

## Exercise Session 7

### GM and VSC

November 9, 2015

#### Problem 1

Show that  $P$  is the weakest failure detector for Group Membership.

##### Solution

In order to show that  $P$  is the weakest failure detector for Group Membership, we need to show that:

- $P$  can be used to implement Group Membership.
- Group Membership can be used to implement  $P$ .

The first direction stems directly from the Group Membership implementation in the class.

For the second direction, we assume that all processes run Group Membership algorithm. Whenever a new view is installed, all processes that are freshly removed from the view are added to the suspected set. This approach satisfies both *Strong Completeness* and *Strong Accuracy* of  $P$ , directly from the corresponding properties of Group Membership.

#### Problem 2

Note: This is a continuation of Problem 3 from Exercise Session 6 (last week).

In this problem we will change the *view-synchronous communication (VSC)* abstraction in order to allow joins of new processes. Answer to the following questions:

1. ~~Are the properties of VSC (as given in the class) suitable to accommodate the joins of new processes. Why / Why not? (done)~~
2. ~~Change the properties of VSC, so that they allow for implementations that support the joins of new processes. (Hint: focus on the properties of group membership) (done)~~
3. Sketch the changes we need to perform on the Consensus-based (Algorithm II) implementation of VSC in order to support joins.

#### Solution 2.3

The solution is described in Algorithm 1, 2 on the last two pages of this document. The changes to the regular algorithm are highlighted in red (note that we used the consensus algorithm that appears in the book — it is similar in spirit to the version in the slides).

We add two new local variables to the algorithm: *joined* and *crashed*. The *joined* variable is a boolean flag that is set to true after the process successfully joins a view (is part of the view members). The *joined* flag differentiates the behavior of processes that are just attempting to join. The *crashed* variable is a local set that keeps track of crash events received from the failure detector. This set is useful in executions where a process  $p$  attempts to join and then crashes. If another correct process  $p_2$  sees the join attempt only after the crash notification, it needs to remember that it has already seen a crash of  $p$  and to disregard the join.

For most events, the only difference to the original algorithm is that we impose the condition  $joined = true$  for event handlers. Recall that such a conditional event handler means that the events are implicitly buffered until the condition becomes true (see the document describing the language used for module specification in “Additional Material” section on the course website). For example, the crash handler is now conditioned by  $joined = true$ . This means that any crash event received by the process while it is still joining will be buffered. The events will, however, be handled right after the process successfully joins a view.

The joining begins when the application emits a *Join* event (line 21). If the process has not joined yet and is not part of the initial set of processes in the view, the process broadcasts a *JoinReq* message to every other process. The *JoinReq* message can be seen as a dual of the *crash* event. It will be the job of the receiving correct processes (that are already view members) to handle the join and propose the addition of the joining process to a view.

Upon receiving a *JoinReq* message (line 23), processes will add the joining process to their correct set. Note that if the receiving process has already seen a crash of the joining process, the correct set will not be changed ( $\{p\} \setminus crashed$  will be  $\emptyset$ ). Changing the *correct* set will trigger the handler at line 37 and initiate a view change. Processes that have seen the broadcast from the joining process will propose it in the new view member set. Since the joining process uses best-effort broadcast, correct processes will eventually receive the *JoinReq* broadcast message (if the joining process is also correct).

Another difference with the initial algorithm is that once a decision is taken in the consensus and a process moves to a new view, every process broadcasts the new view (both its member set and id). This broadcast is useful for joining processes. If a joining process sees that it is part of a new view, it will initialize its view id, member set and correct set accordingly. Finally, the joining process sets the *joined* flag to true (meaning that it will handle all buffered events) and emits a *JoinOk* indication to the application.

