ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

SS 2006   Prof. R. Wattenhofer / Prof. P. Widmayer / T. Locher / R. Flury

Principles of Distributed Computing
Exercise 3: Sample Solution

1   License to Kill

a) We use a variant of the Echo Algorithm 3.7. A node (i.e. an agent in the hierarchy) matches up all (except for at most one) of his children. If the node or one of its children is left out, then the node sends a request to "match" upwards in the hierarchy. Otherwise, it sends a "no match" and that subtree is done. We give an asynchronous, uniform algorithm below.

Algorithm 1 Edge-Disjoint Matching

1: **wait** until received message from all children
2: **while** remain ≥ 2 requests (including myself) **do**
3:    **match** any two requests
4: **end while**
5: **if** exists leftover request **then**
6:    **send up** "match"
7: **else**
8:    **send up** "no match"
9: **end if**

When a node v sends a "match" request to its parent u, then the edge $\{u, v\}$ will be used only once since there will be only one request in the subtree rooted at v.

b) Let T be the tree with n nodes. Assuming each message takes at most 1 time unit, then the time complexity of Algorithm 1 is in $O(depth(T))$ since all the requests travel to the root (and back down if we inform the agents of their assigned partners). We can assume that local computation (matching) is negligible. On each link, there are at most 2 messages: 1 that informs the parent whether a match is needed and optionally 1 more to be informed by the parent of the match partner. So there are a total of at most $2(n - 1)$ messages.

2   License to Distribute

a) Again we apply an echo-style algorithm where a node locally balances the documents as much as possible. For each node v we can define

$$balance(v) := have(v) - need(v)$$

where $have(v)$ is the total number of documents and $need(v)$ is the number of nodes in the subtree rooted at v. Each node then computes and sends this balance to its parent. Again the algorithm is asynchronous and uniform. Abstractly we refer to a document as a token. When we gather the balance information from the children, they implicitly also send along any extra tokens they might have.

Algorithm 2 Token Distribution for node v

1: **wait** until received balance from all children
2: balance(v) := 0
3: **for** each child c **do**
4:     balance(v) := balance(v) + balance(c)
5: **end for**
6: balance(v) := balance(v) + tokens(v) - 1
7: **send up** balance(v)
8: **if** balance(v) < 0 **then**
9:     **wait** to receive needed tokens from parent
10: **end if**
11: redistribute tokens among children

b) The time complexity is in $\mathrm{O}(depth(T))$ analogous to Exercise 1. If we assume that we can (physically) send all the tokens in one message, then again there is one message upwards for each link and optionally one downwards with the missing tokens. Thus there are at most $2(n-1)$ messages.