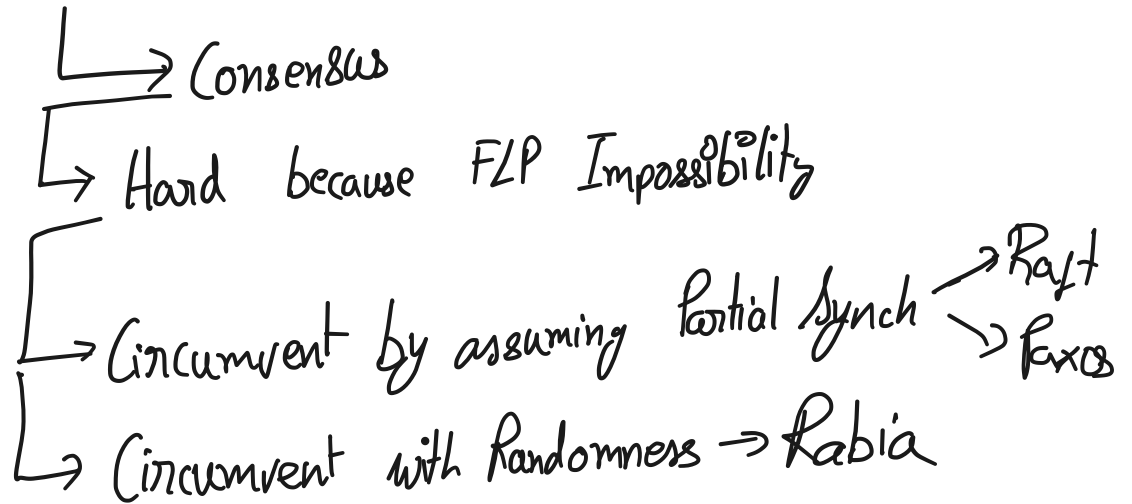


# FLP + Partial Synchrony

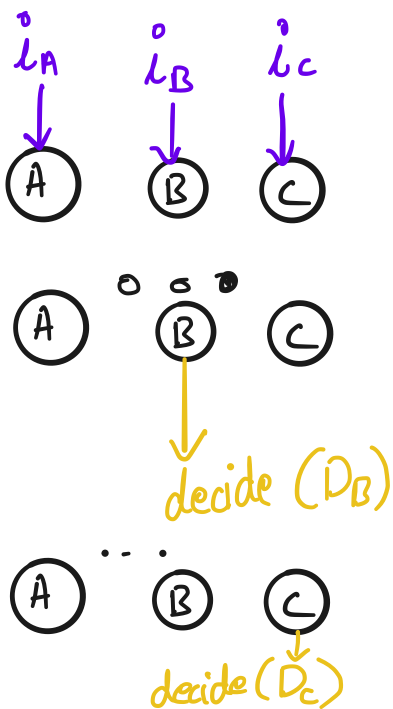
Where we are

RSMs



Today's FLP, Partial Synchrony.

- Defining Consensus



Agreement

→ Process  $p_0$  decides  $d_0$  &  
process  $p_i$  decides  $d_i$   
 $\Rightarrow d_0 = d_i$

Validity

- Process  $p$  decides  $d$   
 $\Rightarrow$  some process  $q$   
had input =  $d$

T. A. Lam

termination

- Eventually some process decides.

## Impossibility Result:

No fault-tolerant, deterministic algorithm in the asynchronous model.

Aside: From an award speech in 2001

- Lynch wanted to prove an asynch consensus protocol (designed by Lamport) correct in '82.
- Kept running into trouble, started working with Fischer & Paterson. No luck
- led to argument in this paper.

