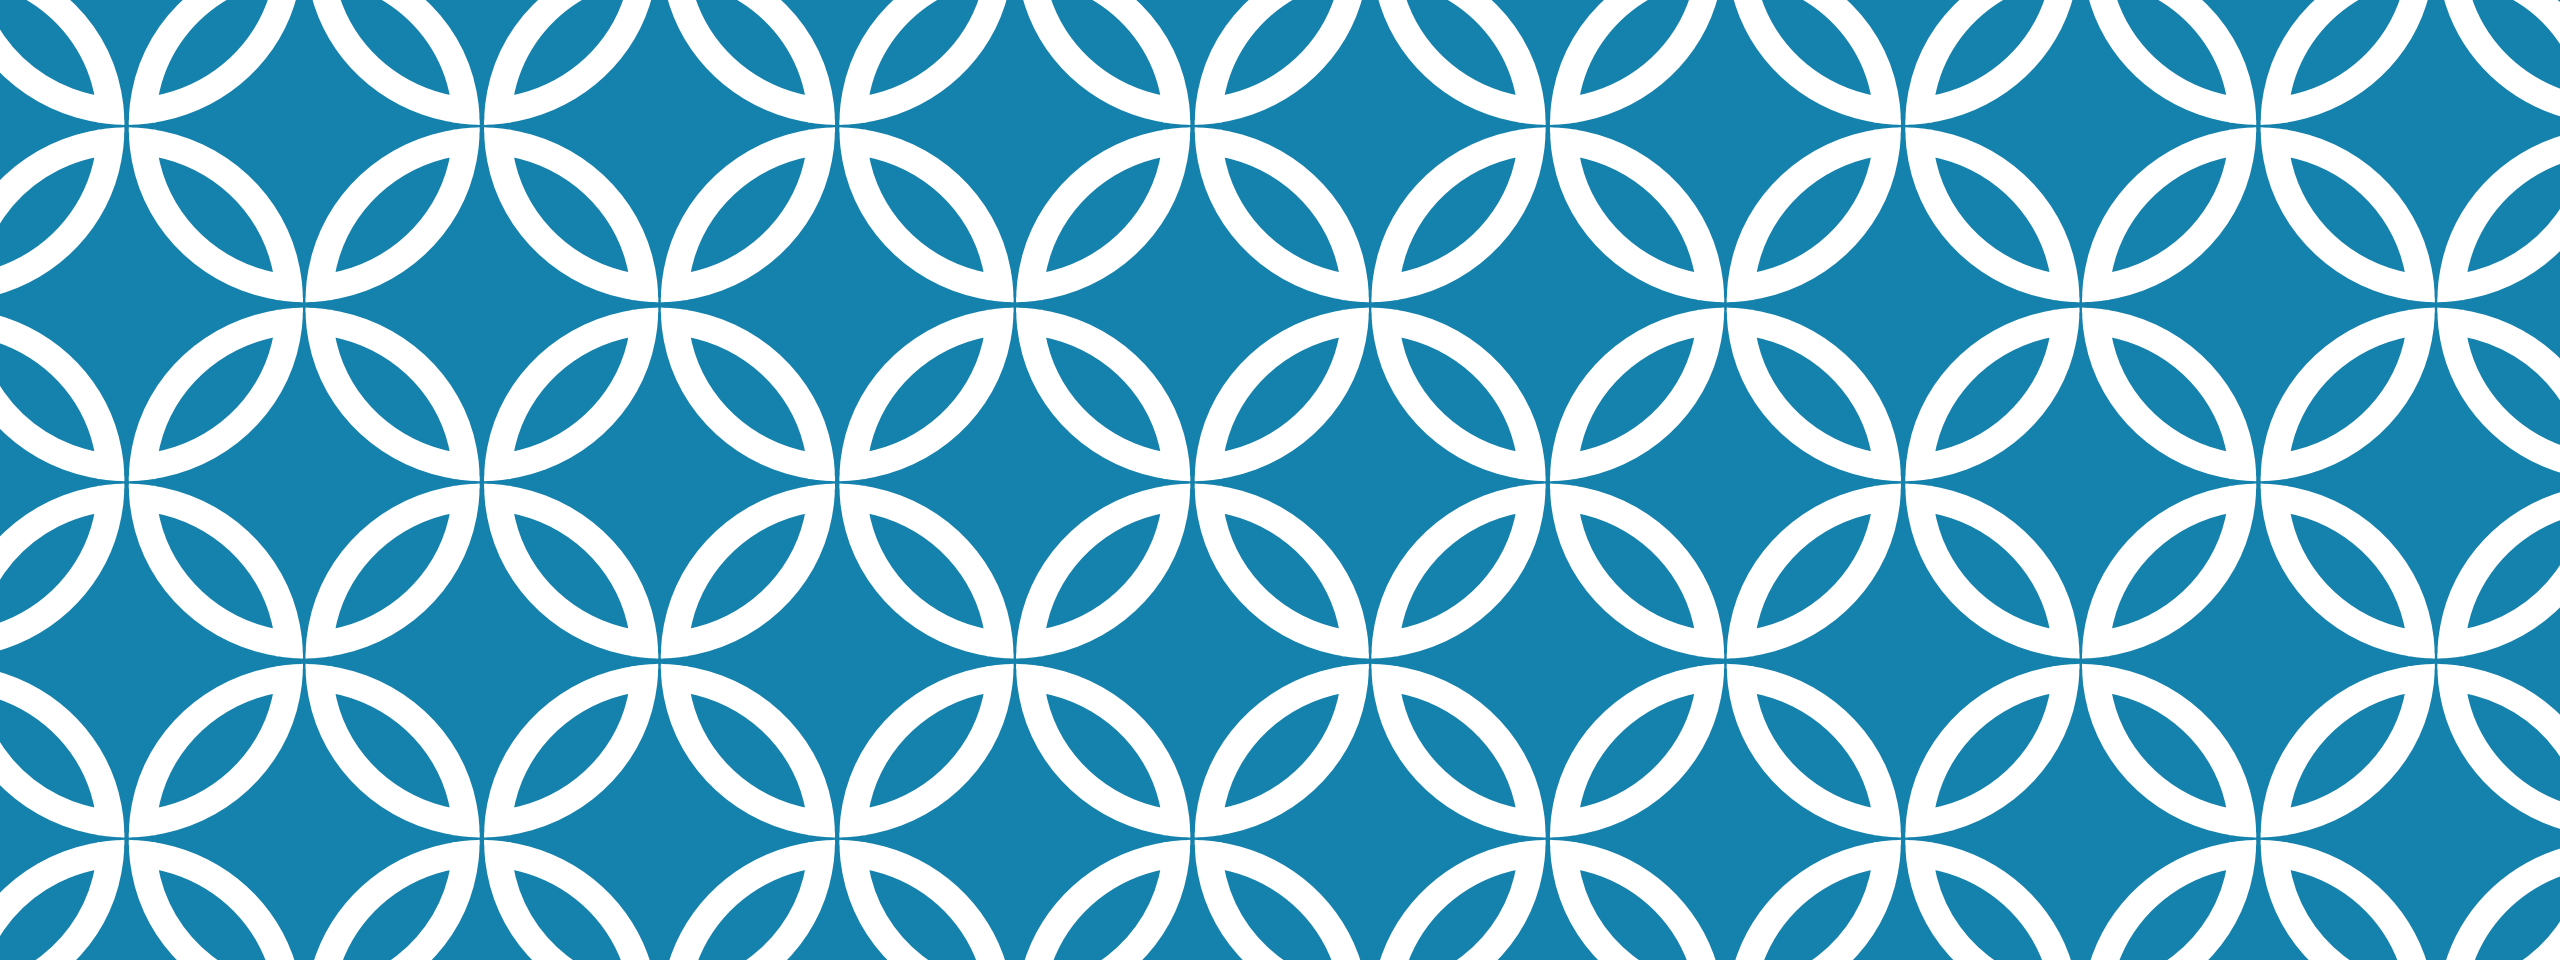

CS5412/LECTURE 15 BLOCKCHAINS FOR IOT (PART 2)

Ken Birman

CS5412 Fall 2022

SUMMARY OF BLOCKCHAIN CONCERNS

- Permissioned [case 1: data center; case 2: geoscale] or Permissionless? Energy cost of permissionless block mining a major issue, but permissioned means >trust, >control by authorities (taxes!), >traceability, <anonymity.
- If the BlockChain gets really large, costs of downloading a copy and verifying become significant.
- National-scale "disruption" scenarios can cause massive rollbacks, chaos.
- Accidental loss of some chunk of the chain, making verification impossible.
- Smart contracts might be too smart for their own good.

>: short for "implies/requires more", <: short for "reduces/barrier to"

MORE CONCERNS

https://www.joe.ie/news/pics-this-pile-of-cash-worth-22bn-was-found-inside-the-insane-home-of-a-mexican-drug-lord-409313

Today's most enthusiastic Blockchain use cases seem to center on a mix of illegal transactions, money laundering, and a gigantic technology boom but without much of a "market" for the associated products.

The model also depends on some hardness assumptions: finding a nonce, factoring RSA key. Quantum computers could shake up these assumptions.

Unresolved privacy concerns: "everything is on the table."

WALKING THROUGH THE ISSUES

Which issues would arise on a “smart farm”?

Would a BlockChain solve those issues? What new risks would it introduce?

What limits the speed of new technology adoption?

BLOCKCHAIN ON A FARM

Main uses seem to be for audit trails of various kinds:

- Capture data about something we are supposed to trace or record.
- Write it digitally into the ledger, securely. Tamperproof and automatic
- Auditors given access to the record.

But they will want to know:

- Why should I trust this BlockChain record? Is the underlying data valid?

BLOCKCHAIN FOR SENSORS

A smart farm would store BlockChain records that include sensor data.

Very likely the sensors themselves would be securely connected to their host machines, for example using Azure's IoT Hub or the AWS equivalent.

Reminder: With an IoT Hub model, the sensor itself uses security keys, and will only connect and talk to the hub. This enables a *digital twin* model.

TRUST WITH SENSORS

… so we can assume the connection to the sensor is secure. But how would a digital twin for a farm work?

