Principles of Distributed Computing
Sample Solution to Exercise 2

1 License to Match

a) We use a variant of the Echo algorithm (Algorithm 2.10). A node (i.e. an agent in the hierarchy) matches up all (except for at most one) of its children. If one participating child remains and the node itself also participates, it matches itself with that child. If either the node or a one of its children remain, 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 matching algorithm below.

Algorithm 1 Edge-Disjoint Matching

1: **wait** until received message from all children
2: **while** at least 2 requests remain (including myself) **do**
3:    **match** any two requests
4: **end while**
5: **if** exists leftover request **then**
6:    **send** "match" to parent (= superior)
7: **else**
8:    **send** "no match" to parent
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. Along with the messages of the algorithm, the required path information is sent; we left this out in the pseudocode to improve readability.

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). 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

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

---

the algorithm is asynchronous and uniform. Abstractly we refer to a document as a token. When children send up the balance information to their parent, they also send along any extra tokens they might have.

b) The time complexity is in O($depth(T)$) analogous to Exercise 1. 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.

3 License to Connect

a) The algorithm terminates in two phases. In phase 1, the following fragments are created:

- N={(N,V), (N,G)}
- K={(K,H), (H,P)}
- M={(M,R)}
- A={(A,C), (A,S), (S,T)}

And in phase 2:

- N={(N,V), (N,G), (N,P), (P,H), (H,K), (K,R), (R,M), (M,C), (C,A), (A,S), (S,T)}

b) In phase 1, the edges added as a merge are (N,V), (M,R), (K,H) and (S,A). The absorbed edges are (N,G), (P,H), (C,A) and (S,T). In phase two, the only merged edge is (N,P) and the absorbed edges are (R,K) and (M,C).

(a) Fragments after phase 1.

(b) Fragment after phase 2.

(c) The resulting MST.

Figure 1: Red nodes are fragment roots, red edges are newly merged edges and purple edges are newly absorbed edges.