

A Lower Bound


Theorem

There is no algorithm that solves the consensus problem in fewer than $f+1$ rounds in the presence of f crash failures, if $n \geq f+2$

We consider a special case ($f=1$) to study the proof technique


Views

Let α be an execution. The **view** of process p_i in α , denoted by $\alpha|p_i$, is the subsequence of computation and message receive events that occur in p_i together with the state of p_i in the initial configuration of α



Views

Let α be an execution. The **view** of process p_i in α , denoted by $\alpha|p_i$, is the subsequence of computation and message receive events that occur in p_i together with the state of p_i in the initial configuration of α



Similarity

Definition Let α_1 and α_2 be two executions of consensus and let p_i be a correct process in both α_1 and α_2 .

α_1 is **similar** to α_2 with respect to p_i , denoted $\alpha_1 \sim_{p_i} \alpha_2$ if

$$\alpha_1|p_i = \alpha_2|p_i$$

Similarity

Definition Let α_1 and α_2 be two executions of consensus and let p_i be a correct process in both α_1 and α_2 .

α_1 is **similar** to α_2 with respect to p_i , denoted $\alpha_1 \sim_{p_i} \alpha_2$ if

$$\alpha_1|_{p_i} = \alpha_2|_{p_i}$$

Note If $\alpha_1 \sim_{p_i} \alpha_2$ then p_i decides the same value in both executions

Similarity

Definition Let α_1 and α_2 be two executions of consensus and let p_i be a correct process in both α_1 and α_2 .

α_1 is **similar** to α_2 with respect to p_i , denoted $\alpha_1 \sim_{p_i} \alpha_2$ if

$$\alpha_1|_{p_i} = \alpha_2|_{p_i}$$

Note If $\alpha_1 \sim_{p_i} \alpha_2$ then p_i decides the same value in both executions

Lemma If $\alpha_1 \sim_{p_i} \alpha_2$ and p_i is correct, then $\text{dec}(\alpha_1) = \text{dec}(\alpha_2)$

Similarity

Definition Let α_1 and α_2 be two executions of consensus and let p_i be a correct process in both α_1 and α_2 .

α_1 is **similar** to α_2 with respect to p_i , denoted $\alpha_1 \sim_{p_i} \alpha_2$ if

$$\alpha_1|_{p_i} = \alpha_2|_{p_i}$$

Note If $\alpha_1 \sim_{p_i} \alpha_2$ then p_i decides the same value in both executions

Lemma If $\alpha_1 \sim_{p_i} \alpha_2$ and p_i is correct, then $\text{dec}(\alpha_1) = \text{dec}(\alpha_2)$

The transitive closure of $\alpha_1 \sim_{p_i} \alpha_2$ is denoted $\alpha_1 \approx \alpha_2$.

We say that $\alpha_1 \approx \alpha_2$ if there exist executions $\beta_1, \beta_2, \dots, \beta_{k+1}$ such that

$$\alpha_1 = \beta_1 \sim_{p_{i_1}} \beta_2 \sim_{p_{i_2}} \dots \sim_{p_{i_k}} \beta_{k+1} = \alpha_2$$

Similarity

Definition Let α_1 and α_2 be two executions of consensus and let p_i be a correct process in both α_1 and α_2 .

α_1 is **similar** to α_2 with respect to p_i , denoted $\alpha_1 \sim_{p_i} \alpha_2$ if

$$\alpha_1|_{p_i} = \alpha_2|_{p_i}$$

Note If $\alpha_1 \sim_{p_i} \alpha_2$ then p_i decides the same value in both executions

Lemma If $\alpha_1 \sim_{p_i} \alpha_2$ and p_i is correct, then $\text{dec}(\alpha_1) = \text{dec}(\alpha_2)$

The transitive closure of $\alpha_1 \sim_{p_i} \alpha_2$ is denoted $\alpha_1 \approx \alpha_2$.

