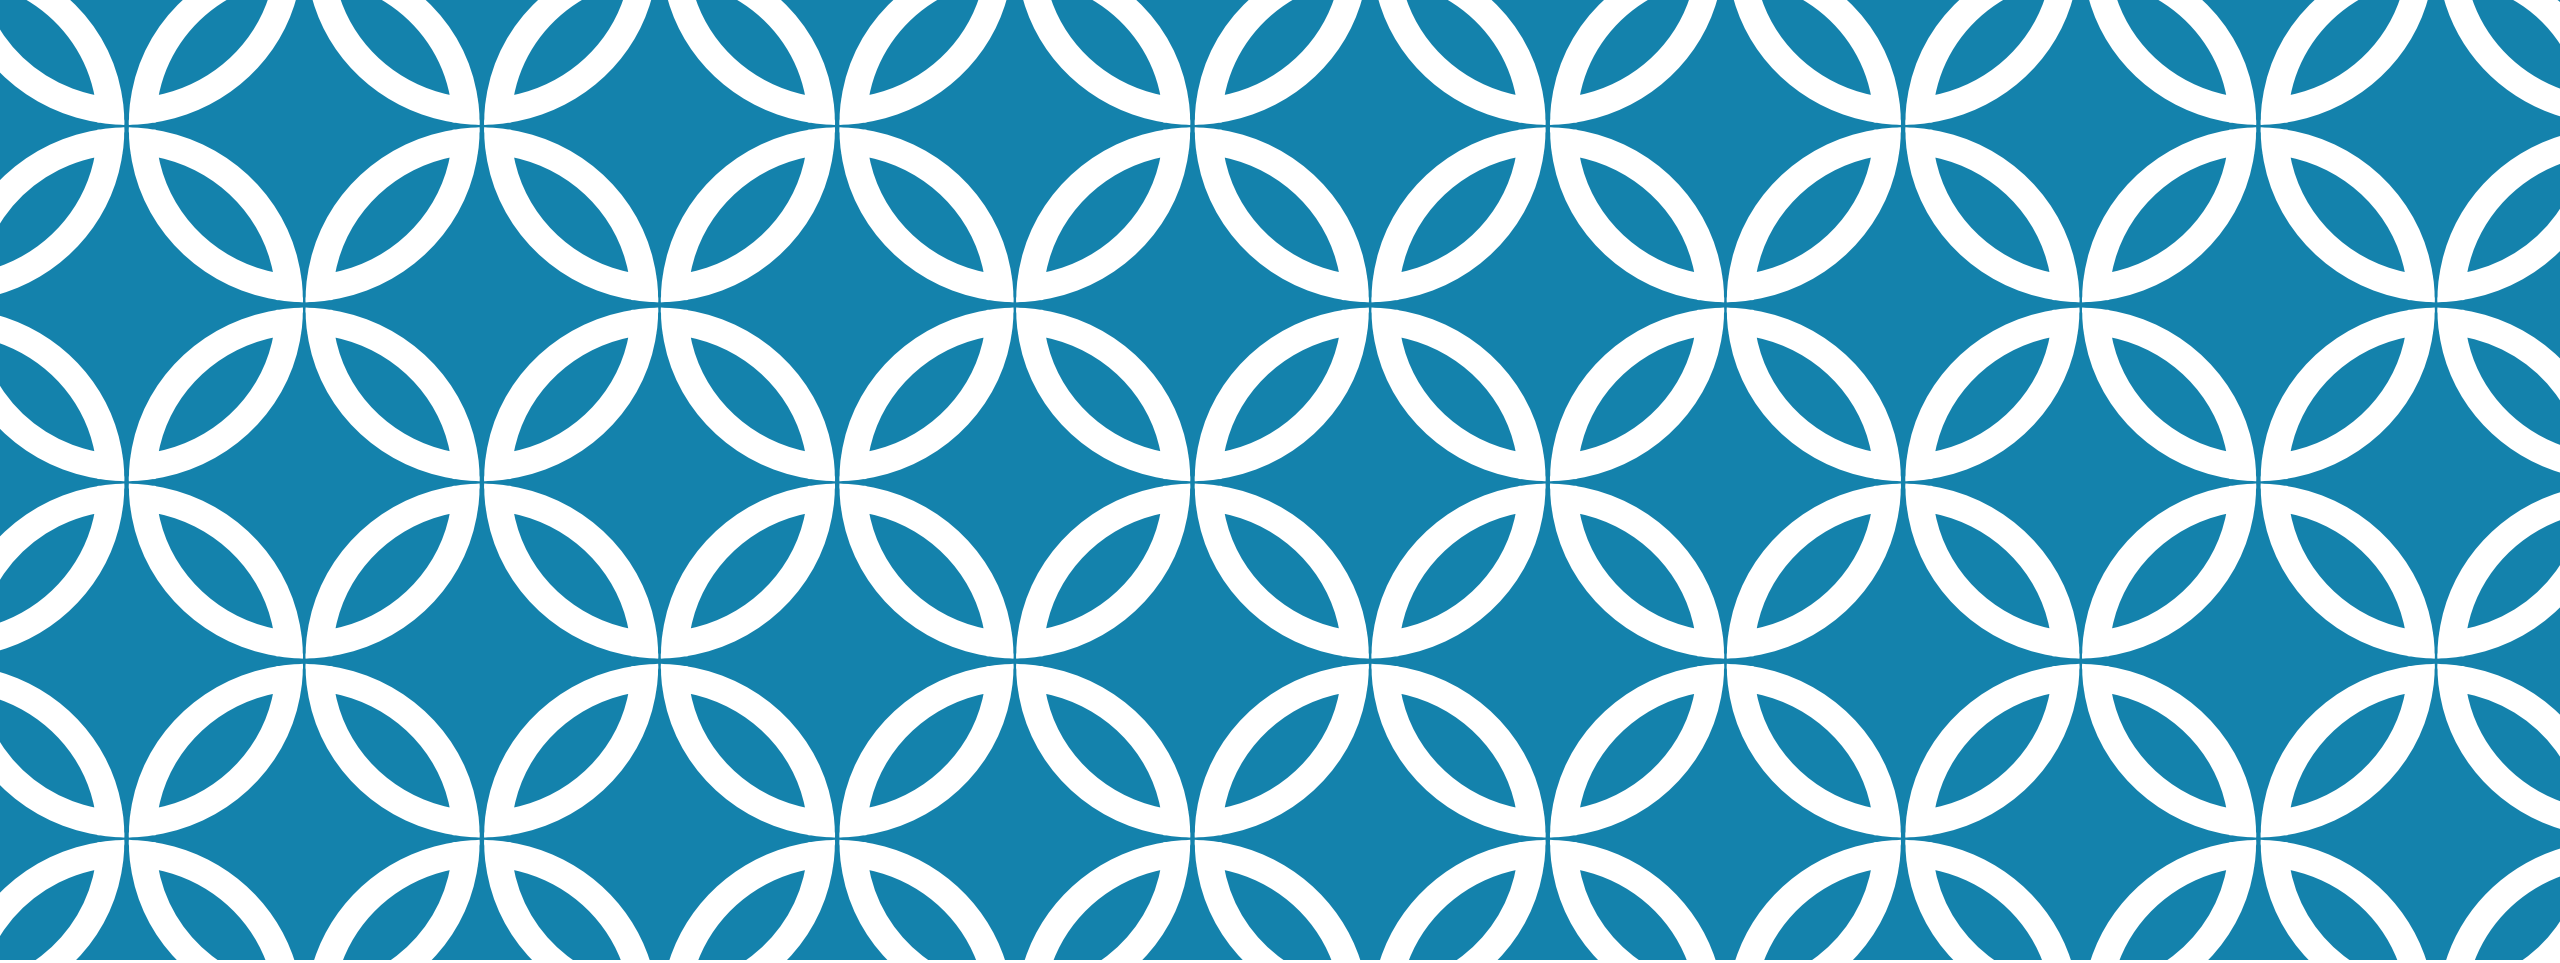

CS5412/LECTURE 14 BLOCKCHAINS FOR IOT (PART 2)

Ken Birman

CS5412 Spring 2019

SUMMARY OF BLOCKCHAIN CONCERNS

Permissioned or Permissionless? Energy cost of permissionless block mining.

How do I know that the sensor data recorded into my blockchain is "legit"?

Cost of verifying that the BlockChain hasn't been tampered with.

National-scale "disruption" scenarios that 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.

WALKING THROUGH THE ISSUES

Which issues would arise on a real "smart farm"?

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

What limits the speed of new technology adoption?

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."

BIGGEST CONCERN OF ALL

What Rumsfeld called the “Unknown unknowns”.

It ain’t what you know that’ll get you. It’s what you know that just ain’t so.

-- Attributed to Samuel Clemens (Mark Twain)

ONE OF KEN'S "TALES OF WOE"

We take client-server computing for granted… but it was nearly on one of those gravestones!

In the earliest days, Digital Equipment Corporation invented client-server with the introduction of their VaxClusters architecture. A huge advance!

The market was exponential – DEC was on track to become the largest computer company ever. *Then it suddenly imploded and was acquired!*

WHY DID EARLY CLIENT-SERVER FAIL?

