

Distributed Computing

Xidian 2021

Prof. Mirosław Kutylowski

8: Byzantine agreement

Consensus with adversarial nodes

A Byzantine node can behave in an arbitrary way, not following the protocol.

Consensus value – must be the value proposed by ONE of correct nodes.

One Byzantine node only

Algorithm 17.9 Byzantine Agreement with $f = 1$.

1: Code for node u , with input value x :

Round 1

2: Send $\text{tuple}(u, x)$ to all other nodes

3: Receive $\text{tuple}(v, y)$ from all other nodes v

4: Store all received $\text{tuple}(v, y)$ in a set S_u

Round 2

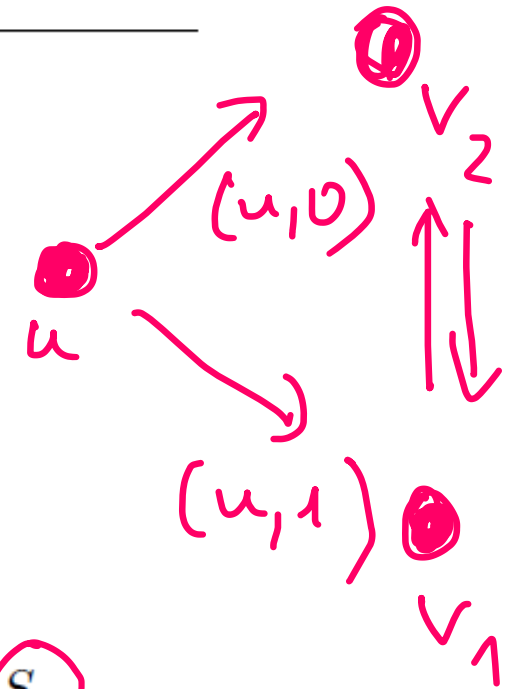
5: Send set S_u to all other nodes

6: Receive sets S_v from all nodes v

7: $T =$ set of $\text{tuple}(v, y)$ seen in at least two sets S_v , including own S_u

