



FS 2016

Prof. Dr. Roger Wattenhofer
G. Bachmeier, P. Bissig, P. Brandes, S. Brandt,
C. Decker, P. Khanchandani, M. König,
D. Melnyk, L. Peer, Y. Wang

Exam

Principles of Distributed Computing

Friday, August 19, 2016
12:00 – 14:00**Do not open or turn until told to by the supervisor!**

The exam lasts 120 minutes, and there is a total of 120 points. The maximal number of points for each question is indicated in parentheses. Your answers must be in English. Be sure to always justify your answers. Algorithms can be specified in high-level pseudocode or as a verbal description. You do not need to give every last detail, but the main aspects need to be there. Big-O notation is acceptable when giving algorithmic complexities. Please write legibly. If we cannot read your answers, we cannot grade them.

Please write down your name and Legi number (your student ID) in the following fields.

| | |
|------|----------|
| Name | Legi-Nr. |
| | |

Points

| Question Nr. | Achieved Points | Maximal Points |
|--------------|-----------------|----------------|
| 1 | | 23 |
| 2 | | 26 |
| 3 | | 35 |
| 4 | | 27 |
| 5 | | 9 |
| Total | | 120 |

1 Multiple Choice

(23 Points)

Evaluate each of the following statements in terms of correctness. Indicate whether a statement is true or not by ticking the corresponding box. Each correct answer gets 1 point, **each wrong answer gets -1 point**. An unanswered statement gets 0 points. If the sum of collected points is negative, then you get 0 points for this question set.

| Statement | true | false |
|---|--------------------------|--------------------------|
| It is possible to do uniform initialization of a wireless network of n nodes without collision detection in $O(n)$ time in expectation. | <input type="checkbox"/> | <input type="checkbox"/> |
| Uniform leader election in an n -node wireless network with collision detection is possible in $O(\log \log n)$ time w.h.p. | <input type="checkbox"/> | <input type="checkbox"/> |
| A minimum dominating set can be larger than a maximal independent set. | <input type="checkbox"/> | <input type="checkbox"/> |
| If a minimum dominating set is not an independent set, then nodes can be removed from it to get a maximal independent set. | <input type="checkbox"/> | <input type="checkbox"/> |
| For the function $\text{ONES} : \{0, 1\}^k \times \{0, 1\}^k \rightarrow \{0, 1\}$ that evaluates to 1 if the two input k -bit strings contain the same number of ones and to 0 otherwise, it holds that $CC(\text{ONES}) \in \Omega(\log k)$. | <input type="checkbox"/> | <input type="checkbox"/> |
| If for a graph G it holds that $m \in \Omega(n^2)$ (where m is the number of edges and n the number of nodes of G), then the diameter of G can be computed in $O(\sqrt{n})$ time. | <input type="checkbox"/> | <input type="checkbox"/> |
| Given a weighted graph and its MST. If we add a single edge to the graph, a distributed algorithm can always update the MST in constant time. | <input type="checkbox"/> | <input type="checkbox"/> |
| Given a weighted graph and its MST. If we add a node and connect it with a single edge to any node previously in the graph, a distributed algorithm can always update the spanning tree in constant time. | <input type="checkbox"/> | <input type="checkbox"/> |
| In any anonymous graph in which we cannot deterministically find a leader, we also cannot deterministically construct a spanning tree. | <input type="checkbox"/> | <input type="checkbox"/> |
| Given an undirected graph with n nodes, if we need to query whether there exists a path between two nodes using a labeling scheme, the lower bound of the label size is $\Omega((\log n)^2)$. | <input type="checkbox"/> | <input type="checkbox"/> |
| Consider three graphs $G_1 = (V, E_1)$, $G_2 = (V, E_2)$ and $G = (V, E)$ on the same node set, where E_1 and E_2 are two edge sets, and E is the union of E_1 and E_2 . If there are adjacency labeling schemes of size k_i for G_i (for $i = 1, 2$), then there exists an adjacency labeling scheme of size $k_1 + k_2$ for G . | <input type="checkbox"/> | <input type="checkbox"/> |
| Synchronizer α may add a substantial overhead to the message complexity. | <input type="checkbox"/> | <input type="checkbox"/> |
| There exists some synchronous distributed algorithm that cannot be executed in an asynchronous environment without changing the asymptotic time complexity of the algorithm. | <input type="checkbox"/> | <input type="checkbox"/> |

