

# Byzantine Failures

## Failure Models

- Showed up in Partial Synchrony
- Fail stop :- Failed process takes no further steps
- Omission :- some messages might never be processed

[Aside: DISKSTRA PRIZE THIS YEAR

Impossibility of consensus in  
Synch. model under enough  
omissions]

Today's  
Focus !!

- Byzantine [Lamport: Oral Messages]

- Byzantine w/  
Authentication [Lamport: Signed Messages]

## Why care about Byzantine Failures

① Systems where some participants are malicious

- PBFT. Malicious clients?

Not our core

- Cryptocurrency?

..

) focus today.

PBFT handling of malicious clients

## ② Bugs

- Fail stop (and # of likely failures) are assumptions.

But these assumptions might not hold

↳ Failed process might not stop sending  
msgs

→ Network might corrupt messages, replay  
previously sent messages, etc.

- Byzantine makes fewer assumptions

- safe (& hopefully live) even under  
arbitrary behavior

Where

- Space (the final frontier)

- Airplanes (Boeing 777, 787 flight control)


[The Real Byzantine Generals - Driscoll  
2004]

- Submarines
- some cloud services

What faulty nodes can do under Byzantine failures

- Pretty much anything
- : But - visibly

- Not send messages
- send malformed messages
- Equivocate



- send messages to a subset of processes

Desirable Properties

- Agreement: All non-faulty nodes agree on value (condition  $|I| < 1$ )

- Validity is a bit weird.

- Assume single proposer [Leader]

If proposer is correct  $\Rightarrow$

value that correct nodes agree on  
is proposed value

- Alt: Weak unanimity [Part. Synchronous]

$\hookrightarrow$  No Byz. failure  $\Rightarrow$

Validity from fail-stop, etc.

- Termination / Liveness

Of particular importance in BFT

Want to ensure that Byz.

processes cannot impede progress.

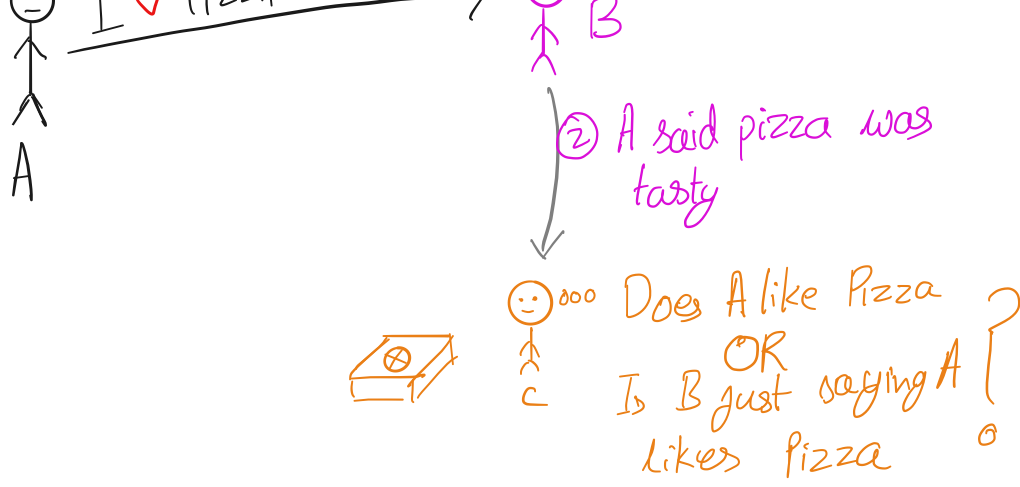
Authentication vs Not

[Note, some papers also refer to this as transferable  
auth]

Core Problem







Auth  $\Rightarrow$  Can prove who sent a message.


