

CONSENSUS

IMPOSSIBLE
OR
NOT?

Admin: Midterm Next Week

→ DURING CLASS.

→ OPEN BOOK: Feel free to bring papers, notes, etc.

→ Focus on testing understanding of the papers & ^{anything additional} concepts covered in class

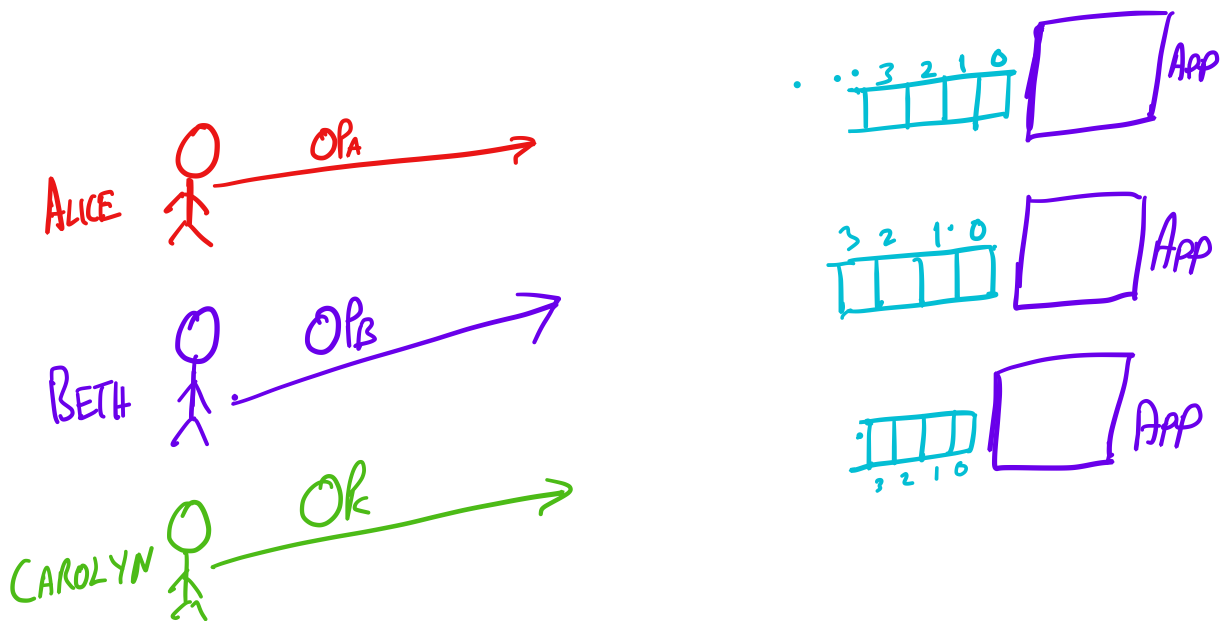
(Please make sure you know what the asynchronous model entails!)

→ Questions in the pre-class notes are good examples of what might be asked

RECAP

WANT TO BUILD A FAULT TOLERANT APPLICATION

↳ IS A STATE MACHINE



REQUIREMENT

→ All LOGS AGREE ON ORDER OF OPERATIONS



All LOGS AGREE ON OPERATION AT
Idx 0, 1, 2, ...

↳ CONSENSUS ALGORITHMS

CORE-RESULT

No DETERMINISTIC FAULT-TOLERANT CONSENSUS PROTOCOL
① ② ③

FOR THE ASYNCHRONOUS SETTING

④

① Deterministic What?



Where used?

② Fault Tolerant

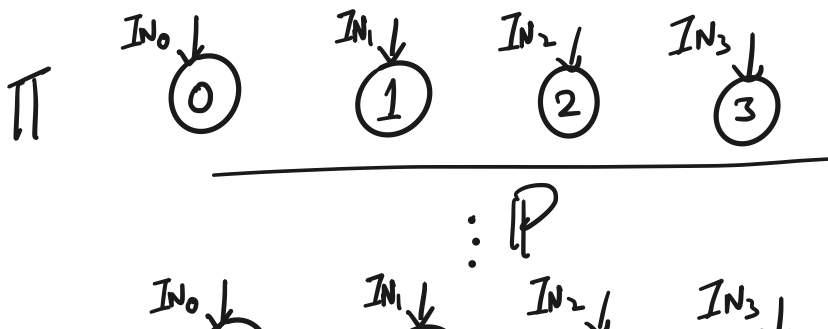
- Want to consider as many protocols as possible