Azure IoT Hub will only allow authorized sensors to be part of the system, and it patches the software and configuration automatically. Feeds events to the Azure IoT function server, where functions consume them.

So presumably these functions log the event records to the Blockchain.

QUESTIONS THIS MIGHT NOT ANSWER

Is this sensor the correct one for the information the application claims to have captured?

Is the sensor working correctly?

Has the data the sensor generated been modified before it was logged?

CAN WE ANSWER THE QUESTIONS AN AUDITOR MIGHT ASK?

A sensor records shows that cow 2143 was milked on Tuesday at 10am. Later the milk turned out to have a dangerous bacteria in it, like Listeria. It got through and a consumer became quite sick.

Was she properly clean when she was milked? Had she been evaluated for mastitis as required by the health department? Was she periodically checked for overall health? Did she receive any "off records" meds?

CAN WE ANSWER THE QUESTIONS AN AUDITOR MIGHT ASK?

Realistically, an audit using sensor data would provide "evidence" but can't answer these questions. Non-technical people might find this surprising, but in cloud computing we have seen why many either *cannot* be answered, or *we can't have confidence our answers will be correct.*

BlockChain does protect against tampering, and does record what the sensors reported, and when. This is already valuable, as long as we are honest about the capabilities and limitations.

TO HAVE REAL CONFIDENCE…

Azure IoT Hub would need to log **management** events too.

The audit-trail examination tool would need to visualize this information and be able to convince the auditor that yes, this is the proper sensor, it seems to be properly calibrated, the data wasn't tampered with…

The mention of time suggests that we might also need to log events related to the way the system tracks time, or at least have a "story" there.

BLOCKCHAIN-SPECIFIC FORM THIS TAKES?

We noted that early adopters are people with transactions to carry out anonymously, or maybe with money to launder. Strongly motivated mostly because for now, Blockchain feels like a way to evade oversight and taxes.

Business community has many people keen to adopt the next new thing. Startup frenzy and huge fortunes made on ICOs adds fuel to the flames.

But farming is a case for mainstream use and these questions need to be answered. This is why the mainstream technology community is more cautious.

LET'S LOOK AT A BLOCKCHAIN CREATED SPECIFICALLY FOR IOT

Cornell "smart farms" research effort (CIDA) is highly visible.

Led by Susan McCouch, Hakim Weatherspoon, Steve Wolf and Abe Strouck.

One early accomplishment: Vegvisir, a BlockChain specifically for agriculture.

SOME ISSUES THEY THOUGHT ABOUT

A lot of the "events" that matter in an agriculture or farming setting are in remote places, disconnected from the main system.

The BlockChain would probably be used primarily as an audit trace, to track events in the food chain from farm to table.

So this raises issues like intermittent connectivity, how we know that the sensor that generated a record is the "correct one" for that role, etc.

CONNECTIVITY: JUST ONE ISSUE OF MANY!

Vegvisir is a research project and a proof of concept, but not deeply integrated with Azure IoT Edge.

Any real product will need more ties to the Azure infrastructure.

But an Azure Blockchain would also benefit: as a part of the official Azure ecosystem, we might gain better answers to some of the trust issues!

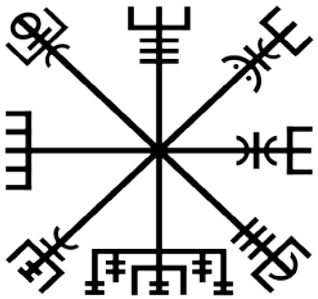

VEGVISIR

Slides from Robbert van Renesse

Talk presented at ICDCS 2018

A BLOCKCHAIN FOR THE FOOD SUPPLY CHAIN

Robbert van Renesse

joint work with Hakim Weatherspoon, Danny Adams,
Kolbeinn Karlsson, and Stephen B. Wicker

Initiative for Crypto-Currencies and Contracts (IC3)
Cornell Digital Agriculture Initiative

BLOCKCHAIN'S PROMISE

Promises

- Global currency
- Smart contracts
- Notarization
- *Accountability*
- …

BlockChain background: Repeats things we've discussed.

A REPLICATED LEDGER OF TRANSACTIONS

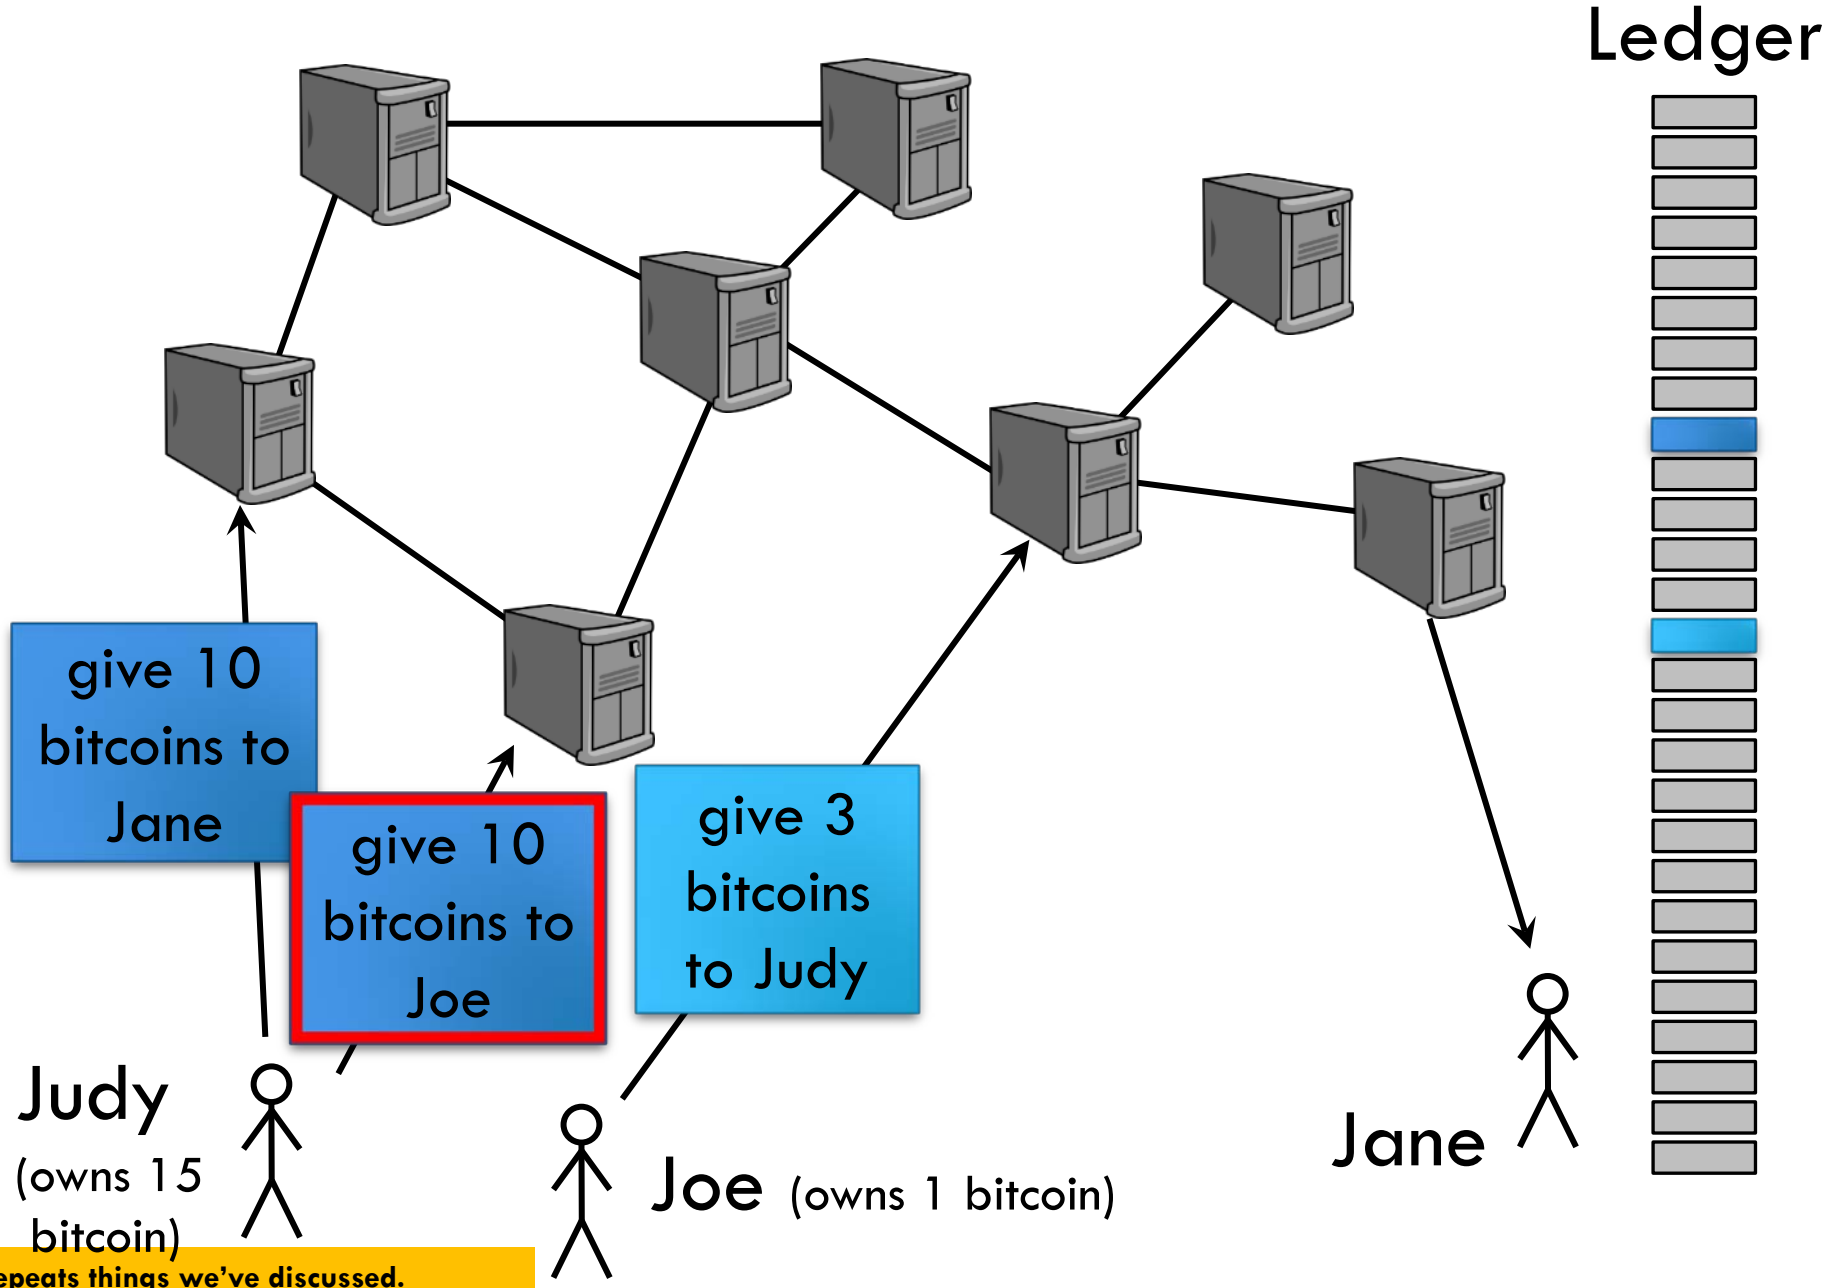

Ittay Eyal, May '16

BlockChain background: Repeats things we've discussed.