**Algorithm 1** View synchrony with joins, first part

---

```
1: Implements:
2:   VSCJ (vscj)
3: Uses:
4:   UniformConsensus (ucons)
5:   BestEffortBroadcast (beb)
6:   PerfectFailureDetector (P)

7: upon event  $\langle vscj, Init \rangle$  do
8:    $(vid, M) := (0, \Pi)$ 
9:    $correct := \Pi$ 
10:   $flushing := false; blocked := false; wait := false;$ 
11:   $pending := \emptyset; delivered := \emptyset; crashed := \emptyset$ 
12:  forall  $m$  do  $ack[m] := \emptyset$ 
13:   $seen := \perp^N$ 
14:  trigger  $\langle vscj, View \mid (vid, M) \rangle$ 
15:  if  $self \in \Pi$  then
16:     $joined := true$ 
17:  else
18:     $joined := false$ 
19:  end if

20:  upon event  $\langle vscj, Broadcast \mid m \rangle$  such that  $blocked = false \wedge joined = true$  do
21:     $pending := pending \cup \{self, m\}$ 
22:    trigger  $\langle beb, Broadcast \mid [DATA, vid, self, m] \rangle$ 

23:  upon event  $\langle vscj, Deliver \mid p, [DATA, id, s, m] \rangle$  such that  $joined = true$  do
24:    if  $id = vid \wedge blocked = false$  then
25:       $ack[m] := ack[m] \cup \{p\}$ 
26:      if  $(s, m) \notin pending$  then
27:         $pending := pending \cup (s, m)$ 
28:        trigger  $\langle beb, Broadcast \mid [DATA, vid, s, m] \rangle$ 
29:      end if
30:    end if

31:  upon  $\exists (s, m) \in pending : M \subseteq ack[m] \wedge m \notin delivered \wedge joined = true$  do
32:     $delivered := delivered \cup \{m\}$ 
33:    trigger  $\langle vscj, Deliver \mid s, m \rangle$ 

34:  upon event  $\langle P, Crash \mid p \rangle$  such that  $joined = true$  do
35:     $correct := correct \setminus \{p\}$ 
36:     $crashed := crashed \cup \{p\}$ 

37:  upon  $correct \neq M \wedge flushing = false \wedge joined = true$  do
38:     $flushing := true$ 
39:    trigger  $\langle vscj, Block \rangle$ 

40:  upon event  $\langle vscj, BlockOk \rangle$  such that  $joined = true$  do
41:     $blocked := true$ 
42:    trigger  $\langle beb, Broadcast \mid [PENDING, vid, pending] \rangle$ 

43:  upon event  $\langle beb, Deliver \mid p, [PENDING, id, pd] \rangle$  such that  $id = vid \wedge joined = true$  do
44:     $seen[p] := pd$ 

45:  upon  $\forall p \in correct : seen[p] \neq \perp \wedge wait = false$  do
46:     $wait := true$ 
47:     $vid := vid + 1$ 
48:    initialize a new instance  $uc.vid$  of uniform consensus
49:    trigger  $\langle uc.vid, Propose \mid (correct, seen) \rangle$ 
```

---

---

**Algorithm 2** View synchrony with joins, second part

---

```
1: upon event  $\langle uc.id, Decide \mid M', S \rangle$  do
2:    $\forall p \in M' : S[p] \neq \perp$  do
3:      $\forall (s, m) \in S[p] : m \notin delivered$  do
4:        $delivered := delivered \cup \{m\}$ 
5:       trigger  $\langle vscj, Deliver \mid s, m \rangle$ 
6:    $flushing := false; blocked := false; wait := false$ 
7:    $pending = \emptyset$ 
8:    $\forall m$  do  $ack[m] := \emptyset$ 
9:    $seen := [\perp]^N$ 
10:   $M := M'$ 
11:  trigger  $\langle vscj, View \mid (vid, M) \rangle$ 
12:   $\forall p \in M$  do
13:    trigger  $\langle beb, Broadcast \mid [NewView, vid, M] \rangle$ 

14:  upon event  $\langle beb, Deliver \mid [NewView, vid', M'] \rangle$  such that  $joined = false$  do
15:    if  $self \in M'$  then
16:       $(vid, M) := (vid', M')$ 
17:       $correct := M$ 
18:       $joined := true$ 
19:      trigger  $\langle vscj, JoinOk \rangle$ 
20:    end if

21:  upon event  $\langle vscj, Join \mid self \rangle$  such that  $joined = false$  do
22:    trigger  $\langle beb, Broadcast \mid [JoinReq, self] \rangle$ 

23:  upon event  $\langle beb, Deliver \mid [JoinReq, p] \rangle$  such that  $joined = true$  do
24:     $correct := correct \cup \{p\} \setminus crashed$ 
```

---