First Year Project Report

University of Liverpool

Thomas Gorry

1st Supervisor: Russell Martin
2nd Supervisor: Leszek Gąsieniec

Advisors: Martin Gairing, Igor Potapov and Prudence Wong

May 21, 2012

Contents

1 Introduction

My chosen topic of study is in the field of Algorithms with a focus on Distributed Computing. This Ph.D project aims to carry out research on performance analysis of networks populated by mobile agents. During the initial stages the emphasis has been on agent location discovery, network exploration and search and rendezvous problems. We have already achieved some results for the location discovery of agents and also on a type of group search problem. Our future work will be to extend this research to other network topologies with applications to more complex network communication and search and rendezvous problems. In particular we intend to look at two types of network environments including graph based and geometric settings. It should also be mentioned that there has been some work we have carried out on Social Network Analysis and there may be some area of research that we may be able to link in with the current plan for the Ph.D.

2 Ph.D aims

The area that this research is concerned with is that of search and rendezvous and location discovery problems with strong links to network patrolling problems. The aim of this Ph.D is to carry out research into provable theoretic bounds for these types of problems in different settings and models.

There are a variety of interesting models that can be considered when looking at these problems. These variations can be categorised into two main groups. They are as follows:

2.1 Variations to the model regarding the mobile agents

- Different maximum speeds for the mobile agents.
- Each mobile entity has unique Ids.
- The memory that agents have is limited.
- Agents are able to leave markers on the graph.
- Markers agents leave may be more complex and actually be more like messages or instructions left behind.
- Communication may be in the form of signals agents can send. This may be at certain intervals or with limited range.

2.2 Variations to the model regarding the environment

- Topological variations to the network graphs studied.
- Different geometric environments.

- Vertices in graphs may have Ids.

These are just a few examples and there are an endless number of different settings that can be looked at within this area of research. Furthermore, the algorithms that we work on will cover both deterministic and randomised approaches to these types of problems.

2.3 Network monitoring

Monitoring is the problem where, in a graph or geometric environment (e.g. a simple polygon), agents are arranged in stationary positions to constantly survey the graph/region, given that each agent has a limited field of vision. One model is where there is an associated cost of deployment for each stationary agent within the network and so, much like the Art Gallery Problem [1], maximising visual range with minimum agents would be the goal.

In a distributed setting, we want to design a procedure where the agents self-organise into a suitable configuration for efficient monitoring. A solution to this problem would involve some initial exploration of the environment, followed by distribution of the agents into appropriate locations for the monitoring to occur. In recent work efficient organising in this way has been looked at by [7] in relation to vehicles that can communicate with each other to position themselves effectively on road systems to minimise congestion with reduced impact on travel time. Also [2] has done work on neighbour discovery in a sensor network with directional antennae, there may be some work around this problem of efficient network surveillance that we can incorporate into our own research.

2.4 Network patrolling

In some circumstances, there may not be enough agents to constantly monitor the environment and so the agents would have to realise this and so then must adapt some type of patrolling strategy instead. Network patrolling will be the main aspect of our research and will cover questions such as, in a variety of different models, how can we minimize the time between visits to nodes in the network. This becomes more interesting when you consider models that may specify that some regions in the network are more important than others and so therefore agents in the network may choose to leave these areas unattended for longer periods in order to ensure maximum frequency of visits to the more important areas of the network. It is worth noting that the problem in a general graph is NP-Hard as if you have only one agent then the problem becomes one of finding a Hamiltonian cycle in the graph if all the edges have the same length. Therefore, in future work we will be considering specialised classes of graphs and possibly also in a geometric setting.

For each model we study it would also be logical to consider what the most efficient patrolling strategies are that apply to the topology and the model being investigated are and so our research will also cover these questions.

2.5 Social network analysis

As mentioned earlier there has been some collaboration with industrial partners, who operate in the media and marketing sector, on the analysis of social networks. Our partnership with this company has thrived and so there may be opportunity to blend our current research with research into this area. For example, we could look to research into practical applications for our findings in the area of Social Networks. Their ambition is to analyse Social Networks with the goal of finding nodes with "topological significance ". While we have been investigating the best approaches to this area we have discovered there are ways we can adopt what we learn here to fit nicely with the research we have already carried out. For example, as Social Networks are inherently dynamic and, like any other network, have nodes that are more significant in certain ways than others then having an equally dynamic set of agents that could patrol these networks while putting more emphasis on the frequency of visits to these significant nodes would be an interesting problem to look at. This is also useful in the real world in terms of aiming at preventing or quickly stopping adversaries that would see entities such as worms proliferate across these networks stealing users′ personal data and passwords.