SMART CONTRACTS

Smart contracts are executable programs on the BlockChain, take input from the BlockChain, and produce output on the BlockChain

Main use: *automated escrow*, where disbursement depends on agreed upon conditions

Caution: Smart Contracts have been found to be prone to (very expensive) bugs

BlockChain background: Repeats things we've discussed.

POTENTIAL USE CASES

Killer app: cryptocurrencies

Other potential uses:

- Reduce opaqueness of supply chains
  - One "trustless" place for all transactions along the way
  - Improvements over paper-based systems and many disjoint databases
- Eliminate middlemen
  - Why does farmer make so little and consumer pay so much?
- Reduce fraud
  - India, Russia, Sweden, Georgia… are building blockchain-based land registries to fight "land fraud" and simplify international property transactions

BlockChain background: Repeats things we've discussed.

FOR THE FOOD SUPPLY CHAIN?

Supply chain management

- Walmart is building one for the food supply chain
  - Food safety: fast identification of tainted foods
  - Consumers are demanding more information about the products they buy (organic, fair trade, …)
- Simplify international transactions

Help farmers

- Want to know what happens to their products for fair pricing
- What products should they be producing?

Reduce food scandals

- illegal production, misrepresentation, loss and waste, …

BlockChain background: Repeats things we've discussed.

INDUSTRIAL UPTAKE?

ripe.io:

- A company that is building a "blockchain of food" with IoT interfaces

Walmart:

- partnered with IBM and Tsinghua to identify sources of contaminated products and speed up recall

But today's blockchain technology may not be appropriate for all use cases

- too dependent on availability of plentiful power, networking, and storage

DESIRED BLOCKCHAIN PROPERTIES

Performance:
- High Throughput, Low Latency
- Energy-Efficient

Security:
- Always available for reading (verifying) and appending
- Fair
- Tamperproof (Integrity)
- Possibly confidentiality as well

No Single Administrative Domain
- *no need to trust a single provider*

Open membership (or not)

OPEN MEMBERSHIP IS HARD

Traditional secure logs are based on voting

Members vote on which transactions to add to the log and in what order

Problem: "Sybil" or impersonation attacks

- a participant may try to vote multiple times
- with closed membership, cryptographic signatures can identify the source of a vote
- with open membership, anybody can create identities and that way vote many times

PERMISSIONLESS VS PERMISSIONED BLOCKCHAINS

	Permissionless	Permissioned
Approach	Competitive	Cooperative
Basic technique	Proof-of-Resource	Voting
Membership	Open	Closed
Energy-efficiency	Often terrible	Excellent
Transaction rate	At best hundreds / sec	Many thousands per second
Transaction latency	As high as many minutes	Less than a second

BITCOIN BLOCKCHAIN

Permissionless, open membership

Proof-of-Work

There are thousands of Bitcoin miners

- they use ASIC hardware to compute SHA256 hashes
- use about more energy than the country of Denmark

Overall rate is a few transactions per second

THE BLOCKCHAIN

Ledger

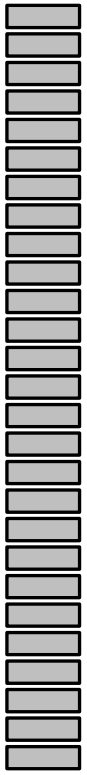

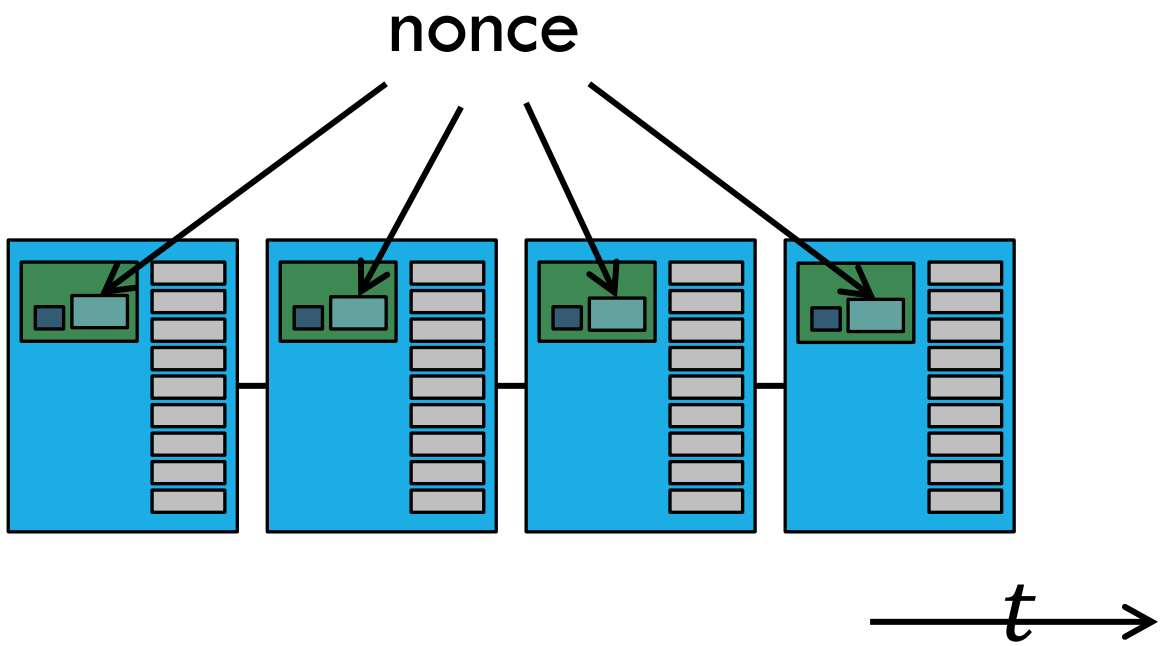

HASH(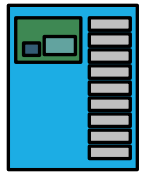) < ***target***

"cryptopuzzle"

