Algorithms for and against the Cloud

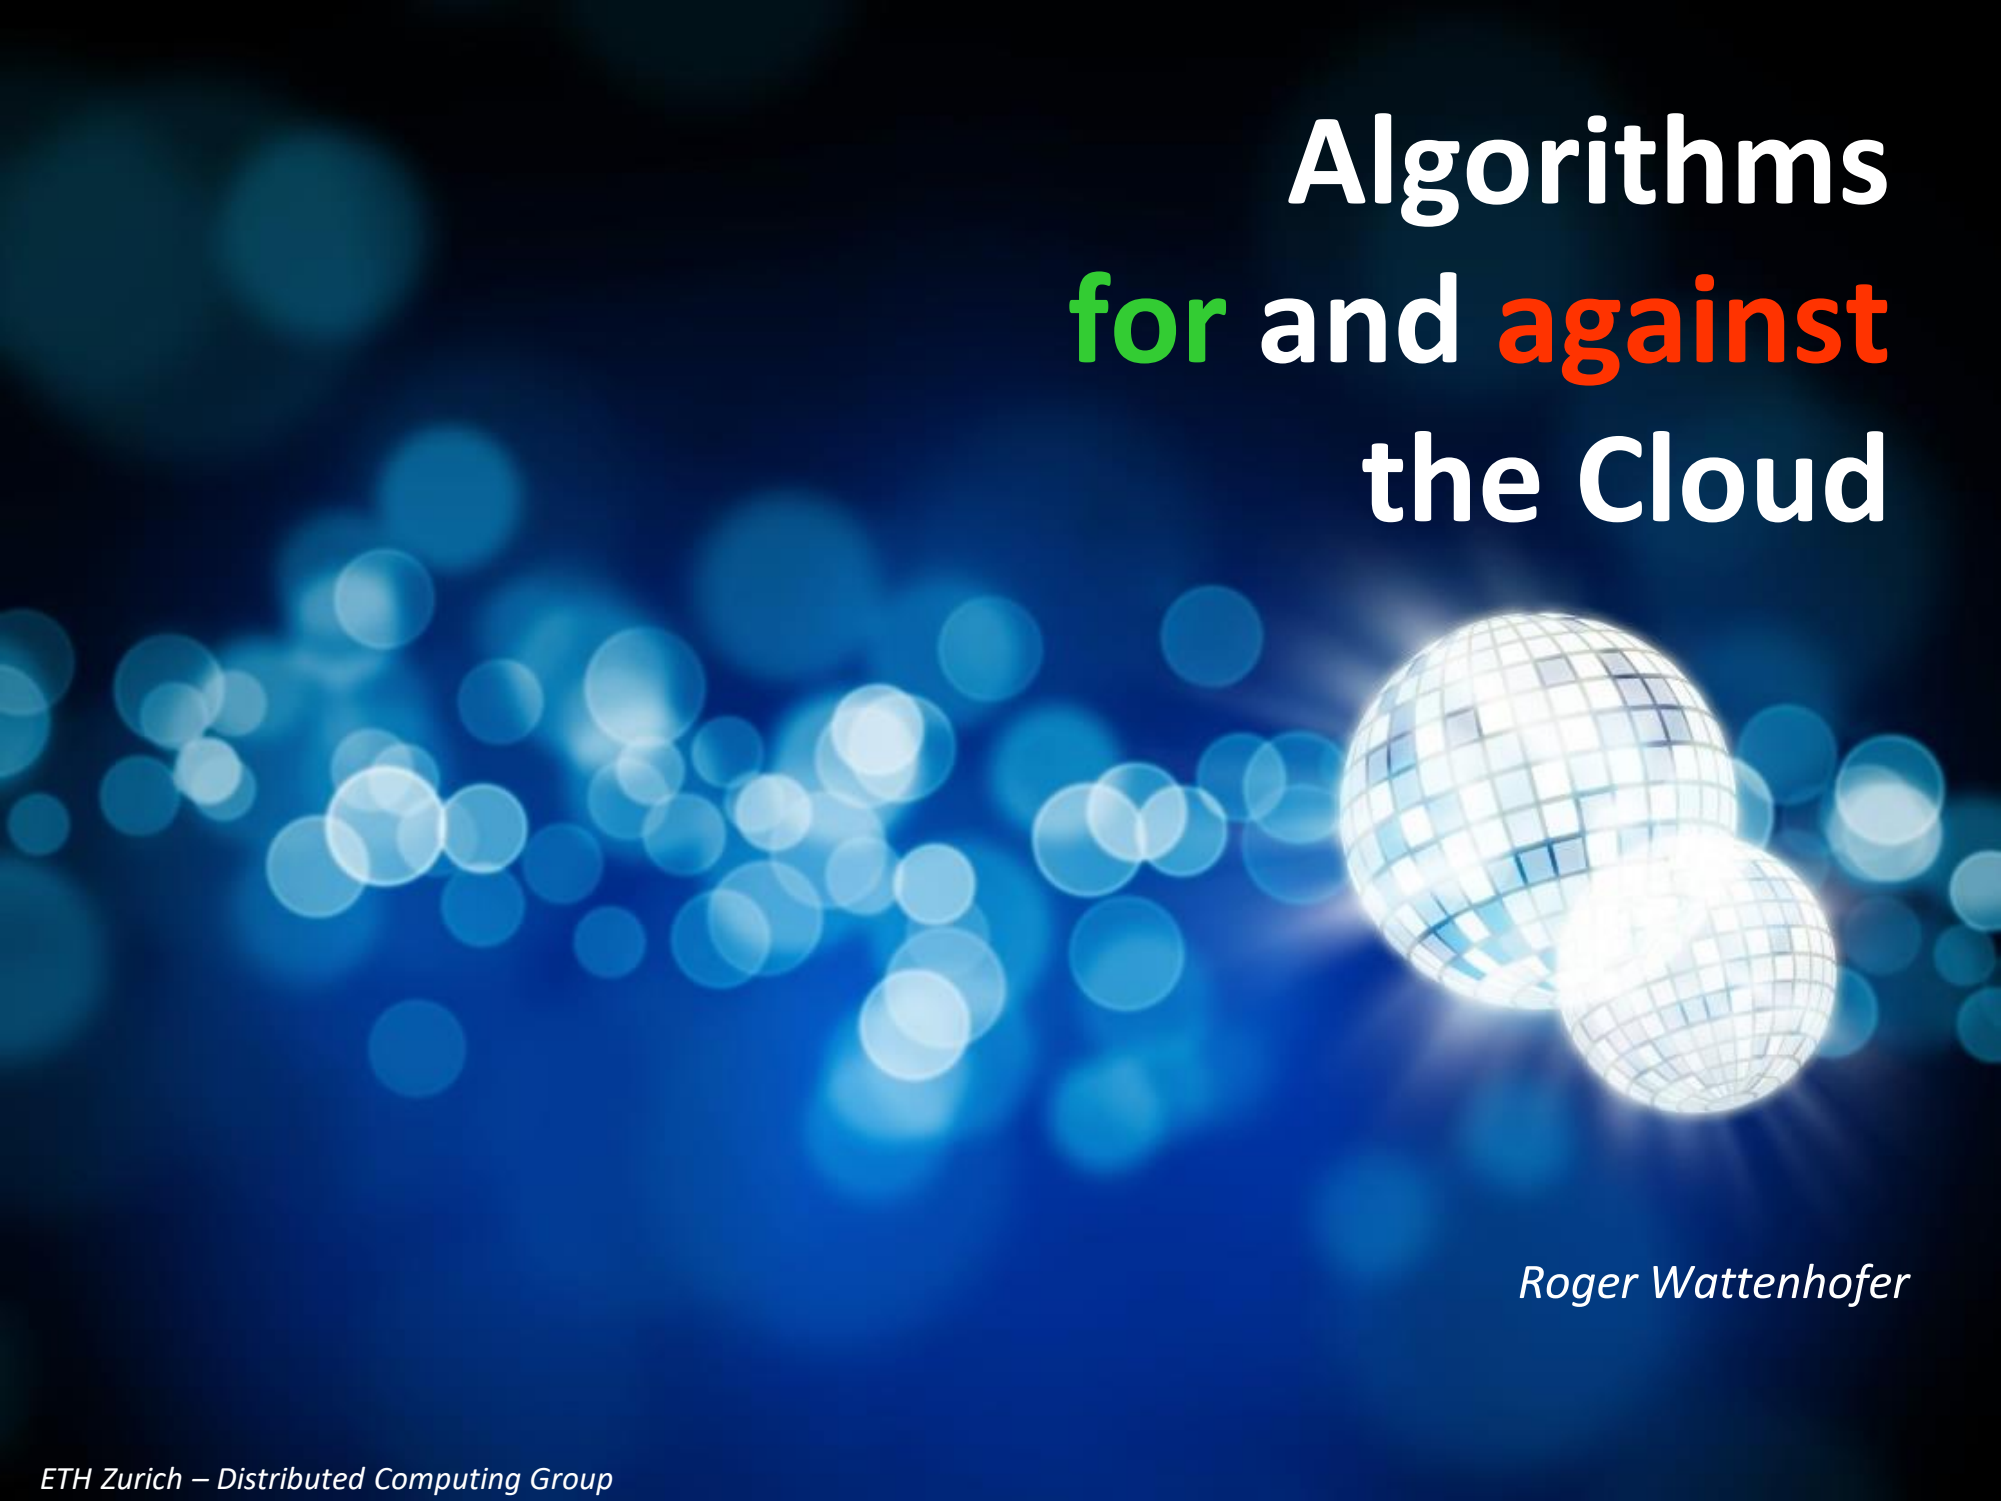

Roger Wattenhofer

ETH Zurich – Distributed Computing Group

Disclaimer


SenSys
OSDI
HotNets
AAAI
PODC
Mobicom
STOC
FOCS
ICALP
SIGCOMM
SPAA
SODA
EC

Algorithms for the Cloud

Algorithms for the Cloud

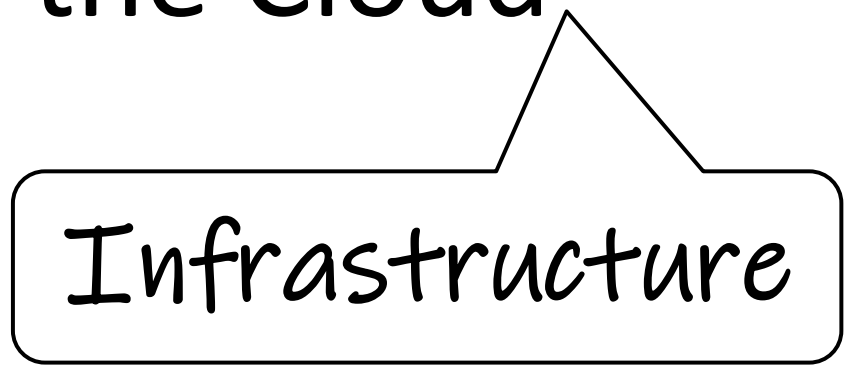

Algorithms for the Cloud

Algorithms for the Cloud

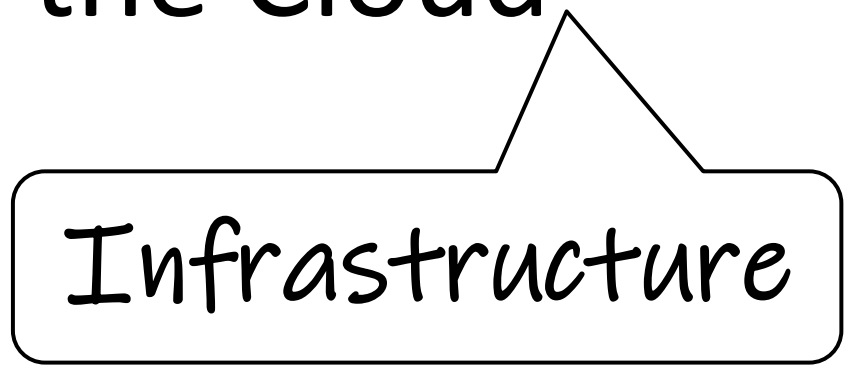

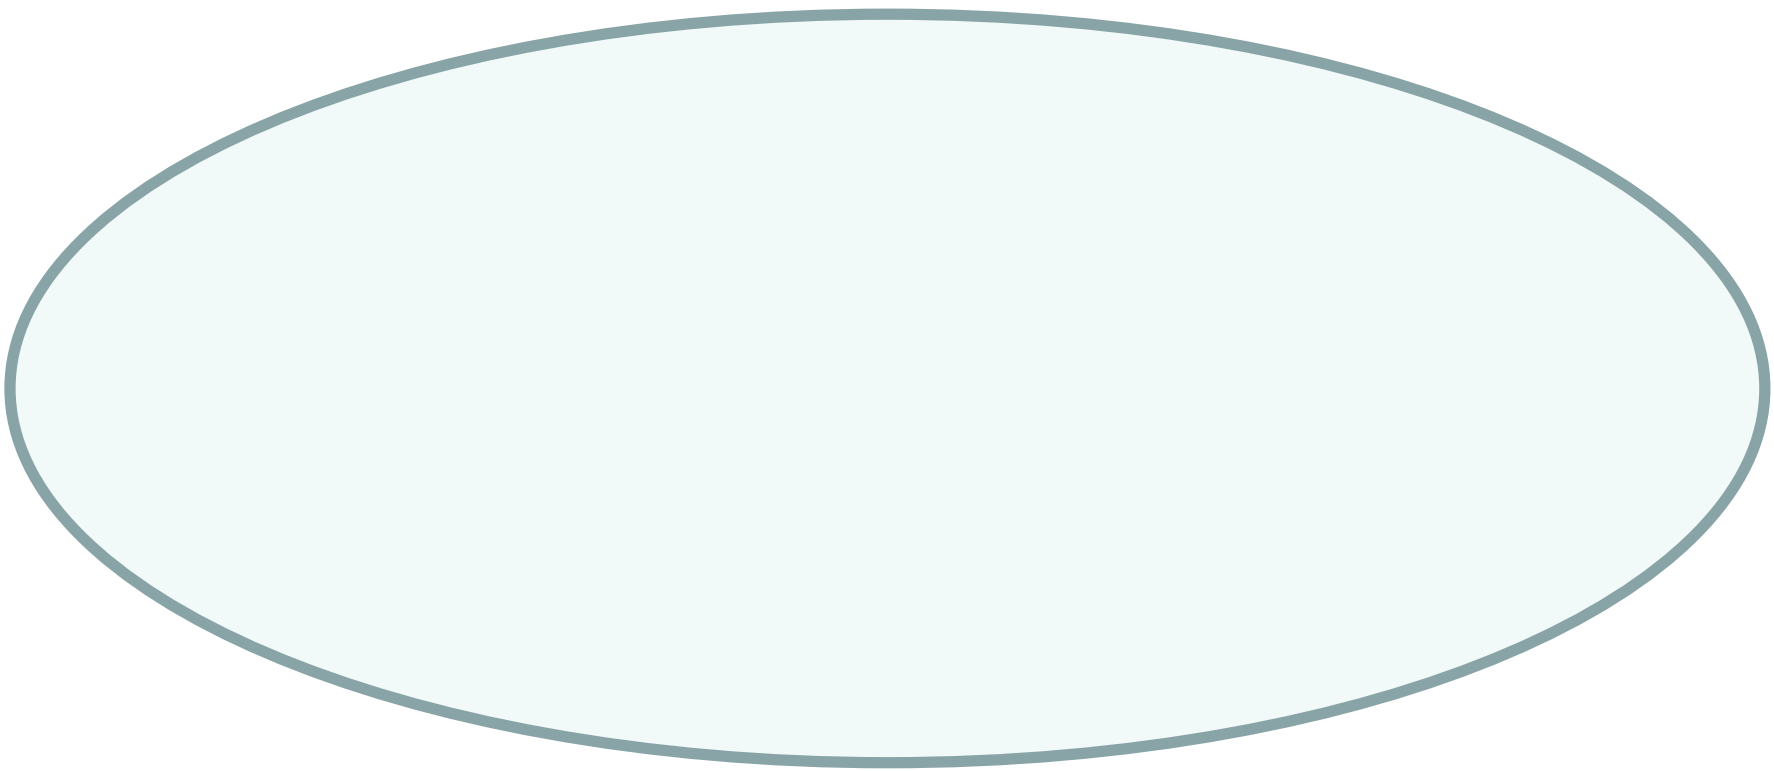

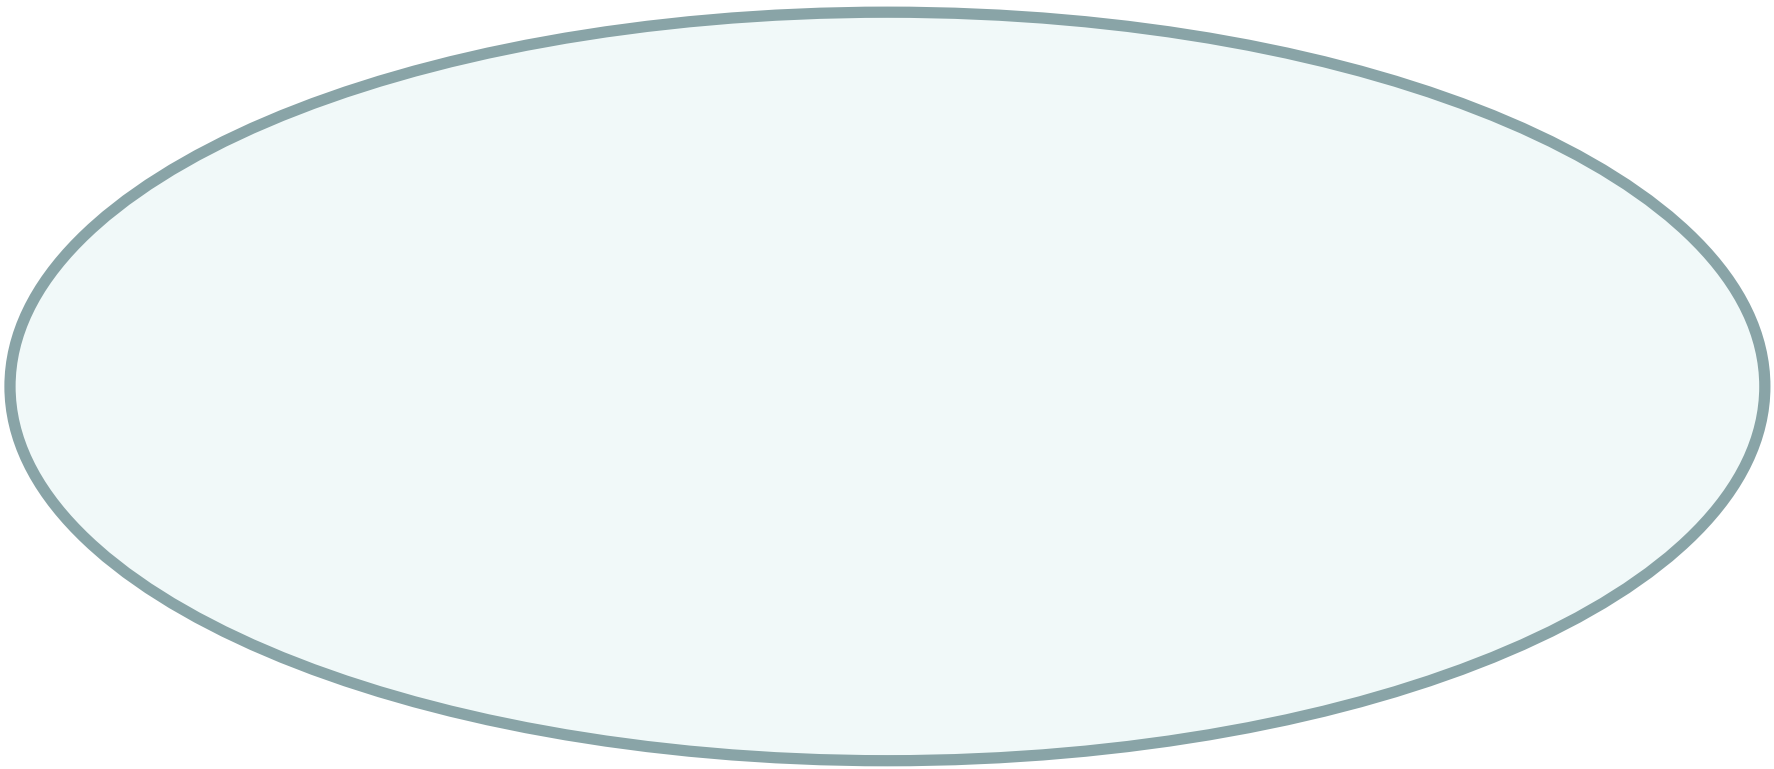

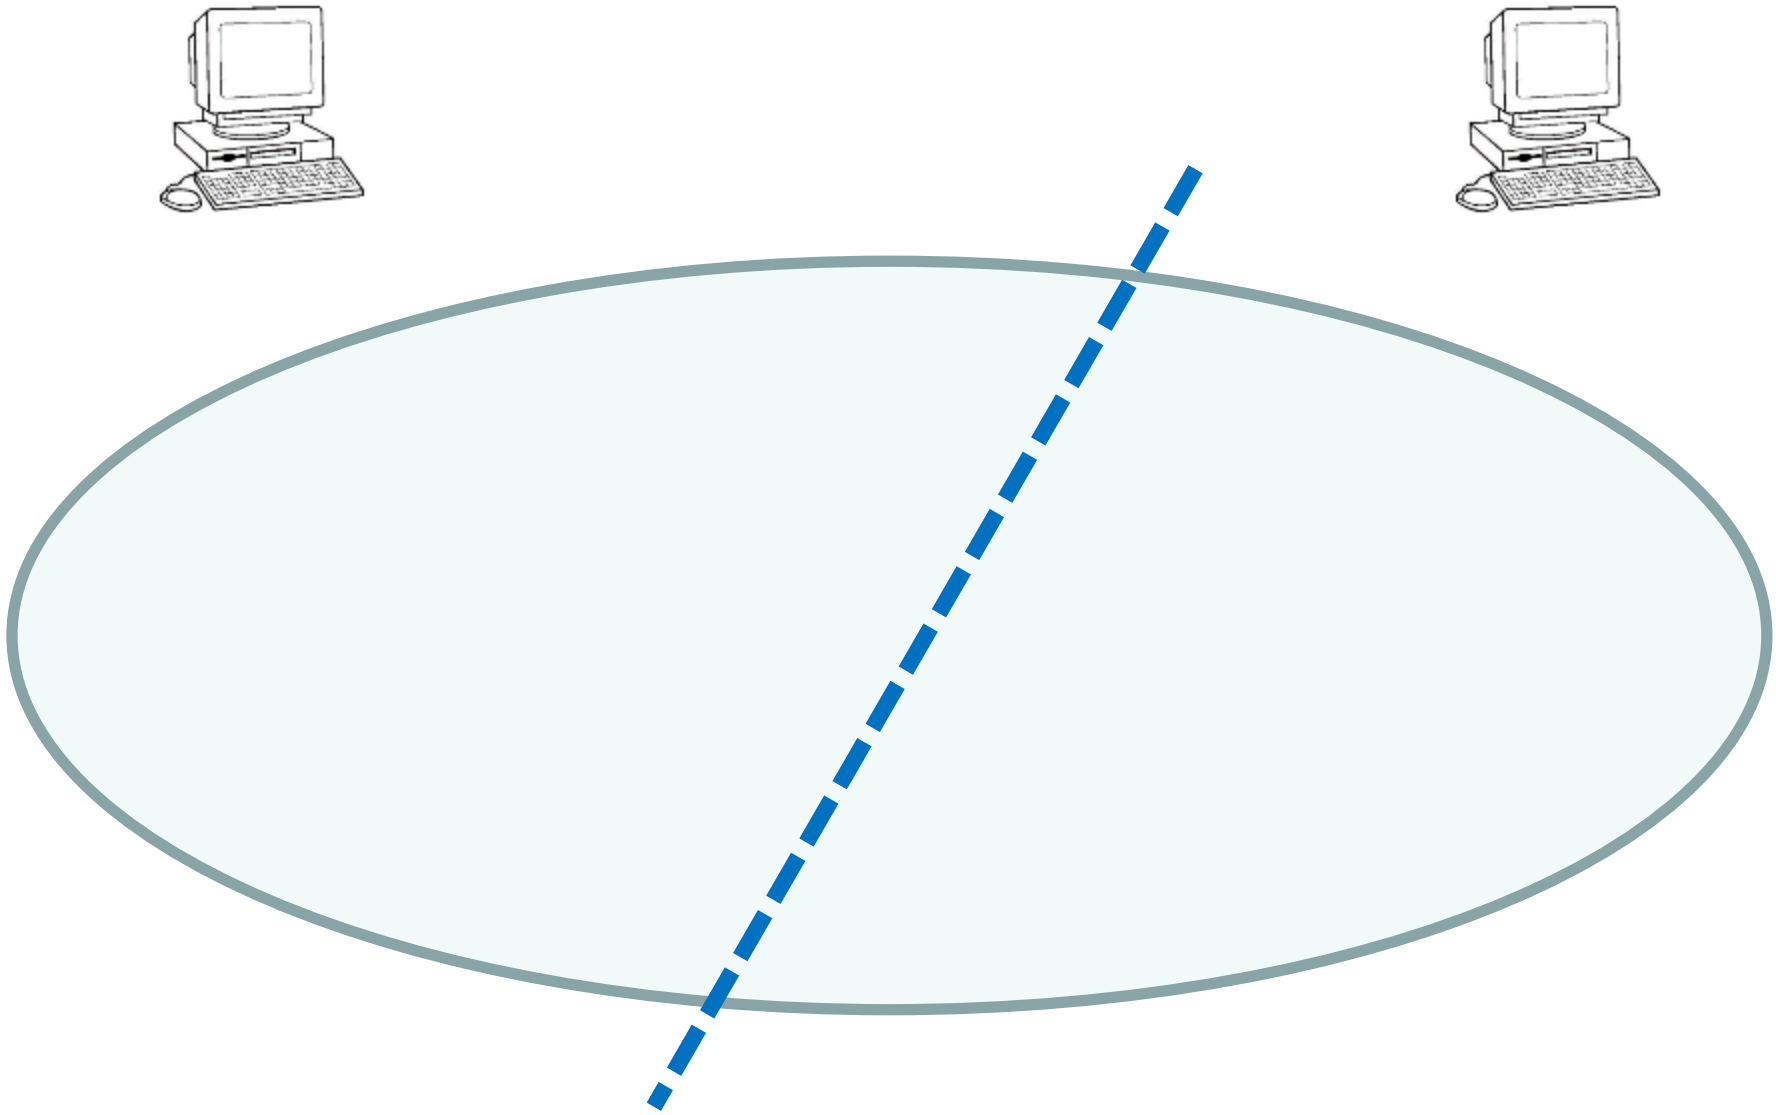

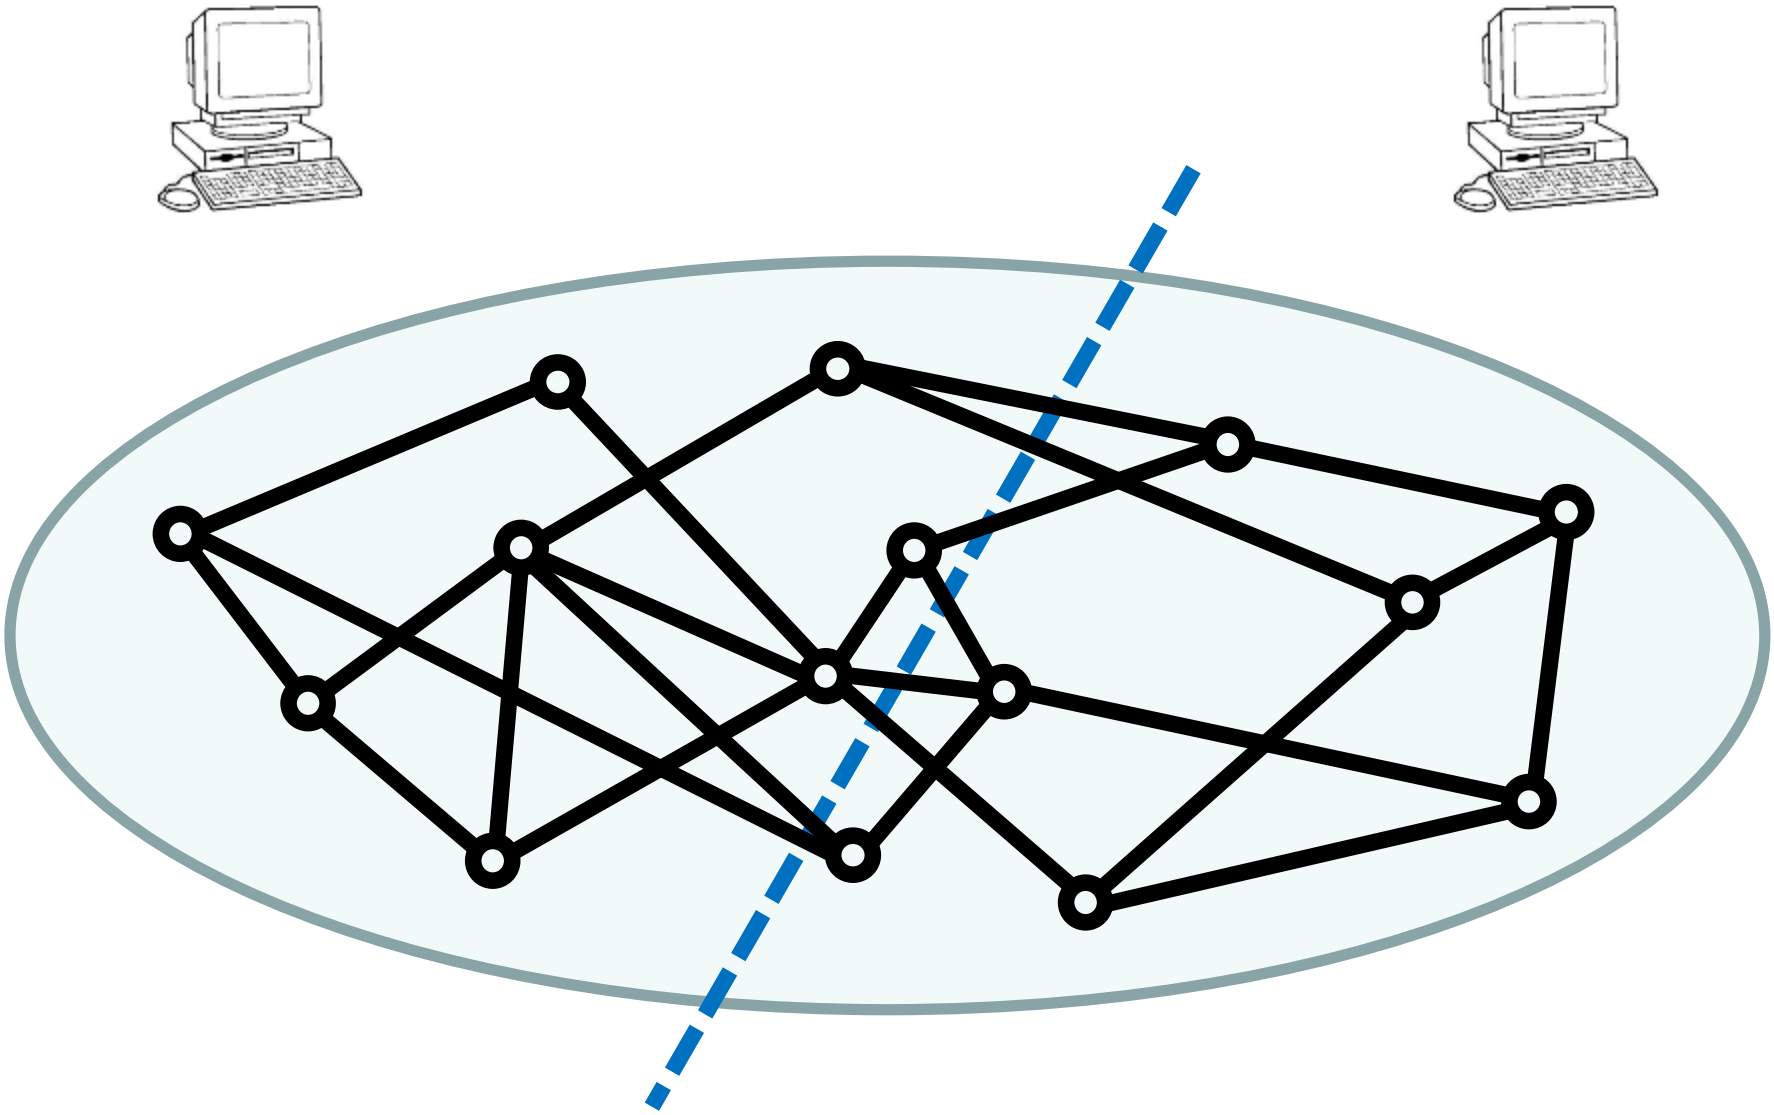

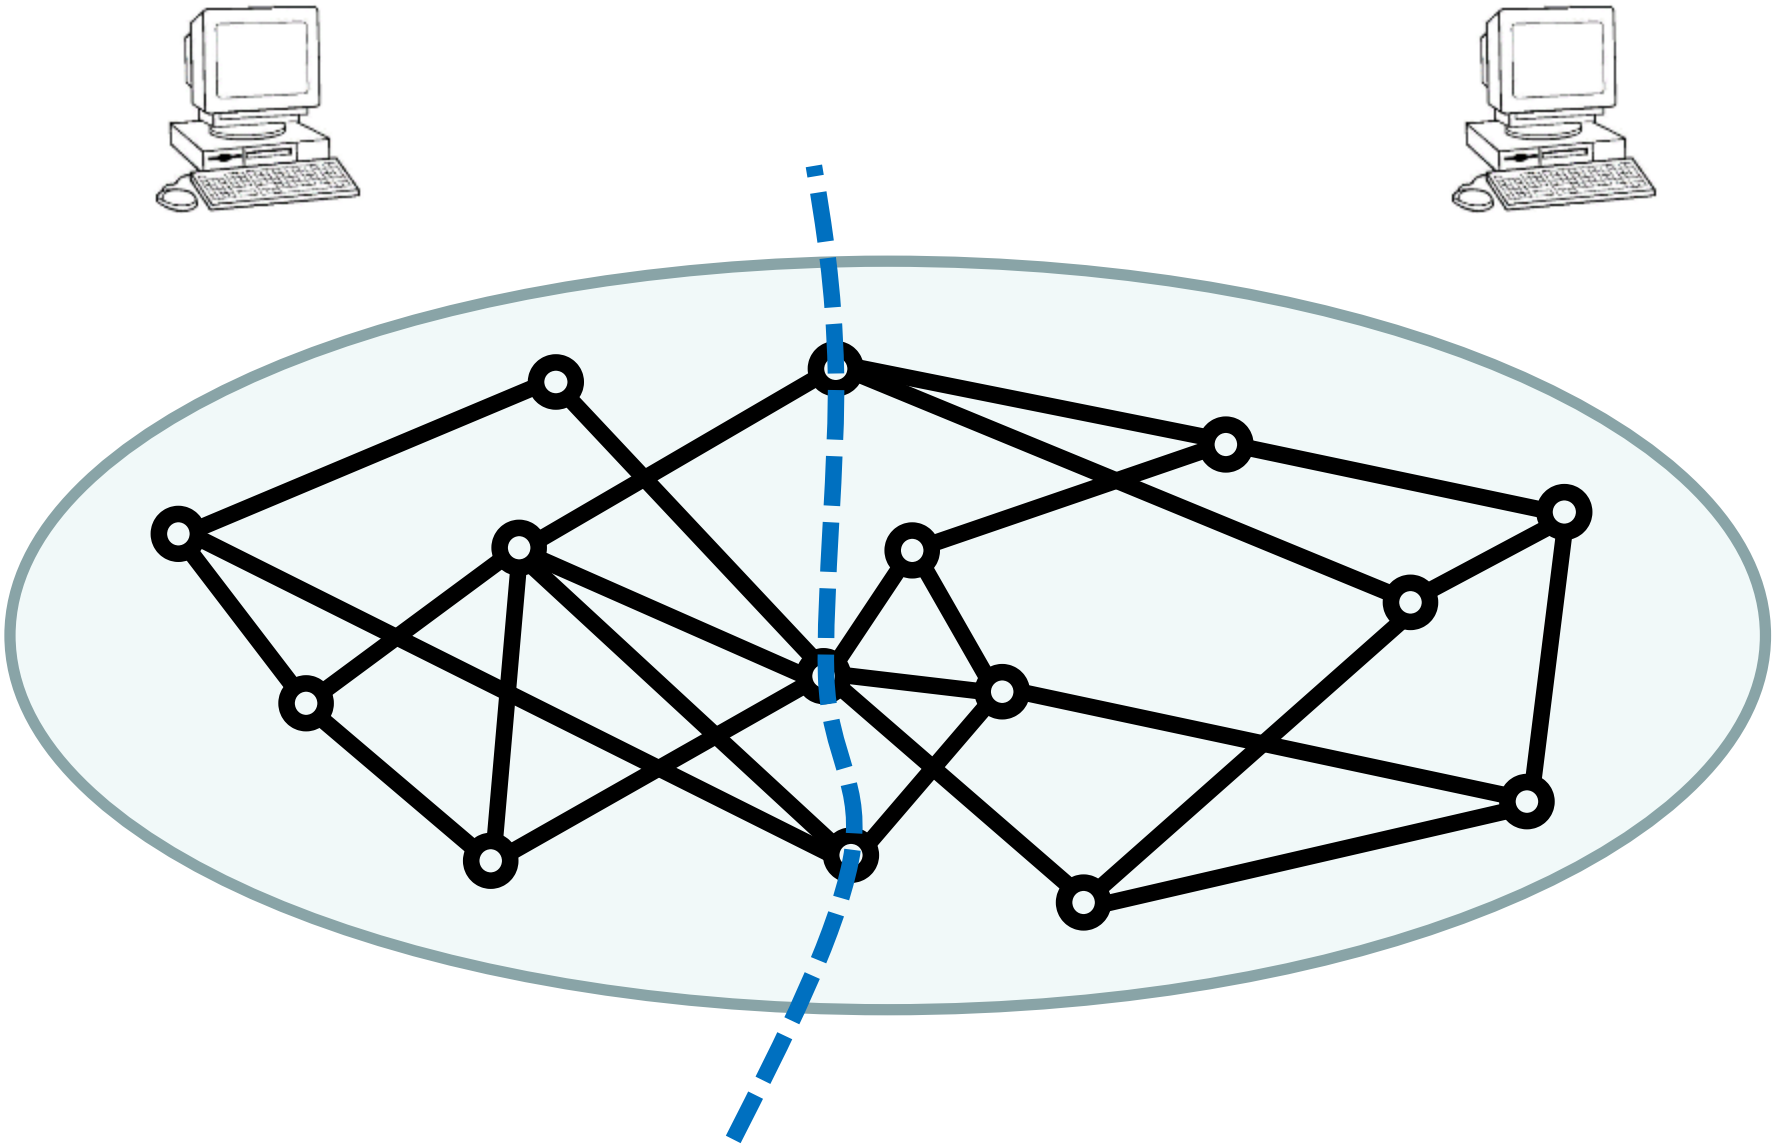

Find balanced separator of minimum size K.