We believe that this industrial work can be a source of good research problems and that exploration of these types of models will allow us to come up with efficient algorithms for the monitoring and patrolling of these continuously shifting types of networks.

3 Background research

Initially I began by sharpening my probability skills by reading Probability and Computing: Randomized Algorithms and Probabilistic Analysis [**?**]. This gave me a good foundation to begin research into the randomised algorithms we would be using in some of our work.

We studied the first algorithmic paper on boundary patrolling that appeared in ESA'11 proceedings [**?**]. They analysed performance of teams of agents with different maximum speeds on the segment and circle. Their solution is fully centralised and their results refer only to a small, constant number of agents. The difficulty in dealing with a larger number of agents is a direct consequence of different maximum speeds. In our approach, Location discovery on the ring and the segment, we assume uniform maximum speeds and this allows us not only to deal with arbitrary number of agents but also allows us to design a fully distributed solution in which agents do not exchange messages.

As mentioned there are many different settings that we can look at regarding any of the problems we propose to study. Once such setting is one where the agents have limited memory, as in [5]. The model they adopt in this paper is a tree graph structure of unknown size. The idea is that with the limited amount of memory the agent must explore and learn the topological structure of the graph. However, as the graph is of unknown size it is possible that it could be too large for the agent to store as memory is limited. In the event this is the case then the agent simply admits that the graph is too large to map and returns to the start.

Another avenue of our recent research explorations refers to algorithmic and visual analysis of Social Networks. This work includes study of basic analysis methods as well as development of software in collaboration with a local industrial partner. There are numerous ways that this could be linked to the agent patrolling research mentioned earlier. One such way is mentioned by [?], they propose that through investigation of significant nodes within networks by a variety of different measures, one is able to predict where an intruder may wish to penetrate a network. Using this knowledge one of the things to work towards is a method whereby the agents in the network can determine these significant nodes in a distributed fashion and then visit these nodes more frequently than others while patrolling as they would be more likely the target of intruders to the network.

- Evangelous (Search Games) - caterina (Patrolling) - S Gal (Search Games)

4 Results obtained so far

4.1 Location discovery and patrolling on the ring

We have already made a significant start into this area with a paper we have produced that deals with agent patrolling on the ring. This work is a natural continuation of my MSc dissertation [4]. The model contains n uniform anonymous mobile agents located on a ring with circumference 1. The agents perform their actions in synchronised rounds. At the beginning of each round an agent chooses the direction of its movement from *clockwise* and *anticlockwise*, as well as its speed $0 \leq v \leq 1$ during this round. We assume that the agents can not pass over one another on the circle and when an agent collides with another it instantly starts moving with the same speed towards the opposite direction. The agents are unable to leave marks on the ring, they have zero visibility and they cannot exchange messages. However, on the conclusion of each round each agent has access to information about it's current position relative to where it started. This information can be processed and stored by the agent for further analysis.

The agents are initially located at arbitrary, but distinct, positions unknown to other agents. The main *location discovery task* to be performed by each agent is to determine the initial position of every other agent and eventually to stop at its initial position, or proceed to another task, in a fully synchronised manner.

Our main result constitutes of a fully distributed randomised algorithm that solves the *location discovery problem* with high probability in $O(n \log^2 n)$ rounds. We also show how this mechanism can be adopted to distribute the agents equidistantly, and how to coordinate their joint effort in patrolling the circle. Note that our result also holds if initially the agents do not know the value of n and they have no coherent sense of direction.

4.2 Location discovery on the segment

Following on from our initial work with the ring we looked at patrolling on the segment as this would be a nice further addition to our paper.

In the model that we adopt this time there are n agents situated on a segment of length 0.5. At either end of the segment there are immovable objects that we will refer to as walls. Each agent moves at speed 1 and does not have the ability to communicate or overpass. Instead, when agents collide with either another agent or one of the walls they immediately begin travelling in the opposite direction. Each agent has the ability to measure time and distance and it is assumed that they know n. However, there is no common sense of direction. As in [3] the notion of a round is defined by one unit of time. At the start of each round the agent chooses its initial direction at random with equal probability.

Using the algorithm we propose the agents first discover the positions of both the walls and/or their neighbours, depending on the agents starting position on the segment. Following on from this agents are then able to learn their ordering on the segment in relation to the walls. Once this is done our procedure allows them to learn the length of the segment and with this information the agents should be able to space themselves as with the ring example above and then patrol along the segment efficiently. We believe that the location and segment length discovery processes can be completed in $O(2n + \frac{1}{2}n \log n)$ rounds.