THE BLOCKCHAIN

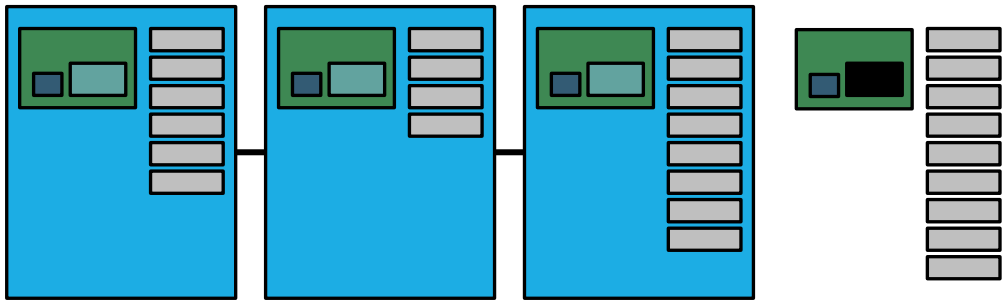

THE BLOCKCHAIN

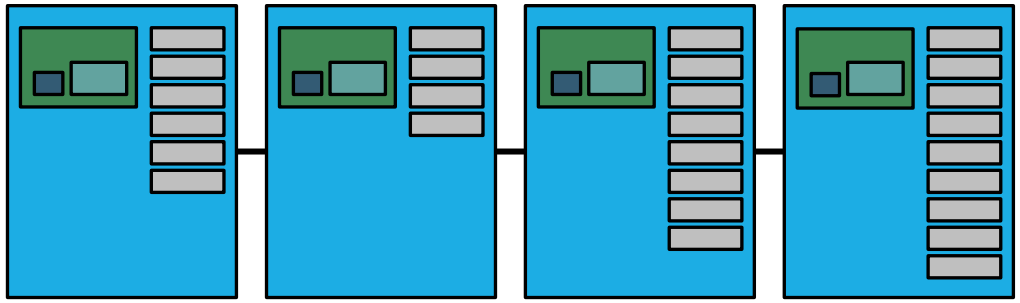

THE BLOCKCHAIN

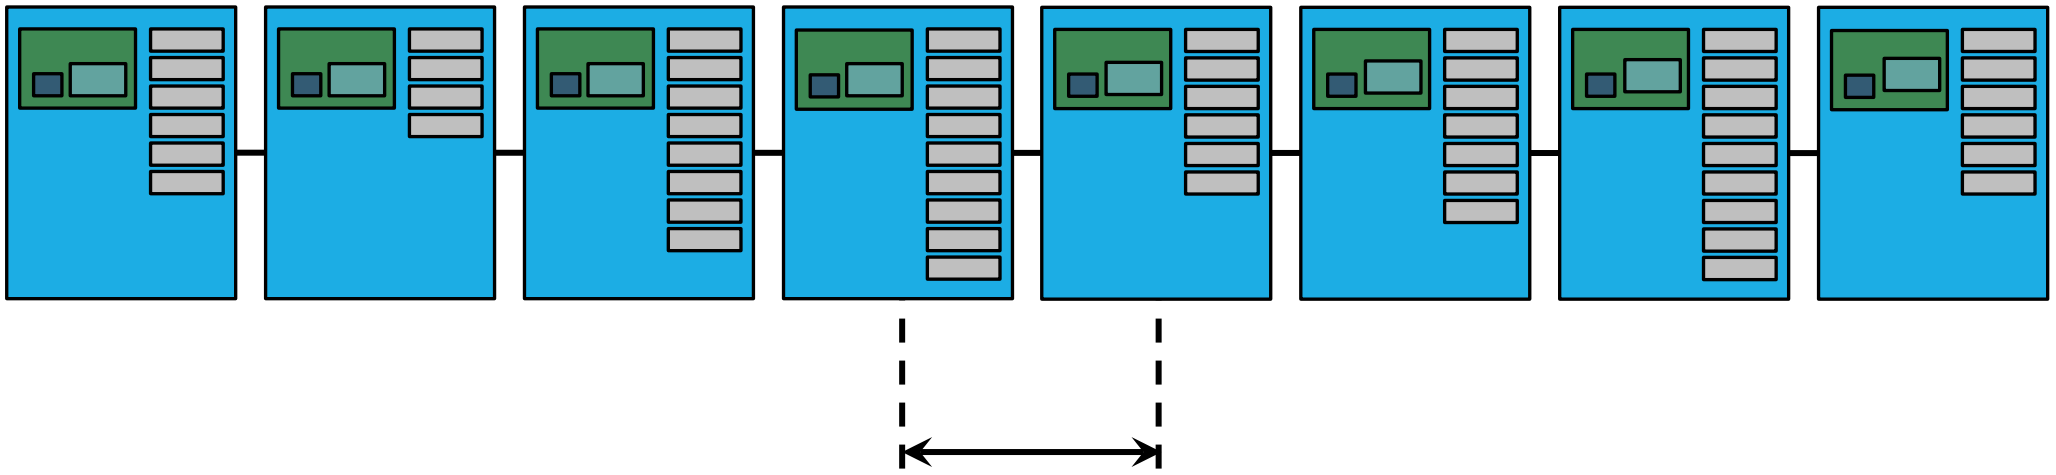

Exponentially distributed rate of new blocks, with constant mean interval

target automatically adjusted every 2016 blocks so that mean interval is **10 minutes**

INCENTIVES FOR MINING

Prize:

