Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

FS 2017

Prof. Dr. Mohsen Ghaffari,
Prof. Dr. Roger Wattenhofer

Exam
Principles of Distributed Computing

Monday, August 14, 2017
14:00 – 16:00

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 (prove) 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.

Exercise	Achieved Points	Maximal Points
1 - Multiple Choice		12
2 - Arrow and Ivy		12
3 - Sorting Networks		14
4 - Very Small Radius		26
5 - Global Computations		16
6 - Special Coloring		20
7 - Combinatorics of Radio Transmissions		20
Total		**120**

1 Multiple Choice (12 points)

Evaluate each of the following statements in terms of correctness. Indicate whether a statement is true or false by ticking the corresponding box. Each correct answer gives one point. Each wrong answer and each unanswered question gives 0 points.

Statement	**true**	**false**
Any rooted binary tree with n nodes can be 1-colored in time $O(n)$.	□	□
Any rooted binary tree with n nodes can be 2-colored in time $O(\log n)$.	□	□
Any rooted binary tree with n nodes can be 3-colored in time $O(\log^* n)$.	□	□
Any rooted binary tree with n nodes can be 4-colored in time $O(1)$.	□	□
The flooding algorithm can be used to determine if a graph is a tree.	□	□
The Gallager-Humblet-Spira algorithm can be modified to compute a spanning tree of maximum weight by defining a blue edge to be the maximum weight outgoing edge.	□	□
In Luby's Maximal Independent Set (MIS) algorithm, in each round each node is removed with probability at least 1/10.	□	□
Any planar graph with n nodes can be 7-colored in time $O(\log^* n)$.	□	□
In any graph with maximum degree at most 5, any maximal independent set has size at least 1/5 of the maximum independent set.	□	□
Suppose that Alice knows 5 message each with 100-bits, Bob knows 4 of these messages, and Alice doesn't know which ones are known to Bob. Alice needs to send at least 200 bits to Bob so that he also knows all the messages.	□	□
Any labeling scheme for distance in cycles needs to use labels of size at least $\Omega(\log^2 n)$.	□	□
There exists at least one task on unweighted simple graphs such that any labeling scheme needs to use labels of size at least $\Omega(n^3)$.	□	□

2 Arrow and Ivy (12 points)

In this task we want to compare the two main algorithms for shared objects, Arrow and Ivy. Figure 1 depicts a precomputed spanning tree of a complete graph G. The nodes v_1, v_2 and v_3 want to access the object which is stored at the node w. Assume that the requests are filed sequentially, i.e. a node files a request after the previous node has already received the object.

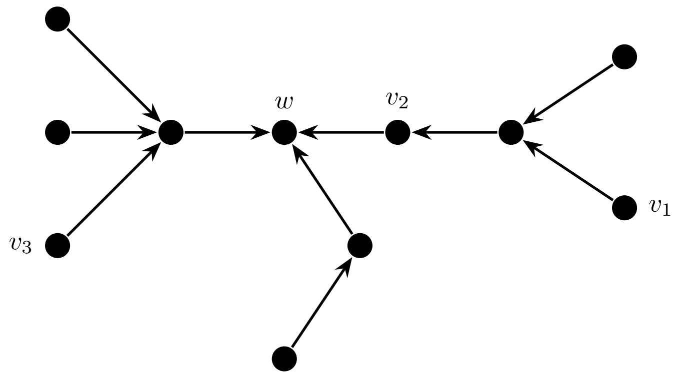

Figure 1: A precomputed spanning tree of G

A) [4] Give the worst-case ordering of the requests v_1, v_2, v_3 if you apply Arrow.

B) [4] Give the best-case ordering of the requests if you apply Ivy.

C) [4] Your friend claims that on average the Arrow algorithm is strictly better than the Ivy algorithm for G with this precomputed spanning tree and any sequential request sequence. Is your friend correct?

3 Sorting Networks (14 points)

A) [4] Consider a correct sorting network. Show that for each pair of adjacent wires i, $i+1$, the network contains a comparator that compares wires i and $i+1$.

B) [10] For each integer $n \geq 2$ determine whether there exists a correct sorting network of width n that contains exactly one comparator for each pair of adjacent wires. (There is no restriction on the number of comparators for non-adjacent wires.) If the answer is yes, provide a construction of a sorting network as described. If the answer is no, show that a sorting network as described cannot exist.

4 Very Small Radius (26 points)

In the lecture, we studied the diameter of a graph. A very related notion is the radius of a graph. Throughout this exam question, we assume that all considered graphs are connected.

The radius $r(u)$ of a node u in a graph G is the maximum of the distances from u to all other nodes, i.e., $r(u) = \max_{v \in G}(\text{dist}(u, v))$. The radius $r(G)$ of a graph G is the minimum of the radii of all nodes in G, i.e., $r(G) = \min_{u \in G}(r(u))$.

