and large-scale distributed systems

Petr Kuznetsov Télécom Paris IP Paris



Journées Partenaires Entreprises Télécom Paris – 20-21 juin 2019



# Chronology

**1982** Byzantine Generals **1990** Paxos 1992 "ProofOfWork" 1999 PBFT 1995 Hashcash 2002 Sybil attack 2009 Bitcoin

[Narayanan, CACM, Dec 2017]

. . .

Linked Public Byzantine Timestamping. Fault Keys as Proof Digital Smart Verifiable Logs Identities Tolerance Cash of work Contracts 1980 Merkle Chaum Ecash<sup>10</sup> Tree<sup>33</sup> Byzantine Anonymous Generals<sup>25</sup> Communication Chaum Security w/o 1985 Identification<sup>11</sup> Offline Paxos<sup>28</sup> Haber & Ecash<sup>32</sup> stornetta22 1990 Digicash Benaloh & de mare<sup>6</sup> Anti-spam<sup>15</sup> Bayer, haber, Szabo stornetta<sup>5</sup> Essav<sup>41</sup> 1995 Micro-Mint<sup>40</sup> Haber & B-money<sup>13</sup> Hashcash<sup>2</sup> stornetta23 Client Pbft<sup>8</sup> Goldberg Puzzles25 2000 Dissertation Paxos made Simple<sup>29</sup> Sybil attack<sup>14</sup> Bit Gold<sup>42</sup> -2005 Computational Bitcoin<sup>34</sup> Impostors1 • Private 2010 Blockchains . Ethereum v 2015 Nakamoto Consensus

# **Distributed ledger**

Shared data structure: linear record of (blocks of) transactions

- Append-only
- Backtrack verifiable
- Consistent: total order



Open environment:

- No static membership
- No identities (public keys)





# Sybil-resistant consistency?

- Sybil attack: the adversary can own an arbitrarily large fraction of participants
   ✓ Why don't good guys do the same? ☺
- Classical (BFT) protocols don't work
   ✓ Bounds on faulty fraction (e.g., <1/3)</li>
- Bitcoin:
  - ✓ Assume a synchronous system
  - $\checkmark\,$  Message delays are bounded by  $\delta\,$
  - ✓ Need to "slow down" updates (wrt  $\delta$ )
  - ✓ Solve a puzzle before updating (PoW)

# (Bitcoin) blockchain

- Clients broadcast an update
- Dedicated clients (miners) collect updates solve puzzles, update and broadcast their local ledgers
- Clients always choose the longest (verifiable) ledger
- Old enough blocks are considered consistent



#### When it works



- Expected time to solve the puzzle >>  $\delta$
- The adversary does not possess most of computing power

The probability of a fork drops exponentially with the staleness of blocks

# When it does not work

- Asynchronous/eventually synchronous communication
- An adversary controls half of computing resources
- Even a small probability of error cannot be tolerated
- Energy consumption and throughput is an issue

Bitcoin Consumes More Electricity Than Iceland



The work of bitcoin miners all over the world are contributing to a massive rise in electricity consumption. Recent data reveals that current levels of consumption surpass those of the country of iceland.<

#### When it is not needed



- No Sybil attacks
  - ✓ Participation under control
- No need for total order
  - ✓ Some form of causality is enough?

#### **Bitcoin: A Peer-to-Peer Electronic Cash System**

Satoshi Nakamoto satoshin@gmx.com www.bitcoin.org

**Abstract.** A purely peer-to-peer version of electronic cash would allow online payments to be sent directly from one party to another without going through a financial institution. Digital signatures provide part of the solution, but the main benefits are lost if a trusted third party is still required to prevent double-spending. We propose a solution to the double-spending problem using a peer-to-peer network.

...

The only way to confirm the absence of a transaction is to be aware of all transactions. In the mint based model, the mint was aware of all transactions and decided which arrived first. To accomplish this without a trusted party, transactions must be publicly announced [1], and we need a system for participants to agree on a single history of the order in which they were received.



#### Cryptocurrency without consensus

[Guerraoui et al., PODC'19]

Consensus number of the asset transfer data type:

 $\checkmark$ k-owned (smart contracts with k parties) – k

- Asynchronous asset transfer algorithm
   ✓1-owned: secure broadcast
   ✓k-owned: k-consensus + secure broadcast
- No need of total order on transactions

#### Commutativity and causality

- T0: \$100 from Alice to Carole
- T1: \$100 from Bob to Alice
- T2: \$100 from Drake to Alice

T0 causally depends on T1 (not enough funds otherwise) T1 and T2 commute (T0 succeeds regardless of the order)



#### What about double-pending?

- T0: \$100 from Bob to Alice
- T1: \$100 from Alice to Carole
- T2: \$100 from Alice to Drake

Alice's initial balance is 0, but it claims to both beneficiaries to have received money from Bob



#### Asset transfer implementation

Message-passing, Byzantine failures

- Each transfer is equipped with its causal past (a set of incoming transactions)
- Make sure that a faulty account holder cannot lie about its causal past
- Secure broadcast [Bracha, 1987, Malkhi-Reiter, 1997]
   ✓ Source-order: messages by the same source are delivered in the same order

#### Modular approach: private and public



#### Cryptocurrency without consensus

- Asset transfers do not always require total order
   ✓ Source order is sufficient for consistency
   ✓ (Asynchronous) secure broadcast
- Can be generalized to (limited-scope) "smart contracts"
  - ✓ only account owners need consensus, but still no global total order
- Coming: probabilistic and Sybil-tolerant secure broadcast can be implemented (coming)

✓ Permissionless asset transfer

## Other algorithmic challenges

- Maintaining the system evolution

   Decentralized updates [Tezos]
   Local views, federated quorums [Stellar]
   Asynchronous reconfigurations [Dynastore]
- Concurrency in smart contracts
  - Sequential programs run in a concurrent environment
- Blockchain ecosystems
  - ✓Cross-chain transactions

✓ Fair exchange/atomic commitment

### Take-aways

- Blockchains do solve a new problem
  - ✓ Maintaining a total order in an open system
  - ✓ With a brute-force approach
  - $\checkmark$  Scalability is a challenge
- The primary application of blockchains do not need blockchains
  - ✓A weaker abstraction may suffice
- Do not go for a technology
   ✓Go for a problem



# Thank you!