8: Let $\text{tuple}(v, y) \in T$ be the tuple with the smallest value y

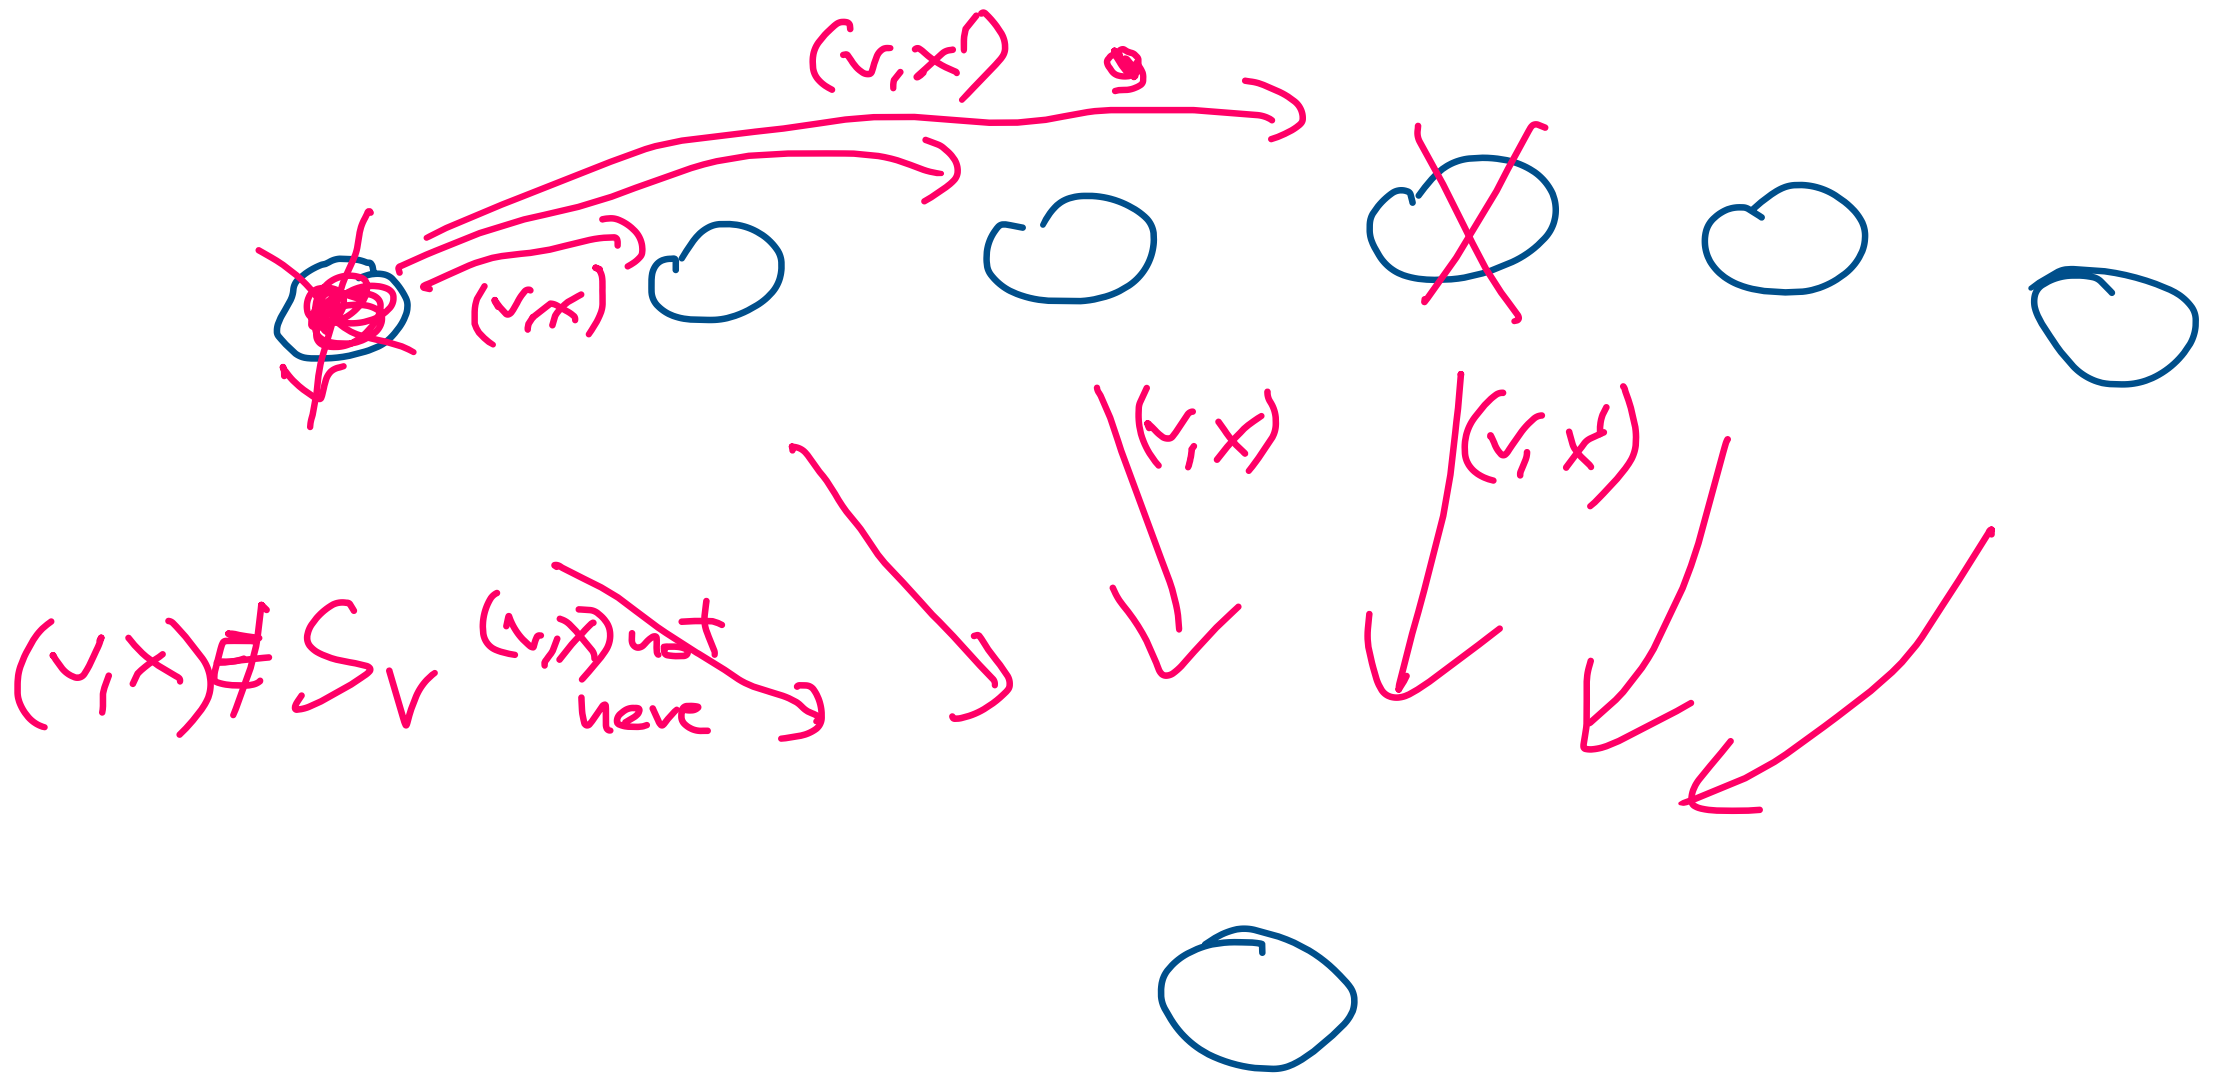
9: Decide on value y





tuple
 $g = (v, x)$

honestv: but not in two S_u



Property for $n \geq 4$

All nodes hold the same T :

- If there are 3 correct nodes, then 2 correct values appear in 2 sets S_u
- The values sent by the Byzantine node are not in T if it sends different values.

