Ch 14 – Consensus

- The problem and its relevance
- The binary Consensus
- Generalization to multivalued Consensus
- Basic lower bounds

Theorem 14.29. *If $3 \leq n \leq 3f$, Consensus with Byzantine faults cannot be solved.*

Theorem 14.35. *Consensus with f faults cannot be solved in fewer than $f + 1$ rounds, even if faults are restricted to crashing nodes.*

Def - Consensus

Definition 14.1 (Consensus). *Each node $v \in V_g$ is given an input $x_v \in X$. To solve* Consensus, *an algorithm must compute output values $o_v \in X$ at all correct nodes $v \in V_g$ meeting the following conditions:*

- ***Agreement:*** *There is $o \in X$ so that $o_v = o$ for all $v \in V_g$. We refer to o as the output of the Consensus algorithm.*
- ***Validity:*** *If there is $x \in X$ so that for all $v \in V_g$ it holds that $x_v = x$, then $o = x$.*
- ***Termination:*** *There is $r \in \mathbb{N}$ satisfying that each $v \in V_g$ terminates and outputs o_v by the end of round r.*

The algorithm has round complexity *$R \in \mathbb{N}$, if it terminates in R rounds in all executions.*

If $X = \{0, 1\}$, we refer to the task as Binary Consensus, *and denote the input of node v by b_v (to indicate that it is a bit).*

n < 3f+1 lower bound

Theorem 14.29. *If $3 \leq n \leq 3f$, Consensus with Byzantine faults cannot be solved.*

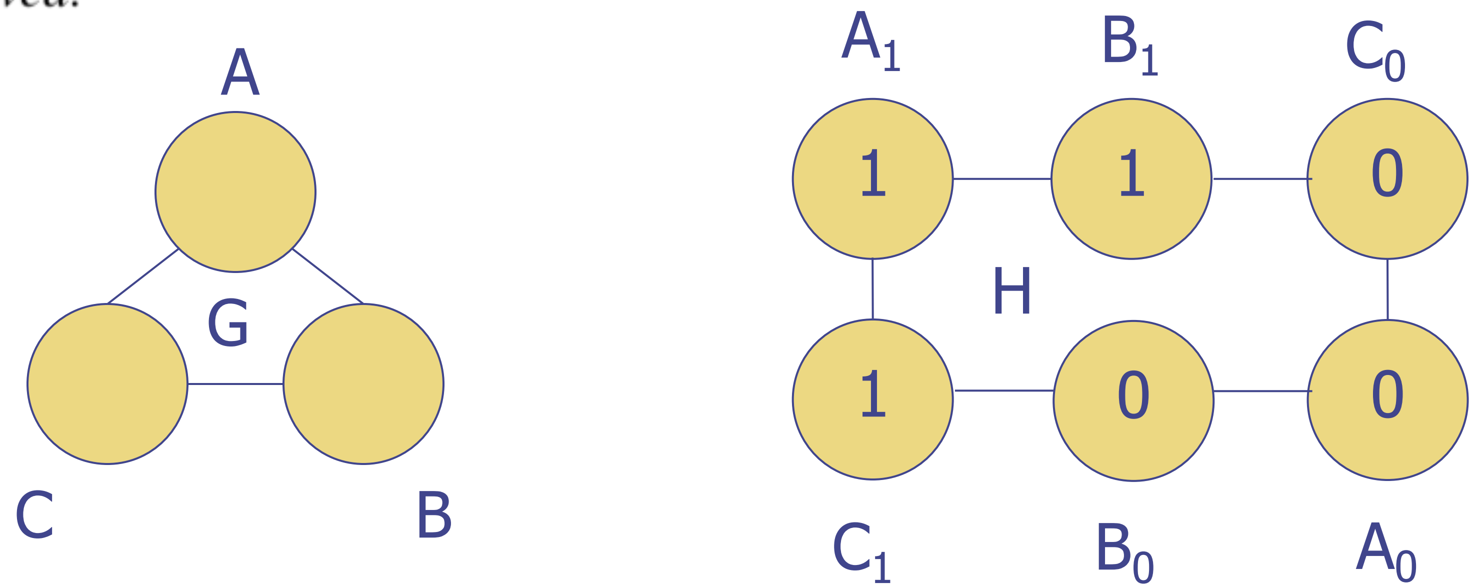

- Divide G to 3 sets of size up to f each.
- Assume there is an algorithm $\mathcal{A}$ that ensures consensus on G for any inputs.
- Construct H as described.
- We can execute $\mathcal{A}$ on H – since each node in H sees the same structure as in G. It's state machine determined by $\mathcal{A}$ functions the same in both graphs.
- E will be the execution on H with the specific inputs (it may detect an error)

n < 3f+1 lower bound

Corollary 14.27. *In $\mathcal{E}$, each node terminates and outputs the same value as it would output with $\mathcal{A}$ on G.*

