

# Demystifying Bitcoin



**Journey to the center of  
distributed computing**



## 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

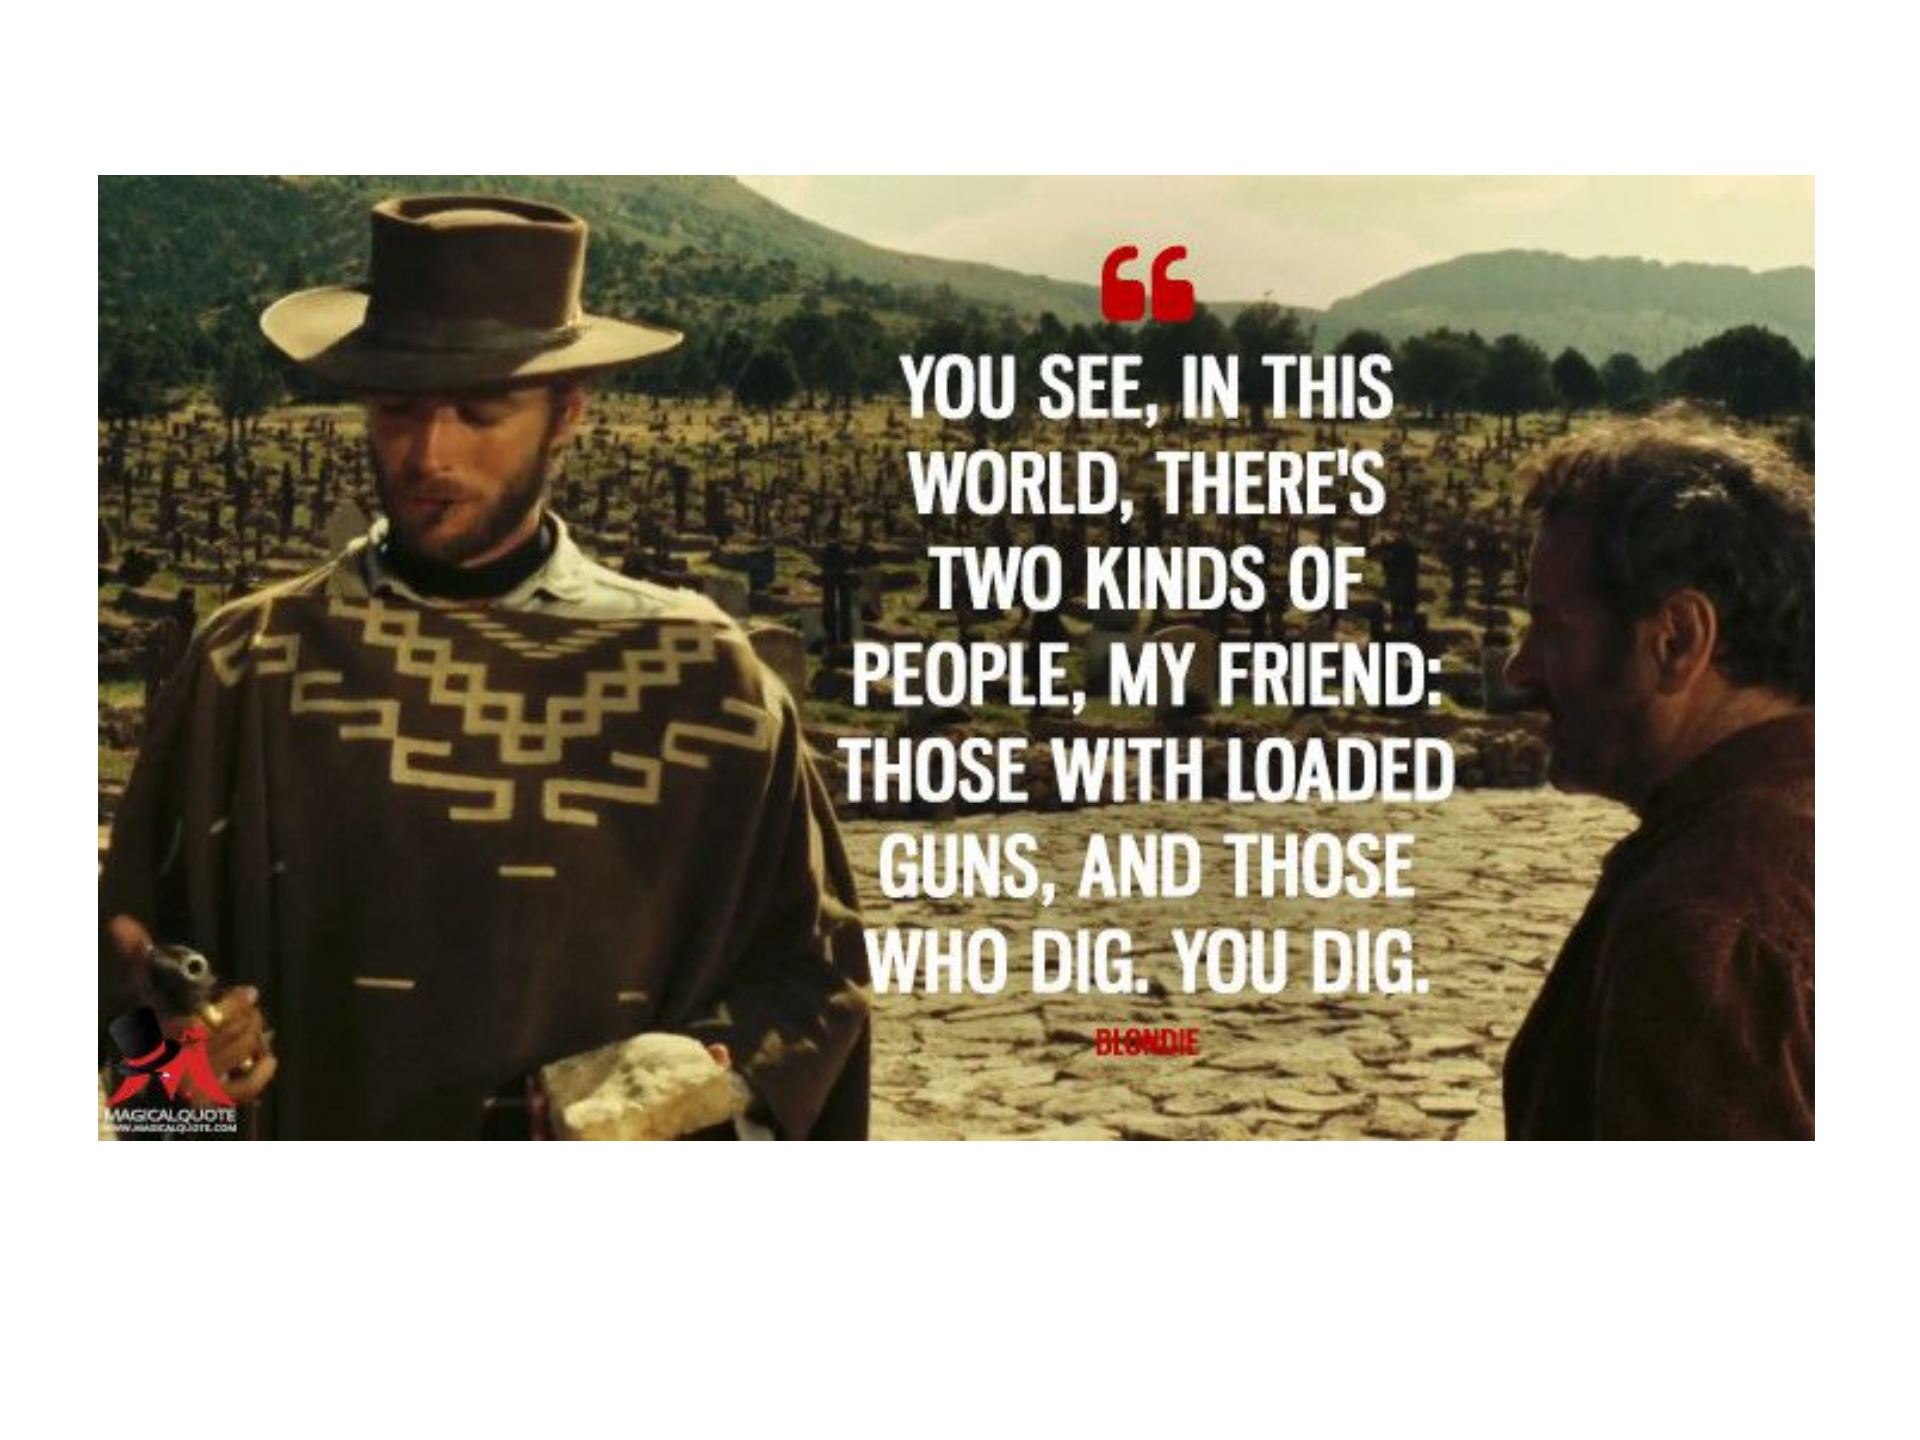
# Roadmap

- ☛ **(1) The main question**
  - ☛ **(2) The bitcoin problem**
    - ☛ **(3) The main answer**

☛ **X000 implementations**



☛ « Computing's central challenge is how not to make a mess of it ...» E. Dijkstra

