Distributed Computing

The Computer

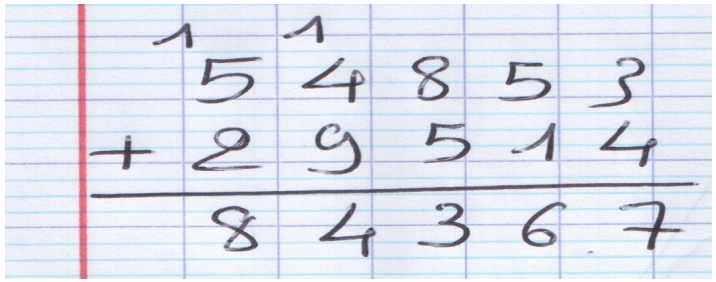

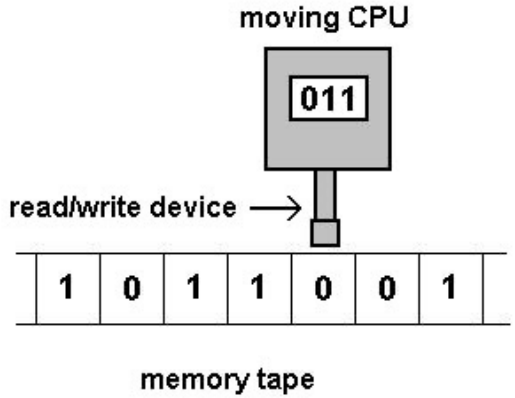

Algorithmi

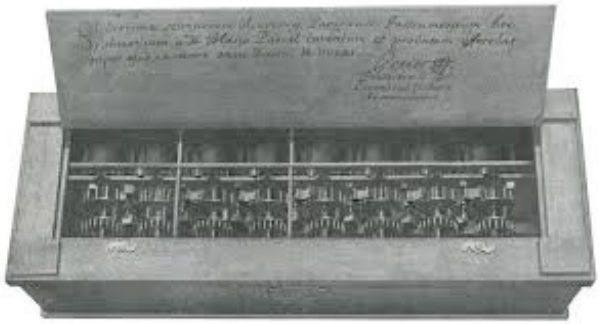

Turing

The Network

The March of the Penguins

The March of the Penguins

- The group is threatened if more than a threshold dies on the way back to the sea
- If a penguin starts its trip with a very low temperature, the probability that it reaches the sea is very low

The Escort

- Provide each penguin with a computing device to:
  - measure its temperature;
  - trigger an alert if a threshold (say 5) has a very low temperature

The Assumptions

- Every device holds a finite counter (<6)
  - Its initial value is 1 if the penguin has a low temperature and 0 otherwise
- A pair of devices communicate if they get close enough
  - Every pair of devices eventually meet

The Problem

- All devices eventually output “alert” iff at least 5 initial values are 1

Algorithm ?