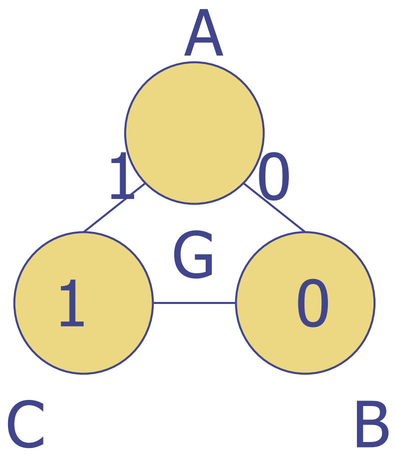

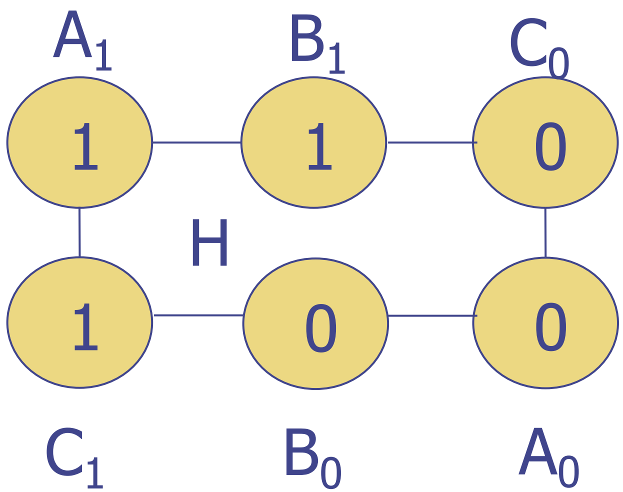

- Since nodes in any pair can't tell the difference their state machines produce the same output in H as it would in G
- Thus, specifically both C_1 and B_0 output the same value, say b

n < 3f+1 lower bound

Theorem 14.29. *If $3 \le n \le 3f$, Consensus with Byzantine faults cannot be solved.*