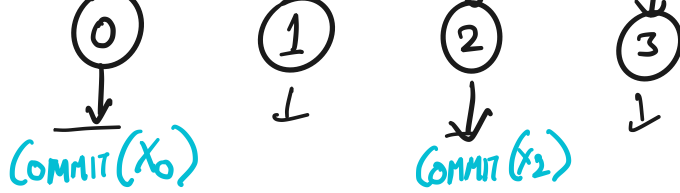
↳ Choose a weak failure model

At most 1 process fails

Fail-stop failures

③ Consensus Protocol





REQUIREMENTS

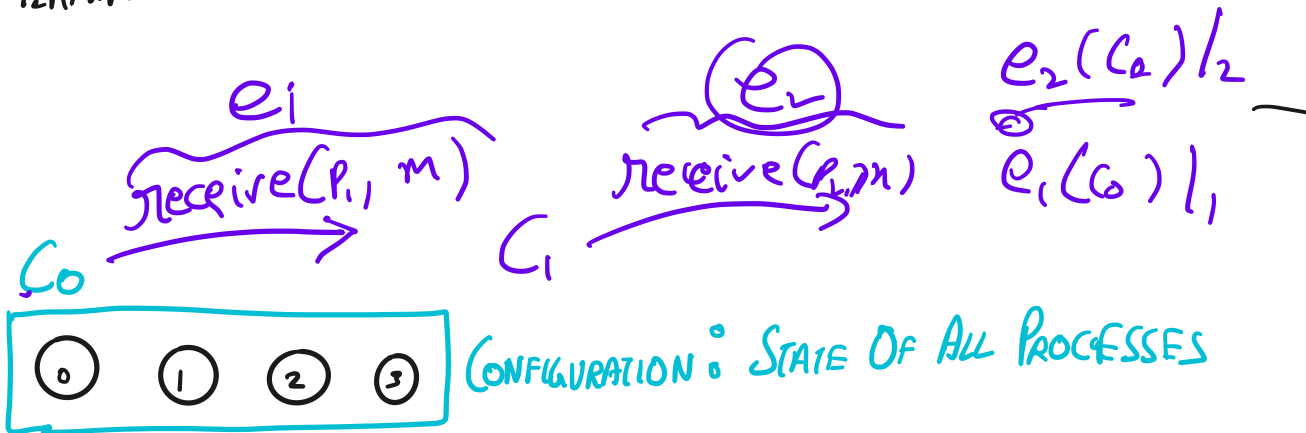
AGREEMENT: IF PROCESS P_i DECIDES X_i & P_j DECIDES X_j
outputs $\text{commit}(x_i)$

Then $X_i = X_j$

VALIDITY: IF PROCESS P_i DECIDES X_i THEN $\exists j$ s.t. $\underline{I_{N_j}} = X_i$
(EQUIVALENT TO STRONG UNANIMITY)

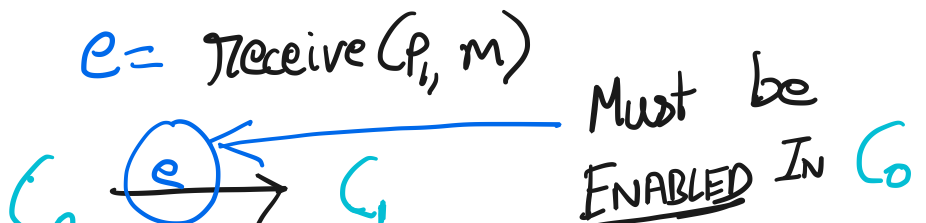
TERMINATION: ALL PROCESSES EVENTUALLY DECIDE

TERMINOLOGY



EVENTS/STEPS: $\text{receive}(p, m)$

Receive message m at process p
 (changes the state of process p .
 m can be \emptyset the empty message



SCHEDULE (σ): Sequence of steps.