The concept was absolutely right! Yet the customers were unhappy.

The early products felt like a research prototype, not professional.

The “life cycle” of a full client-server deployment had not yet been thought through, hence there were a huge number of missing tools and features.

REMINDER: CROSSING THE CHASM (MOORE, 2009)

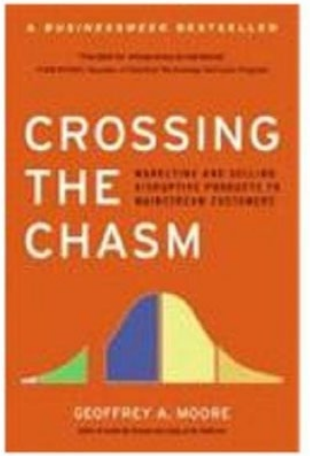

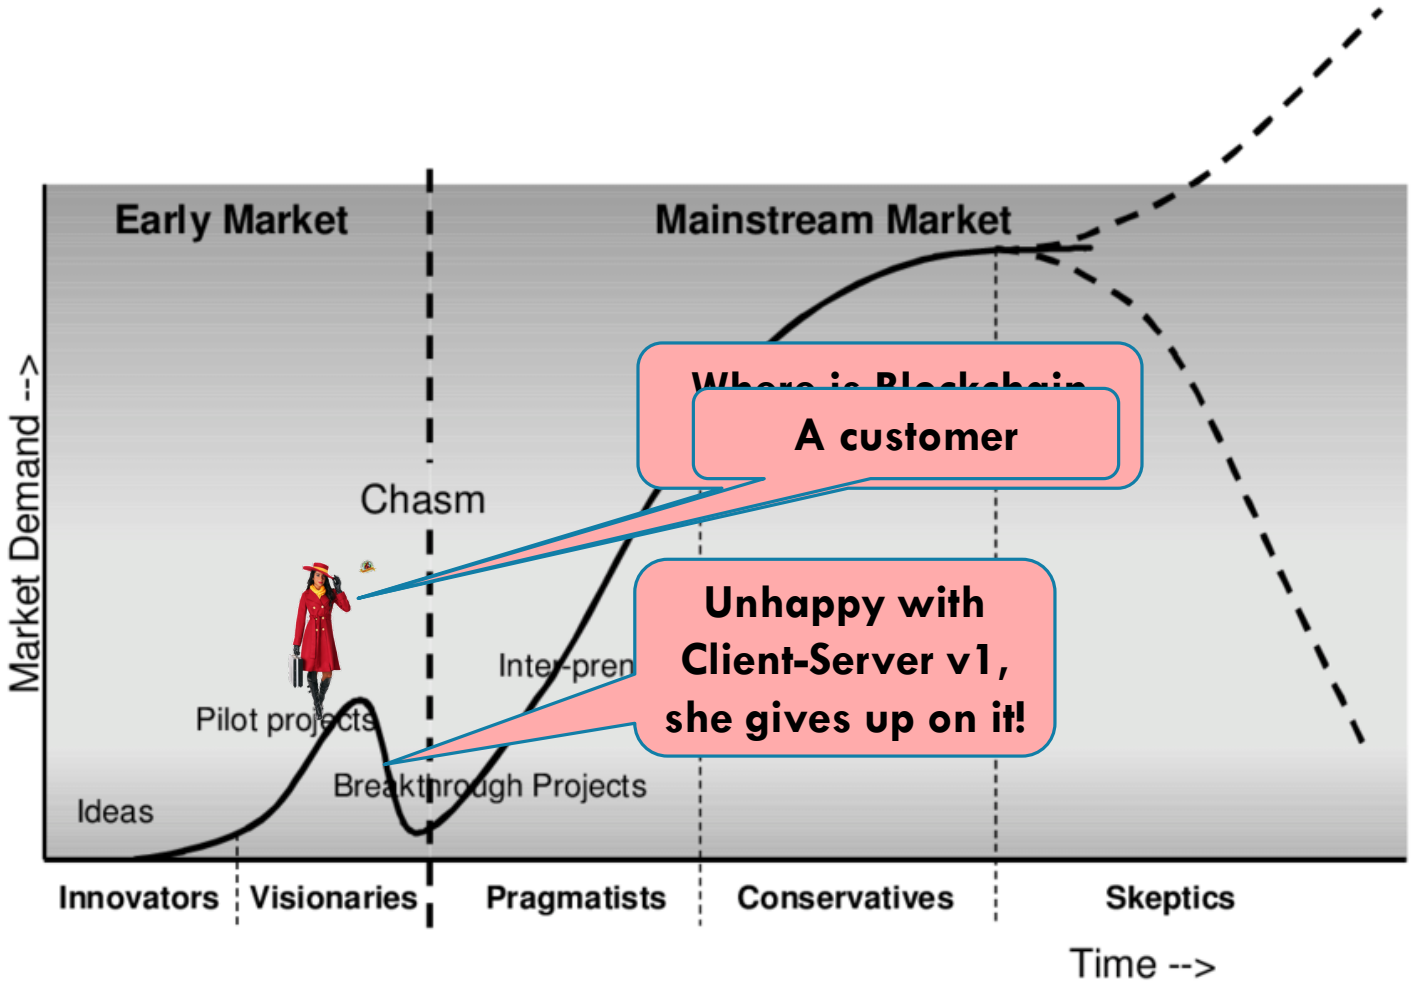

FACTORS THAT LIMIT UPTAKE

Early enthusiasm / bleeding edge always moves on to the next better thing, so the first adopters are certain to wander away.

Conservative customers want to be "the first to be last" and wait for the mainstream uptake to occur.

If you move too quickly, you simply overextend and fall into the chasm.

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 sensor record?

TRUST WITH SENSORS

We saw that Azure IoT Hub is a professional-quality sensor security and management solution as of 2019. But what can it actually do?

It 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.

CAN WE ANSWER THE MAIN QUESTIONS?

Is this the proper sensor to measure pH or photograph the milking station at 10:02am on Monday morning on Smith Farm?

Was the sensor working properly?

Was the record later modified in some way?

CAN WE ANSWER THE MAIN QUESTIONS?

Is this the proper sensor to measure pH or photograph the milking station at 10:02am on Monday morning on Smith Farm?

- Perhaps not. Sensors are often taken out of service, or upgraded…

Was the sensor working properly?

- Perhaps not. Think back to Meta, the sensor fault-tolerance technology.

Was the record later modified in some way?

- Blockchain could detect this, if you trust the operator of the service.

REQUIREMENT?

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 AbeStrouck.

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!