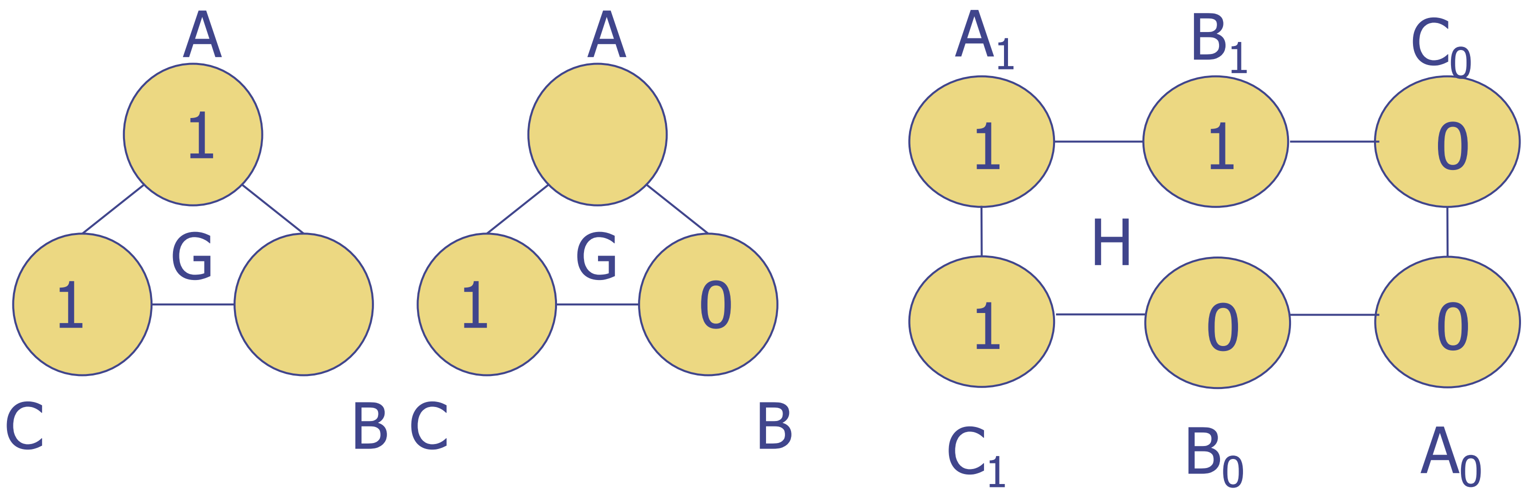

- As we claimed before, both C_1 and B_0 output the same value, say b
- Assume b=0.
- But, the left graph shows that both A_1 and C_1 should output "1"
- A contradiction

R < f+1 lower bound

Theorem 14.35. *Consensus with f faults cannot be solved in fewer than $f + 1$ rounds, even if faults are restricted to crashing nodes.*

Definition 14.30 (Crash Faults)**.** *If node $v \in V$* crashes *in round $r \in \mathbb{N}_{>0}$, it operates like a non-faulty node in rounds $1, \ldots, r - 1$, does nothing at all in rounds $r + 1, r + 2, \ldots$, and in round r sends an arbitrary subset of the messages it would send according to the algorithm.*

R < f+1 lower bound

- BREAKOUT ROOM

- Think of a binary consensus algorithm for the crash fault model
- fully synchronous system
- start with n=5 and various number of faults
- what can we conclude?

R < f+1 lower bound

Definition14.31(Pivotal Nodes)**.** *Observe that an execution in the synchronous model with crash faults is fully determined by*
1) specifying the node inputs and,
2) for each node, whether it crashes
3) and, if so, in which round and which of its messages of this round get sent.

Given an execution E *of a Consensus algorithm with at most* $n-2$ *crash faults and a node* $v \in V$ *that does not crash in* E, *we call* v **pivotal in round r** (of E) *if changing* E *by crashing* v *in round* r *of* E *without* v *sending any messages results in an execution with a* ***different output*** *(the execution does have an output, because at least one node does not crash).*

R < f+1 lower bound

Lemma 14.32. *There is a fault-free execution with a node that is pivotal in round* 1.

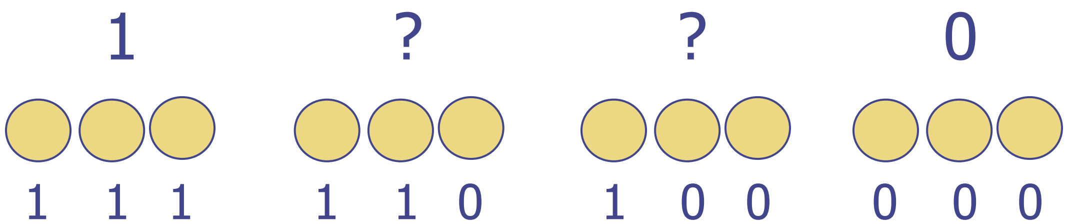

- Two possible decisions are forced by consistency
- The other two depends on the protocol.
- somewhere we move from 1 to 0.
- that determines the pivotal node.

R < f+1 lower bound

Lemma 14.32. *There is a fault-free execution with a node that is pivotal in round* 1.

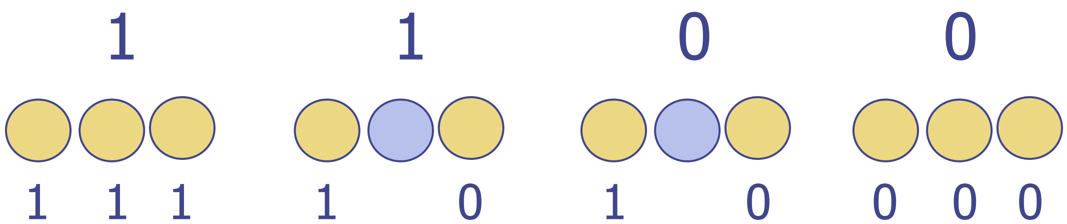

- Two possible decisions are forced by consistency
- The output of each of the other two depends on the protocol
- somewhere we move from outputting 1 to 0

- that determines the pivotal node.

