

## Exercise 1: Global vs. Local

### Task 1: Even more Globally Optimal (60 points)

In the lecture, we saw how to achieve an upper bound of  $\mathcal{O}((\vartheta - 1)d + u)D$  on the global skew. Now, we will essentially remove the front term.

- a) Consider a node  $v$  and one of its neighbors  $w \in N_v$ . Let  $v$  and  $w$  bounce messages back and forth, i.e.,  $v$  sends a timestamped message at time 0, which  $w$  sends back with the same timestamp. Upon reception,  $v$  sends another timestamped message, and so on. Use this to derive a “logical hardware clock”  $\hat{H}_v$  satisfying  $\hat{H}_v(t) - \hat{H}_v(t') \in (1 + u/d)(t - t') + \mathcal{O}(u + (\vartheta - 1)d)$  for all  $t > t'$  and  $H_v(0) + t \leq \hat{H}_v(t)$ , but which can be set back. (25 points)
- b) Consider the “Min Algorithm,” in which nodes reduce their logical clock value to  $L + d$  when receiving a message  $\langle L \rangle$ , but use  $\hat{H}_v$  as “hardware” clock at  $v$ . (The message sent are specified in c)). Define a “virtual clock”  $L_{\min}$  that lower bounds all actual clock values, but also takes into account messages that are in transit, which might cause nodes to set their clocks back. This clock should satisfy that it increases at least at rate 1. Conclude that the algorithm guarantees amortized 1-progress. (10 points)
- c) Bound the global skew of the algorithm. We are not concerned with the number of sent messages, so nodes will send a message immediately after receiving one, as well as send one every  $d$  local time (w.r.t.  $H_v$ ). (Remark: If it helps with notation, feel free to show the statement only on a graph that is a simple path.) (25 points)

### Task 2: The Least Possible Locally Optimal (40 points)

The *local skew* is defined in a similar way as the global skew except that we consider the logical clock difference between the neighboring nodes only:

$$\mathcal{L} := \sup_{t \in \mathbb{R}_0^+} \{\mathcal{L}(t)\},$$

where

$$\mathcal{L}(t) := \max_{\{v,w\} \in E} \{|L_v(t) - L_w(t)|\}$$

In this exercise, we analyze the local skew of the Refined Max Algorithm.

- a) Let  $P = (v_0, \dots, v_D)$  be a path of length  $D$ . Construct an execution in which (i)  $H_v(0) = 0$  for all  $v \in V$ , (ii)  $H_{v_{i+1}}(t_0) - H_{v_i}(t_0) = u$  for all  $i \in [D]$ , (iii)  $L_v(t) = H_v(t)$  for all  $v \in V$  and times  $t$ , and (iv) each hardware clock runs at rate 1 after time  $t_0$ . (Hint: Simply set all message delays to  $d$  and ramp hardware clock speeds. Then show that  $L_v(t) = H_v(t)$  for all  $v \in V$  and  $t$ , because no node  $v$  receives a message with  $L > H_v(t) - d + u$  at any time  $t$ .) (20 points)
- b) Take the above execution and modify it such that a skew of  $u(D - 1)$  occurs at some time between nodes  $v_0$  and  $v_1$ . (Hint: Only change message delays after time  $t_0$ , nothing else. “Pull” the nodes to the current maximum clock value one by one, starting with  $v_{D-1}$ .) (20 points)

### Task 3\*: Global and Local Happiness

- a) Find out what the term synchronizers refers to in distributed computing!
- b) Use basic techniques for synchronizers to devise an algorithm that satisfies constant amortized progress, has asymptotically optimal global skew, and has local skew  $\mathcal{O}(uD)$ . You may assume that  $(\vartheta - 1)d \leq u$ , so you don't have to worry about removing  $(\vartheta - 1)d$  terms.
- c) Synchronize (i.e., communicate your findings) with the other students in the TA session!