| Statement | true | false |
|---|--------------------------|--------------------------|
| If a graph with n nodes can be colored with $c < n$ colors, then it can also be colored with $c + 1$ colors. | <input type="checkbox"/> | <input type="checkbox"/> |
| When using Peterson's Algorithm deadlocks can occur. | <input type="checkbox"/> | <input type="checkbox"/> |
| Compare-and-Swap is considered more powerful than Test-and-Set. | <input type="checkbox"/> | <input type="checkbox"/> |
| When using the Arrow algorithm: If there is a quiescent moment, then all the arrows point towards the node holding the variable. | <input type="checkbox"/> | <input type="checkbox"/> |
| In the Arrow algorithm, the chosen spanning tree does not influence the runtime. | <input type="checkbox"/> | <input type="checkbox"/> |
| An advantage of the Ivy algorithm over the Arrow algorithm is that it also works in an asynchronous setting. | <input type="checkbox"/> | <input type="checkbox"/> |
| The self-stabilizing MIS (Algorithm 12.5) may fail to compute a locally stable solution if the adversary is allowed to modify the network topology. | <input type="checkbox"/> | <input type="checkbox"/> |
| In the Democrats and Republicans problem, if initially the majority of citizens votes for Democrats, there will eventually be no citizen who switches her opinion every other day. | <input type="checkbox"/> | <input type="checkbox"/> |
| If a distributed problem on an n -node graph can be solved in $O(n)$ time when the message size is restricted to $O(\log n)$ bits, then it can also be solved in $O(\log n)$ time when the message size is restricted to $O(n)$ bits. | <input type="checkbox"/> | <input type="checkbox"/> |
| There is a bipartite graph in which deterministic anonymous leader election is possible. | <input type="checkbox"/> | <input type="checkbox"/> |