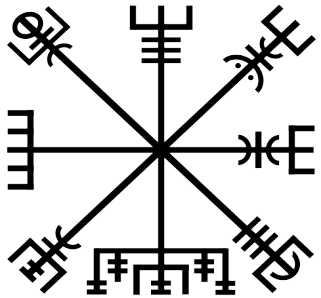

VEGVISIR

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

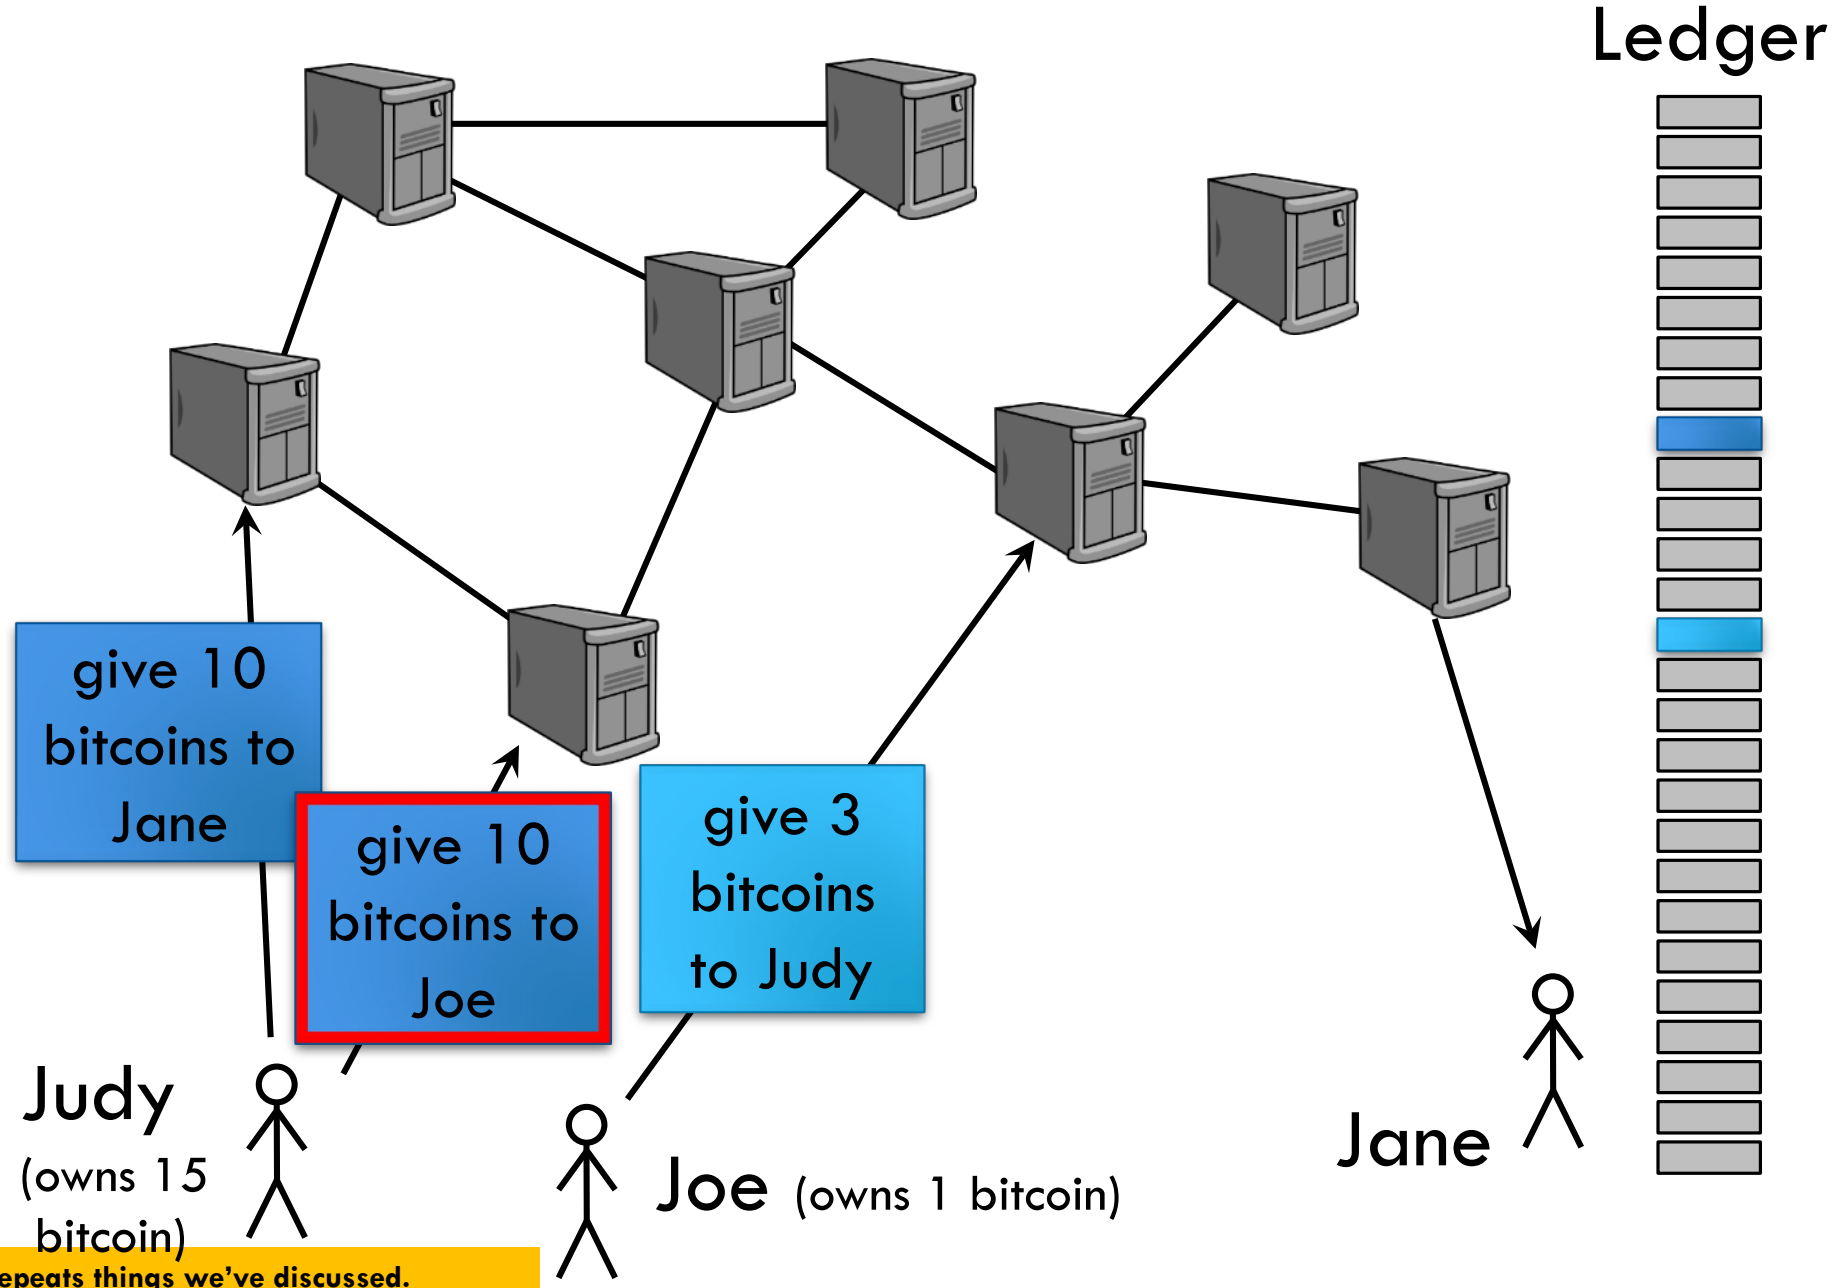

Ittay Eyal, May '16

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

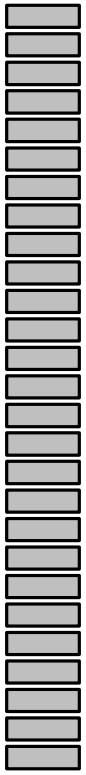

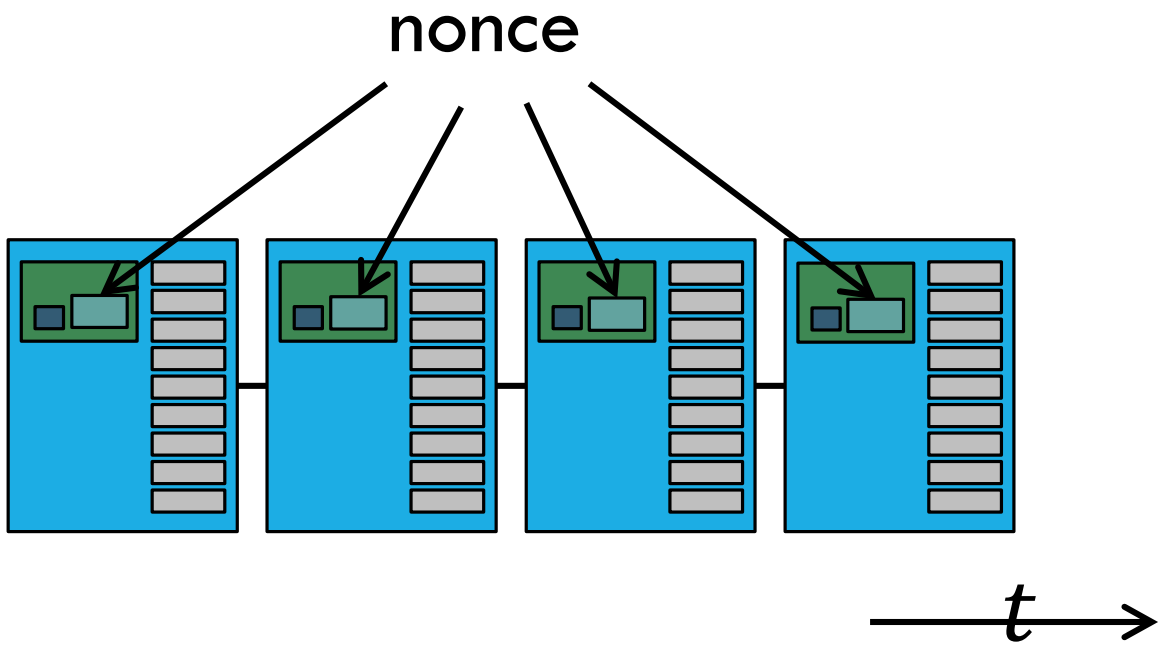

HASH(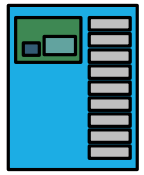) < ***target***

"cryptopuzzle"

