

Distributed Computing


Xidian 2021

Prof. Mirosław Kutylowski

7: Consensus

Consensus problem:

- number of nodes is n , at most f of them faulty
- each node gets an input value $\in \{0, 1\}$
- finally, the correct nodes decide upon a common value
- the common value must be one of the input values



Impossibility of deterministic consensus in asynchronous model with failures

- arbitrary delays but not longer than 1 time unit
- each node executes a deterministic algorithm
- for simplicity – binary inputs only

Configuration

- 1.) States of all nodes
- 2.) Messages in transit

univalent configuration: decision value can be only one no matter what happens later

b-configuration: univalent, decision will be b


bivalent: decision will be 0 or 1

Lemma

Let $f > 1$.


There is at least one bivalent initial configuration

e.g. consider inputs ($n=8$):

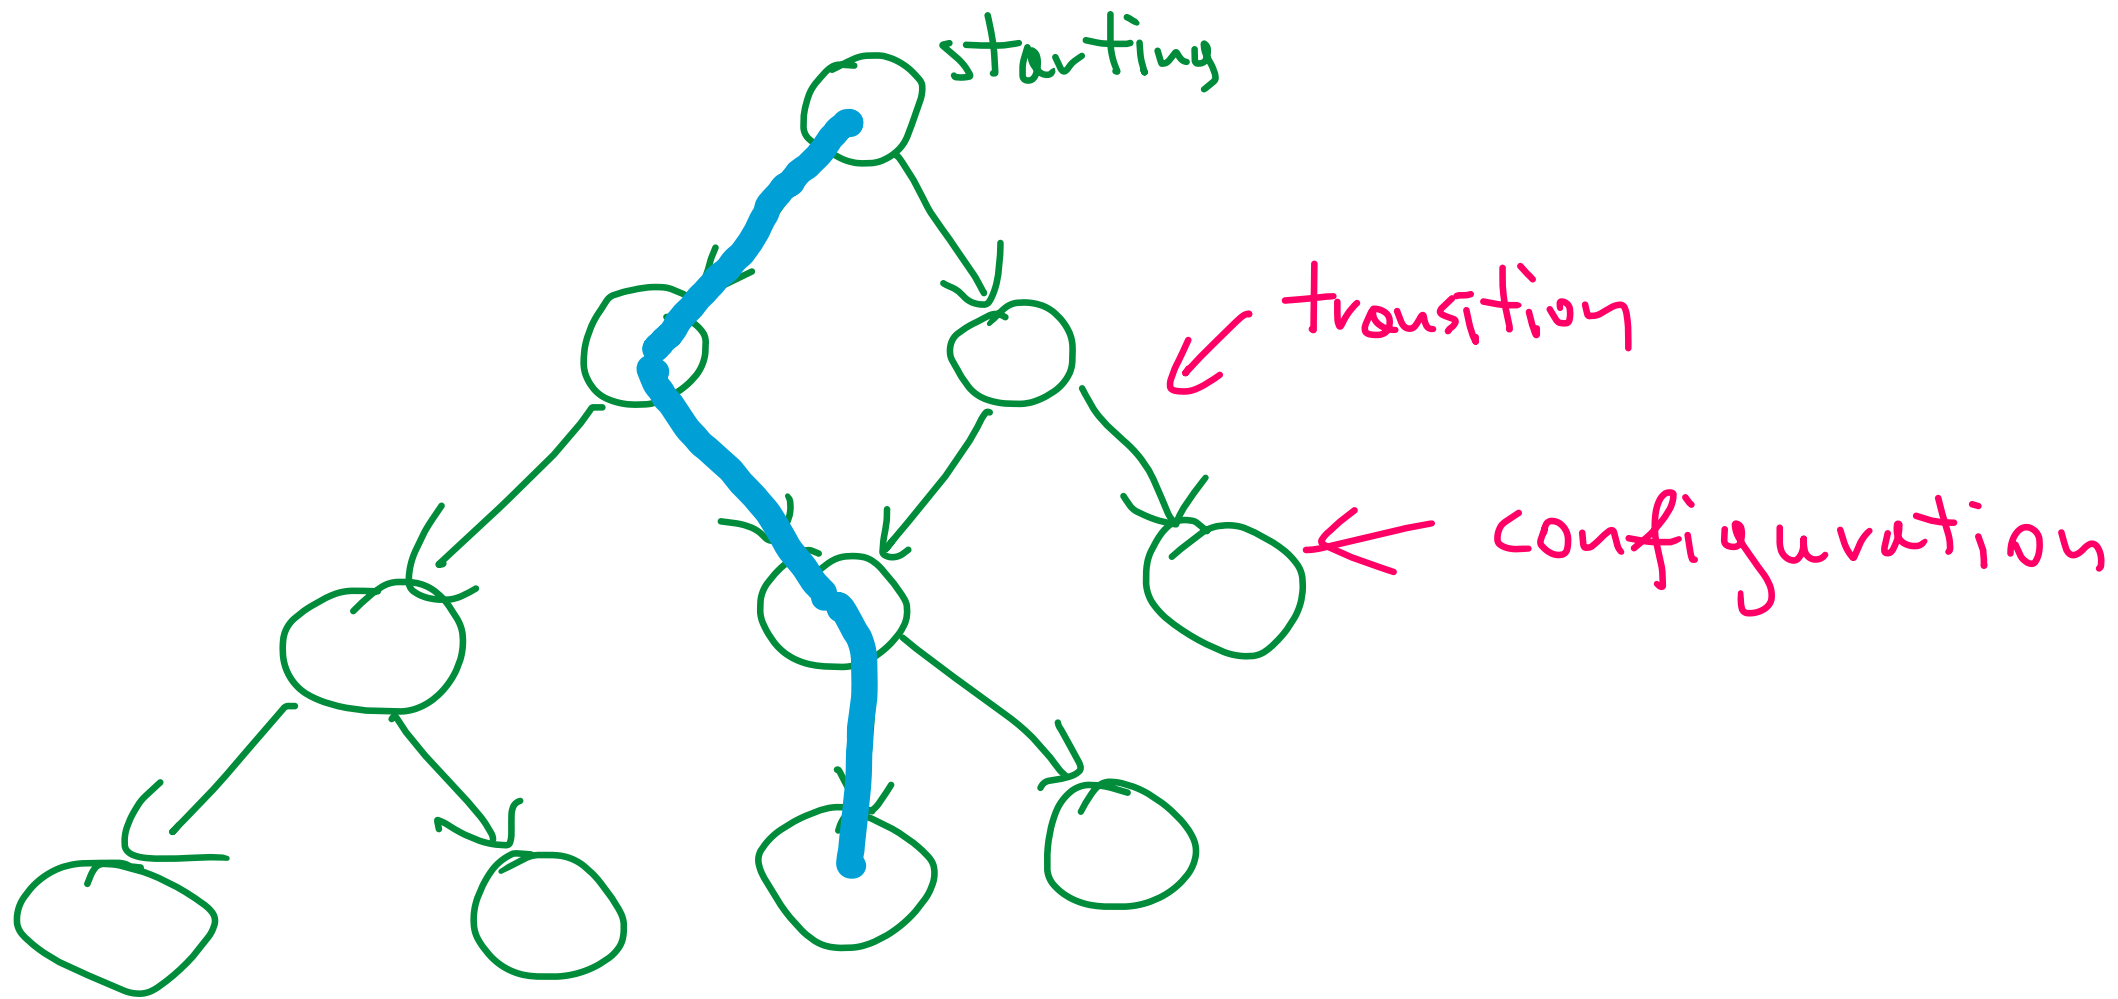


Configuration transition

- One message delivered, state changed, new messages despatched



Configuration directed graph

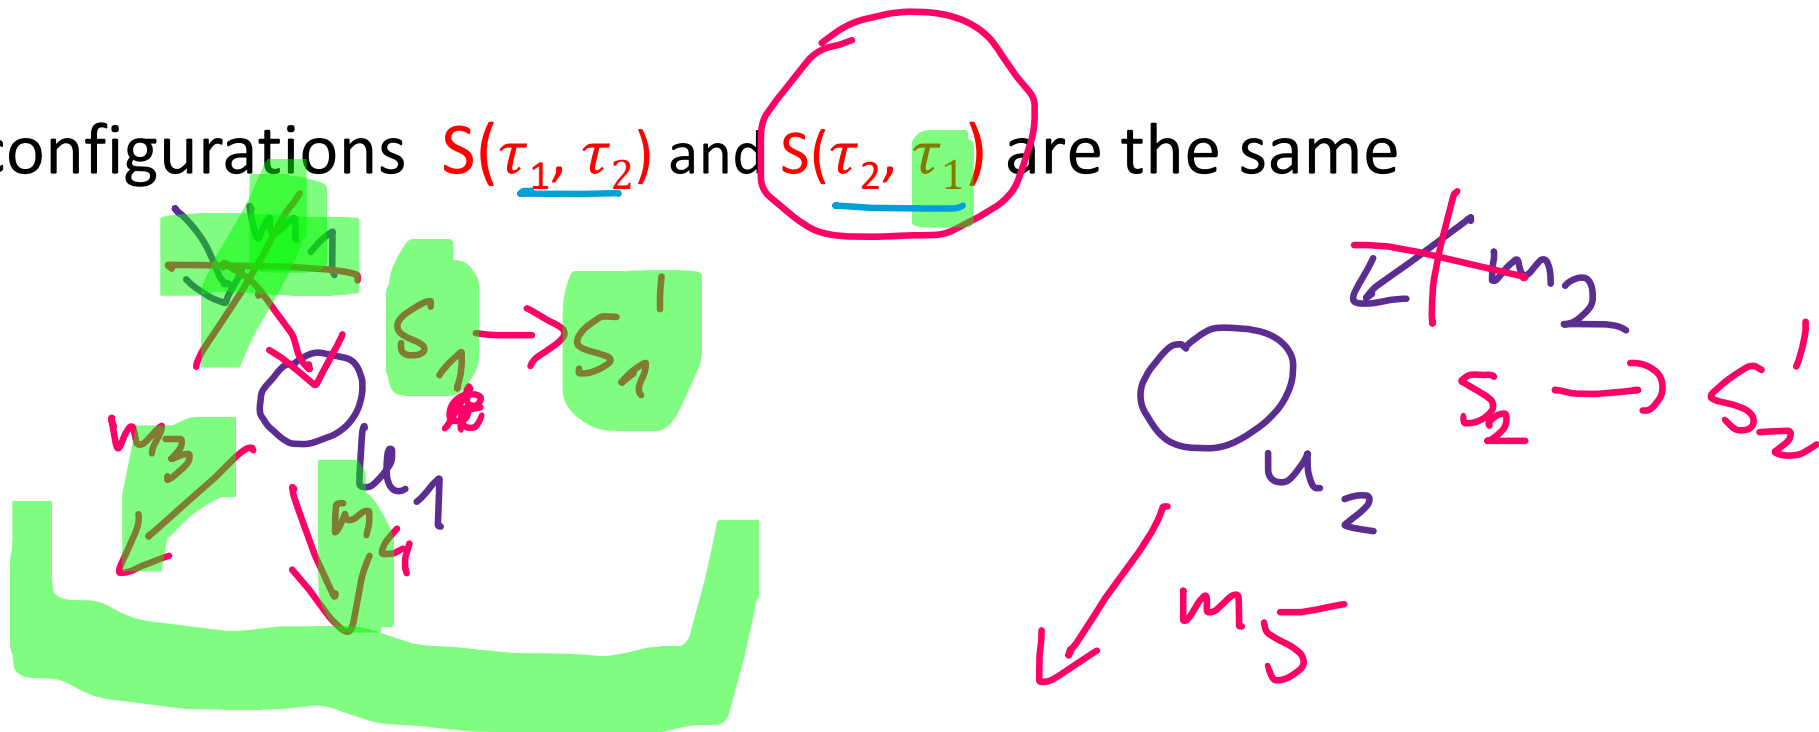


Two transitions

- in a state S two transitions are possible $\tau_1=(u_1,m_1)$ and $\tau_2=(u_2,m_2)$,
- where $u_1 \neq u_2$


