

Distributed Algorithms, Final Exam

January 14, 2013

Solution

1 Multiple Choice Questions (15 points)

Question 1. (2 points) 1, 3

Question 2. (2 points) 1, 2, 3

Question 3. (2 points) 1, 2, 4

Question 4. (2 points) 2

Question 5. (2 points) 3

Question 6. (5 points)

1. L

2. S

3. L

4. S

5. S

2 Reliable Broadcast (13 points)


Question 1. (1 point) Slides or book.

Question 2. (6 points)

1. Uniform reliable broadcast – No



2. Causal broadcast – No



3. Terminating reliable broadcast – No



Question 3. (6 points) This is a best-effort causal broadcast abstraction. Accordingly, on a crash-free execution (all processes are correct) agreement is guaranteed due to validity.

Of course, the crash-free case is not that interesting, so let's discuss what happens in case there are crashes. What happens when a message is broadcast? A broadcast of message m by process p enforces all other processes to receive all the messages that belong to the causal past of m ¹. This, of course, includes the messages that were delivered and the messages that were broadcast by p before message m . It should be clear that this is a direct consequence of the definition of causality:

A message m_1 causally precedes a message m_2 ($m_1 \rightarrow m_2$) when:


1. both are broadcasts of the same process and m_1 was broadcast before m_2

¹if p does not crash while broadcasting m , case that could lead to m not being delivered by every process

2. m_1 is a broadcast of p_1 and m_2 is a broadcast of p_2 and m_2 was sent after p_2 delivered m_1
3. $m_1 \rightarrow m'$ and $m' \rightarrow m_2$ entails $m_1 \rightarrow m_2$ (transitivity)

So, in an execution where the correct processes keep broadcast messages, the causal delivery property ensures that all the delivered messages of a process before m will be delivered before delivering m , ensuring agreement even in the case of crashes.

However, we cannot guarantee that every process will send an infinite number of messages, so the following execution is possible:



As you can see, p_1 delivers message m sent by p_0 just before p_0 crashed. Due to the crash, m was not delivered by p_2 . If p_1 stays inactive (as it happens in the execution above), p_3 is not guaranteed to deliver message m , hence violating agreement.

Consequently, the broadcast algorithm of the question **does not guarantee the agreement property** in executions that there are crashes.

3 View Synchronous Communication (10 points)

Question 1. (4 points) See the lecture's slides or the book.

Question 2. (4 points) See ?? for the solution.

Algorithm 1 Implementation of Stoppable Broadcast using UTRB

```
1: Implements:
2:   StoppableBroadcast, instance sb
3: Uses:
4:   UniformTerminatingReliableBroadcast, instances  $utrb.p_i$  with sender  $p_i \in \Pi$ 
5:   ReliableBroadcast, instance rb
6: upon event  $\langle sb, Init \rangle$  do
7:    $delivered \leftarrow \emptyset$ 
8:    $trbdone \leftarrow \emptyset$ 
9: upon event  $\langle sb, Broadcast \mid m \rangle$  do
10:  trigger  $\langle rb, Broadcast \mid m \rangle$ 
11: upon event  $\langle rb, Deliver \mid p, m \rangle$  do
12:  if  $(p, m) \notin delivered$  then
13:     $delivered \leftarrow delivered \cup \{(p, m)\}$ 
14:    trigger  $\langle sb, Deliver \mid p, m \rangle$ 
15:  end if
16: upon event  $\langle sb, Stop \rangle$  do
17:  trigger  $\langle utrb.p_i, Broadcast \mid delivered \rangle$ 
18: upon event  $\langle utrb.p_i, Deliver \mid p_i, m \rangle$  do
19:   $trbdone \leftarrow trbdone \cup \{p_i\}$ 
20:  if  $m \neq \phi$  then
21:    forall  $(s, m') \in m'$  do
22:      if  $m' \notin delivered$  then
23:         $delivered \leftarrow delivered \cup \{m'\}$ 
24:        trigger  $\langle sb, Deliver \mid s, m' \rangle$ 
25:      end if
26:    end if
27: upon event  $trbdone = \Pi$  do
28:  trigger  $\langle sb, StopOk \rangle$ 
```

Question 3. (4 points) See ?? for the solution.

To understand the algorithm, remember that events that are not triggered because of a “**such that**” construct (lines 12 and 14) are buffered and delivered, in the order in which they arrived, when the conditions becomes true.

