# Incentives behind Consensus in Distributed Ledgers

Davide Grossi

Institute of Artificial Intelligence

university of groningen

www.ankemarijedamdesign.nl

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

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

## 6. Incentive

By convention, the first transaction in a block is a special transaction that starts a new coin owned by the creator of the block. This adds an incentive for nodes to support the network, and provides a way to initially distribute coins into circulation, since there is no central authority to issue them. The steady addition of a constant of amount of new coins is analogous to gold miners expending resources to add gold to circulation. In our case, it is CPU time and electricity that is expended.

The incentive can also be funded with transaction fees. If the output value of a transaction is less than its input value, the difference is a transaction fee that is added to the incentive value of the block containing the transaction. Once a predetermined number of coins have entered circulation, the incentive can transition entirely to transaction fees and be completely inflation free.

The incentive may help encourage nodes to stay honest. If a greedy attacker is able to assemble more CPU power than all the honest nodes, he would have to choose between using it to defraud people by stealing back his payments, or using it to generate new coins. He ought to find it more profitable to play by the rules, such rules that favour him with more new coins than everyone else combined, than to undermine the system and the validity of his own wealth.

**Nick Szabo** @NickSzabo4
Incentives are misleading & dangerous basis for security unless objective math (computer science) done first & much engineerng margin added.

4:08 PM - 5 Jul 2017

university of groningen

# PART I

## Coordination & Common Knowledge

*(or:* ***What are DLs actually for?****)*

*[convention or agreement] that is, a sense of interest, suppos'd to be common to all, and where every single act is perform'd* ***in expectation that others are to perform the like****. Without such a convention, no one wou'd ever […] have been induc'd to conform his actions to it*

D. Hume, Treatise of Human Nature, 1738-40

university of groningen

# Coordinated Attacks & Empty Ledgers

- ☐ General 1 sends "Attack at dawn" to General 2
- ☐ Message arrives. Would General 2 attack?
- ☐ General 2 sends "Acknowledged" to General 1.
- ☐ Message arrives. Would General 1 attack?

| | Attack | Wait |
|---|---|---|
| Attack | Win | Lose |
| Wait | Lose | Wait |

A. Rubinstein. The Electronic Mail Game: Strategic Behavior under 'Almost CK'. AER, 1989

J. Halpern. Reasoning About Knowledge: an Overview. 1986

university of groningen

# Type of coordination enabled in DLs

Each honest agent knows that within a bound, from that point on …

Each honest agent knows that within a bound, from that point on …

Each honest agent knows that within a bound, from that point on …

…

T is a prefix in all honest nodes' ledgers

*If* $i$ is honest and $T$ is a prefix of $i$'s ledger *then* there is $CK$ among all honest nodes that in $\Delta$ steps $T$ will be a prefix of all honest nodes' ledgers

- Nakamoto consensus (under some assumptions) suffices to ensure coordination within a given time window

… *based on cryptographic proof instead of trust, allowing any two willing parties to transact directly with each other without the need for a trusted third party*

S. Nakamoto, Bitcoin: A Peer-to-Peer Electronic Cash System 2008

proof-of-work

fault-tolerant replication

J. Halpern, R. Pass. A Knowledge-Based Analysis of the Blockchain Protocol. TARK'17, 2017

university of groningen

# PART II

Why would I **build** a block? And why would I **check** it?