| Solutions | | |
|---|------|-------|
| Statement | true | false |
| <p>It is possible to do uniform initialization of a wireless network of n nodes without collision detection in $O(n)$ time in expectation. <i>Reason: The leader election algorithm without CD finishes in $O(n)$ time in expectation and then $O(n)$ expected time uniform initialization with CD can be used.</i></p> | ✓ | |
| <p>Uniform leader election in an n-node wireless network with collision detection is possible in $O(\log \log n)$ time w.h.p. <i>Reason: If it finished w.h.p, then it must take $\log n$ time slots.</i></p> | | ✓ |
| <p>A minimum dominating set can be larger than a maximal independent set. <i>Reason: Every maximal independent set is a dominating set so a minimum dominating set cannot be larger.</i></p> | | ✓ |
| <p>If a minimum dominating set is not an independent set, then nodes can be removed from it to get a maximal independent set. <i>Reason: From above, size of minimum dominating set is at most the size of minimum maximal independent set, say x. Assume the statement is true, then we get a maximal independent set of size less than x, a contradiction.</i></p> | | ✓ |
| <p>For the function $\text{ONES} : \{0, 1\}^k \times \{0, 1\}^k \rightarrow \{0, 1\}$ that evaluates to 1 if the two input k-bit strings contain the same number of ones and to 0 otherwise, it holds that $CC(\text{ONES}) \in \Omega(\log k)$. <i>Reason: This function ONES is essentially EQ on the number of ones which is essentially a $(\log k)$-bit number.</i></p> | ✓ | |
| <p>If for a graph G it holds that $m \in \Omega(n^2)$ (where m is the number of edges and n the number of nodes of G), then the diameter of G can be computed in $O(\sqrt{n})$ time. <i>Reason: The lower bound example from the lecture still holds if $m \in \Omega(n^2)$. Also, there are graphs with $m \in \Omega(n^2)$ and diameter $\Omega(n)$.</i></p> | | ✓ |
| <p>Given a weighted graph and its MST. If we add a single edge to the graph, a distributed algorithm can always update the MST in constant time. <i>Reason: We need to find the blue edge for either one of the two new fragments. This will require messages to propagate through at least one of the two fragments which is not possible in constant time.</i></p> | | ✓ |
| <p>Given a weighted graph and its MST. If we add a node and connect it with a single edge to any node previously in the graph, a distributed algorithm can always update the spanning tree in constant time. <i>Reason: If there is only one edge connecting a node, it has to be part of the spanning tree.</i></p> | ✓ | |
| <p>In any anonymous graph in which we cannot deterministically find a leader, we also cannot deterministically construct a spanning tree. <i>Reason: Counter example: linked list with 4 nodes.</i></p> | | ✓ |
| <p>Given an undirected graph with n nodes, if we need to query whether there exists a path between two nodes using a labeling scheme, the lower bound of the label size is $\Omega((\log n)^2)$. <i>Reason: $\log n$ is sufficient. Just label a node by the id of the connected component it belongs to.</i></p> | | ✓ |
| <p>Consider three graphs $G_1 = (V, E_1), G_2 = (V, E_2)$ and $G = (V, E)$ on the same node set, where E_1 and E_2 are two edge sets, and E is the union of E_1 and E_2. If there are adjacency labeling schemes of size k_i for G_i (for $i = 1, 2$), then there exists an adjacency labeling scheme of size $k_1 + k_2$ for G. <i>Reason: Let l_i be the labeling of graph G_i ($i = 1, 2$). Combine these two labelings to label G.</i></p> | ✓ | |
| <p>Synchronizer α may add a substantial overhead to the message complexity. <i>Reason: $O(m)$</i></p> | ✓ | |
| <p>There exists some synchronous distributed algorithm that cannot be executed in an asynchronous environment without changing the asymptotic time complexity of the algorithm. <i>Reason: Just use synchronizer α.</i></p> | | ✓ |

| Statement | true | false |
|--|------|-------|
| If a graph with n nodes can be colored with $c < n$ colors, then it can also be colored with $c + 1$ colors. <i>Reason: coloring definition</i> | ✓ | |
| When using Peterson's Algorithm deadlocks can occur. <i>Reason: see Theorem 4.6</i> | | ✓ |
| Compare-and-Swap is considered more powerful than Test-and-Set. <i>Reason: consensus number ∞ vs 2</i> | ✓ | |
| When using the Arrow algorithm: If there is a quiescent moment, then all the arrows point towards the node holding the variable. <i>Reason: see remarks above Theorem 6.4</i> | ✓ | |
| In the Arrow algorithm, the chosen spanning tree does not influence the runtime. <i>Reason: counter example: line vs star</i> | | ✓ |
| An advantage of the Ivy algorithm over the Arrow algorithm is that it also works in an asynchronous setting. <i>Reason: both do</i> | | ✓ |
| The self-stabilizing MIS (Algorithm 12.5) may fail to compute a locally stable solution if the adversary is allowed to modify the network topology. <i>Reason: if the adversary does not change the k-neighborhood within k time steps, the self-stabilizing MIS converges to a locally stable solution</i> | | ✓ |
| In the Democrats and Republicans problem, if initially the majority of citizens votes for Democrats, there will eventually be no citizen who switches her opinion every other day. <i>Reason: counter example: a star graph, where only the central node votes for republicans</i> | | ✓ |
| If a distributed problem on an n -node graph can be solved in $O(n)$ time when the message size is restricted to $O(\log n)$ bits, then it can also be solved in $O(\log n)$ time when the message size is restricted to $O(n)$ bits. <i>Reason: Take for instance the problem of leader election on a path.</i> | | ✓ |
| There is a bipartite graph in which deterministic anonymous leader election is possible. <i>Reason: A three-node path works.</i> | ✓ | |

2 Low Cost Sorting

(26 Points)

Consider a network of four nodes: v_1 , v_2 , v_3 and v_4 . Each node is given an input number. An algorithm on this network executes in synchronous rounds. In each round, each node can participate in a compare-and-exchange operation with at most one neighbor. See Figure 1 for an example.