R < f+1 lower bound

Lemma 14.33. *Suppose* $0 \le k \le n - 3$ *and* E *is an execution with* k *failing nodes, one in each round* 1,..., k*, that has a pivotal node in round* k +1*. Then there is an execution* E′ *which differs from* E *only in that this pivotal node crashes in round* k + 1 *and satisfies that there is a pivotal node in round* k + 2*.*

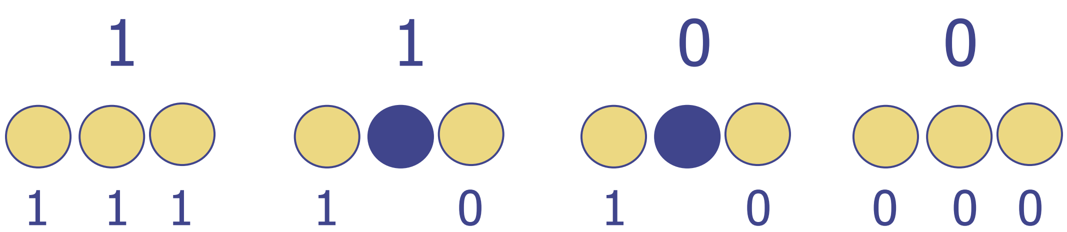

- Two possible decisions are forced by consistency
- The other two depends on the protocol.
- somewhere we move from 1 to 0.
- that determines the pivotal node.

R < f+1 lower bound

Lemma 14.33. *Suppose* $0 \leq k \leq n - 3$ *and* E *is an execution with* k *failing nodes, one in each round* 1,..., k, *that has a pivotal node in round* k +1. *Then there is an execution* E′ *which differs from* E *only in that this pivotal node crashes in round* k + 1 *and satisfies that there is a pivotal node in round* k + 2.

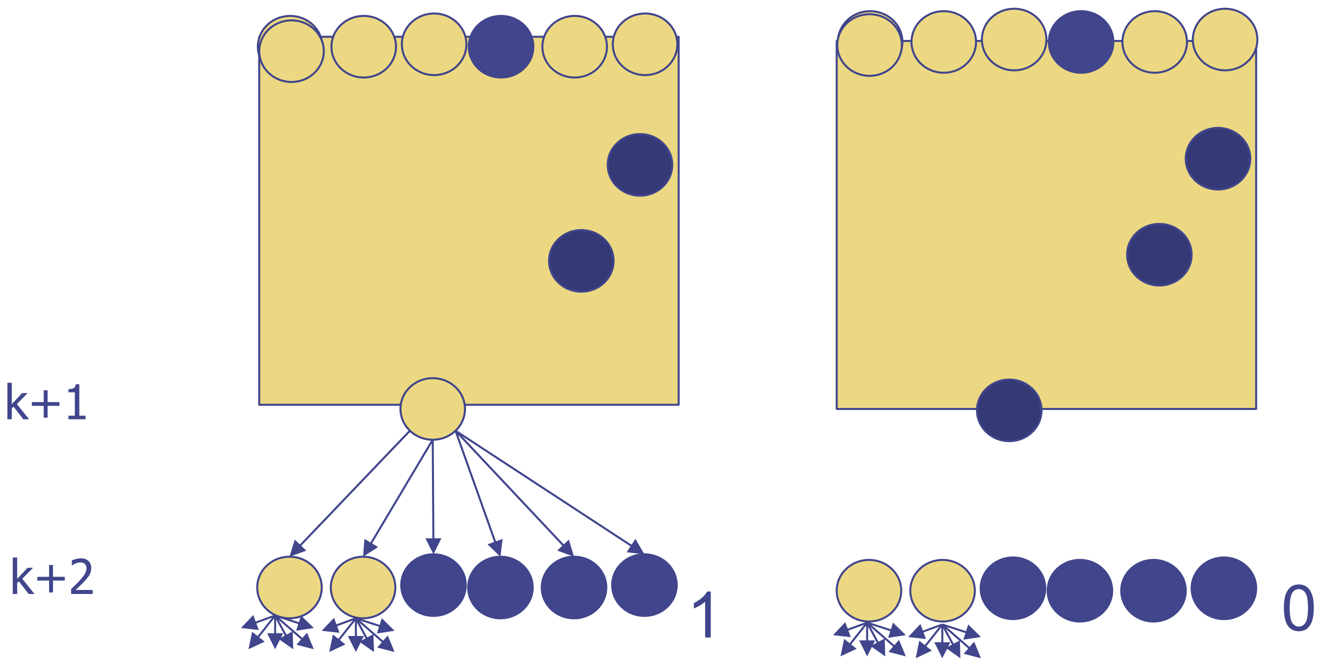

R < f+1 lower bound

Lemma 14.33. *Suppose* $0 \le k \le n-3$ *and* E *is an execution with* k *failing nodes, one in each round* 1,..., k*, that has a pivotal node in round* k +1*. Then there is an execution* E′ *which differs from* E *only in that this pivotal node crashes in round* k + 1 *and satisfies that there is a pivotal node in round* k + 2*.*