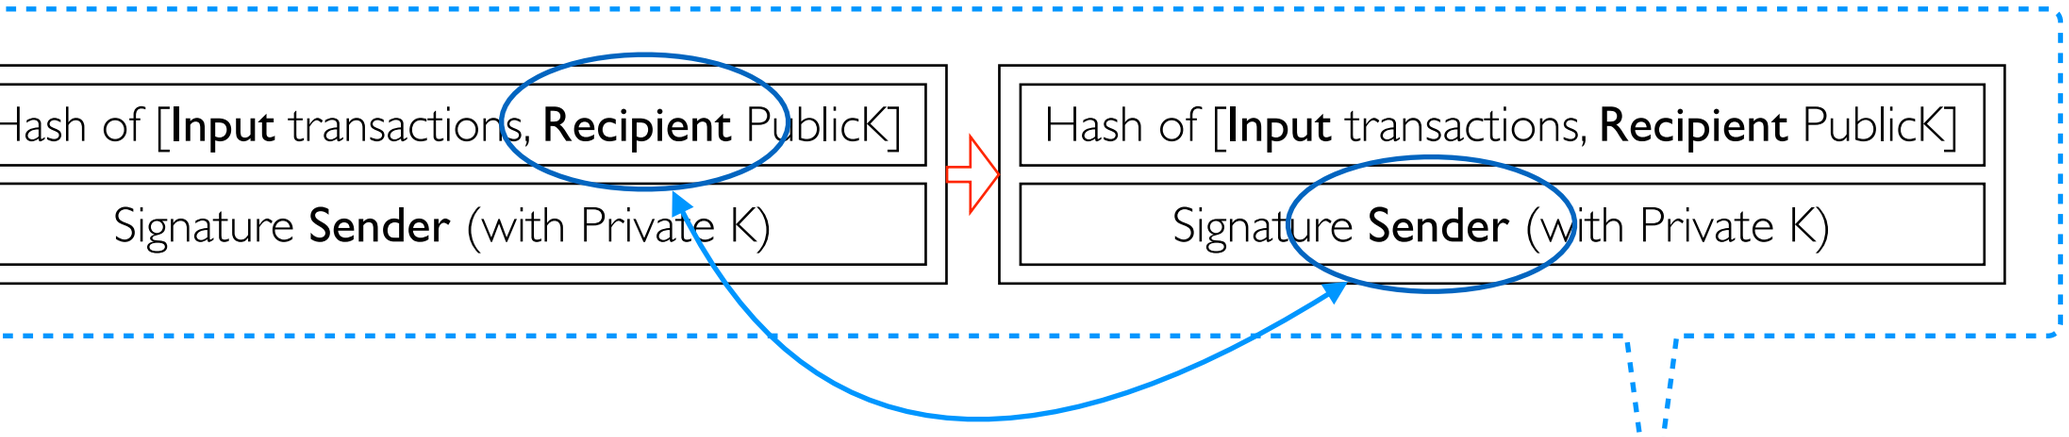


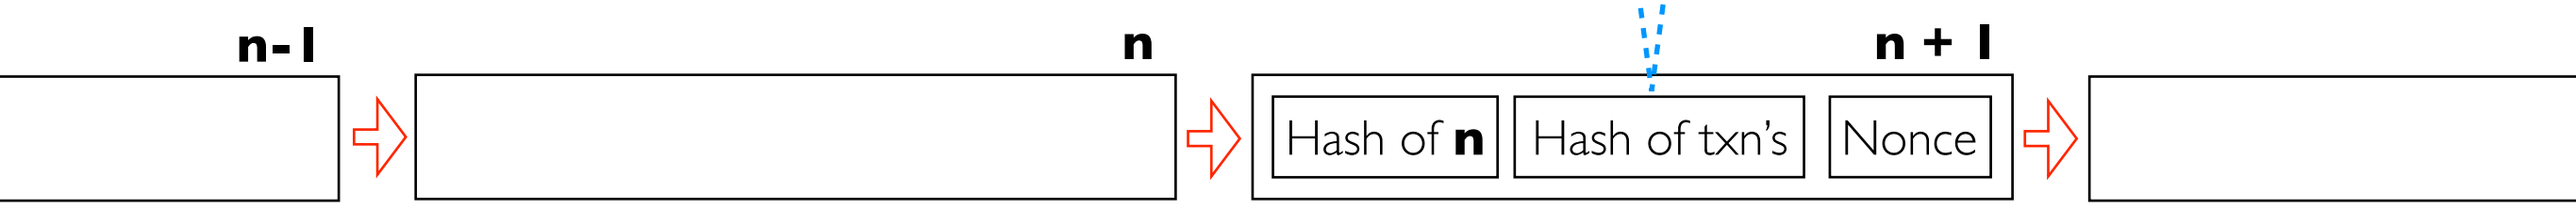


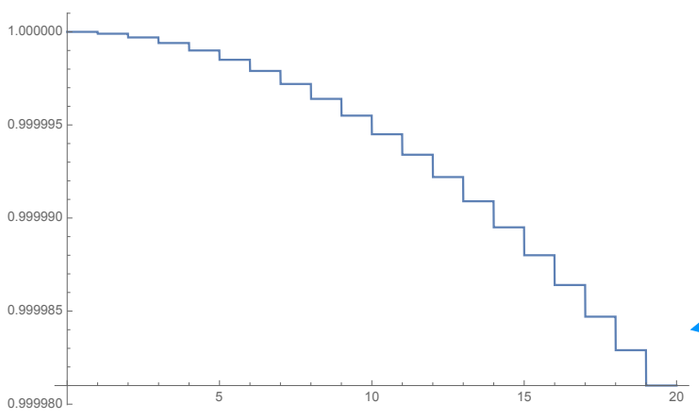


Is the Hash of **n + 1** < **V** ?