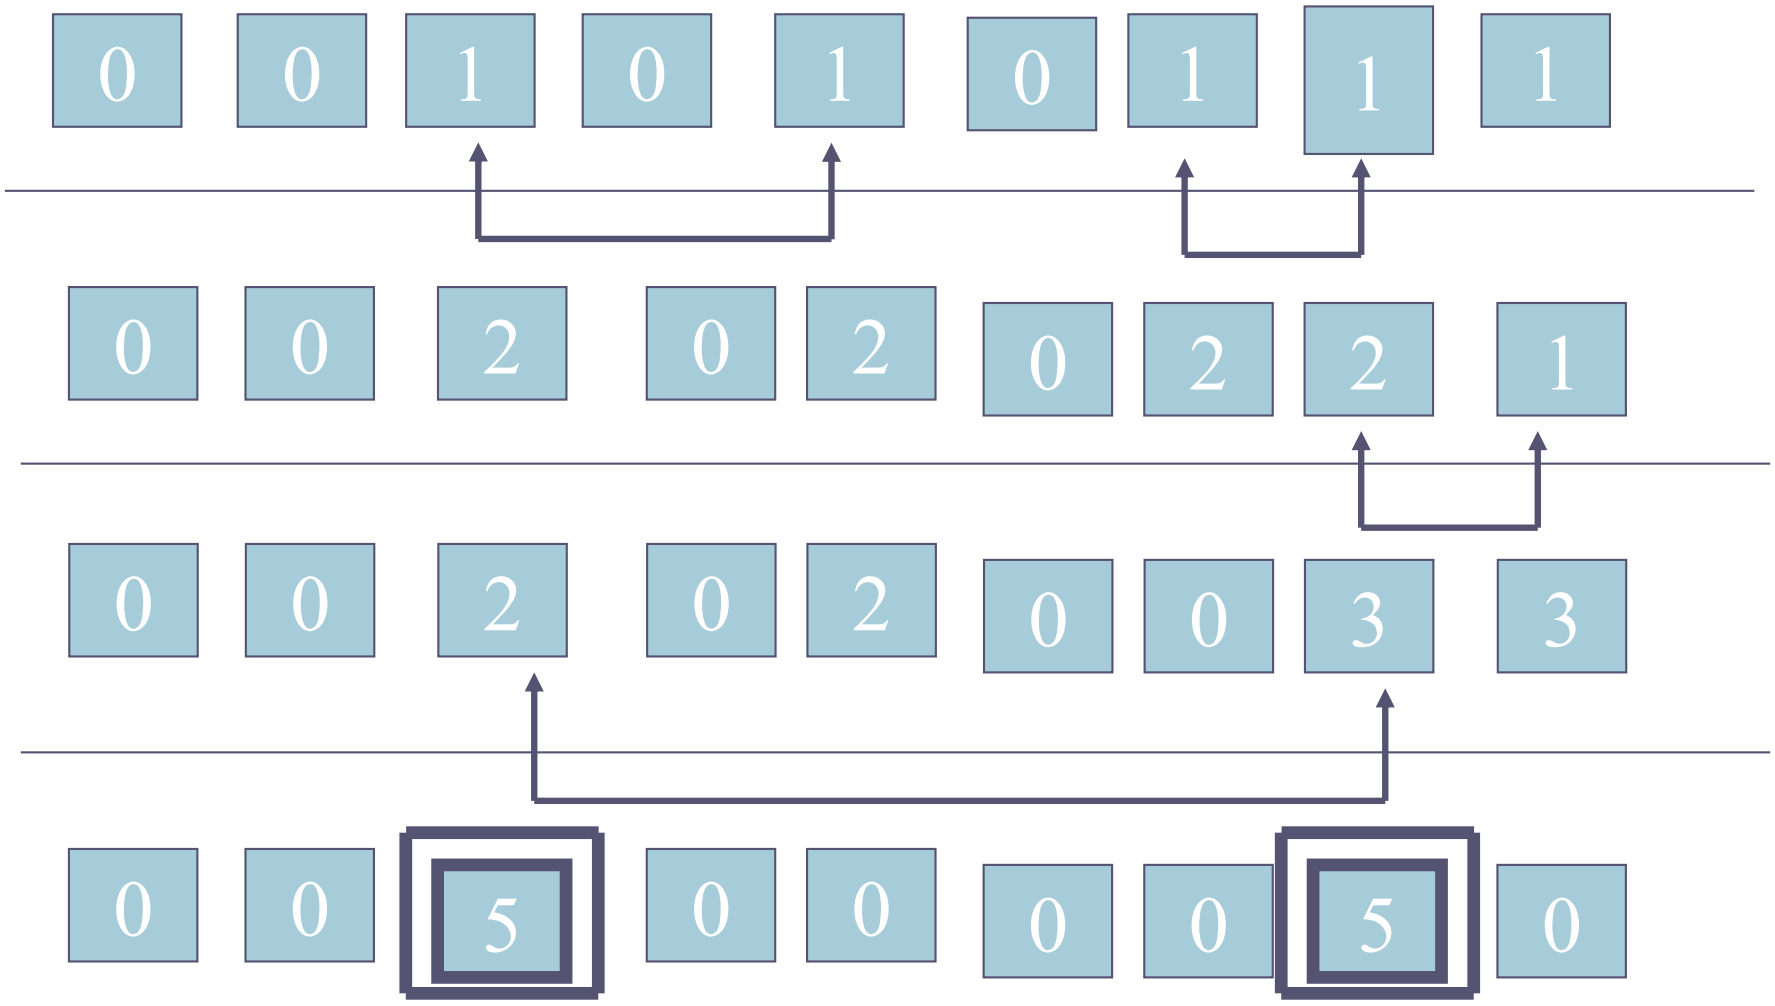

Algorithm?