e_0 DISJOINT FROM e_1 : e_0 is a receive from a different process than

σ_0 DISJOINT FROM σ_1

NO DETERMINISTIC FAULT-TOLERANT CONSENSUS PROTOCOL FOR THE ASYNCHRONOUS SETTING

OVERALL IDEA:

① FOCUS ON BINARY CONSENSUS [INPUTS & OUTPUTS ARE 0 OR 1]

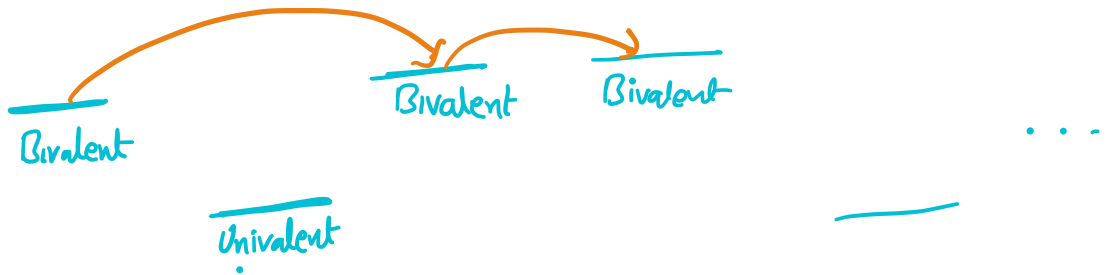
↳ ASSUME A DET, FT CONSENSUS PROTOCOL EXISTS P

② SHOW THAT \exists CONFIGURATIONS FOR WHICH P
CAN DECIDE BOTH 0 OR 1 [BIVALENT CONFIGS]

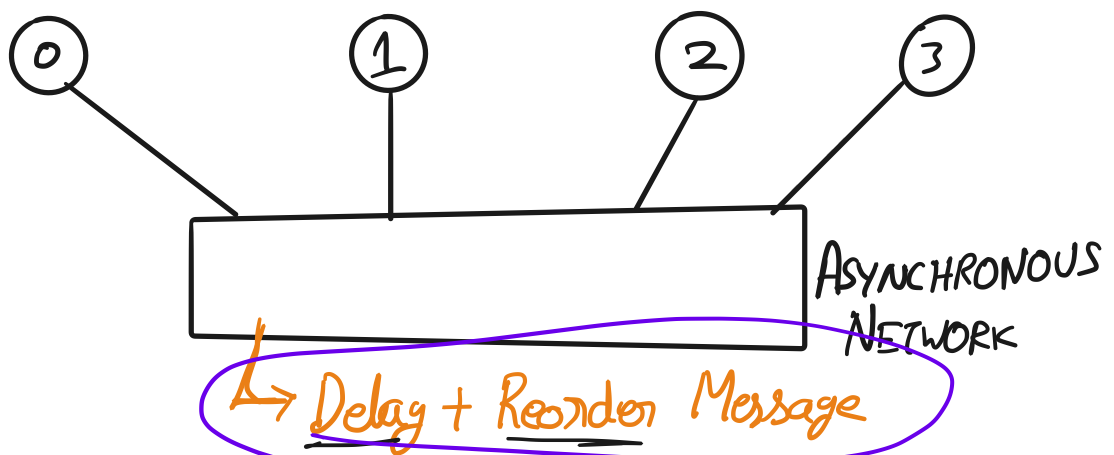
$\exists \sigma_0 \dots$
 $\sigma_1 \dots$

③ SHOW THAT \exists INITIAL CONFIGURATIONS THAT
(INPUTS)
ARE BIVALENT

④ SHOW THAT ONE CAN ALWAYS FIND A SCHEDULE THAT GOES FROM BIVALENT C TO ANOTHER BIVALENT C



"CONSTRUCT A SCHEDULE"



WHY DOES THIS WORK?