Corollary: one value will be chosen

For $n=3$ one cannot reach agreement

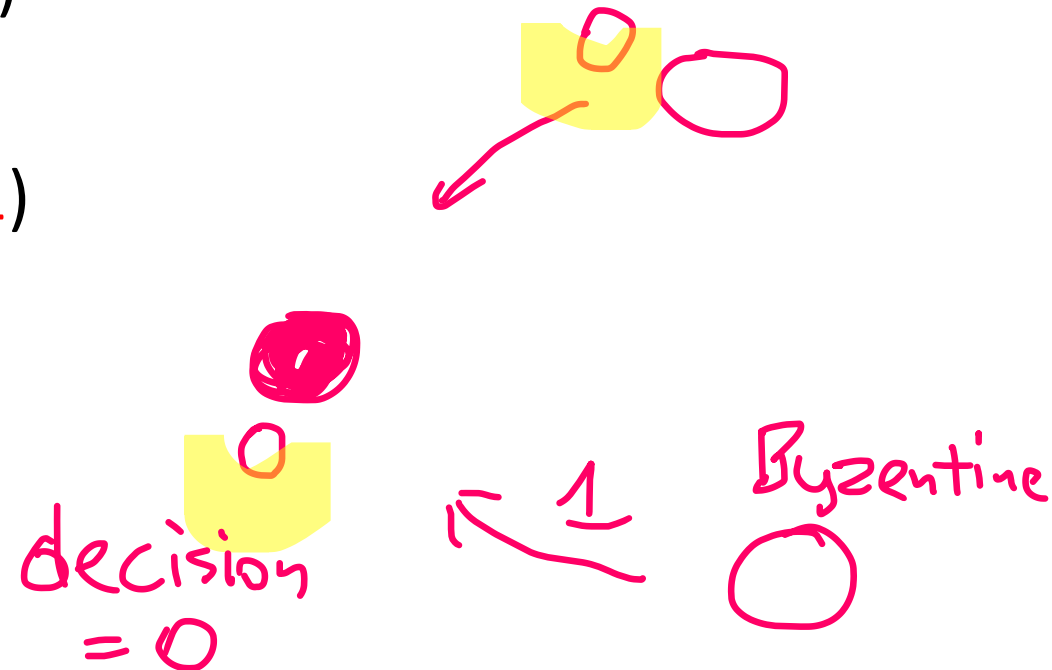
Property. A correct node u must decide on its value if node w supports it (while v disagrees – v might be byzantine)

Honest: u (with input 0) and v (with input 1)


Byzantine: w says 0 to u

w says 1 to w

u decides on 0 , v decides on 1



Byzantine?