**Yes**

**No** Try again!

university of groningen

# Why mining?

- All pay - one wins
- R&D race
- NE exists and is unique

$$u_i(\mathbf{h}) = (R - c_i h_i) \cdot \frac{h_i}{\sum_{j \in N} h_j} - c_i h_i \cdot \frac{h_{-i}}{\sum_{j \in N} h_j}$$

Investments in hashing power

Reward for solving puzzle

i's hashing power

i's cost of hashing

probability that i fails solving the puzzle first

probability that i solves the puzzle first


J. Ma, J. Gans, R. Tourky. Market Structure in Bitcoin Mining. NBER Working Paper, 2018

N. Dimitri. Bitcoin Mining as a Contest. Ledger, 2017


university of groningen

# Why Verifying?

- In Bitcoin verification work is negligible compared to mining, but that's not the case in general (see Ethereum)
- Miners are aware that non-valid transactions have the potential to decrease Bitcoin's value
- But this is ultimately a **public good** game and there is potential for 'tragedy of the commons' scenario

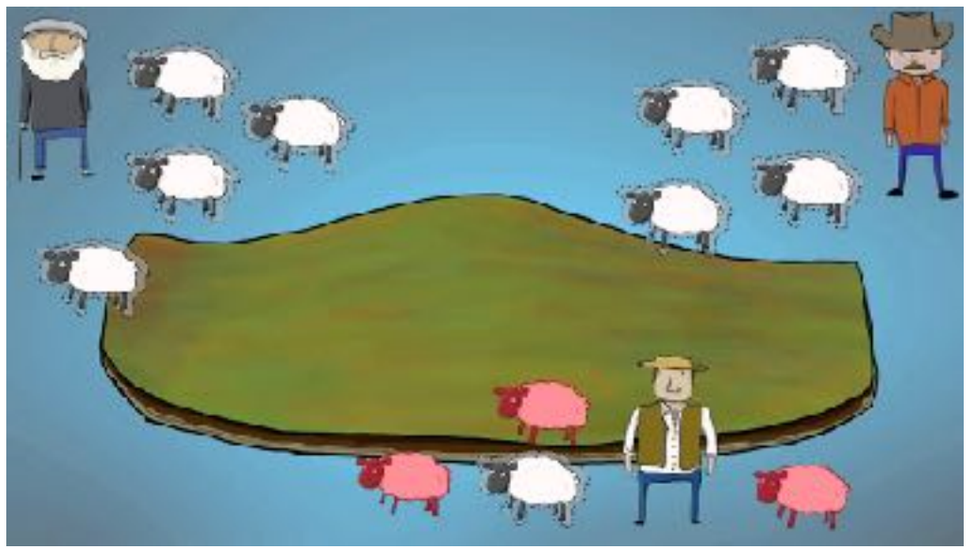

L. Luu, J. Teusch, R. Kulkarni, P. Saxena. Demistifying Incentives in the Consensus Computer, CCS'15, 2015

university of groningen

# PART III

## Why should I **fork** (or rather not)?


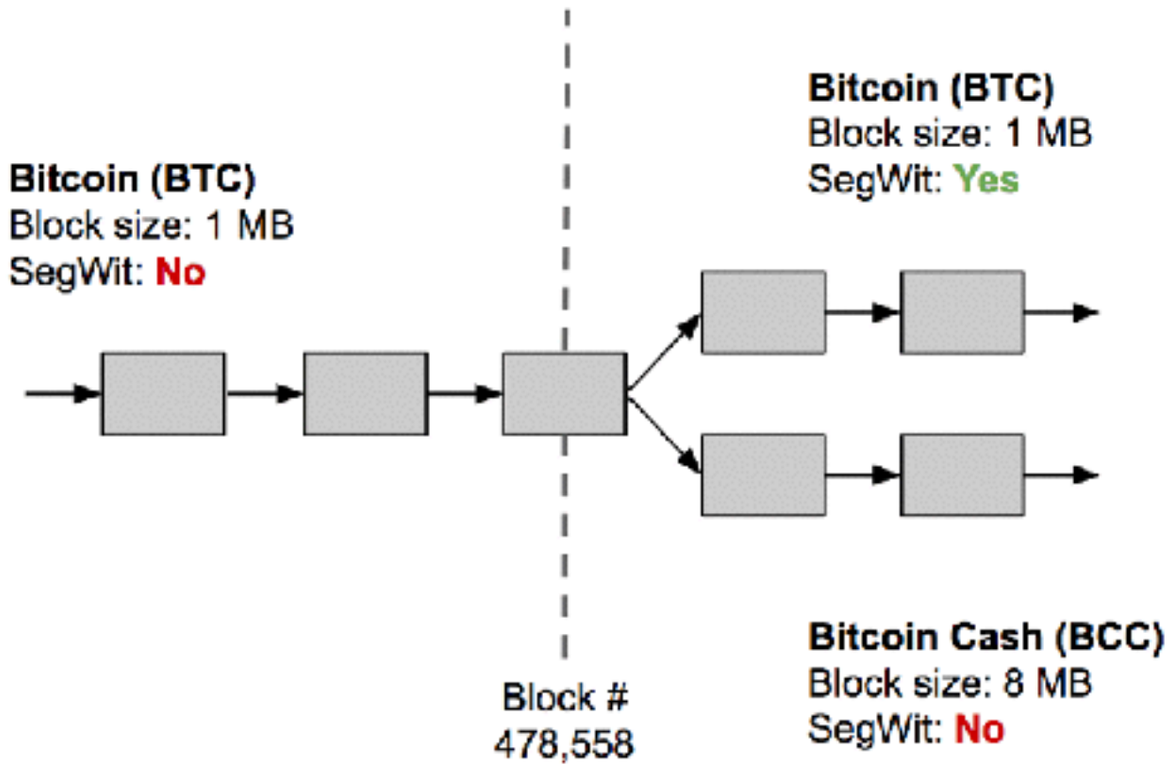


Other reasons:

*Signal delays*

*Double-spending attacks*

*Software upgrades (11 March 2013)*

university of groningen

# Blockchain Folk-Theorem

*Nakamoto Consensus rules out the occurrence of forks*

- ☐ True, at certain levels of abstraction
- ☐ But …