P Is 1-fault tolerant

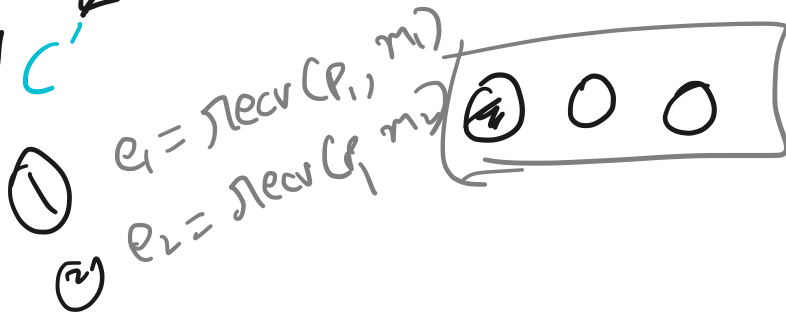
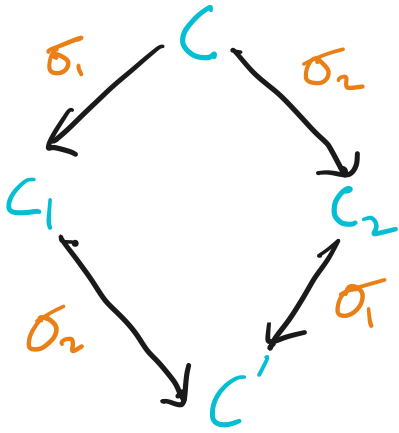
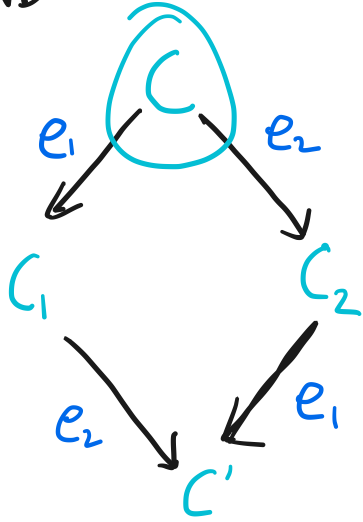
But works in async networks

↳ CANNOT DISTINGUISH B/W DELAYED MESSAGES
↳ FAILED PROCESS

Must eventually suspect quiet processes have failed

Also must deal with quiet processes sending messages

DIAMOND LEMMA

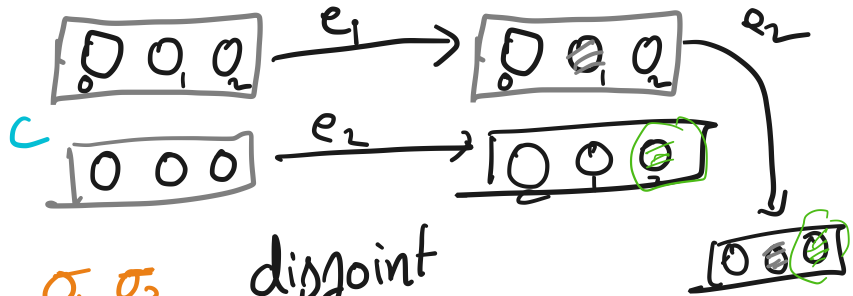


$$P_1 = 1$$

$$P_2 = 2$$

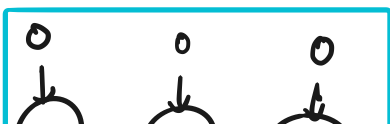
e_1, e_2 disjoint $\Rightarrow e_1 := \text{recv}(P_1, m)$
 $e_2 := \text{recv}(P_2, m')$

e_1, e_2 enabled in C $P_1 \neq P_2$



σ_1, σ_2 disjoint
enabled in C
 ??