A still from the movie 'Blondie' showing a man in a brown hat and poncho (Blondie) holding a gun and a rock, standing in a cemetery. Another man is partially visible on the right. The background shows a landscape with hills and trees.

“  
YOU SEE, IN THIS  
WORLD, THERE'S  
TWO KINDS OF  
PEOPLE, MY FRIEND:  
THOSE WITH LOADED  
GUNS, AND THOSE  
WHO DIG. YOU DIG.

BLONDIE

# P vs NP

$$7 * 13 = ?$$


$$? * ? = 91$$

## Asynchronous vs Synchronous

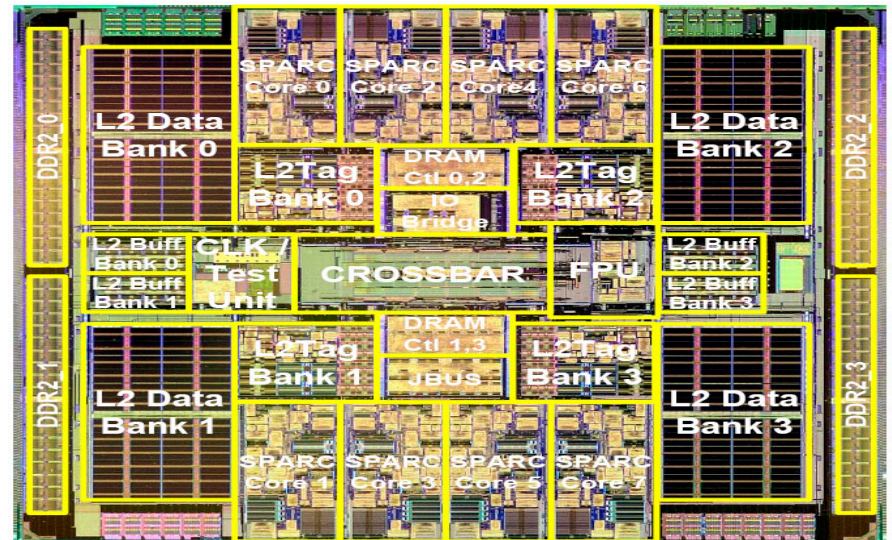
Is payment an asynchronous problem?

