University of Freiburg
Dept. of Computer Science
Prof. Dr. F. Kuhn
M. Fuchs, A. Mályusz, Z. Parsaeian

Theory of Distributed Systems
Sample Solution Exercise Sheet 1

Exercise 1: Schedules

Consider three nodes, v_1, v_2, and v_3, which are connected via FIFO channels, that is, messages between any two nodes are received in the same order they are sent. For example, if node v_1 sends first message m_1 then m_2 to node v_2, then v_2 will first receive m_1 and then m_2.
Devise **one** possible schedule S which is consistent with the following local restrictions to the three nodes.

- $S|1 = s_{1,3}\ s_{1,3}\ r_{1,2}\ r_{1,3}\ s_{1,2}\ r_{1,2}\ s_{1,3}$,
- $S|2 = s_{2,3}\ s_{2,1}\ r_{2,1}\ s_{2,1}$,
- $S|3 = r_{3,2}\ r_{3,1}\ s_{3,1}\ r_{3,1}\ r_{3,1}$.

$s_{i,j}$ denotes the send event from node i to node j and $r_{j,i}$ denotes the event that node j receives a message from node i.

Sample Solution

There could be more than one possible global schedule S. Two possible ones are the following.

$$s_{2,3}\ s_{1,3}\ r_{3,2}\ r_{3,1}\ s_{2,1}\ s_{1,3}\ s_{3,1}\ r_{1,2}\ r_{1,3}\ r_{3,1}\ s_{1,2}\ r_{2,1}\ s_{2,1}\ r_{1,2}\ s_{1,3}\ r_{3,1}.$$

$$s_{1,3}\ s_{2,3}\ r_{3,2}\ r_{3,1}\ s_{1,3}\ s_{2,1}\ s_{3,1}\ r_{3,1}\ r_{1,2}\ r_{1,3}\ s_{1,2}\ r_{2,1}\ s_{2,1}\ r_{1,2}\ s_{1,3}\ r_{3,1}.$$

One can also get a possible solution by drawing the graphical diagram as in the lecture.

Exercise 2: The Level Algorithm

Consider the following algorithm between two connected nodes u and v:

The two nodes maintain levels ℓ_u and ℓ_v, which are both initialized to 0. One round of the algorithm works as follows:

1. Both nodes send their current level to each other
2. If u receives level ℓ_v from v, u updates its level to $\ell_u := \max\{\ell_u, \ell_v + 1\}$. If the message to node u is lost, node u does not change its level ℓ_u. Node v updates its level ℓ_v in the same (symmetric) way.

If the level algorithm runs for r rounds:

(a) What can you say about the level of the two nodes?

(b) If all messages succeed, what can you say about the level of the two nodes?

(c) In what case the level of a node is at least one?

Sample Solution

a) We argue by induction on the number of rounds r that : after r rounds, the two levels differ by at most one. For a round r, let ℓ_u^r and ℓ_v^r be the levels of nodes u and v after round r. We have that $r \geq 0$. So for the base case: for $r = 0$ we have that $\ell_u^0 = 0 = \ell_v^0$. For the inductive step: assume that the statement holds after round r i.e. , we have $|\ell_u^r - \ell_v^r| \leq 1$, which is equivalent to $\ell_v^r - 1 \leq \ell_u^r \leq \ell_v^r + 1$. We have

$$\ell_u^{r+1} \leq \max\{\ell_u^r, \ell_v^r + 1\} \overset{\text{I.H.}}{=} \ell_v^r + 1 \leq \ell_v^{r+1} + 1$$

where the last inequation holds because levels can only increase.

Analogously, we prove $\ell_v^{r+1} \leq \ell_u^{r+1} + 1$.

b) Induction on the number of rounds: At the beginning (after round 0), we have $\ell_u^0 = \ell_v^0 = 0$. Now assume $\ell_u^r = \ell_v^r = r$ and in round r both messages succeed. Then $\ell_u^{r+1} = \max\{\ell_u^r, \ell_v^r + 1\} = \max\{r, r+1\} = r + 1$ and $\ell_v^{r+1} = \max\{\ell_v^r, \ell_u^r + 1\} = r + 1$.

c) For the forward direction: if a node never receives a message, it never updates its level (which is initially 0). So if its level is at least one, then it must have received a message.
For the backward direction: if node u receives $\ell_v \geq 0$ in some round, then its level becomes $\ell_v + 1 \geq 1$ which never decreases again.

Exercise 3: The Randomized Two Generals Algorithm

Now assume we have two nodes u and v running the following modified protocol:

- Each node has an input $x_u, x_v \in \{0, 1\}$.
- Node u picks a random number $R \in \{1, \ldots, r\}$ uniformly at random.
- The nodes run the **Level Algorithm** (as in Exercise 2) for r rounds (In each message, both nodes also include their inputs and node u also includes the value of R.).
- At the end, a node outputs 1 if:
  - It knows that both inputs are 1,
  - It knows the value of R,
  - Its own level is at least R.
- Otherwise, it outputs 0.