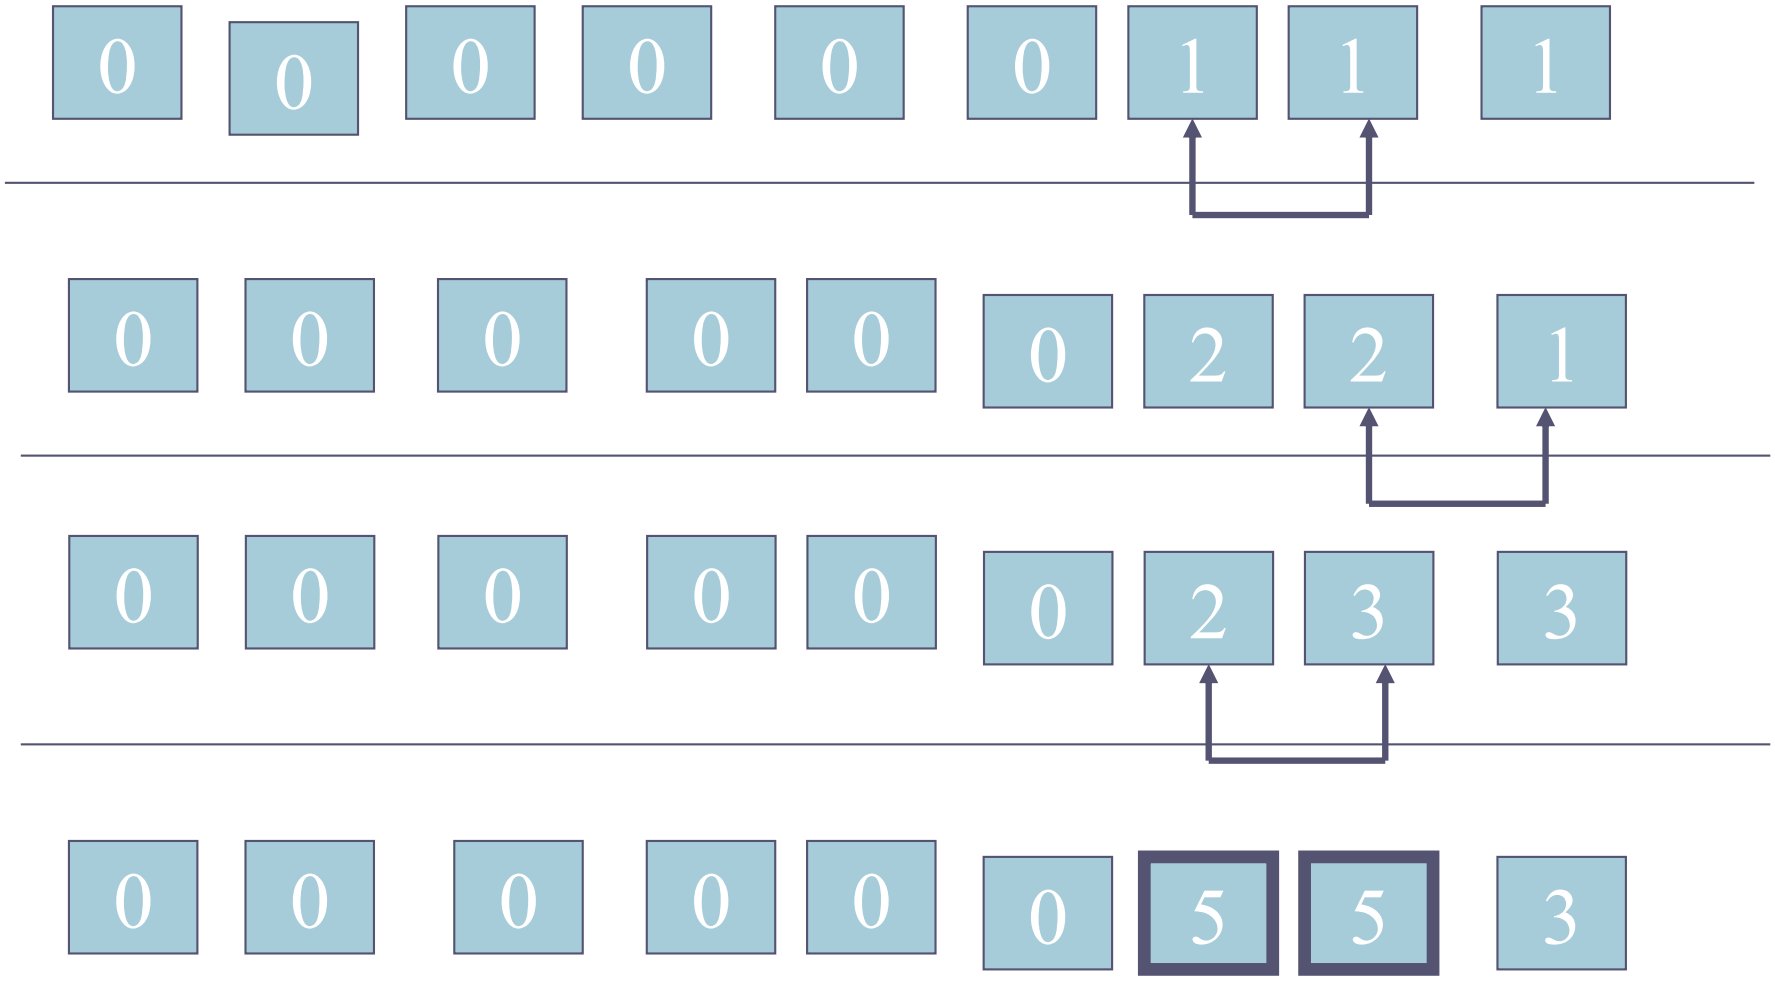

Algorithm

- When two devices meet, one keeps in its counter the sum of the values whereas the other puts it back to 0
- Any device with value 5 triggers the alert