THE BLOCKCHAIN

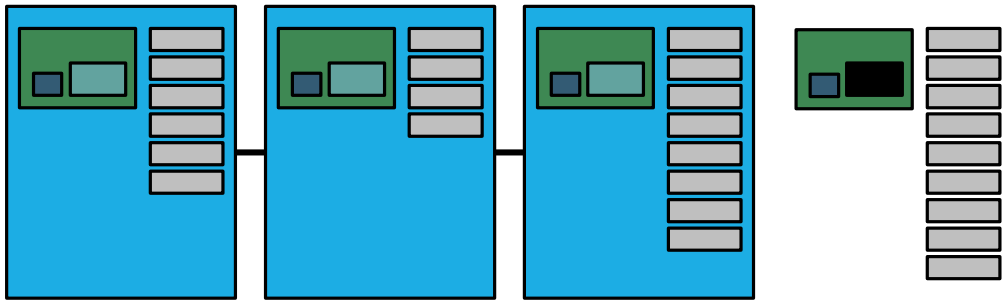

THE BLOCKCHAIN

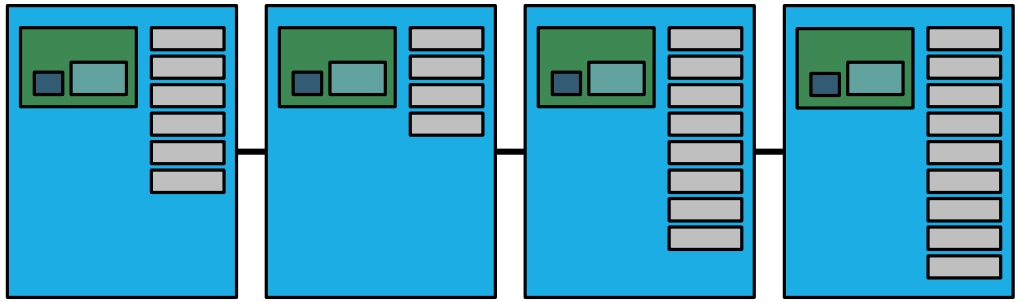

THE BLOCKCHAIN

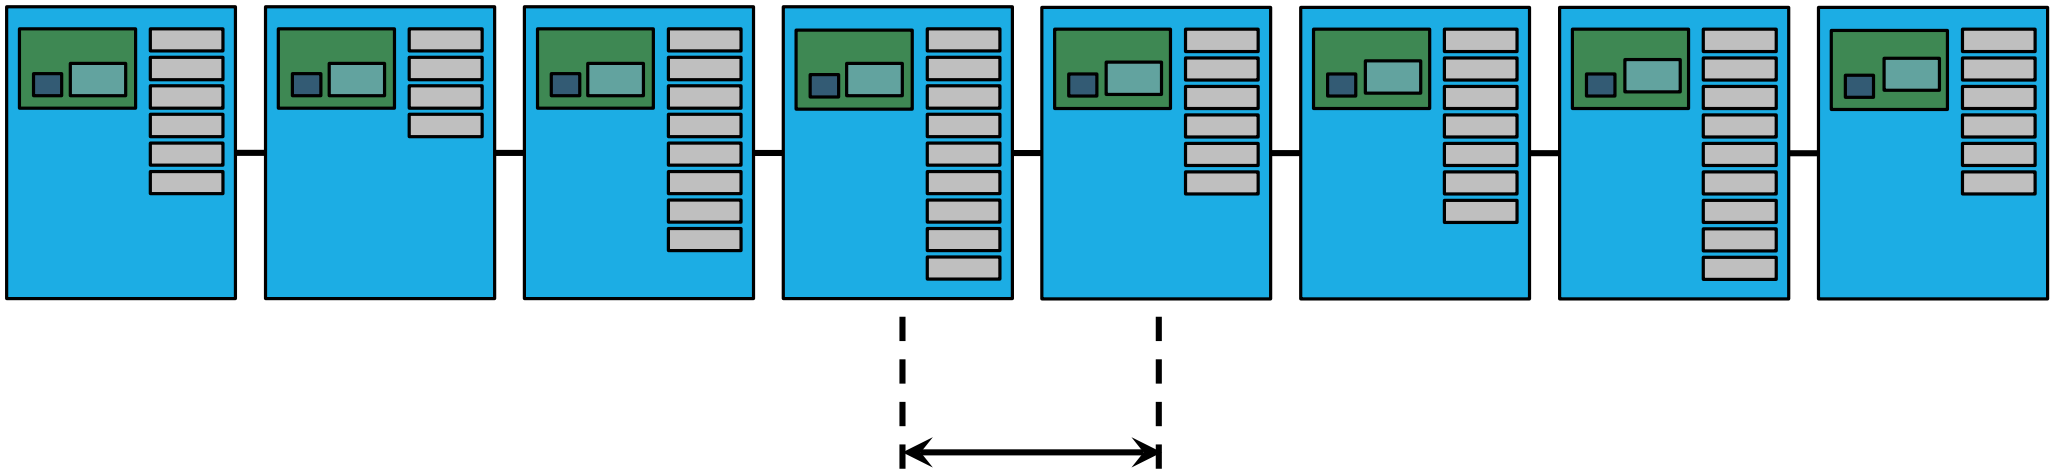

Exponentially distributed rate of new blocks, with constant mean interval

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

INCENTIVES FOR MINING

Prize:

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

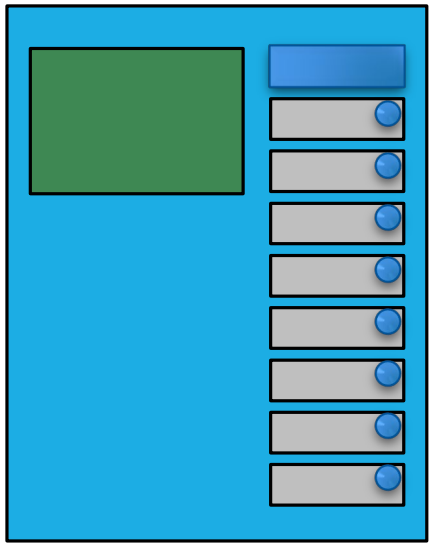

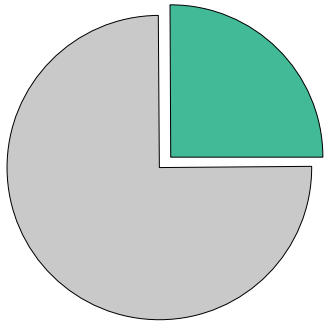

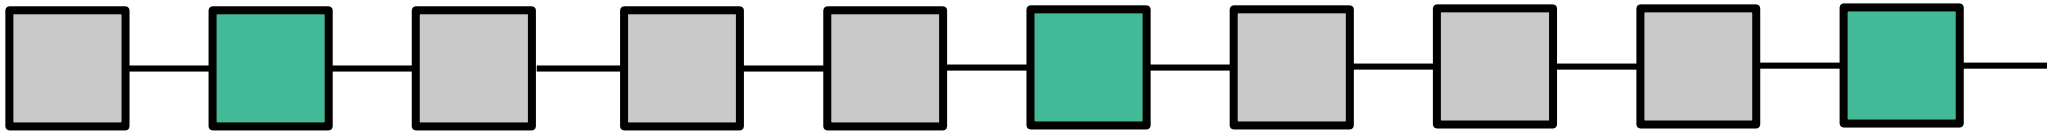

Wins *proportional to computation power*

FORKS

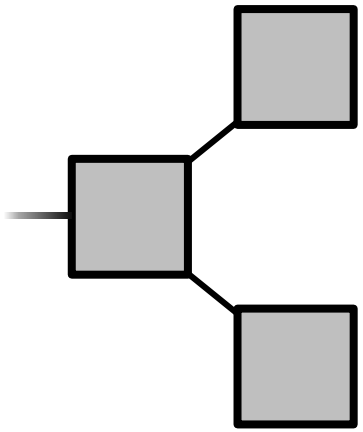

Two blocks “mined” at approximately the same time by two different miners

FORK RESOLUTION

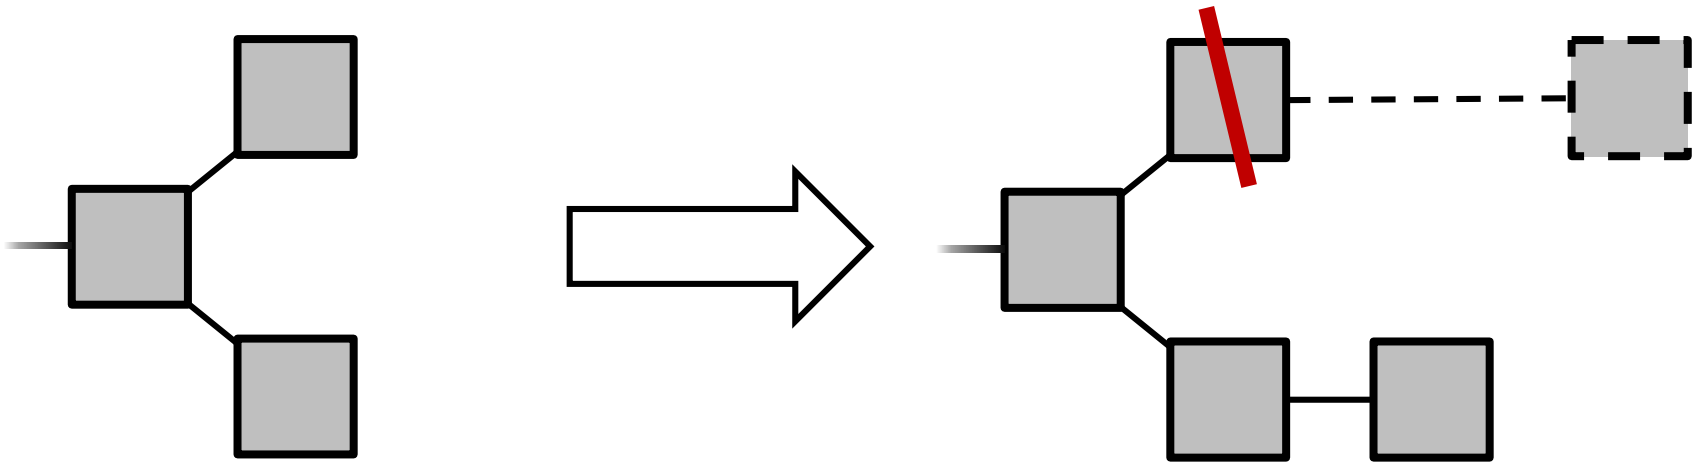

- **Longest** chain wins
- Transactions on short chain are reverted

FORK RESOLUTION

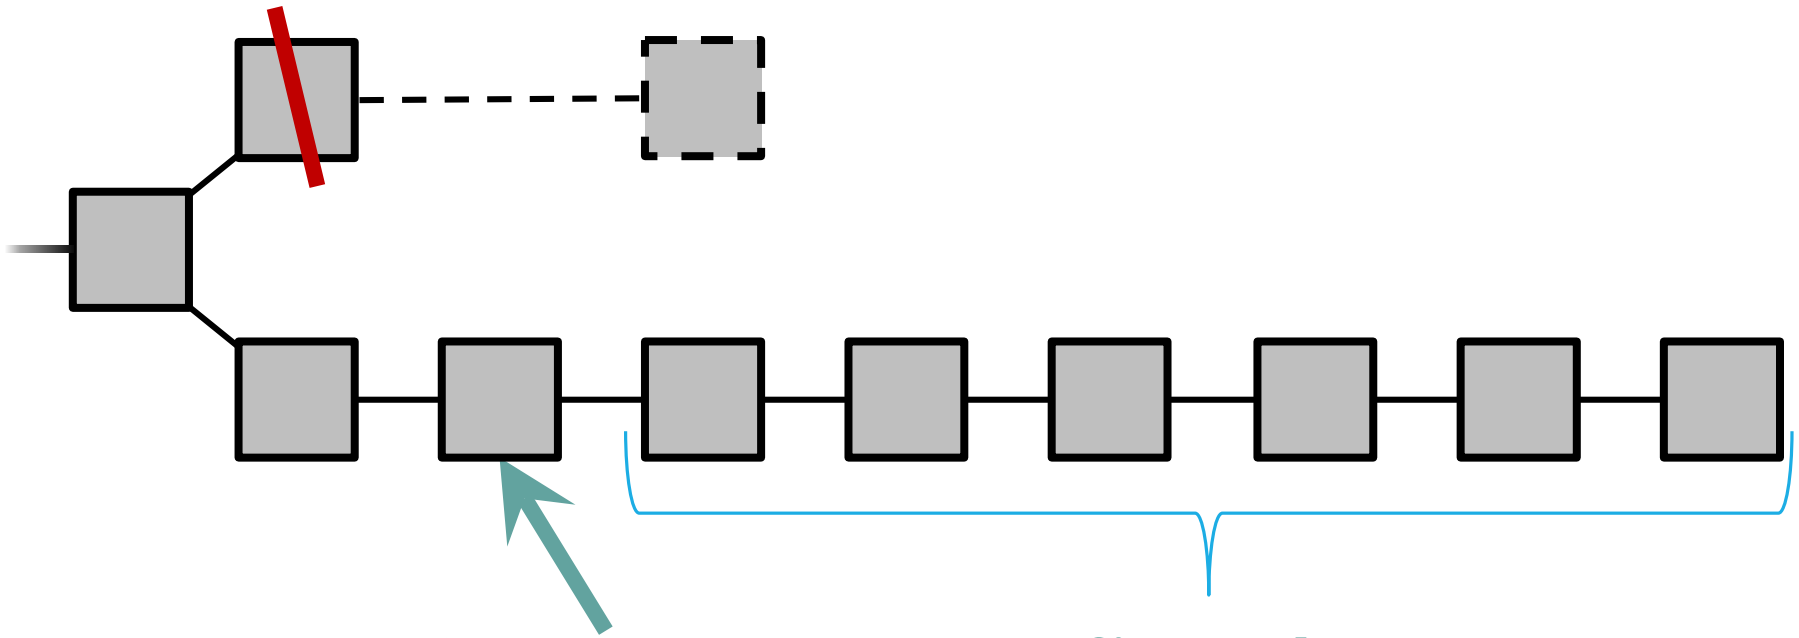

A transaction is **confirmed** when
it is **buried** "deep enough"
(typically 6 blocks – i.e., one hour)

SECURITY THREAT!

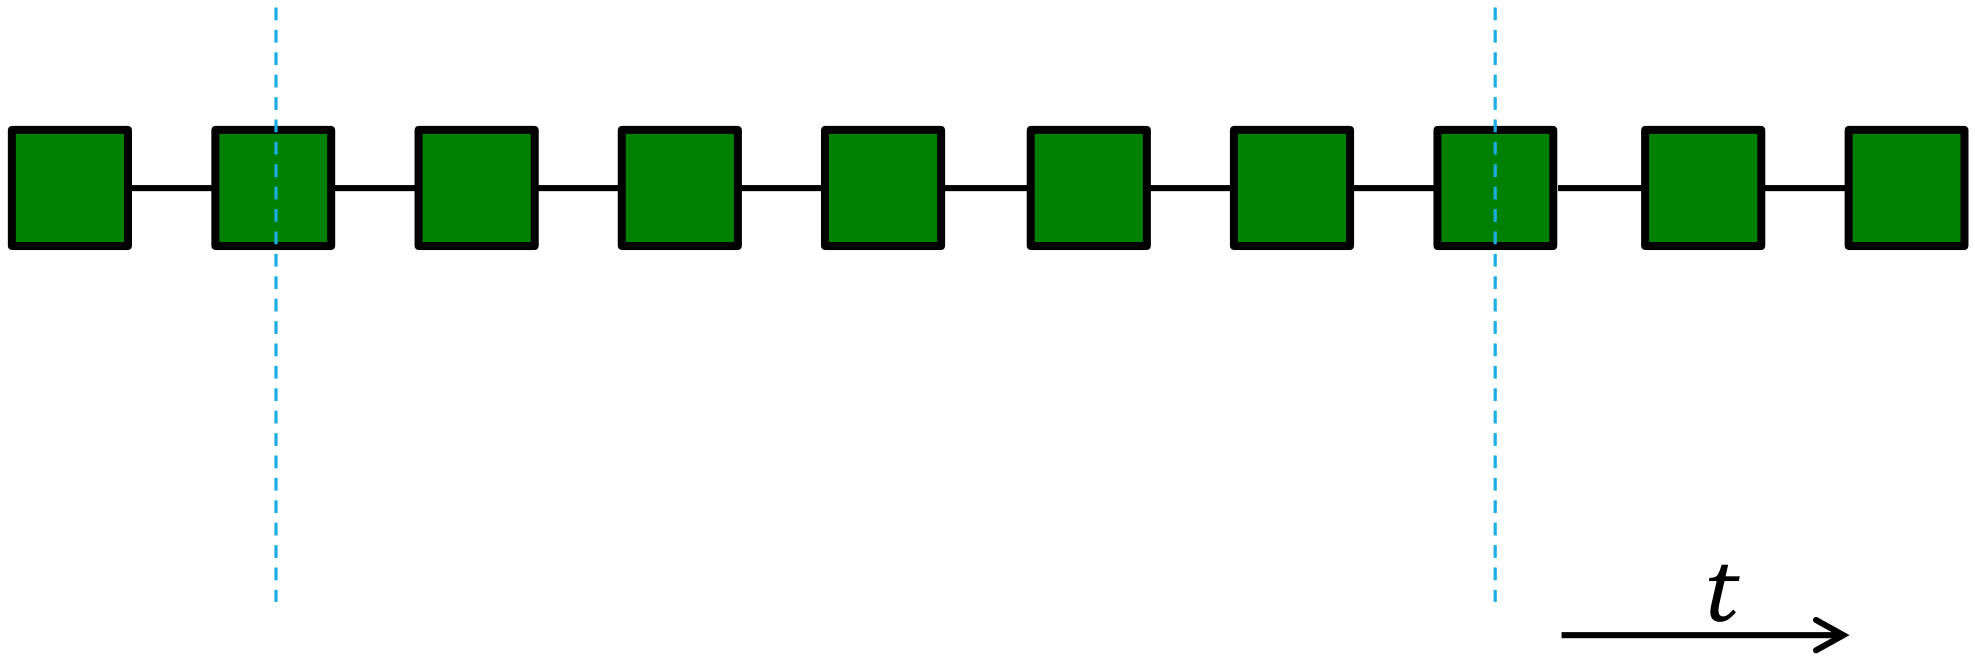

SECURITY THREAT!