- **"Minting"**
- **Transaction Fees**

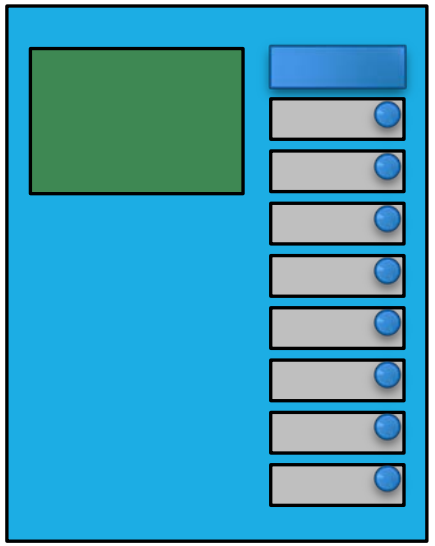

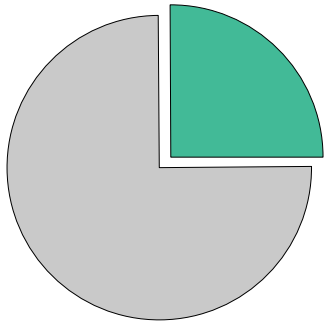

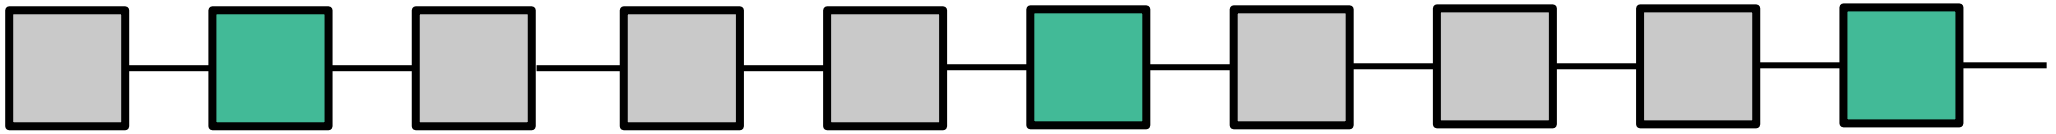

Wins *proportional to computation power*

FORKS

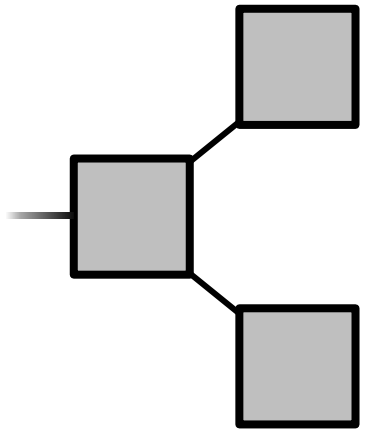

Two blocks “mined” at approximately the same time by two different miners

FORK RESOLUTION

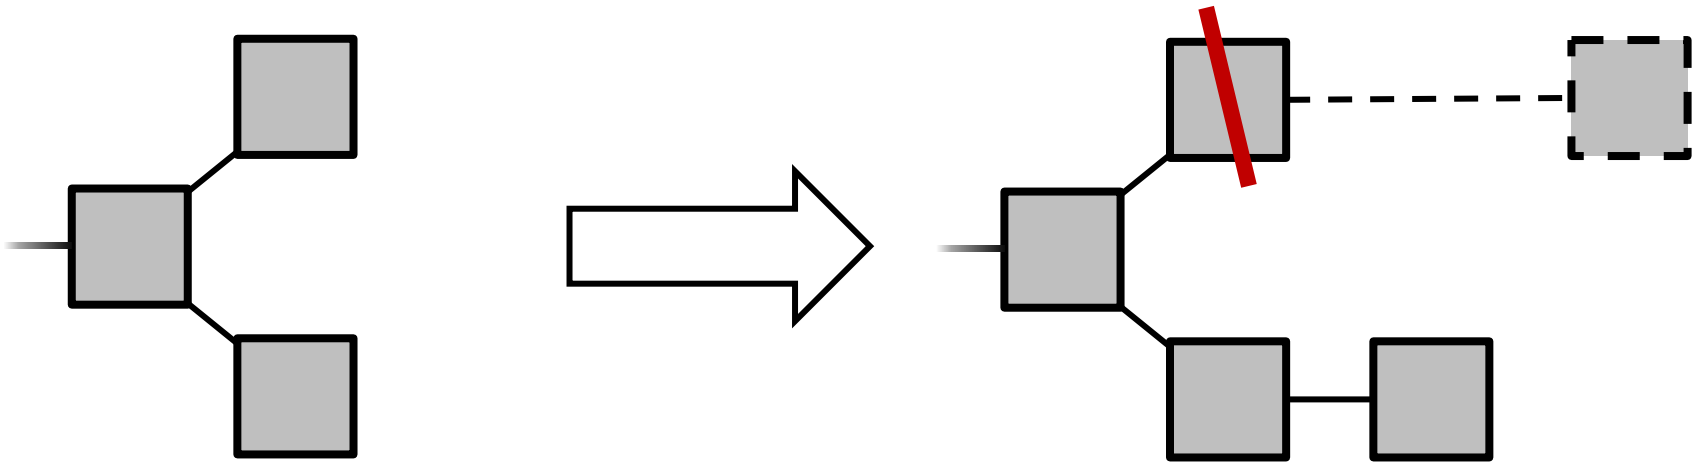

- **Longest** chain wins
- Transactions on short chain are reverted

FORK RESOLUTION

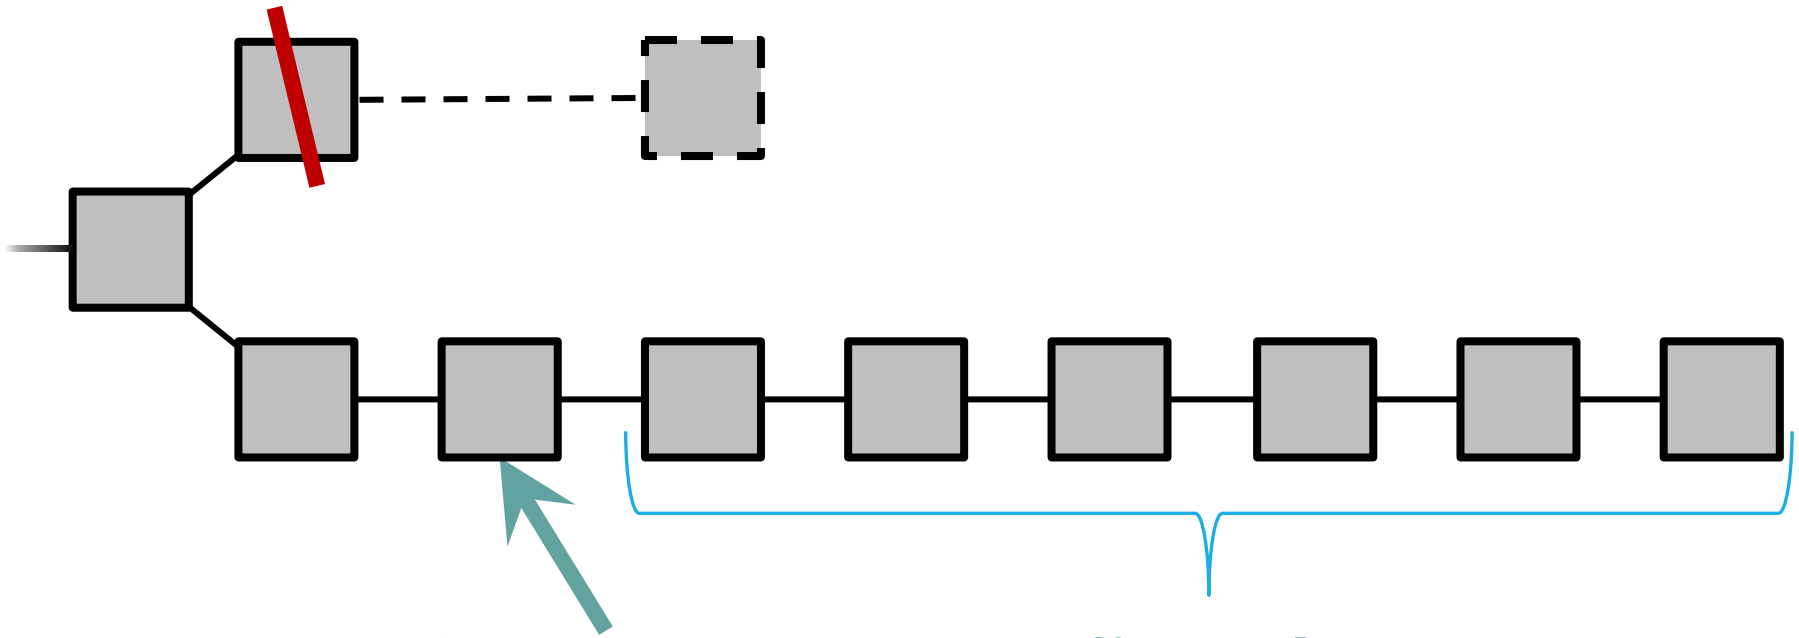

A transaction is **confirmed** when
it is **buried** "deep enough"
(typically 6 blocks – i.e., one hour)

SECURITY THREAT!

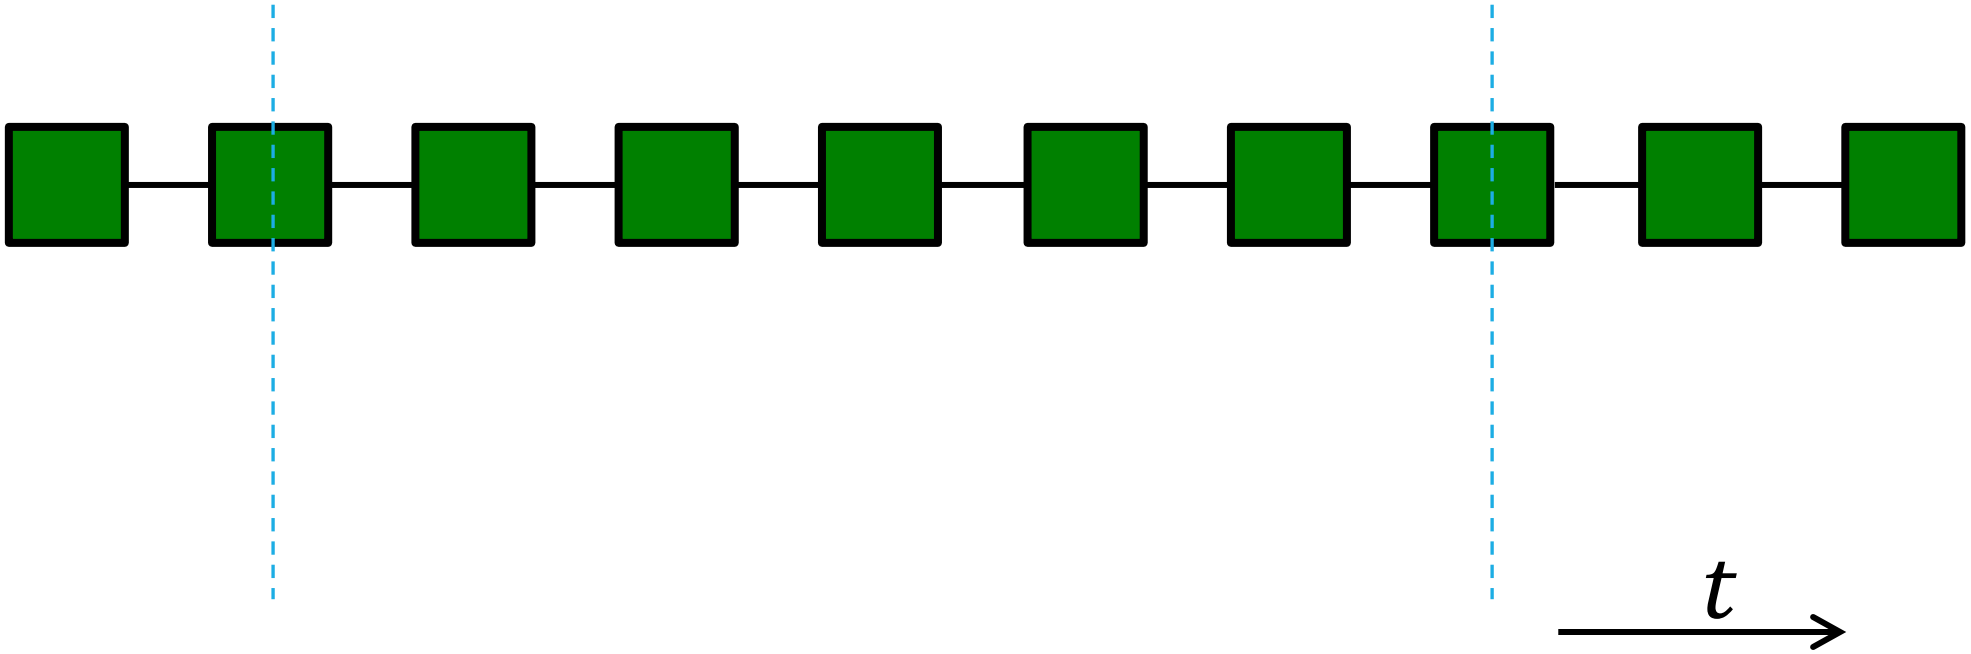

SECURITY THREAT!

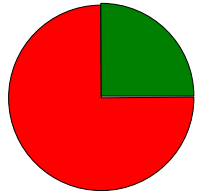

Attacker secretly controls ¾ of the available compute power

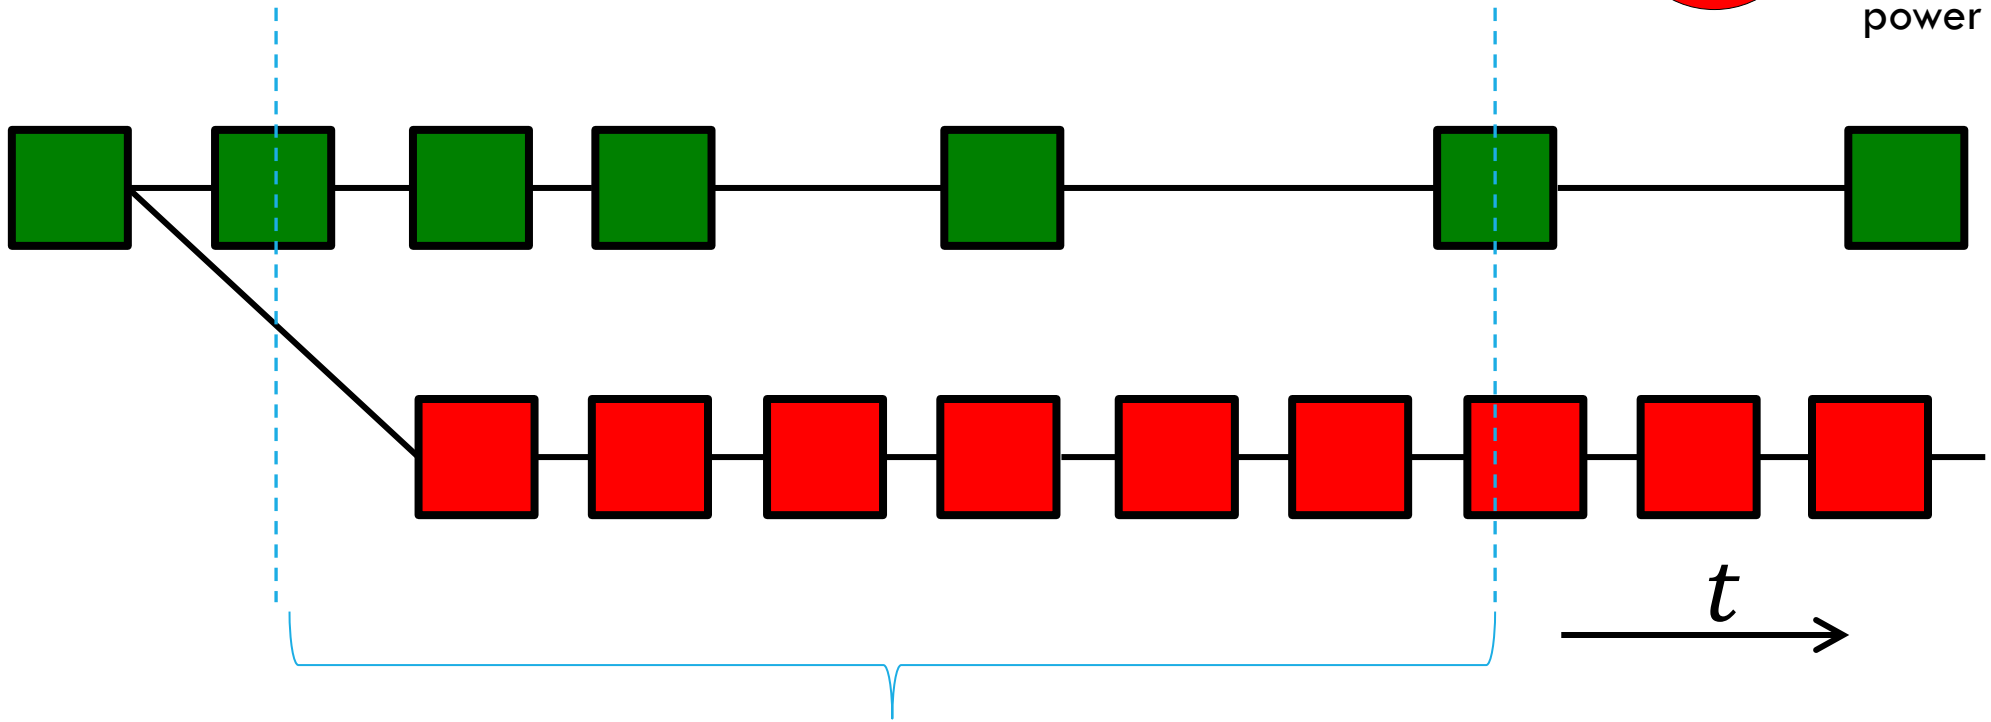

Red (attack) generates 6 blocks in time that the good miners generate 4

Threat: attacker outruns good miners, causing their transactions to roll back

SECURITY THREAT!

