Random Search of a Large Area and Formation Flight in Complex Environments: SURF Report

J. C. Warrington, Cambridge University
SURF Student, Cambridge-Caltech Exchange

Mentor: Prof. Jerrold E. Marsden, CDS
Co-mentor: Sujit Nair, Postdoctoral Scholar, CDS

September 28, 2007

1 Abstract

This project had two parts. Firstly a study was carried out into the surveillance of a large rectangular area with a set of autonomous agents. The agents employed a random search technique with collision avoidance rules; it was shown that an inverse rule for coverage time vs. number of agents applies. Secondly a potential-based formation flight algorithm was implemented, primarily employing a triangular lattice formation. Triangular to square transformations, and the reverse, were observed. A 2D flock was made to navigate through a complex region populated with walls and obstacles while retaining triangular formation and avoiding collisions.

2 Introduction and Motivation

The study of Reactive Collision Avoidance (RCA) algorithms, and their optimization, is relevant in several areas of research. [1] suggests application to the problem of decentralized air traffic control in scenarios where a centralized control scheme has become unwieldy due to high airspace density. [2] describes the application of RCA to control arrays of space satellites collecting astronomical data (for example interferometry). In addition, these schemes can be used in computer simulations modeling flocking behavior in biological organisms [3]. The overarching objective is to avoid collisions while maintaining an ordered formation. In the interests of expanding the scheme to larger and more complex scenarios (i.e. thousands, rather than tens, of agents), local rather than global information must be the primary means of determining the trajectory - that is, while broad strategic decisions regarding flock motion can be made at a higher managerial level, agents must make use of what limited sensing capability they have (e.g. limited sensing radius or line of vision) in order to avoid collisions.

Rimon and Koditschek [4] defined a global potential function $V(\mathbf{x})$ such that under the action of $\mathbf{F} = -\nabla V$ a single agent could be guided to a target location; high potentials repel the agent from obstacles and a low potential attracts it to the target destination. Although collision avoidance was proven under various conditions, this method has several drawbacks: the global potential field needs to be known in order to calculate its gradient at any one time; certain conditions need to be satisfied in order to avoid trapping the agent in a local minimum; the method does not consider multiple agents. It also becomes computationally impractical when modeling complex situations.

Marsden, Shadden and others [5], [6] added gyroscopic forces (defined as forces acting in a direction perpendicular to the velocity vector). Collision avoidance was proven for a pair of agents [5]. More recently, schemes have used a combination of Rimon-Koditschek potential terms, drag forces, and gyroscopic and braking forces to avoid collisions.

3 Project Structure

The summer project became associated with the work of those at Caltech working on the DyNARUM (Dynamic Network Analysis for Robust Uncertainty Management) program, funded by DARPA. Although that program has various goals, this project was of relevance to the collision avoidance and surveillance aspects of the program. It comprised two distinct parts:

- Surveillance performance of multiple agents performing random search of a uniform square area, under the dynamics described in §2.
- Formation flight, using elaborated potential-based methods, in the presence of an obstacle-strewn environment, using only information local to each agent.

These two components can be linked in the sense that, as discussed in project meetings, it may be desirable in the long term to develop a scheme whereby a set of agents can ascertain the best method for searching an area (which could be complex, varying drastically in feature size, or time-evolving), such that only the broadest managerial control is necessary. Examples of this could be a search and rescue with several helicopters, or military surveillance with many cheap or even disposable aircraft.

4 Surveillance

4.1 Description of scheme

Although the best method for surveying a large square area with multiple agents travelling at constant speed is to use a systematic 'lawnmower' sweep (especially if the starting locations can be chosen. In fact any method chosen such that the paths of the agents do not overlap is clearly optimal. However, the performance of this random search technique is of interest because the theory may then be extended to complex or time-evolving regions where the 'lawnmower' breaks down. The following is a description of the scheme used in this project.

- Define lattice of points spaced at $2r_s$ apart.
- Assign n of these points as starting positions.
- Give all vehicles randomly assigned destinations.
- When a vehicle reaches destination, choose a new point, irrespective of whether it has been chosen before.
- Measure T_{90}, time taken to cover 90% of the area (measured by coverage of many other randomly-generated points).

4.2 Results

A typical curve showing Cumulative Area Coverage (CAC) vs. time is shown in Figure 1. The mean values of T_{90} for each n after 30 trials, along with standard deviations, are plotted in Figure 2. The 90% coverage time T_{90} was chosen since for large regions 100% coverage is only reached asymptotically. Clearly there is a very good fit with an inverse k/n relationship. Although effort was made to compute and find a trend in the control effort required, there was not enough time for a satisfactory analysis. Sujit Nair [7] has run similar trials using agents in a circular area moving under 'hovercraft' dynamics and it is hoped that some comparison of the control cost associated with different collision avoidance schemes can be made.