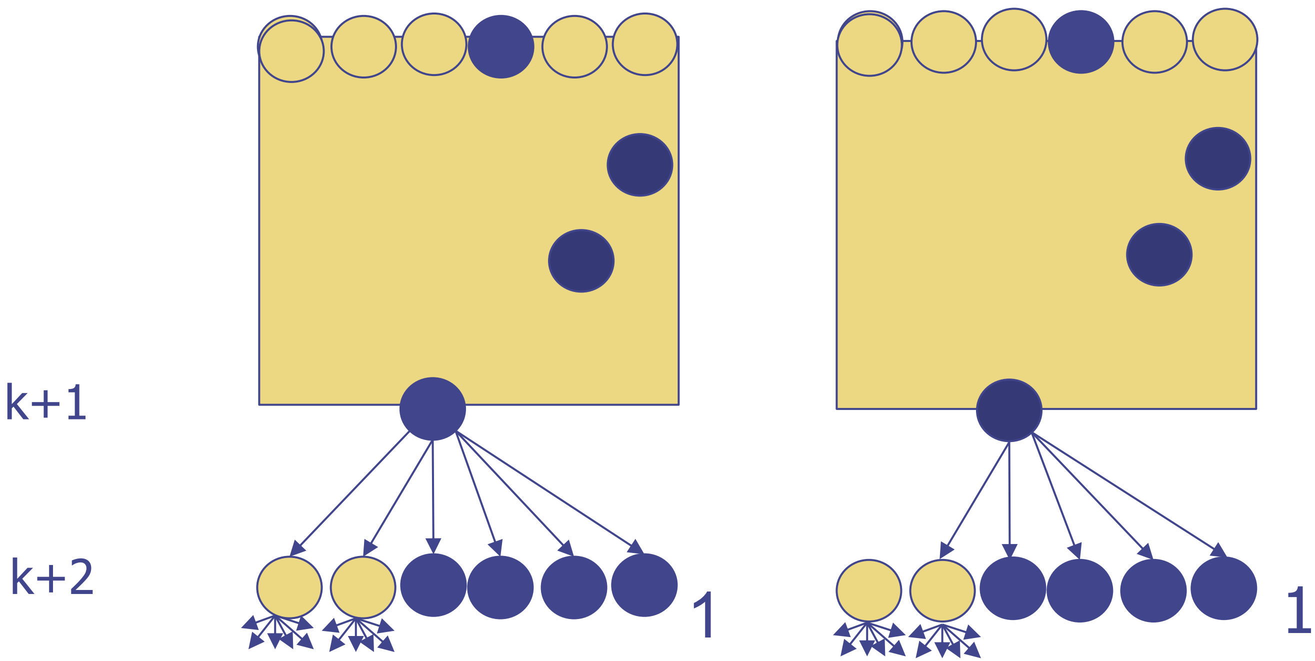

Do not send one at a time

R < f+1 lower bound

Lemma 14.33. *Suppose* 0 ≤ k ≤ n − 3 *and* E *is an execution with* k *failing nodes, one in each round* 1,..., k, *that has a pivotal node in round* k +1. *Then there is an execution* E′ *which differs from* E *only in that this pivotal node crashes in round* k + 1 *and satisfies that there is a pivotal node in round* k + 2.