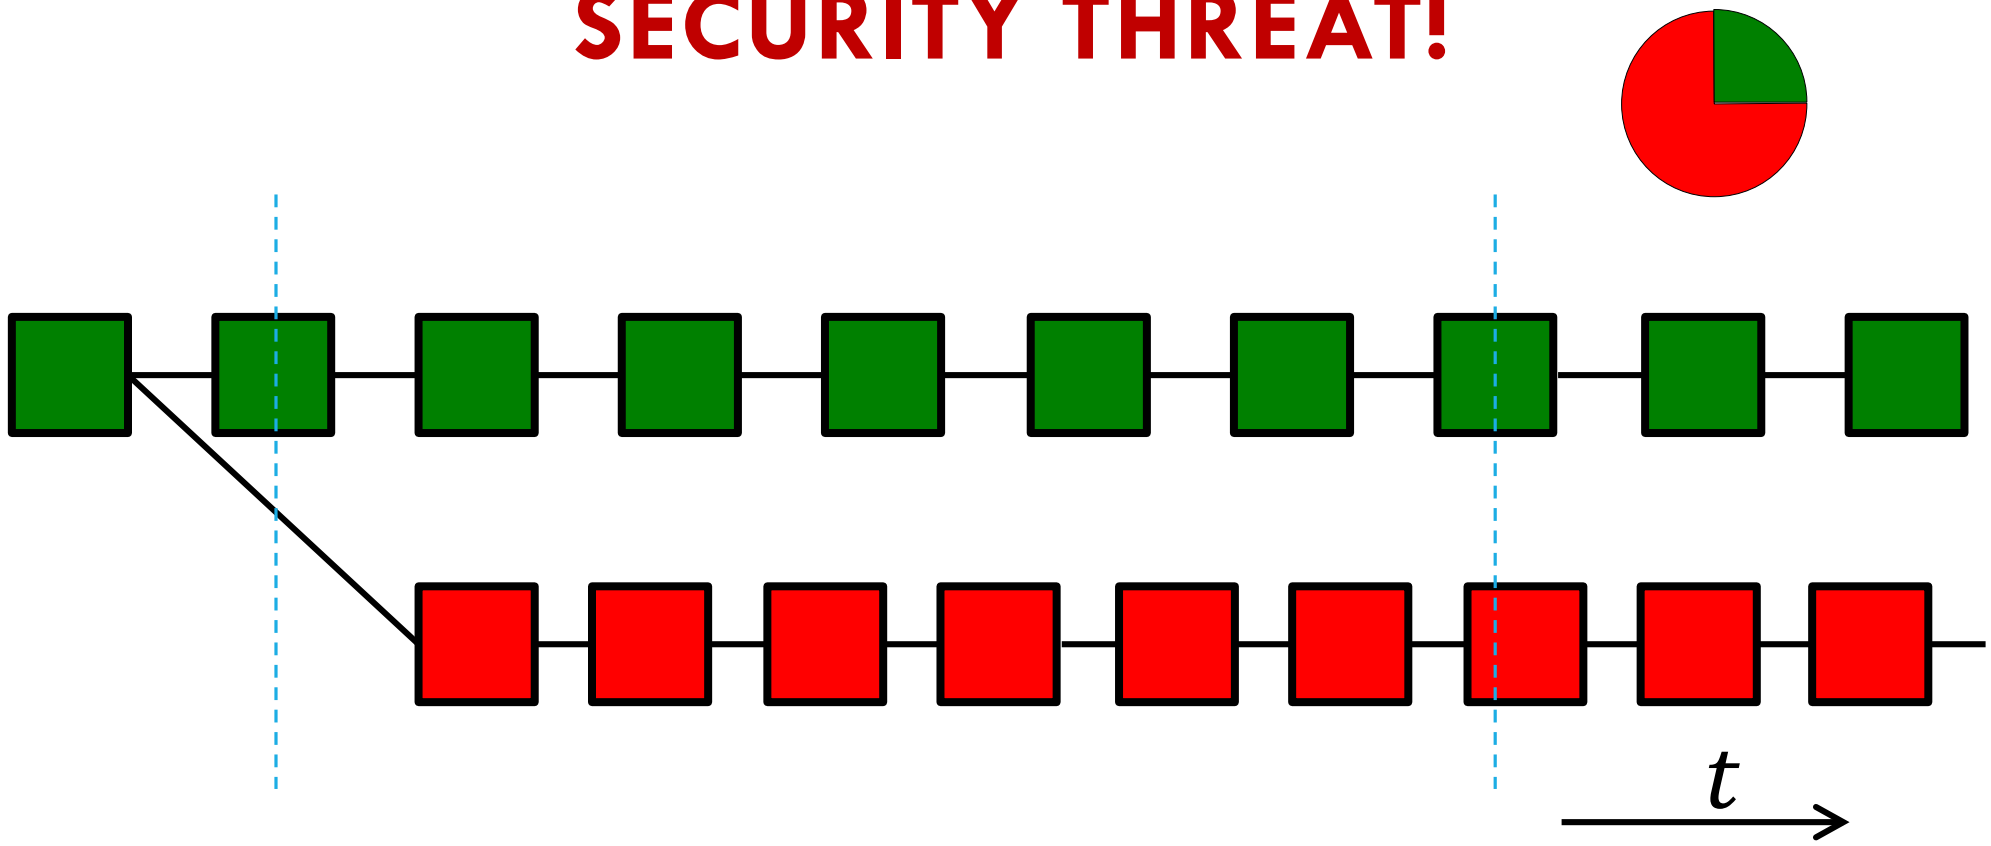

Threat: attacker outruns good miners

SECURITY THREAT!

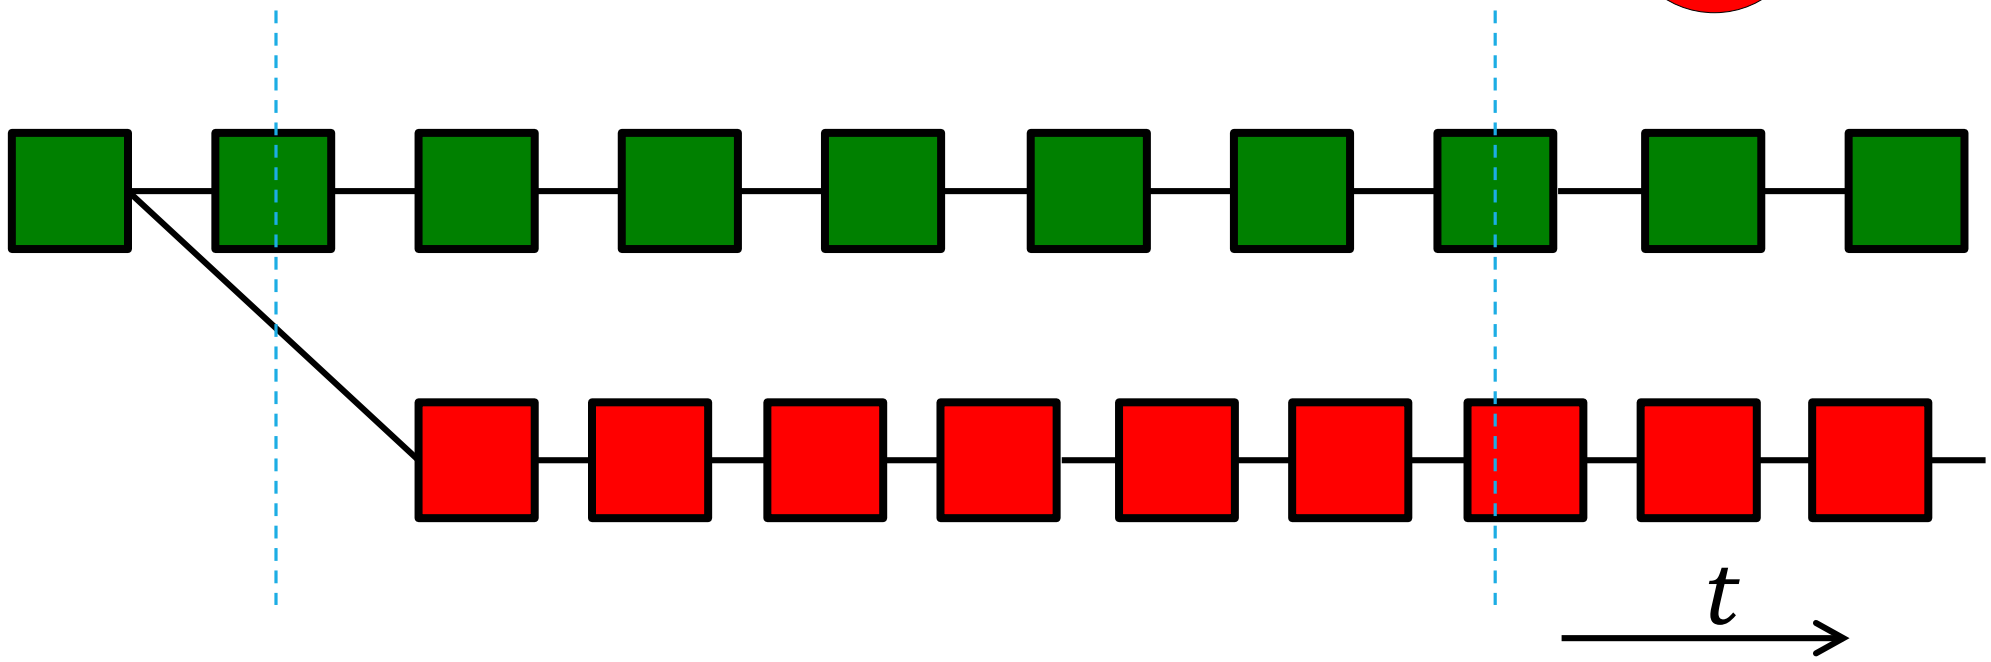

Threat: attacker outruns good miners

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

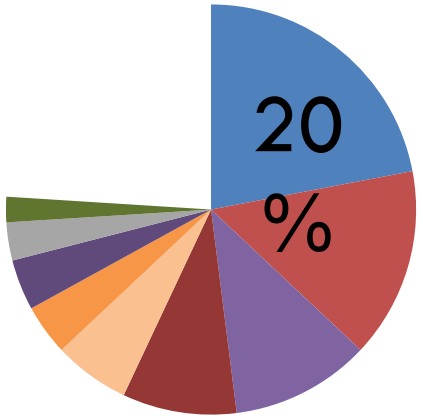

[blockchain.info, April 2015]