- « To understand a distributed computing problem: bring it to shared memory » T. Lannister


# The infinitely big



# The infinitely small

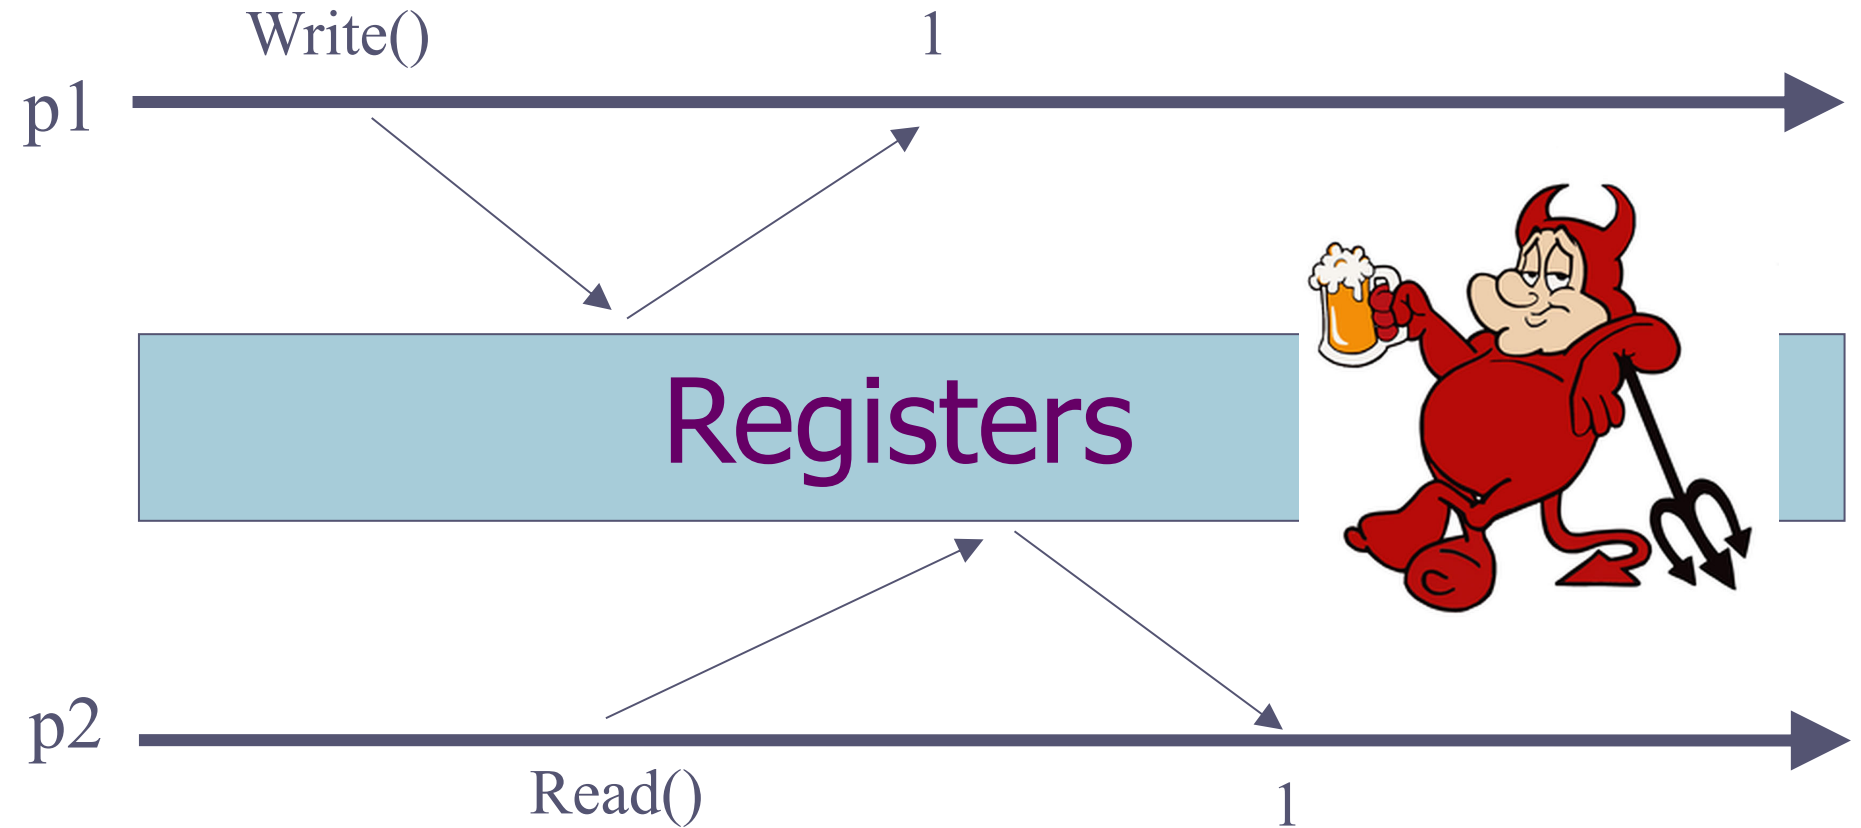


# Message Passing






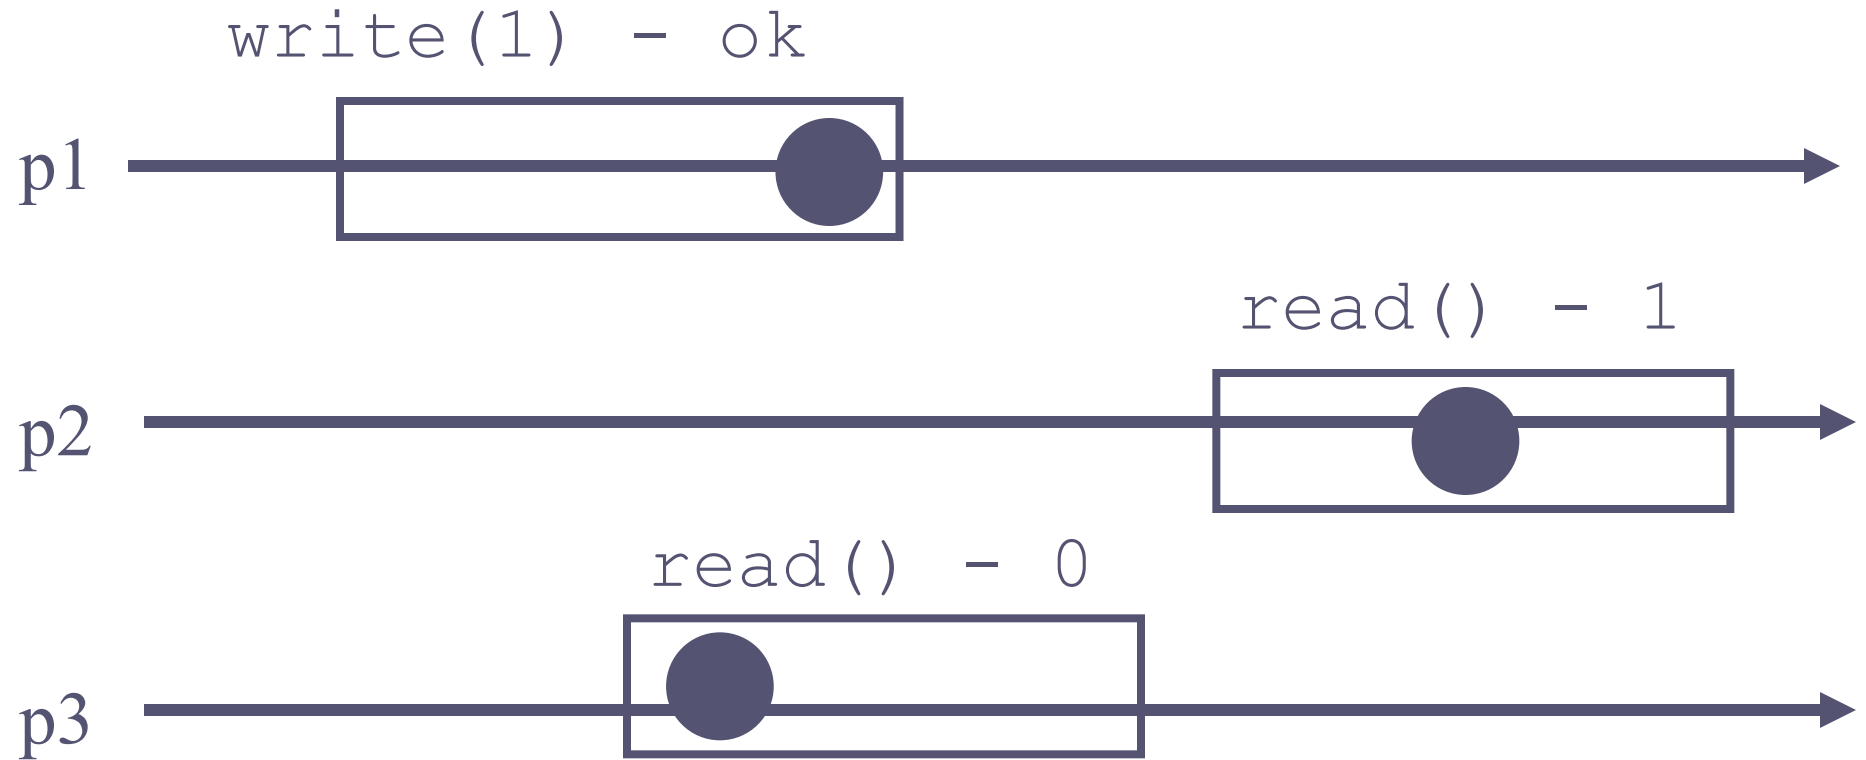
# Shared Memory



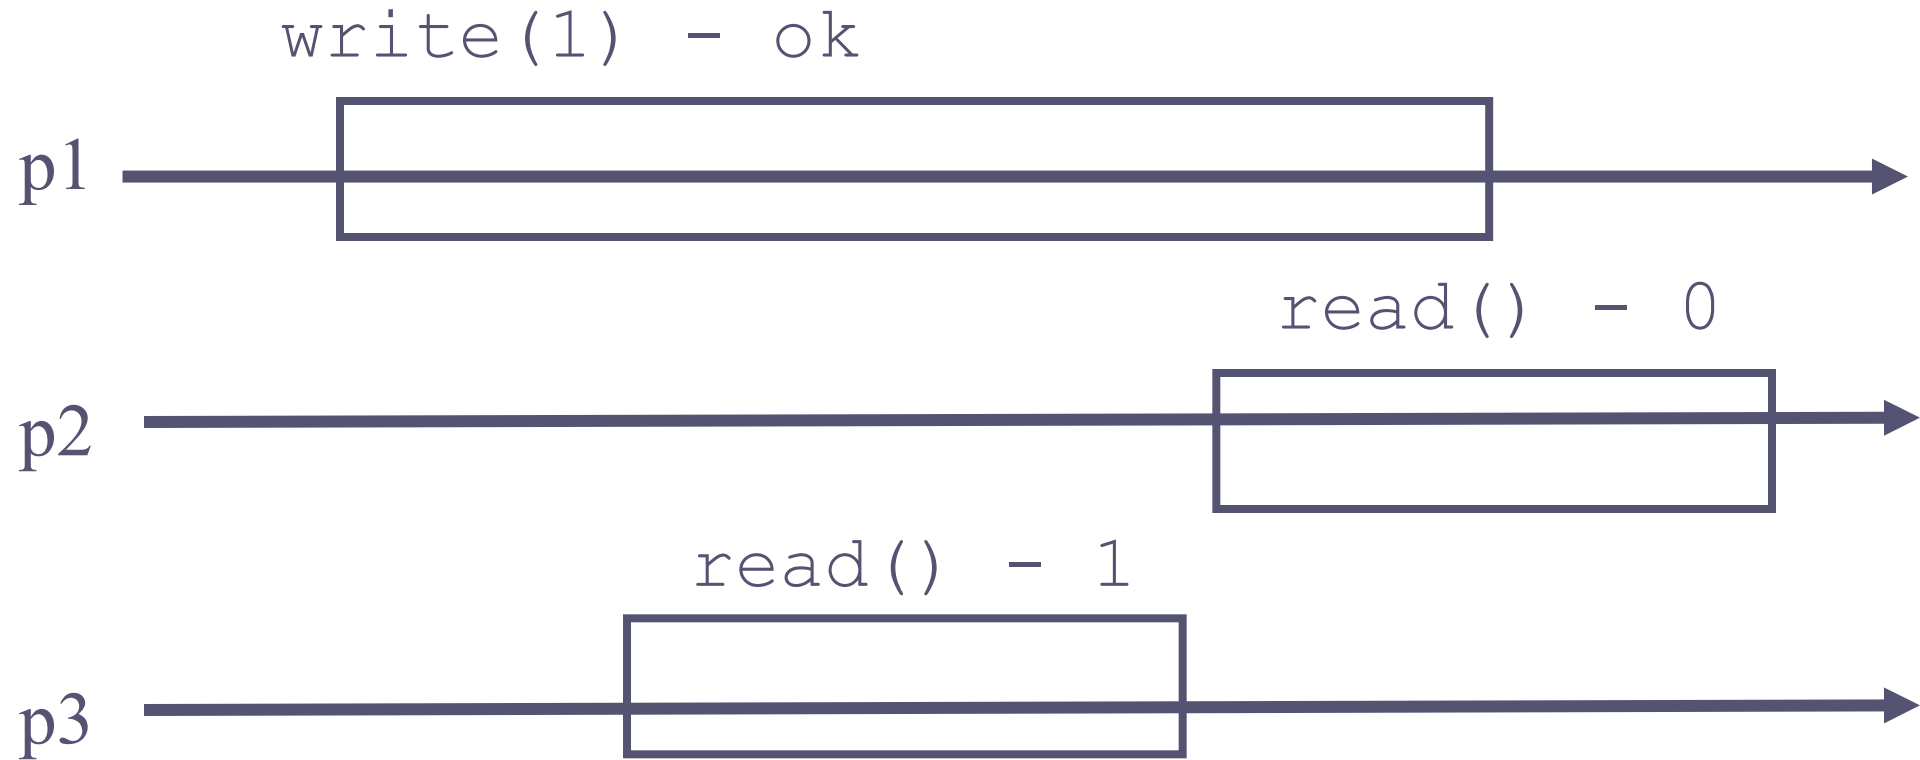
# Atomic Shared Memory



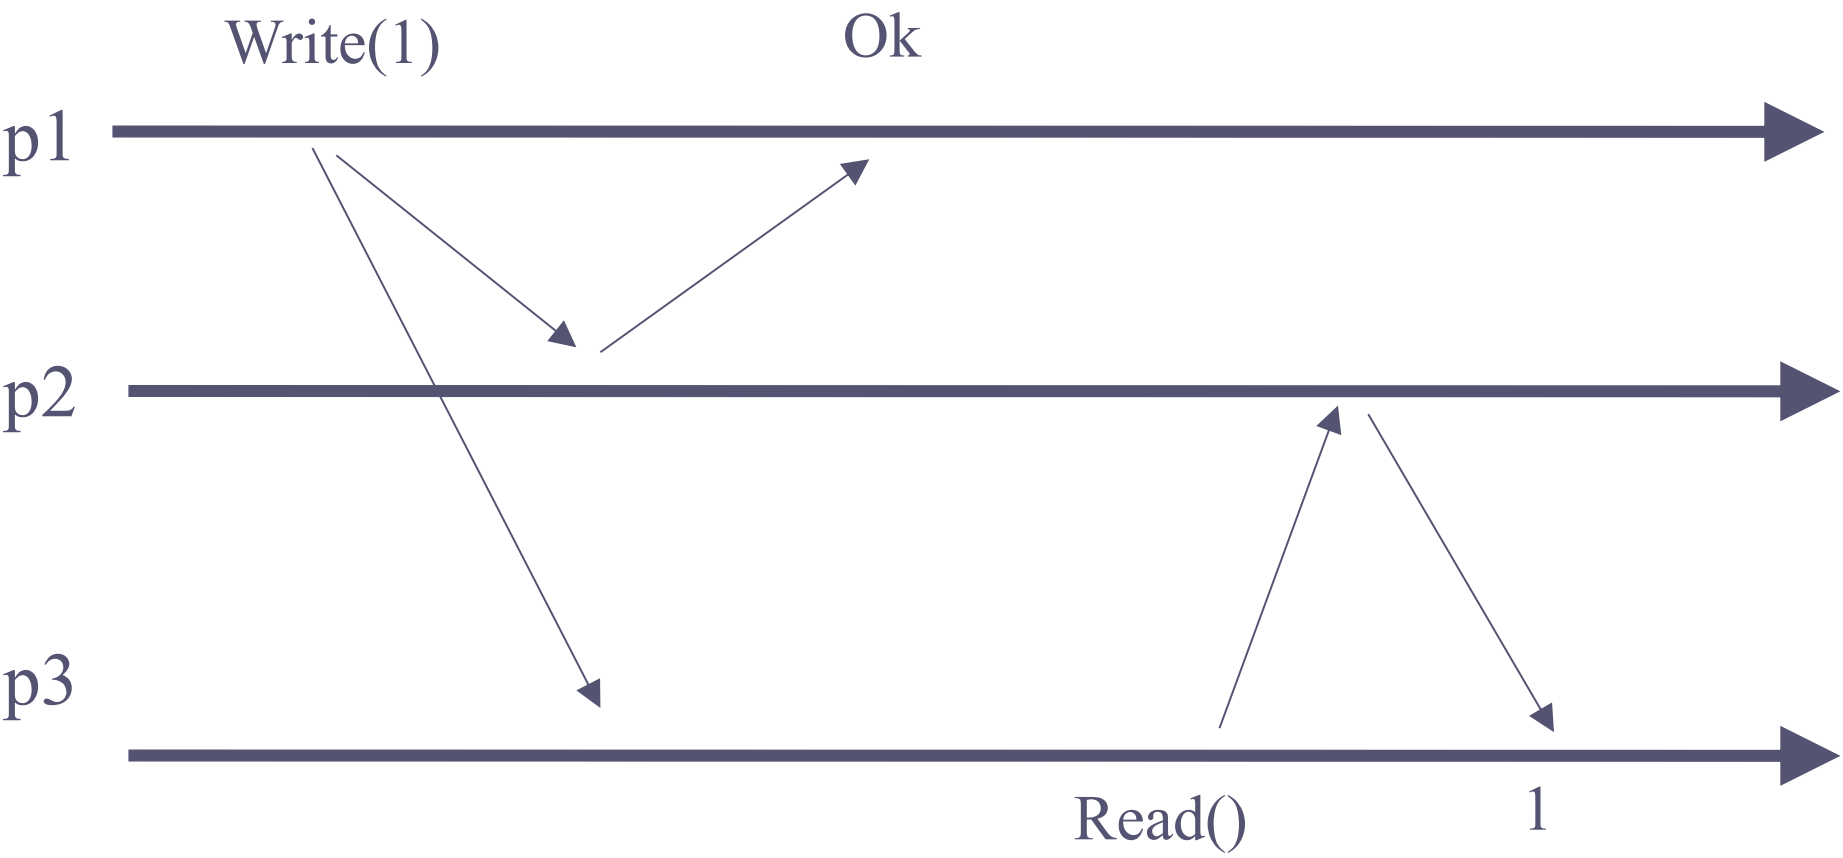
# Atomic Shared Memory



# Non-Atomic Shared Memory

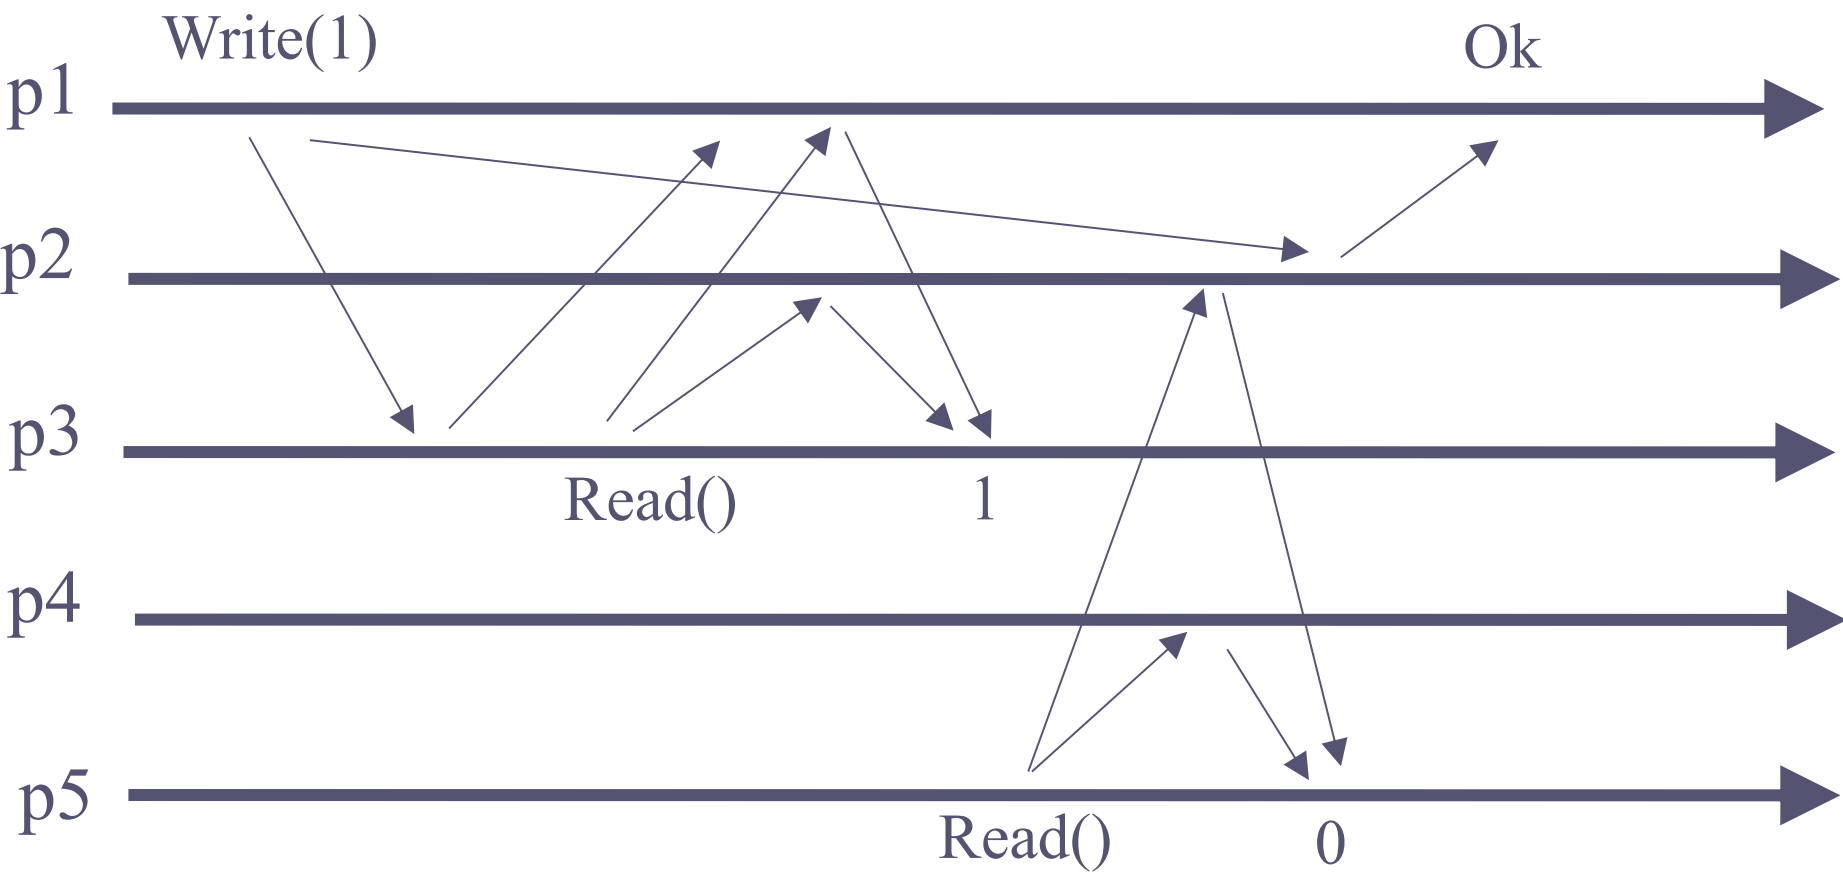


# Message Passing ↔ Shared Memory



***Quorums***

# Message Passing ↔ Shared Memory



# Message Passing $\leftrightarrow$ Shared Memory Modulo Quorums



☛ « Optimization is the source of all evil » D. Knuth

# Is payment an asynchronous problem?

## Payment Object



- Atomicity
- Wait-freedom