If the level algorithm runs for r rounds:

(a) If at least one input is 0, what is the output of the two nodes?

(b) If both inputs are 1:

- what is the output if no message is lost?
- under what circumstances the output of the two nodes is not the same value?

(c) If both inputs are 1, what is the probability that both nodes output the same value?

(d) Using the same technique as in the impossibility proof for the deterministic Two Generals Problem (discussed in the lecture), prove a lower bound for the error probability.

Sample Solution

Throughout the solution let $r \in \mathbb{N}$ be the number of rounds and let $R \in \{1, \dots, r\}$ be the random number chosen by node u. Denote by $\ell_u, \ell_v \in \{0, \dots, r\}$ the (deterministic) *final level* obtained by the Level Algorithm at u and v, respectively.

(a) At least one input is 0. Assume without loss of generality that the input of node u is 0. Since u knows its own input, it will deterministically output 0 regardless of the communication pattern. We now consider the behavior of node$v$.
If no message is delivered from u to v, then v does not learn the value of R, violating the second output condition. Hence, v outputs 0.
If at least one message from u is delivered to v, then v learns that the input of u is 0 (as every message includes the sender's input). Consequently, the first output condition ("both inputs are 1") is violated. Hence, v also outputs 0.
In both cases, the nodes agree and output 0.

(b) Both inputs are 1.

(i) *No message is lost.* When all r messages of every round are delivered: after round t both nodes have learned the other's level $t-1$ and therefore increase their own level to t. Thus $\ell_u = \ell_v = r$. Node v receives R in the very first message, so it also knows R. Every node therefore fulfils the three output conditions and the common output is 1.

(ii) *When can the outputs differ?* Based on the output of level algorithm, the only possible gap is $\ell_u = \ell_v + 1$. Let $L := \min\{\ell_u, \ell_v\}$. If $R \le L$ then both nodes satisfy "level $\ge R$" and output 1; if $R > L + 1$ then neither node satisfies it and both output 0. Hence disagreement can occur *only* when

$$R = L + 1 = \ell_u = \ell_v + 1.$$

In other words, exactly one more round was successfully acknowledged by u than by v, and the random counter happens to hit this delicate level. In this case u outputs 1 and v outputs 0.

(c) Probability of agreement when both inputs are 1. Fix an arbitrary loss pattern. After the r rounds each node has a deterministic level ℓ_u, ℓ_v with the property $|\ell_u - \ell_v| \le 1$. Conditional on the loss pattern the random variable R is still uniformly distributed in $\{1, \dots, r\}$. The nodes disagree exactly if $R = \ell_u$ and $\ell_u = \ell_v + 1$ (the situation described above). Whatever the values of ℓ_u, ℓ_v are, at most

$$\Pr[\text{disagreement} \mid \text{pattern}] \;\le\; \frac{1}{r}.$$

Taking the expectation over all loss patterns gives

$$\Pr[\text{disagreement}] \le \frac{1}{r}, \qquad \text{so} \qquad \boxed{\Pr[\text{agreement}] \ge 1 - \tfrac{1}{r}.}$$

(d) A lower bound on the error probability. We prove a lower bound on the error probability by constructing an indistinguishability sequence in which messages are gradually removed, and we track the cumulative probability of an incorrect output. This mirrors the classical Two Generals impossibility argument, adapted to the randomized setting.
Assume both nodes u and v have input 1. Let the protocol run for r rounds. We define a sequence of executions $\mathcal{E}_0, \mathcal{E}_1, \dots, \mathcal{E}_{2r}$, where in each step we remove one message (either from u to v or from v to u), starting from the case with full delivery:

- $\mathcal{E}_0$: All $2r$ messages (in both directions) are delivered. In this case, both nodes should output 1 with probability at least $1 - \varepsilon$.
- In $\mathcal{E}_1$: The last message from u to v is lost. Since this changes only v's view, the probability that v outputs 0 increases by at most ε.

- In $\mathcal{E}_2$: The last message from v to u is also lost. Now u's view may be affected, so the probability that either node outputs 0 increases to at most 2ε.

- ...

- In general, after t message losses (i.e., in $\mathcal{E}_t$), the probability that at least one node outputs 0 is at most $t \cdot \varepsilon$.

- Finally, in $\mathcal{E}_{2r}$: All messages are lost. In this case, node v receives nothing and cannot distinguish this execution from one where u's input is 0. Hence, v must output 0 with probability 1.

So, starting from $\mathcal{E}_0$ (where both nodes output 1 with high probability) and gradually dropping messages until $\mathcal{E}_{2r}$ (where at least one node must output 0 with probability 1), the total cumulative change in output must be at least 1.
Thus, we conclude:

$$2r \cdot \varepsilon \geq 1 \quad \Rightarrow \quad \boxed{\varepsilon \geq \frac{1}{2r}}.$$

This proves that every r-round randomized protocol has worst-case error at least $1/(2r)$. The randomized protocol analyzed in parts (a)–(c) achieves error at most $1/r$, so it is asymptotically optimal up to a constant factor.