Find balanced separator of minimum size K.

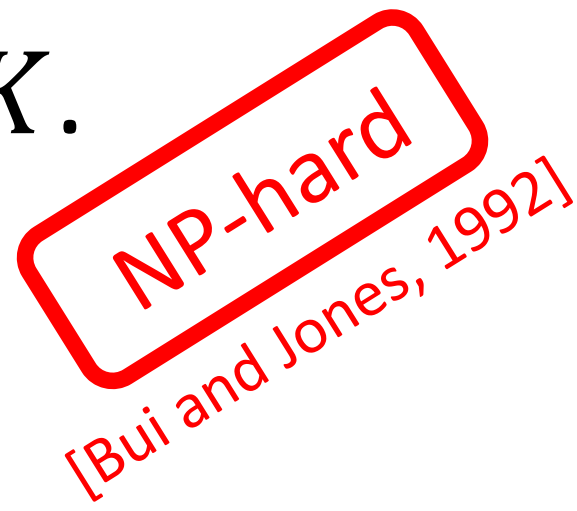

Find balanced separator of minimum size K.

Approx.

e.g.
[Feige, Mahdian, 2006]

NP-hard

[Bui and Jones, 1992]

Find balanced separator of minimum size K.

Our result: almost linear time algorithm for small K.

[Brandt, W., 2017]

Find balanced separator of minimum size K.

Our result: almost linear time algorithm for small K.

...in a boring way

[Brandt, W., 2017]

Algorithms for the Cloud

GPS for the Cloud

Just record 1ms of raw data

Coarse Time Navigation

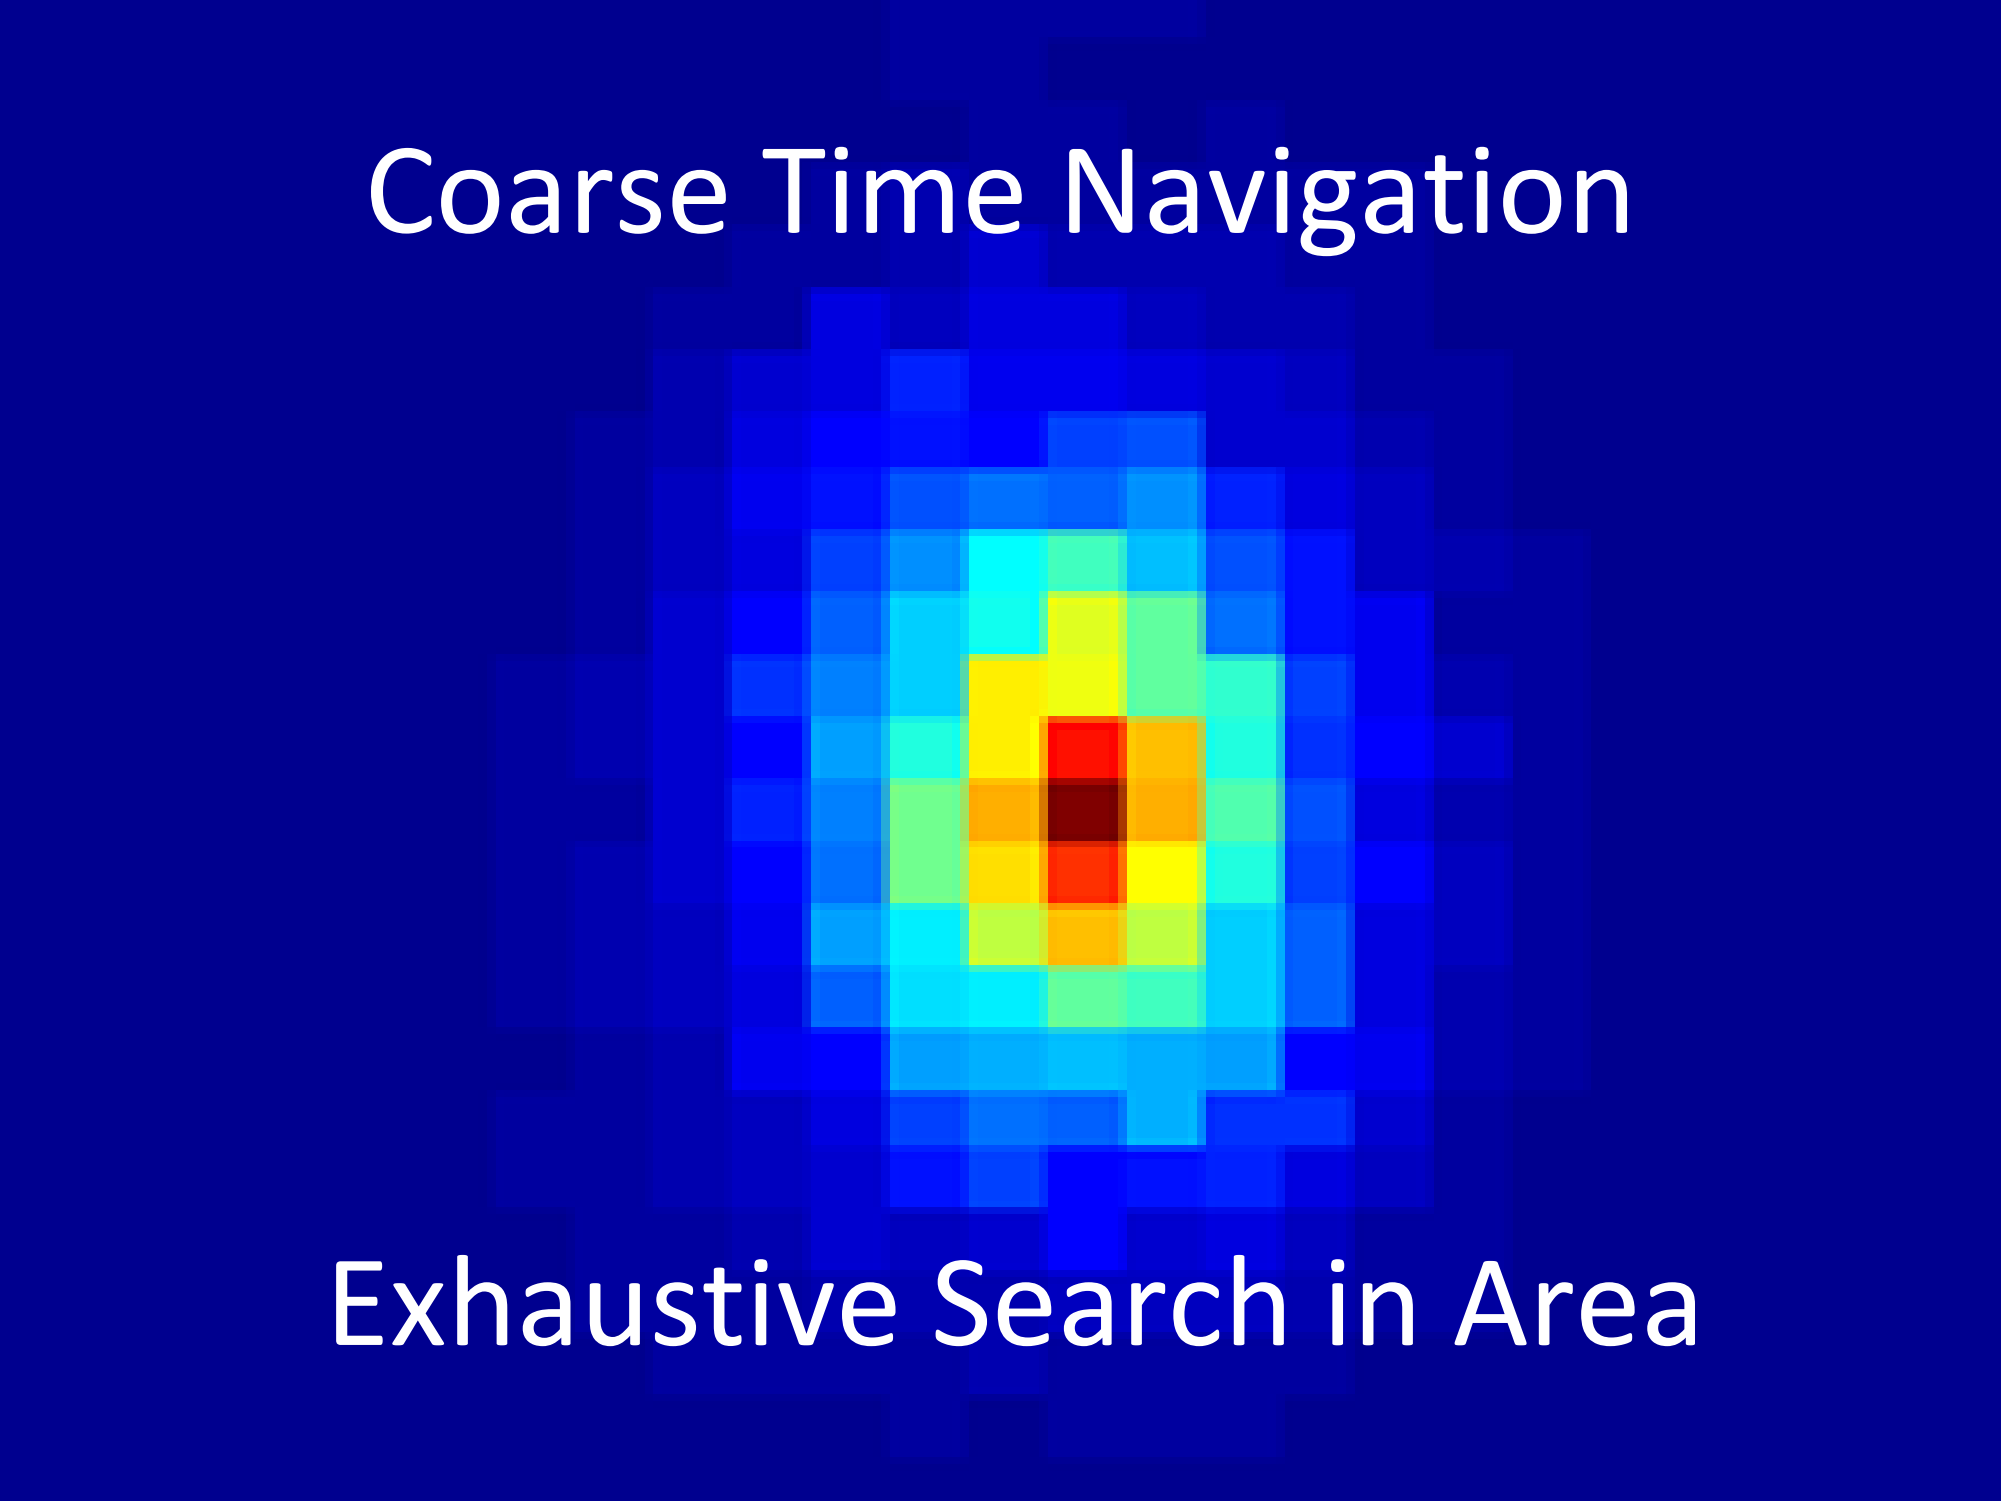

Exhaustive Search in Area

Also Robust to GPS Spoofing

Algorithms for the Cloud

$100B Revenue

¾ Online

Online Two Player Games

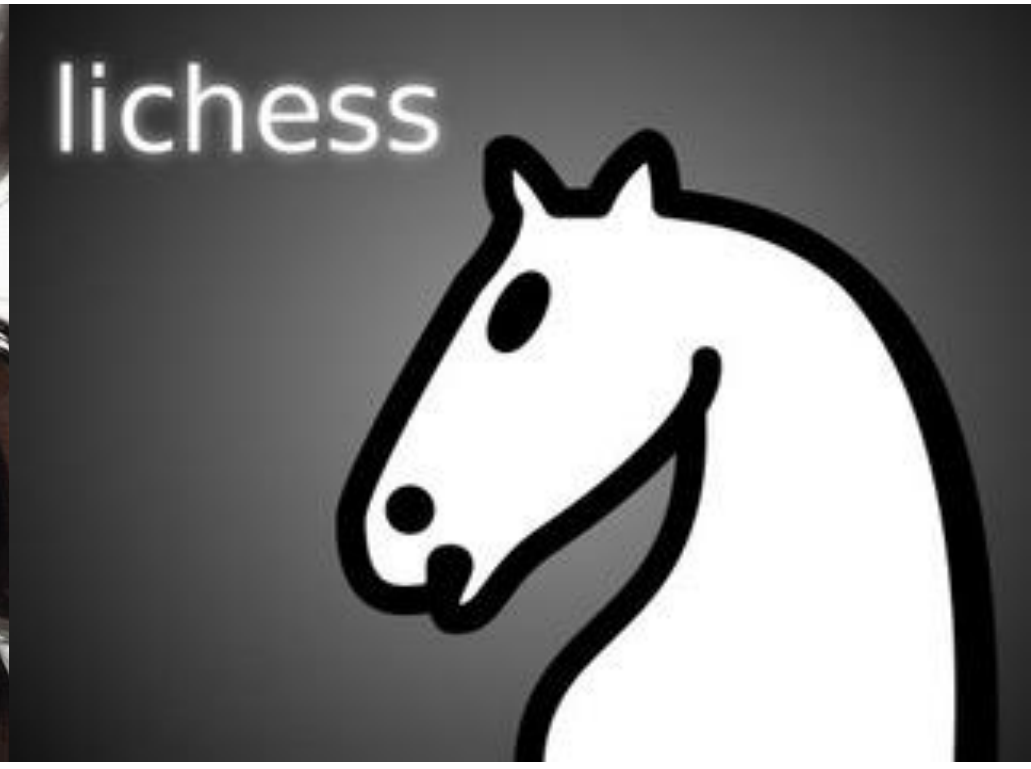

Match Players Fast

Waiting is Booooooring

Match Players Well

Similar Rating, Location, etc.

Min-Cost Perfect Matching With Delays (MPMD)

