

NBAC (Problem 2)

The problem

- Give an algorithm which implements NBAC in an asynchronous environment, using a *Reliable broadcast*, *Best-effort broadcast (beb)*, and failure detectors $\diamond P$, and $\langle \triangleright \rangle S$.

What is NBAC?

- NBAC1. Agreement: No two processes decide differently
- NBAC2. Termination: Every correct process eventually decides
- NBAC3. Commit-Validity: 1 can only be decided if all processes propose 1
- NBAC4. Abort-Validity: 0 can only be decided if *some process crashes* or votes 0

Implementing NBAC

- In the class:
 - $\text{NBAC} = \text{beb} + \text{ucons} + P$
- Is P really necessary?
 - NBAC4 mentions “some process”
 - hence, no (there are weaker FDs which solve NBAC)

?P

- **Anonymous Completeness:** If some process crashes, then there is a time after which every correct process permanently detects a crash.
- **Anonymous Accuracy:** No crash is detected unless some process crashes.

Outline of the idea

- Every process casts its vote
- Waits for all other votes, or a crash from ?P
- If there was a crash, or there is a vote “no”
 - run *consensus* with “abort”
 - else, run *consensus* with “commit”
- Implement consensus

Pseudo code (book)

Implements:

NonBlockingAtomicCommit, instance nbac.

Uses:

BestEffortBroadcast, instance beb;

UniformConsensus, instance uc;

EventuallyPerfectFailureDetector, instance ?P.

```
upon event nbac, Init do
```

```
  voted :=  $\emptyset$ ;
```

```
  proposed := FALSE;
```

Pseudo code (book)

```
upon event <?P,Crash> do
  if proposed = FALSE then
    trigger <uc, Propose | ABORT>;
    proposed := TRUE;
```

```
upon event <nbac, Propose | v> do
  trigger <beb, Broadcast | v>;
```


Pseudo code (book)

```
upon event <beb, Deliver | p, v > do
  if v = ABORT  $\wedge$  proposed = FALSE then
    trigger <uc, Propose | ABORT>;
    proposed := TRUE;
```

else

```
voted := voted  $\cup$  {p};
```

```
if voted =  $\Pi$   $\wedge$  proposed = FALSE do
  trigger <uc, Propose | COMMIT>;
  proposed := TRUE;
```

*we wait for all processes here!
the trick is that we won't wait forever, due
to the eventual failure detector*


```
upon event <uc, Decide | decided> do
  trigger <nbac, Decide | decided>;
```

Consensus

- FLP: Impossible to solve in asynchronous environment, when one process may crash
 - can be circumvented by using Failure Detectors and correct majority
- In the class:
 - solving Consensus using P, and $\langle \rangle P$

Can we use $\diamond S$?

- Short answer: “yes”
- Long answer follows...



- **Strong Completeness:** Eventually, every process that crashes is permanently suspected by every process.
- **Eventual Weak Accuracy:** Eventually, *some correct process is never suspected.*
- we can have a unique correct process recognized by other processes (hint, hint)

Outline of the idea

- There is a correct process which is not suspected by anyone
- hence, if that process is the leader during some time, it can make a decision, and use reliable broadcast to disseminate

Algorithm

- Use rotating coordinator
- Asynchronous “rounds”
 - all messages are either to, or from the coordinator
 - each round has 4 phases

Algorithm

- Phase I: Every node sends the estimate to *the coordinator*, timestamped with the round in which it adopted the value
- Phase II: *the coordinator* waits for majority of estimates
 - picks the one with the highest timestamps
 - broadcasts the new estimate

Algorithm

- Phase III:
 - if a node suspect *the coordinator*, it *nacks* the estimate
 - otherwise, it *adopts* the estimate, and *acks* it
- Phase IV:
 - the coordinator waits for majority of *responses*
 - if the coordinator gets **a majority of acks**, it rbcasts the decision

Pseudo code: Init

Implements:

UniformConsensus, instance uc.

Uses:

EventuallyStrongFailureDetector, instance $\diamond S$;

BestEffortBroadcast, instance beb;

UniformReliableBroadcast, instance urb.

upon event $\langle uc, Init \rangle$ do

round := 1; suspected := \emptyset ;

estimate := \perp ;

votes := $[\perp]^N$; estimates := $[\perp]^N$;

Pseudo code: Phase I

upon event $\langle uc, Propose \mid v \rangle$ such that proposal = \perp
do

 estimate := v;
 begin_round();

function begin_round() do
 trigger $\langle beb, Broadcast \mid [ESTIMATE, tsp, estimate] \rangle$

Pseudo code: Phase II

upon event $\langle \text{beb}, \text{Deliver}|p, [\text{ESTIMATE}, r, ts, v] \rangle$ such that $r = \text{round}$
do //only leader
 $\text{est}[p] = (ts, v)$

upon event $\#(\text{est}) > N/2$ do // only leader
 $(ts, \text{estimate}) := \text{highest}(\text{est});$
 trigger $\langle \text{beb}, \text{Broadcast}|[\text{PROPOSE}, \text{round}, \text{estimate}] \rangle$

Pseudo code: Phase III

```
upon event <beb, Deliver| p, [PROPOSE,r,v]>  
    such that leader(round) = p do  
    estimate := v;  
    tsp := r;  
    trigger <beb, Broadcast|[ACK,r]>;  
    round = round+1;  
    begin_round();
```

Pseudo code: Phase III

```
upon event <◇S, Suspect | p> do
  suspected := suspected ∪ {p};
  if leader(round) = p then
    trigger <beb, Broadcast|[NACK,round]>;
    round = round+1;
    begin_round();
```

```
upon event <◇S, Restore | p> do
  suspected := suspected \ {p};
```

Pseudo code: Phase IV

upon event $\langle \text{beb}, \text{Deliver}|p, [\text{ACK}, r] \rangle$ such that $r = \text{round}$

do // only leader

 votes[p] := ack;

upon event $\langle \text{beb}, \text{Deliver}|p, [\text{NACK}, r] \rangle$ such that $r = \text{round}$

do // only leader

 votes[p] := nack;

upon $\#(\text{votes}) > N/2$ do // only leader

 if $|\{v: v \text{ in votes} \wedge v = \text{ack}\}| > N/2$ then

 trigger $\langle \text{urb}, \text{Broadcast}|[\text{DECIDED}, \text{estimate}] \rangle$;

upon $\langle \text{urb}, \text{Deliver}|[\text{DECIDED}, v] \rangle$ do

 trigger $\langle \text{uc}, \text{Decide}|v \rangle$;