Exercise 3: Control Issues

Task 1: Controlling the Global Skew

In the lecture, we proved that the GCS algorithm achieves local skew $\mathcal{O}(\log \mathcal{G})$, where $\mathcal{G}$ is maximum global skew in the network. However, we did show that GCS maintains any bound on $\mathcal{G}$ itself. Here you will analyze a variant of the GCS algorithm that maintains $\mathcal{G} = \mathcal{O}(D)$, where D is the network diameter. Thus the performance of this variant matches the lower bound proven in the previous lecture.

Recall that our analysis of the GCS algorithm only requires that nodes satisfying **FC** (respectively **SC**) run in fast (respectively slow) mode. The modification of the algorithm you describe here will modify the behavior of GCS only when neither **FC** nor **SC** applies in such a way that the algorithm maintains a bound on the global skew.

(a) Add the condition that any node $v \in V$ satisfying $L_v(t) = \max_{w \in V}\{L_w(t)\}$ is in slow mode at time t and determine a suitable trigger condition. Show that your trigger condition does not conflict with **FT**.

(b) Apply the techniques used in the (refined) Max Algorithm to maintain an estimate $M_v(t)$ of the largest clock value throughout the system at each $v \in V$. Show that $\max_{v \in V}\{L_v(t)\} \geq M_v(t) \geq \max_{v \in V}\{L_v(t)\} - \mathcal{G}_{\max}$ for some $\mathcal{G}_{\max}$. (Hints: Make minimal modifications to the Max Algorithm, so that the reasoning changes very little. This way, you can argue that the proof of the bound is analogous. Note that you need to be slightly more careful regarding the rate at which nodes increase the estimates when $L_v(t) < M_v(t)$: use rate $h_v/\vartheta \leq 1$. You should obtain $\mathcal{G}_{\max} = ((\vartheta - 1/\vartheta)T + (\vartheta - 1)d + u)D$.)

(c) Show that $L_v(t) = \min_{w \in V}\{L_w(t)\}$ implies that v does not satisfy **ST** at time t.

(d) Assume that $\sigma = \mu/(\vartheta - 1) > 1$ and that $\max_{v \in V}\{H_v(0)\} \leq \mathcal{G}_{\max}$. Add the condition that any node $v \in V$ satisfying $L_v(t) < M_v(t)$ such that **ST** does not hold at time t is in fast mode. Conclude that the modified algorithm has global skew $\mathcal{G} \leq \mathcal{G}_{\max}$ and still obeys **FC** and **SC**. What is the resulting local skew, provided that $\max_{\{v,w\} \in E}\{H_v(0) - H_w(0)\} \leq \delta$?

Solution

a) The trigger condition is that $L_v(t) \geq \tilde{L}_w(t)$ for all $w \in N_v$, which is implied by $L_v(t) = \max_{w \in V}\{L_w(t)\}$. If **FT**holds, there are $s \in \mathbb{N}$ and $x \in N_v$ so that $\tilde{L}_x(t) - L_v(t) > (2s-1)\delta > 0$, i.e., the new trigger condition cannot hold.

b) We modify the Max Algorithm as follows.

Observe that (i) $M_v(t) \geq L_v(t)$ at all times, (ii) M_v is never increased beyond $\max_{v \in V}\{L_v(t)\}$ (either $M_v(t) = L_v(t)$ or it increases at rate $h_v/\vartheta \leq 1$, and it is never set to a value larger than $\max_{v \in V}\{L_v(t)\}$, as messages are under way for at least $d - u$ time), and (iii) $M_v(t)$ increases at least at rate $1/\vartheta$ while being stored, while estimates "increase at rate 1" while travelling as a message. Arguing as in the lecture, we thus get a bound of $\mathcal{G}_{\max} \leq ((\vartheta - 1/\vartheta)T + (\vartheta - 1)d + u)D$.

c) If $L_v(t) = \min_{w \in V}\{L_w(t)\}$, then $\tilde{L}_x(t) > L_x(t) - \delta \geq L_v(t) - \delta$ for all $x \in \mathbb{N}_v$. Thus, $L_v(t) - \tilde{L}_x(t) < \delta \leq (2s-1)\delta$ for any $s \in \mathbb{N}$ and $x \in N_v$, so **ST** 1 does not hold.