MPMD Example

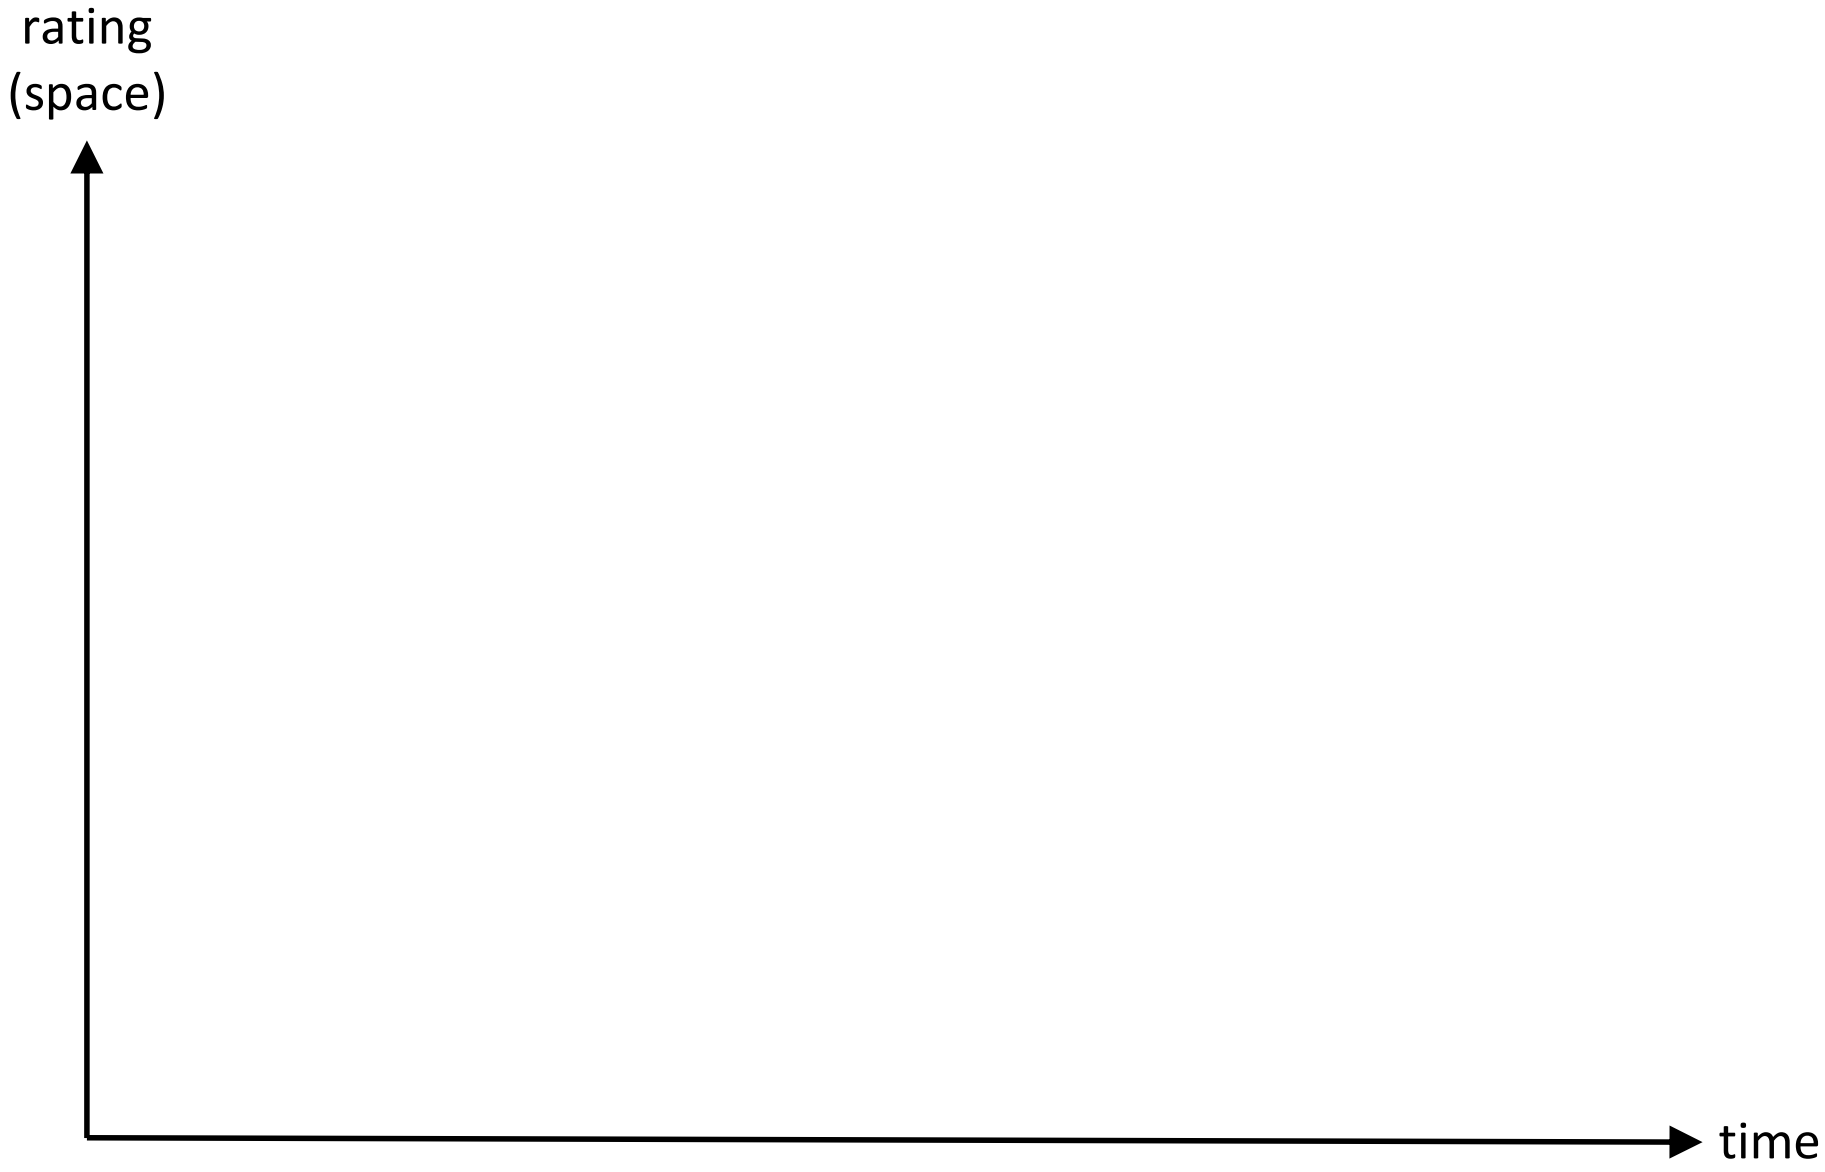

MPMD Example

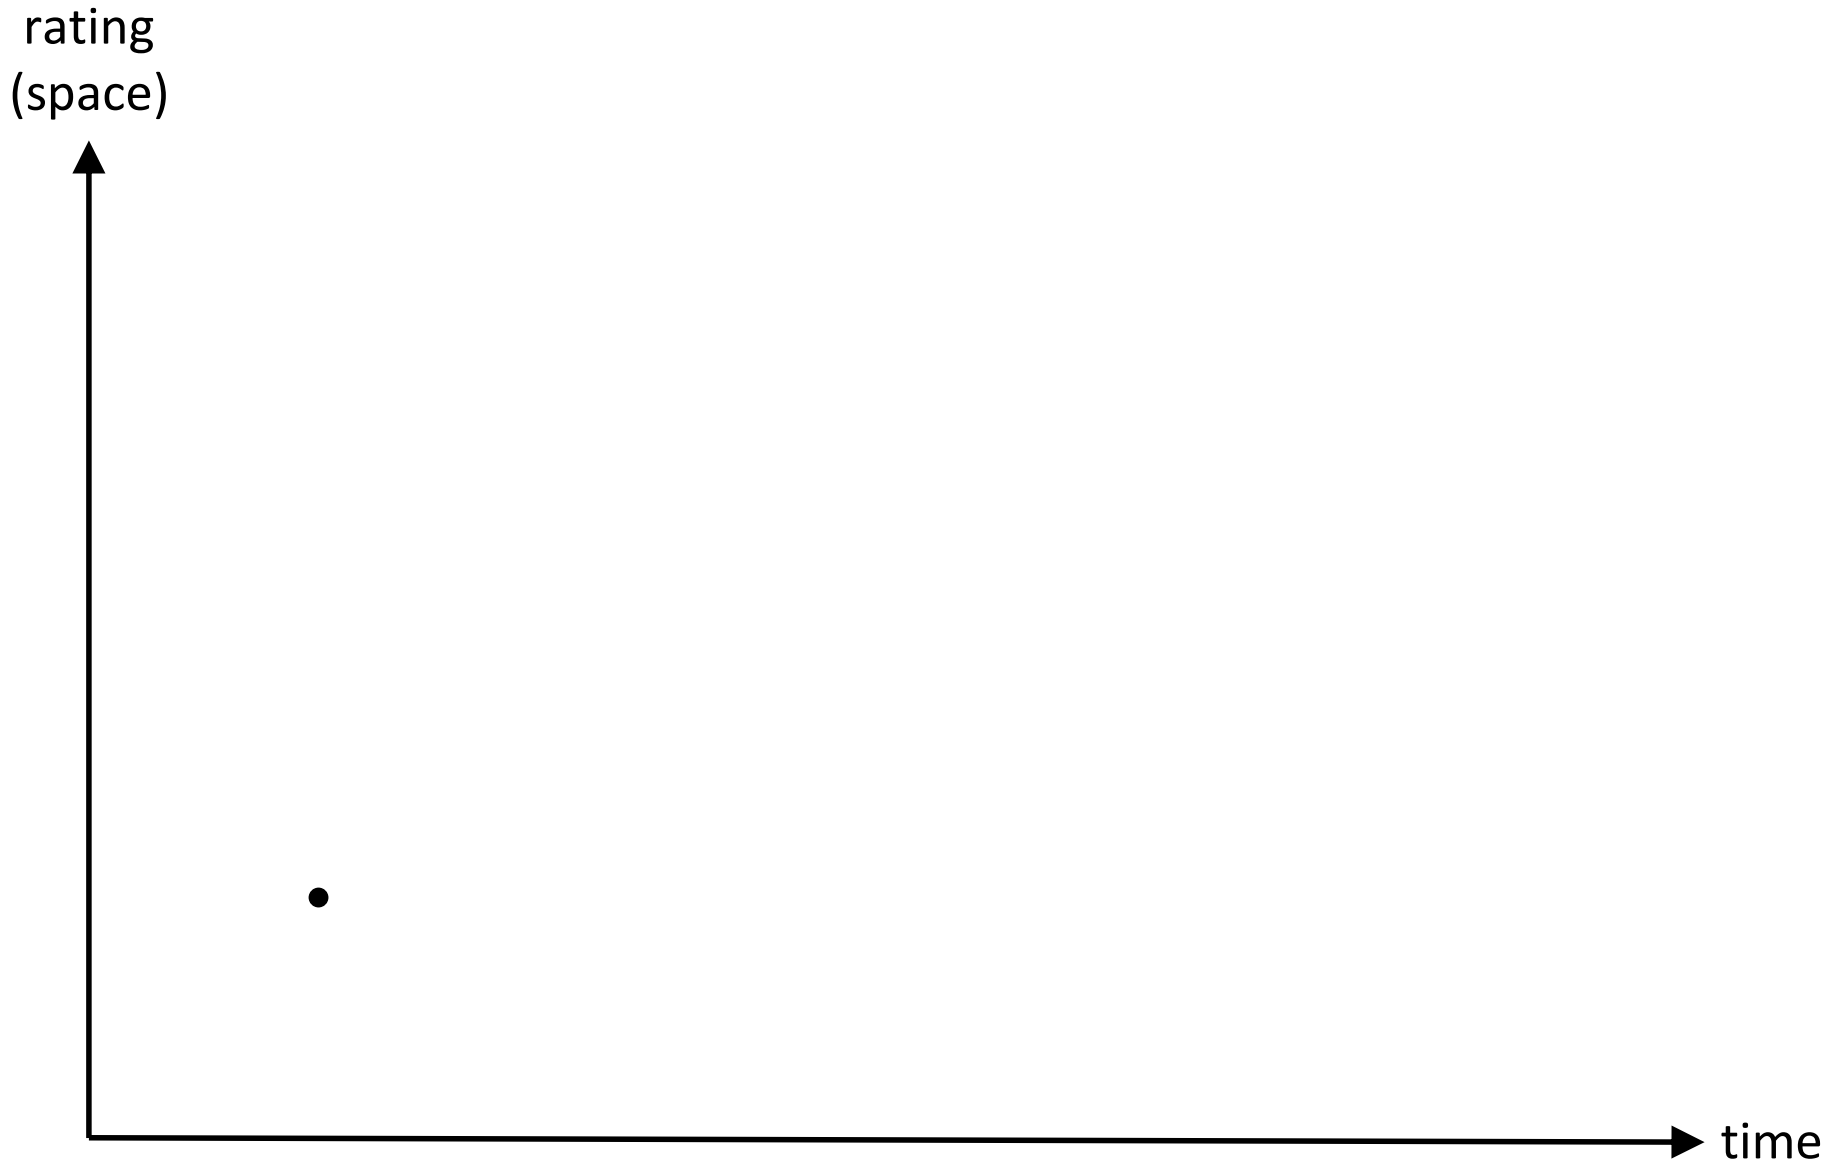

MPMD Example

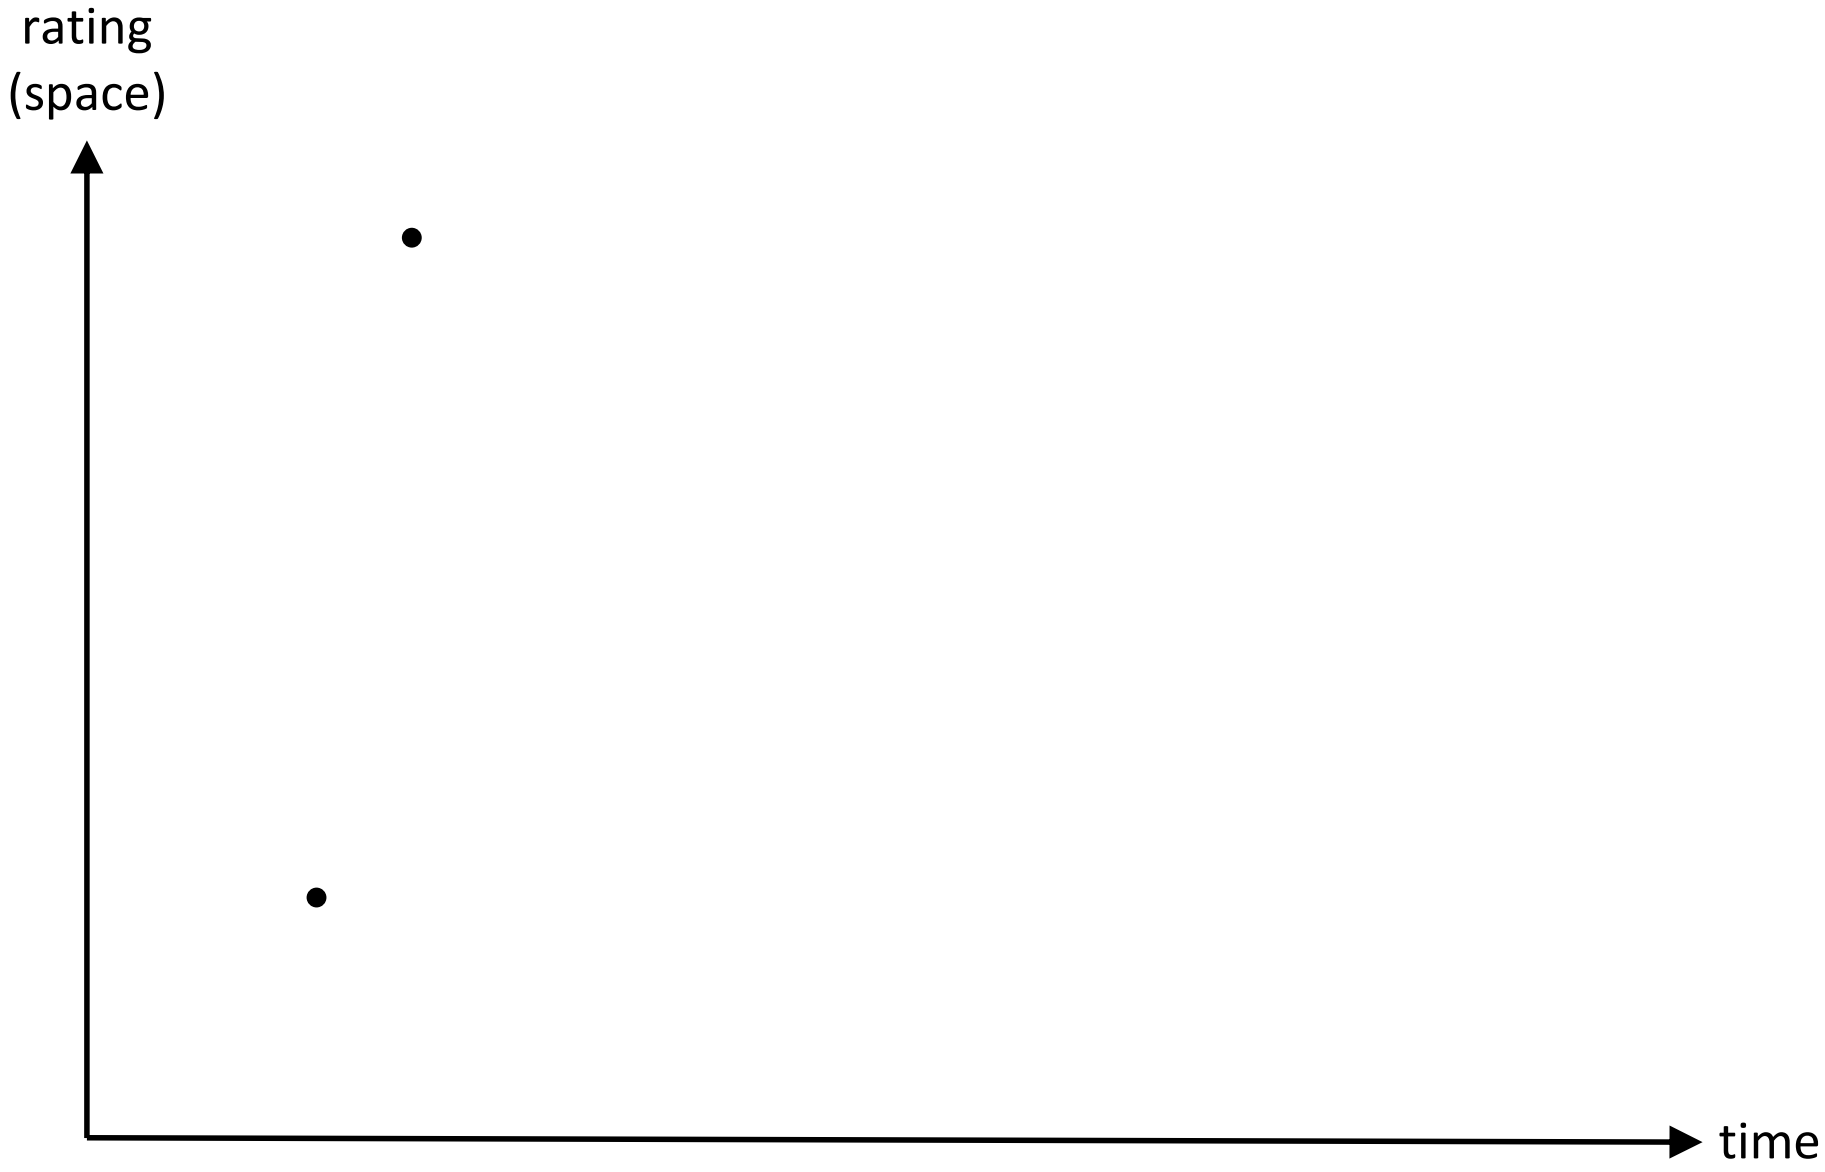

MPMD Example

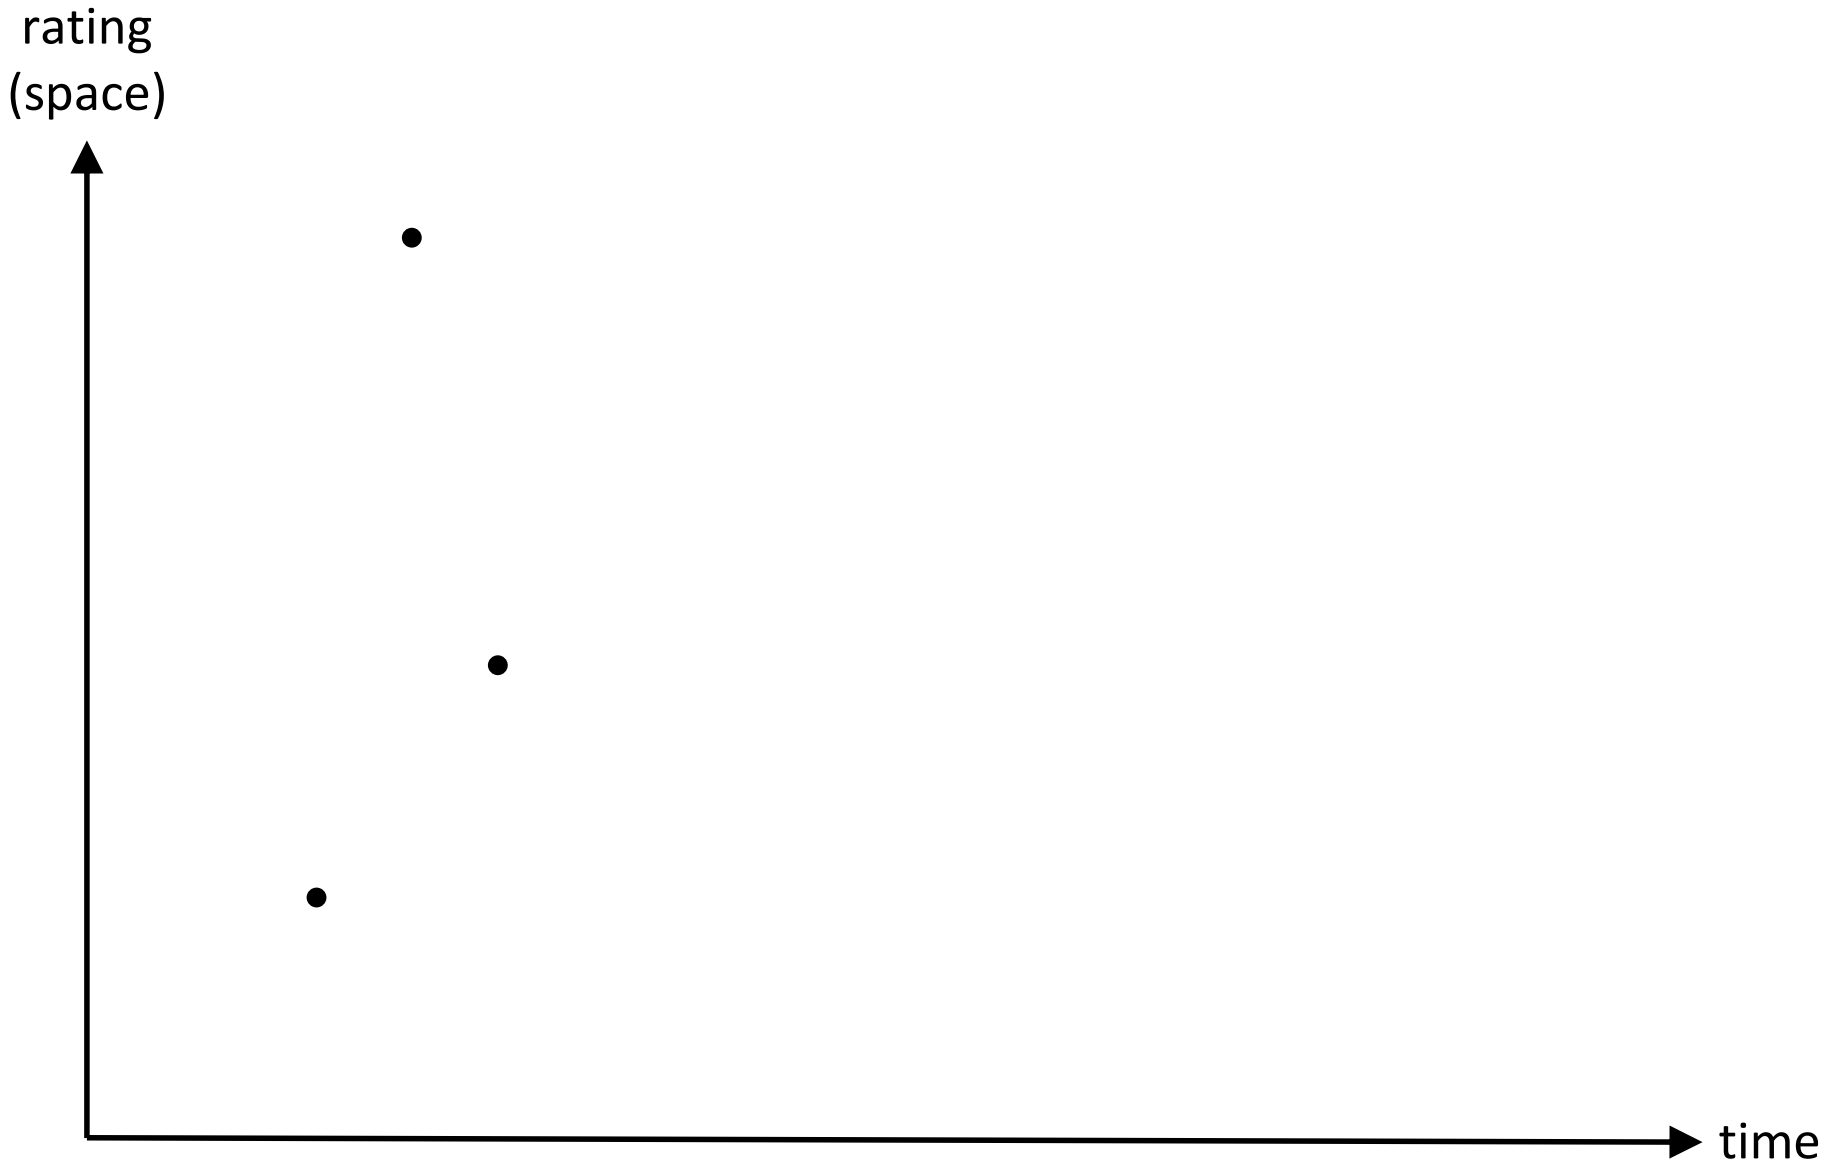

MPMD Example

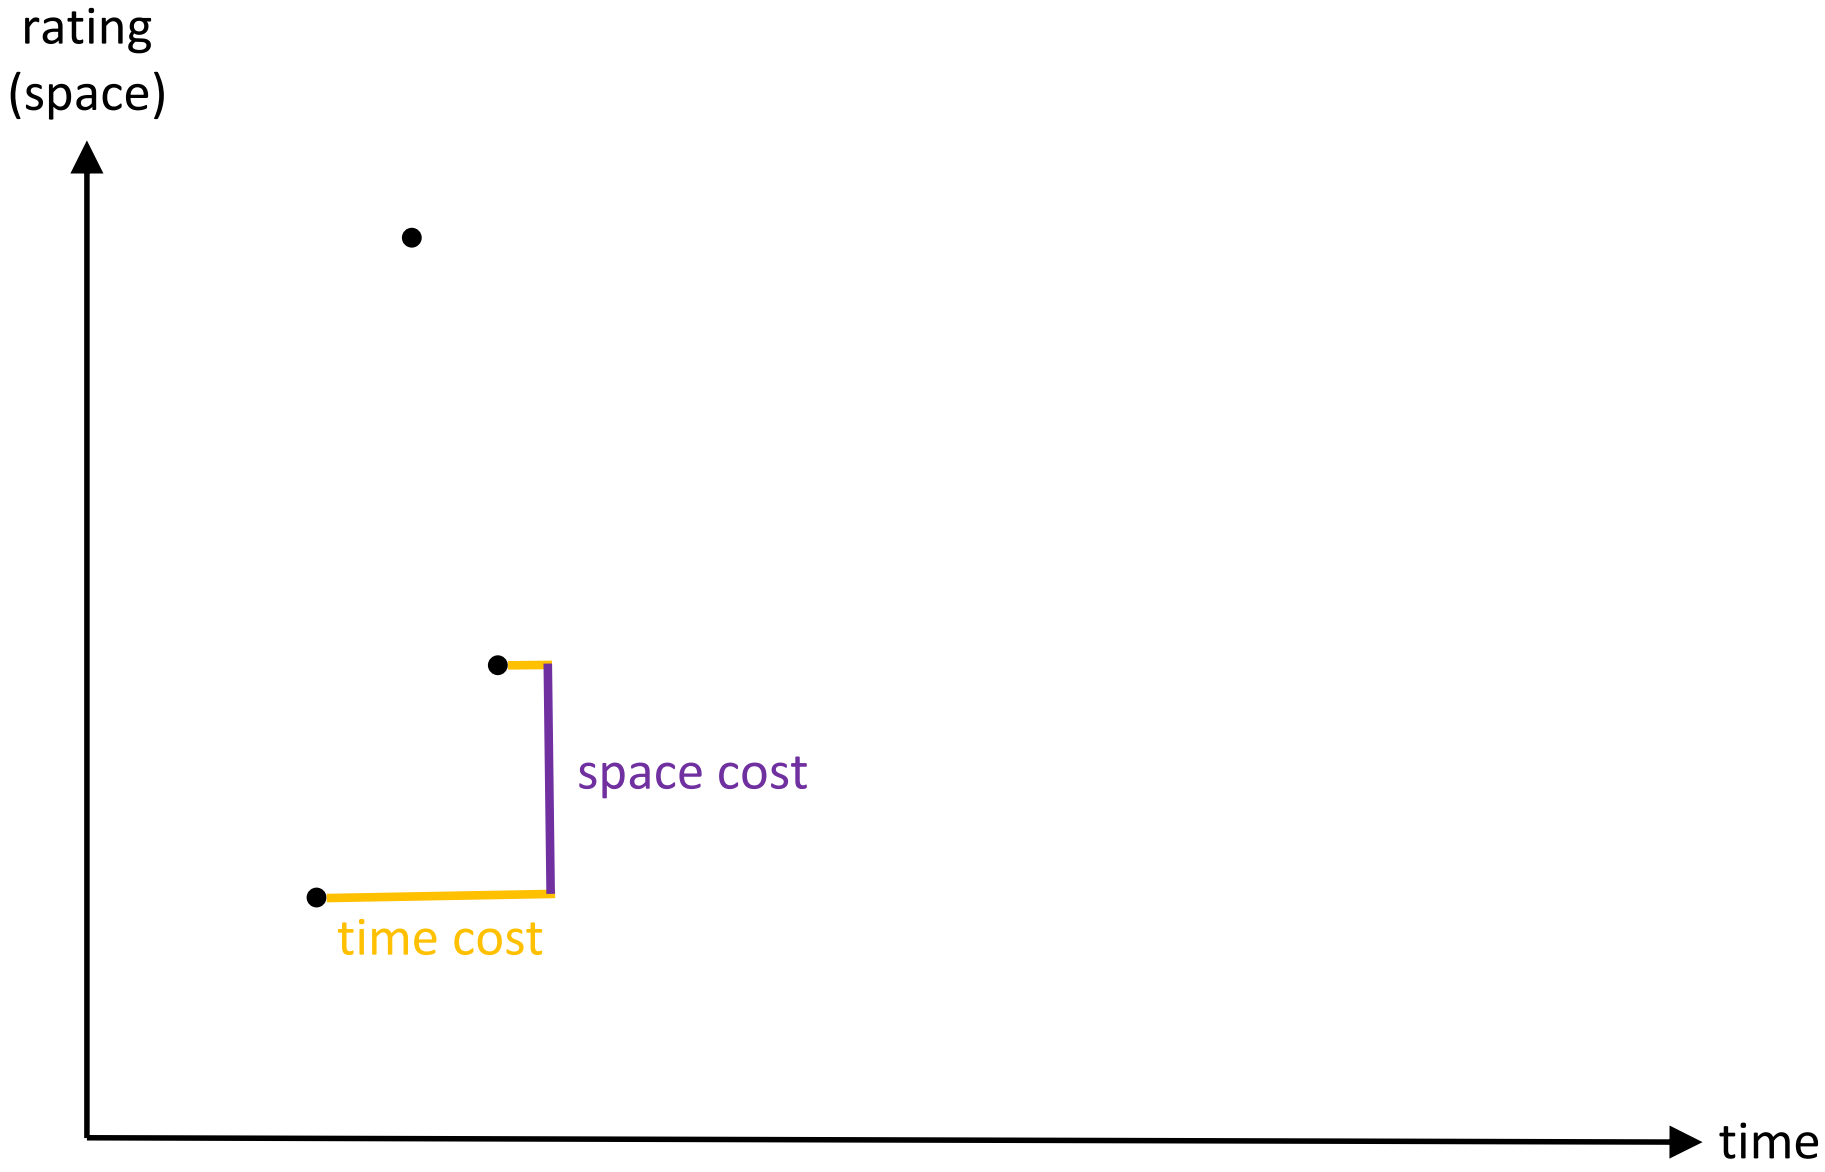

MPMD Example

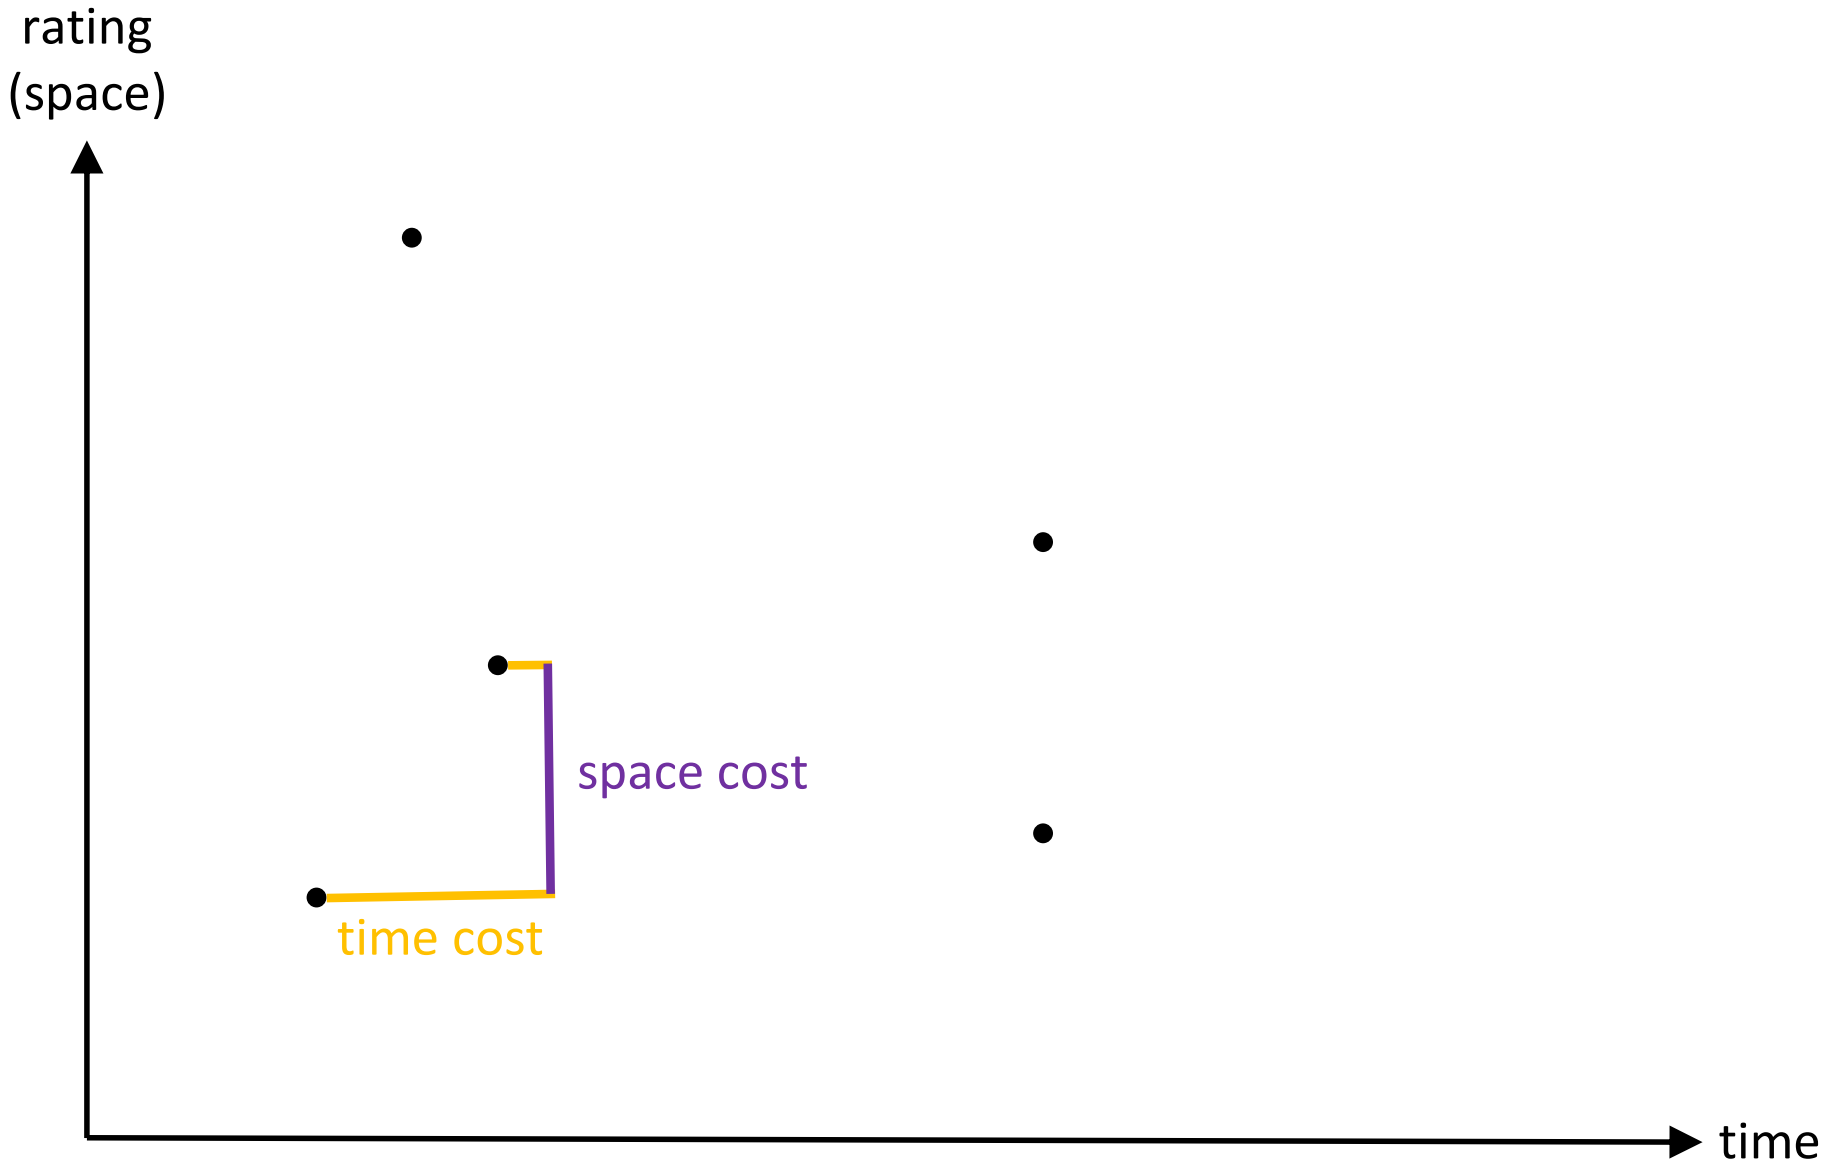

MPMD Example

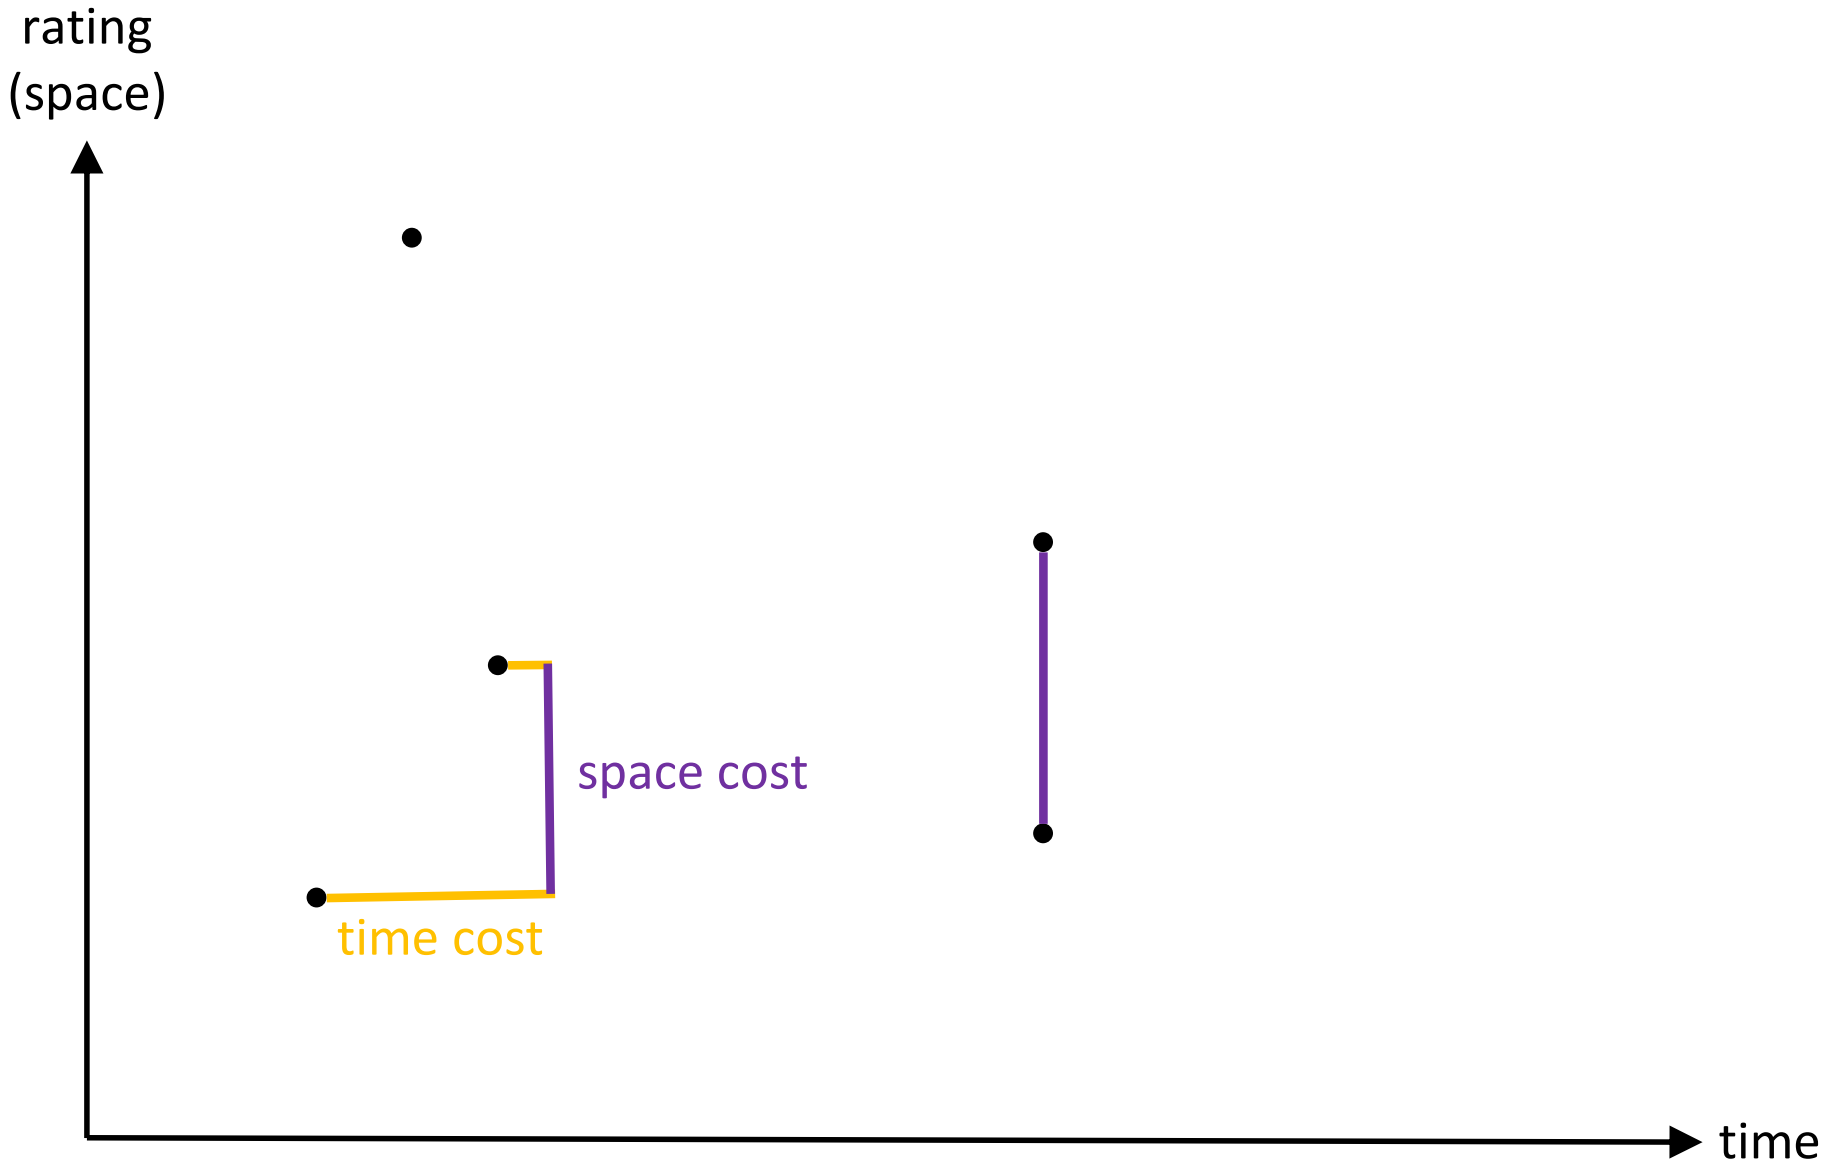

MPMD Example

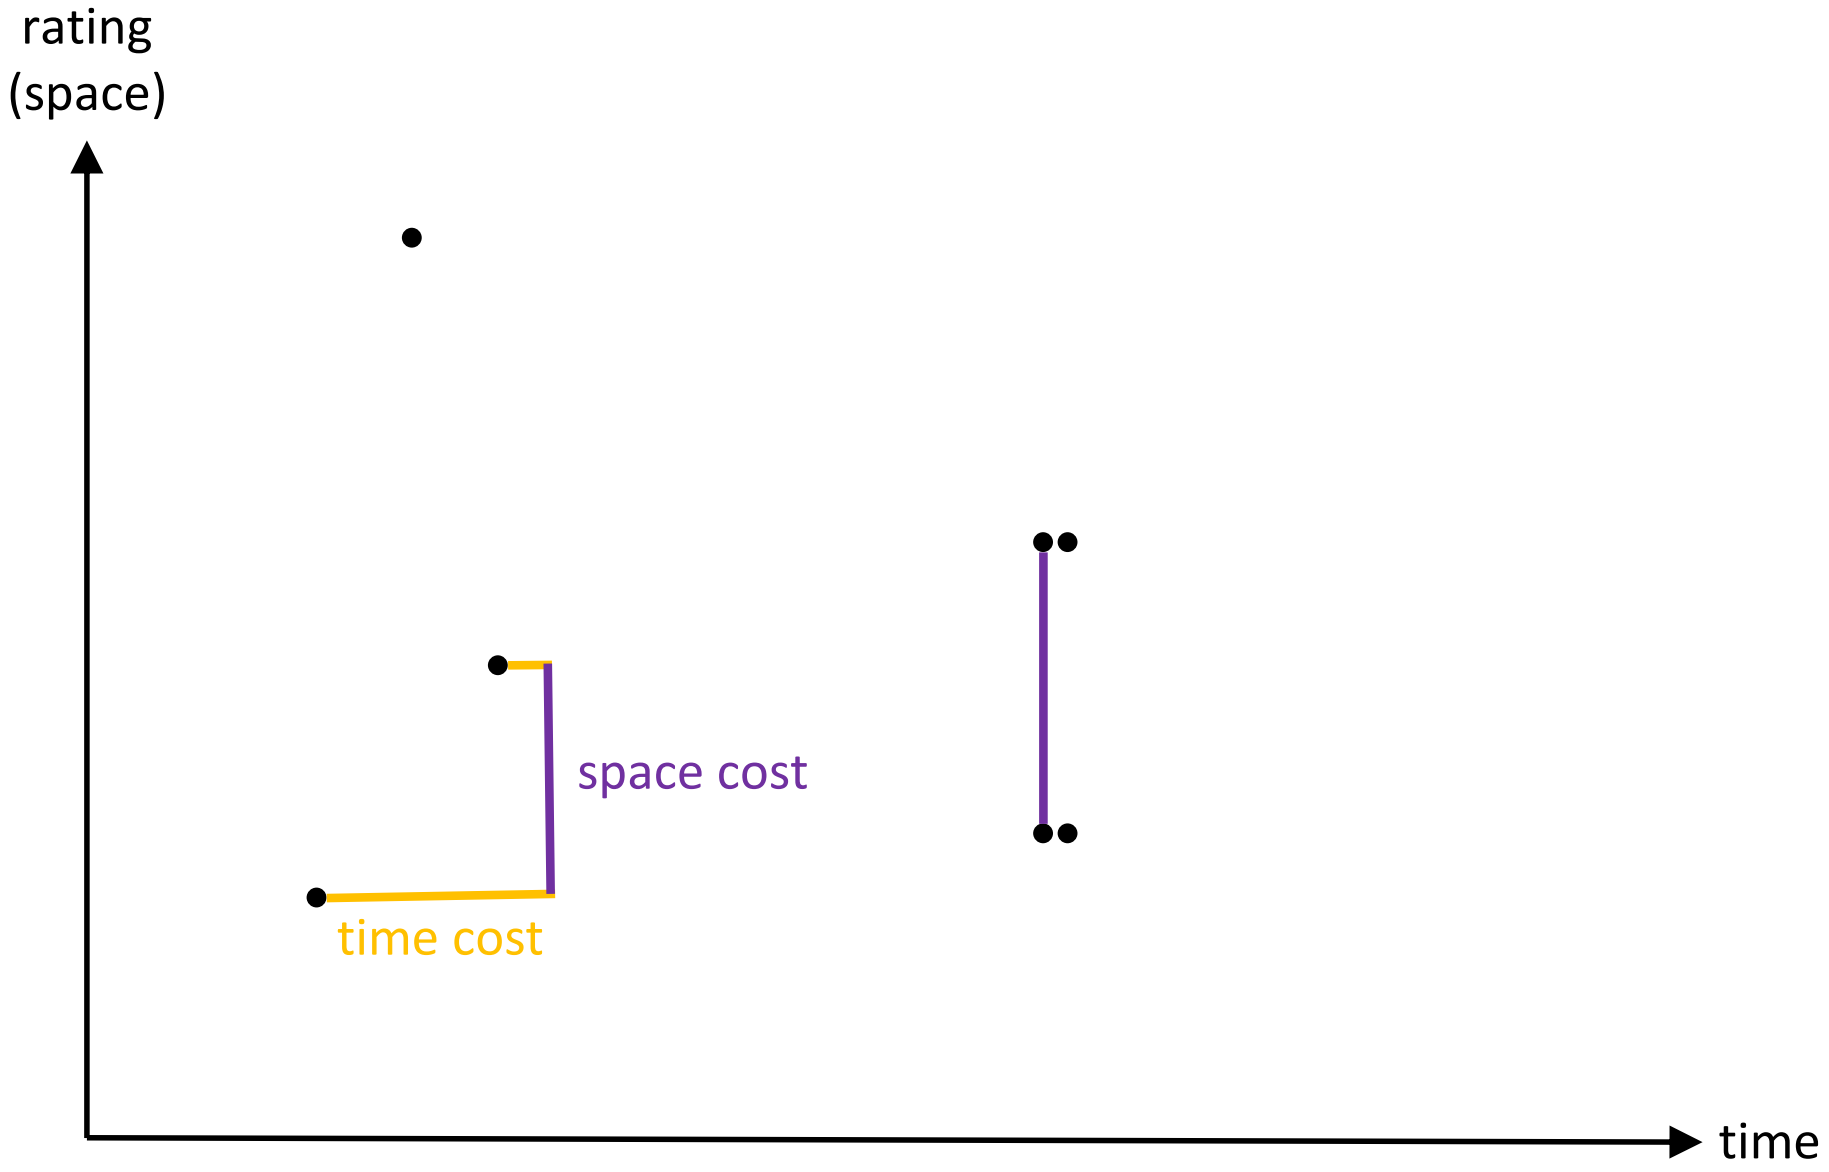

Haste Makes Waste!

