



Distributed Systems Part II

Solution to Exercise Sheet 3

1 Consensus in a Grid

- a) The goal of this protocol is for every process to know the input of all the other processes. Each process internally allocates an array of size $w \cdot h$, each entry of the array represents the input of one process and initially is set to '?. The process fills the array until no '?' are left, then the process can make its decision (e.g. the minimal value received). Whenever a process gets new information, it sends its updated array to all its four neighbors. At the beginning the "new information" is its own input.
- b) The number of required rounds is the length of the longest shortest path between any two processes. In the case of a grid this is $w + h$.
- c) If $w \approx h$ one can arrange the faulty nodes in two diagonals (see figure 1) the longest shortest path has a length of $\approx 3 \cdot (w + h)$. So the number of required rounds is about $3 \cdot (w + h)$. In each case it is possible to arrange the faulty nodes in a pattern like in figure 2. So a longest shortest path of $2 \cdot (w + h) + c$ where c is a small constant can be achieved.




Figure 1: Strategy for $w \approx h$ with faulty nodes marked as black. The longest shortest path leads from t to s




Figure 2: Strategy which can be done in every grid and achieves a longest shortest path of $\approx 2 \cdot (w + h)$

- d) Byzantine processes can catch and modify messages. In a setting like in figure 3 only one of four messages reaching t is correct. Process t cannot find out which messages are wrong, thus there is a high probability for t to make false decisions.




Figure 3: Byzantine processes (dark) can “bombard” process t with wrong information and lead t to a wrong decision

2 Consensus in a Hypercube

- a) We only need to alter how broadcast is implemented, all other parts of the king algorithm remain unchanged.

In a hypercube broadcast is not as easy as in a fully connected network. Processes need to forward the messages, and any byzantine process can alter the messages it has to forward.

We can increase the chance for a message to reach its destination unaltered by sending the message over multiple paths. To be more exact: in a hypercube with dimension m we can always find m non-intersecting paths between two processes.

A process may receive the same message up to m times, but the faulty processes could have altered some of them. The process needs to decide what the original message was. Since there is no authentication available, the process can do nothing but assume the message that is most often received is the original message. This leads to $f < m/2$.

- b) The king algorithm has $f + 1$ phases, in each phase the algorithm sends three broadcast. So the algorithm requires $3 \cdot (f + 1)$ times the duration of a broadcast.

How much time does a broadcast require? Approximately m rounds, because the longest shortest path leading from one corner of the hypercube to the opposite corner has length m .

To be more exact, a broadcast requires $m + 1$ rounds: If process a sends to process b , then b must receive the message up to m times. In a hypercube every process has coordinates which we can express as binary strings of length m . A path between a and b can be seen as sequence of bits that are switched. Example: if $c_a = 000$ and $c_b = 111$ then one path from a to b first visits the processes at 001 (switching bit 3), then the process at 011 (switching bit 2), and finally b (switching bit 1).

The length k of the shortest path between a and b is the number of bits c_a and c_b differ in. There are k non-intersecting paths of length k between a and b . (The i 'th path would start with switching the i 'th different bit, then the $i + 1$ 'th different bit, etc.). This leaves $m - k$ paths with a length greater than k .

These $m - k$ paths need first to switch a bit which is the same in c_a and c_b (otherwise the path would intersect with one of the shortest paths). This bit needs to be switched back at the end of a path. Other than that a path needs to switch k bits that differ, leading to $k + 2$ operations all together.

We set k as big as possible, but ensure $k < m$. This means we set $k = m - 1$ and find the longest path a message must ever travel has length $m + 1$.

3 Consensus with Byzantine Failures

No 2-resilient algorithm for 6 processes without authentication exists.

This is a proof by contradiction: It is known that consensus for 3 processes, one of them Byzantine, cannot be solved. This proof can be found in the lecture notes on slides 108ff. Assume now that consensus for 6 processes could be solved by an algorithm a . We create a new algorithm b which solves consensus for 3 processes using a . The algorithm b works as follows: each of the 3 processes (hosts) simulates 2 processes (guests). The input of the guests is equal to the input of their hosts. Correct hosts simulate correct guests, the Byzantine host simulates Byzantine guests. The 6 guests solve consensus using a . The hosts copy the decision of their guests (all guests make the same decision). This is a contradiction! Since b works if a works, therefore a must be wrong (as it is known, that b cannot exist.) Hence there exists no algorithm a solving consensus for 6 processes.

4 Authentication

a) The new algorithm looks like this:

```

if I am P then
     $values \leftarrow \{input\}$ 
    broadcast "P has input"
else
     $values \leftarrow \{\}$ 
end if
for  $r = 0$  to  $f + 1$  do
    for all received values  $x$  do
        if  $|values| < 2$  and accepted  $r$  messages "P has  $x$ " with  $x \notin values$  then
             $values \leftarrow values \cup \{x\}$ 
            broadcast "P has  $x$ "
        end if
    end for
end for
if  $|values| = 1$  then
    decide item in  $values$ 
else

```

decide “sender faulty”
end if

- b) If P is correct: there is only one message in the system, which is accepted in the first round. There are no other messages, hence for all processes $|values| = 1$.

If P is Byzantine:

- Assume that a correct process p adds x to its value set in a round $r < f + 1$: Process p has accepted r messages including the message from P. Therefore all other correct processes accept the same r messages plus p 's message and add x to their value set as well in round $r + 1$.
- Assume a correct process p adds x to its value set in round $f + 1$: In this case, p accepted $f + 1$ messages. At least one of those is sent by a correct process, which must have added x to its set in an earlier round. We are again in the previous case, i.e., all correct processes added x to its value set.

5 Randomization

- a) The algorithm can handle $f < n/8$ failures. To find this result we check the proofs for the validity condition, agreement, and termination for numbers that change:

Validity condition Nothing changes

Agreement Nothing changes

Termination If some process does not set its value randomly, all processes must set the same value, i.e. there must not be $n - 4f$ proposals for 0 and $n - 4f$ proposals for 1. This means that $2 * (n - 4f) > n$, or $f < n/8$.

The reason why this property changes is that Byzantine processes can create two different messages, while simple crashing processes cannot.

- b) One solution is to replace “if at least $n-4f$ proposals” by “if at least $n-3f$ proposals”. There are other correct solutions which will not be discussed in this master solution.

- c) We modify the proof from the lecture. For the agreement property we replace “Every other correct process must have received x at least $n - 4f$ times.” by “... at least $n - 3f$ ” times.

Why? Because $n - 3f$ correct processes had to send their proposal in order for one process to decide. Meaning $n - 3f$ correct processes sent their proposal to any process.

Rewriting the termination property leads to $2 * (n - 3f) > n$, or $f < n/6$.