d) Assume for contradiction that the global skew reaches $\mathcal{G}_{\max} + \varepsilon$ for some $\varepsilon > 0$. Thus, there is a time t_1 so that $\max_{v \in V}\{L_v(t_1)\} - \min_{v \in V}\{L_v(t_1)\} = \mathcal{G}_{\max} + \delta$. Let t_0 be maximal such that $t_0 \leq t_1$ and $\max_{v \in V}\{L_v(t_0)\} - \min_{v \in V}\{L_v(t_0)\} = \mathcal{G}_{\max}$; as logical

Algorithm 1: Max Estimate Algorithm.

$M_v(0) := L_v(0)$
at all times t:
if $M_v(t) = L_v(t)$ **then**
  increase M_v at rate l_v
else
  increase M_v at rate h_v/ϑ
if *received* $\langle M \rangle$ *at time* t *and* $M + d - u > M_v(t)$ **then**
  $M_v(t) := M + d - u$
if $H_v(t) = kT$ *for some* $k \in \mathbb{N}$ **then**
  send $\langle M_v(t) \rangle$ to all neighbors

clocks are continuous and $L_v(0) - L_w(0) = H_v(0) - H_w(0) \leq \mathcal{G}_{\max}$ for all $v, w \in V$, such a time exists.

Consider $\max_{v \in V}\{L_v(t)\}$. Whenever $L_v(t) = \max_{w \in V}\{L_w(t)\}$, by a) v is in slow mode, i.e., $l_v(t) = h_v(t) \leq \vartheta$. It follows that $\max_{v \in V}\{L_v(t_1)\} - \max_{v \in V}\{L_v(t_0)\} \leq \vartheta(t_1 - t_0)$. Now consider $\min_{v \in V}\{L_v(t)\}$. Whenever $L_v(t) = \min_{w \in V}\{L_w(t)\}$ at a time $t \in (t_0, t_1]$, we have by b) that $M_v(t) \geq \max_{v \in V}\{L_v(t_0)\} - \mathcal{G}_{\max} > L_v(t)$ and by c) that **ST**is not satisfied at v. Hence, v is in fast mode and thus $l_v(t) = (1+\mu)h_v(t) \geq 1+\mu > \vartheta$, as $\sigma > 1$. Hence $\min_{v \in V}\{L_v(t_1)\} - \min_{v \in V}\{L_v(t_0)\} > \vartheta(t_1 - t_0)$. It follows that

$$\mathcal{G}_{\max} + \delta = \max_{v \in V}\{L_v(t_1)\} - \min_{v \in V}\{L_v(t_1)\} < \max_{v \in V}\{L_v(t_0)\} - \min_{v \in V}\{L_v(t_0)\}) = \mathcal{G}_{\max}\,,$$

a contradiction.

By a), adding the new trigger condition for slow mode is not in conflict with **FT**, and the modification from d) is not in conflict with **ST**by construction. Hence, Theorem 2.9 still applies and we can conclude that

$$\mathcal{L} \leq 2\delta \left\lceil \log_\sigma \frac{\mathcal{G}}{\delta} \right\rceil \leq 2\delta \left\lceil \log_\sigma \frac{((\vartheta - 1/\vartheta)T + (\vartheta - 1)d + u)D}{\delta} \right\rceil .$$

Task 2: Controlling Uncertainty

In the lecture, we assumed that $v \in V$ has an estimate $\tilde{L}_w$ of the logical clock L_w of each of its neighbors $w \in N_v$, satisfying that $L_w(t) - \delta < \tilde{L}_w(t) \leq L_w(t)$ at all times t. In this exercise, you determine δ for a straightforward way of deriving such an estimate. You may assume that $\max_{\{v,w\} \in E}\{H_v(0) - H_w(0)\} \leq \delta - (\vartheta(1+\mu) - 1/\vartheta)d$ throughout this exercise and that $\vartheta \in \mathcal{O}(1)$.

a) Suppose $w \in V$ sends a message with its current logical clock value whenever $H_v(t) = kT$ for some $k \in \mathbb{N}$, and also at time 0. Determine a (good) estimate $\tilde{L}_w(t)$ that $v \in V$ can compute based on this information. Bound the resulting δ. (Hint: It's ok to be a bit sloppy with lower order terms or constant factors, as long as you get the asymptotics right.)