MPMD Example

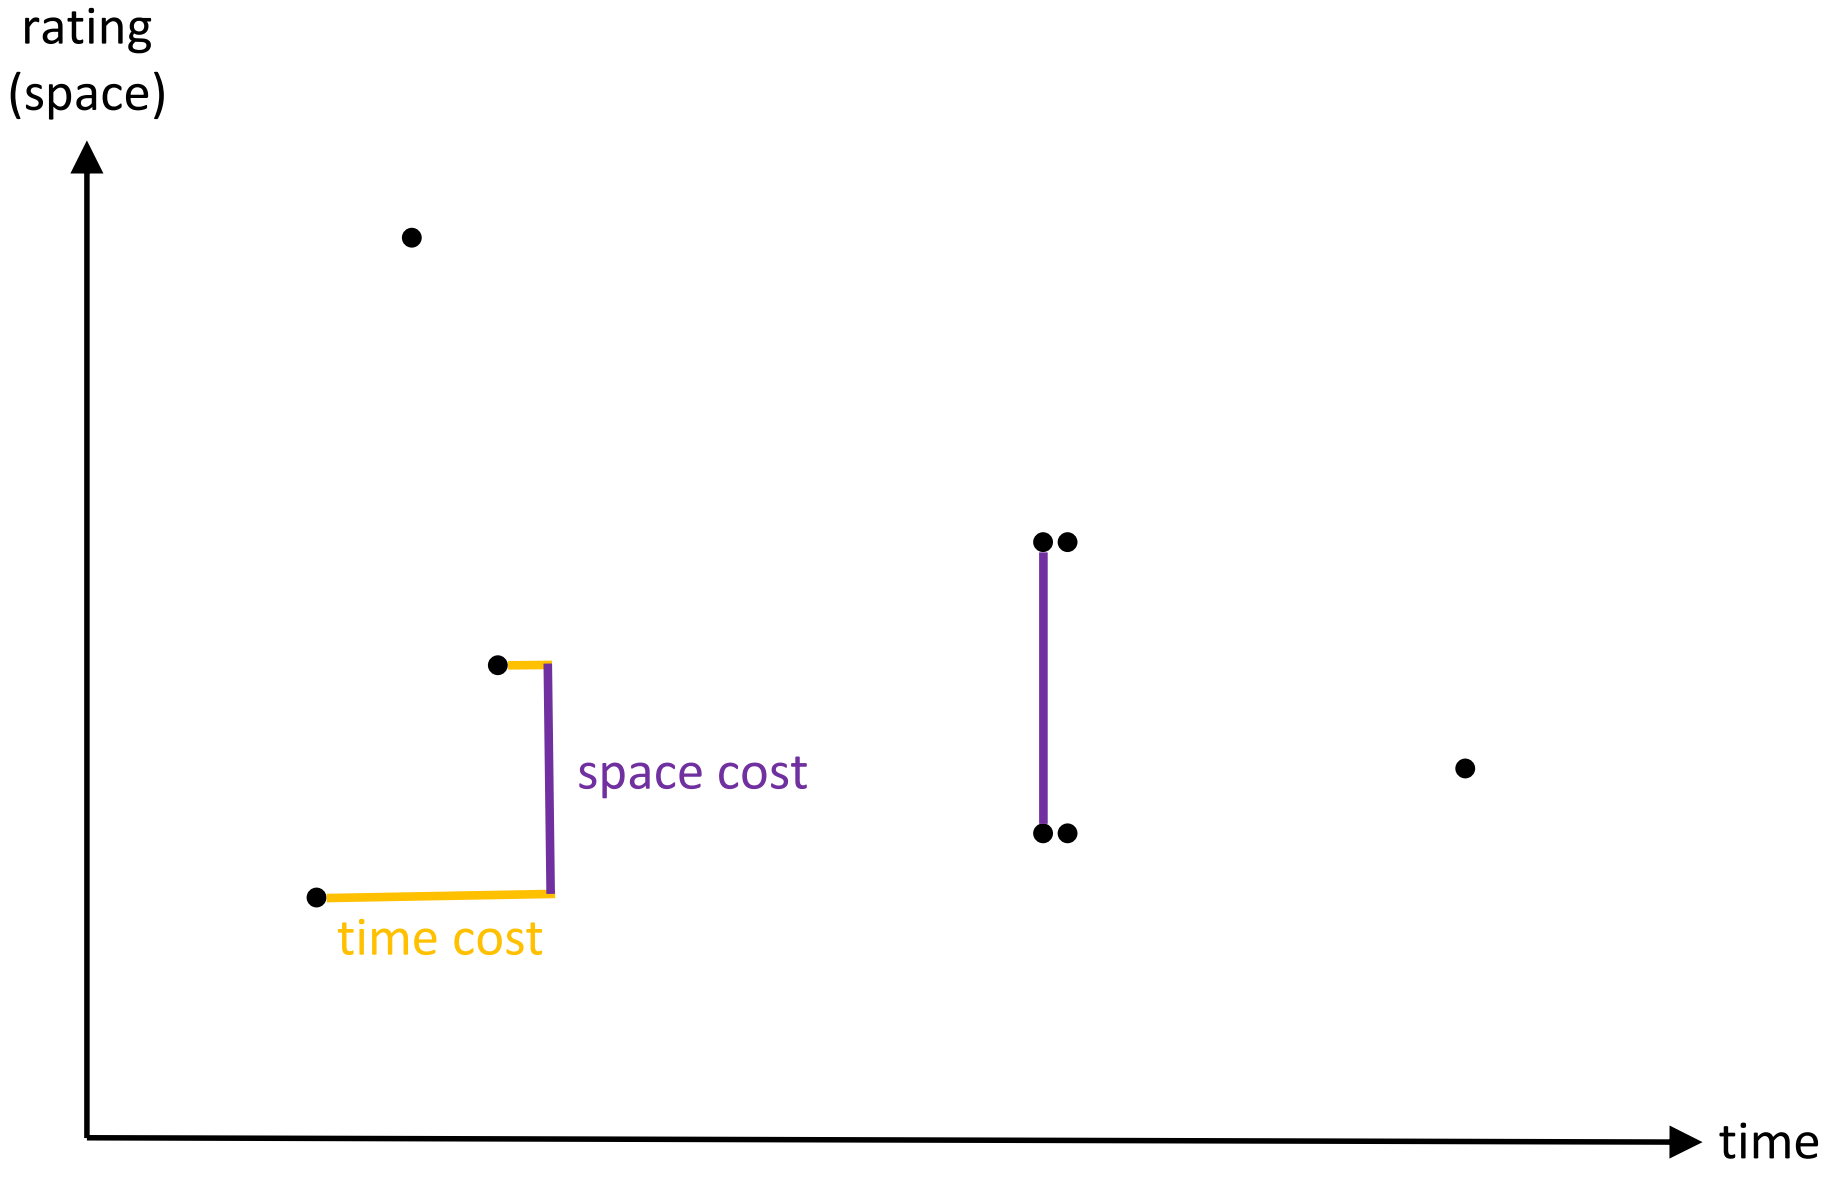

MPMD Example

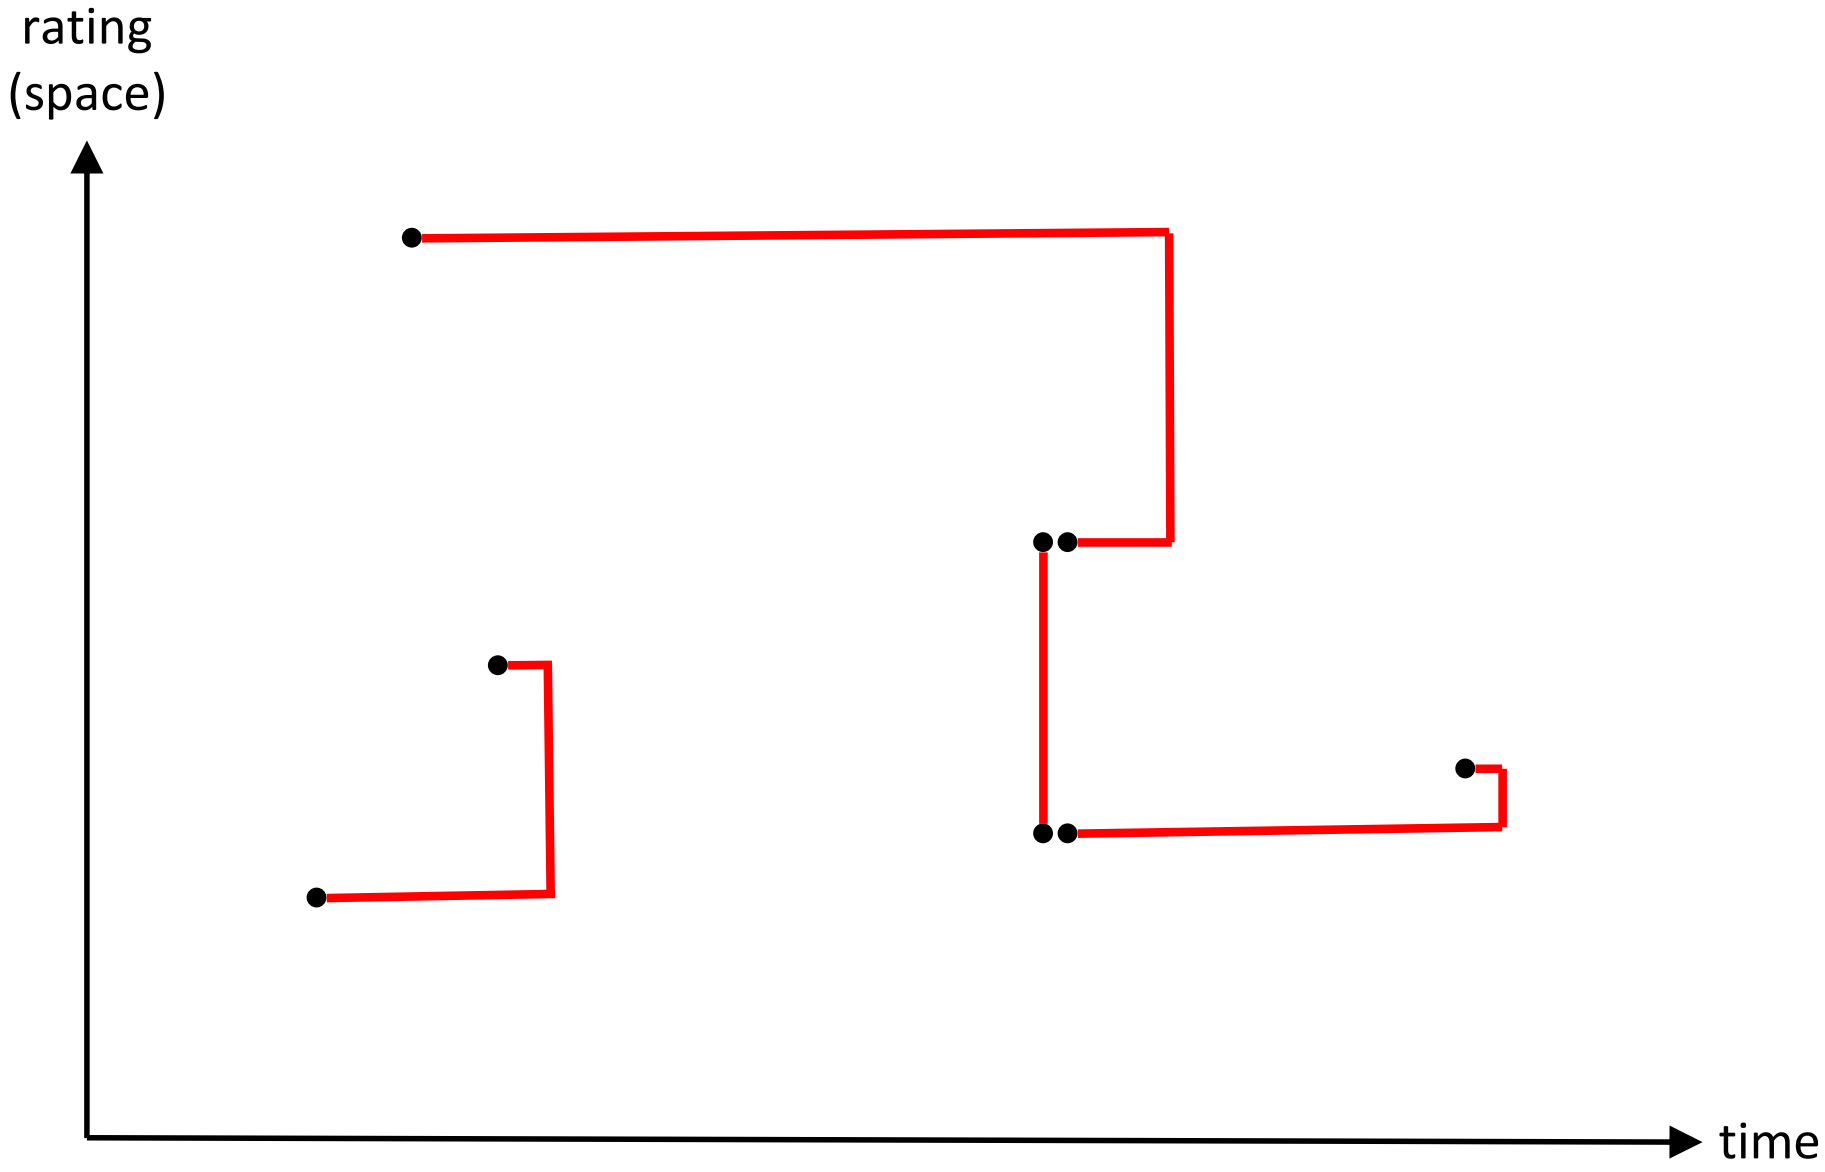

MPMD Example

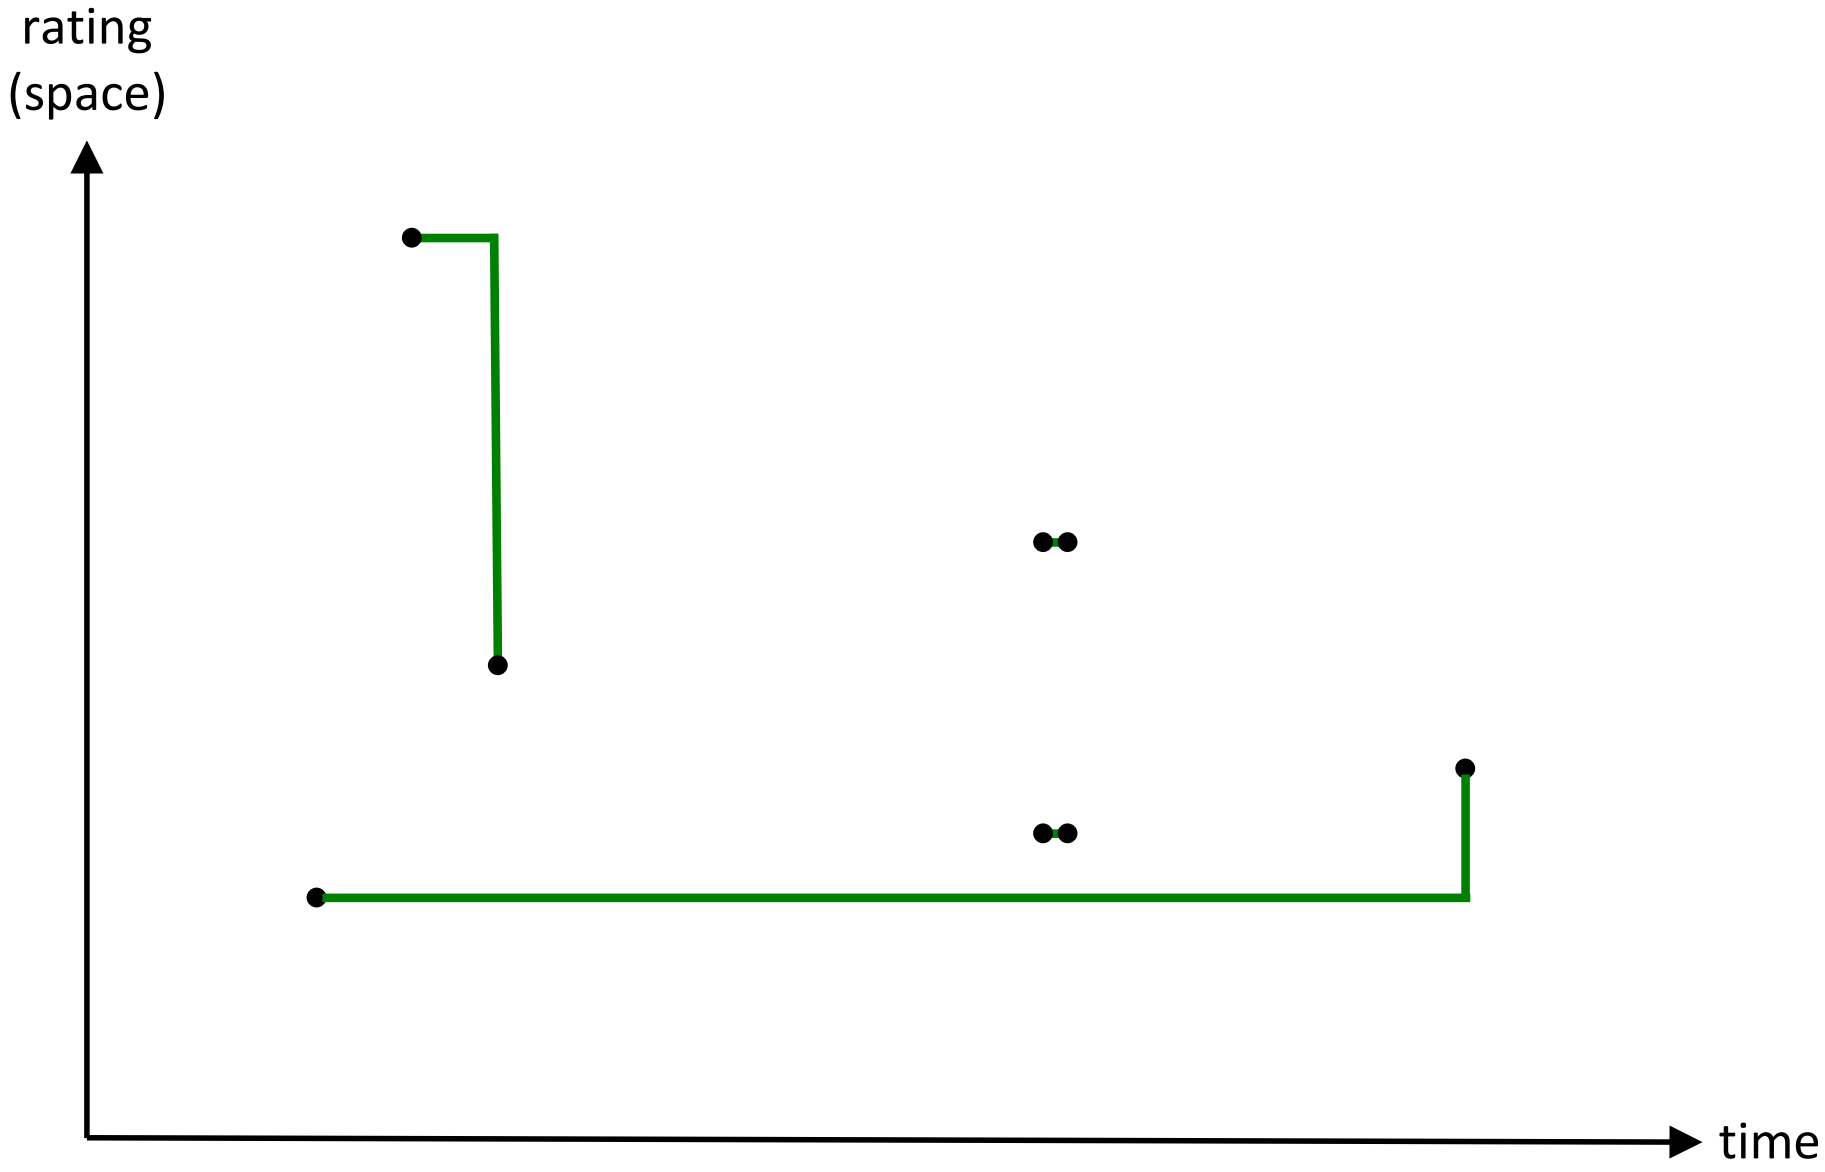

MPMD Example

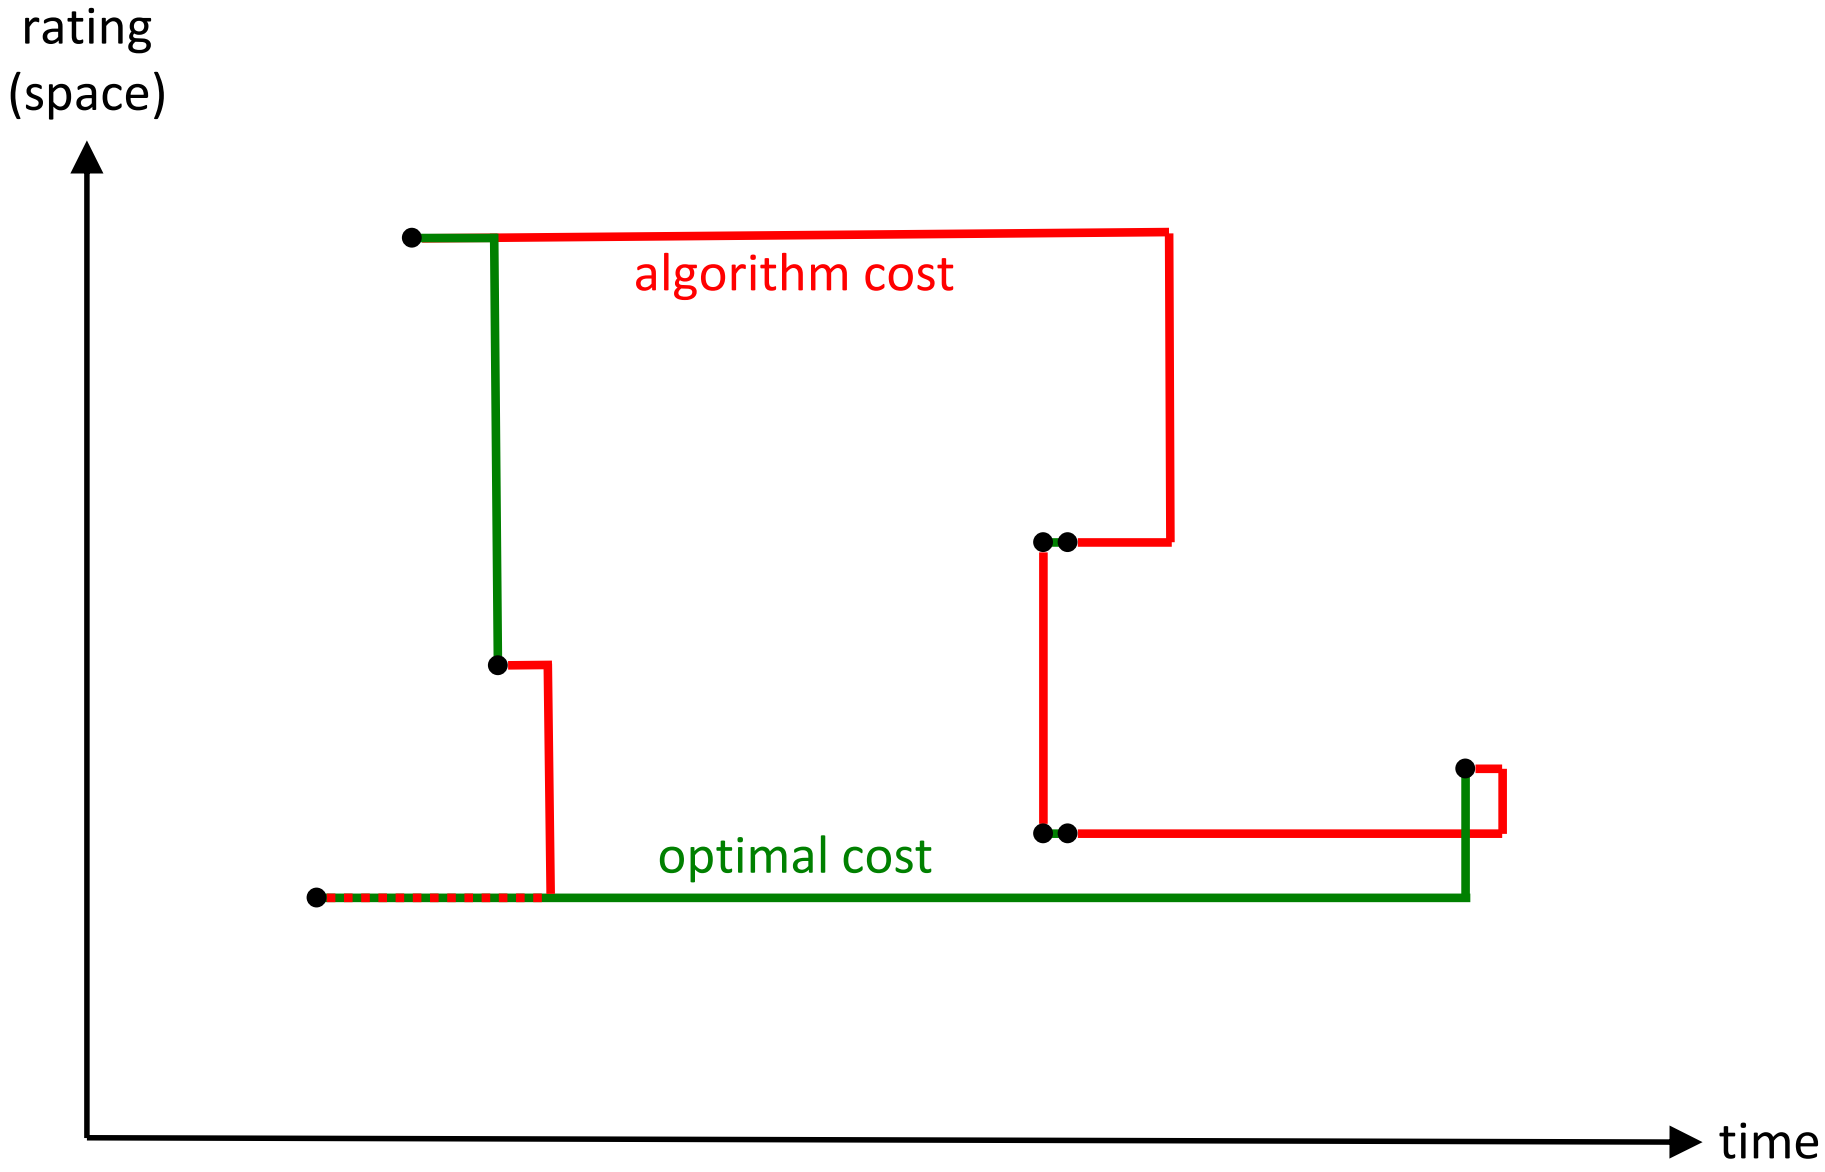

Online Matching Literature

- Bipartite graph, left side is known, right side revealed online
  - Maximum cardinality matching
    [KVV1990, BM2008, GM2008, DJK2013, M2014, NW2015]
  - Maximum vertex weighted matching
    [AGKM2011, DJK2013, NW2015]
  - Maximum capacitated assignment (the AdWords problem)
    [MSVV2005, BJN2007, GM2008, AGKM2011, NW2015]
  - Metric maximum weight matching
    [KP1993, KMV1994]
  - Metric minimum cost perfect matching
    [KP1993, MNP2006, BBGN2014]
  - Metric minimum capacitated assignment (transportation)
    [KP2000]
- MPMD: known graph, both sides revealed online

MPMD Results

- Finite metric space $\mathcal{M} = (V, \delta)$
  - $n = |V|$
  - $\Delta = \frac{\max_{x \neq y \in V} \delta(x,y)}{\min_{x \neq y \in V} \delta(x,y)}$
- $O(\log^2 n + \log \Delta)$-competitive randomized algorithm
  [Emek, Kutten, W 2016]
- $O(\log n)$-competitive (almost) deterministic algorithm
  Lower bound of $\Omega(\sqrt{\log n})$
  [Azar, Chiplunkar, Kaplan 2017]
- $O(\log n)$-competitive (almost) det. bipartite algorithm
  $\Omega(\sqrt{\log n / \log \log n})$ lower bound for bipartite
  $\Omega(\log n / \log \log n)$ lower bound for non-bipartite
  [Wang et al., 2018]
  ...

The $O(\log n)$ Algorithm

Approximate Metric by Tree

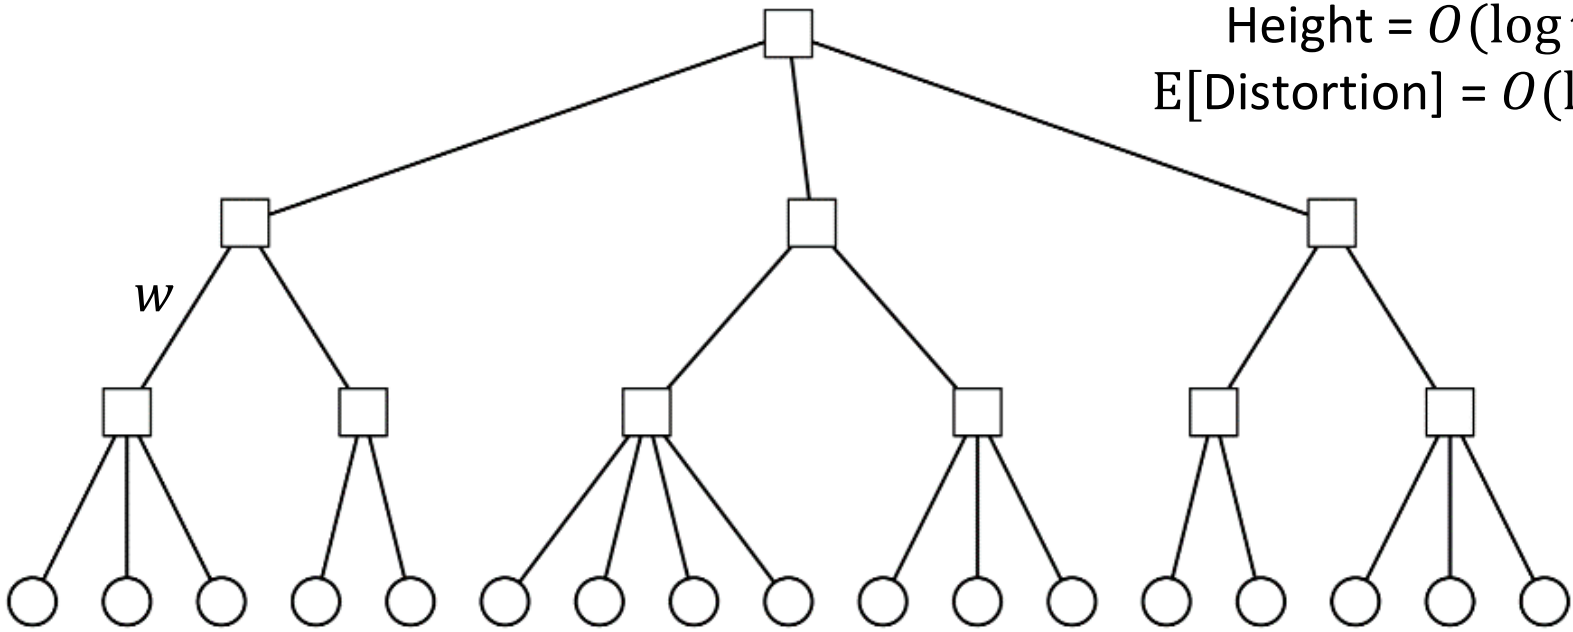

Leaves = Nodes in Metric Space

[Fakcharoenphol, Rao, Talwar 2004], [Bansal, Buchbinder, Gupta, Naor 2015]