Byzantine



decision 1

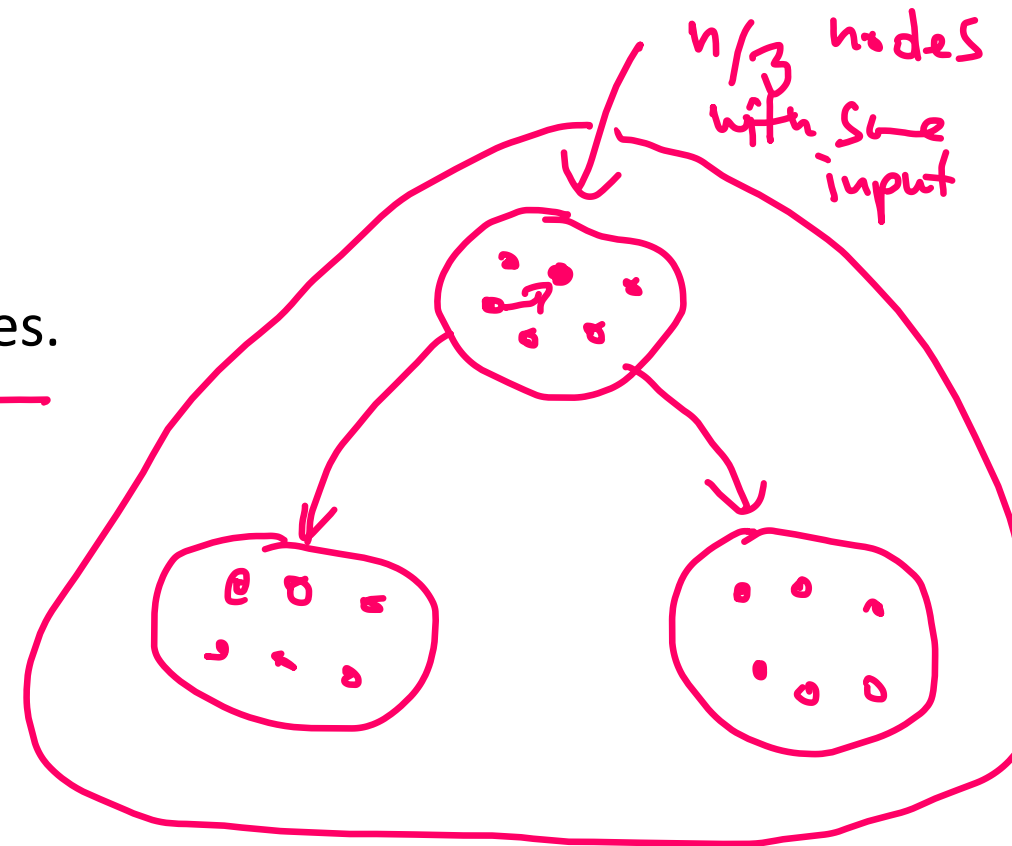
Fundamental theorem

Theorem 17.13. *A network with n nodes cannot reach byzantine agreement with $f \geq n/3$ byzantine nodes.*

Proof.

Assume there is such agreement protocol A for n nodes.

Emulate it by 3 nodes – each emulating $n/3$ nodes



King's algorithm

Algorithm 17.14 King Algorithm (for $f < n/3$)

1: $x =$ my input value

2: for phase = 1 to $f + 1$ do

Round 1

3: Broadcast value(x)

Round 2

4: if some value(y) at least $n - f$ times then

5: Broadcast propose(y)

6: end if

7: if some propose(z) received more than f times then

8: $x = z$

9: end if

Round 3

10: Let node v_i be the predefined king of this phase i

11: The king v_i broadcasts its current value w

12: if received strictly less than $n - f$ propose(x) then

13: $x = w$

14: end if

15: end for

King 1

King 2

⋮

King($f+1$)

Same value input

Algorithm 17.14 King Algorithm (for $f < n/3$)

1: $x =$ my input value

2: for phase = 1 to $f + 1$ do

Round 1

3: Broadcast value(x)

Round 2

4: if some value(y) at least $n - f$ times then

5: Broadcast propose(y)

6: end if

7: if some propose(z) received more than f times then

8: $x =$ z

9: end if

Round 3

10: Let node v_i be the predefined king of this phase i

11: The king v_i broadcasts its current value w

12: if received strictly less than $n - f$ propose(x) then

13: $x =$ w

14: end if

15: end for

The king's value
will never be
accepted.