b) For fixed values of all other parameters, determine a choice of μ asymptotically minimizing our upper bound on the local skew (i.e., up to constant factors). (Hint: Argue that $\delta \in \Omega(\mathcal{G})$ implies that the upper bound is trivial and that it doesn't matter (asymptotically) to choose μ to be at least $\max\{u/(T+d), 8(\vartheta - 1)\}$. Having ruled out these corner cases, check how the bound changes if a value of μ satisfying these constraints is doubled.)

c) For this method of determining estimates, the asymptotically optimal choice of μ you computed, and the global skew bound you obtained in the first exercise, determine the bound on the local skew as function of T (use the same value of T for global and local estimates).

Solution

a) Let t_0 be the time when $v \in V$ receives the first message from $w \in N_v$. For $t \geq t_0$, denote by t_r the most recent time when v received a message from w, and let $t_s \in (t_r - d, t_r - d + u)$ be the time when it was sent. We have that

$$\begin{aligned}
L_w(t) &\geq L_w(t_r) + t - t_r \\
&\geq L_w(t_r) + \frac{H_v(t) - H_v(t_r)}{\vartheta} \\
&\geq L_w(t_s) + (t_r - t_s) + \frac{H_v(t) - H_v(t_r)}{\vartheta} \\
&\geq L_w(t_s) + d - u + \frac{H_v(t) - H_v(t_r)}{\vartheta}\,.
\end{aligned}$$

Thus, setting $\tilde{L}_w(t) := L_w(t_s) + d - u + \frac{H_v(t)-H_v(t_r)}{\vartheta}$ at times $t \geq t_0$ is a valid choice. Now let us check by how much this choice may underestimate $L_w(t)$.

$$\begin{aligned}
&L_w(t) - \left(L_w(t_s) + d - u + \frac{H_v(t) - H_v(t_r)}{\vartheta}\right) \\
&\leq \vartheta(1+\mu)(t - t_s) - d + u - \frac{t - t_r}{\vartheta} \\
&= \left(\vartheta(1+\mu) - \frac{1}{\vartheta}\right)(t - t_r) + \vartheta(1+\mu)(t_r - t_s) - d + u \\
&< \left(\vartheta(1+\mu) - \frac{1}{\vartheta}\right)(t - t_s) + u \\
&< \left(\vartheta(1+\mu) - \frac{1}{\vartheta}\right)(T + d) + u \\
&= \left(\vartheta\mu + \frac{(\vartheta+1)(\vartheta-1)}{\vartheta}\right)(T + d) + u \\
&< (\vartheta\mu + 2(\vartheta - 1))(T + d) + u \\
&< 3\vartheta\mu(T + d) + u\,,
\end{aligned}$$

where the last step exploits that $\sigma > 1$.

Regarding the estimates to be used before the first message is received, simply initialize the estimate to $H_v(0) - \delta + (\vartheta(1+\mu) - 1/\vartheta)d$ and increase it at rate h_v/ϑ; this is a safe lower bound, which at time $t < d$ is off by no more than $\delta - (\vartheta(1+\mu) - 1/\vartheta)(d - t) < \delta$.

To see that the obtained bound is (asymptotically) optimal (given this simple communication scheme), observe that (i) a difference of u is unavoidable by indistinguishability, (ii) $\Omega(T + d)$ time passes between updates, and (iii) in the meantime, the logical clock rate of w may be anywhere between 1 and $\vartheta(1+\mu)$. (ii) and (iii) yield that $\delta \in \Omega((\vartheta - 1 + \vartheta\mu)(T + d)) = \Omega(\vartheta\mu(T + d))$.

b) Abbreviate $c := \vartheta(T + d)$ and recall that

$$\mathcal{L} \leq 2\delta \left\lceil \log_\sigma \frac{\mathcal{G}}{\delta} \right\rceil \in \Theta\left((c\mu + u)\left\lceil \log_{\mu/(\vartheta-1)} \frac{\mathcal{G}}{c\mu + u} \right\rceil\right).$$