```
23:06  Luke Dashjr        so??? yay accidental hardfork? :x
23:06  Jouke Hofman       Holy crap


23:22  Gavin Andresen     the 0.8 fork is longer, yes? So majority hashpower is 0.8....
23:22  Luke Dashjr        Gavin Andresen: but 0.8 fork is not compatible earlier will be accepted by
all versions

23:23  Gavin Andresen     first rule of bitcoin: majority hashpower wins
23:23  Luke Dashjr        if we go with 0.8, we are hardforking


23:24  Luke Dashjr        so it's either 1) lose 6 blocks, or 2) hardfork for no benefit
23:25  BTC Guild      We'll lose more than 6

23:43  BTC Guild      I can single handedly put 0.7 back to the majority hash power I just need
confirmation

23:44  Pieter Wuille      BTC Guild: imho, that is was you should do, but we should have consensus
first
```


A. Narayanan. Analysing the 2013 Bitcoin Fork: Centralized Decision Making Saved the Day, 2015

A. Miller, J. LaViola. Anonymous Byzantine Consensus from Moderately-Hard Puzzles: A Model for Bitcoin, 2014

B. Biais, C. Bisiere, M. Bouvard, C. Casamatta. The Blockchain Folk Theorem. TSE Working Papers, 17-187, 2018


# Blockchain Folk-Theorem

*Nakamoto Consensus rules out the occurrence of forks*

- With no centralised solution:
  - Gradual consensus towards 0.8 branch (vs 0.7)
  - Coordination on which branch to mine harder/slower
  - Double spending attacks more possible
  - Fork would survive longer (than 8hrs), likely because of vested interest of miners on 0.7 fork