Figure 1: An example network where all pairs of nodes have edges between them except v_2 and v_3 . After compare-and-exchange between a pair of nodes (represented by an arrow), the smaller value moves to the node at the head of the arrow, the larger to the node at the tail. In the first round, compare-and-exchange occurs between v_1 and v_3 , the smaller value moving to node v_1 . In the second round, compare-and-exchange occurs between v_2 and v_1 , and between v_3 and v_4 .

The cost of running an algorithm for t rounds on a network of m edges is $m \cdot t^2$. For questions **A)** and **B)**, give exact instead of asymptotic bounds.

- A)** Consider the problem of finding the minimum of the input numbers in a network with four nodes v_1 , v_2 , v_3 and v_4 .
- 1) [2] Give a lower bound on the number of *edges* any network needs to put the minimum number at v_1 .
 - 2) [4] Give a lower bound on the number of *rounds* any algorithm needs to put the minimum number at v_1 .
 - 3) [3] Give a network and an algorithm for placing the minimum number at v_1 so that the cost $m \cdot t^2$ is minimum.
- B)** Now consider the problem of sorting the input numbers in a network with four nodes v_1 , v_2 , v_3 and v_4 so that node v_i gets the i^{th} smallest number.
- 1) [7] Give a lower bound on the number of *rounds* any algorithm needs to sort the numbers.
 - 2) [5] Give a network and an algorithm to sort the numbers minimizing the cost $m \cdot t^2$ (better cost gives more points).
- C)** [5] Now you are given n nodes. Describe how to design a network and a sorting algorithm minimizing the cost $m \cdot t^2$. Please use the big-O notation for this part (better asymptotic cost gives more points).

Solutions

- A) 1) The graph must be connected otherwise v_1 and the minimum number can be in different components. Thus, 3 edges are needed.
- 2) In one round, max two compare-and-exchange can happen. The minimum could be in a node other than v_1 and the node involved in compare and exchange with v_1 . Thus, another round is needed to move the minimum to v_1 .
- 3)




Figure 2: Algorithm with $m = 3$, $t = 2$ and cost = 12: After Round 1 (left), minimum is either in v_1 or v_2 . After Round 2 (right) minimum is in v_1 .

- B) 1) In each round, there are at most 2 swaps possible. Thus, there are 2^4 swaps possible in 2 rounds. The algorithm must be able to sort each of the $4!$ inputs. As $16 < 24 = 4!$, one more round than two is needed.
- 2)




Figure 3: Algorithm with $m = 4$, $t = 3$ and cost = 36: After Round 1 (left), minimum is either in v_1 or v_2 and maximum is either in v_3 or v_4 . After Round 2 (middle) minimum is in v_1 and maximum in v_4 . After Round 3 (right), the second largest and third largest input is at v_2 and v_3 respectively.

- C) One can simulate batcher's sorting network in this model where each subsequent level defines the compare-and-exchange to be done in each subsequent round. This is valid as at any given level a node is incident to at most one comparator.

As there are $O(\log^2 n)$ levels, it takes $O(\log^2 n)$ rounds. As there can be only $n/2$ edges per level in the sorting network, we need $O(n \log^2 n)$ edges in total. Thus, the total cost is $O(n \log^6 n)$. (Using AKS sorting network is even better but batcher's is sufficient for full points.)

3 1D Land Postal System

(35 Points)

1D Land occupies a thin strip of land in the middle of the ocean. Recently, its population decided to implement a postal system. To do so the land is divided into n segments of equal length. For a packet to travel from one segment to a neighboring segment requires one time unit. Travel time within a segment is negligible.

To compete with modern technology such as email and delivery drones, a cannon is installed in each segment, which can safely and rapidly transport packets k segments in one direction within a single time unit. The cannon installations are fixed, i.e., cannons cannot turn or adjust their angle or power – each cannon always delivers to the same segment in distance k .