# Counter: Specification

- A *counter* has two operations *inc()* and *read()*; it maintains an integer *x* *init to 0*
- *read()*:
  - return(x)
- *inc()*:
  - $x := x + 1;$
  - return(ok)

# Counter: Algorithm

- The processes share an array of registers  $\text{Reg}[1, \dots, N]$
- *inc()*:
  - $\text{Reg}[i].\text{write}(\text{Reg}[i].\text{read}() + 1);$
  - $\text{return}(\text{ok})$
- *read()*:
  - $\text{sum} := 0;$
  - for  $j = 1$  to  $N$  do
    - $\text{sum} := \text{sum} + \text{Reg}[j].\text{read}();$
  - $\text{return}(\text{sum})$

# Counter\*: Specification

- *Counter\** has, in addition, operation *dec()*
- *dec()*:
  - if  $x > 0$  then  $x := x - 1$ ; return(ok)
  - else return(no)

Can we implement Counter\* asynchronously?

# Consensus



Agreement on a single value among multiple

**Safety:** No two processes must choose different values.

The chosen value must have been proposed by a process.

**Liveness:** Each process must eventually choose a value.

# 2-Consensus with Counter\*

- Registers R0 and R1 and Counter\* C - initialized to 1
- Process pI:
  - propose(vI)
  - RI.write(vI)
  - res := C.dec()
  - if(res = ok) then
    - ✓ return(vI)
    - ✓ else return(R{1-I}.read())

# Impossibility [FLP85,LA87]

- ***Theorem:*** no *asynchronous* algorithm implements *consensus* among two processes using *registers*
  
- ***Corollary:*** no asynchronous algorithm implements Counter\* among two processes using *registers*

The **consensus number** of an object is the maximum number of processes than can solve consensus with it

Group	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
Period 1	1 H																	2 He