Going to walk through proof a bit


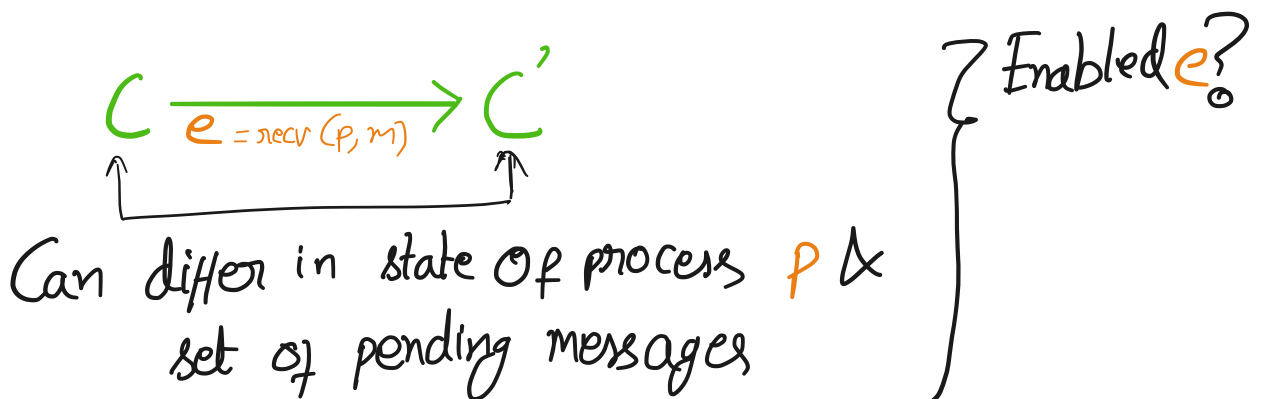



- Why?
- Ideas are interesting & at the core of how we think of many distributed algorithms
  - Origin of the model we have been using.


How?

Assume a simple setting

- Binary Consensus
- No message loss (only delays)
- At most 1 failure

## System Model





"Partially" Correctness (or safe correctness)

- Agreement
- Validity

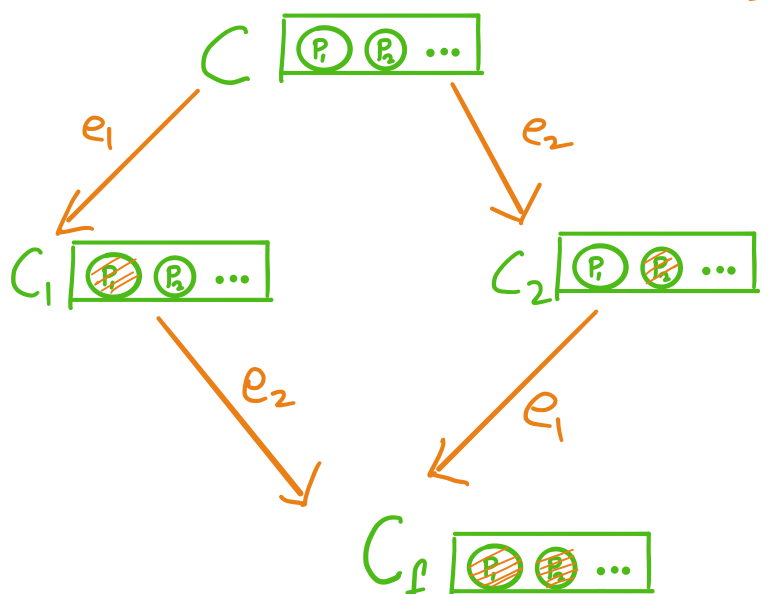
"Totally" Correct = partial correct + Liveness  
Termination

totally correct

No  $n$ -fault-tolerant, deterministic algorithm  
in the asynchronous model.

Some tools

- ①  $C \mid e_1, e_2$  |  $e_1, e_2$  enabled in  $C$   
 $e_1 = \text{next}(P_1, m)$     $e_2 = \text{next}(P_2, m')$   
 $P_1 \neq P_2$



Diamond Lemma

② Valence: Assume  $\exists$  partially correct consensus protocol  $P$  implemented by system

$C$   $\emptyset$ -valent iff  $\forall$  valid schedules  $\sigma$  from  $C$

$C \xrightarrow{\sigma} \begin{cases} \circ \emptyset\text{-decided} \\ \text{(some process } p \text{ decided } \emptyset) \\ \text{OR} \\ \circ \text{Undecided} \end{cases}$

Similarly 1-valent.