Algorithm

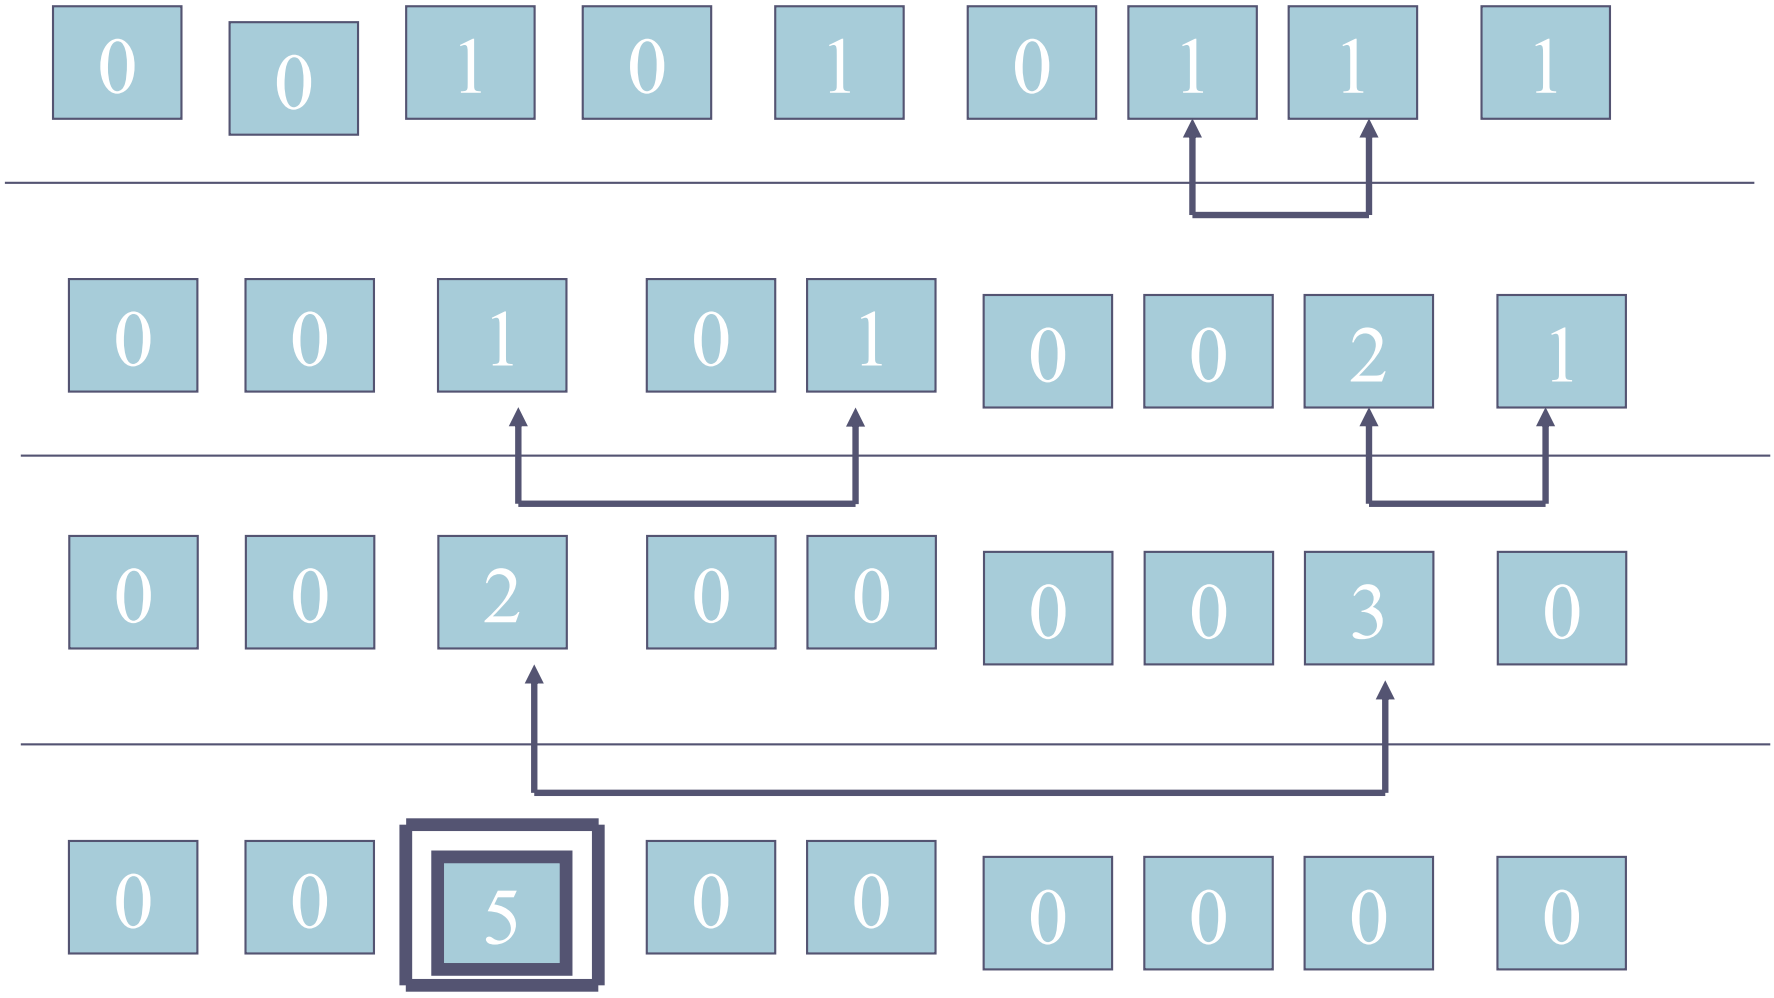

Algorithm

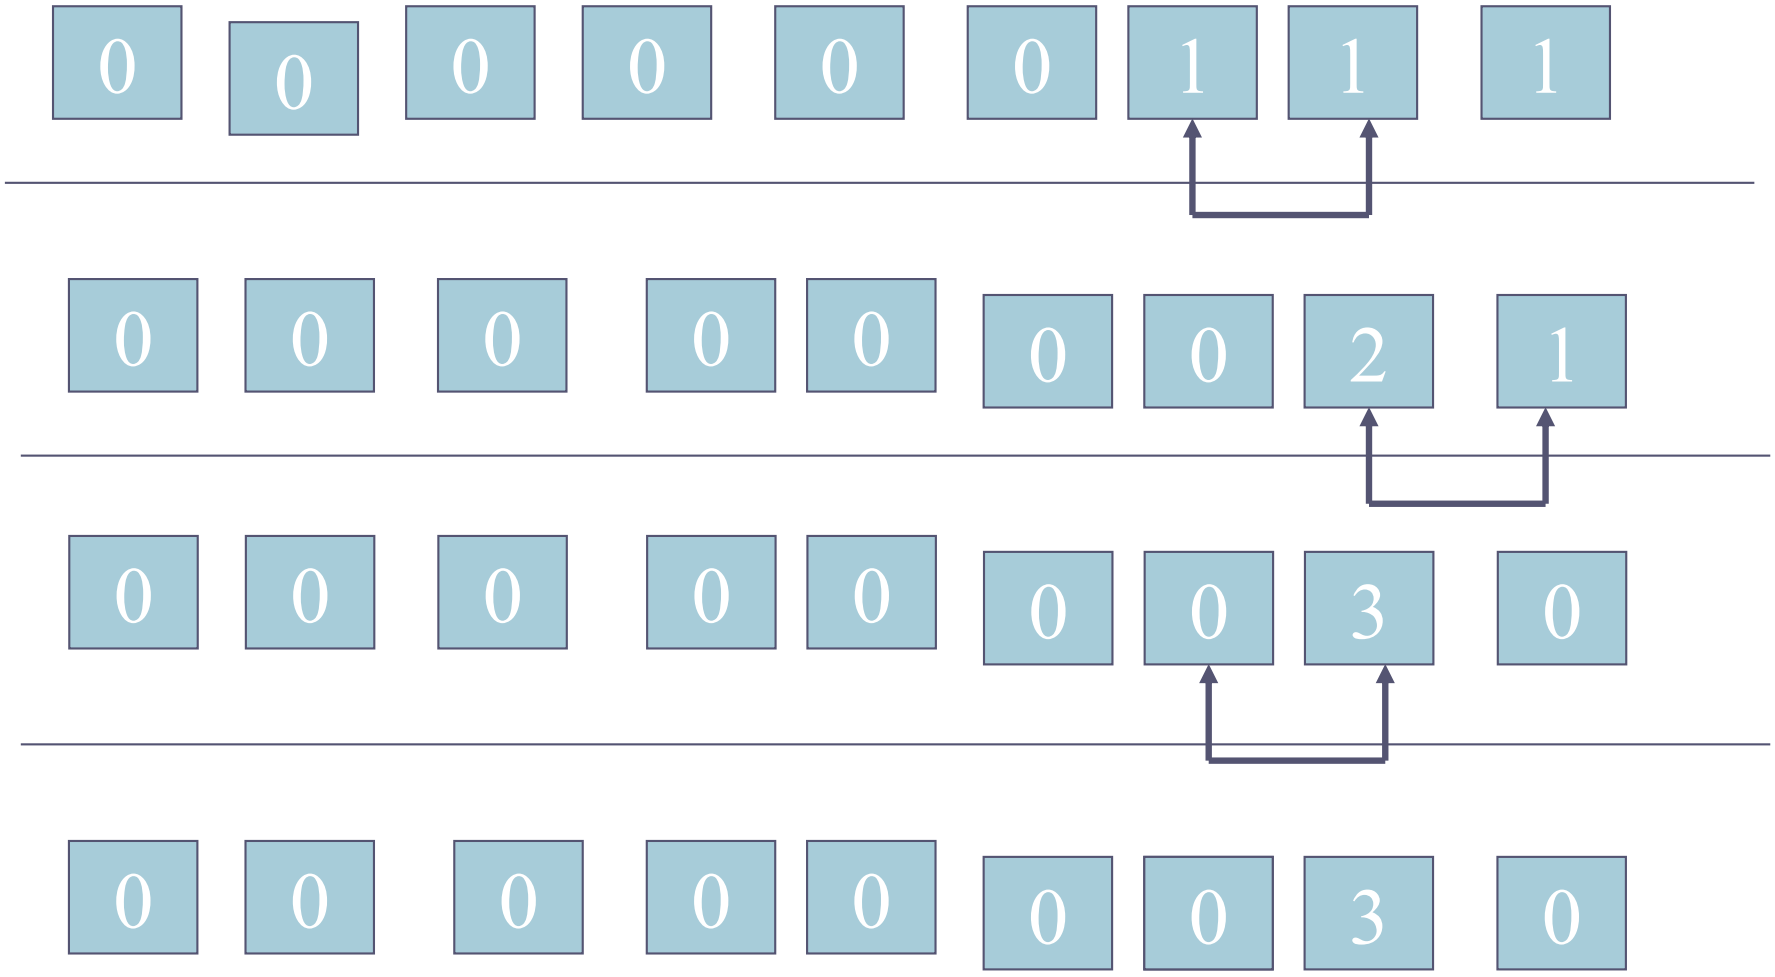

What If?

- One of the agents fail (say by crashing at some inappropriate time)
- An agent might crash exactly when it reaches value 4