We want to study the problem of determining whether a graph has radius 1 or not.

A) [4] Radoslav claims that if the radius of a graph is 1, then its diameter is always 2. Diana claims that if the diameter of a graph is 2, then its radius is always 1. Who is right, who is wrong?

B) [22] Assume that the nodes do not know n, i.e. the total number of nodes in the graph. Design a synchronous deterministic distributed algorithm that determines if the input graph has radius 1 or not. Each node has a unique $O(\log n)$-bit identifier. In each round each node can send a message that has $O(\log n)$ bits over each incident edge.

You can obtain up to 5 points if your algorithm has the following properties:

1) If the radius is 1, then all nodes output "YES".

2) If the radius is not 1, then at least one node outputs "NO".

You can obtain up to 5 **additional** points if the runtime of your algorithm is **constant**.

You can obtain up to 12 **additional** points if the runtime of your algorithm is **constant** and the algorithm has the following properties:

1) If the radius is 1, then all nodes have to output "YES".

2) If the radius is not 1, then all nodes have to output "NO".

5 Global Computations (16 points)

Model: The network is abstracted as an n-node undirected graph $G = (V, E)$, with diameter D. Each node has a unique $\Theta(\log n)$-bit ID. Initially, each node only knows its own edges, as well as the value of n. We consider the CONGEST model where per round each node can send $O(\log n)$ bits to each of its neighbors.

Problem: Suppose that each node $v \in V$ receives an input value $x_v \in \{1, 2, \ldots, n^{10}\}$. The objective is for all nodes to learn the average of the values of the $k = \lceil\sqrt{n^{0.4}}\rceil$ largest inputs. Suppose that the diameter is $D = \lceil\sqrt{n^{0.6}}\rceil$. Devise a fast algorithm for the problem. Provide a correctness argument for your algorithm and analyze its round complexity.

A correct algorithm receives **6 points**. Proving the algorithm's correctness receives another **5 points**, and obtaining a good round complexity (with analysis) receives the last **5 points**.

6 Special Coloring (20 points)

Model: The network is abstracted as an n-node undirected graph $G = (V, E)$. Initially, each node only knows its own edges, as well as the values of n and the maximum degree Δ. We consider the CONGEST model where per round each node can send $O(\log n)$ bits to each of its neighbors.

Problem: Each node v must choose a color value $c_v \in \{\Delta^{3/2}, \Delta^{3/2} + 1, \ldots, 2\Delta^{3/2}\}$. Each edge $e = \{v, u\}$ is annotated with a *conflict description* (a_e, b_e) for some values $a_e \in \{10\Delta, 10\Delta + 1, \ldots, 20\Delta\}$ and $b_e \in \{0, 1, \ldots, a_e - 1\}$; we say that the coloring is invalid and has a conflict on edge e if we have $c_v \equiv b_e + c_u \pmod{a_e}$. The conflict description values (a_e, b_e) are known to both endpoints v and u of the edge e. Notice that the coloring is invalid if there is a conflict on any edge.

Devise a fast randomized distributed algorithm that computes a valid coloring, with high probability. The algorithm should terminate within your claimed round complexity and the computed coloring must be valid with probability at least $1-1/n$. Analyze your algorithm's round complexity and prove its correctness.

A correct algorithm receives **8 points**. Proving the algorithm's correctness receives another **4 points**, and obtaining a good round complexity (with analysis) receives the last **8 points**.

7 Combinatorics of Radio Transmissions (20 points)

Model: We assume the *radio networks* model, which works as follows: The network is abstracted as an n-node undirected graph $G = (V, E)$. Each node has a unique $\Theta(\log n)$-bit ID. Initially, each node knows only the values of n and Δ, and its own identifier. Per round, each node either transmits a message or listens. A node receives a message only if it is listening and exactly one of its neighbors is transmitting a message, and in that case, the node receives the message of that single transmitting neighbor. If a node is not listening, or two or more of its neighbors are transmitting, then the node does not receive anything.

Problem: Suppose that each node v has a message m_v, which should be delivered to all of its neighbors. Devise a short and deterministic schedule for transmissions so that at the end, each node u receives the messages of all of its neighbors. The schedule describes for each node v what it should do in each round — i.e., whether it should transmit its message m_v or it should listen — as a function of n, Δ, and the node's identifier ID_v. The schedule must be independent of G and should work for any n-node graph.You should also argue about the schedule's correctness, and analyze its length as a function of n and Δ.

A correct algorithm receives **8 points**. Arguing about the algorithm's correctness receives another **4 points**, and obtaining a good schedule length receives the last **8 points**.

Hint: Think about *cover-free families*, from lecture 5.