Period 2	3 Li	4 Be											5 B	6 C	7 N	8 O	9 F	10 Ne
Period 3	11 Na	12 Mg											13 Al	14 Si	15 P	16 S	17 Cl	18 Ar
Period 4	19 K	20 Ca	21 Sc	22 Ti	23 V	24 Cr	25 Mn	26 Fe	27 Co	28 Ni	29 Cu	30 Zn	31 Ga	32 Ge	33 As	34 Se	35 Br	36 Kr
Period 5	37 Rb	38 Sr	39 Y	40 Zr	41 Nb	42 Mo	43 Tc	44 Ru	45 Rh	46 Pd	47 Ag	48 Cd	49 In	50 Sn	51 Sb	52 Te	53 I	54 Xe
Period 6	55 Cs	56 Ba	57 La	* 72 Hf	73 Ta	74 W	75 Re	76 Os	77 Ir	78 Pt	79 Au	80 Hg	81 Tl	82 Pb	83 Bi	84 Po	85 At	86 Rn
Period 7	87 Fr	88 Ra	89 Ac	* 104 Rf	105 Db	106 Sg	107 Bh	108 Hs	109 Mt	110 Ds	111 Rg	112 Cn	113 Nh	114 Fl	115 Mc	116 Lv	117 Ts	118 Og
				* 58 Ce	59 Pr	60 Nd	61 Pm	62 Sm	63 Eu	64 Gd	65 Tb	66 Dy	67 Ho	68 Er	69 Tm	70 Yb	71 Lu	
				* 90 Th	91 Pa	92 U	93 Np	94 Pu	95 Am	96 Cm	97 Bk	98 Cf	99 Es	100 Fm	101 Md	102 No	103 Lr	

# Roadmap

- ☛ **(1) The main question**
- ☛ **(2) The bitcoin problem/object**
- ☛ **(3) The main answer**



# Payment Object (PO): Specification

- ☞  $\text{Pay}(a,b,x)$ : transfer amount  $x$  from  $a$  to  $b$  if  $a > x$   
(return ok; else return no)
- ☞ **Important.** Only the owner of  $a$  invokes  $\text{Pay}(a,*,*)$
- **Questions:**
  - - can PO be implemented asynchronously?
  - - what is the consensus number of PO?

# Snapshot: Specification


- A *snapshot* has operations *update()* and *scan()*; it maintains an array  $x$  of size  $N$
- *scan()*:
  - return( $x$ )
- *update( $i, v$ )*:
  - $x[i] := v$ ;
  - return(ok)

# Algorithm?

- The processes share one array of N registers  
Reg[1,..,N]
- *scan()*:
  - for j = 1 to N do
    - x[j] := Reg[j].read();
  - return(x)
- *update(i, v)*:
  - Reg[i].write(v); return(ok)