$C$  bivalent  $\rightarrow$   $C$  is not  $\emptyset$ -valent &  $C$  is not 1-valent & some configuration reachable from  $C$  by an enabled schedule decides

F/P proof  $\rightarrow$  Proof by contradiction

Assume  $\exists$  totally correct protocol  $P$  that is 1-fault tolerant.

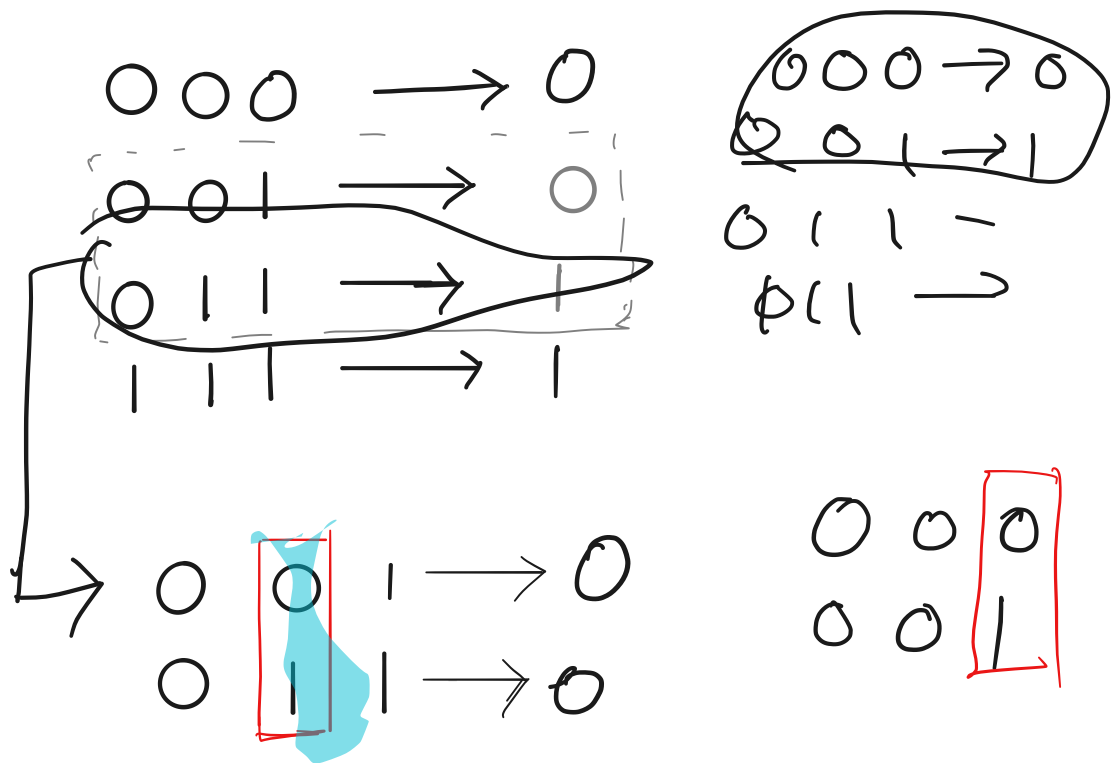
①  $\exists$  0- and 1-valent initial configurations

inputs                      Decision } Validity

0 0 0 → 0

1 1 1 → 1

②  $\exists$  Bivalent initial configuration





③ For any enabled event  $e$  in bivalent configuration  $C$  can find valid  $\sigma$  (where at most 1 process is silent) s.t.

$C \xrightarrow{\sigma} C' \xrightarrow{e} C''$

$C''$  is Bivalent.

④ Start from initial bivalent configuration

Why not pick consistent  $e$ ?




Schedule of unbounded length going from one BIVALENT CONFIG to another  
 $\Rightarrow$  Unbounded schedule where no process decides  
 $\Rightarrow$  No termination.

Look at (3)

Given bivalent configuration  $C$ , event  $e$  want to show  $\exists \sigma$  s.t.  $C \xrightarrow{\sigma} e \rightarrow$  Bivalent

- If  $C \xrightarrow{e}$  bivalent : Done  $\sigma = []$
- Otherwise,





Observation: Claim is equivalent to claim that  $D$  contains bivalent configuration.