Algorithm

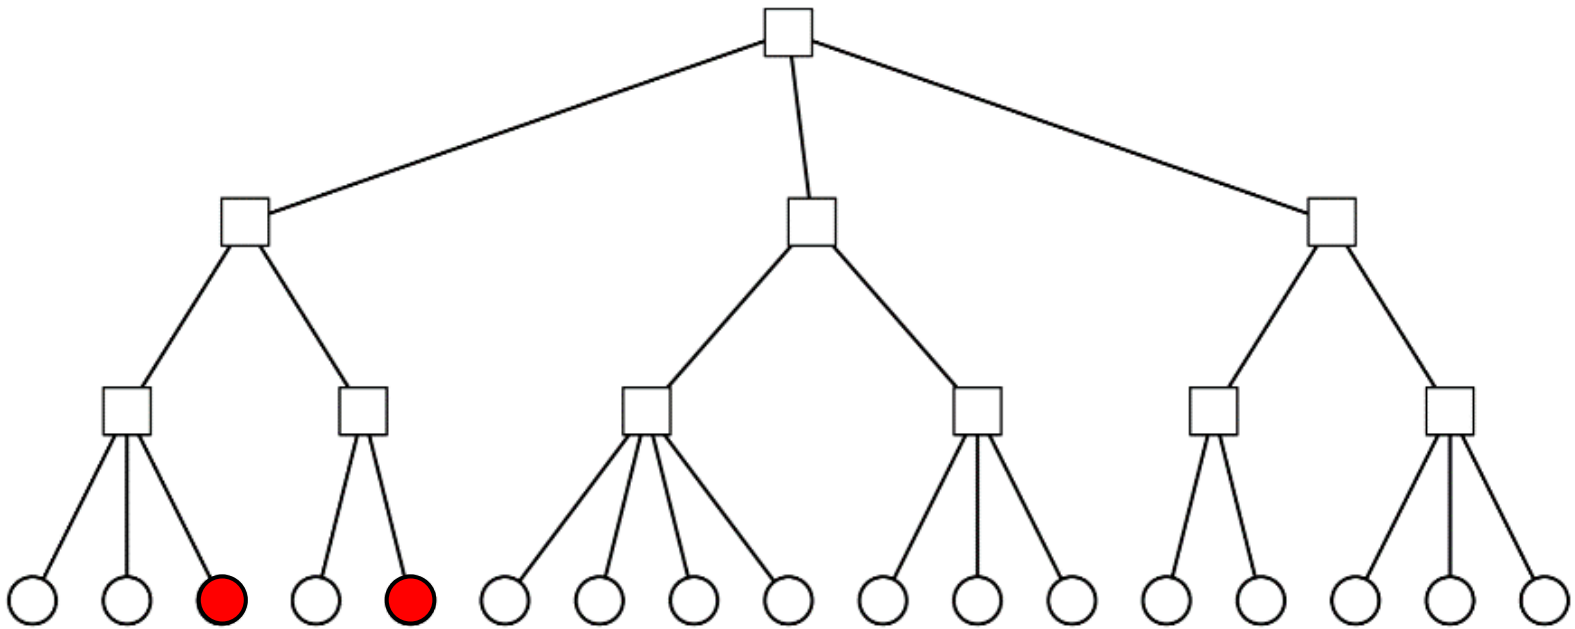

Algorithm

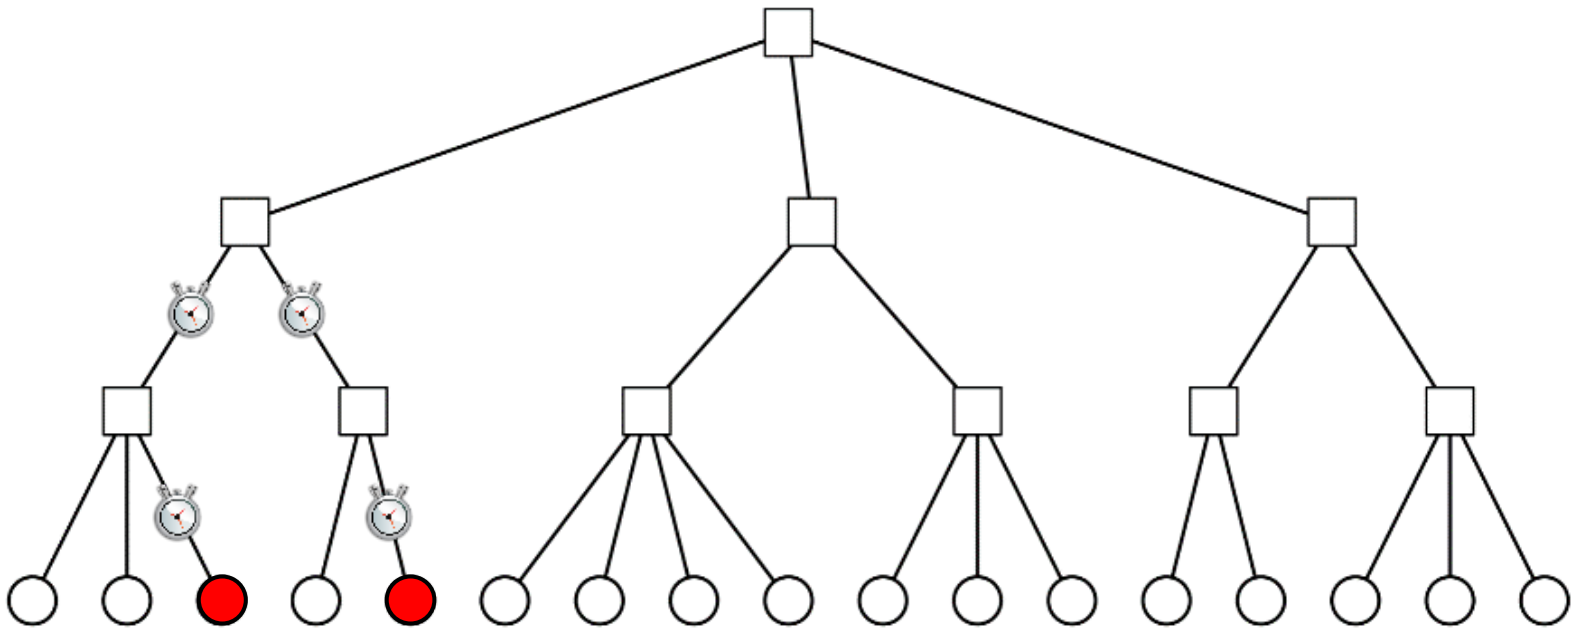

Algorithm

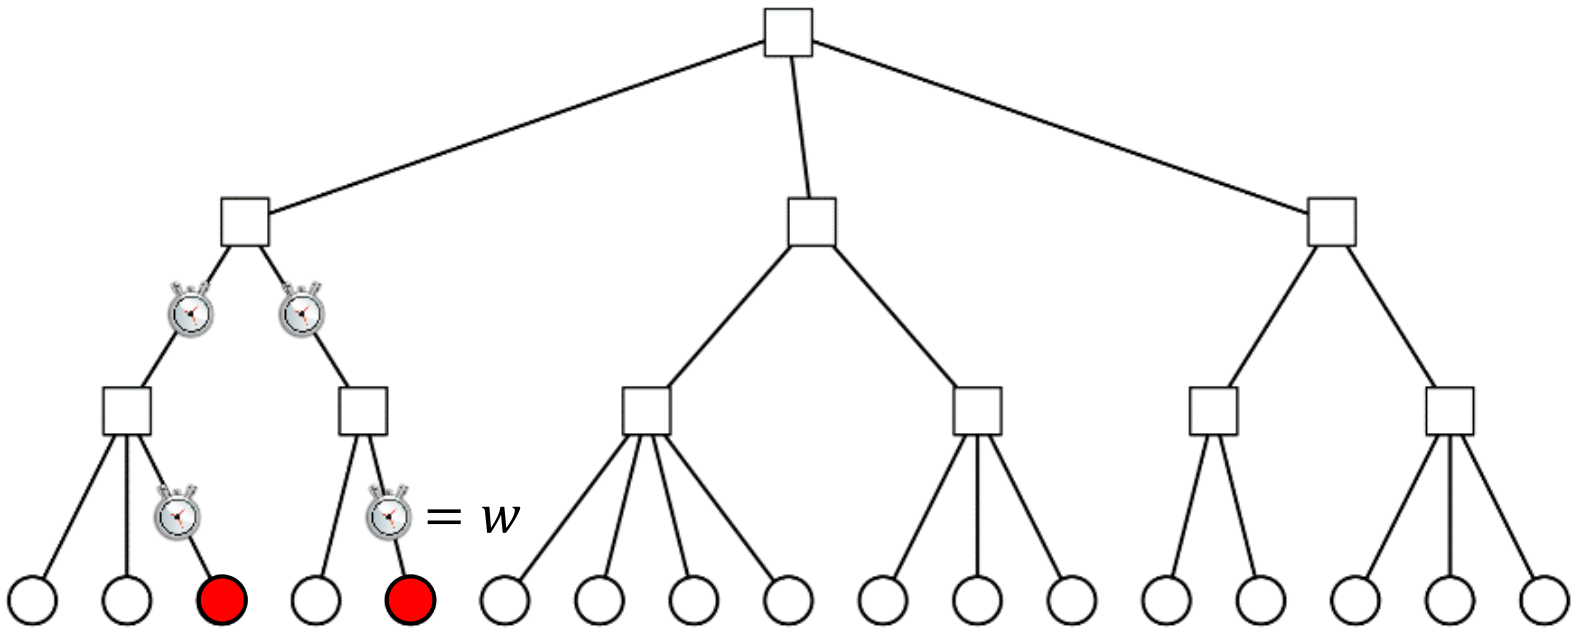

Algorithm

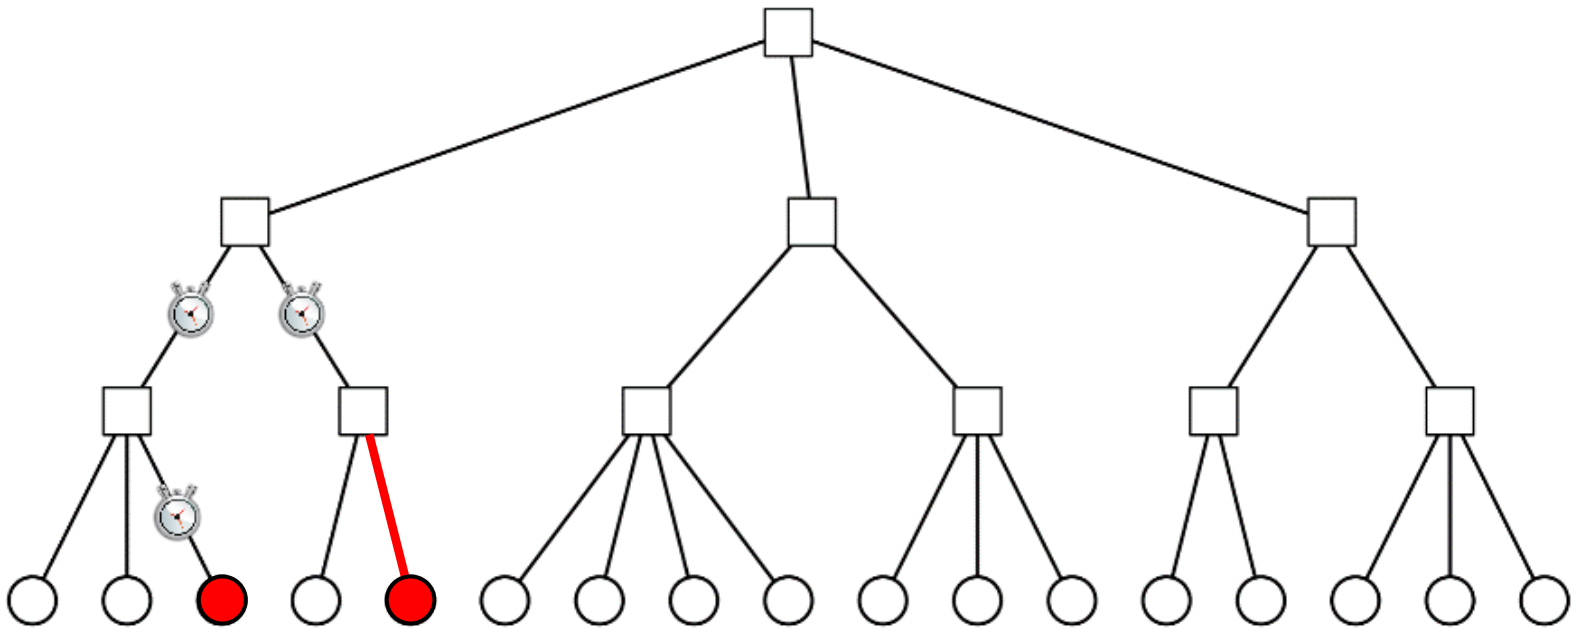

Algorithm

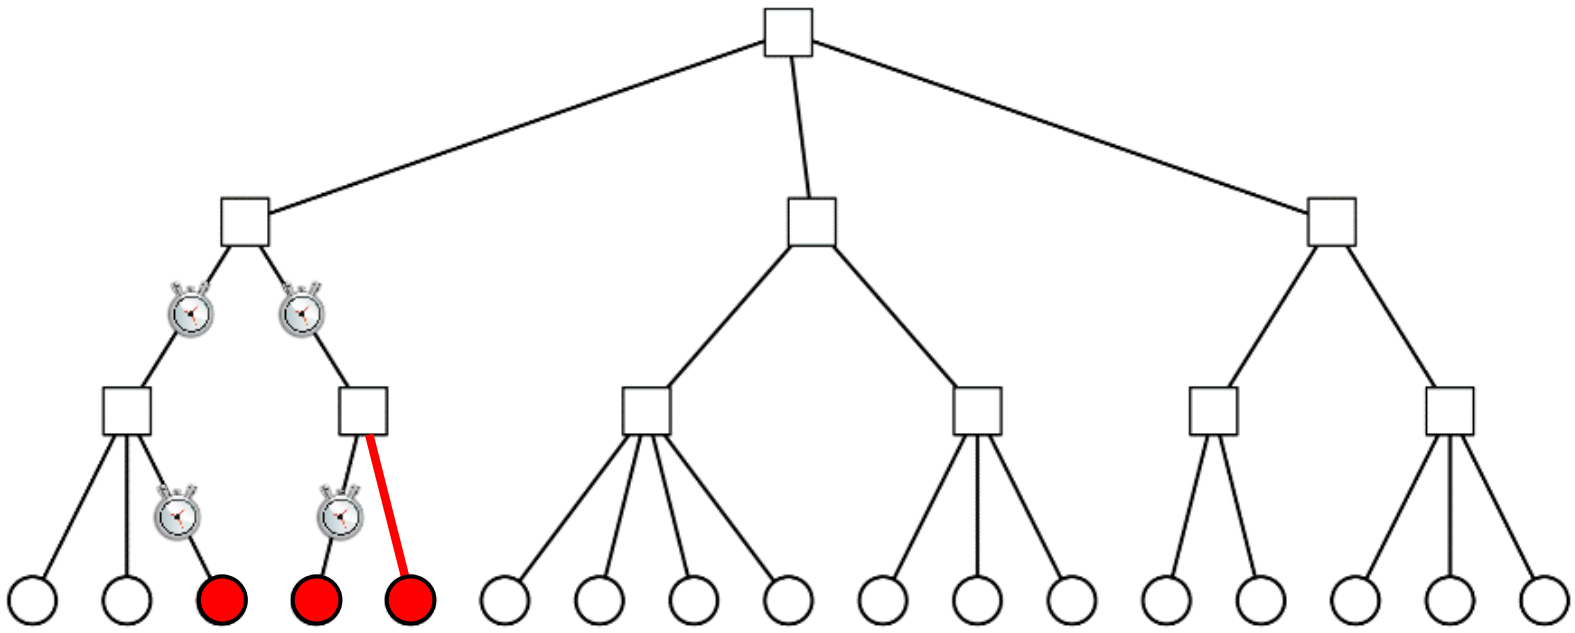

Algorithm

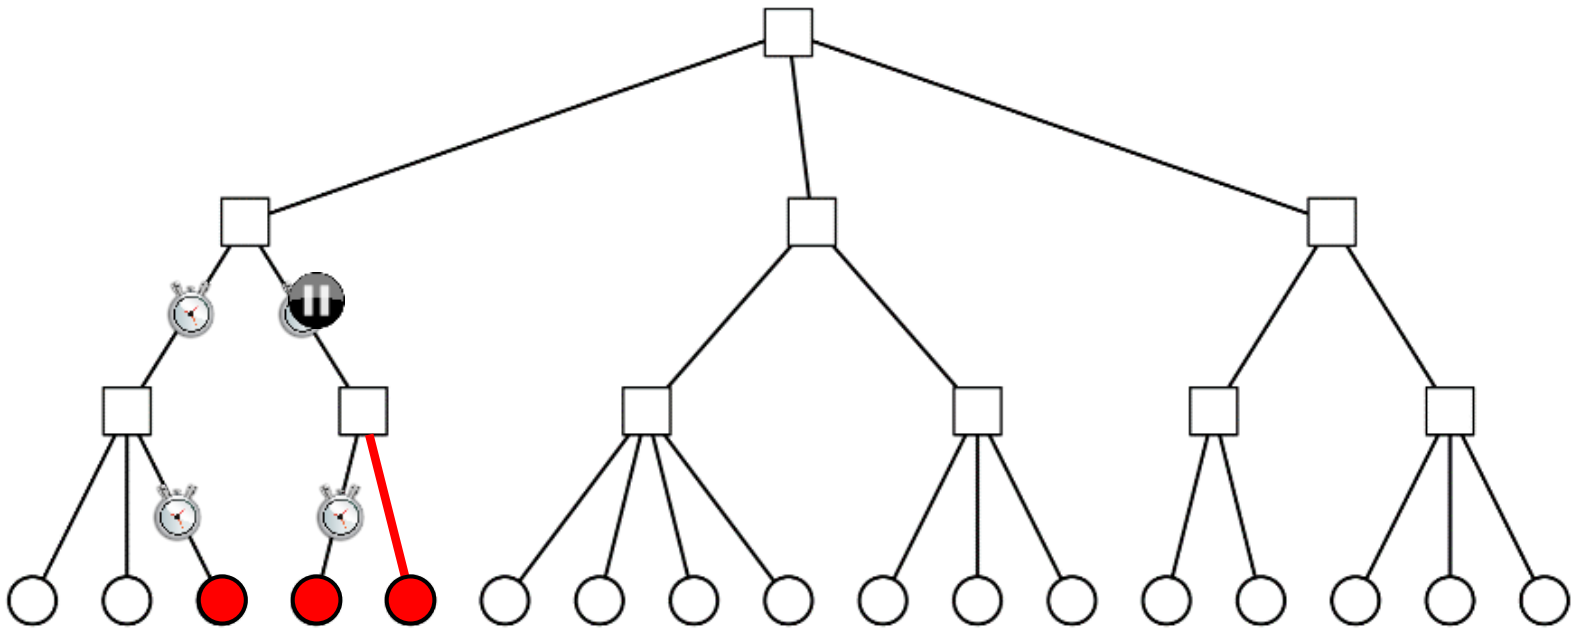

Algorithm

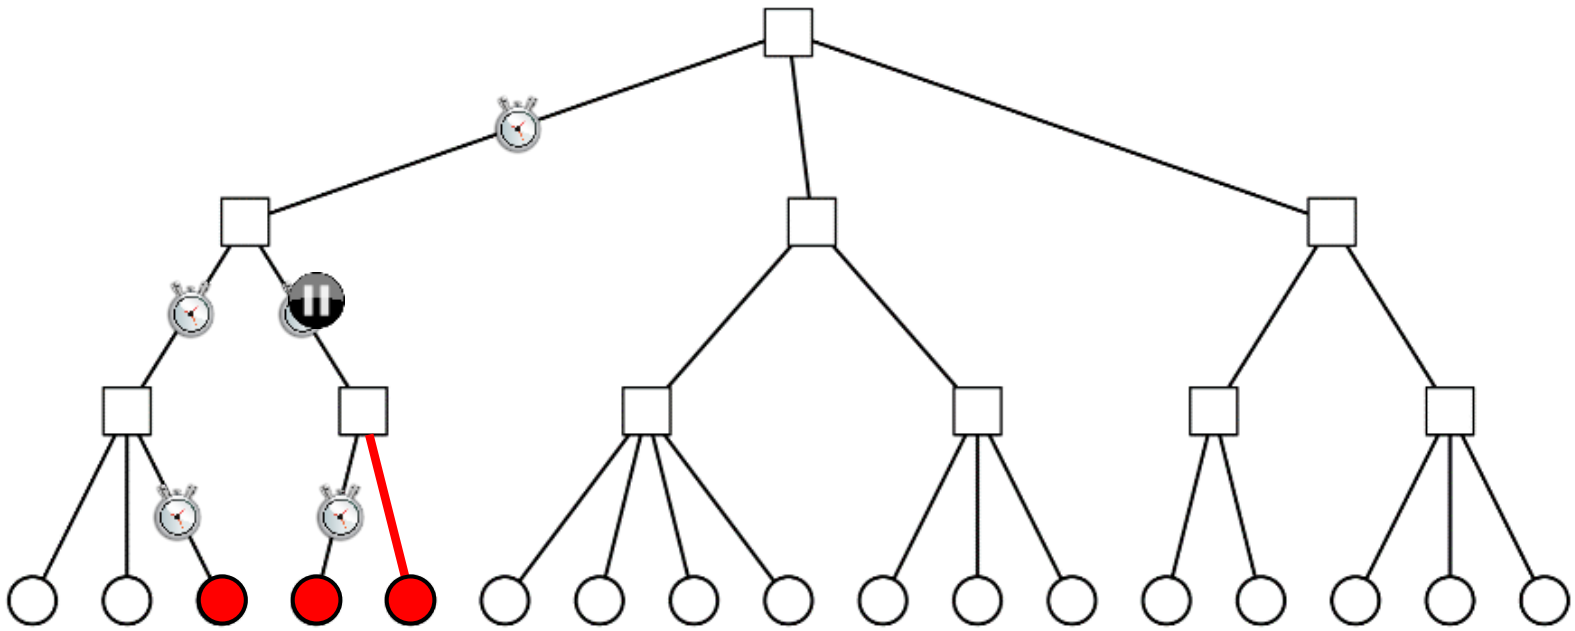

Algorithm

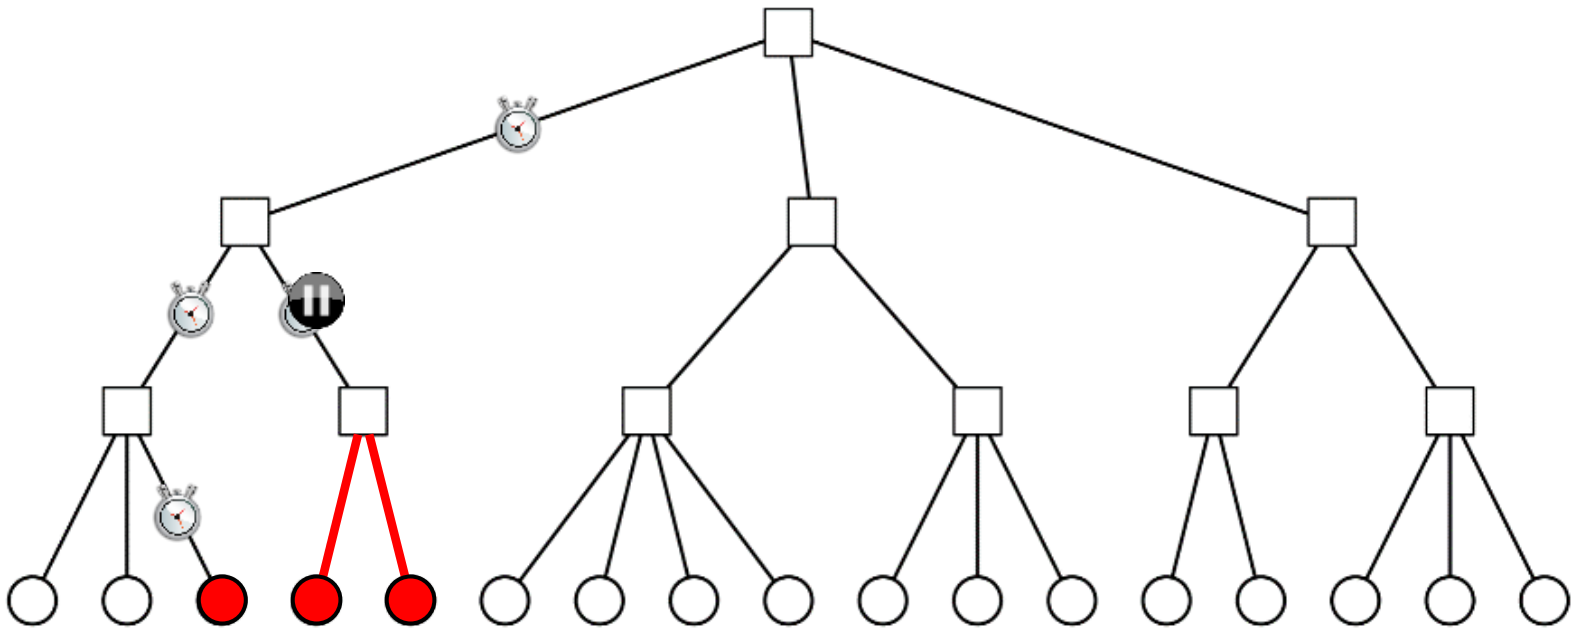

Algorithm

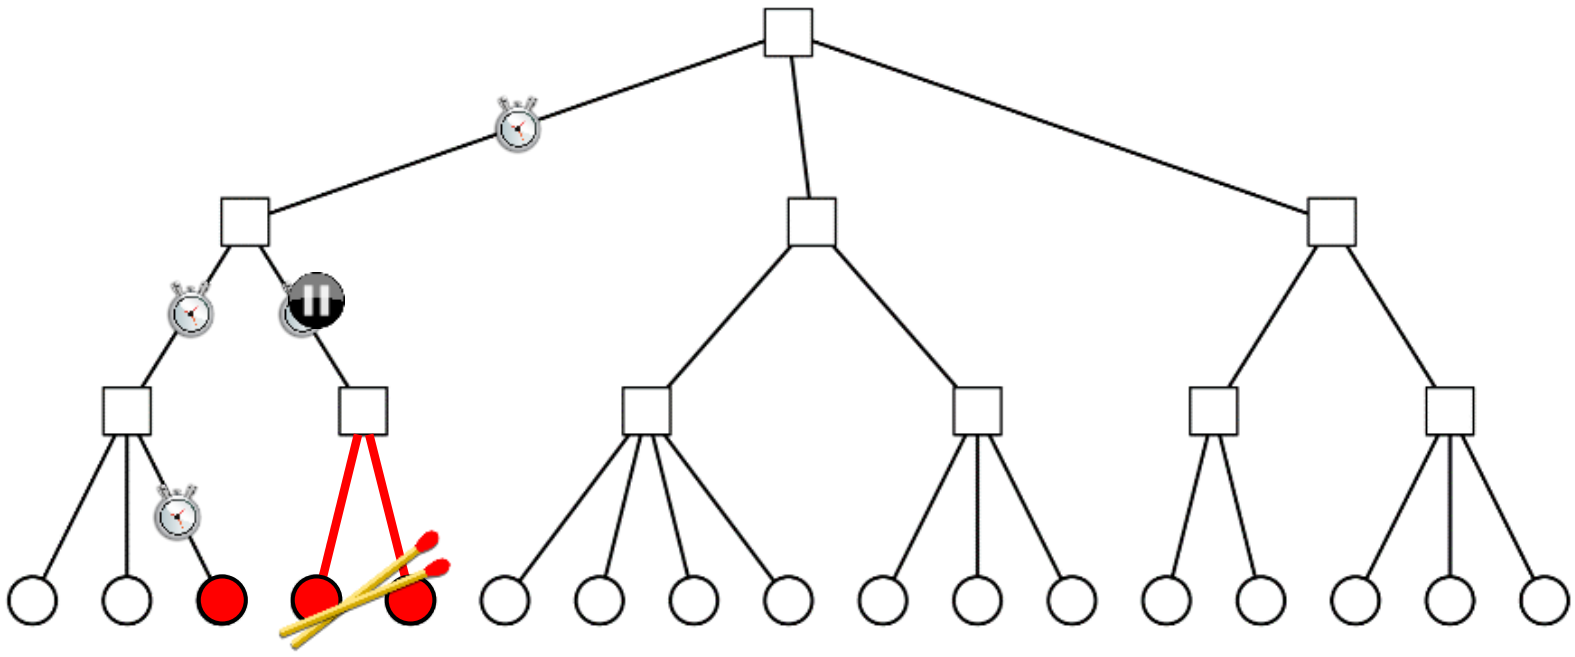

Algorithm

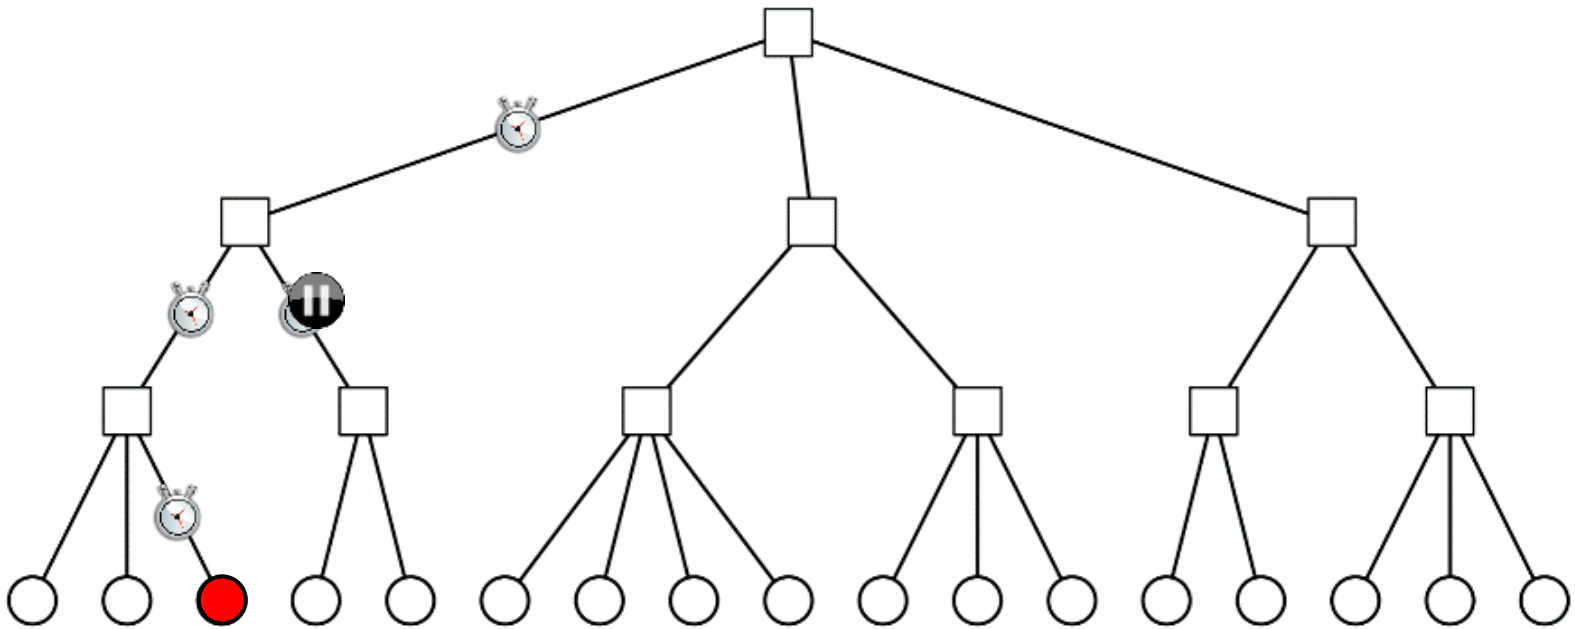

Proof

Proof

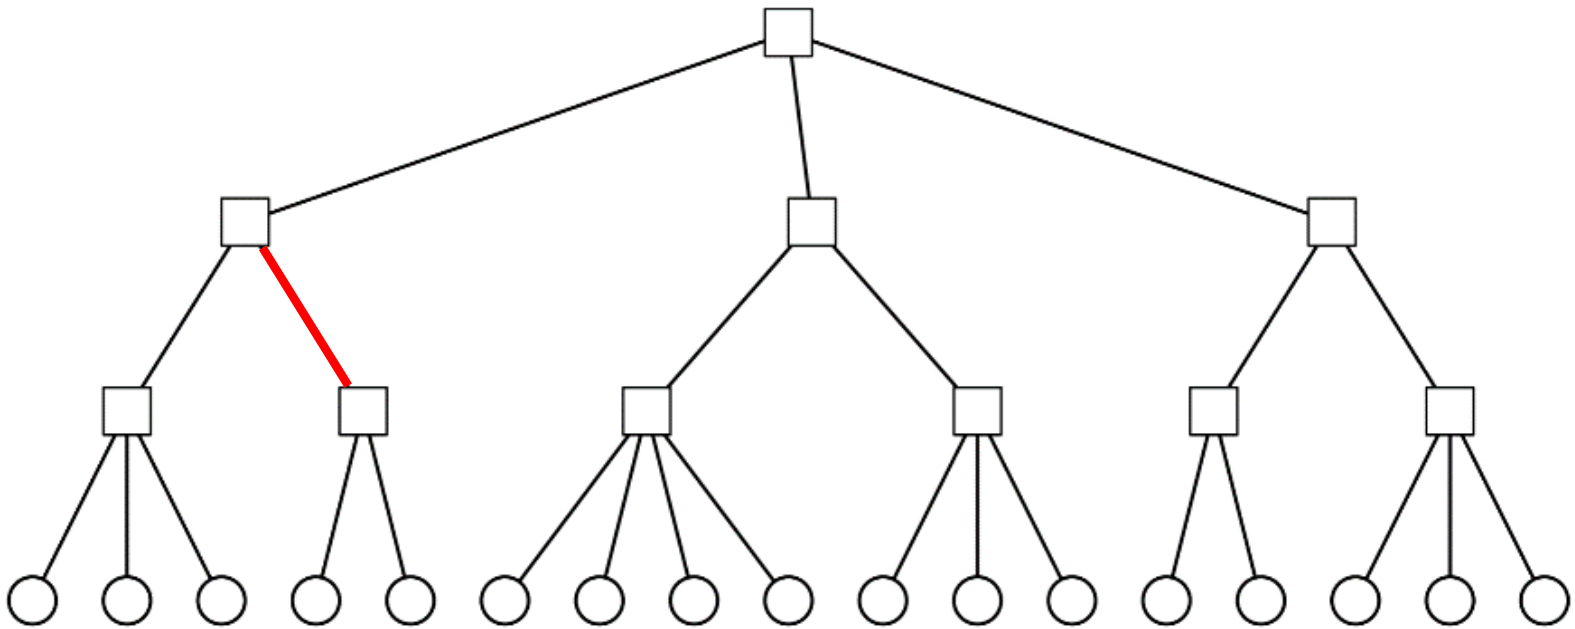

Proof

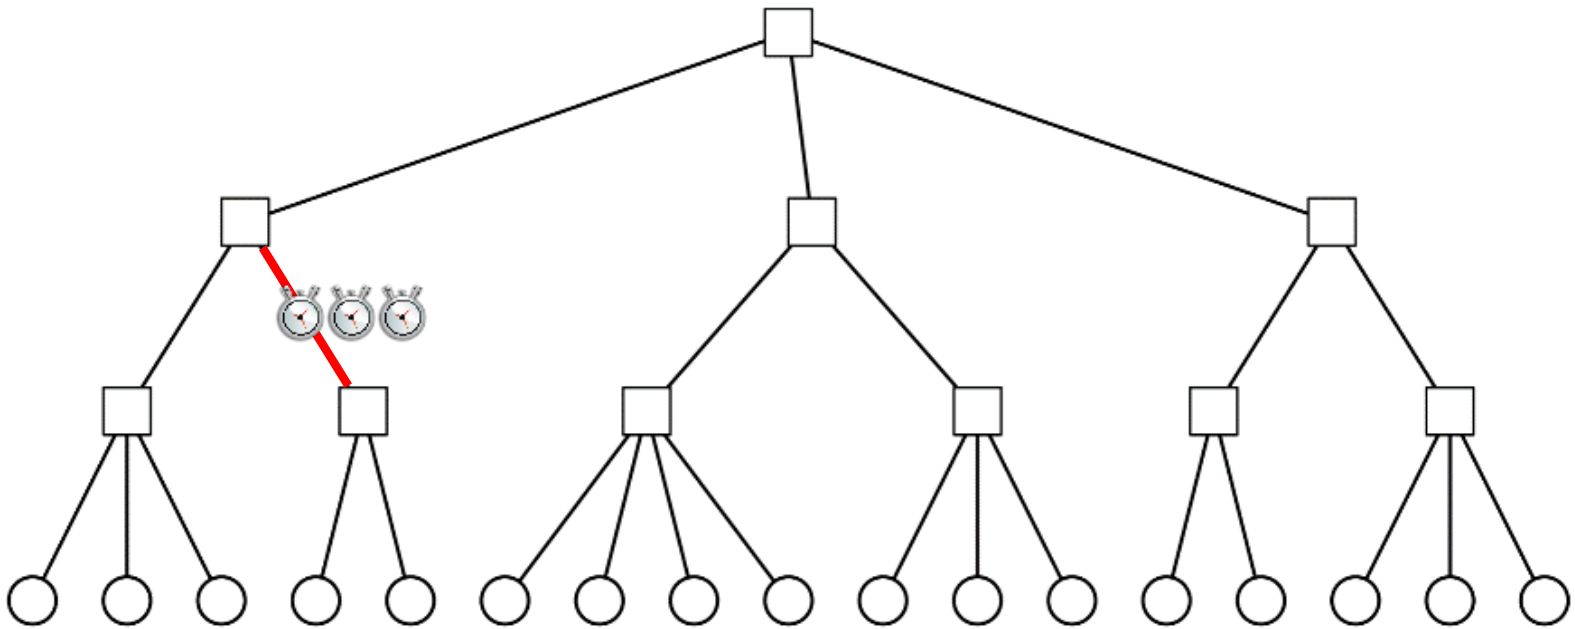

Proof

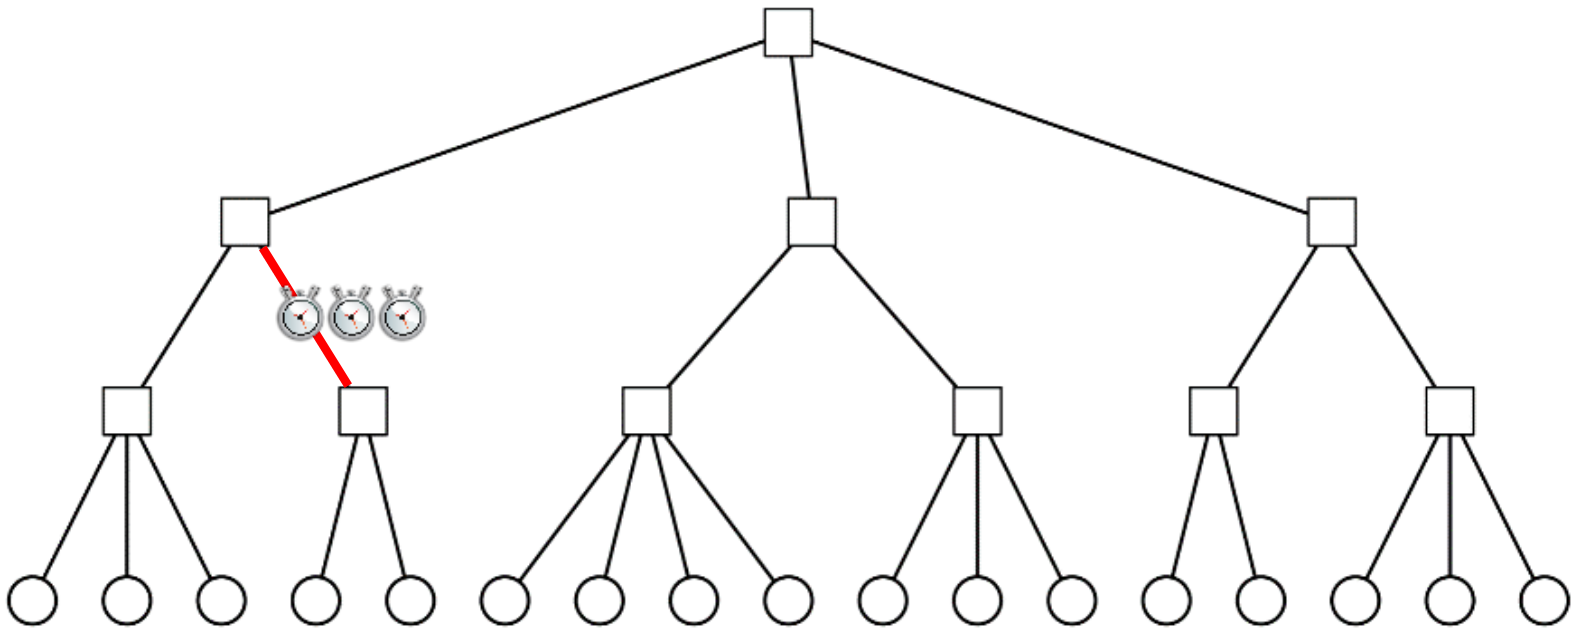

Total space cost = $\sum$ ⏱

Proof

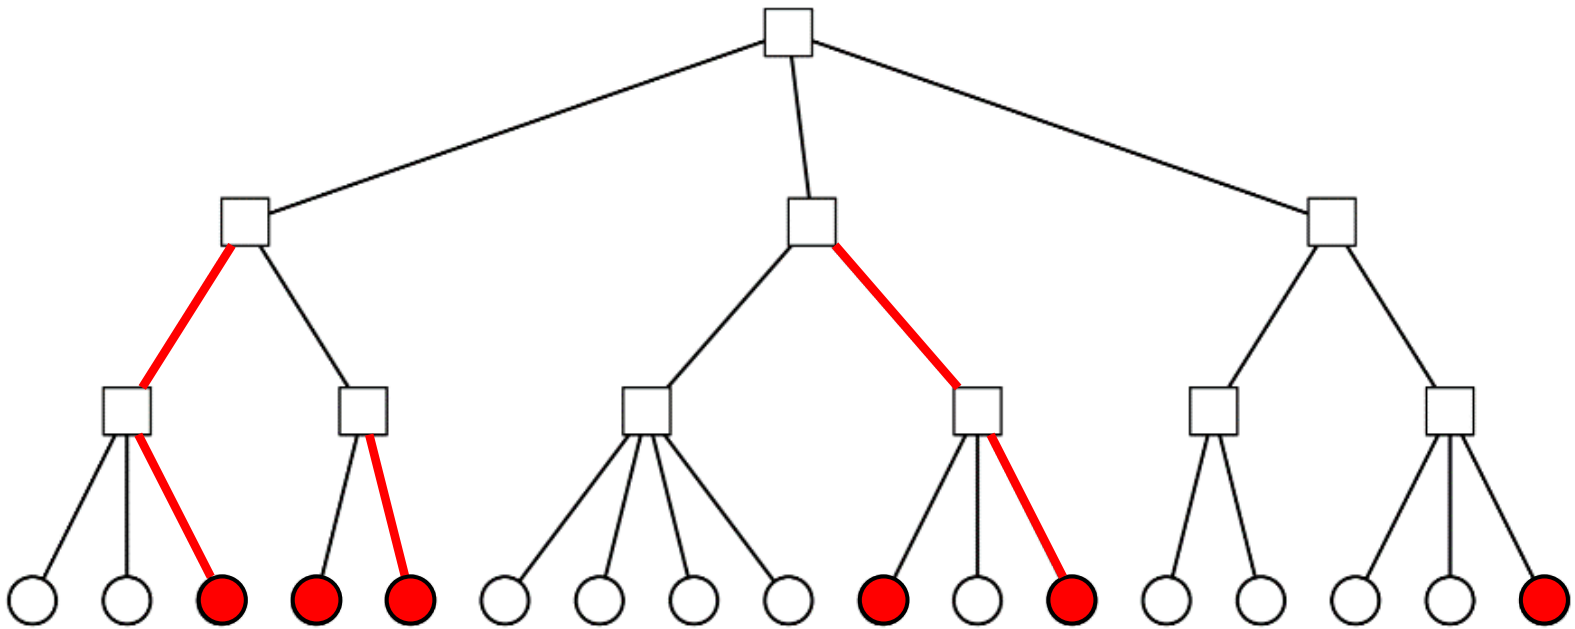

Proof

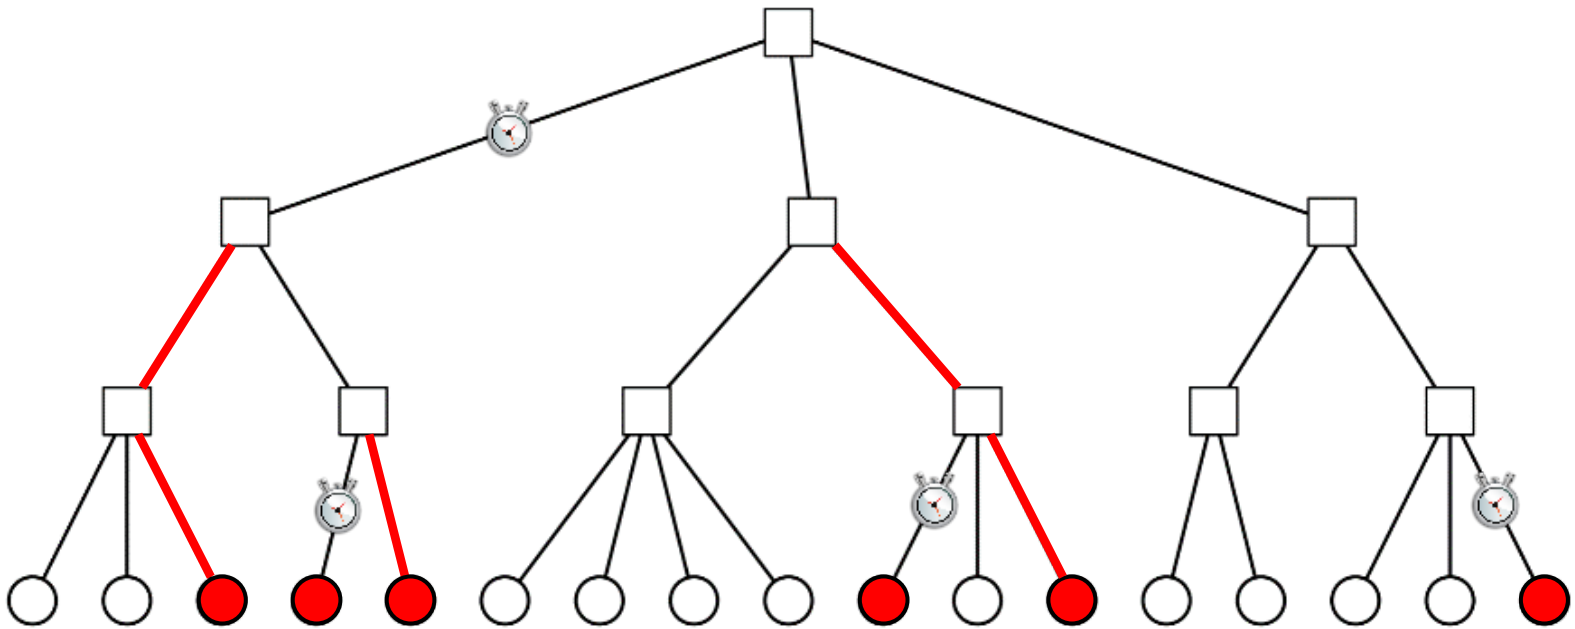

Proof

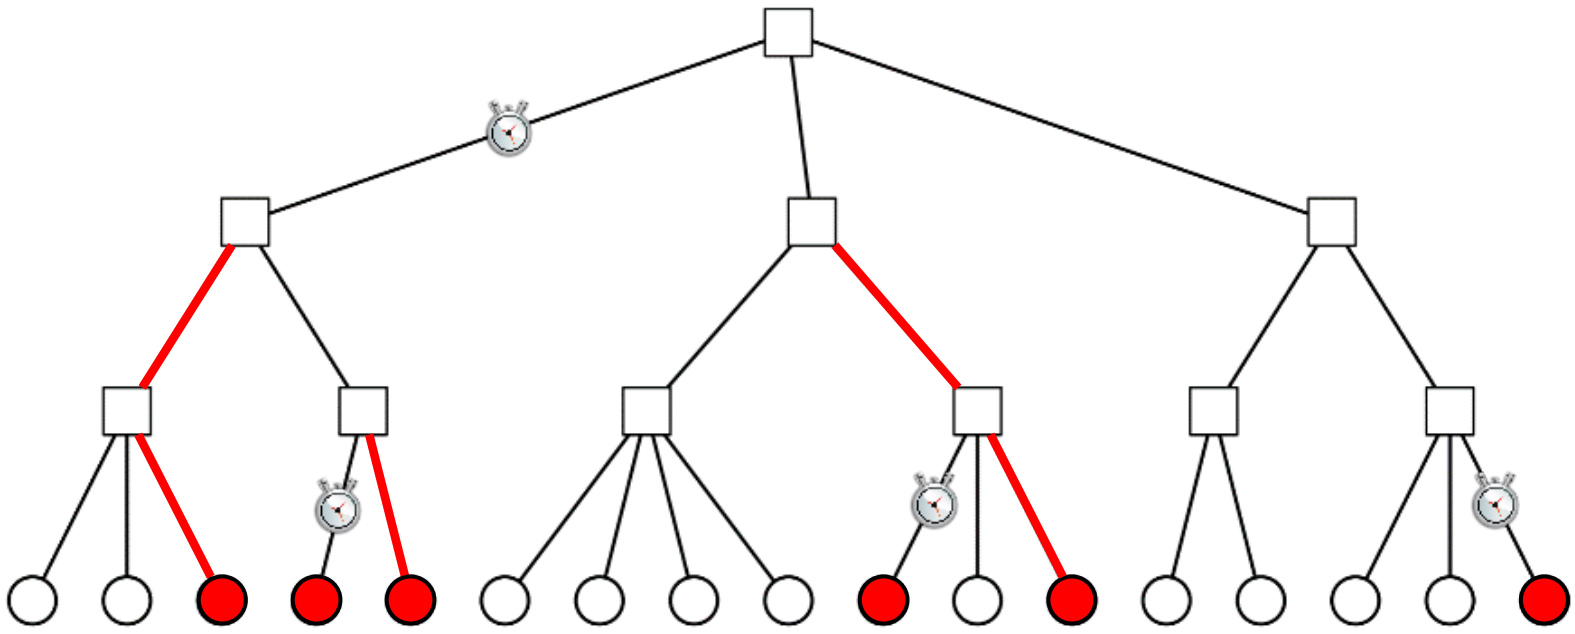

For each pair at least one timer running

Total time cost $\leq 2 \sum$ ⏱

Total Algorithm Cost = $O(\sum$  $)$

What about OPT?