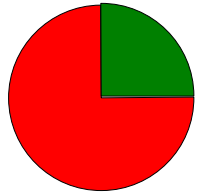

Attacker secretly controls ¾ of the available compute power

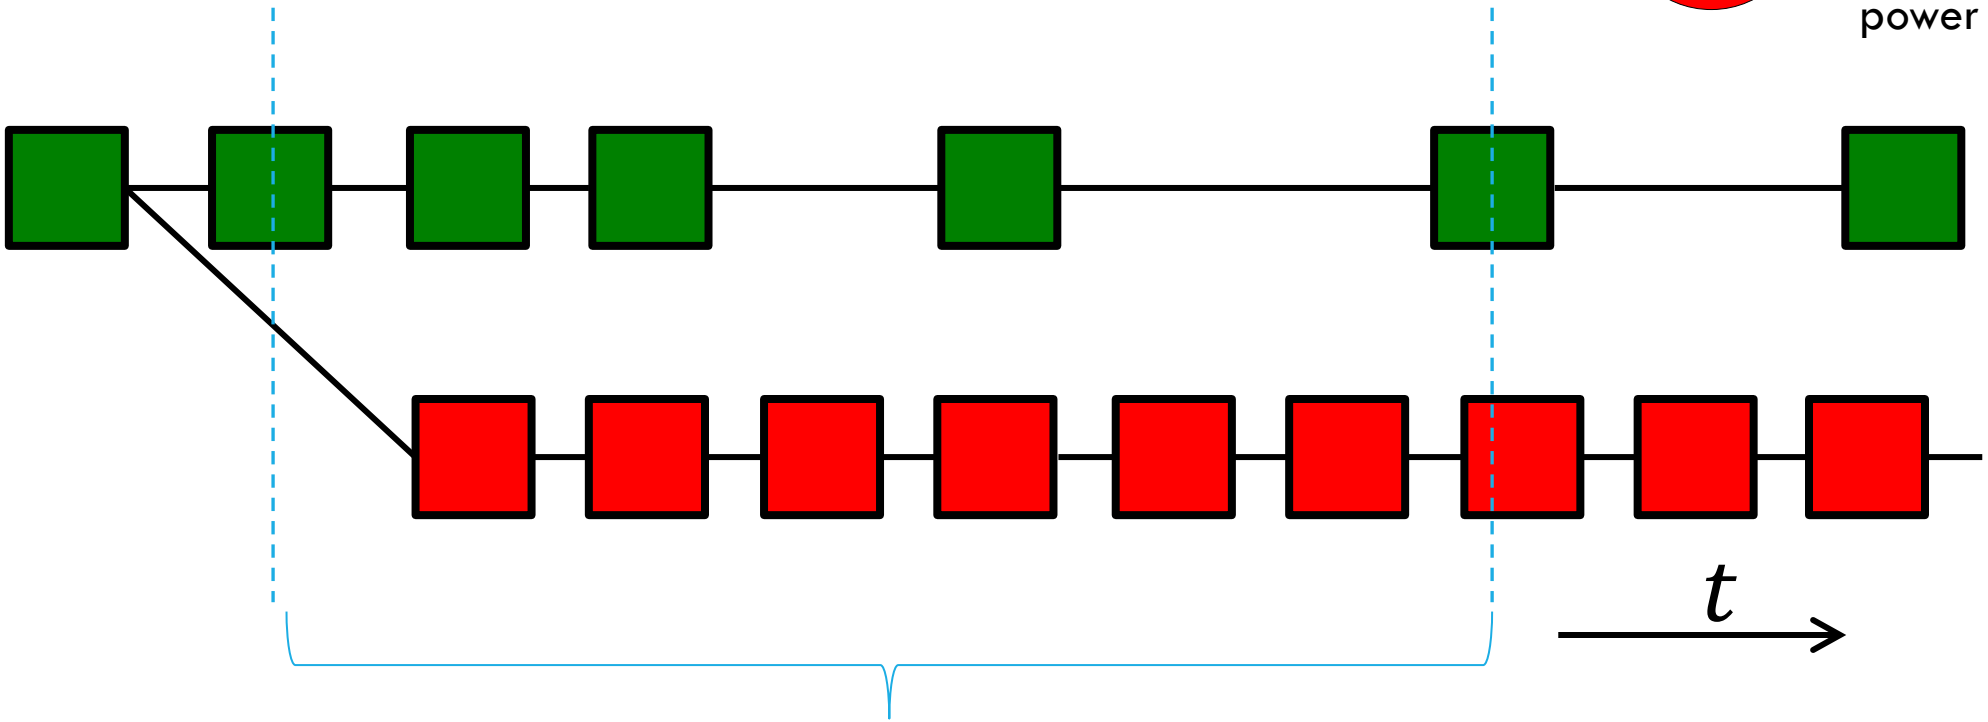

Red (attack) generates 6 blocks in time that the good miners generate 4

Threat: attacker outruns good miners, causing their transactions to roll back

RESPONSE: AN ASSUMPTION!

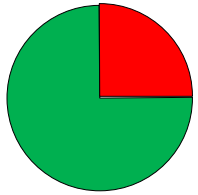

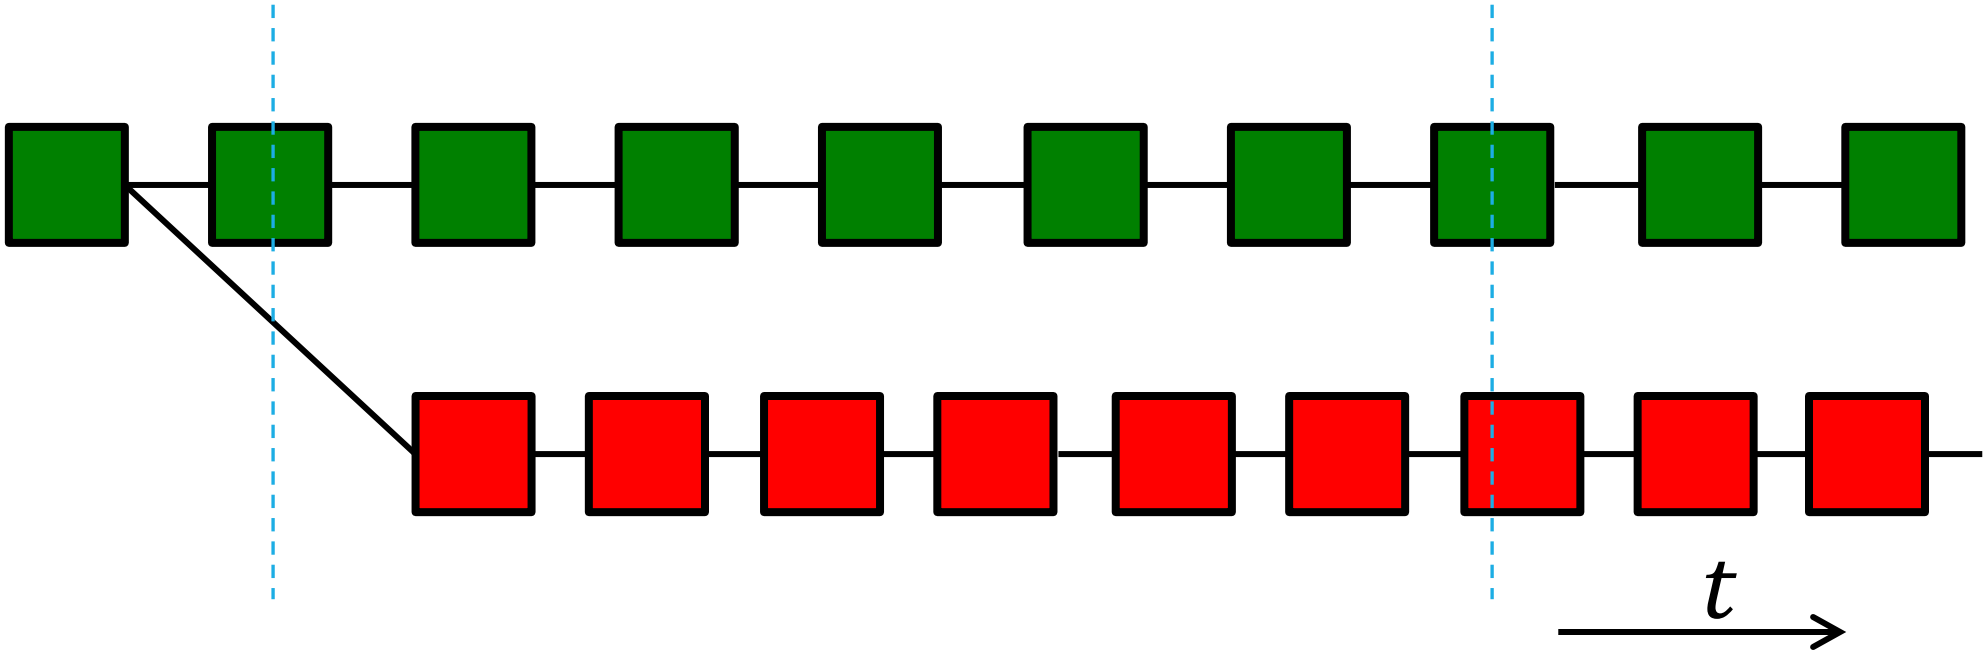

➔ **Security Assumption:** *good miners own >.5 of the total compute power*

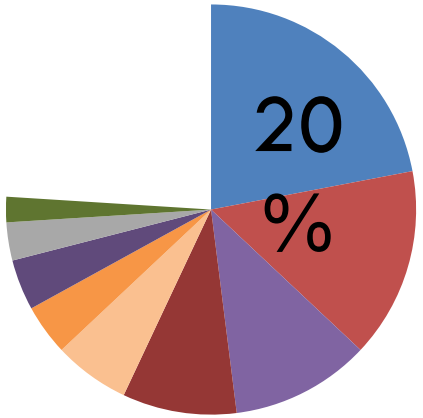

[blockchain.info, April 2015]

PERMISSIONLESS BLOCKCHAINS

Open membership, but inefficient

Vulnerable to 50% attacks

Examples include Bitcoin, Ethereum, IOTA

PERMISSIONED BLOCKCHAINS

Performance:

- High Throughput, Low Latency
- Energy-efficient

Security:

- No forks!