Nodes stick to
their input values

$x = 0$ for all

value(0) for lowest nodes

$$n - f > f$$

$$n > 2f$$

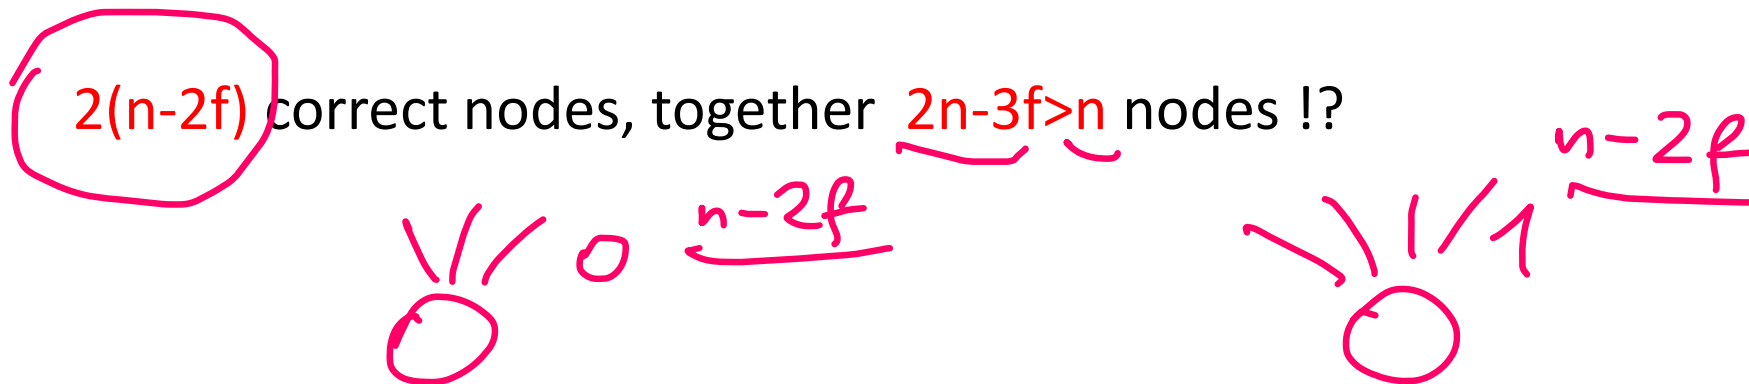
Proposing values

Lemma 17.16. *If a correct node proposes x , no other correct node proposes y , with $y \neq x$, if $n > 3f$.*

Proof

A correct node must receive $n-f$ messages, so at least $n-2f$ from correct nodes

Two different proposals require the following number of nodes:



Kings

At least one king correct

Lemma: after the round with the correct king the nodes do not change their values anymore.

Proof

Easy case: all correct nodes take the king's value

Crucial case: line 12 executed by only SOME correct nodes
is the King's value the same as for other correct nodes??

Crucial case

- thereby we assume that some node received a proposal at least $n-f$ times
- ... so from at least $n-2f$ correct nodes
- as $n-2f > f$ every node received a proposal from more than f nodes
- ... and therefore every correct node will change the value according to the proposal

Lower bound – no algorithm with f rounds determining minimum

- u_1 sends to u_2 at round 1 and then crashes
- u_2 sends to u_3 at round 2 and then crashes
- ...
- u_{f-1} sends to u_f at round f and then crashes

- Only u_f will see the minimum value

Byzantine agreement, asynchronous model

Algorithm 17.21 Asynchronous Byzantine Agreement (Ben-Or, for $f < n/9$)

1: $x_i \in \{0, 1\}$ \triangleleft input bit
2: $r = 1$ \triangleleft round
3: $\text{decided} = \text{false}$
4: Broadcast **propose**(x_i, r)
5: **repeat**
6: Wait until $n - f$ **propose** messages of current round r arrived
7: **if** at least $n - 2f$ **propose** messages contain the same value x **then**
8: $x_i = x$, $\text{decided} = \text{true}$
9: **else if** at least $n - 4f$ **propose** messages contain the same value x **then**
10: $x_i = x$
11: **else**
12: choose x_i randomly, with $Pr[x_i = 0] = Pr[x_i = 1] = 1/2$
13: **end if**
14: $r = r + 1$
15: Broadcast **propose**(x_i, r)
16: **until** decided (see Line 8)
17: $\text{decision} = x_i$