Proof by contradiction:

(i)  $\exists$  reachable  $\phi$  & 1-valent configuration from  $C$

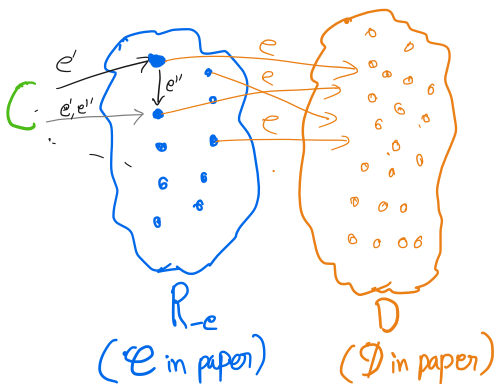
[Reachable:  $\exists$  valid  $\sigma$  from  $C$  to config]

-  $\mathcal{P}$  totally correct  $\Rightarrow$  terminates

-  $C$  bivalent  $\Rightarrow \exists \sigma_0, \sigma_1$


$C \xrightarrow{\sigma_0} \text{decide } \phi$

$\xrightarrow{\sigma_1} \text{decide } 1$





(ii)  $\exists \emptyset$  1-valent configurations in  $D$



Let  $C_0$  be  $\emptyset$  valent config reachable from  $C$

- If  $C_0 \in R_e$ , then

$e(C_0) \in D$  is a  $\emptyset$ -valent config.

- If  $C_0 \notin R_e$ , then

$e \in$  schedule  $\sigma_0$  from  $C \xrightarrow{\sigma_0} C_0$

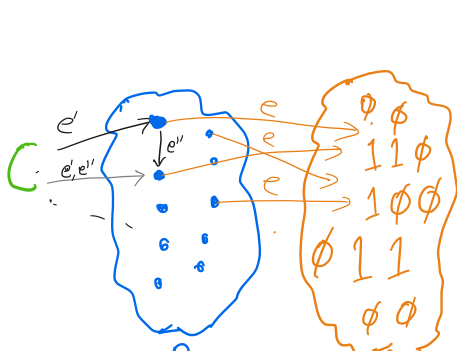
$\Rightarrow \exists C_0'$  s.t.

$C \xrightarrow{\sigma_0 \text{ (before } e)} C' \xrightarrow{e} C_0' \xrightarrow{\sigma_0 \text{ (after)}} C_0$

$\Delta C_0' \in D$

But by assumption no bivalent config in  $D \Rightarrow C_0'$  is  $\emptyset$ -valent

By symmetry for 1-valent



(iii) Hand-way:

$\exists C_0, C_1 \in R_e$  s.t.

(a)  $\exists e' C_0 \xrightarrow{e'} C_1$

$P_e$   
( $\mathcal{E}$  in paper)

$\mathcal{D}$   
( $\mathcal{D}$  in paper)

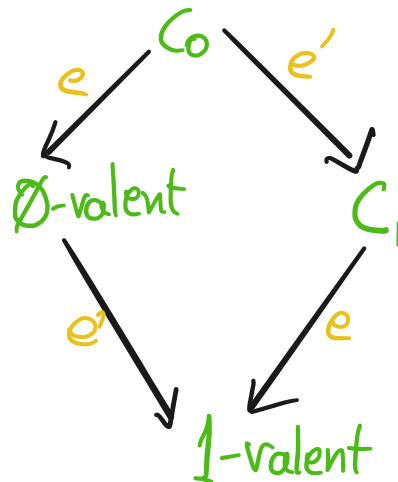
(b)  $e(C_0) - \emptyset$ -valent &  
 $e(C_1) - 1$ -valent