Proof

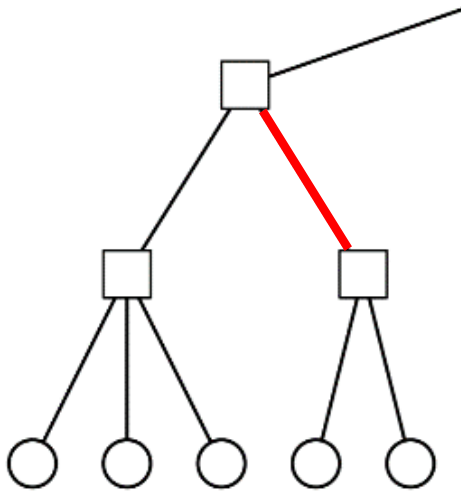

Proof

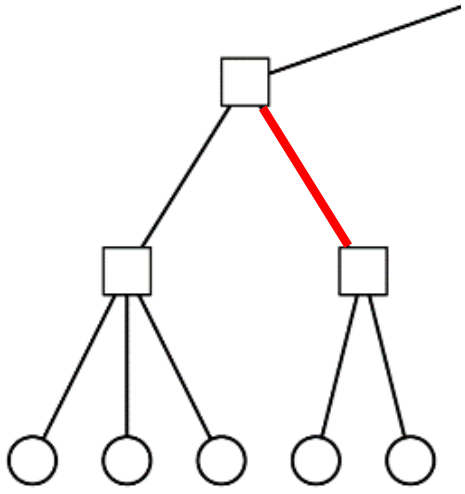

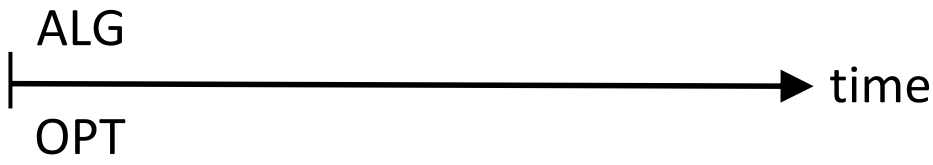

Proof

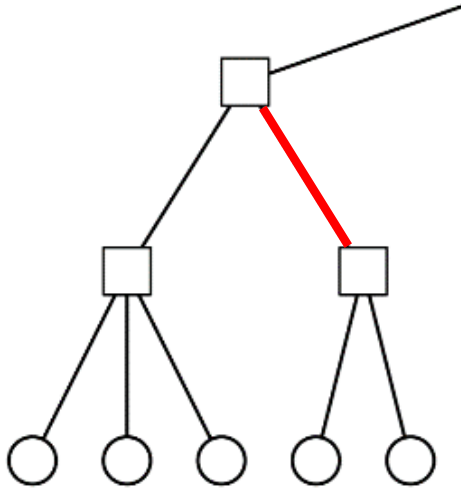

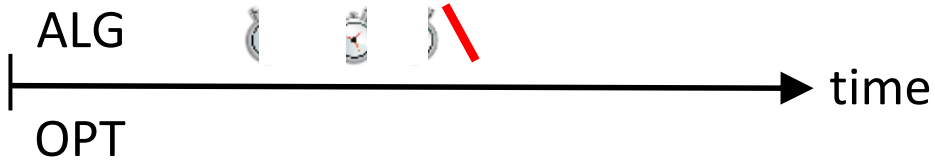

Proof

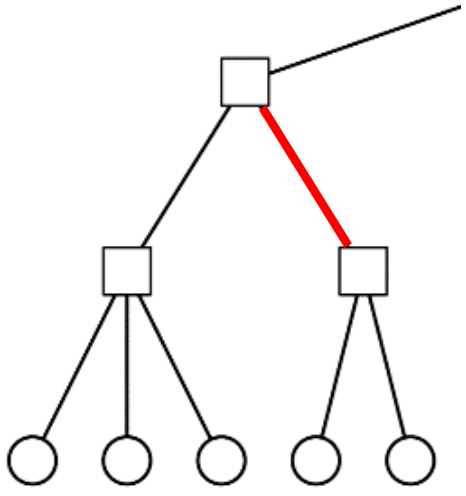

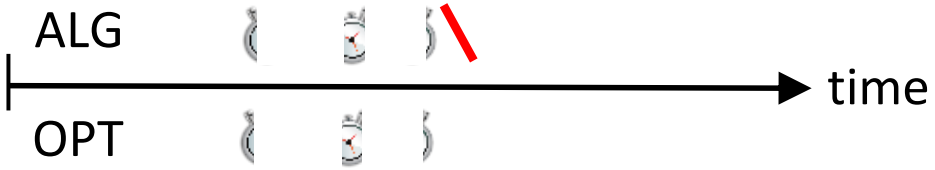

Proof

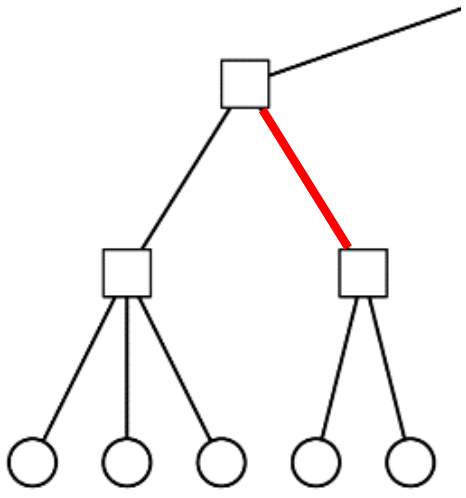

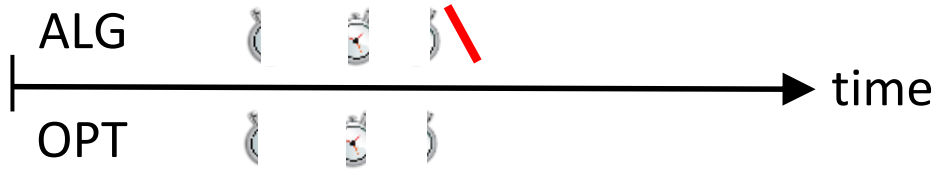

or

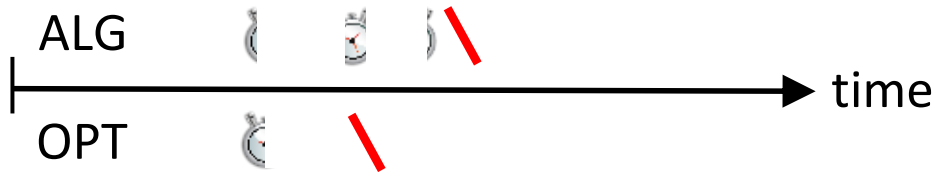

Proof

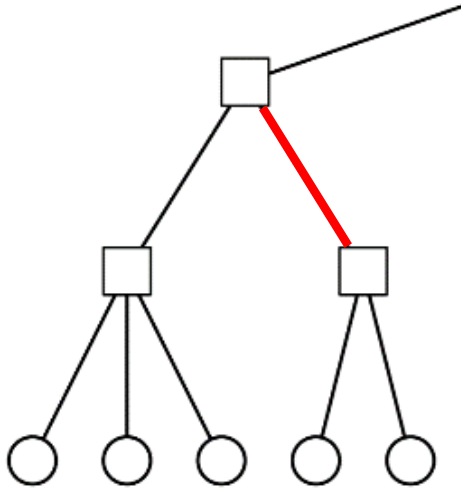

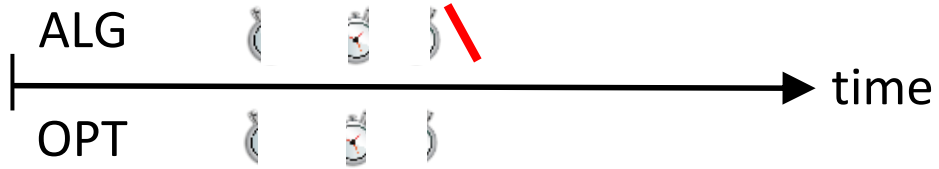

or

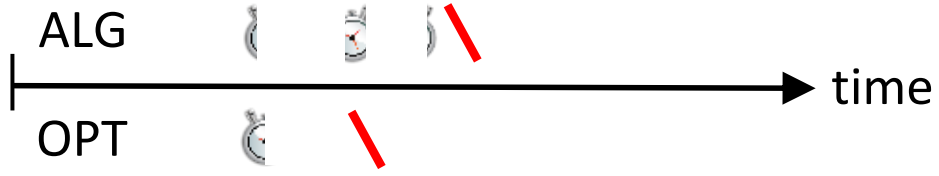

cost = cost

Done?

Just One Little Thing…

Proof

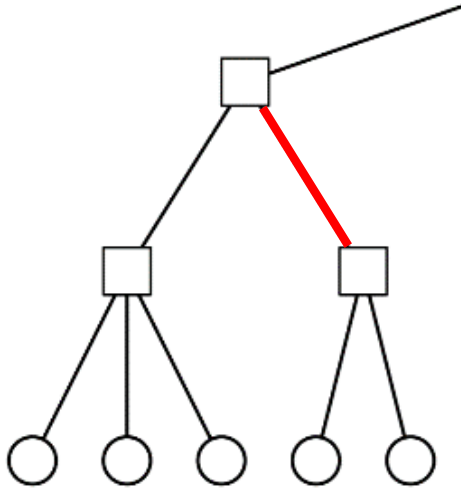

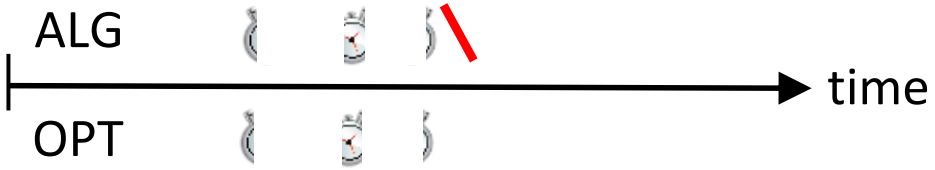

Proof

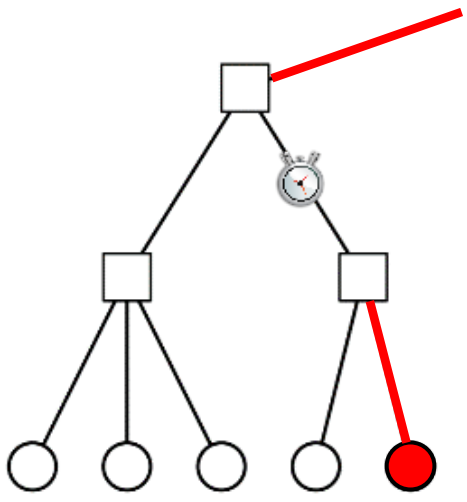

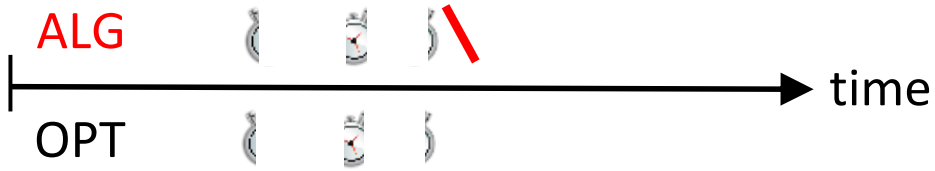

Proof

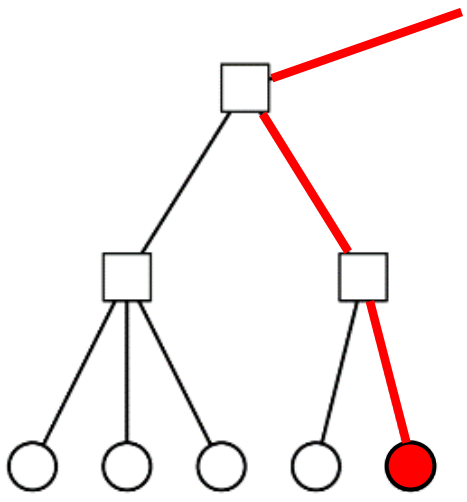

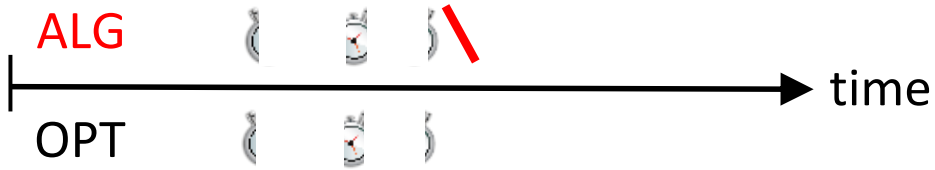

Proof

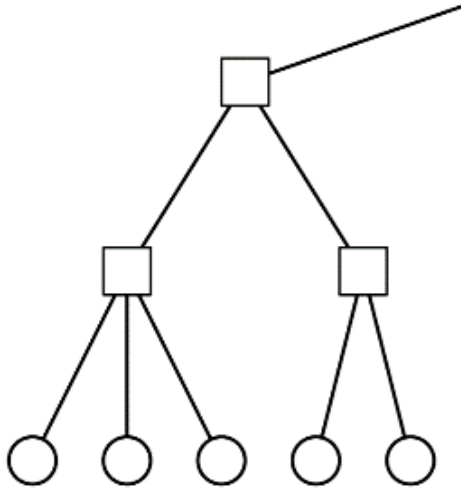

Proof

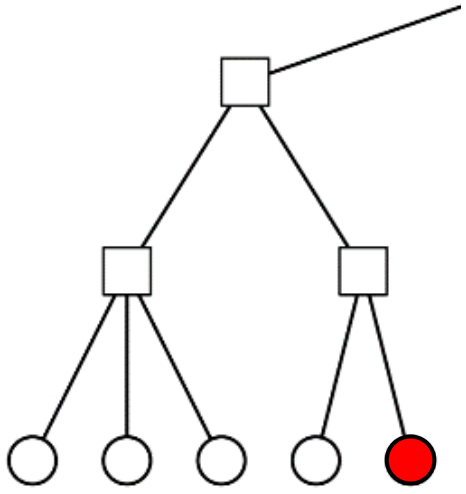

Proof

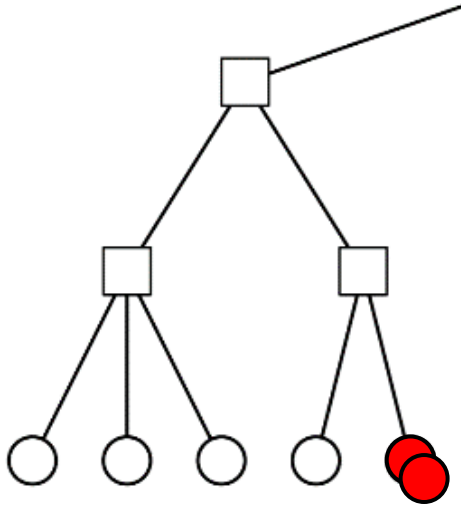

Proof

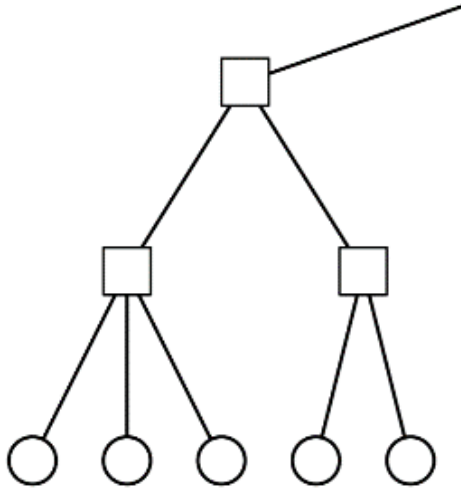

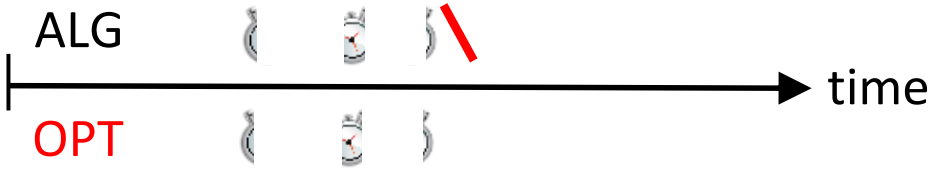

Proof

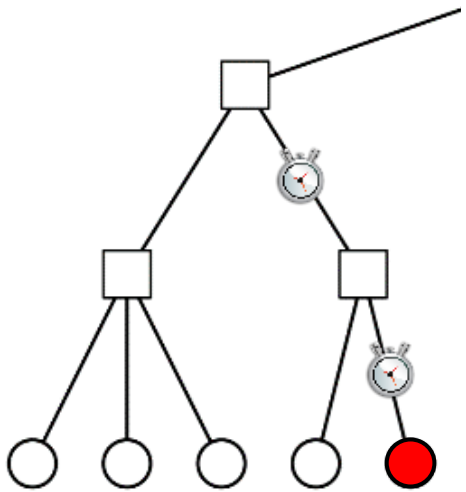

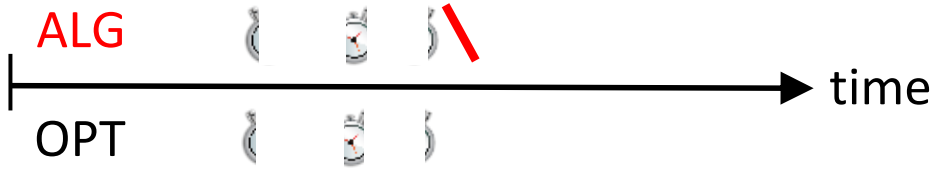

Proof

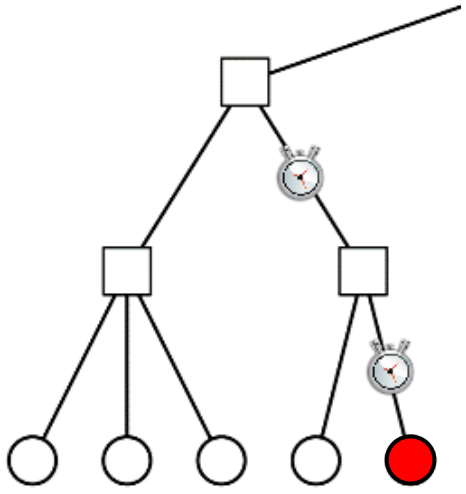

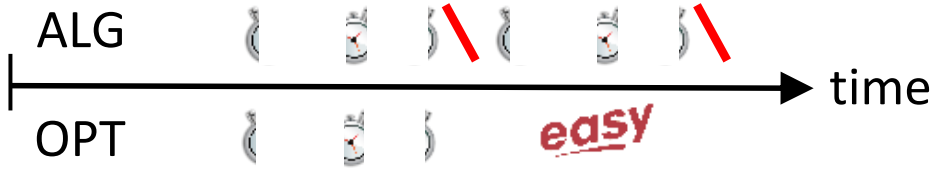

OPT has an easy time…

… but only every other phase!

Total OPT Cost = $\Omega(\sum$  $)$

Where is the $\log n$ coming from?

Height = $O(\log n)$ for time

$\mathrm{E}[\text{Distortion}] = O(\log n)$ for space


AND NOW FOR
SOMETHING
COMPLETELY
DIFFERENT

Algorithms against the Cloud

Will Blockchain Kill the Cloud?

SAP

Launch: Introducing Oracle Autonomous Blockchain Cloud Service

ORACLE Cloud

The blockchain is here to make cloud computing better

Information Age

Why Blockchain is Cloud 2.0

COINTELEGRAPH
The future of money

BlockCloud: Re-inventing Cloud with Blockchains

guardtime

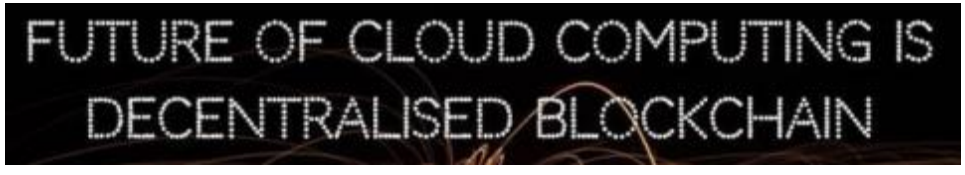

TECHSUTRAM


STEVE FORBES
Chairman, Forbes Media
CURRENCY OF THE FUTURE?

2008

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 network timestamps transactions by hashing them into an ongoing chain of hash-based proof-of-work, forming a record that cannot be changed without redoing the proof-of-work. The longest chain not only serves as proof of the sequence of events witnessed, but proof that it came from the largest pool of CPU power. As long as a majority of CPU power is controlled by nodes that are not cooperating to attack the network, they'll generate the longest chain and outpace attackers. The network itself requires minimal structure. Messages are broadcast on a best effort

Blockchain

Figure 9-3 Manual Journal Voucher.

Page 1 of 1

MANUAL JOURNAL VOUCHER

PREPARED BY	DATE
WLR	2/2/X5
APPROVED	DATE

Batch 1101	Batch Line 9	Total Amount 11,200.20	
Description ACCRUED INTEREST INCOME		Effective Date 1/31/X5	Type A
Reference J43-JAN INTEREST		Accounting Company 10 - CORPORATE	

Seq.	Account Number	Description	Debit Amount	Credit Amount
01	1280-000	INTEREST RECEIVABLE	11,200.20	
02	8050-010	FIRST NATIONAL - CD		1,330.10
03	8050-020	MUNICIPAL BONDS		6,220.80
04	8050-010	OTHER INVESTMENTS		3,649.30

FinTech developers and managers understand that the *blockchain* has the potential to disrupt the financial world. The blockchain allows the participants of a distributed system to agree on a common view of the system, to track changes in the system, in a reliable way. In the distributed systems community, agreement techniques have been known long before cryptocurrencies such as Bitcoin (where the term blockchain is borrowed) emerged. Various concepts and protocols exist, each with its own advantages and disadvantages. This book introduces the basic techniques when building fault-tolerant distributed systems, in a *scientific* way. We will present different protocols and algorithms that allow for fault-tolerant operation, and we will discuss practical systems that implement these techniques.

About the author

Roger Wattenhofer is a professor at ETH Zurich. Before joining ETH Zurich, he was at Brown University and Microsoft Research. His research interests include fault-tolerant distributed systems, efficient network algorithms, and cryptocurrencies such as Bitcoin. He has published more than 250 scientific articles.

Inverted Forest Publishing
First Edition, 2016
ISBN-13 978-1522751830
ISBN-10 1522751831

ISBN 9781522751830
90000 >
9 781522 751830

Blockchain Basics

Transaction

Transaction

Transaction

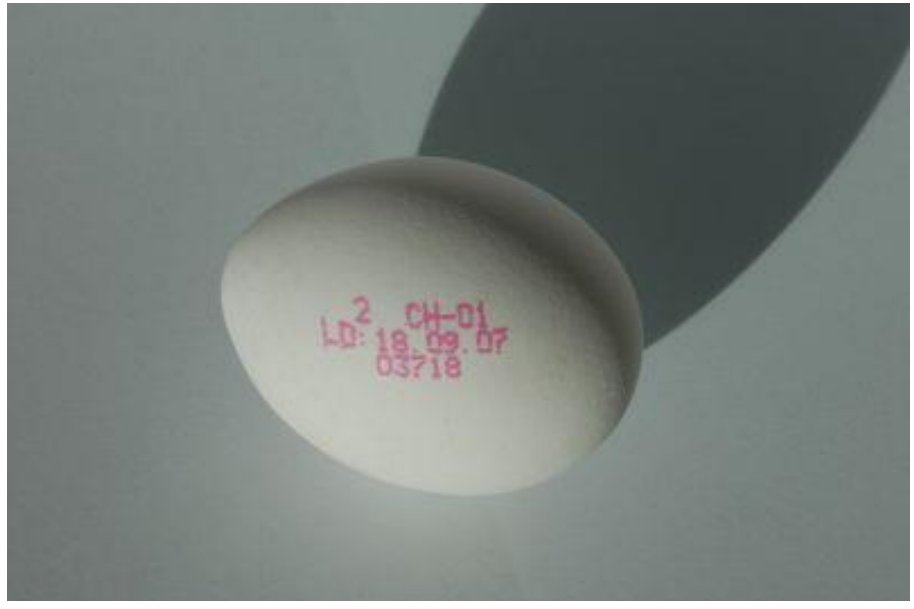

Transaction

Block

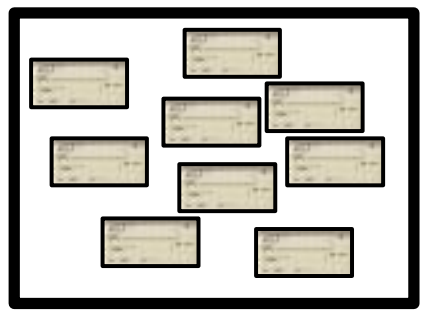

Blockchain

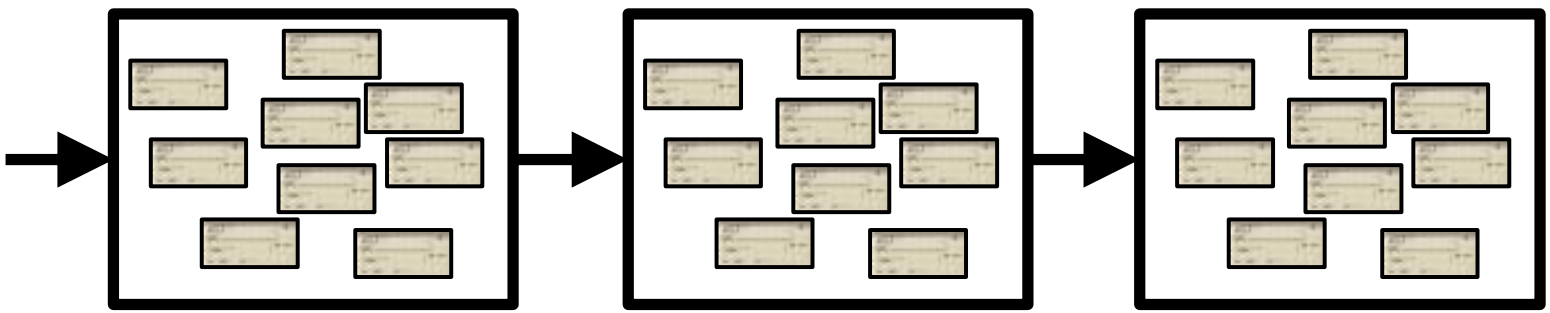

Blockchain is Replicated

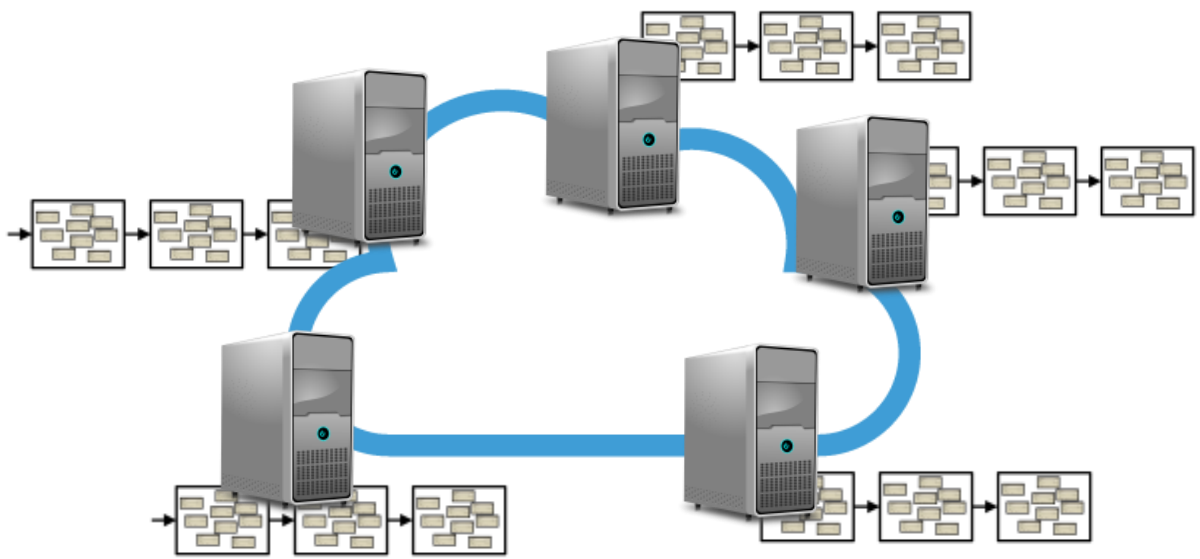

Blockchain

Distributed Systems (1982) & Cryptography (1976)

Blockchain

Distributed Systems & Cryptography

Fault-Tolerance & Digital Signatures

Rule of Thumb

Blockchains* may disrupt your business if you use signatures.

*or blockchain-like tech

Blockchain Variants

Ledger of Transactions

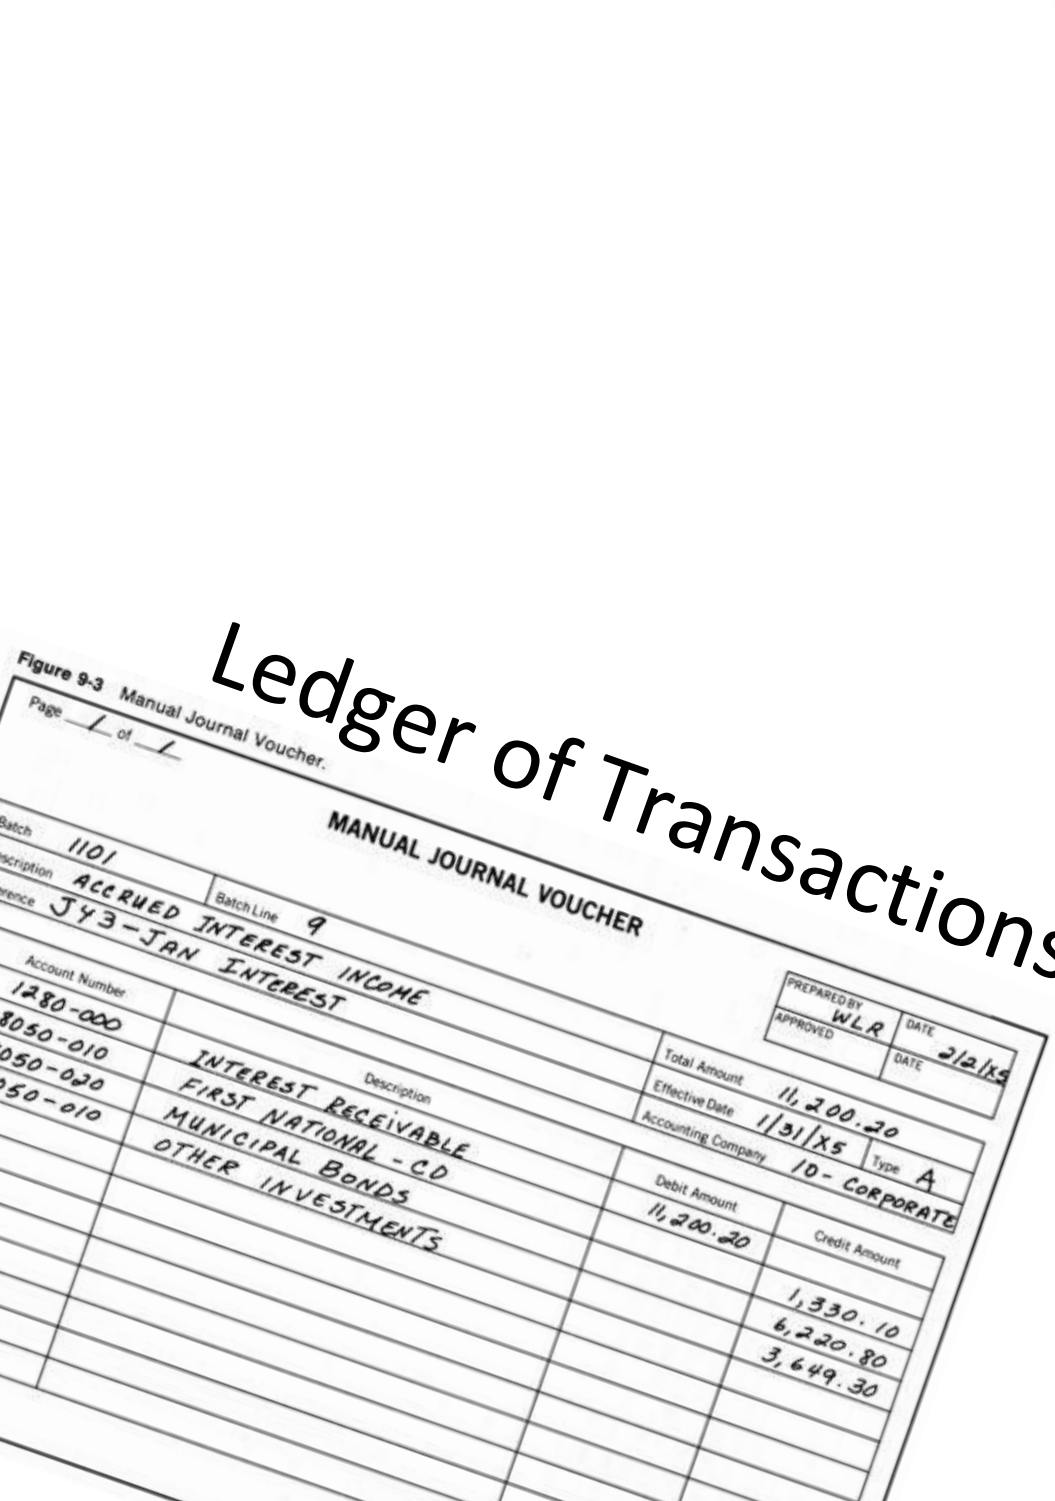

Figure 9-3 Manual Journal Voucher.

Page 1 of 1

MANUAL JOURNAL VOUCHER

Batch 1101 | Batch Line 9

Description ACCRUED INTEREST INCOME

Reference JY3-JAN INTEREST

PREPARED BY WLR | DATE 2/2/X5

APPROVED | DATE

Total Amount 11,200.20

Effective Date 1/31/X5 | Type A

Accounting Company 10 - CORPORATE

Account Number	Description	Debit Amount	Credit Amount
1280-000	INTEREST RECEIVABLE	11,200.20	
8050-010	FIRST NATIONAL - CD		1,330.10
050-030	MUNICIPAL BONDS		6,220.80
050-010	OTHER INVESTMENTS		3,649.30

Bitcoin

Permissionless / Open

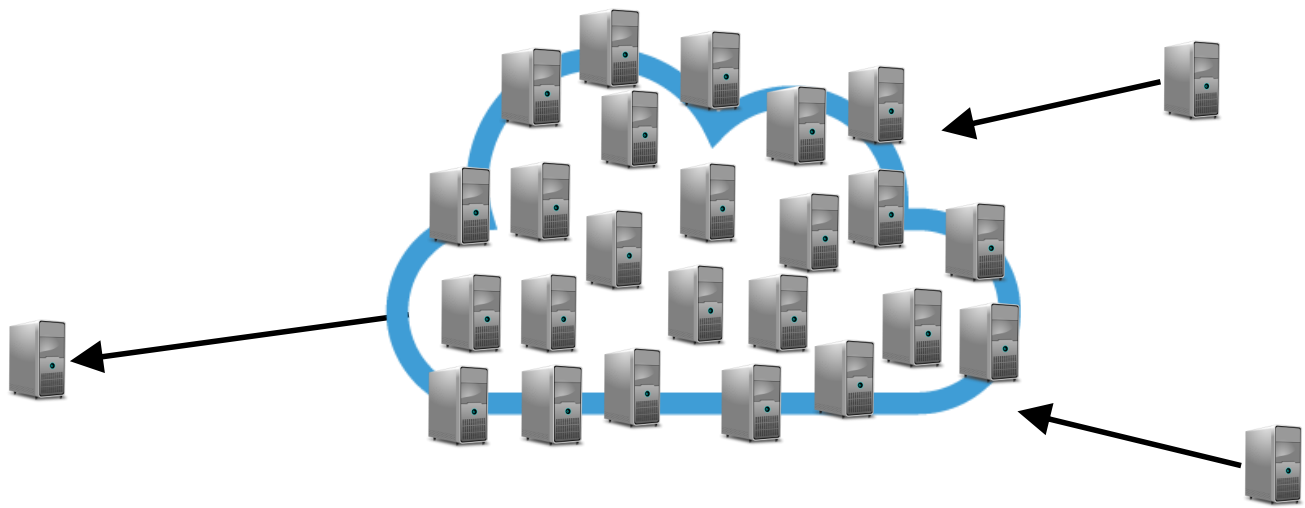

Permissioned / Closed

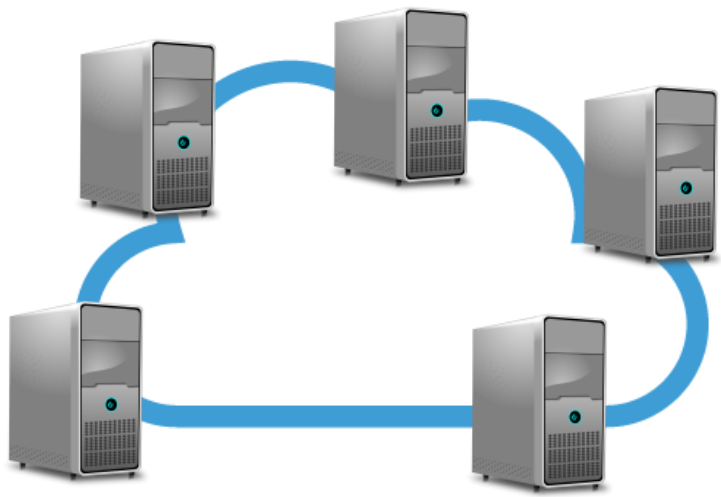

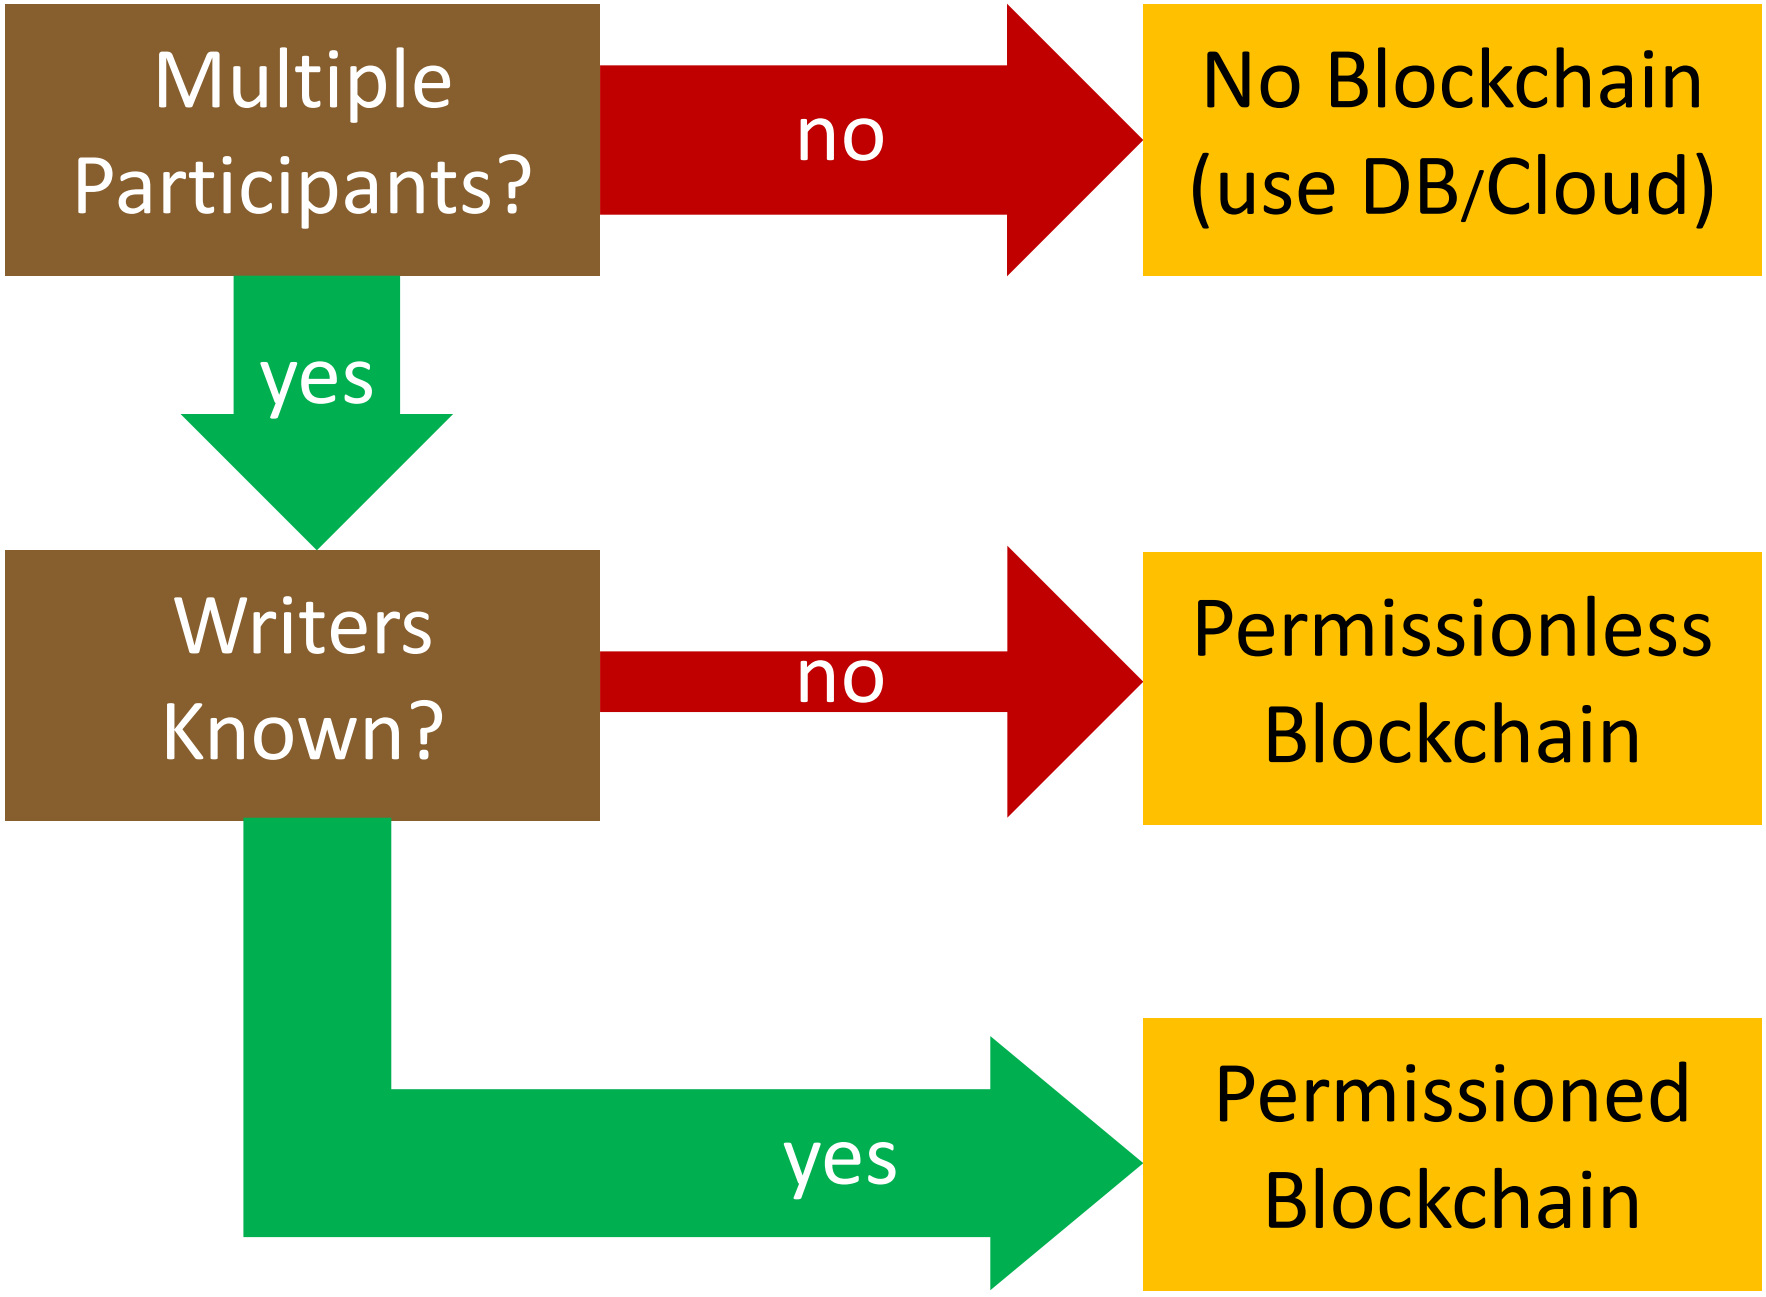
Multiple Participants?
no
No Blockchain (use DB/Cloud)
yes
Writers Known?
no
Permissionless Blockchain
yes
Permissioned Blockchain