In the following we are interested in the MPTT (maximum packet travel time) which is the maximum over the shortest delivery times for every possible packet source and destination. All results and intermediate steps may be specified asymptotically using big O notation.

- A) [5] Specify a k and cannon directions achieving a low MPTT. Lower MPTTs give more points.

Now suppose the cannon directions are chosen uniformly at random between *left* and *right*.

- B) [3] For your algorithm (for your k): What are the worst possible cannon directions, and what is the MPTT in this scenario?

We are now interested in the PTTLWHP (packet travel time limit with high probability): a time limit in which a packet will reach its destination w.h.p., when sampling over the random cannon directions and packet source and destination.

- C) [10] Show that the PTTLWHP is in $O(\frac{n}{k} \log n + k)$ for any k .

- D) [10] Show a lower bound of $\Omega(\sqrt{n})$ for the PTTLWHP, for any k .

Hint: You may assume that if some predicate A does not hold in expectation, A also does not hold w.h.p..

You may now choose both the direction and the power of each cannon freely, i.e., you may choose any one target segment for every cannon.

- E) [7] Specify a cannon configuration reaching a low MPTT (for any n). Lower MPTTs give more points.

| |
|------------------|
| Solutions |
|------------------|

- A) $k = \sqrt{n}$. Alternatingly point the cannons left and right. The maximum travel time is $O(\sqrt{n})$: at most $2\sqrt{n}$ hops to reach the \sqrt{n} -neighborhood of the target segment, plus at most \sqrt{n} hops within that neighborhood.
- B) k doesn't matter. Worst case travel time is $n - 1$. Example: all cannons point to the left, but our packet starts at the left and needs to go to the right.
- C) The bound consists of two terms:
- $O(\frac{n}{k} \log n)$ represents the number of steps to reach the k -neighborhood. We need to take at most $\frac{n}{k}$ cannons facing in the correct direction. To find a cannon facing in the correct direction takes at most $\log n$ steps "on foot" w.h.p. (proof below).
 - $O(k)$ represents the number of steps within the k -neighborhood to the final destination.

Proof:

Starting anywhere, let $X = \sum_{i=1}^{\log n} X_i$ denote the number of cannons pointing the correct direction found within the next $\log n$ segments, and let X_i be 1 if there is a cannon pointing in the correct direction at segment i , and 0 otherwise. Note that $\mathbb{E}[X_i = 1] = \frac{1}{2}$ and $\mathbb{E}[X] = \frac{\log n}{2}$. Using the Chernoff bound with $\delta = \frac{1}{2}$ we obtain:

$$\begin{aligned} Pr[X \leq (1 - \delta)\mathbb{E}[X]] &\leq e^{-\mathbb{E}[X]\delta^2/2} \\ Pr[X \leq \frac{\log n}{4}] &\leq e^{-\log n/16} = \frac{1}{n^c} \text{ (for some constant } c) \\ \Rightarrow Pr[X > \frac{\log n}{4}] &\geq 1 - \frac{1}{n^c} \end{aligned}$$

Hence the number of cannons in the correct direction within $\log n$ segments is at least $\frac{\log n}{4}$ w.h.p..

Alternative shorter proof:

Let $p = \frac{1}{2}$ be the probability that the next cannon points in the wrong direction. Then, when inspecting the next $c \log_2 n$ cannons, for some constant c , the probability that not all of them point in the wrong direction is $1 - p^{c \log_2 n} = 1 - \frac{1}{n^c}$.

- D) We show the lower bound $\Omega(\sqrt{n})$ for any k for the expected value. From the lower bound for the expected value, the lower bound for any k w.h.p. follows.

Randomly sampled start-destination pairs have an expected distance of $\Theta(n)$. Hence, no matter the k and no matter the cannon configurations, the expected number of steps to reach the goal is always at least $\Omega(\frac{n}{k} + k)$: $\Omega(\frac{n}{k})$ cannon shots to reach the k -neighborhood of the goal, plus k steps within that neighborhood.