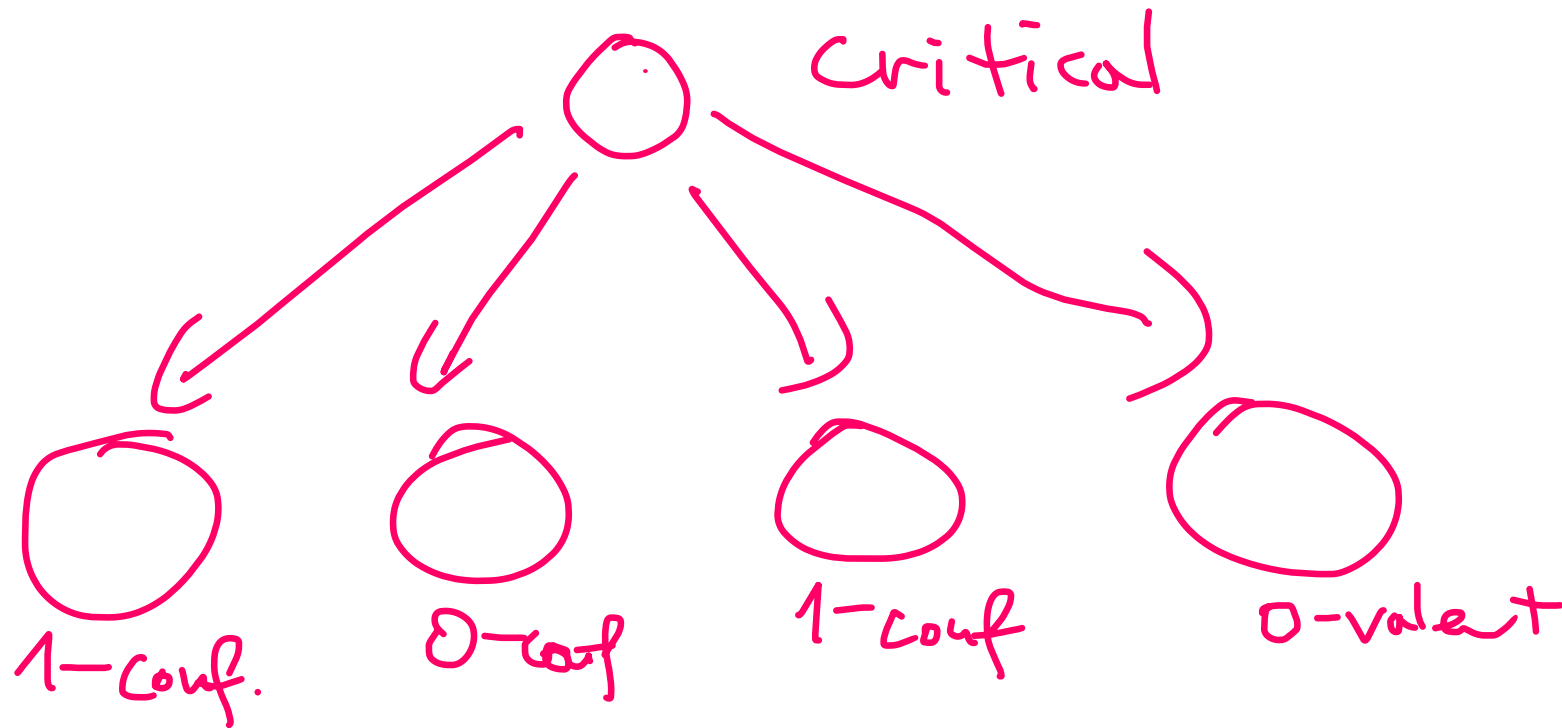
then

configurations $S(\tau_1, \tau_2)$ and $S(\tau_2, \tau_1)$ are the same

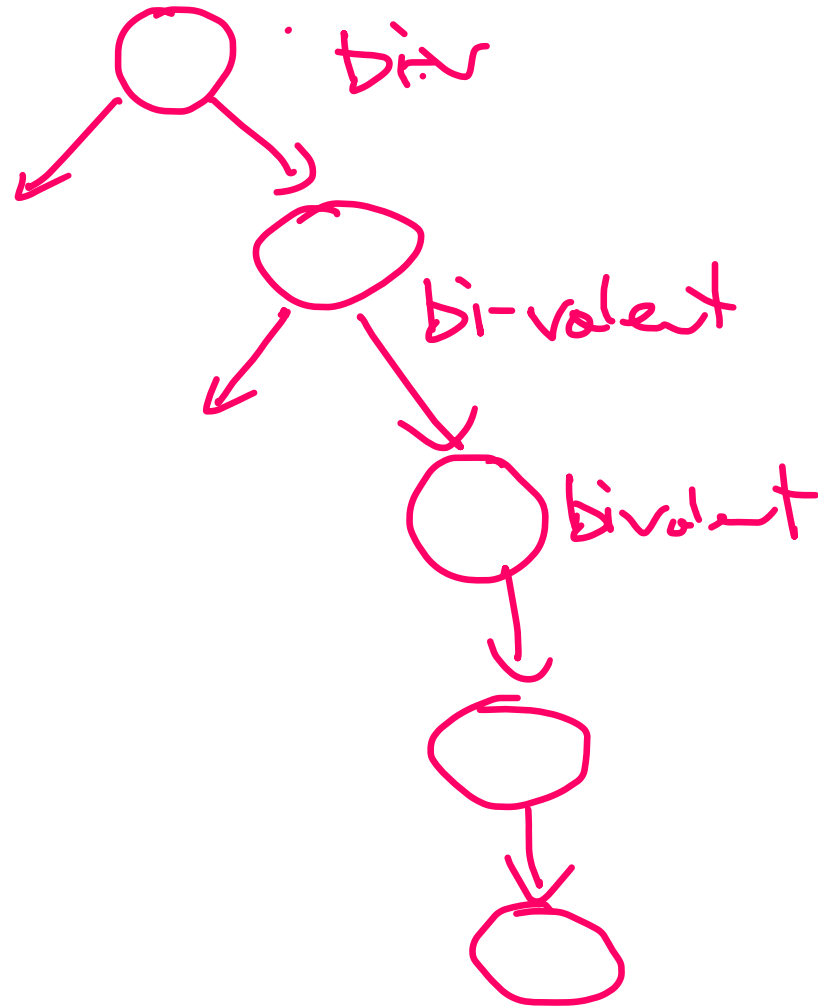


Critical configuration

A bivalent configuration such that each child configuration is univalent




Lemma 16.12. *If a system is in a bivalent configuration, it must reach a critical configuration within finite time, or it does not always solve consensus.*



Lemma 16.13. *If a configuration tree contains a critical configuration, crashing a single node can create a bivalent leaf; i.e., a crash prevents the algorithm from reaching agreement.*


- Let C be a critical configuration
- Let $\tau_1 = (u_1, m_1)$ and $\tau_2 = (u_2, m_2)$ be transitions such that
 - $C(\tau_1)$ is 1-valent
 - $C(\tau_2)$ is 0-valent
- If $u_1 \neq u_2$ then $C(\tau_1 \tau_2) = C(\tau_2 \tau_1)$, so it is 1-valent and 0-valent at the same time !?
- Consequently: $u_1 = u_2$



Corollary

Crash
it

All transitions from a critical configuration C must involve the same node



FINAL OBSERVATION

crush this node at configuration C

No consensus can be reached

Theorem 16.14. *There is no deterministic algorithm which always achieves consensus in the asynchronous model, with $f > 0$.*

Randomized consensus

Deterministic solution is impossible, but ...