**Keynes' Beauty Contest**

**Shubik's dollar auction**


A. Narayanan. Analysing the 2013 Bitcoin Fork: Centralized Decision Making Saved the Day, 2015

A. Miller, J. LaViola. Anonymous Byzantine Consensus from Moderately-Hard Puzzles: A Model for Bitcoin, 2014

B. Biais, C. Bisiere, M. Bouvard, C. Casamatta. The Blockchain Folk Theorem. TSE Working Papers, 17-187, 2018


university of groningen

# PART IV

## Why not just voting?

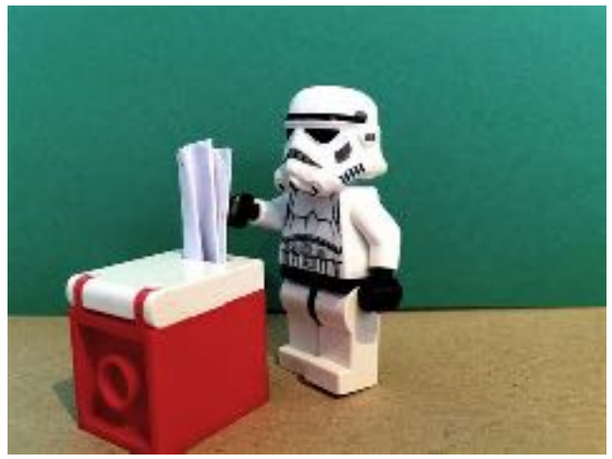

… *If the majority were based on one-IP-address-onve-vote, it could be subverted by anyone able to allocate many IPs.*

S. Nakamoto, Bitcoin: A Peer-to-Peer Electronic Cash System 2008

university of groningen

# The Byzantine Generals Problem

LESLIE LAMPORT, ROBERT SHOSTAK, and MARSHALL PEASE
SRI International

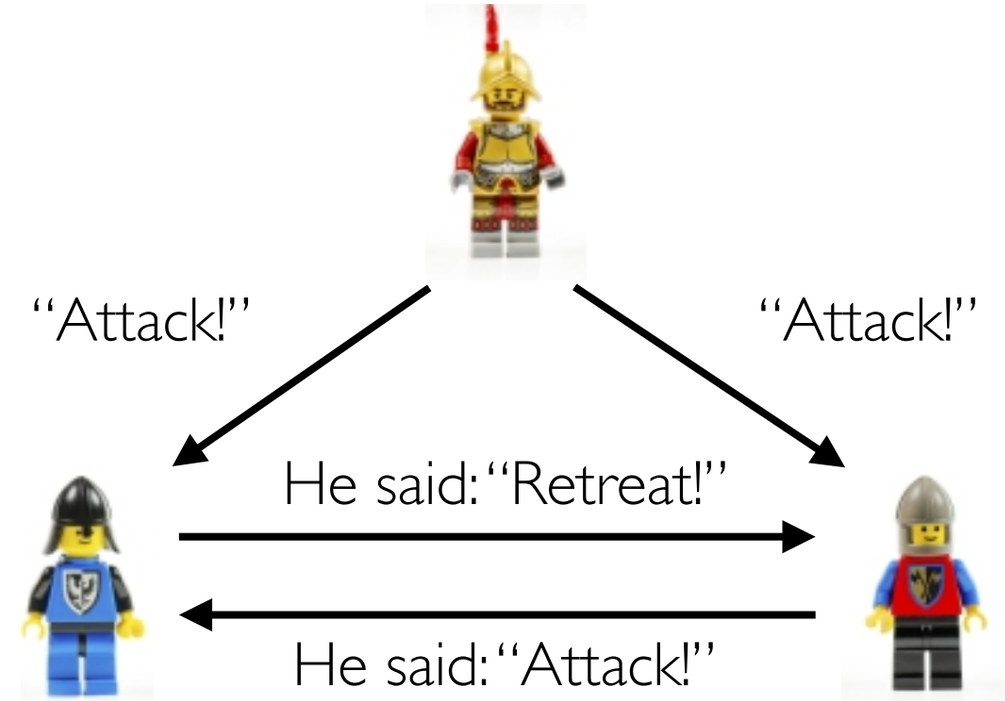


- *If the general is loyal*, then every *loyal lieutenant* obey the same order
- Solvable with private messages if: **|Loyals| > 3|Non-Loyals|**
- **BUT**: Centralized identity management (prevents Sybil Attacks)