Closed membership

Examples include Ripple, Hyperledger

BLOCKCHAIN FOR THE FARM?

Blockchains require strong network connectivity and lots of storage

Permissionless blockchain are power-hungry

Sensors have limited resources

- Sensors for growing conditions, storage conditions, shipping conditions, …

*Blockchain for a farm will generate records in a decentralized way, and hence it *must* work in a network-partitioned or -challenged environment*

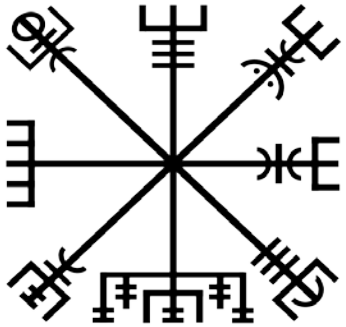

Vegvisir: tolerate branches

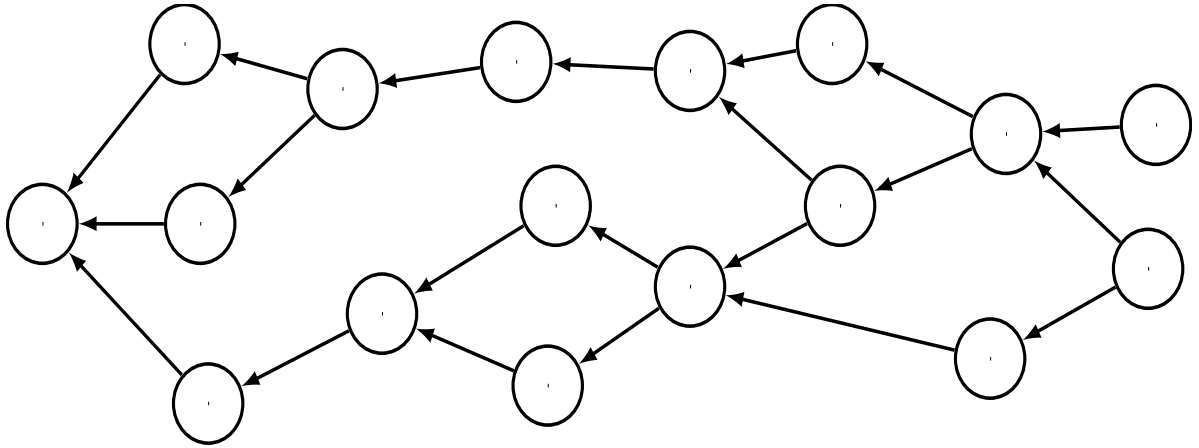

- The key innovation is to allow branching as a feature.
- Leads to DAG structure instead of linear blockchain
- DAG is a full causal history of events (respect's Lamport's $\rightarrow$), but note that the DAG arrows run in the opposite direction than Lamport's $\rightarrow$

WHAT ABOUT THE ← NOTATION?

When Lamport talks about a → b, this is a statement about causality.

In the Vegvisir DAG, nodes represent operations but the notation is reversed: a ← b, meaning "b came after a." They made this choice because b "countersigns" a, in the actual blockchain

As long as you understand the notation it is easy to understand their examples.

PROOF-OF-WITNESS TO PERSIST BLOCKS SECURELY

No more than *k* malicious nodes in any neighborhood

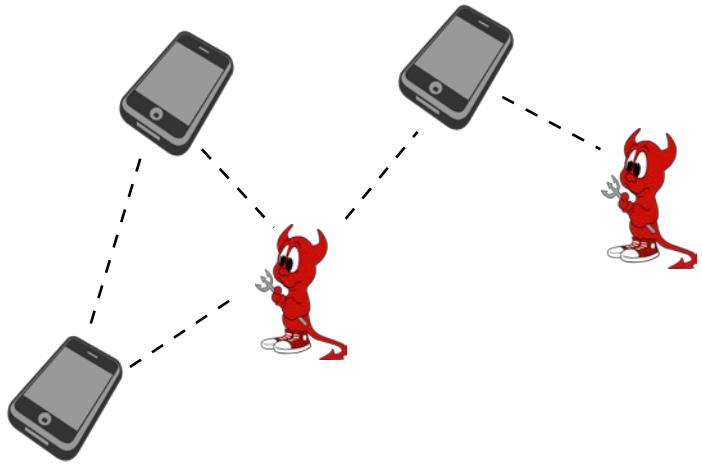

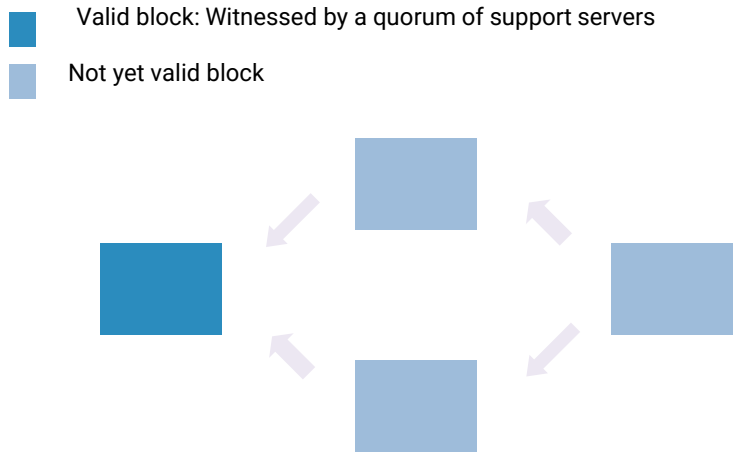

At least one copy of a valid block will survive if <*k* malicious peers

The Support Blockchain reduces sensor storage needs

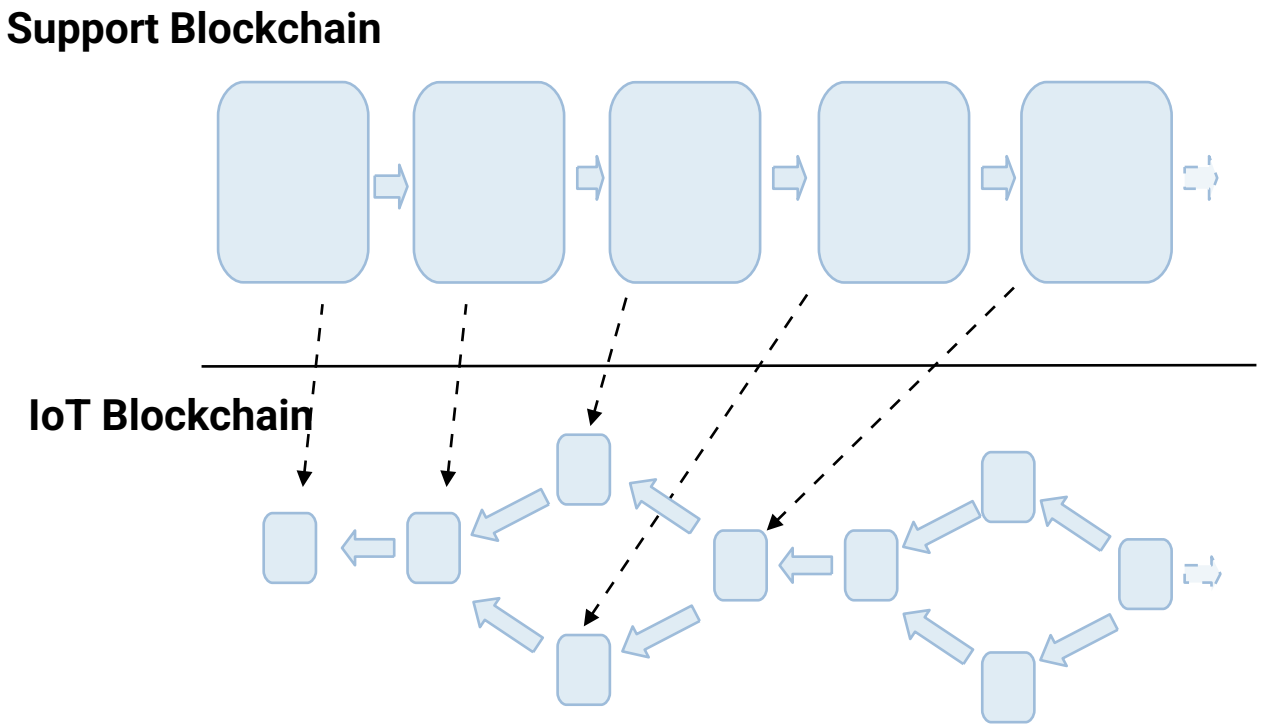

Allows regular peers to discard old blocks when storage space is low

Blocks are *gossiped* over ad hoc network

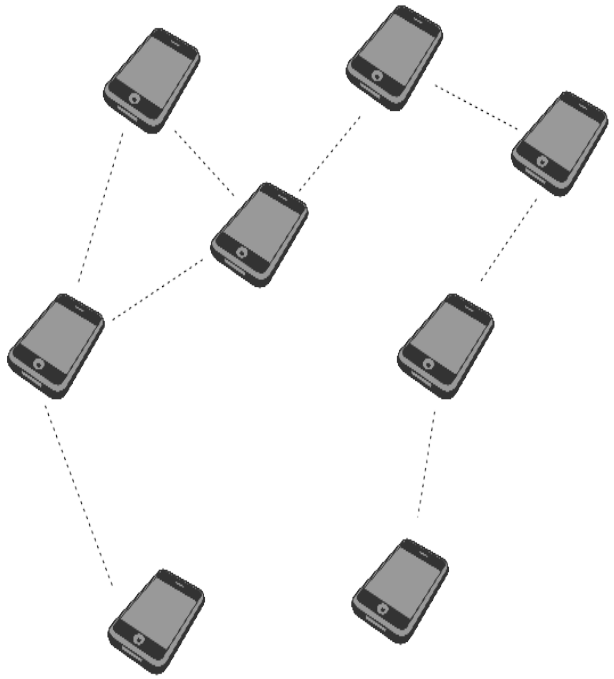

Heterogeneous, opportunistic networking

WHAT WOULD YOU FIND INSIDE A RECORD?