... what about a randomized algorithm with random execution time
(and always correct output)? (Las Vegas algorithm)

Monte Carlo - fixed time
results sometimes are
wrong

Algorithm 16.15 Randomized Consensus (Ben-Or)

- 1: $v_i \in \{0, 1\}$ \triangleleft input bit
- 2: round = 1
- 3: decided = false
- 4: Broadcast myValue(v_i , round)
- 5: while true do


myValue(0, 1)


Propose

```
6: Wait until a majority of myValue messages of current round arrived
7: if all messages contain the same value  $v$  then
8:   Broadcast propose( $v$ , round)
9: else
10:   Broadcast propose( $\perp$ , round)
11: end if

12: if decided then
13:   Broadcast myValue( $v_i$ , round+1)
14:   Decide for  $v_i$  and terminate
15: end if
```

nodes
Myvalue
~~propose~~





Adapt

```

16: Wait until a majority of propose messages of current round arrived
17: if all messages propose the same value v then
18:    $v_i = v$ 
19:   decide = true
20: else if there is at least one proposal for v then
21:    $v_i = v$ 
22:   else
23:     Choose  $v_i$  randomly, with  $Pr[v_i = 0] = Pr[v_i = 1] = 1/2$ 
24:   end if
25:   round = round + 1
26:   Broadcast myValue( $v_i$ , round)
27: end while
  
```

w

shared coin algo

u-decided, majority with v
w-another majority, they intersect!