university of groningen

# Distributed Fault-Tolerant Replication

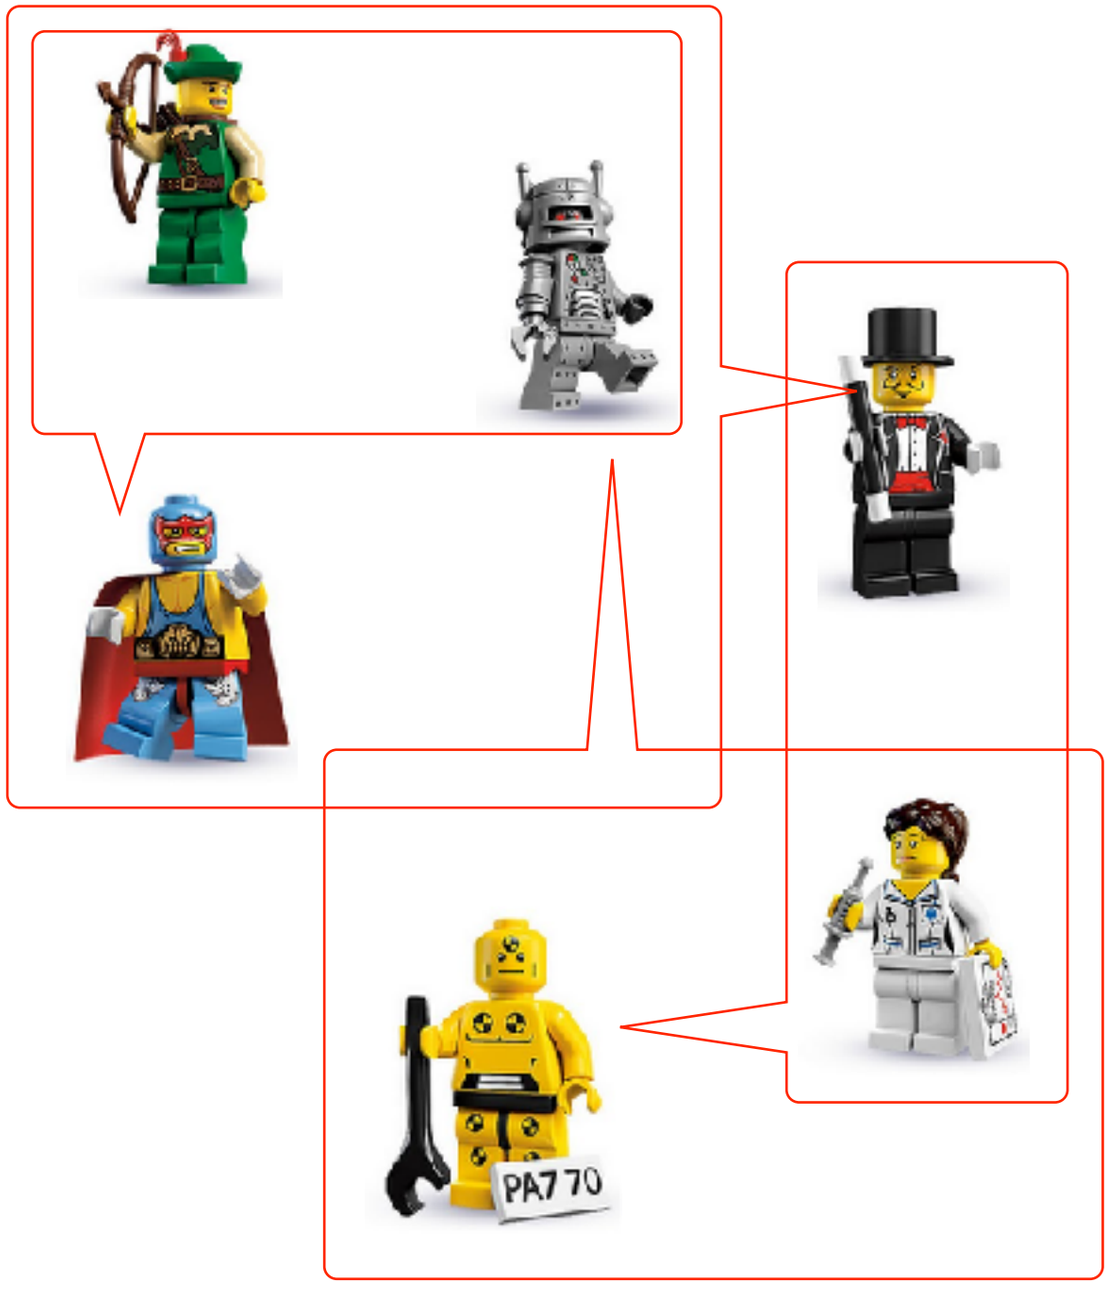


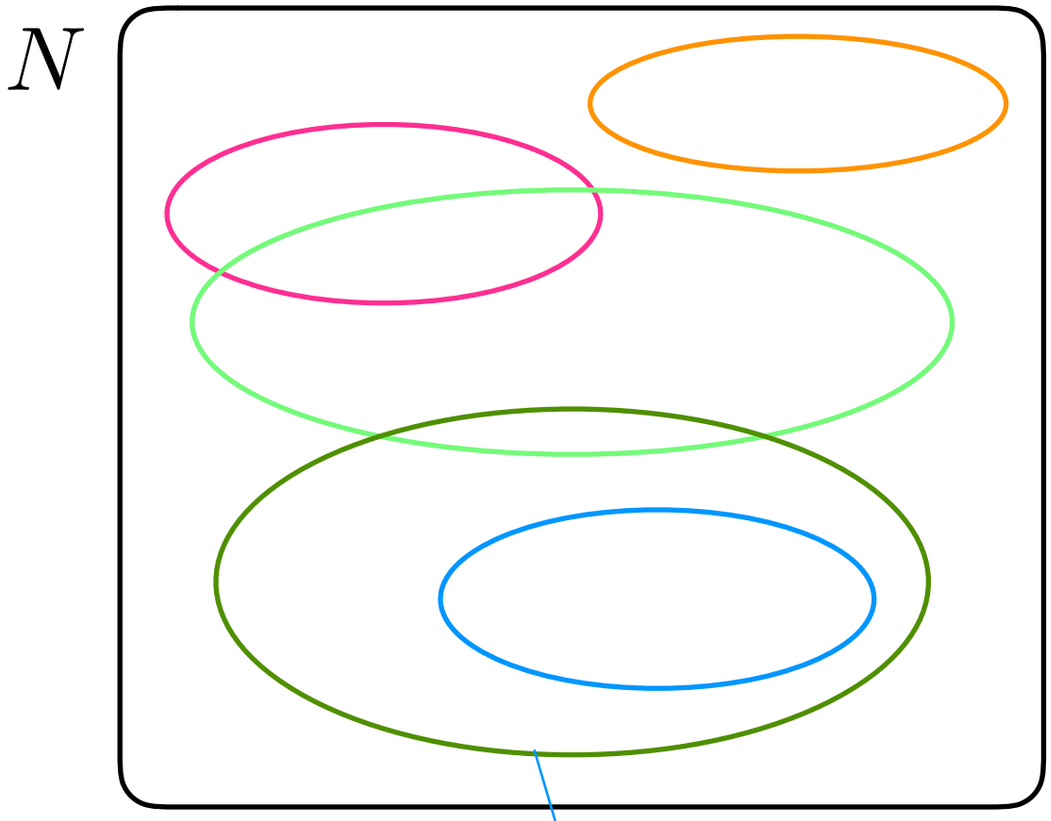


Sets of nodes which, once they agree on a value, they stabilise on that value (*Simple Games*)

No mining! Cryptocurrency has **no role** in consensus

university of groningen

# Do we stabilize on one value or not?

Simple game induced by $\mathcal{C}$

Largest consensus forms

Veto nodes
(they belong to all coalitions)

**Proposition** Let $\mathcal{C}$ be a coalitional system:

- $\mathbf{Q}^{\mathcal{C}}$ has a non-empty core if and only if $\mathcal{V}^{\mathcal{C}} \neq \emptyset$;
- if $\mathcal{V}^{\mathcal{C}} \neq \emptyset$, then the core consists of all imputations in which non-veto nodes get value 0.