How? Cryptography

$$\boxed{f_{\text{sig}_A}(m) \rightarrow s} \quad \text{Secret: known only to A}$$

$$f_{\text{verif}_A}(s) \rightarrow \text{Yes or No} : \text{Public}$$

Hardness assumption:

Guessing/computing  $f_{\text{sig}_A}(m)$  impossible for anyone other than A



A

tasty, <I ♥ PIZZA>\_BA



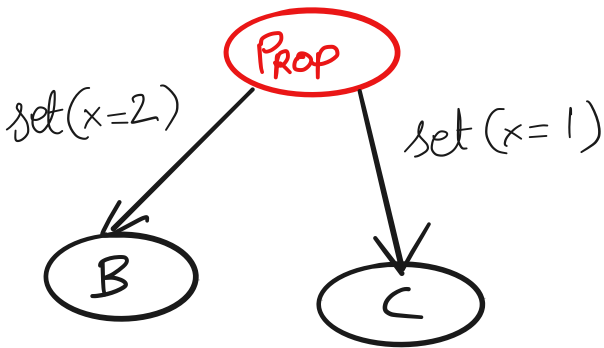
A likes Pizza!!

Lamport: Core result:

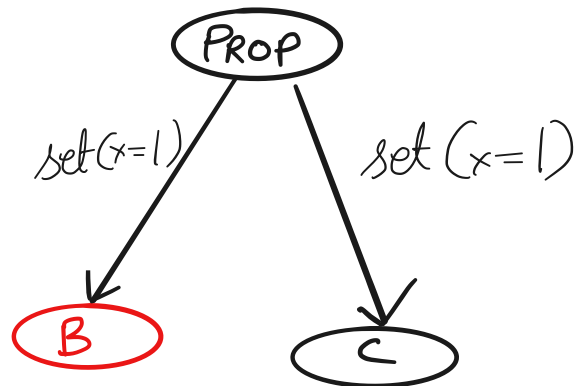
Without authentication, need  $\geq 3f+1$  nodes to tolerate  $f$  failures

Q1 · Asynch OR Partially Synchronous OR Synch.?

Why? Start with 3 nodes, possibly 1 faulty



Case P





Case B

# Focus on Node C



- Agreement: All honest nodes agree.

B & C

↳ Case P: B & C must agree / communicate



- Validity: If Prop is honest, must accept its proposal




Situations appear indistinguishable to C.

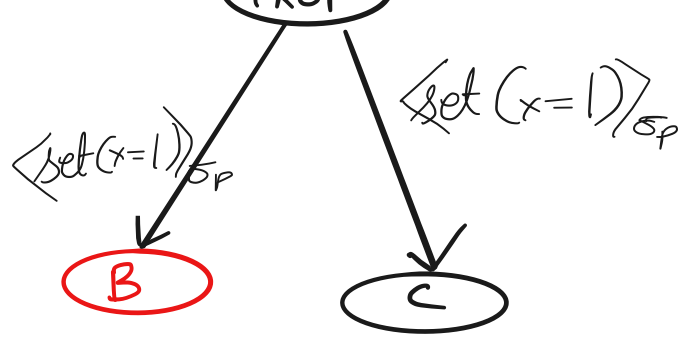
Aside: Why does auth help?

Prop

Prop



Case P




Case B

$OM(n)$ : Protocol for agreement in Byz. w/o authentication w/ up to  $n$  failures.

- Core idea: Reach agreement on what was proposed


Agreement how? Majority vote.


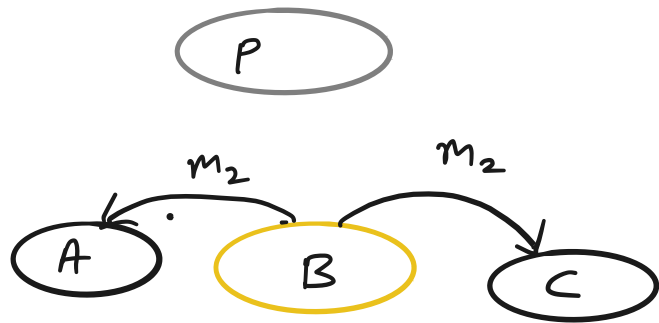

$OM(0)$



No failures, hence safe!