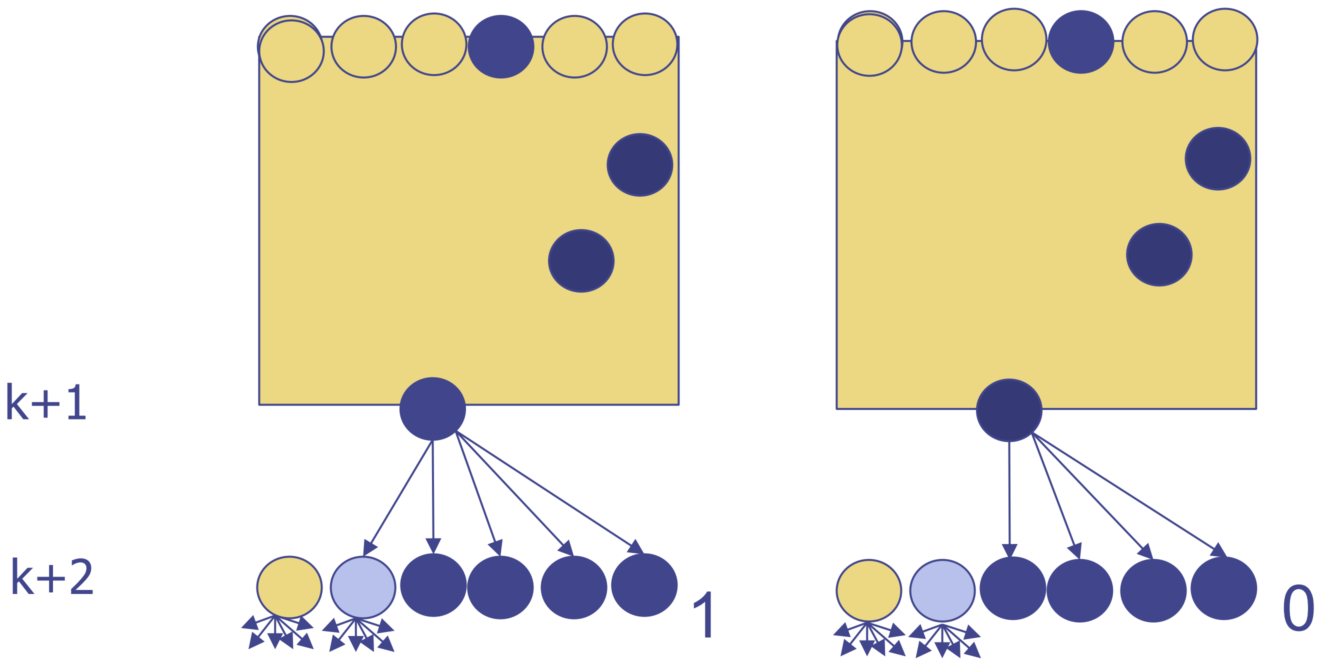

Do not send one at a time

it changes to 0 when a correct node sees the difference

R < f+1 lower bound

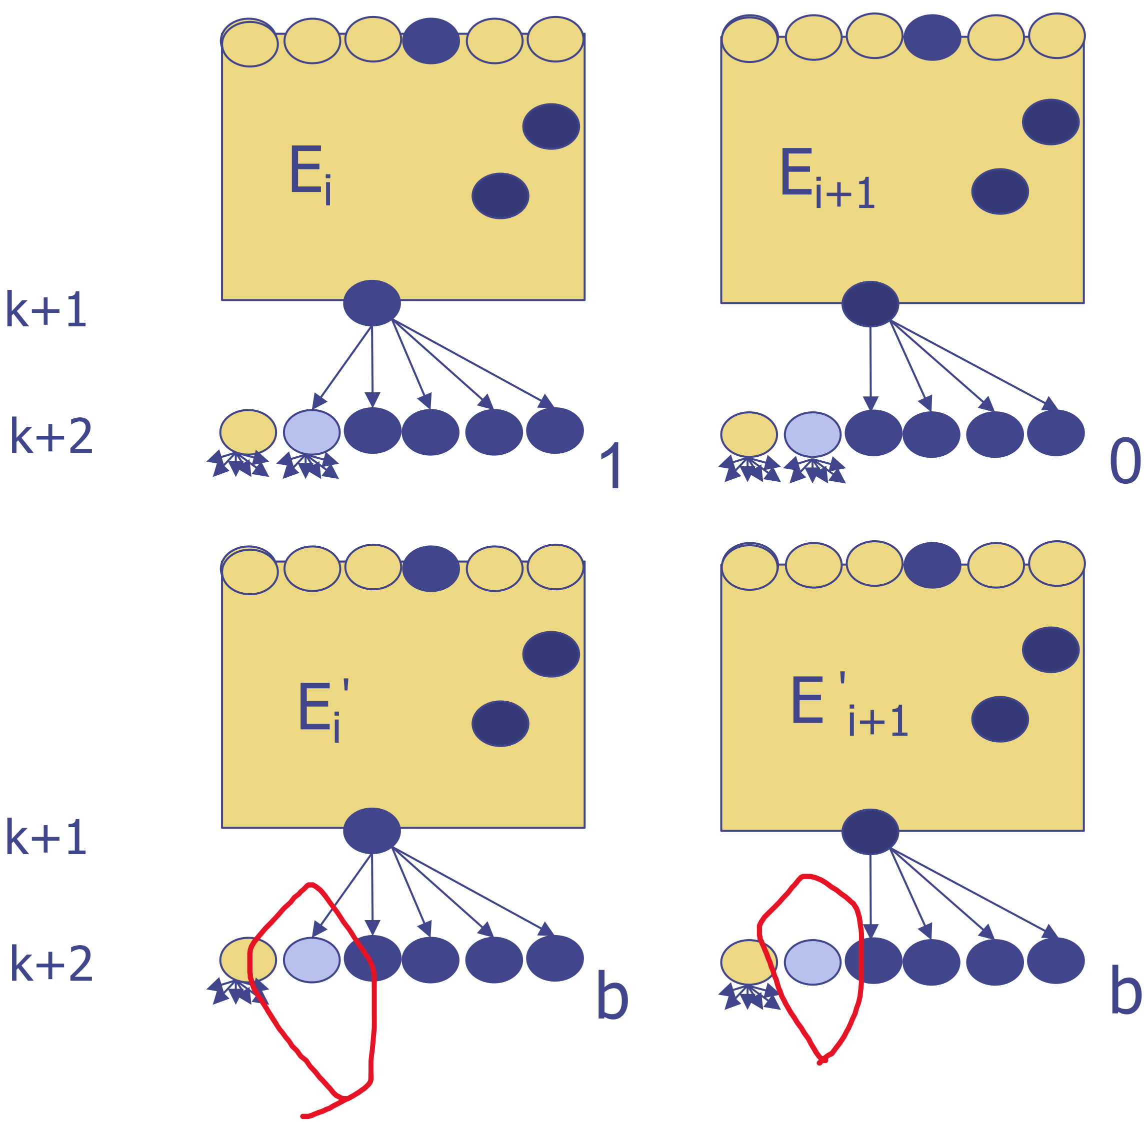

Output of E_i' and E'_{i+1} are the same

Thus the node is either pivotal of E_i or E_{i+1}

R < f+1 lower bound

Lemma 14.33. *Suppose* $0 \leq k \leq n - 3$ *and* E *is an execution with* k *failing nodes, one in each round* 1,..., k, *that has a pivotal node in round* k +1. *Then there is an execution* E′ *which differs from* E *only in that this pivotal node crashes in round* k + 1 *and satisfies that there is a pivotal node in round* k + 2.

