



Distributed

Systems ◦ Randomized Consensus

Randomized Consensus Protocols

Why? Because of FLP

There is ^① ~~No~~ deterministic, ^② fault tolerant
CONSENSUS PROTOCOL (AGREEMENT, VALIDITY,
TERMINATION) in the asynchronous model.

Partial Synchrony:

There is ~~No~~ deterministic, fault tolerant
CONSENSUS PROTOCOL (AGREEMENT, VALIDITY,
TERMINATION) in the ~~asynchronous model~~.

Today's goal: Relax a different constraint

There is ^{RANDOMIZED} ~~No~~ deterministic, fault tolerant
CONSENSUS PROTOCOL (AGREEMENT, VALIDITY,
TERMINATION) in the asynchronous model.

Why?

- Revisit binary consensus; Binary consensus \rightarrow RSM
- May remove some assumptions about deployment - How to SET TIMEOUTS
- Intellectually interesting
- From people here ☺

BEN-OR CONSENSUS

- Binary Consensus



PROTOCOL RUNS



↓
DECIDE \emptyset
OR 1

- Agreement: If process p decides v and process q decides v' then $v = v'$

- Termination: Eventually some process p decides.

- Validity: The decided value v must have been some processes input

↳ If all inputs are \emptyset then protocol decides \emptyset

$\emptyset \ \emptyset \ \emptyset \Rightarrow$ Decide \emptyset

If all inputs are 1 then protocol decides 1


$1 \ 1 \ 1 \Rightarrow$ Decide 1

$\emptyset \quad 1 \quad \emptyset \Rightarrow \text{DECIDE } \emptyset$ } Both are OK!
 $\emptyset \quad 1 \quad \emptyset \Rightarrow \text{Decide } 1$ }

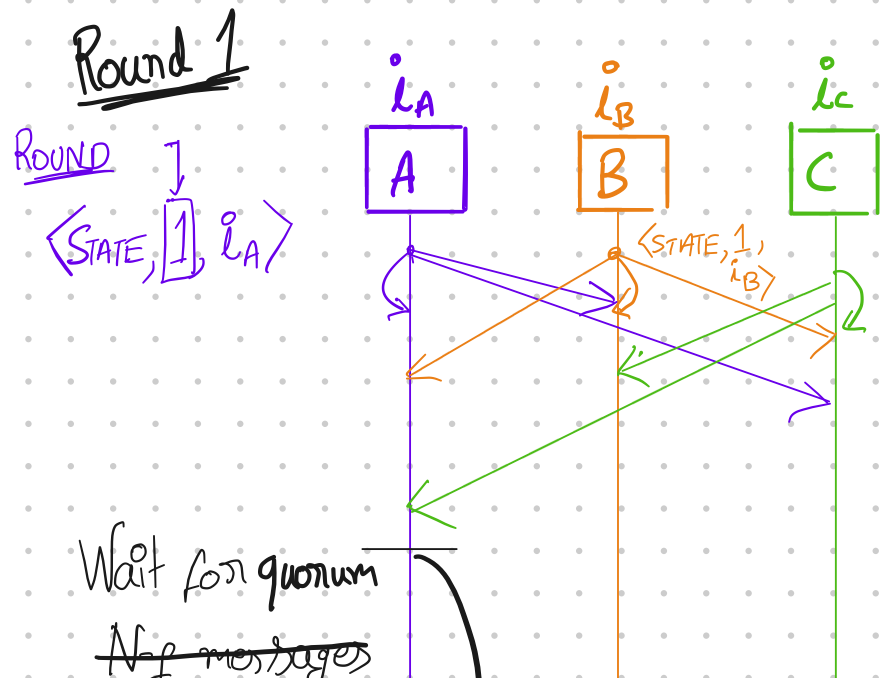
BEN-OR: Binary consensus with

- Agreement
- Validity
- Termination w/ probability 1

CORE TECHNIQUES



- QUORUM INTERSECTION
 - RANDOM COIN FLIPS TO BREAK CONFLICTS
- ↳ Think back to Paxos discussion
 no termination with competing
 proposers Section 2.4