The Seven Blockchain Dimensions

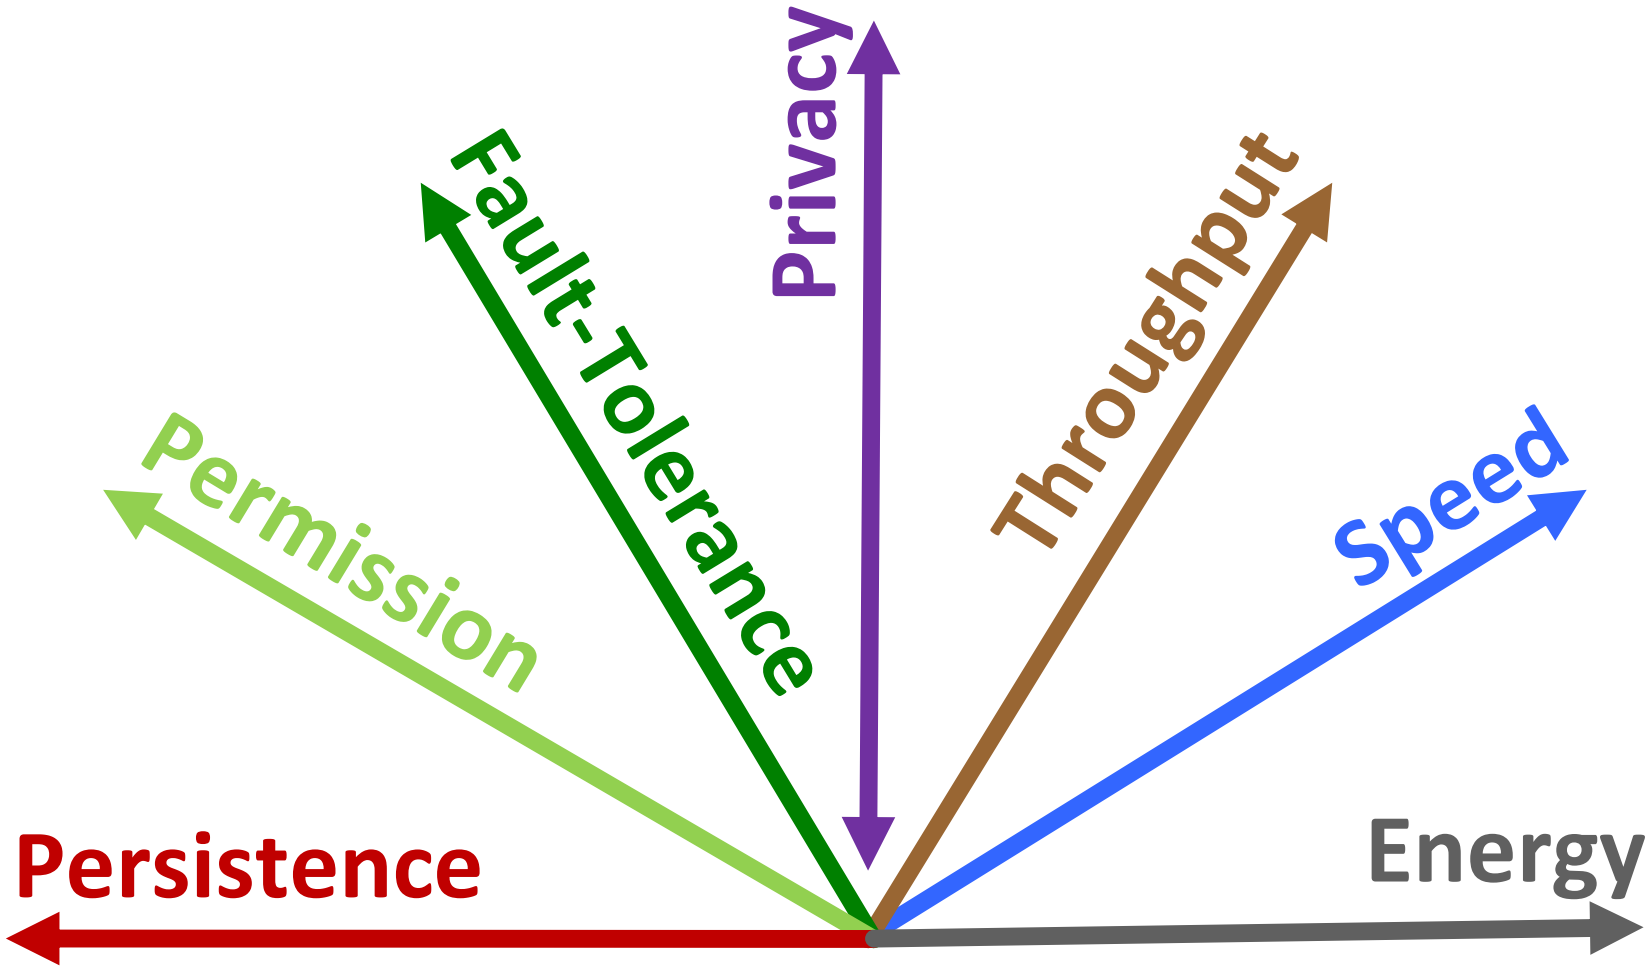

Blockchain

Persistence

Database

↓

Immutable

↓

Provable 

Fault-Tolerance

Correct

↓

Crash

↓

Byzantine 

Blockchain

Speed

1 hour 

↓

1 minute

↓

1 second

Throughput

10 tx/s 

↓

10k tx/s

↓

10m tx/s

Blockchain

Scalability

10 nodes

↓

100 nodes

↓

1000 nodes 

Energy Consumption

«Ich wäre nicht überrascht, wenn Bitcoin verboten würde»

ETH-Informationstechnologe Roger Wattenhofer über den Energiebedarf der Kryptowährung und bessere Alternativen

Prof. Dr. Roger Wattenhofer vom Departement Informationstechnologie und Elektrotechnik der ETH Zürich

Economic Incentives

Market / Energy Value ≈ 12 GW

$1M/h $0.08/kWh

Proof of Work

Hashrate · Energy/Hash ≈ 1.3 GW

$13 \cdot 10^9$ GH/s 0.1 J/GH

The Seven Blockchain Dimensions

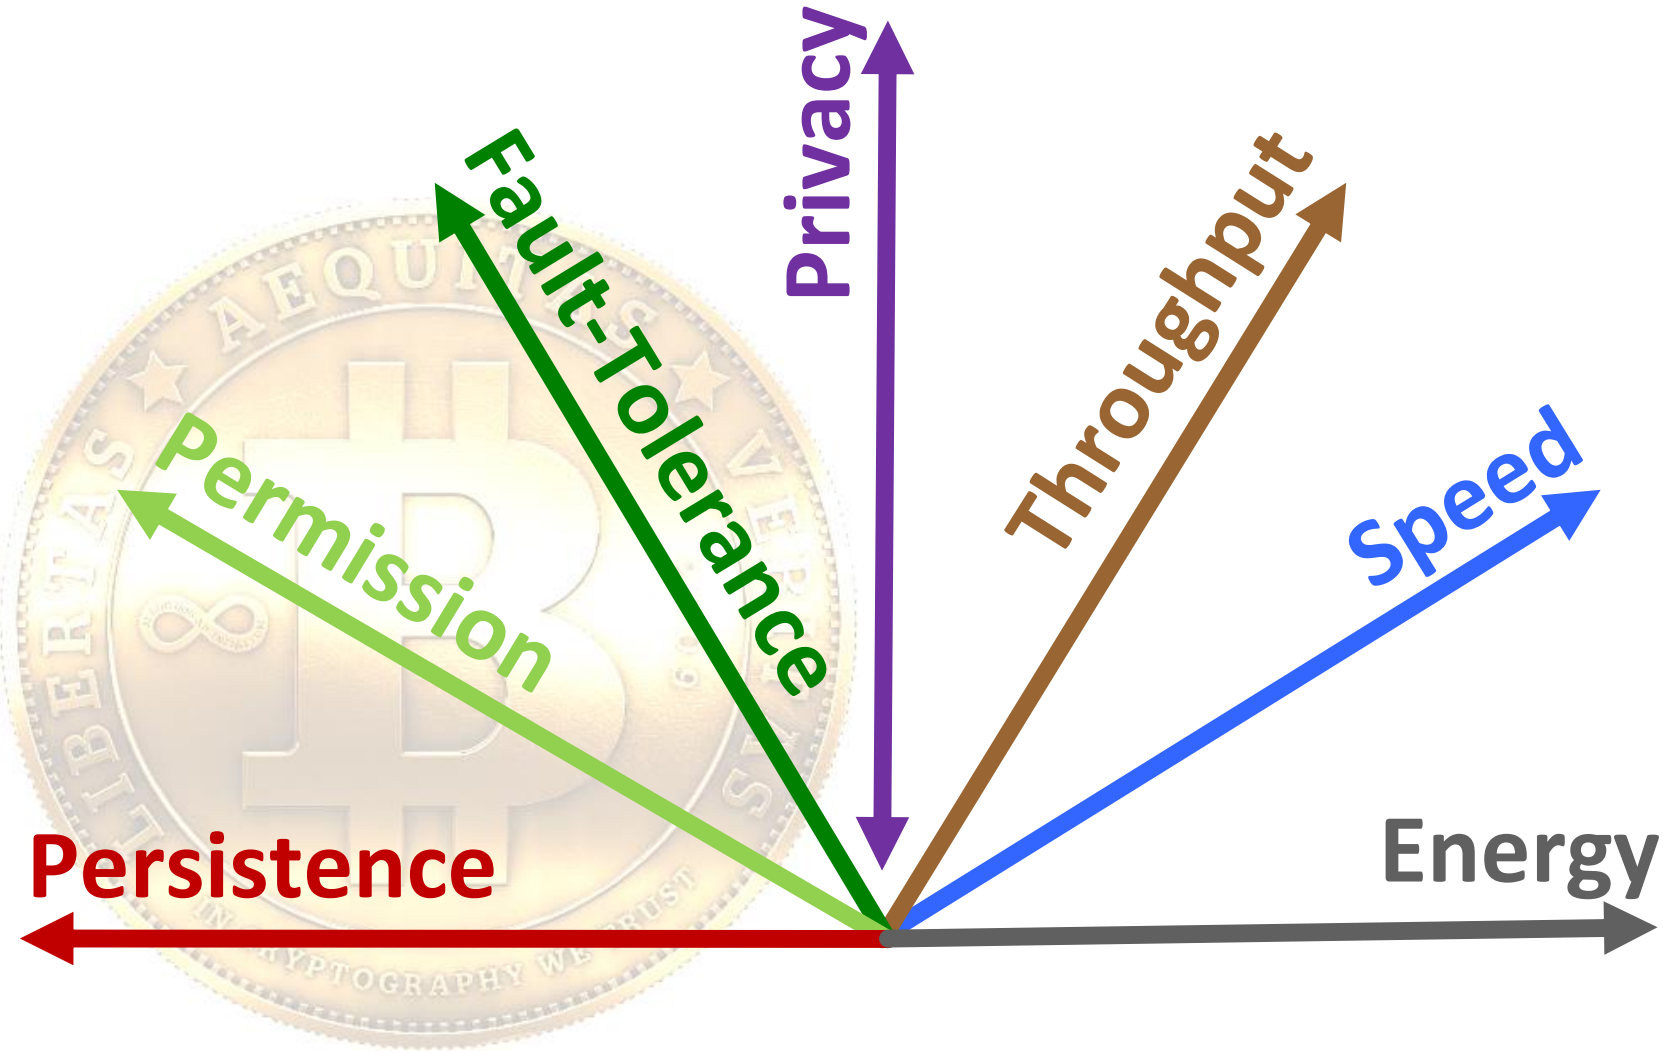

What About Privacy?

It’s Complicated.

Privacy

Anonymity/Public ⟷ Identity/Private

Research Issues

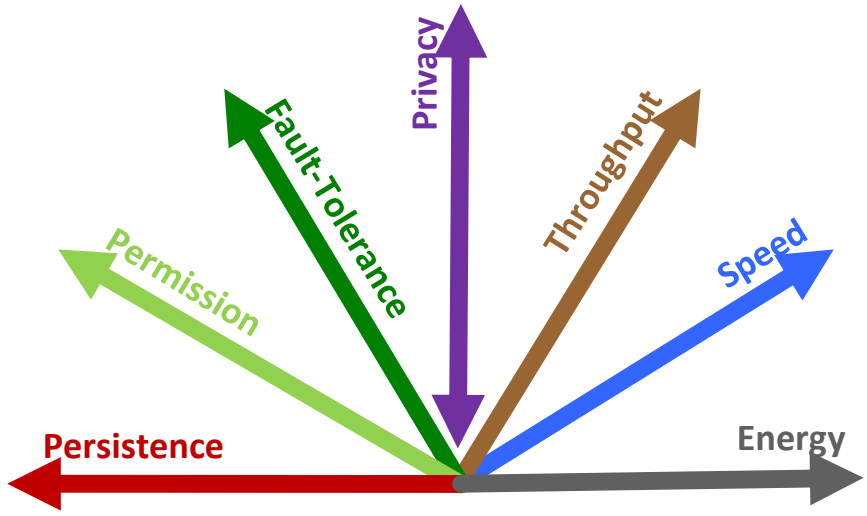

Solution to “many” problems: “Layer 2”

Plus: crypto, language (smart contracts), game theory, measurements, …

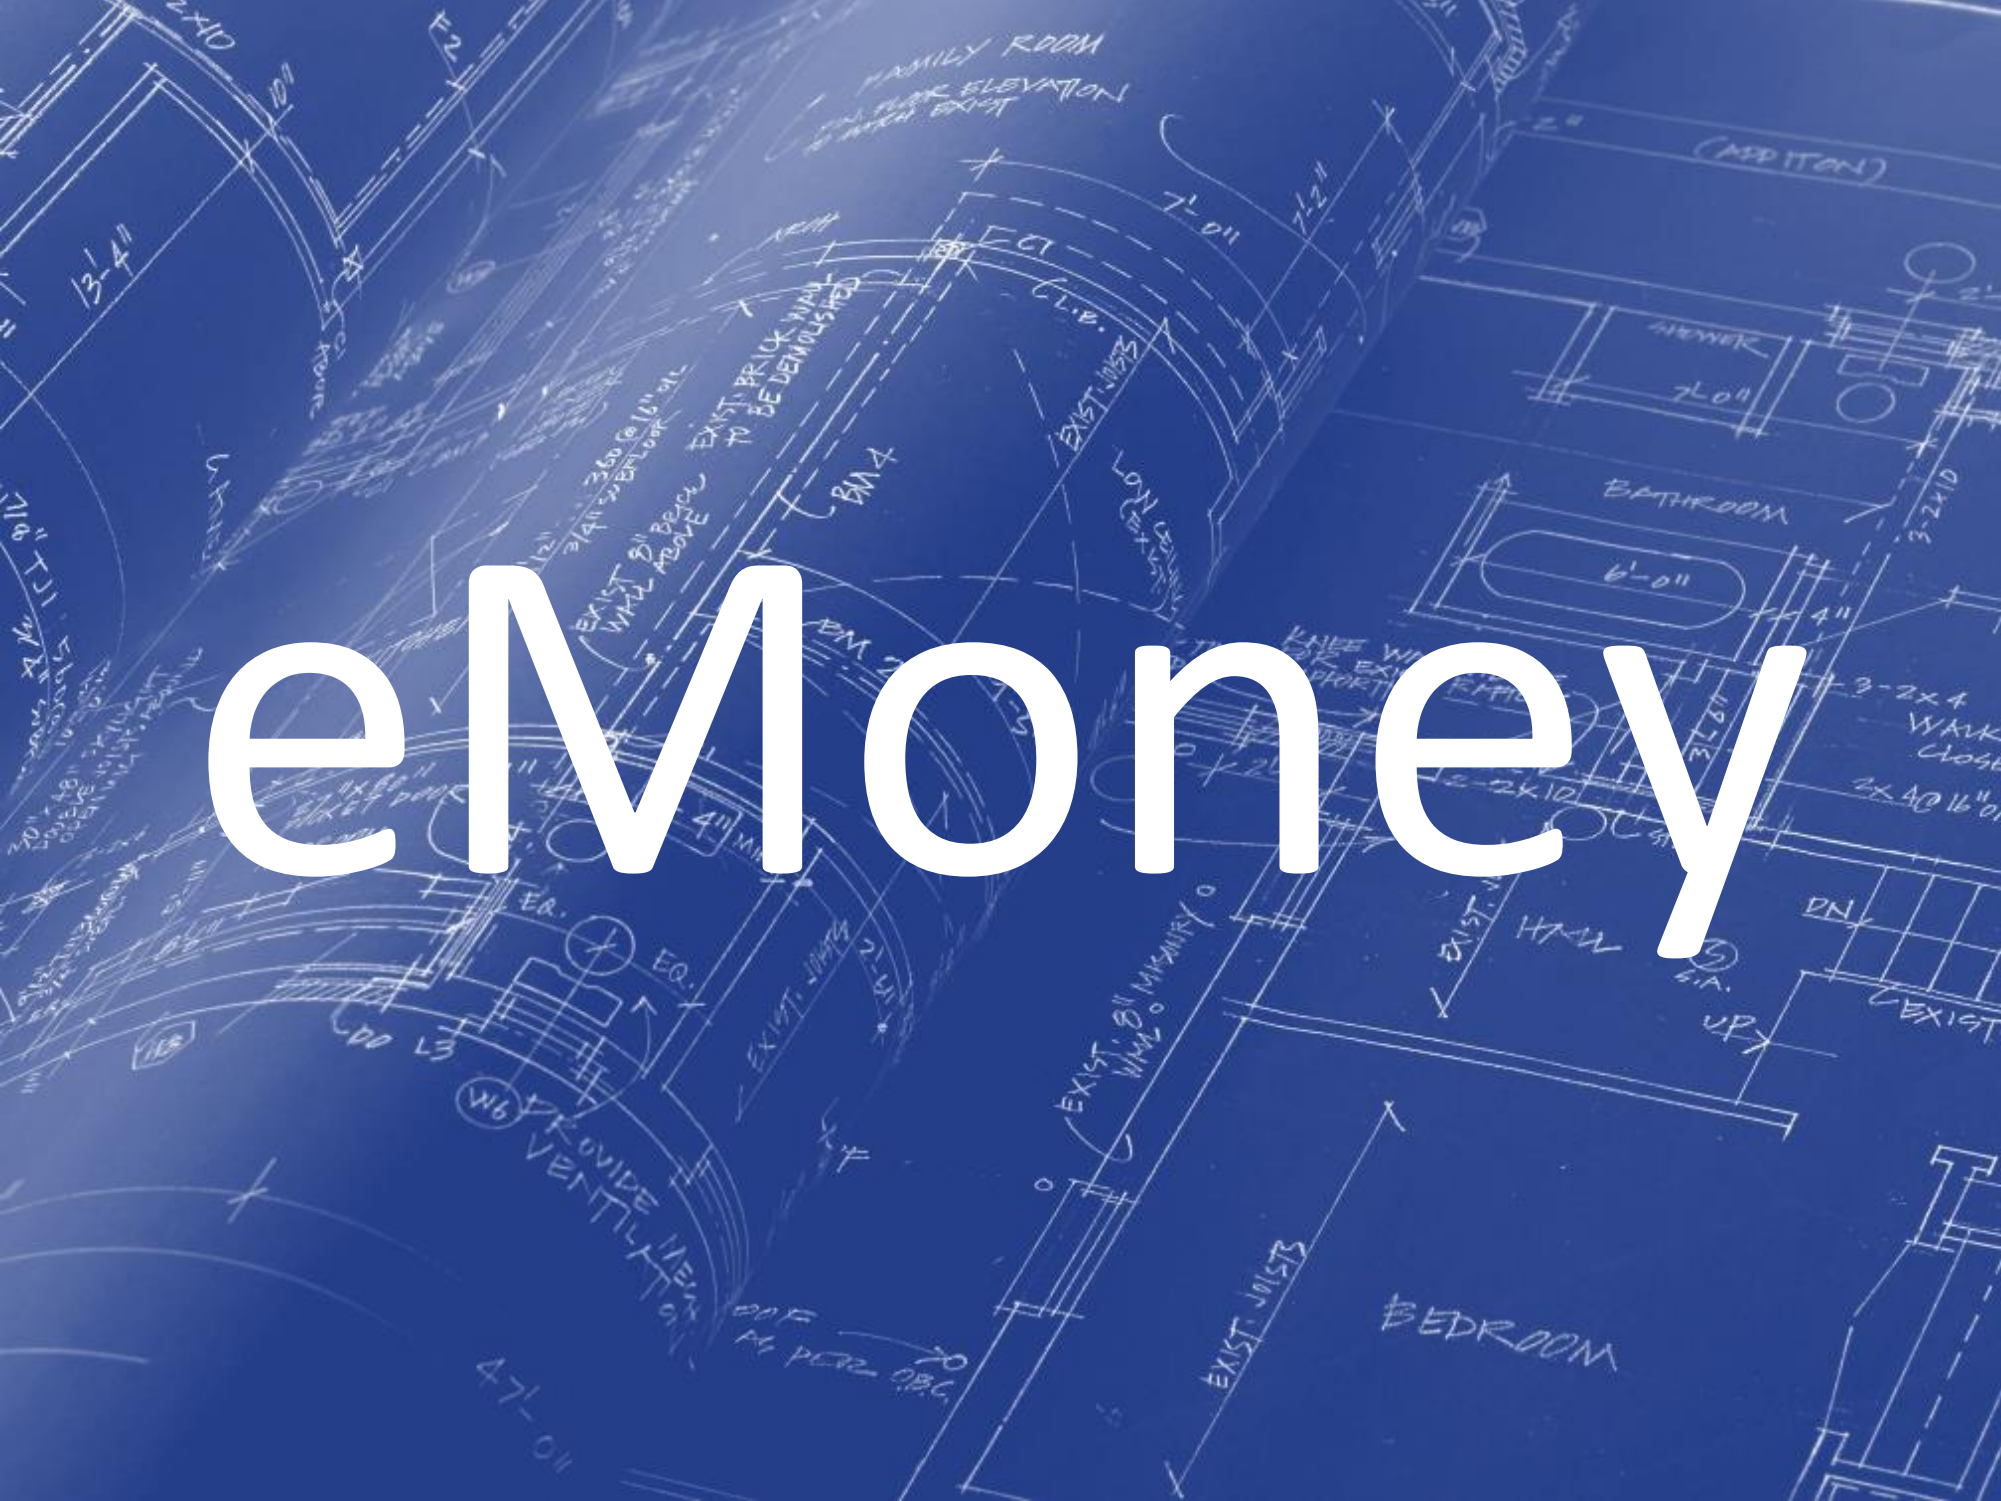
eMoney

Permissioned Blockchain

&

Payment Network

Permissioned Blockchain

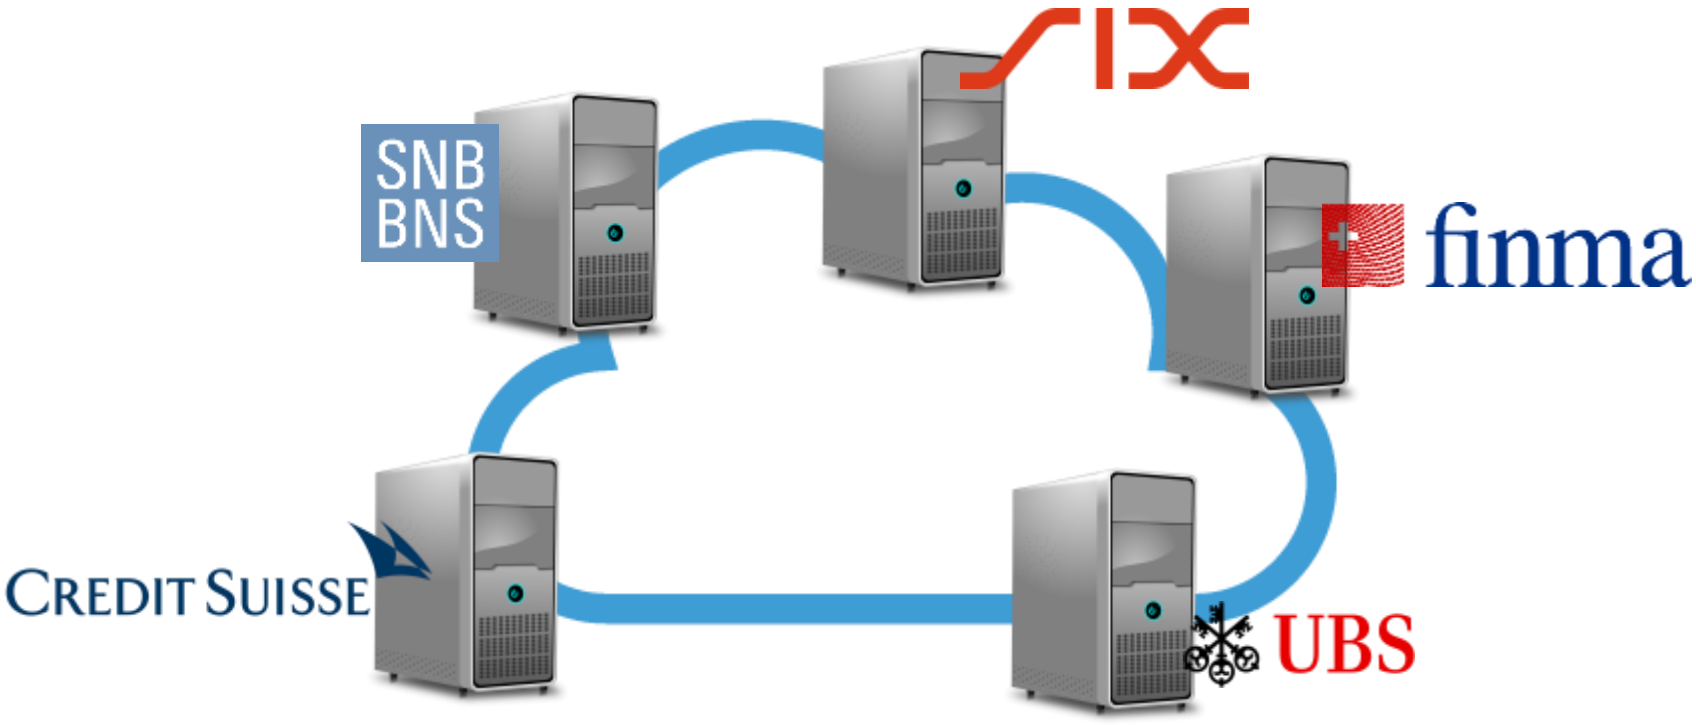

Payment Network

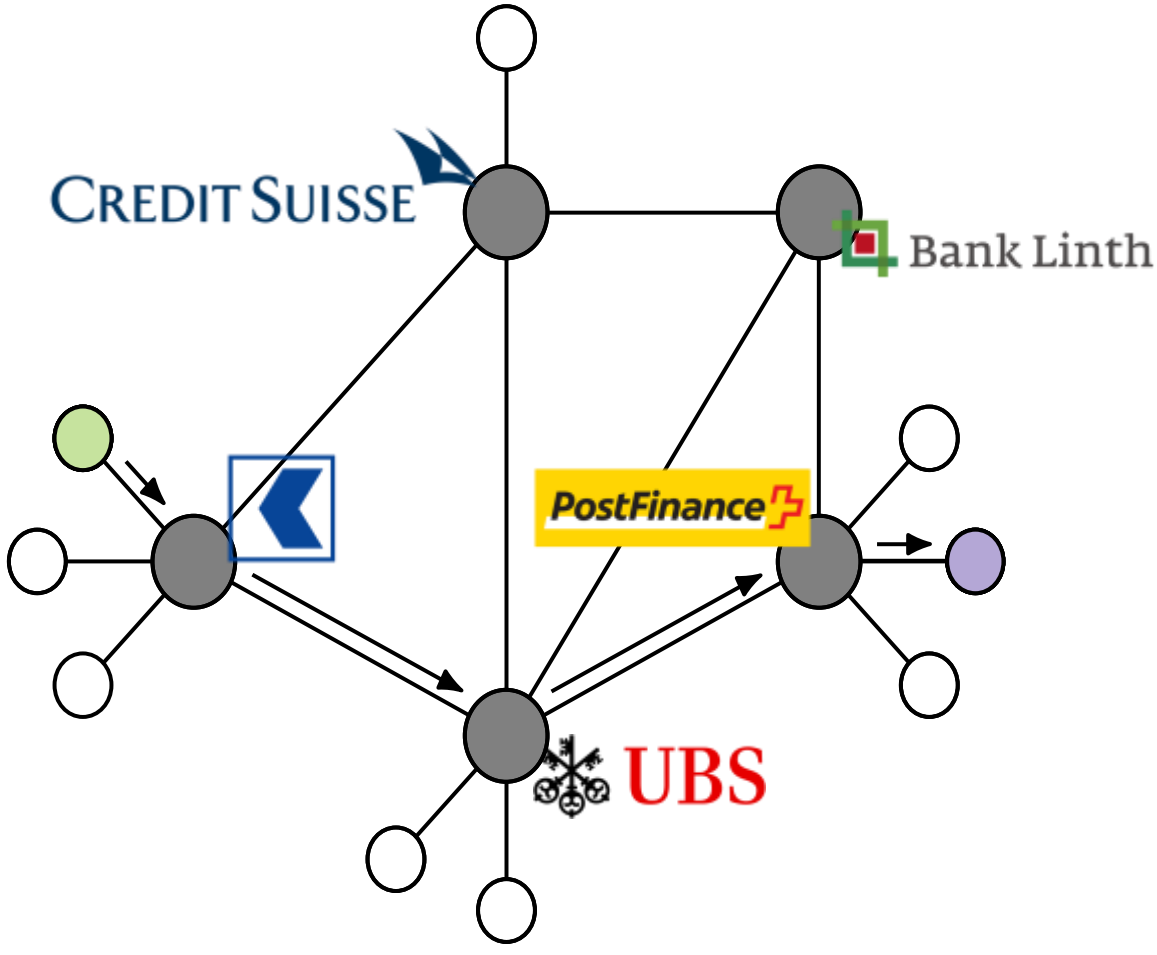

Bitcoin	eMoney
Anonymity	Accountability
Open/Anarchic	Closed/Private
Blockchain	Paxos, PBFT, …
Eventual Consistency	Strong Consistency
Proof-of-Work	Central Banks

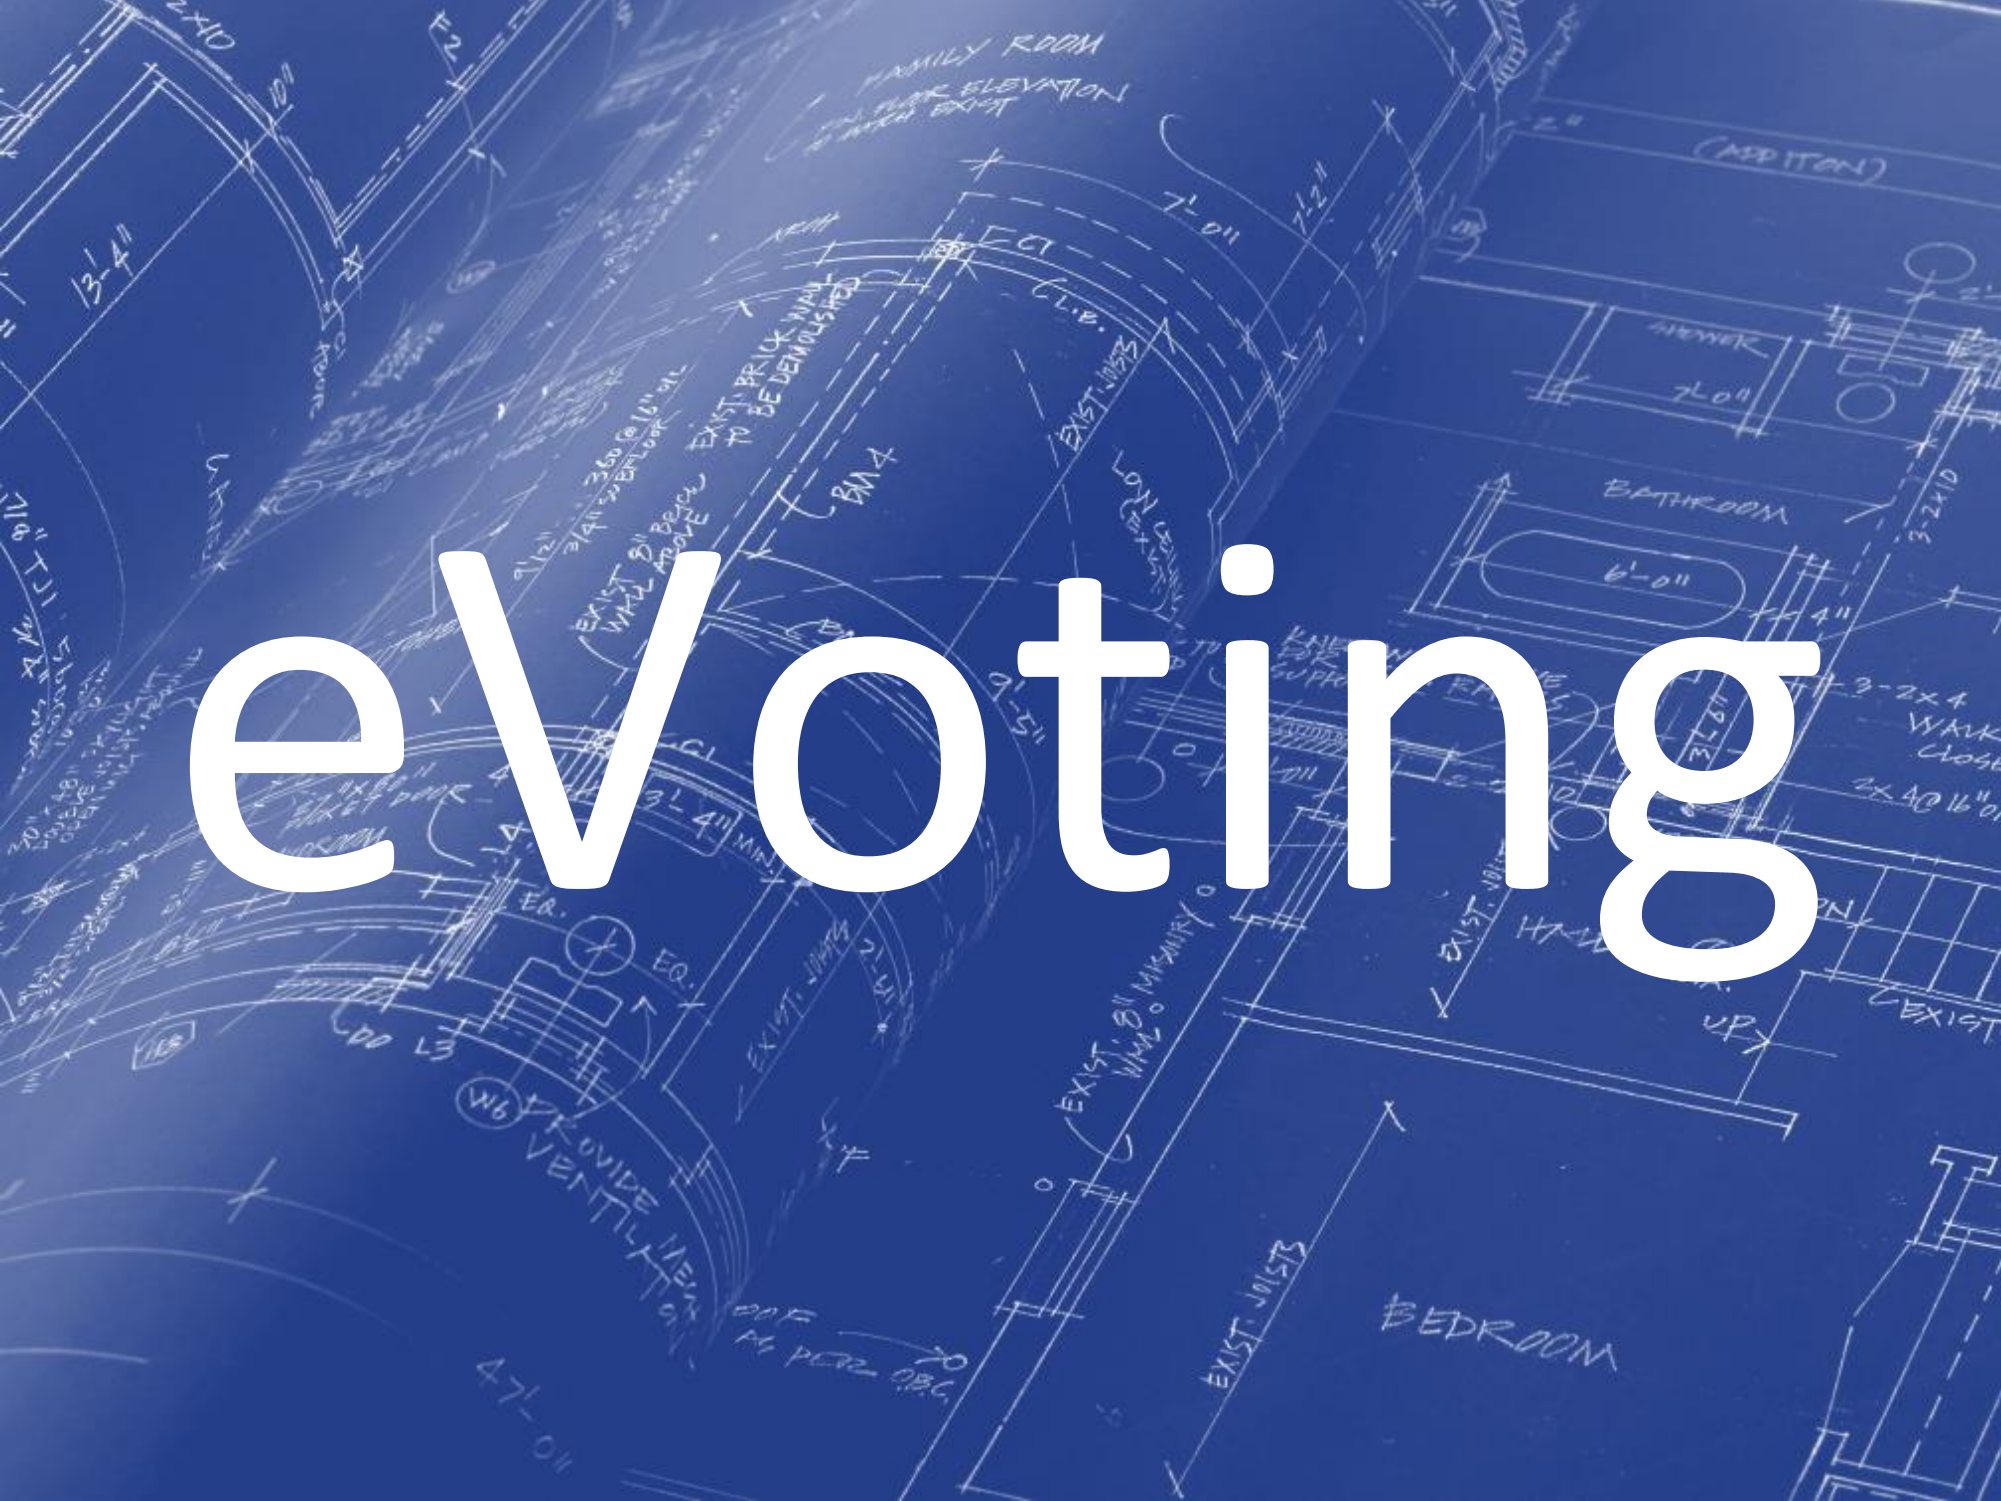
eVoting

What’s Wrong with Paper?