VALENCE

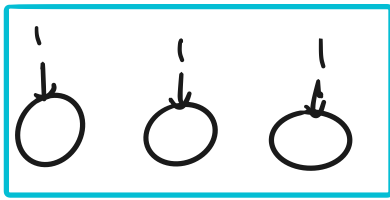


\Rightarrow Choose 0 / Decide 0



INITIAL CONFIG

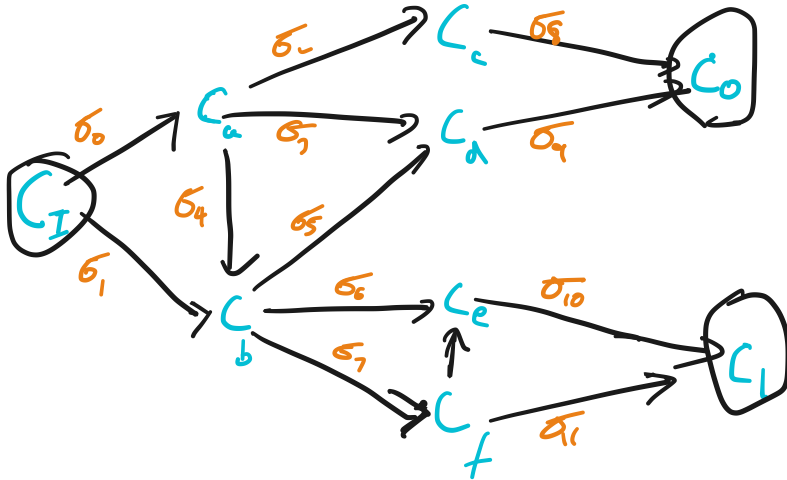
P



INITIAL CONFIG



Choose 1 / Decide 1

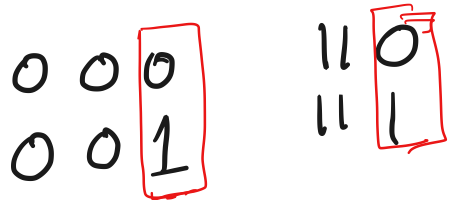


BIVALENT INITIAL CONFIGURATION

INPUT	OUTPUT
0 0 0	0
0 0 1	0
0 1 0	0
1 0 0	1
1 1 0	1
1 0 1	1
0 1 1	1
1 1 1	1

Observation

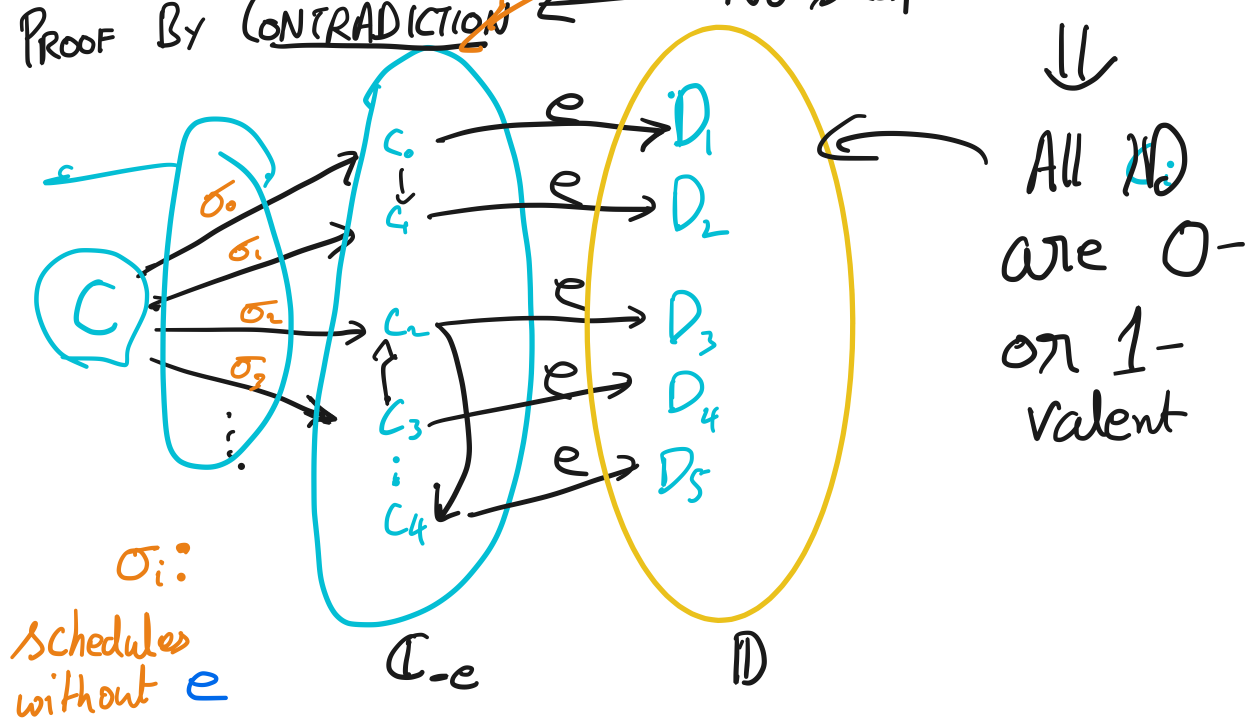
P is 1-Fault tolerant



s.t. $e \in \sigma(C)$ bivalent

PROOF BY CONTRADICTION \leftarrow No such σ exists \Downarrow

$e \notin \sigma_0 \rightarrow \sigma_0$
 σ_i



What does D contain?