4.3 Group search on the line (Evacuation Problem)

Here we considered a *group search problem* in which k mobile agents located on the line perform a search for a specific destination. The agents are initially placed at the same point (origin) on the line and the target is located at unknown distance d either to the left or to the right from the origin. The problem with $k = 1$ is known as the *cow-path* problem, and the complexity of this problem is known to be $9d$ in the worst case, it is also known that this is the case for $k \geq 1$ agents. In this work we presented a clear argument for this claim by showing a rather counter-intuitive result. This complexity also applies to any metric. Namely, in any metric independently from the number of agents group search cannot be performed faster than in time $9d$, where d is the distance between the origin and the destination. We also look at the case when there are $k = 2$ mobile agents with different speeds. It is assumed that the search is accomplished when all agents eventually arrive at the destination.

From our initial work on this problem we have discovered that an algorithm exists that allows two agents of different speeds, one with a speed of 1 and the other with a speed of at least $\frac{1}{3}$, to still complete this search problem with a time of $9d$. This work has been accepted to appear in a publication that will be produced for Search and Rendezvous '12. We are now finalising this paper by investigating the lower bounds for this problem.

5 Research plan

As mentioned earlier there are many avenues that our work could take us down. However, we know that over the course of the next month we will be looking at proving the lower bounds for our work on group searching on the line. We have also identified the following areas that we will be looking at in the future for this project:

- With regards to the evacuation problem, we will be looking at what can be achieved when the slowest agent has a speed of $< \frac{1}{3}$.
- Staying with the evacuation problem we will then consider what happens if there are more that two agents with different maximum speeds.
- We may also adapt the evacuation problem to a model that moves the setting to multiple paths rather than a single line.
- With regards to our work on the location discovery problem on the circle and the segment collaboration has been revived from Durham University and we are in the initial stages of collaboration with researchers from Carleton University, The University of Quebec in Ottawa and also The University of Bordeaux. We expect these collaborations to continue and result in joint publications in this field.
- We will be looking at a variety of search games that are listed as open problems in the field.
- We will also be looking more into the areas of network patrolling than what we have previously done.
- As mentioned earlier we have been working with a local business on social network analysis techniques and so there may be some blend that can be achieved between our research into network patrolling and the type of social network analysis we are carrying out.

6 Events and Activities

- November 2011 - Attended and won the University of Lincoln Library Prize at DEVxS '11 development conference, Lincoln, UK.
- January 2012 - Research trip to Bordeaux University, France to discuss possible collaboration and present current research.
- April 2012 - Attended and presented a talk on our research at the BCTCS '12 workshop, Manchester, UK.
- May 2012 - Attended and presented a poster on our research at the Search and Rendezvous '12 workshop, Lorentz Center, Leiden, Netherlands.

References

[1] V. Chvátal, A combinatorial theorem in plane geometry, *Journal of Combinatorial Theory, Series B*, vol. 18, pp. 39–41, 1975.

[2] J. Du, E. Kranakis, O. Ponce and S. Rajsbaum, Neighbor Discovery in a Sensor Network with Directional Antennae, *Algorithms for Sensor Systems*, isbn. 978-3-642-28208-9, vol. 7111, pp. 57–71, 2012.

[3] T. Friedetzky, L. Gąsieniec, T. Gorry and R. Martin, Observe and Remain Silent (Communication-less Agent Locaiton Discovery), *Submitted to MFCS '12*, url. http://www.csc.liv.ac.uk/ tgorry/publications/T.%20Friedetzky,%20L.%20Gasieniec,%20T.%20Gorry,%20and%20R.%20Martin%20-%20Observe%20and%20Remain%20Silent.pdf, 2011.

[4] T. Gorry, Distributed Agents Location Discovery on the Ring, *MSc Dissertation, Liverpool University Computer Science*, url. http://www.csc.liv.ac.uk/ tgorry/publications/Thomas%20Gorry%20-%20MSc%20Dissertation.pdf , 2011.

[5] L. Gąsieniec, A. Pelc, T. Radzik and X. Zhang, Tree exploration with logarithmic memory, *Proc. 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 585-594, 2007.

[6] T. Temple and E. Frazzoli, Whittle-indexability of the Cow Path Problem, *American Control Conference (ACC)*, pp. 4152–4158, 2010.

[7] L. Wischhof, A. Ebner, H. Rohling, M. Lott and R. Halfmann, SOTIS - A Self-Organizing Traffic Information System, *VTC 2003-Spring: Proceedings of the 57th IEEE Vehicular Technology Conference*, pp. 2442–2446, 2003.