Cost

Verifiability

Neue Zürcher Zeitung

Rund 26 Prozent der Zürcher Wahlzettel waren nicht gültig

Anonymity

Identity Swapper

Identity Mixer

…

Election Help

Democracy Beyond Yes or No

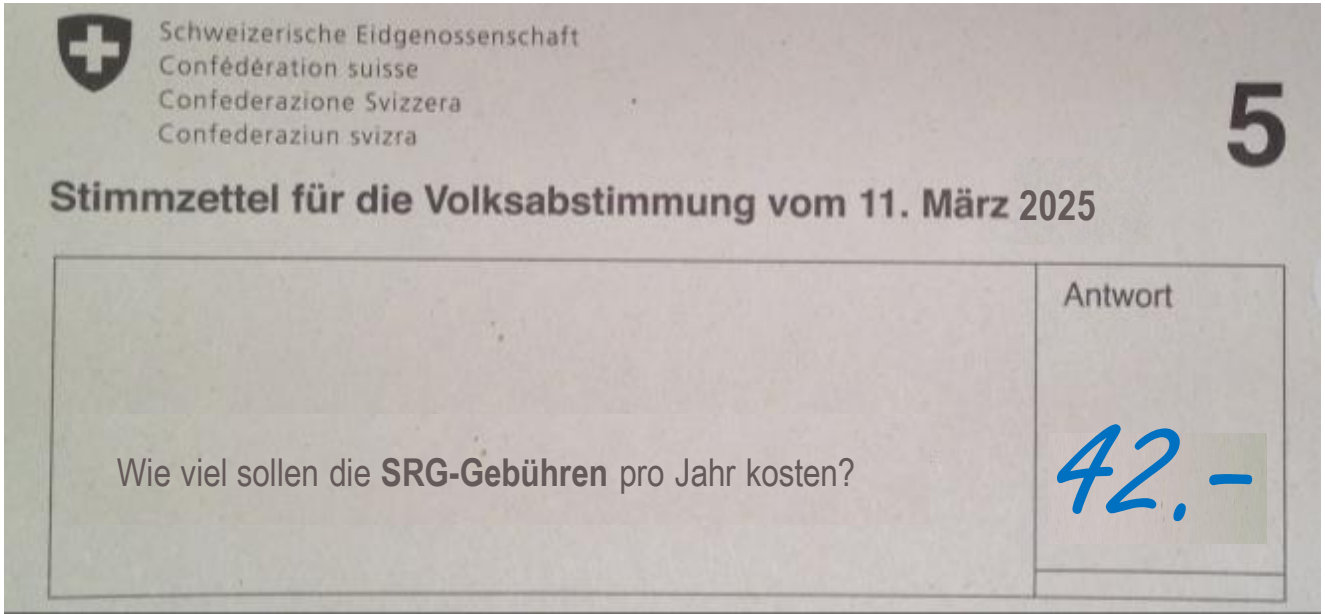

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

5

Stimmzettel für die Volksabstimmung vom 11. März 2025

	Antwort
Wie viel sollen die **SRG-Gebühren** pro Jahr kosten?	42.–

Don’t bring a Blockchain to a Gunfight

So what’s new, really?

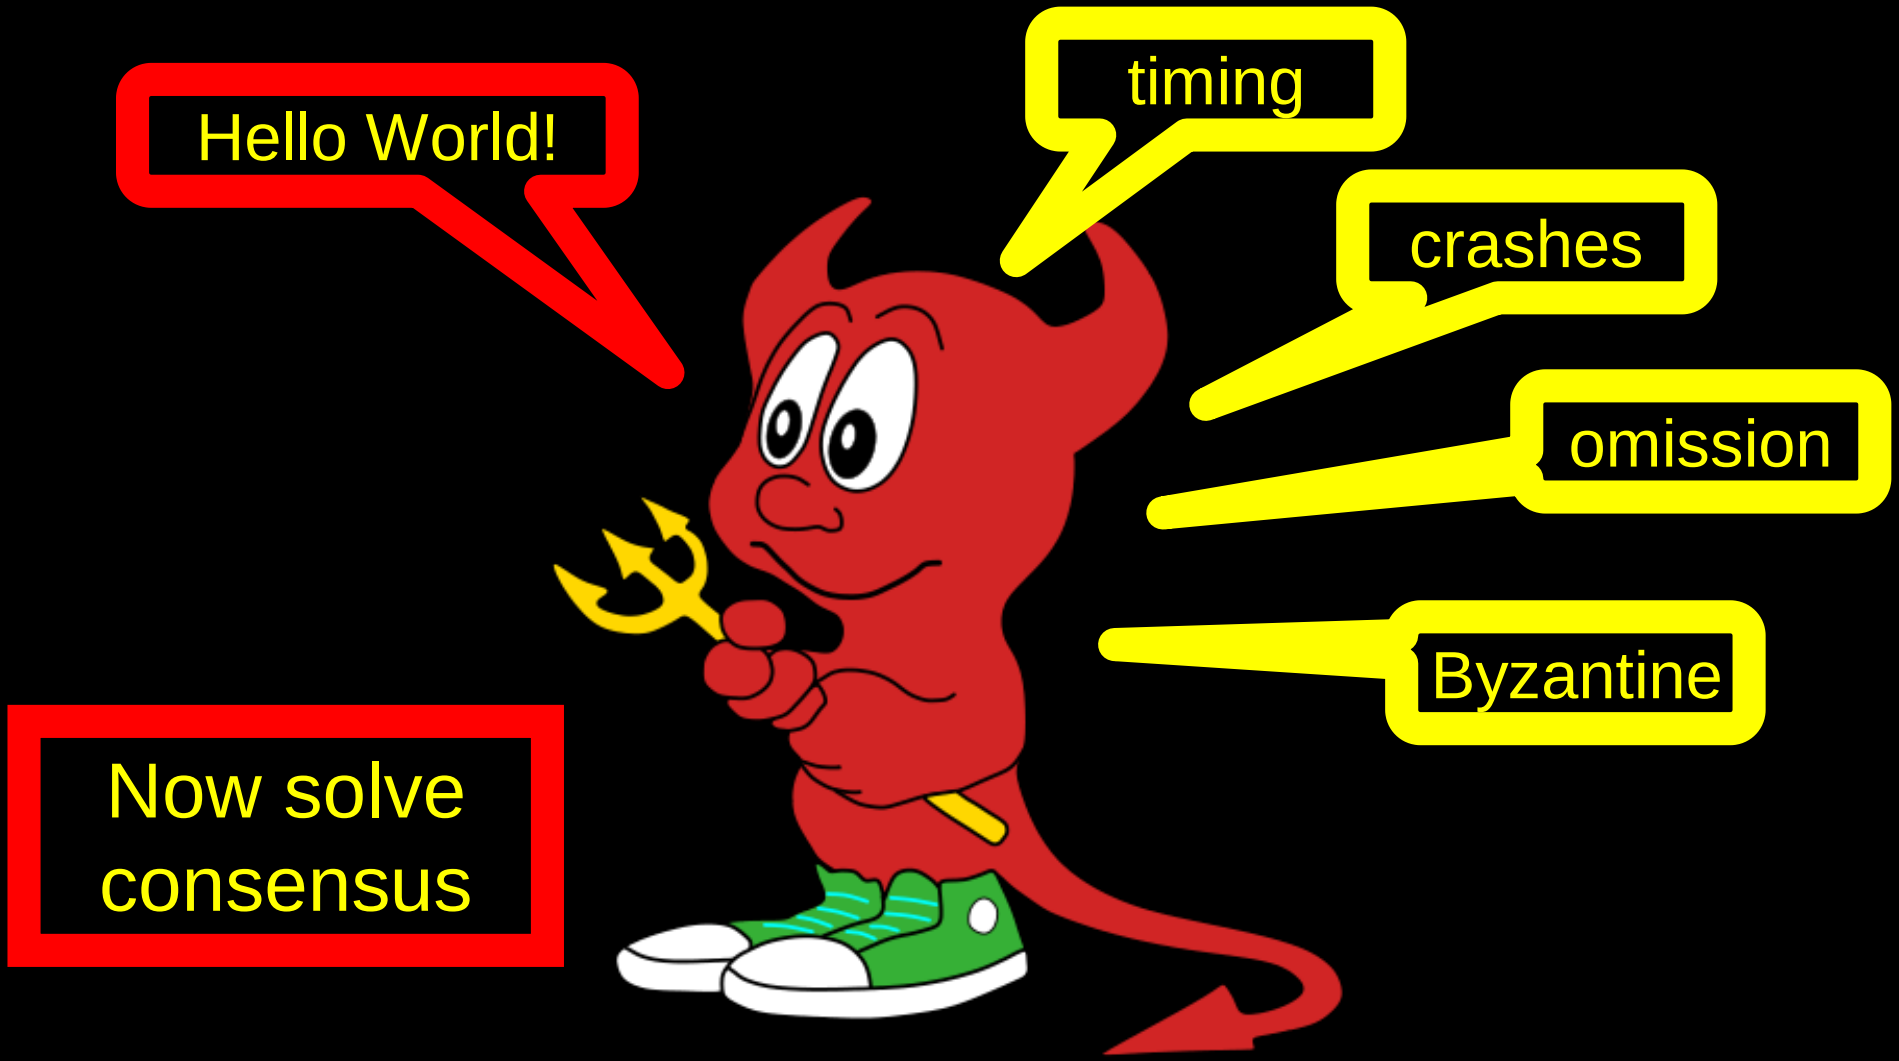

Classical Adversary

Modern Adversary

Hype

"First practical solution to a longstanding problem in computer science, Byzantine Generals."

"Satoshi solved a problem that academic computer scientists thought was impossible"

"Bitcoin is digital gold, it will put us back onto a sound monetary policy"

"Bitcoin will end wars"

… and Criticism

"A non-deliberate Ponzi scheme"

"It's yet another eventually consistent database"

"Flawed technology, inherently limited in scale and performance"

"Unlikely to impact the finance sector"

Would you rather fight…?

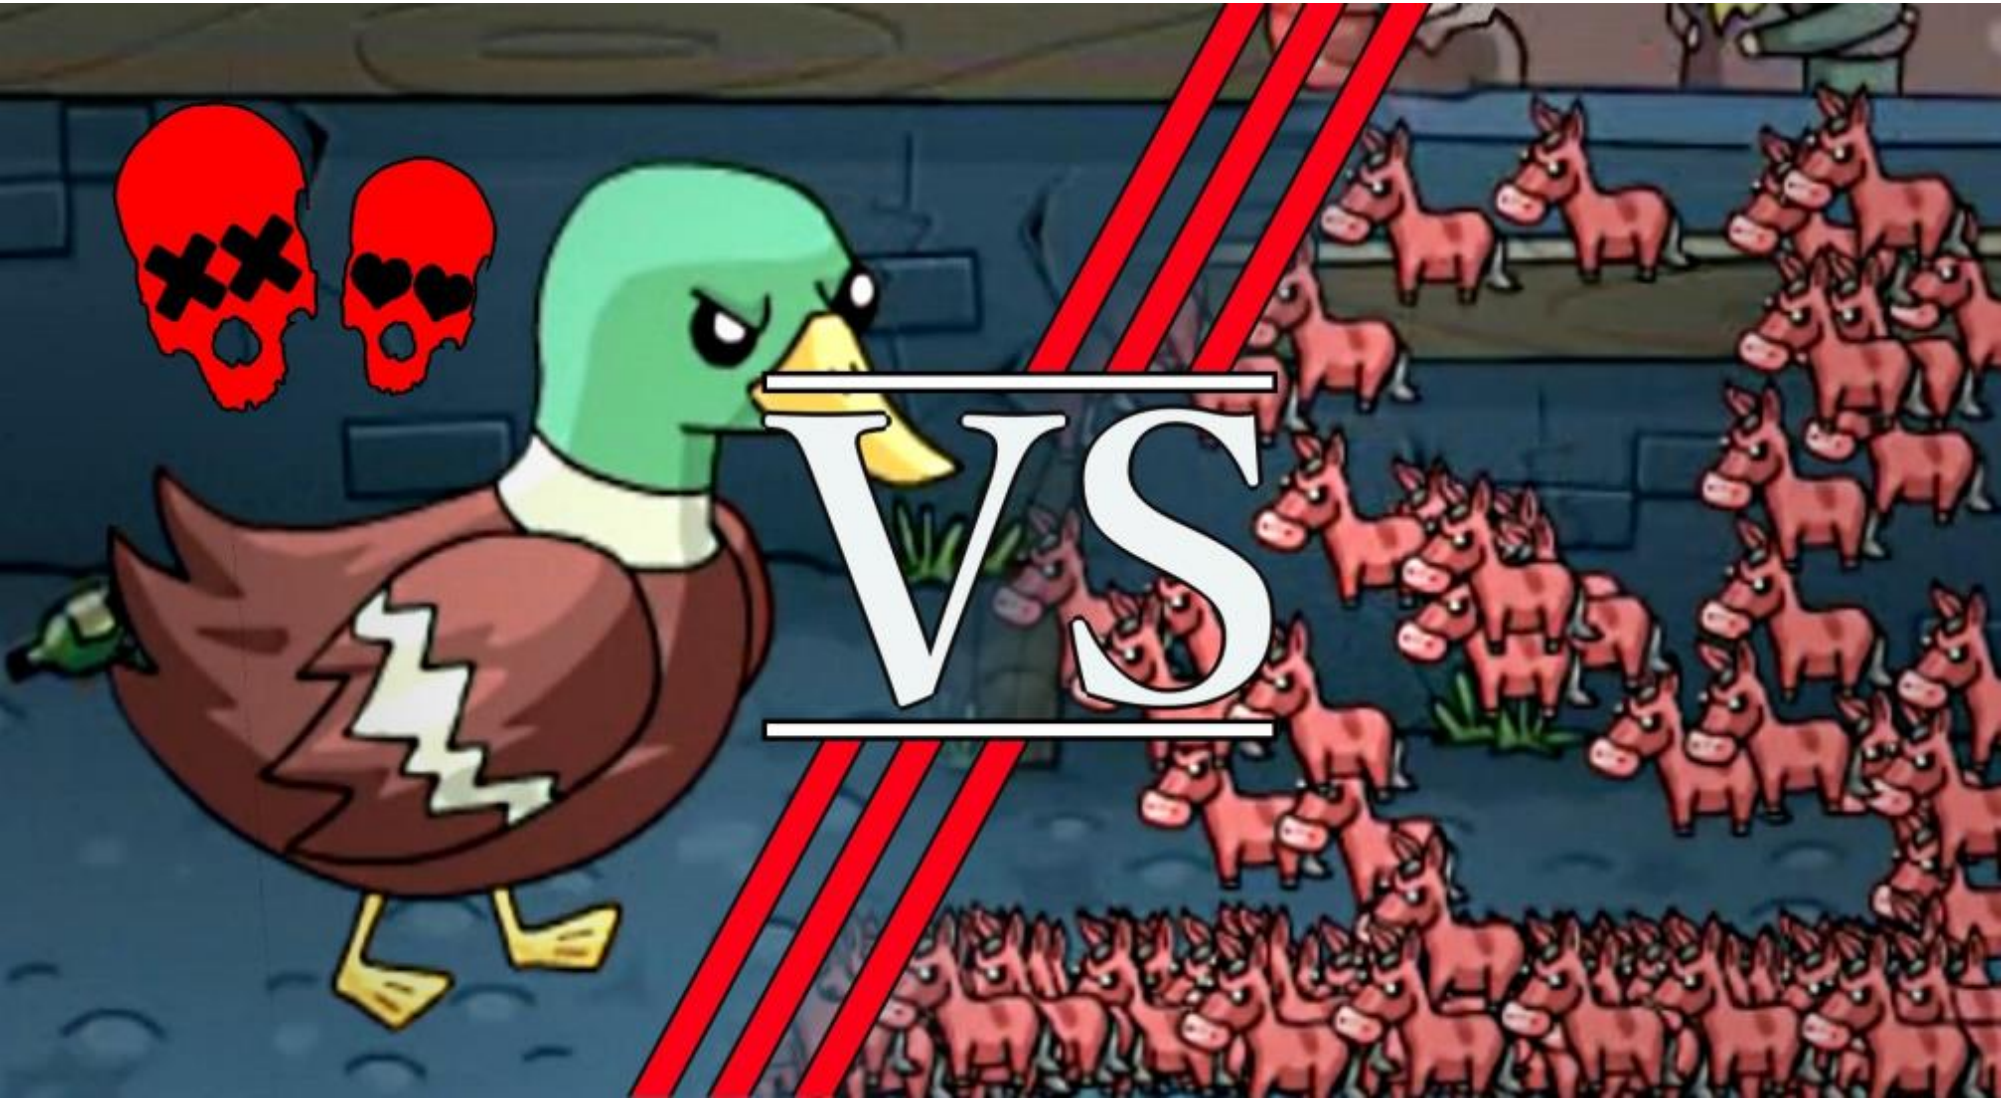

Cloud vs. Blockchain

Would you rather trust…?

Cloud vs. Blockchain

Thanks to lots of hardware…

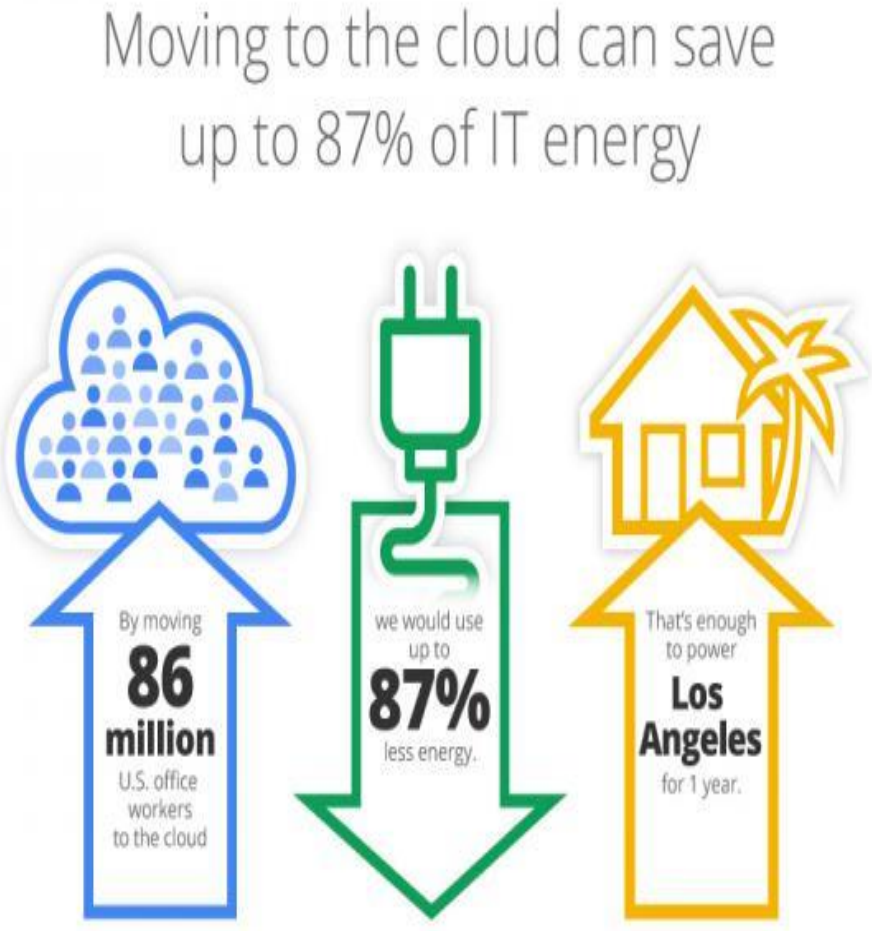

Cloud vs. Blockchain

“We at big corp will run your blockchain in our cloud!”

What's this?

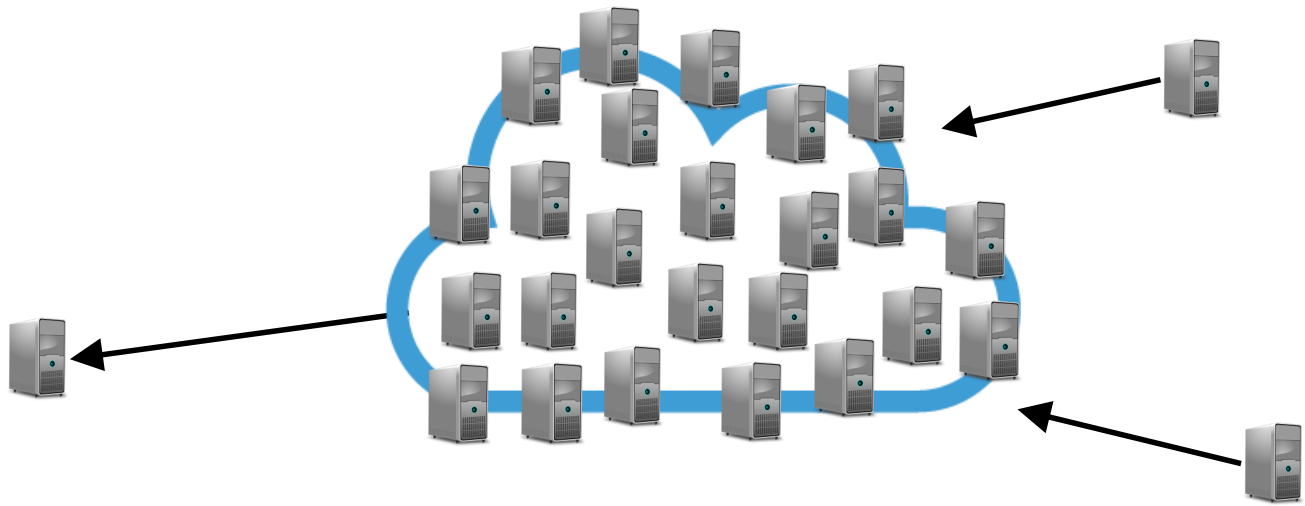

A Blockchain?

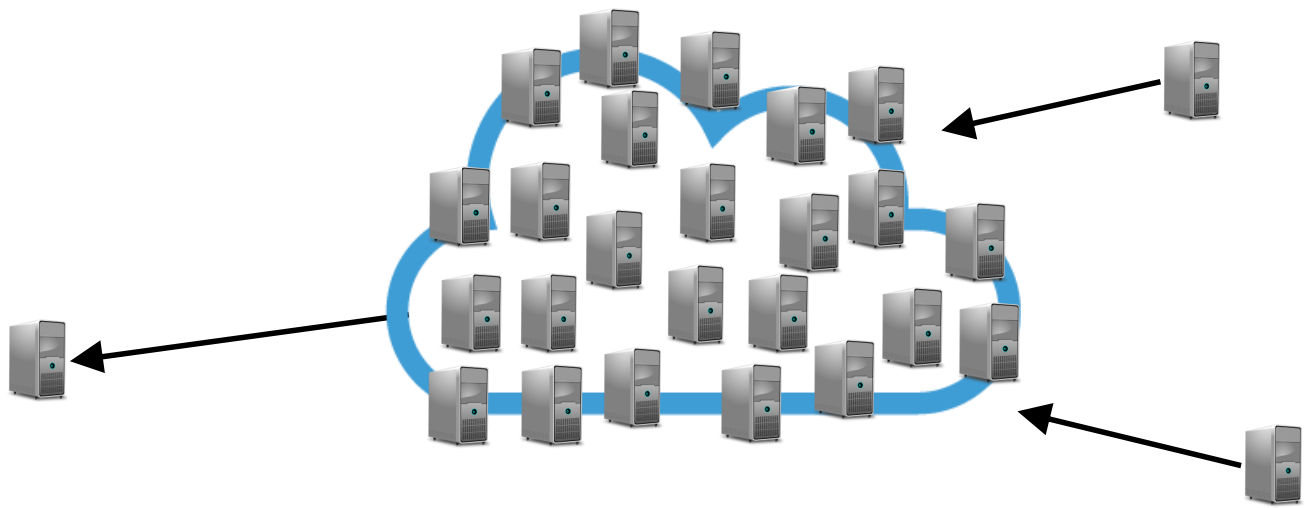

A Cloud?

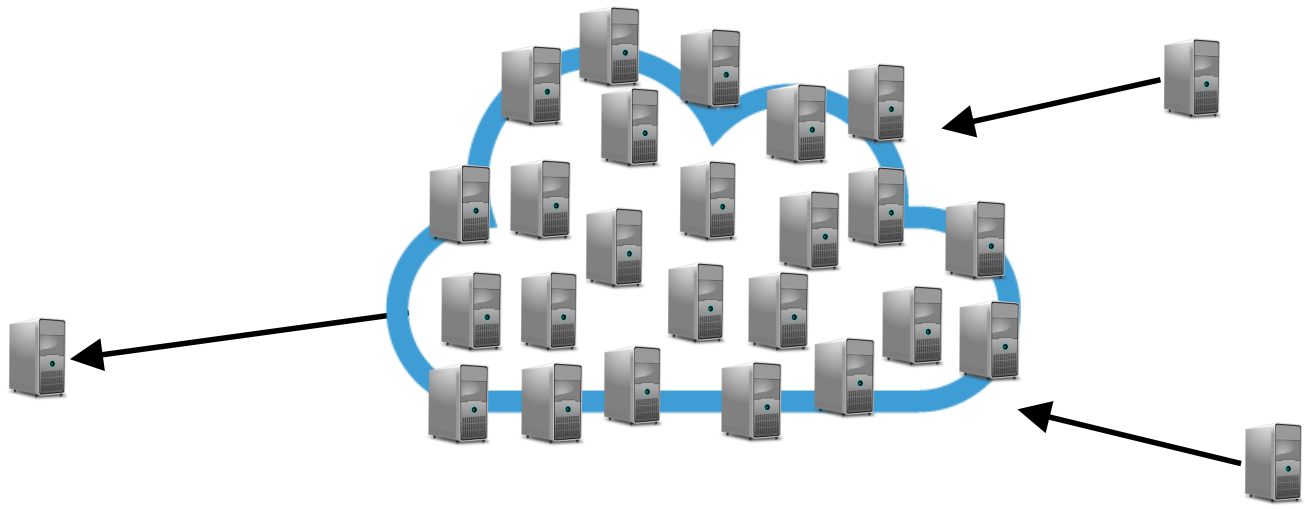

A Distributed System!

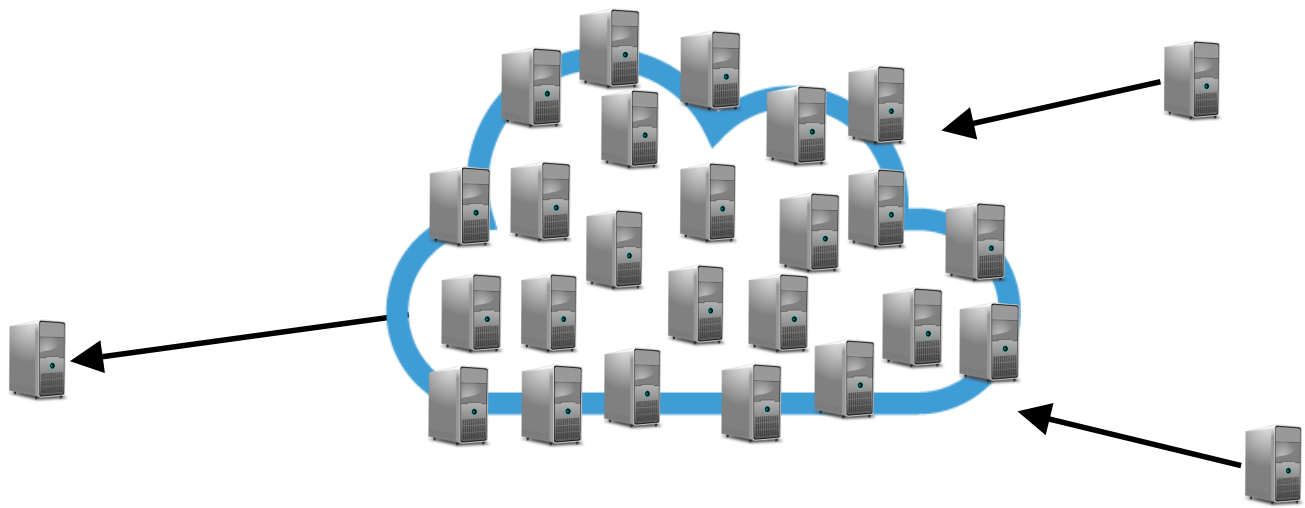

Summary

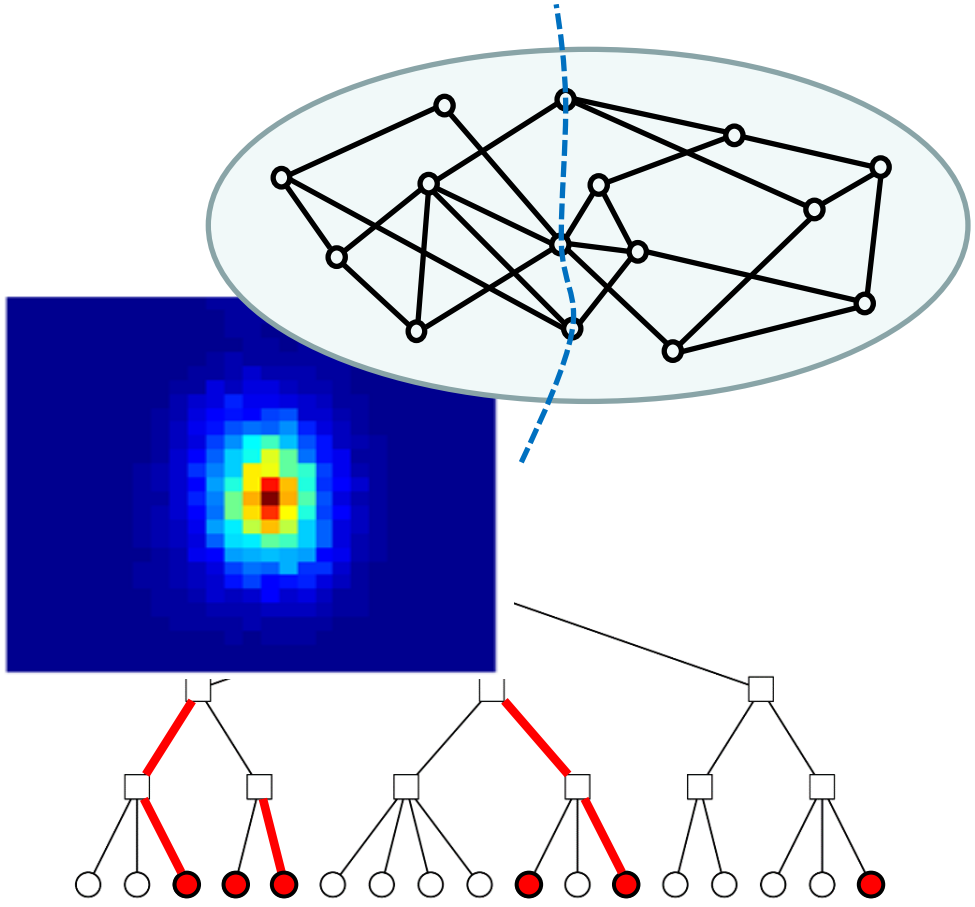

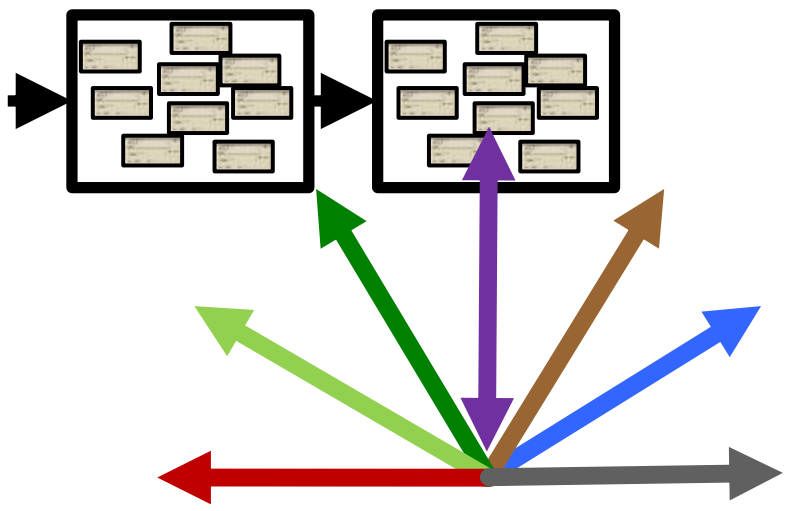


Thank You!
Questions & Comments?
Thanks to my co-authors
Vertex Separators: Sebastian Brandt
Online With Delay: Yuval Emek, Shay Kutten
Cloud GPS: Manuel Eichelberger
www.disco.ethz.ch

Abstract:

Algorithms interact in two main ways with the cloud. There exist algorithms which are tailored for the cloud, for which the cloud is the perfect environment. Moreover, the cloud may also benefit from optimization algorithms, algorithms that make the cloud more efficient. The AlgoCloud program features papers which roughly fit one of the two, and I will also give a few examples in the first part of my talk. Apart from these algorithms *for* the cloud, I will also talk about algorithms *against* the cloud. Recently, blockchains are hyped to be a cloud competitor, sometimes even a cloud killer. In the second part of my talk we will discuss whether there is some truth to whether blockchains are going to threaten the successful cloud paradigm.