Must contain at least one 0-valent & one 1-valent configuration.

Why

C is bivalent

$\Rightarrow \exists$ enabled σ_0 s.t.

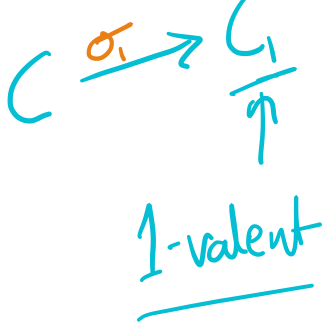
$C \xrightarrow{\sigma_0} C_0 \leftarrow 0\text{-valent}$

If $e \notin \sigma_0$ then $C_0 \in D$
0-valent

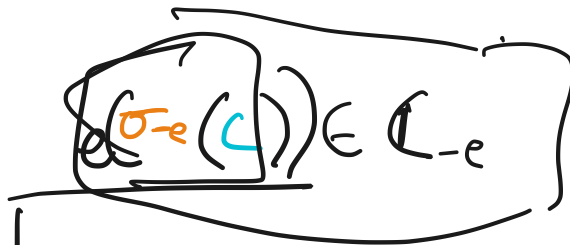
If $e \in \sigma_0$ then

$\sigma_0 = \underbrace{e_0; e_1; e_2; \dots}_{\sigma_{-e}} \underbrace{e; \dots}_{\sigma_{Te}}$

Imagine same thing with



But



↳ can be

0-valent
 or bivalent [Disallowed by

assumption that \mathcal{C}

has no bivalent configurations]

\mathcal{C}_e contains both 0- & 1-valent configurations

\mathcal{C}_e contains configurations C_A, C_B that are

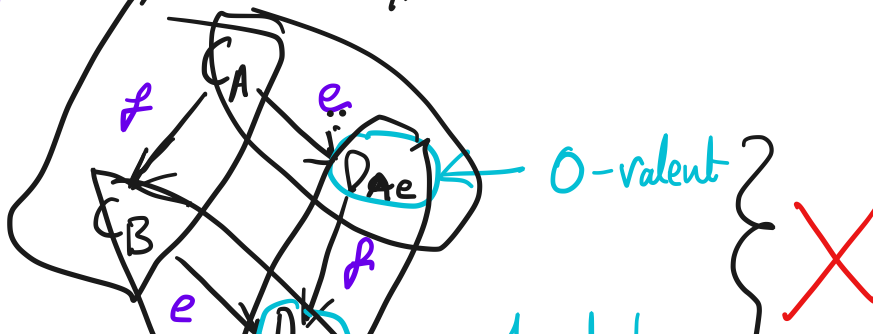
(a) Neighbors $C_A \xrightarrow{f} C_B$

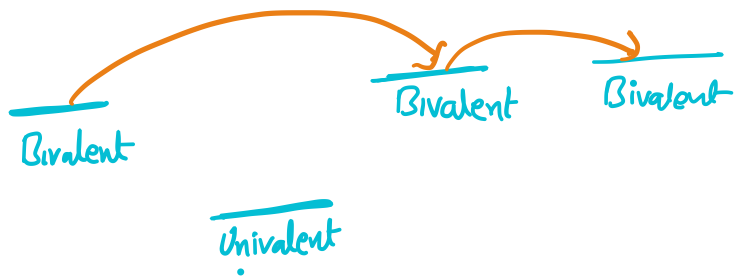
(b) $e(C_A)$ is 0-valent, $e(C_B)$ is 1-valent.

How? Hand-wavy but use initial bivalent argument.

Finally: arrive at contradiction: no such f can exist

(1) If f and e are different processes





o o o

No DETERMINISTIC ^① FAULT-TOLERANT CONSENSUS PROTOCOL
FOR THE ASYNCHRONOUS ^② SETTING

What happens if we relax assumptions?

① Deterministic?

② PARTIAL SYNCHRONY