$OM(1)$





	P	A	B	C
A	$m_1$		$m_2$	$m_3$
B	$m_2$	$m_1$		$m_3$
C	$m_3$	$m_1$	$m_2$	

Does this work?

- Validity :- if P is honest

$$m_1 = m_2 = m_3 = m$$


Since at most one of A, B, C dishonest  
m is majority

- Agreement :- Easier: Pick deterministically  
if no majority

- Termination?

- OM(2)... OM(n)

Simple induction proof



Moving on to partial synch/asynch.

- Quorums for fail stop

↳ Assume  $f$  failures

↳ At any step

↳ No failures so far

↳ **Safety**: Enough nodes to

ensure at least 1 correct node is involved

↳  $f+1$  participants

...

→  $f$  failures so far

**Liveness:** Must have enough nodes to ensure progress.

Need  $n = f + f + 1$  participants =  $2f + 1$  participants

- Quorums for BFT: Up to  $f$  nodes

- A failed node can participate or not.

**Safety**: Failed nodes should not be able to decide next protocol state

Decision requires  $f + f + 1$  nodes

[Correct nodes always out-vote faulty ones]

**Liveness**: Must make progress if failed nodes do not participate

$$N \geq f + 2f + 1 = 3f + 1$$

## PBFT

- System assumptions
  - Partial synchrony
  - Authentication [kind of]
- Core technique: Quorum intersection
  - ↳ Multipaxos / Viewstamp replication like
    - w/ DIFFERENCES for BFT

## DIFFERENCES


- Larger quorum
- Leader is less special
  - ↳ + Decides what to propose
  - Protocol is symmetric otherwise
- Changing leader requires  $2f + 1$  nodes to suspect leader
  - [Leader stability]

- Clients




- Gather proof that commands are committed [safety]
- Can trigger leadership change [liveness]

## Replication




Client request

### Prepare



- $\langle C \rangle_\sigma$
- $\langle \text{View (term), } n \text{ (slot), } d \rangle_A$
- Hash of  $\langle C \rangle_\sigma$




### Prepare

- $\langle \text{view, } n, d, B \rangle_B$


$n$  (majority)  $\geq 2c$

OK



On receiving  $\geq 2f$   
 Prepare for  
 a prev. prepared  
 command: Commit  
 $\langle \text{view},$   
 $n,$   
 $d,$   
 $c \rangle_c$

OK



On receiving  $\geq 2f$   
 Commit for command  
 that we have seen  $\geq 2f$   
 prepare messages for  
 Respond to client.

Client counts and decides.

## Leader Election/View Change

- Goal: Faulty nodes should not disrupt correct leader

How's

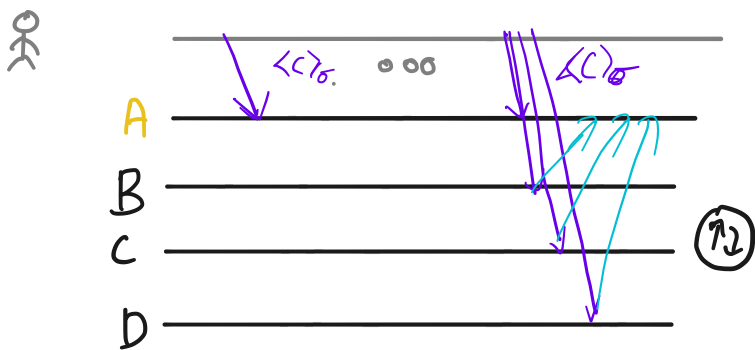
Current view (term) remains active until  
 $2f + 1$  nodes vote to move to new  
 view

- Goal: Faulty leaders should not be able to monopolize  
 leadership

How's: Associate views with leaders

Beyond this, view change is Multi Paxos like with the new leader merging logs from the  $2f$  nodes proposing view change.

## Client Triggered View Change



Clients can notify other nodes about commands which have not received responses

- If node has not seen command it forwards it to leader & sets a timer.

## PBFT in Practice

- Not widely used
- Performance
- # machines

of machines  
- "Independence" assumption

- Tradeoff

- effect of bugs

- costs