

Distributed Algorithms, Bonus Exam

xx

December 14, 2012

Name:

Sciper number:

Assume that we work in the fail-stop model and consider Algorithm 1.

Recall that in the fail-stop model we consider a set of n processes $\{p_1, p_2, \dots, p_n\}$ and we suppose that any process may stop taking steps at any time and that all processes have access to a perfect failure detector. Moreover, when a process stops taking steps, it never takes any step again.

Observe that Algorithm 1 is similar to an Algorithm you saw in the course and that both solve non-uniform consensus in the fail-stop model.

Your task is to answer the following question:

Is it possible to obtain a uniform consensus algorithm by modifying Algorithm 1 as follows?

- Adding new variables to the algorithm and possibly initializing them in the $\langle cons \mid Init \rangle$ handler.
- Replacing the pseudo-code of the $\langle urb, Deliver \mid round, v \rangle$ handler (lines 18 to 22) by a new block of pseudo-code that may do any local computation but that must not trigger any event other than the *Decide* event.

To clarify, by local computation we mean variable assignment, “if” blocks, and “while” blocks. Your modifications should not contain any trigger event except for the *Decide* event.

Write a detailed explanation of your answer, focusing on both the Agreement and Termination properties.

Algorithm 1 Non-Uniform consensus algorithm, code for process p_i

```
1: Implements:  
2:   Consensus ( $cons$ )  
  
3: Uses:  
4:   UniformReliableBroadcast ( $urb$ )  
5:   PerfectFailureDetector ( $P$ )  
  
6: upon event  $\langle cons, Init \rangle$  do  
7:    $suspected \leftarrow \emptyset$   
8:    $round \leftarrow 1$   
9:    $currentProposal \leftarrow \perp$   
10:   $delivered \leftarrow [false]^N$   
  
11: upon event  $\langle P, Crash \mid p \rangle$  do  
12:   $suspected \leftarrow suspected \cup \{p\}$   
  
13: upon event  $\langle cons, Propose \mid v \rangle$  do  
14:   if  $currentProposal = \perp$  then  
15:      $currentProposal \leftarrow v$   
16:   end if  
  
17: upon event  $\langle urb, Deliver \mid round, v \rangle$  do  
18:   if  $round = i$  then  
19:     trigger  $\langle cons, Decide \mid v \rangle$   
20:   end if  
21:    $currentProposal \leftarrow v$   
22:    $delivered[round] \leftarrow true$   
  
23: upon event  $delivered[round] = true$  or  $p_{round} \in suspected$  do  
24:    $round \leftarrow round + 1$   
  
25: upon event  $round = i$  and  $currentProposal \neq \perp$  do  
26:   trigger  $\langle urb, Broadcast \mid p_i, currentProposal \rangle$ 
```

Solution

Whatever the modifications, the resulting algorithm would not solve uniform consensus.

Agreement

First, let us consider agreement.

Suppose that a process decides before hearing from everybody else (either receiving a message or a crash notification). In other words, suppose that p_i decides after receiving a message or the crash notification from a set of processes S , where $S \subset \Pi$, $S \neq \Pi$. Suppose there is a correct process p_j such that p_i did not deliver any message from p_j before deciding.

Suppose now that p_i crashes, along with all processes in S . Suppose that p_j receives the crash notifications for p_i and all processes in S before it delivers the *urb* broadcasts. p_j cannot know whether the processes in S broadcast anything before crashing, therefore p_j must decide on a value without waiting for a deliver from the processes in S . Since p_j did not receive the value proposed by p_i or by the processes in S , p_j may decide on a different value. Therefore, p_i and p_j may decide differently, violating uniform agreement. This means that p_i cannot decide before hearing a message or a crash notification from every other process. (Basically, in the round-based algorithm, it means that processes must either use acknowledgements from every correct process or wait until the final round before deciding.)

Termination

Now let us consider termination.

We proved previously that we need to wait for a message or crash notification from every process before deciding.

Consider now a scenario where there is a set S of correct processes and a faulty process p_i , $p_i \notin S$, where i is greater than the ids of all processes in S . Suppose that every process $p \in S$ receives a message from every other process in S , but not from p_i . This means that all processes are waiting to hear a message from p_i or to hear of p_i 's crash before deciding. Since the *urbDeliver* function only accepts messages from a particular sender (the sender id must be equal to the local *round* variable) and all processes are waiting only for a message from p_i , it means that the local *round* of at least one process is equal to i (p_i 's id). We denote that process by p_j .

Suppose that p_i crashes before broadcasting any message. Notice that the crash handler of the correct processes will add p_i to *suspected* and increment the local round number. Therefore, the local *round* of p_j will become $i + 1$. However, since in our scenario i is the highest id, there is no process p_{i+1} . This means that p_j will now wait forever for messages from a process that does not exist. The *urbDeliver* of p_j will never be triggered again, meaning that p_j will never decide and termination will be violated.

Elegant proof by one of your colleagues.

Suppose that there is a modification of algorithm 1 that, respecting the rules of the problem, solves uniform consensus.

First make the following two observations:

(1) The Decide event can only be triggered in the urbDeliver handler. This is because you are not allowed to modify the other handlers and they do not trigger any Decide event.

(2) The urbDeliver handler is executed by a process p_i in a round r only if p_r successfully broadcasts its message to p_i .

Now Consider an execution e in which all processes but p_i are faulty and never send any message.

(3) By (2), in execution e , p_i will execute the urbDeliver handler only in round i .

Suppose that p_i doesn't decide in round i . Since, by (1), p_i can only decide in the urbDeliver handler and, by (3), this handler is only executed in round i , we conclude that p_i will never decide and termination will be violated.

This conclusion is in contradiction with the assumption that the algorithm solves uniform consensus. Thus we have:

(4) In execution e , p_i decides in round i .

Now consider execution e' that is the same as e up to round i but where

- p_i crashes in round $i + 1$,
- all messages from p_i to other processes are delayed,
- processes p_j with $j > i$ do not crash.

To p_i , e' is indistinguishable from e up to round i . Hence, by (4), p_i decides a value in round i .

However, since all messages from p_i are delayed, the other processes do not now about p_i 's decision before deciding and may violate uniform agreement.

This is in contradiction with the fact that the modified algorithm solves uniform consensus. Hence this assumption cannot hold and we have proved that no modification of algorithm 1 that respects the constraints of the problem can solve uniform consensus.