The term $\frac{n}{k} + k$ becomes minimal for $k = \sqrt{n}$, in which case it evaluates to $\frac{n}{\sqrt{n}} + \sqrt{n} \in \Omega(\sqrt{n})$, the sought lower bound.

- E) For a segment interval (starting with the whole island): make sure there are cannons at each end which shoot to the middle of the segment, and there are two cannons in the middle that shoot to both ends. Then repeat this for both the left and the right half. It helps to always pair left and right cannons of equal strength.

The result is a kind of balanced binary tree with its root in the middle of the island. The MPTT is in $O(\log n)$.

4 Galactic Neighborhood

(27 Points)


In the lecture, the concept of the neighborhood graph was used to prove a lower bound for coloring certain classes of graphs. We will look at another class of graphs here: stars. A star consists of a central node and a number of leaves (can be 0, 1, or more), each of which is connected with the central node, but no other nodes. Let \mathcal{G}_k be the set of labeled stars with *up to and including* k nodes. We will look at the case of one round of communication where every node will learn exactly the labels of its immediate neighbors and nothing else, in particular not its neighbors' degrees.

- A) [5] Draw the 1-neighborhood graph $\mathcal{N}_1(\mathcal{G}_3)$ for labeled stars with up to 3 nodes where each star can use labels from $\{1, 2, 3\}$.
- B) [5] Give a deterministic distributed algorithm that colors a star with k nodes where $k \leq 3$ with labels from $\{1, 2, 3\}$ with two colors in 1 round, or prove that such an algorithm does not exist.


Now we restrict the labels that can occur in a star with k nodes to $\{1, \dots, k\}$, i.e. every label from 1 to k occurs exactly once.

- C) [5] Draw the 1-neighborhood graph $\mathcal{N}_1(\mathcal{G}_3)$ for labeled stars with up to 3 nodes where each star with k nodes uses exactly the labels $\{1, \dots, k\}$.
- D) [5] Give a deterministic distributed algorithm that colors a star with $1 \leq k \leq 3$ nodes with labels from $\{1, \dots, k\}$ with two colors in 1 round, or prove that such an algorithm does not exist.
- E) [7] Prove or disprove that the result of D) generalizes to stars with k nodes and labels $\{1, \dots, k\}$ for arbitrarily large k .

Solutions



B) The neighborhood graph has an odd-length cycle and thus has no 2-coloring. According to Lemma 8.6, no deterministic 1-round 2-coloring algorithm exists.



Instead of drawing this, you could also mark in your solution for A) which edges and nodes are missing here.

D) The simplest option is to give a 2-coloring of the solution to part C), and let the algorithm hard-code this coloring.

E) A node that sees exactly one neighbor j uses the color $j \bmod 2$, and a node i that sees no neighbor or at least two neighbors uses the color $i + 1 \bmod 2$. **Proof:**

- If the star consists of 1 node, that node will be colored $1 + 1 \bmod 2 = 0$.
- If the star consists of 2 nodes, node 1 will be colored $2 \bmod 2 = 0$ and node 2 will be colored $1 \bmod 2 = 1$.
- If the star consists of at least 3 nodes, assume the center has label j . All leaves see exactly one neighbor, namely the center j , and will get color $j \bmod 2$. The center will see at least two neighbors and will thus get color $j + 1 \bmod 2$.

5 Reading Assignment

(9 Points)

- A) [3] Show that there is an algorithm for online graph exploration that is 1-competitive on weighted trees.
- B) [3] Does the hDFS algorithm have a constant competitive ratio on weighted paths?
- C) [3] Is there an algorithm for online graph exploration that has a competitive ratio of strictly less than 16 on planar graphs with unit-weight edges?

Solutions

- A) The optimal offline algorithm has to traverse each edge twice since the graph is a tree. Since DFS traverses each edge exactly twice, it is 1-competitive.
- B) No: The lower bound in Theorem 5 holds also for paths since the used lower bound graph is a path and it holds also for hDFS without rounded weights since the used lower bound graph has only powers of 2 as weights.
- C) Yes: hDFS (or DFS) is 2-competitive (see Theorem 4 for hDFS or introduction of reading assignment for DFS).