# Atomicity?


update(1,1) - ok



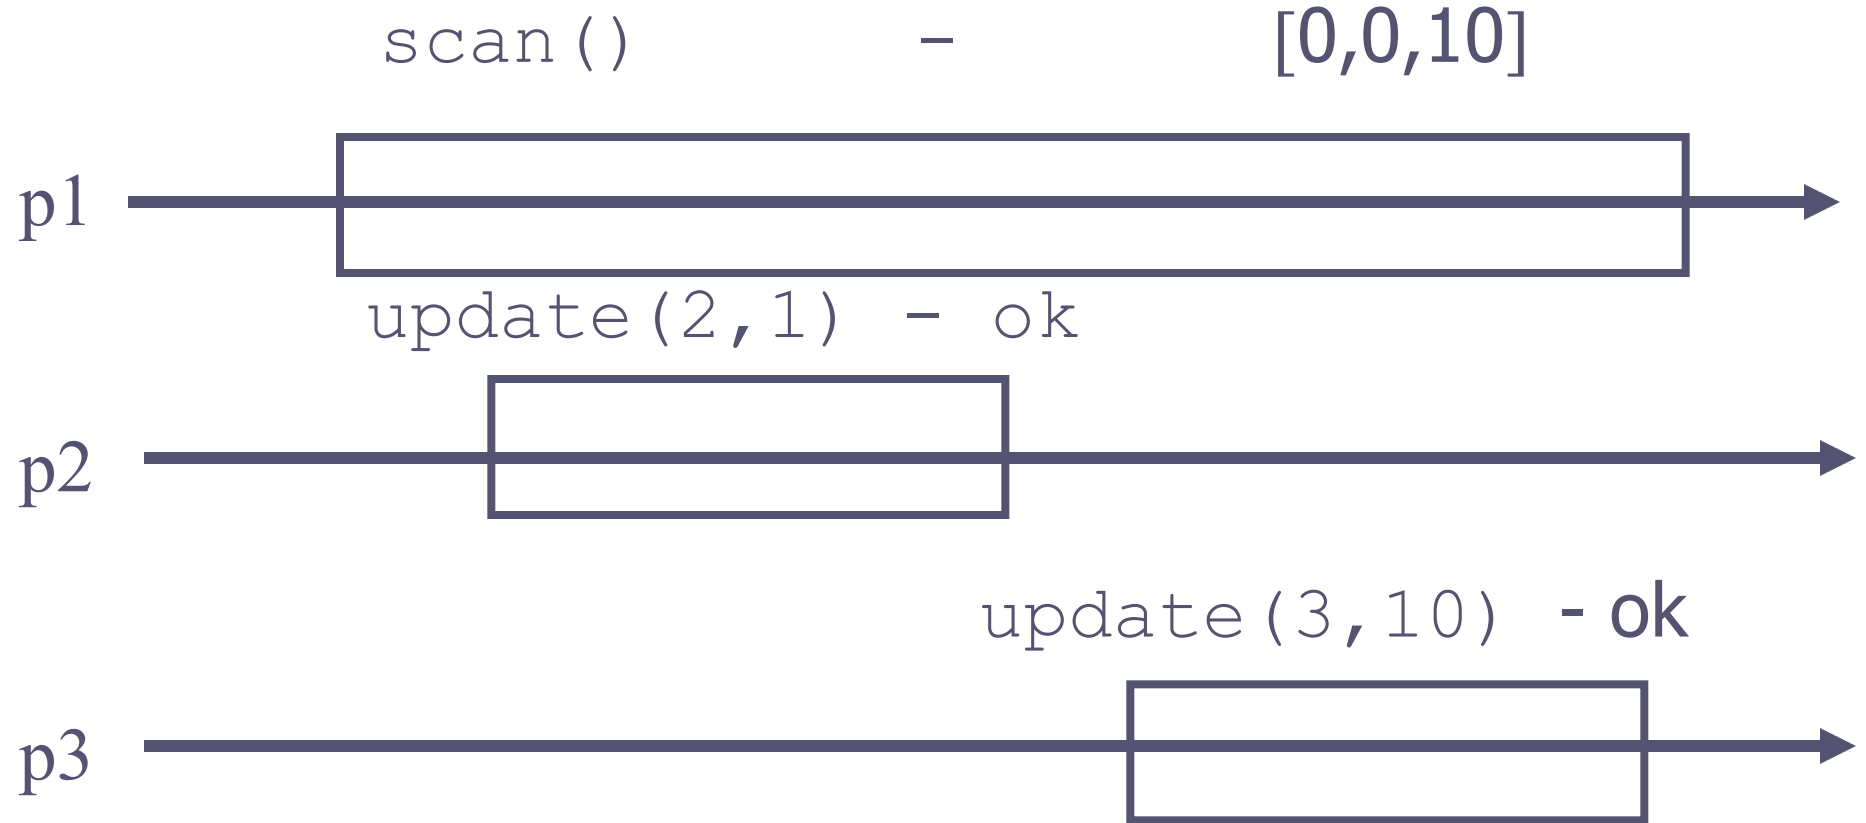
scan() - [1,0,2]



update(3,2) - ok



# Atomicity?



# Key idea for atomicity

- To *scan*, a process keeps reading the entire snapshot (i.e., *collecting*), until two arrays are the same

# Key idea for wait-freedom

- To update, scan then write the value and the scan
- To *scan*, a process keeps collecting and returns a collect if it did not change, or some collect returned by a concurrent *scan*

# The Payment Object: Algorithm

- Every process stores the sequence of its outgoing payments in its snapshot location
- To *pay*, the process scans, computes its current balance: if bigger than the transfer, updates and returns ok, otherwise returns no
- To *read*, scan and return the current balance

# **PO can be implemented asynchronously**

Consensus number of PO is 1

Consensus number of PO(k) is k



# Message Passing $\leftrightarrow$ Shared Memory Modulo Samples



# Roadmap

- ☛ **(1) The main question**
- ☛ **(2) The bitcoin problem/object**
- ☛ **(3) The main answer**

# Payment System (AT2)

☛ AT2\_S

☛ AT2\_D

☛ AT2\_R


- ☛ « The price of software reliability is the pursuit of the utmost simplicity » T. Hoare
- ☛ Number of lines of code: one order of magnitude less

# References

- **(1) The Consensus Number of a Cryptocurrency – ACM PODC 2019– Guerraoui et al**
- **(2) Scalable Byzantine Reliable Broadcast - DISC 2019 - Guerraoui et al**
- **(3) Online Payments by Merely Broadcasting Messages IEEE DSN 2020 - Collin et al**

# Journey to the center of distributed computing

- Bitcoin
- Blockchain
- Proof of work
- Smart contracts
- Ethereum



- Atomicity
- Wait-freedom
- Snapshot
- Consensus
- Quorums
- Secure Broadcast