All 0 input

- It finishes quite quickly!

Lemma 16.16. *As long as no node sets decided to true, Algorithm 16.15 always makes progress, independent of which nodes crash.*

What happens if
all input bits are
the same?

```
6:  Wait until a majority of myValue messages of current round arrived
7:  if all messages contain the same value  $v$  then
8:    Broadcast propose( $v$ , round)
9:  else
10:   Broadcast propose( $\perp$ , round)
11: end if

12: if decided then
13:   Broadcast myValue( $v_i$ , round+1)
14:   Decide for  $v_i$  and terminate
15: end if
```

Adapt

```
16: Wait until a majority of propose messages of current round arrived
17: if all messages propose the same value  $v$  then
18:    $v_i = v$ 
19:   decide = true
20: else if there is at least one proposal for  $v$  then
21:    $v_i = v$ 
22: else
23:   Choose  $v_i$  randomly, with  $Pr[v_i = 0] = Pr[v_i = 1] = 1/2$ 
24: end if
25: round = round + 1
26: Broadcast myValue( $v_i$ , round)
27: end while
```

Essential case – both 0 and 1 as input

- any consensus value is ok
- no proposals for different bits in the same round
- u – the first node that decides, say: for v , at round r
 - no decide at this round to a different v'
 - u terminates at the round $r+1$

case: $\text{Propose}(v, r)$
 $\text{Propose}(1, r)$

Termination

- u – the first node that decides, say: for v , at round r
 - no decide at this round to a different v'
 - u terminates at the round $r+1$
- An other node u'
 - If decides, then only for the same v
 - otherwise, it has heard at least one $\text{propose}(v,r)$ and sets $v_{u'}$ to v
- So all nodes broadcast v at the end of round r

Runtime


quite bad, since the choices in line 23 must be almost the same

□ exponential time

It would help to have a public coin and to toss it together!

(it seems to be a hopeless task in a distributed system)

input



Shared coin

Random variable that is equal to 0 with probability $> \frac{1}{4}$ and $\frac{1}{8}$ and 1 with probability $> \frac{1}{4}$

(the value of the shared coin is the same for all nodes)

Algorithm 16.22 Shared Coin (code for node u)

- 1: Choose local coin $c_u = 0$ with probability $1/n$, else $c_u = 1$
 - 2: Broadcast $\text{myCoin}(c_u)$
 - 3: Wait for $n - f$ coins and store them in the local coin set C_u
 - 4: Broadcast $\text{mySet}(C_u)$
 - 5: Wait for $n - f$ coin sets
 - 6: if at least one coin is 0 among all coins in the coin sets then
 - 7: return 0
 - 8: else
 - 9: return 1
 - 10: end if
-

We show the correctness of the algorithm for $f < n/3$. To simplify the proof we assume that $n = 3f + 1$, i.e., we assume the worst case.

tossing own coin:

$$\left(1 - \frac{1}{n}\right)^n$$

SS

$$\frac{1}{e}$$

$$> \frac{1}{4}$$

no zero

\Rightarrow

no coin set
contains 0

\Rightarrow decisions = 1

Lemma 16.23. *Let u be a node, and let W be the set of coins that u received in at least $f + 1$ different coin sets. It holds that $|W| \geq f + 1$.*

C be the multiset of coins received by u .

$$|C| = (n - f)^2$$

$$|C| \leq f \cdot (n - f) + (n - f) \cdot f = 2f(n - f)$$

$$, n - f > 2f$$

$$|C| \leq 2f(n - f) < (n - f)^2 = |C|$$

Assuming that the lemma is false

Lemma 16.24. *All coins in W are seen by all correct nodes.*

Theorem 16.25. *If $f < n/3$ nodes crash, Algorithm 16.22 implements a shared coin.*

With probability $(1 - 1/n)^n \approx 1/e \approx 0.37$ all nodes chose their local coin equal to $\bar{1}$ (Line 1), and in that case $\bar{1}$ will be decided.

With probability $1 - (1 - 1/n)^{|W|}$ there is at least one 0 in W

$$|W| \geq f + 1 \approx n/3,$$

$$1 - (1 - 1/n)^{n/3} \approx 1 - (1/e)^{1/3} \approx 0.28$$