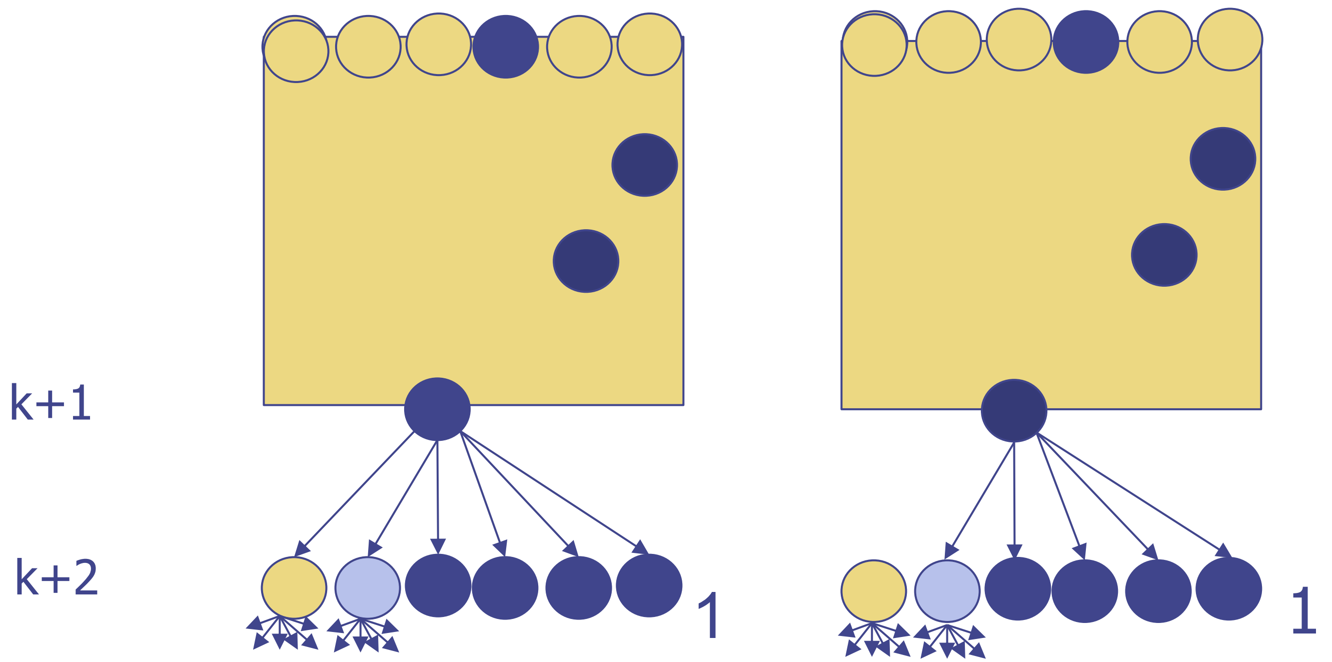

R < f+1 lower bound

Corollary 14.34. *Any Consensus algorithm has an execution with a pivotal node in round* $\min\{f, n-2\}$.

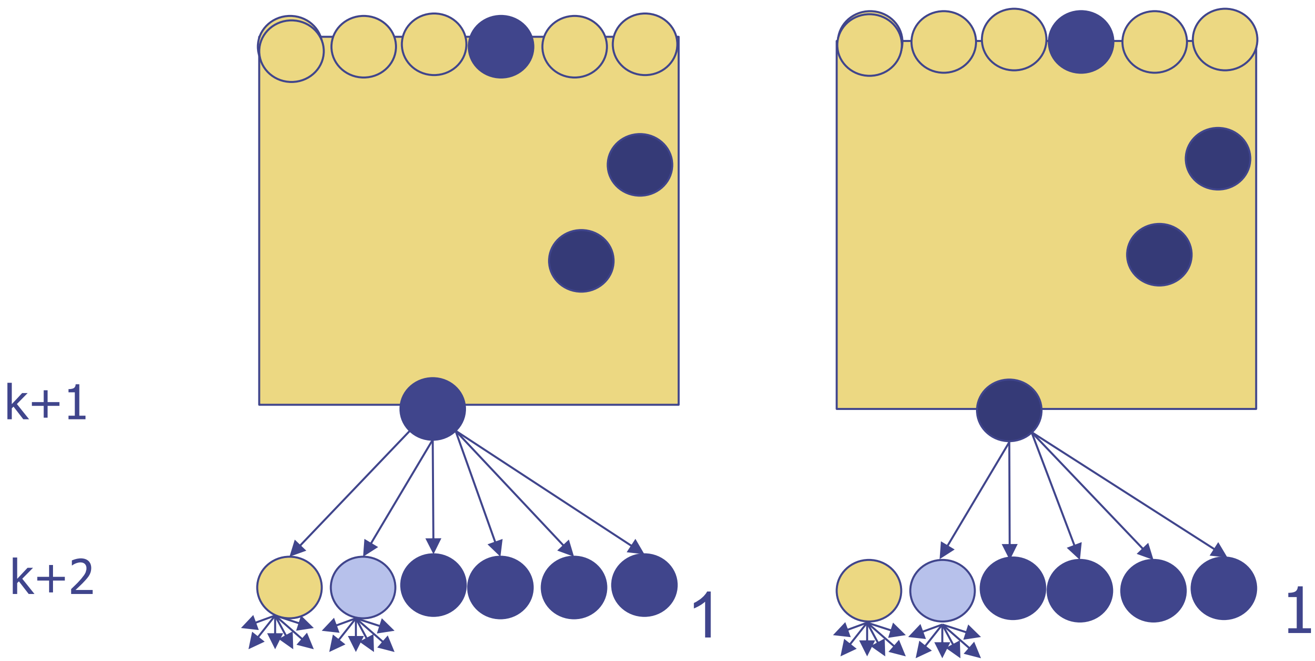

R < f+1 lower bound

Theorem 14.35. *Consensus with f faults cannot be solved in fewer than $f + 1$ rounds, even if faults are restricted to crashing nodes.*