Original Algorithm O

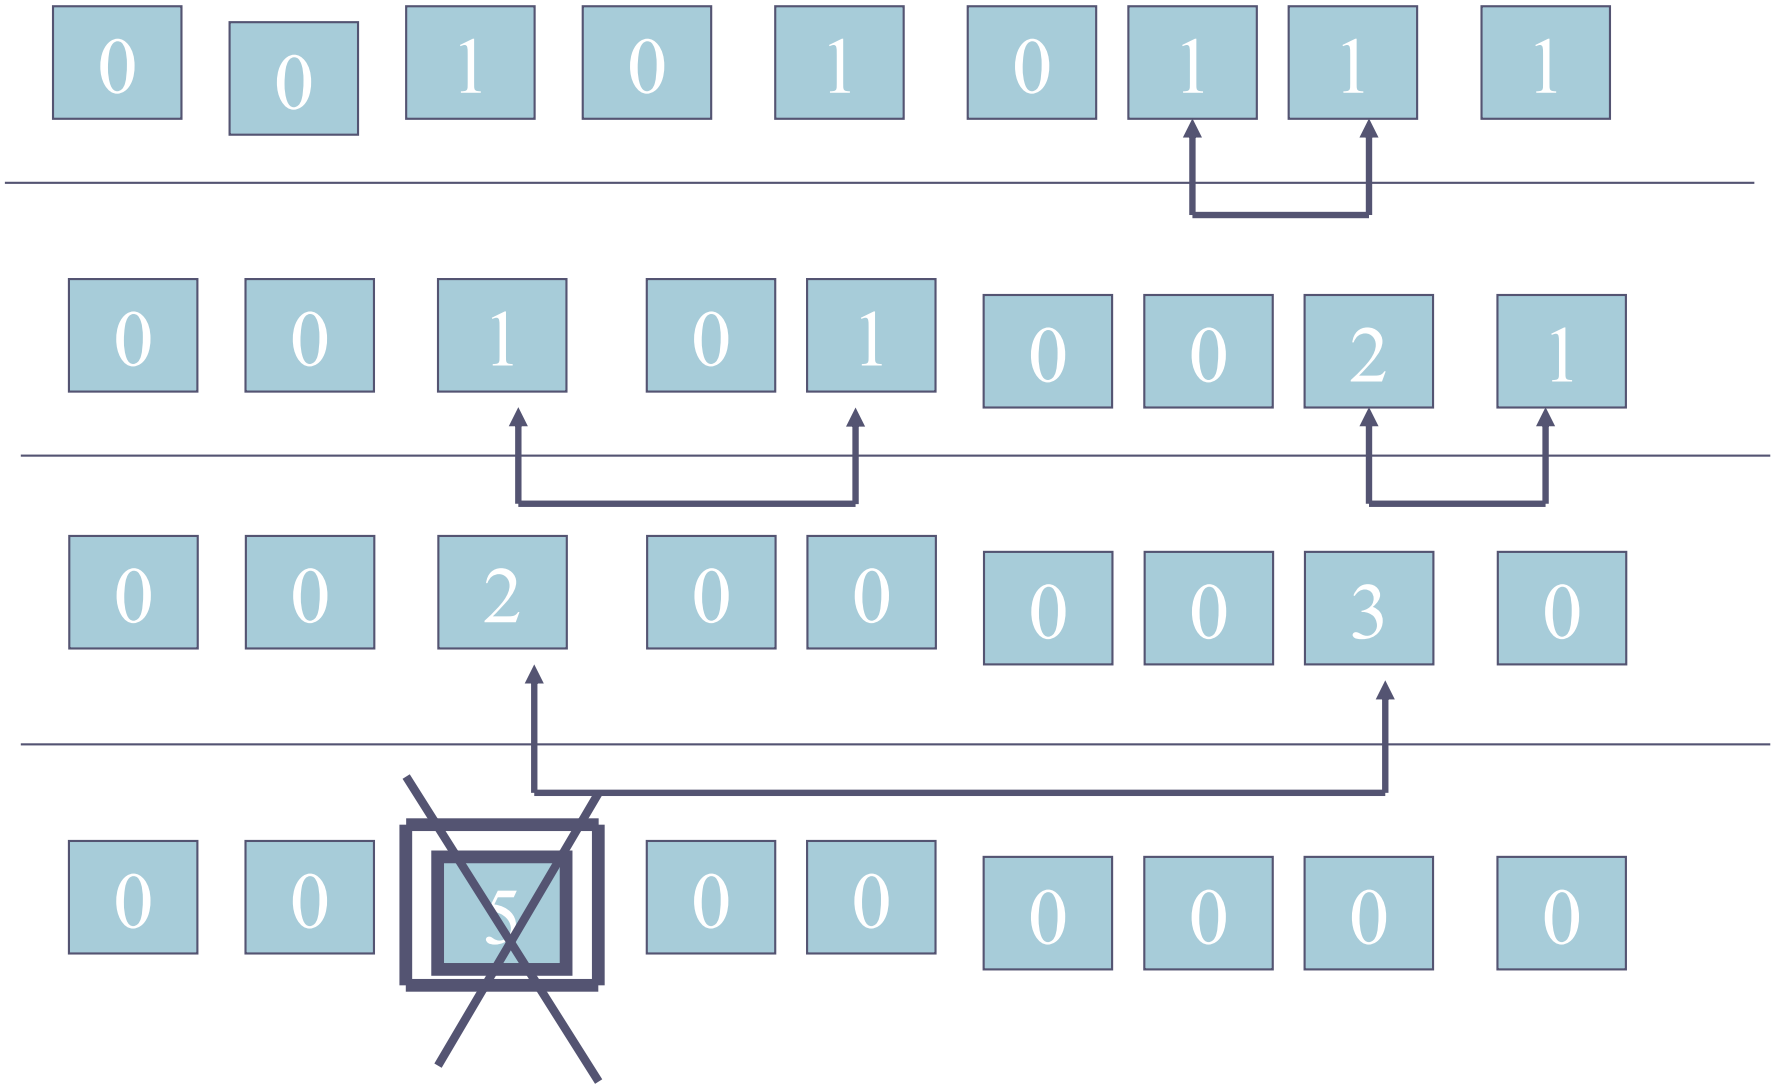

Robust Algorithm

- Every agent performs ***twice*** the original algorithm: O1 and O2
- When two agent communicate, one acts as the initiator for O1 and the other as the initiator for O2

What About Privacy?

- How can we ***hide*** the initial values from curious agents?

- How can we compute a result while preventing any agent from figuring out, at any point in time, any information besides its own input and the result of the computation?

How to ensure Privacy?

- An agent cannot use crypto (not even signatures because of anonymity)
- An agent can see the entire state of the agent it is interacting with: hence no secret keys are possible