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*

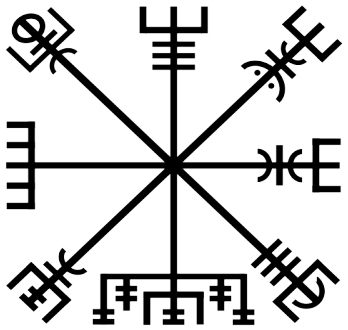

Vegvisir: tolerate branches

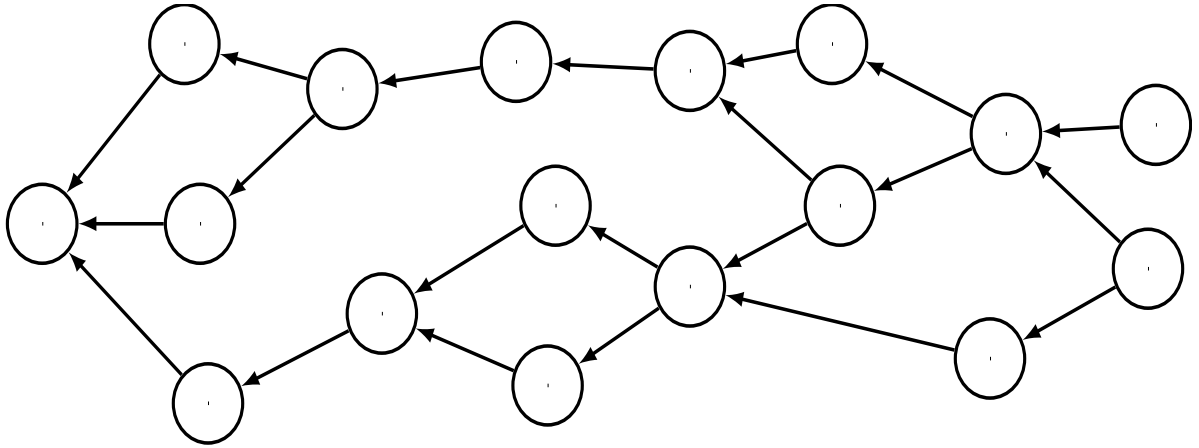

- The key innovation is to allow branching as a feature.
- Leads to DAG structure instead of linear blockchain
- Still maintains full causal history of events (respect's Lamport's $\rightarrow$)