- ☐ Realising optimal consensus requires nodes with veto power
- ☐ Which should be rewarded
- ☐ … but these should be few (centralization) for otherwise consensus would be hard (possibility of deadlocks)

D.B. Gillies. Some Theorems on n-Person Games. PhD thesis, Department of Mathematics, Princeton, 1959

university of groningen

# Conclusions

- The variety of incentive structures (games) behind consensus protocols for DL is extremely rich: many challenges!
- Some incentive structures still unclear (verification in PoW)
- Models of ***parallel*** game-playing?
- Do we have the right ***solution concepts*** for interaction in DLs? They should be validated by data


university of groningen


university of
groningen

# Stellar Consensus (*in a nutshell*)

- **N** nodes holds (binary) opinions on the value of a slot (in the ledger), and some may misbehave (Byzantine failure)
  - All 'good' nodes should be able to stabilize their opinion on one value of the slot (**liveness**)
  - No two 'good' nodes should stabilize on opposite opinions (**safety**)
- Stellar exploits the notion of **trust**: *when all nodes I trust agree on a value, then I accept that value and stabilize on it*
- Who to trust is an individual choice
  - Trust should exhibit structural properties which make **safety** and **liveness** possible (Federated Byzantine Agreement Systems)
  - Nodes should be able to recognise agreement among trusted nodes (Federated Voting)

# Safety & Liveness in FBASs

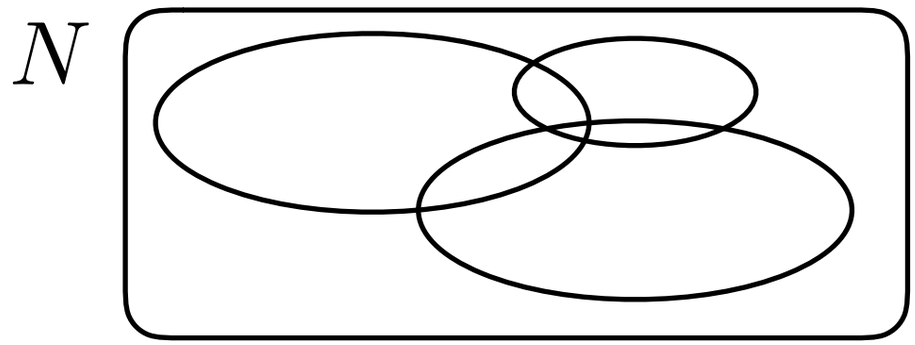


**Theorem** Let a coalitional system $\mathcal{C}$ be given. If $\mathcal{C}$ satisfies quorum intersection, then for any coalition $X \subseteq N$ (of ill-behaved nodes), the set of nodes that are befouled (by $X$) is a dispensible coalition.

X plus the 'good' nodes whose ability for correct agreement depends on X

A set of nodes that can be eliminated without jeopardising safety and liveness of the remaining nodes

university of groningen

# Is Quorum Intersection feasible?



systems thanks to the duplicity of the ill-behaved nodes. In short, FBAS $\langle \mathbf{V}, \mathbf{Q} \rangle$ can survive Byzantine failure by a set of nodes $B \subseteq \mathbf{V}$ iff $\langle \mathbf{V}, \mathbf{Q} \rangle$ enjoys quorum intersection after deleting the nodes in $B$ from $\mathbf{V}$ and from all slices in $\mathbf{Q}$. More formally:

*Definition* (*delete*). If $\langle \mathbf{V}, \mathbf{Q} \rangle$ is an FBAS and $B \subseteq \mathbf{V}$ is a set of nodes, then to *delete* $B$ from $\langle \mathbf{V}, \mathbf{Q} \rangle$, written $\langle \mathbf{V}, \mathbf{Q} \rangle^B$, means to compute the modified FBAS $\langle \mathbf{V} \setminus B, \mathbf{Q}^B \rangle$ where $\mathbf{Q}^B(v) = \{ q \setminus B \mid q \in \mathbf{Q}(v) \}$.

It is the responsibility of each node $v$ to ensure $\mathbf{Q}(v)$ does not violate quorum intersection. One way to do so is to pick conservative slices that lead to large quorums. Of course, a malicious $v$ may intentionally pick $\mathbf{Q}(v)$ to violate quorum intersection. But a malicious $v$ can also lie about the value of $\mathbf{Q}(v)$ or ignore $\mathbf{Q}(v)$ to make arbitrary as-

**Proposition** QUORUM-INTERSECTION is co-NP-complete.

- It is implausible to leave the responsibility to individual nodes to guarantee **QI**
- So how is **QI** enforced in Stellar?

university of groningen