Obfuscation

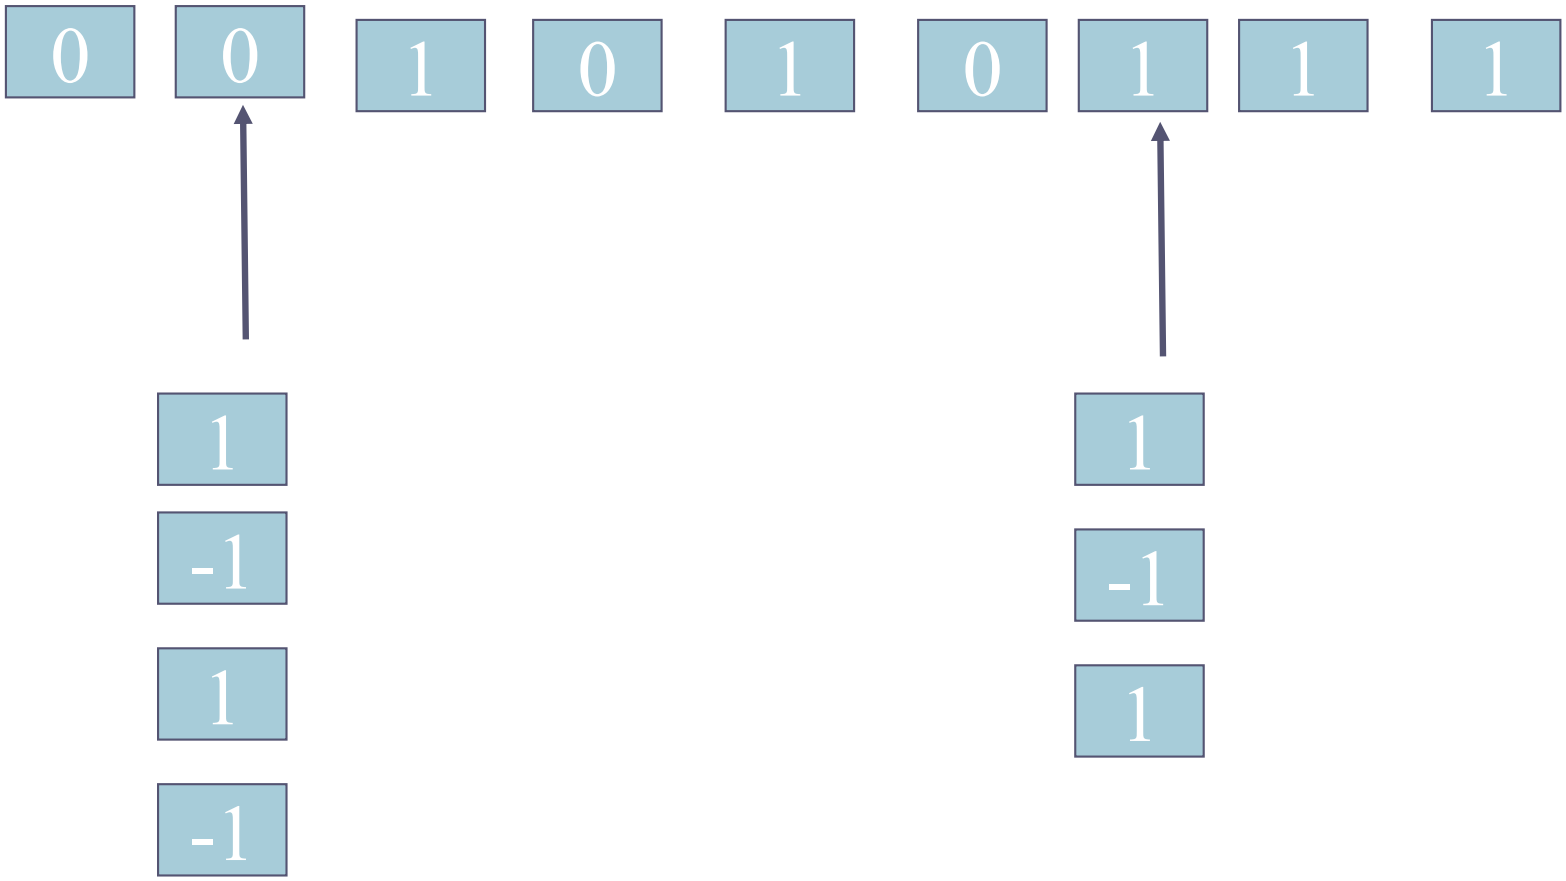

What if Agents can be Malicious?

- What can we compute with one malicious agent? (arbitrary transitions)

1. RDMA

- Remote shared / protected memory
- Consensus with 2f+1 and f+1 (vs 3f+1 and 2f+1) and 2 steps (vs 4 steps) –
- µ: SMR in 1µs / 1ms

[PODC / OSDI]

2. Cryptocurrencies

X000 implementations

[PODC / DSN]

3. Learning

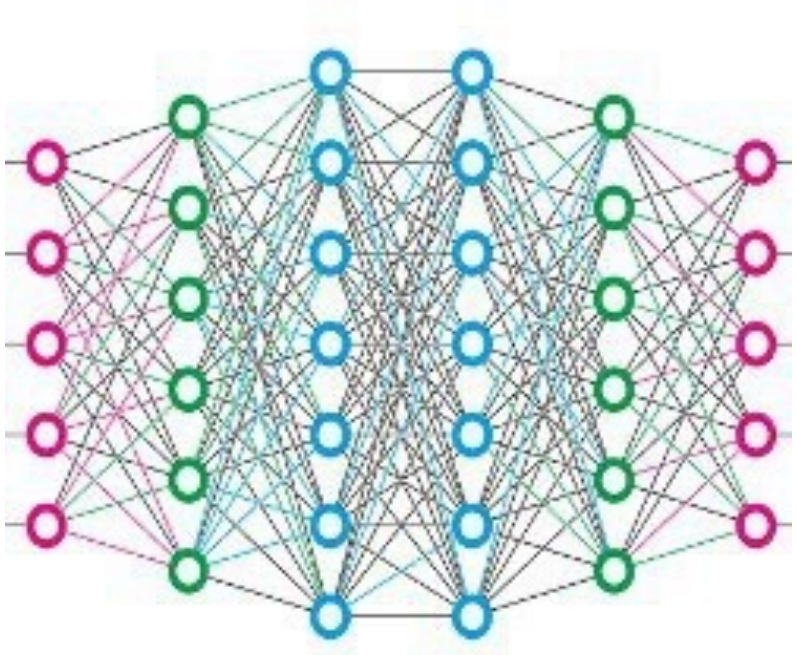

[ICML / NeurIPS / PODC]