1 Objectives

Many practical problems require more than one entity performing some work or some computation. For example search engines refresh the World Wide Web index by repeatedly crawing the network of websites in parallel with multiple software agents. Another example is a network of sensors monitoring some terrain. Using only a single entity would not be enough because of the time constrant (we want to refresh the index every day) or because of some physical limitations (a single sensor cannot cover the whole terrain). Using multiple entities solves these problems but creates new. How to organize the work in order to avoid every entity working on the same task? How to disseminate information between the entities? How to design the entities and the algorithm for them in order to keep the entities simple? Answering these questions is the goal of distributed computing. In this project we will focus on three models of distributed computing that are inspired by practical applications.

In the first model we will study agents trying to explore unknown networks. The agent can visit the nodes and traverse the edges with the goal of visiting the whole network. It the agent had large memory (proportional to the size of the whole graph) then the problem could be solved using classical algorithms like BFS. But clearly nowadays many interesting and important networks cannot fit in the memory of a single machine so we need to come up with a new approach. The idea is to use smaller memory at the agent accompanied by either use of randomness or some number of bits at every node of the network.

In the model of radio networks, each node is a station capable of transmitting and receiving messages. However when multiple nodes are transmitting to the same receiver, their signals collide and it becomes impossible for the receiver to decode either of the transmissions. This is obviously a simplified model of a communication network but designing good algorithms in such a model can lead to useful practical solutions. We would like to focus on solutions that use small energy. Assuming that each node is a battery-powered device and each transmission costs some energy we would like to preserve the energy (and increase the lifespan of the device) by minimizing the number of transmissions. Our second objective will be to design algorithm that are robust to some changes in the topology of network. It is obvious that in practice many communication networks are not static hence it is crucial to have algorithms that can work in the dynamic setting.

Our third model are the insect colonies. We want to find out how is it possible that some social insects (like ants or bees) achieve effective collaboration in their work on some tasks of the colony despite lack of any central control and limited communication. Both ants and bees are simple animals hence their algorithm also cannot be too complicated. Our main focus will be task allocation – how to split the workers between the tasks so that each task is being handled well. We want to find out how much memory is necessary for efficient task allocation. Secondly we want to find a task allocation algorithm that will work even if the demands of the tasks will change.

2 Methodology

Over the course of this project we will propose algorithms for mobile agents, radio networks and ant colonies in different scenarios. We will analyze the proposed solutions using formal mathematical proofs (usually preceeded by simulations). Such analysis will require different tools depending on the scenario. In the randomized case we will study biased random walks, bounded memory computations, Markov Chains. In the deterministic case we will use combinatorics, graph theory.

3 Impact of the research project

The problems of exploration, broadcast and task allocation are well-studied problems in their respective domains. However we will study approaches and generalisations of the original problem that have not received enough attention. The main outcome of the project will be new, efficient algorithms for graph exploration, broadcast and task allocation. We will propose practical solutions (fast, using small memory and robust to changes). We want to design at least two new algorithms for each of the three problems. Our goal will be to design simple algorithms since we believe that such algorithms are naturally robust and are better suited to be applied in practice.