There are two kinds of transactions.

One kind spends coins – these need to validated using a total ordering

The others are commutative and the DAG ordering is sufficient. These are composed entirely from CRDT operations, explained on the next slide

CRDTS: STRONG SEMANTICS IN DAGS

- *Conflict-Free Replicated Datatype*
- Updates must be associative, commutative, idempotent
- Replicas can be updated independently and concurrently
- Basic CRDTs form registers, counters, sets

TRANSACTIONS MANIPULATE CRDTS

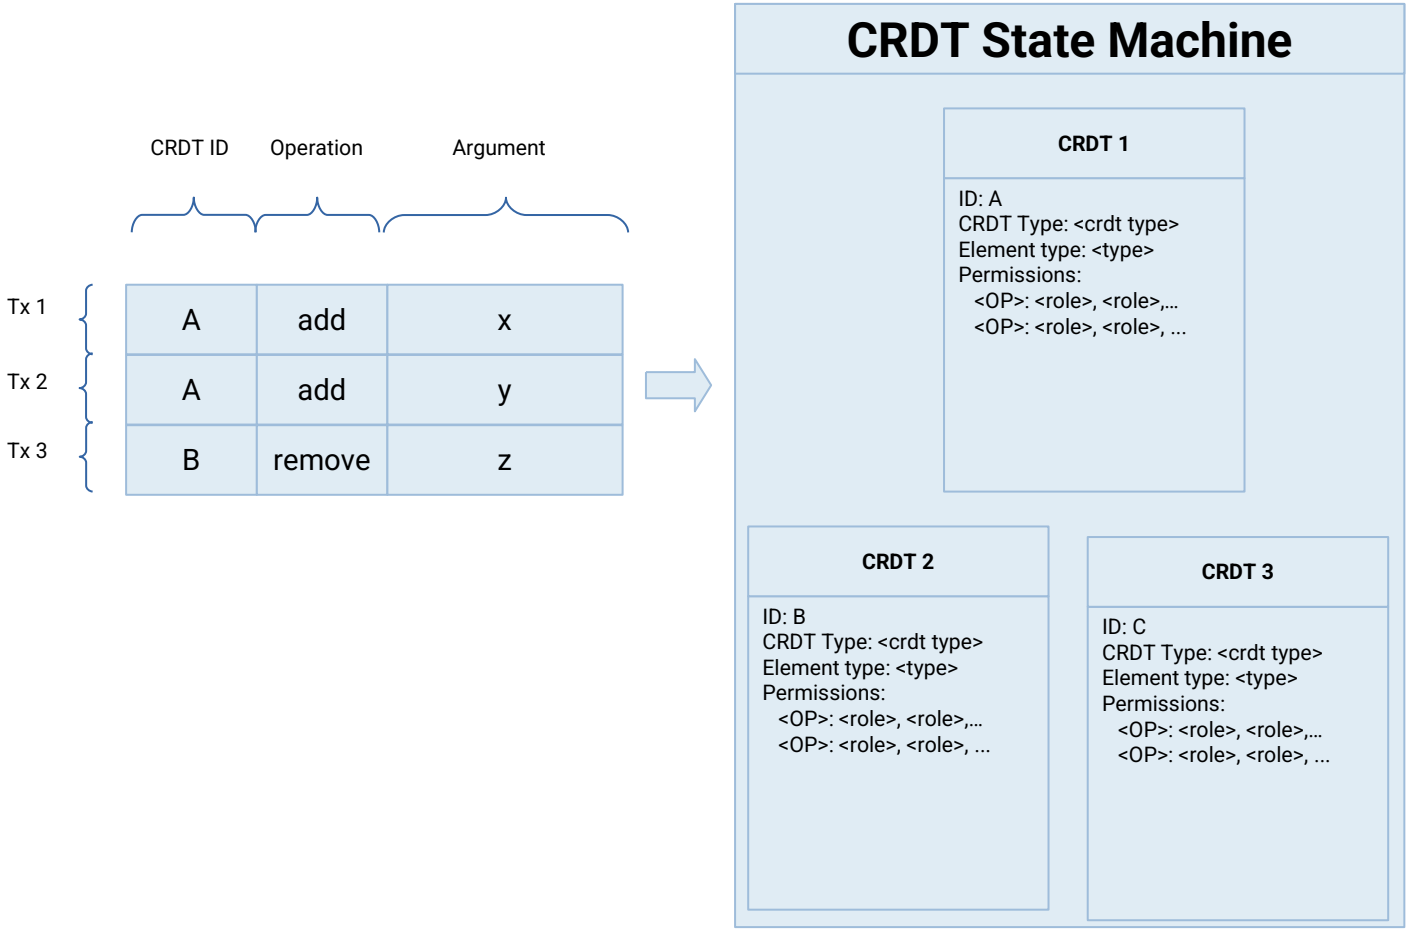

DEMONSTRATION VIDEO

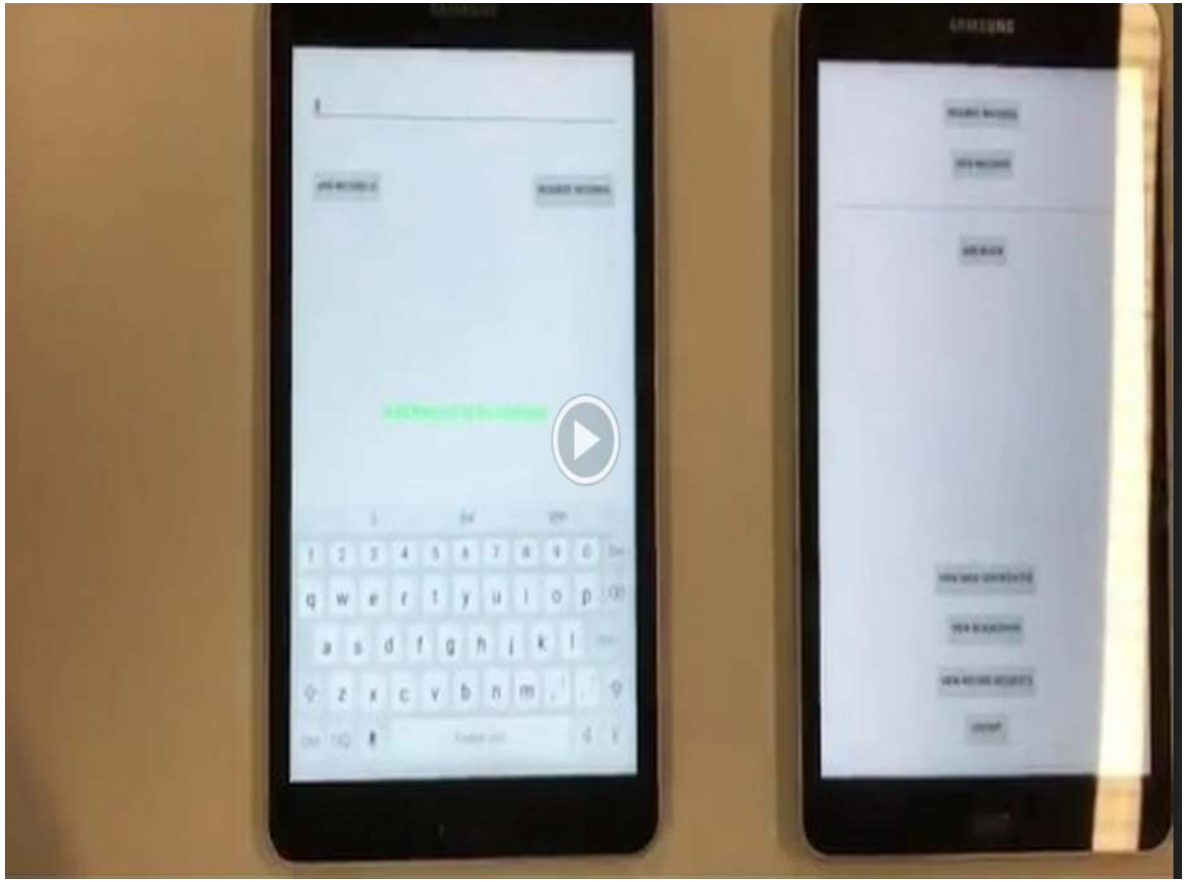

Proof of concept system

BUT SOME QUESTIONS REMAIN OPEN

When we merge the DAG to form a total order we might discover double-spending of cybercurrency. Vegvisir doesn't currently offer a solution here: it is easy to impose a total order, but only during periods of connectivity.

Vegvisir also leaves open many of the other questions we touched on, that an auditor might want to see answered.

WRAPUP: SOME QUESTIONS ABOUT ETHICS

We have seen some exciting uses of blockchain, yet we also know that cybercurrency is popular for illegal purposes too, and even legitimate speculators may seem to be evading taxes (these speculators disagree!)

So, is it acceptable to work with a technology that has these issues, or are blockchains and bitcoins tainted by the bad uses?

Should we distinguish smart contracts from cybercurrency? They often live in the same chains…

CONCLUSION

Exciting possibilities for blockchains in the food supply chain

But current blockchain designs may not be compatible with some deployment scenarios in the food supply chain

Vegvisir supports partitioned operation and has low power/networking/storage requirements