The solution assumes that processes do not receive StopOk before calling Stop. To be precise, the specification of Stoppable Broadcast does not enforce this. However the algorithm is simpler with this assumption and it would be easy to modify it so that it would work without the assumption.

Algorithm 2 View-Synchronous Communication using Stoppable Broadcast

```
1: Implements:
2:   ViewSynchronousCommunication, instance vs
3: Uses:
4:   StoppableBroadcast, instances sb.i,  $i \in \mathbb{N}$ 
5:   GroupMembership, instance gm
6: upon event  $\langle vs, Init \rangle$  do
7:    $viewId \leftarrow 0$ 
8:    $changingView \leftarrow false$ 
9:    $nextView \leftarrow \perp$ 
10: upon event  $\langle vs, Broadcast \mid m \rangle$  do
11:   trigger  $\langle sb.viewId, Broadcast \mid m \rangle$ 
12: upon event  $\langle sb.i, Deliver \mid p, m \rangle$  such that  $i = viewId$  do
13:   trigger  $\langle vs, Deliver \mid p, m \rangle$ 
14: upon event  $\langle gm, View \mid v \rangle$  such that  $nextView = \perp$  do
15:    $nextView \leftarrow v$ 
16: upon event  $nextView \neq \perp$  and  $changingView = false$  do
17:    $changingView \leftarrow true$ 
18:   trigger  $\langle vs, Block \rangle$ 
19: upon event  $\langle vs, BlockOk \rangle$  do
20:   trigger  $\langle sb.viewId, Stop \rangle$ 
21: upon event  $\langle sb.viewId, StopOk \rangle$  do
22:    $viewId \leftarrow nextView.id$ 
23:    $changingView \leftarrow false$ 
24:   trigger  $\langle vs, View \mid nextView \rangle$ 
25:    $nextView \leftarrow \perp$ 
```

4 Shared Memory (8 points)

Question 1. (2 points)

1. not safe

```
P1 [ W(1) ]
P2          [ R()->0 ]
```

2. safe, but not regular

```
P1 [ W(1) ] [ W(2) ]
P2          [ R()->0 ]
```

3. regular, but not atomic

```
P1 [ W(1) ] [ W(2) ]
P2          [ R()->2 ]
P3          [ R()->1 ]
```

4. atomic

```
P1 [ W(1) ] [ W(2) ]
P2          [ R()->1 ]
P3          [ R()->2 ]
```

Question 2. (6 points)

1. The solution in the course has each reader also write its value to every process before returning its value. This ensures that nobody has an older version than the reader. However, due to the fact that readers can also issue writes, the processes need to check the timestamp before accepting a write.

The algorithm in the exam has the readers only check the timestamp of other processes before returning. Since only the Writer process can issue write commands, there is no longer the need to check the timestamp when accepting a write command — the timestamp of new writes is guaranteed to be newer, since the writer ensures increasing timestamps. However, as shown below, the algorithm in the exam does not solve 1-N atomic registers.

2. The algorithm does not solve the 1-N atomic register problem.

By reading the timestamps from everyone, the algorithm can indeed detect potential concurrent writes. However, the versioning implemented by the algorithm is not complete. The algorithm only holds two versions of data, "new" and "old", which is not enough. Since only a successful Read sets val_{old} , it might happen that the value of val_{old} is very old. The issue is shown in the execution below.

Writer	[W(1)]		[W(2)]
Reader1			[R()->0]
Reader2	[R()->1]		[R()->1]

Suppose all registers are initialized to 0 ($val_{new} = val_{old} = 0$). There are two writes issued. The first write is not concurrent with any other operation and terminates successfully. Thus, both *Reader1* and *Reader2* will store 1 in val_{new} and 1 in ts_{new} . The first read by *Reader2* completes successfully and set val_{old} , val_{new} , ts_{old} and ts_{new} to 1. Then, during the second write, *Reader1* receives the new value before attempting the read, thus val_{new} is 2 and ts_{new} is 2. Notice that, val_{old} and ts_{old} are still 0 for *Reader1*. However, *Reader2*, does not yet receive the new values by the time *Reader1* starts reading. *Reader1* issues a read, detects that *Reader2* has old values and thus returns val_{old} . However, val_{old} is still 0 at this point, since *Reader1* never had a successful non-concurrent read. This execution does not respect an atomic register specification (and not even a regular register specification).