We say that $\alpha_1 \approx \alpha_2$ if there exist executions $\beta_1, \beta_2, \dots, \beta_{k+1}$ such that

$$\alpha_1 = \beta_1 \sim_{p_{i_1}} \beta_2 \sim_{p_{i_2}} \dots \sim_{p_{i_k}} \beta_{k+1} = \alpha_2$$

Lemma If $\alpha_1 \approx \alpha_2$ then $\text{dec}(\alpha_1) = \text{dec}(\alpha_2)$

Single-Failure Case

There is no algorithm that solves consensus in fewer than two rounds in the presence of one crash failure, if $n \geq 3$

The Idea

By contradiction

- Consider a one-round execution in which each process proposes 0. What is the decision value?
- Consider another one-round execution in which each process proposes 1. What is the decision value?
- Show that there is a chain of similar executions that relate the two executions.


So what?

α^i 's

Definition

α^i is the execution of the algorithm in which

- no failures occur
- only processes p_0, \dots, p_{i-1} propose 1




Adjacent α^i 's are similar!


Starting from α^i , we build a set of executions α_j^i where $0 \leq j \leq n-1$ as follows:

α_j^i is obtained from α^i after removing the messages that p_i sends to the j -th highest numbered processors (excluding itself)


The executions




Indistinguishability




Indistinguishability




Indistinguishability



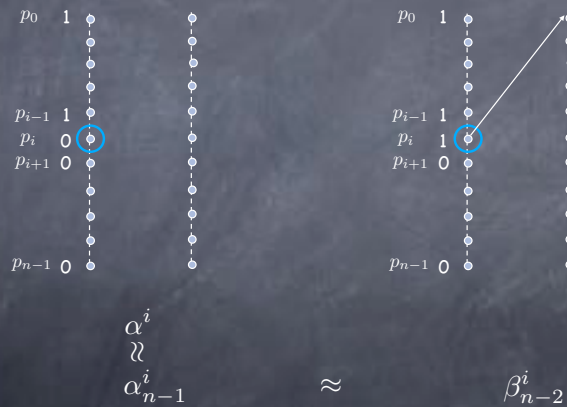
Indistinguishability



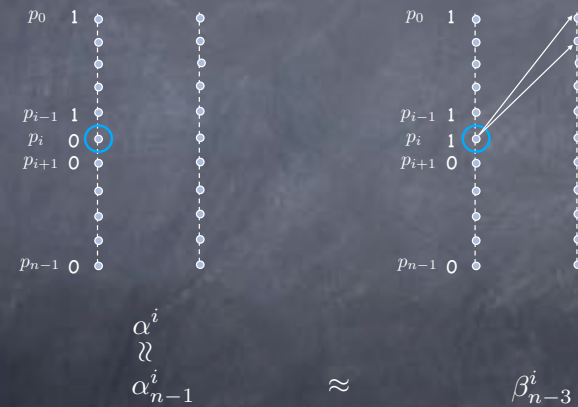
Indistinguishability



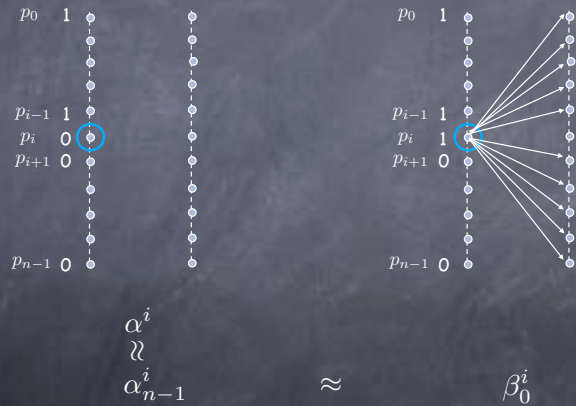
Indistinguishability




Indistinguishability



Indistinguishability



Indistinguishability



Indistinguishability