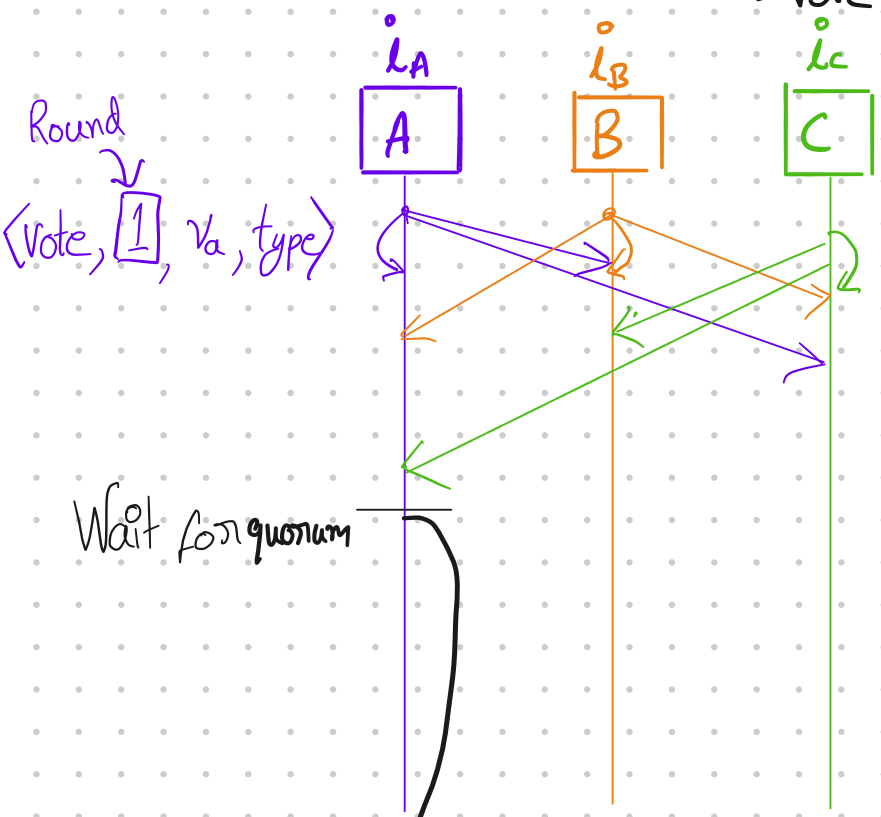


PAPER USES
 MESSAGES OF
 TYPE 1 $(1, \sigma, \dots)$
 & TYPE 2 $(2, \nu, \dots)$

Going to use
 type 1 \equiv state
 type 2 \equiv vote

→ Cases

- Quorum has the same value v
 - vote value $v_a = v$ type = D
- Else
 - vote value $v_a = ?$ type = L



→ Cases

① One message

$\langle \text{vote}, 1, v, D \rangle$

} Some node saw evidence that a majority had the same value v

$i_A = v$

② $\geq f+1$ $\langle \text{vote}, 1, v, D \rangle$ } A quorum saw
message } that a majority
had the same value
 v

Decide v

[Note, this also means for all
future rounds $i_A = v$]

③ No message of type D

$i_A =$ Randomly choose \emptyset & 1

Go to Round 2

- State

- Vote

...

Why does it work

- What happens if

$$i_A = i_B = i_C = \emptyset \text{ (or } 1)$$

at round R?

Dec. \emptyset

- What happens if a quorum has the same input, e.g., $i_A = i_B = 1$?

1 or ?

- Claim: If node A decides in round R, then other nodes decide in round R+1. Why?

$\geq f+1$ $\langle \text{vote}, R, v, D \rangle$ } A quorum saw
message } that a majority
had the same value
 v

Decide v

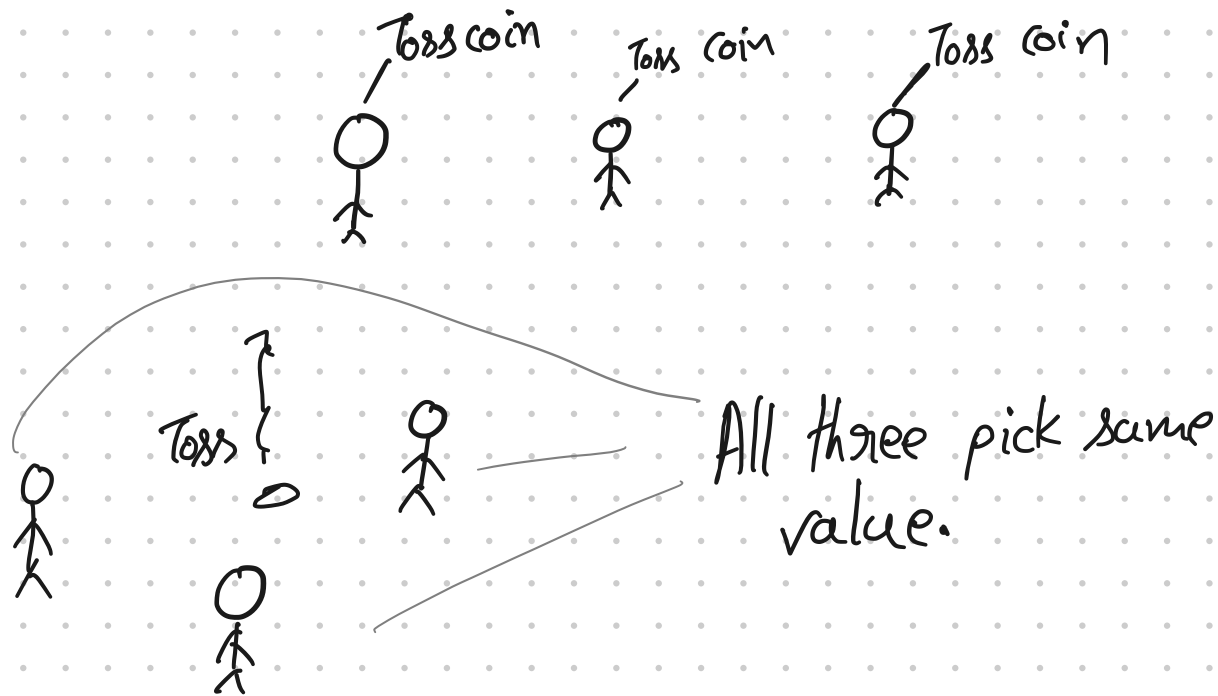
[Note, this also means for all
future rounds $i_A = v$]