(iv)  $e = \text{recv}(p, m)$

$e' = \text{recv}(p', m')$

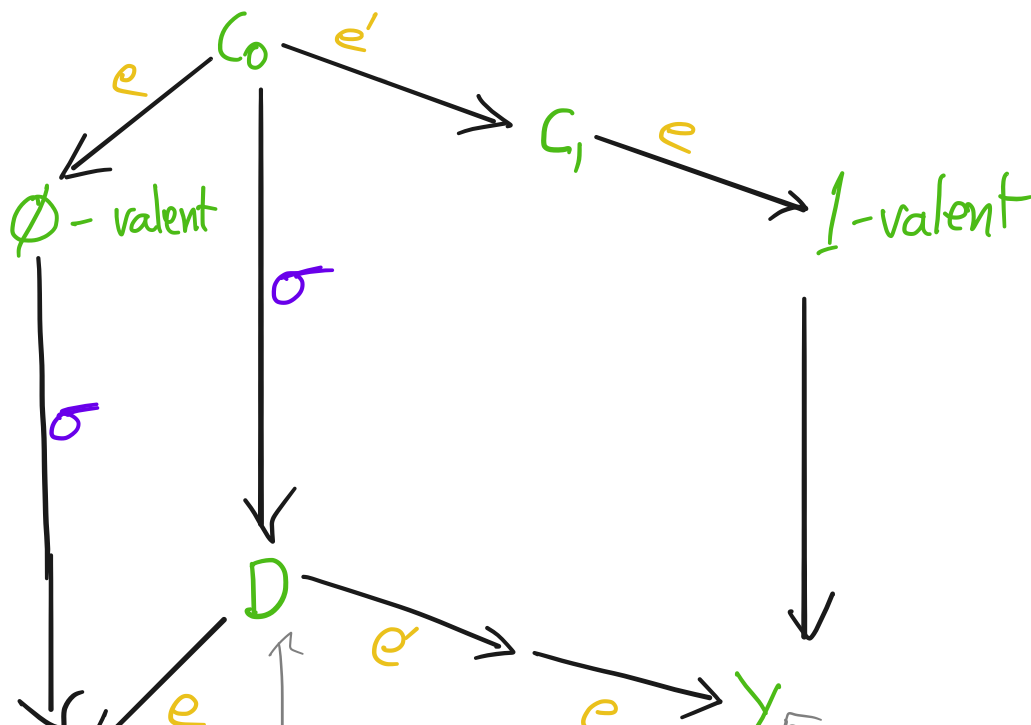
If  $p \neq p'$

X



If  $p = p'$

$\exists \sigma$  where  $C_0 \xrightarrow{\sigma}$  some process decides;  
 $\sigma$  finite &  $p$  takes no steps in  $\sigma$   
[Because F.T.]




$\times$   
 must be  $\emptyset$ ?

must be 1?

$\times$

Reached a decision.  
 Agreement  $\Rightarrow$  must hold

Thus  $D$  must contain a bivalent config.



Given bivalent configuration  $C$ , event  $e$

$\exists \sigma$  s.t.  $C \xrightarrow{\sigma} \xrightarrow{e}$  Bivalent

$\rightarrow$  For any consensus protocol that is 1-FT,  $\exists$  unbounded schedule w/o decision

$\Rightarrow$  No totally correct protocol.

So where does that leave us

$\frac{\text{FT-Consensus}}{\text{synchronous}}$  ( Real world?? )  $\frac{\text{No consensus}}{\text{asynch}}$

(Fail-stop,

(1-fail stop)

Omission,

Byz.,


Auth. Byz.)

Partial synchrony

- Separates out comm. & processing

Talks about two variants

(a) Known  $\triangle$



(b) Unknown  $\triangle$ , but it always holds.

Results are the same, but most people mean (a)

Approach

- Design protocols for a model where

- Communication & processing are synchronous

- Messages can be lost

... But  $\exists$  time (GST) after which message from correct process  $p$  to correct process  $q$  is received

- Show they are totally correct

→ Design

- Rounds where processes send & receive messages

- Decisions require quorum of CORRECT processes to be involved

Why termination

- Each round takes bounded time (synchrony assumption)

= After GST, all correct processes can communicate w/ each other

Assume large enough set of processes to ensure  $t$ -failures still allow decisions

→ Bounded rounds to reach consensus

→ bounded time to reach consensus.

- Show that synchronous rounds can be emulated in partially synchronous model

↳ When  $\Delta$  is known, use that to decide how long to wait for communication

→ When  $\Delta$  is unknown, keep increasing until it hits the right value.