Proof-of-Witness to persist blocks securely

No more than *k* malicious nodes in any neighborhood

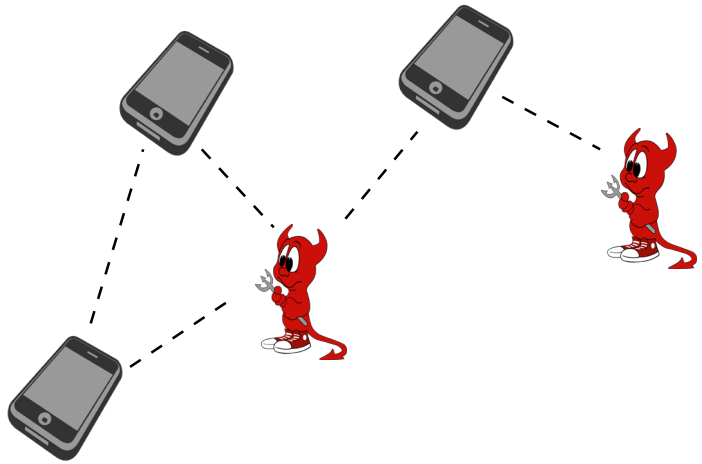

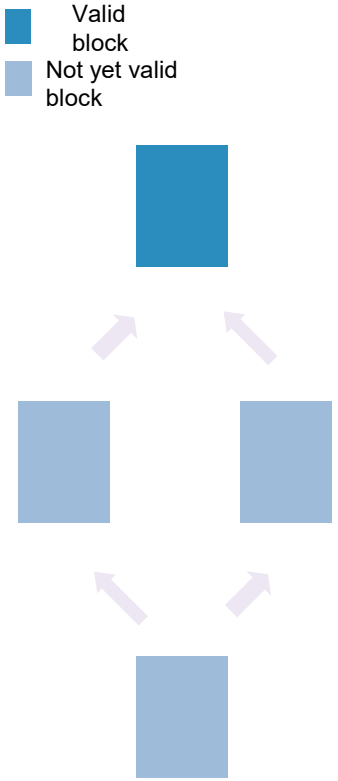

At least one copy of a valid block will survive if <*k* malicious peers

The Support Blockchain reduces sensor storage needs

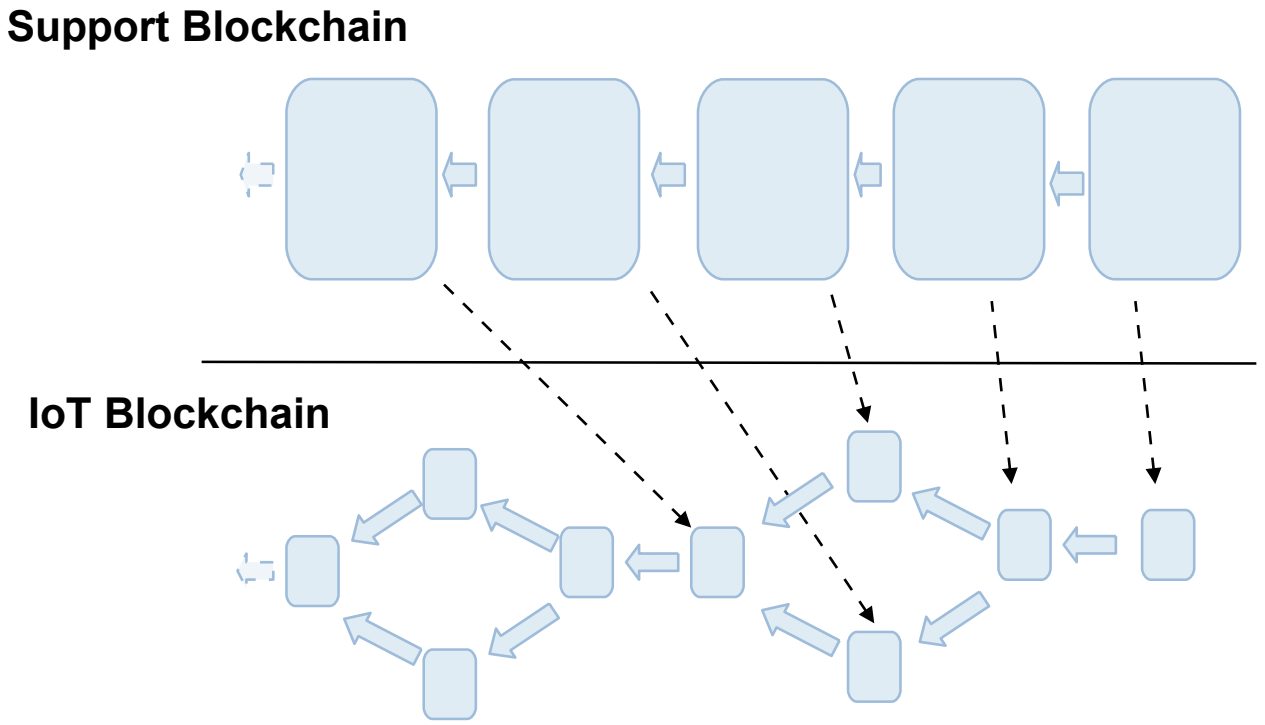

Allows regular peers to discard old blocks when storage space is low

Blocks are *gossiped* over ad hoc network

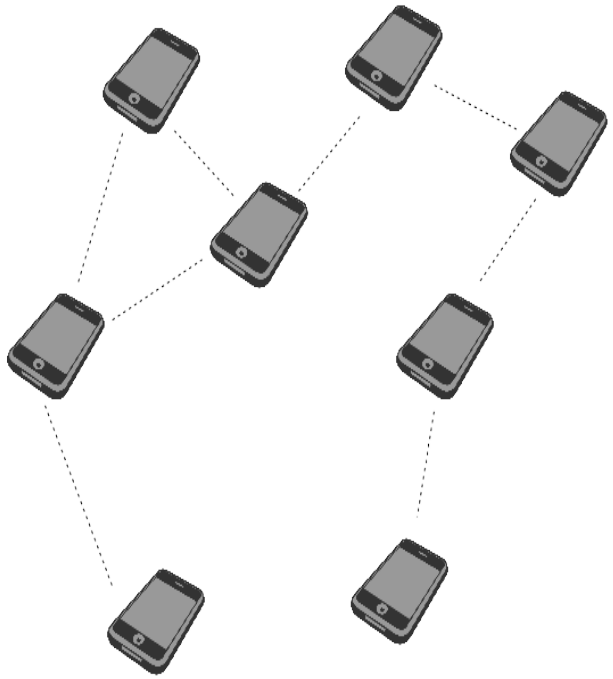

Heterogeneous, opportunistic networking

CRDTs for strong semantics in partitioned world

- *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

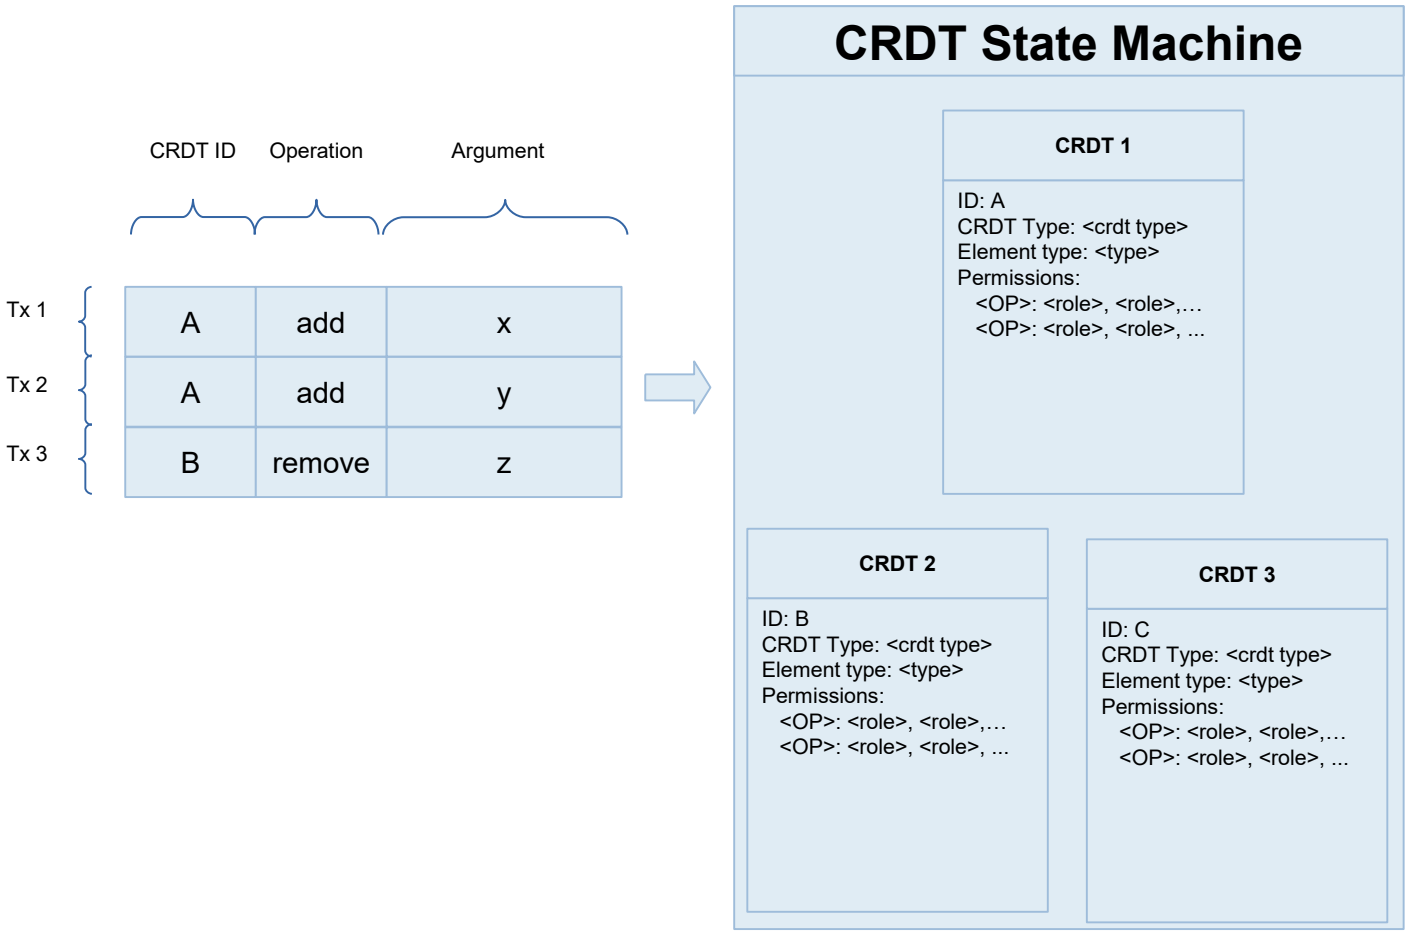

BUT SOME QUESTIONS REMAIN OPEN

How would Vegsivir handle "double spending" (same coin spent in both branches), or other kinds of semantic conflicts that might not involve coins?

- The actual meaning of the operation changes, or it becomes invalid.
- This could cascade to impact subsequent operations, too.
- We can't simply merge the chains and walk away…

Also, although smart farms have many sensors, Vegvisir lacks an answer to the issue of trust: we need the IoT hub to log enough information to know why we should trust a sensor, but this topic is out of scope for the paper.

DEMONSTRATION VIDEO

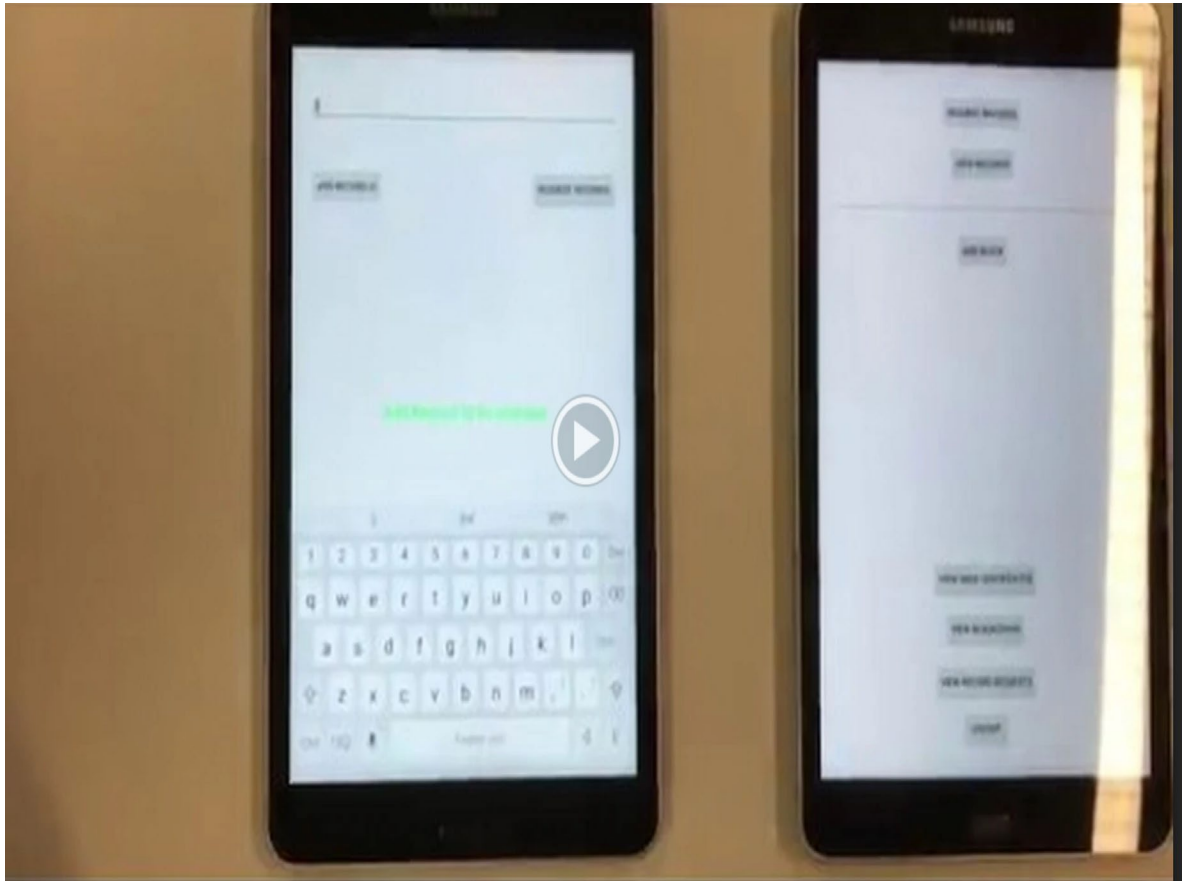

Proof of concept system

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