If $c\mu < u$, we increase μ such that $c\mu = u$, which increases the bound by at most a factor of 2. Thus, we can assume $c\mu \geq u$ in the following, which implies that

$$2\delta \left\lceil \log_\sigma \frac{\mathcal{G}}{\delta} \right\rceil \in \Theta\left(c\mu \left\lceil \log_{\mu/(\vartheta-1)} \frac{\mathcal{G}}{c\mu} \right\rceil\right).$$

Next, if $\log(\mu/(\vartheta-1)) < 3$, we increase μ such that $\mu = 8(\vartheta-1)$, which increases the bound by no more than factor 8. If $\log \mathcal{G}/(c\mu) < 3$, the upper bound is in $\Omega(\mathcal{G})$ and thus trivial; in this case, reducing μ can't hurt. Having dealt with the corner cases, write $\ell_1 := \log \mathcal{G}/(c\mu)$ and $\ell_2 := \log \mu/(\vartheta-1)$. Doubling μ thus changes the term in the asymptotic expression from

$$c\mu \left\lceil \frac{\ln \mathcal{G}/(c\mu)}{\ln \mu/(\vartheta-1)} \right\rceil = c\mu \left\lceil \frac{\ell_1}{\ell_2} \right\rceil \qquad \text{to} \qquad 2c\mu \left\lceil \frac{\ell_1 - 1}{\ell_2 + 1} \right\rceil \geq c\mu \left\lceil \frac{\ell_1}{\ell_2} \right\rceil,$$

where the inequality uses the assumption that $\ell_1 \geq 3$ and $\ell_2 \geq 3$. Hence, increasing μ cannot improve the bound by more than a constant factor. As $\vartheta \in \mathcal{O}(1)$, we conclude that the choice $\mu := \vartheta u/c + 8(\vartheta-1) = u/(T+d) + 8(\vartheta-1)$ is optimal up to constant factors.

c) According to b), we can safely set $\mu := u/(T+d) + 8(\vartheta-1)$. Plugging in the bounds from a) and Task 1 d), we get that

$$\begin{aligned}
\mathcal{L} &\leq 2\delta \left\lceil \log_\sigma \frac{((\vartheta - 1/\vartheta)T + (\vartheta-1)d + u)D}{\delta} \right\rceil \\
&< 2\delta \left\lceil \log_\sigma \frac{(2(\vartheta-1)(T+d) + u)D}{\delta} \right\rceil \\
&< ((6\vartheta+2)u + 48\vartheta(\vartheta-1)(T+d)) \lceil \log_\sigma D \rceil \\
&= ((6\vartheta+2)u + 48\vartheta(\vartheta-1)(T+d)) \left\lceil \frac{\log D}{\log(u/((\vartheta-1)(T+d)) + 8)} \right\rceil \\
&\in \mathcal{O}\left((u + (\vartheta-1)(T+d)) \left\lceil \frac{\log D}{\log(u/((\vartheta-1)(T+d)) + 8)} \right\rceil\right).
\end{aligned}$$

Task 3*: Control Right from the Beginning

So far, we have largely ignored the issue of network initialization. It is unrealistic to assume that all nodes start executing the algorithm precisely at time 0. Indeed, this would require perfect synchronization! Instead, we now assume that nodes can spontanteously wake up and execute the algorithm at any time, and that a node wakes up when it receives its first message. The hardware clock of a node is 0 at the time when it wakes up. W.l.o.g., assume that at least one node wakes up at time 0.

a) Initialize the network by flooding, i.e., on wake-up, a node broadcasts a message to all its neighbors. Adapt the clock estimation technique from Task 2 to account for the modified initialization.

b) Extend the hardware clock functions, logical clock functions, and clock estimates to be defined from time 0 on such that (i) $1 \leq h_v(t) \leq \vartheta$ for all t, (ii) $L_v(t) = H_v(t)$ at times t when v has not yet woken up, (iii) $L_w(t) - \delta < \tilde{L}_w(t) \leq L_w(t)$ at all times t, and (iv) node v is in slow mode at times t when it has not woken up yet according to the (modified) GCS algorithm.

c) Convince Will and your fellow students that this approach yields the same skew bounds you computed in Tasks 1 and 2!