Optimization: Common Coin

③ No message of type D



i_n = Randomly choose 0 & 1

Remember: Want to get all nodes to propose the same value




USING BEN-OR IN AN RSM


- Kind of saw this when building multipaxos



- Similar approach



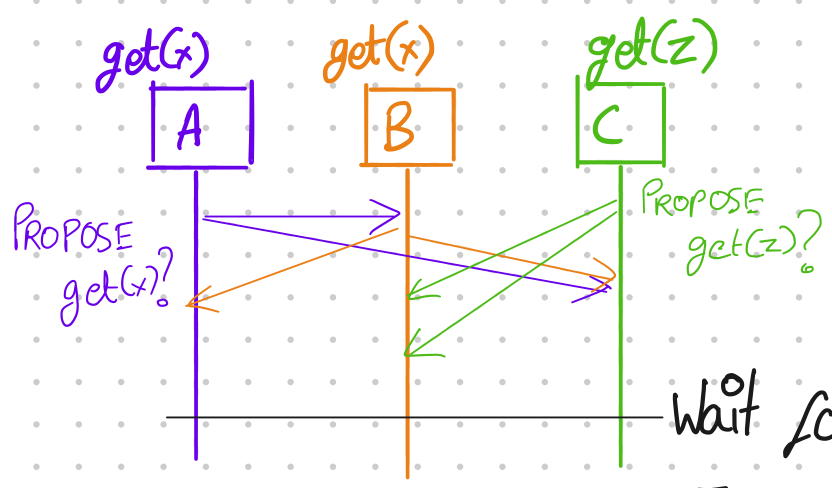
Unique challenge: Binary consensus → Arbitrary commands




WHAT DO A, B & C USE AS
INPUT TO BEN-OR?

Possibilities — Leader chooses command;
 ↳ That is what we did for multi-paxos.

PABIA — USE A PROTOCOL TO CHOOSE ONE (OR TWO) PROPOSALS
 ↳ Be wary: Why isn't this just consensus?






- wait for $n-f$ (quorum)
- If quorum has some command
 - Ben On input [for round 0] = C [or COMMON COMMAND]
- Else - Ben On input = \perp



How to deal with this situation?

Purpose of timestamps?

$T_0, \text{get}(x)$ $T_1, \text{set}(y, 1)$ $T_2, \text{get}(z)$



Priority queue with a consistent ordering

Propose from queue head

Propose from queue head

- Any node (A) that sees the same command C from a quorum enters Ben-On with C.

- Other nodes use \perp

What to do with coin toss?

Remember:

③ No message of type D

$i_A = \text{Randomly choose } 0 \& 1$

Want coin to randomly choose b/w
 c & \perp

But, getting to this stage \Rightarrow

(a) Some command c was agreed as
a proposal by a quorum of
nodes.

Why?

(b) ≥ 1 node did not see a
quorum support the command c

Why?

Might not know c ?

Errata: Any node arriving at
coin toss can recover c

Implementing Rabin

- Might be instructive to see what an implementation
looks like

(\rightarrow Will make one available later this week

Old
Lab 4

(in Elixir, using the framework
you used for labs)

→ Might help further explore ideas

Prep for next week — BFT

- On faults (remember discussion about failure
models last week)

→ fail-stop / crash

→ fail-recover / crash-recover

Assumptions?

What if the assumptions do not hold?

Why would they not hold?

Effects?

- Part (midterm question)

- Rubia

BFT

What can we do with no assumptions about the behavior of faulty (failed) nodes.

Note, we are still making assumptions about the FAILURE MODEL

↳ # of nodes failed
(often at a cost)