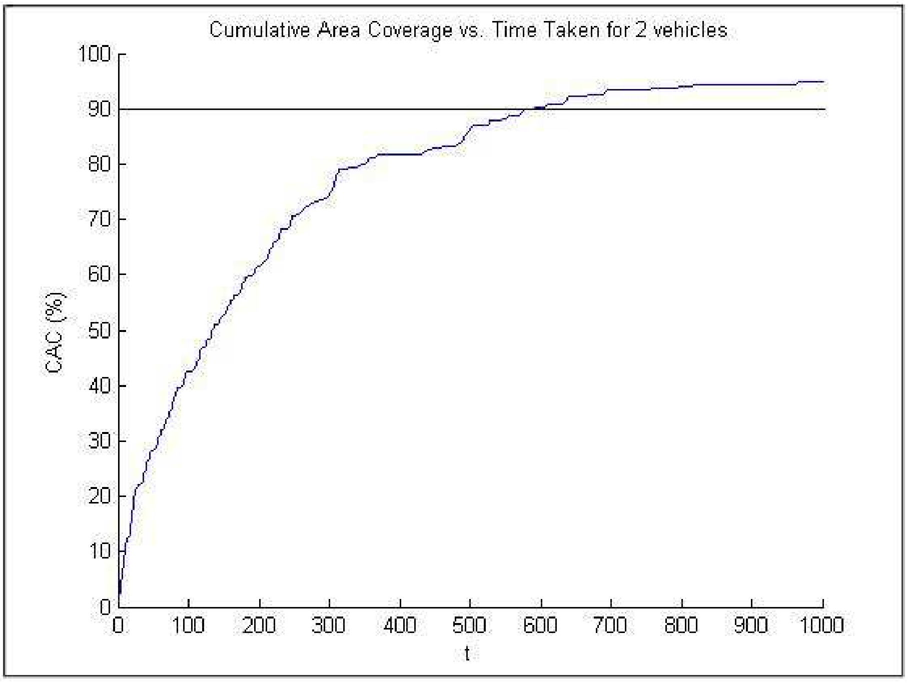

Figure 1: Typical Cumulative Area Coverage curve for $n = 2$.

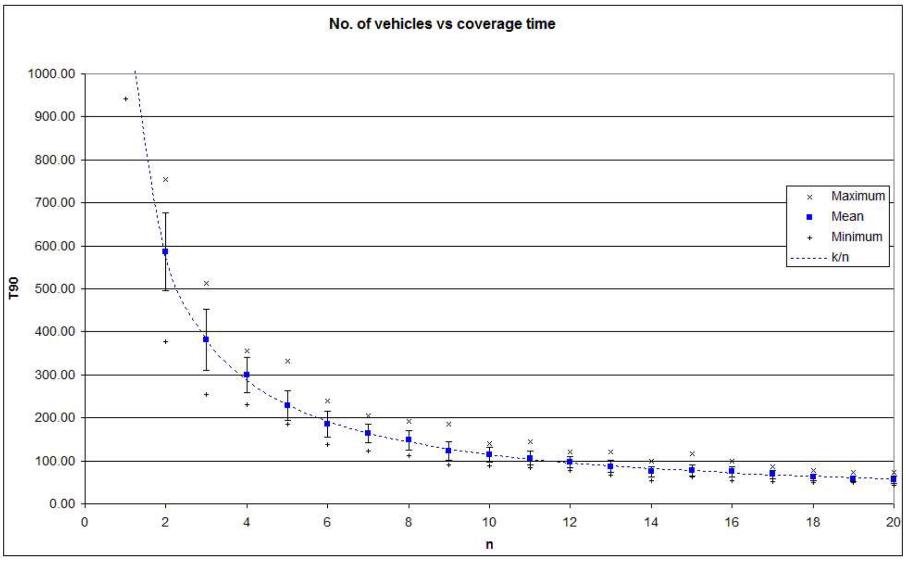

Figure 2: Mean T_{90} vs. n for 30 iterations. Error bars show standard deviations.

5 Formation flight

5.1 Description of scheme

The agents in the 'flock' act under the following scheme:

- Agents may only sense up to a sensing radius r_s.
- The lattice constant, or equilibrium spacing, r_l is set such that $\frac{1}{2}r_s < r_l < r_s$.
- A 1D (i.e. varying with radius only) potential acting between all pairs of agents is defined such that the minima are found at desired spacings, depending on the configuration desired.

5.2 Choice of potential functions

To avoid collisions with walls, all that was necessary was a repulsive force that increases with proximity with the walls, up to some maximum based on the largest possible acceleration of an agent. However, behavior of the flock was very sensitive to the potential functions defined between agents. Much is owed to Philip Du Toit [8] for his advice on choosing suitable potential functions.

5.2.1 Triangular formation

For a triangular lattice to be stable, it should be the only low-energy configuration possible; therefore the inter-agent potential shape was chosen such that the only minimum was found at r_l. The function chosen was piecewise-quadratic for simplicity, and since many such functions exist a certain amount of tuning was necessary in order to achieve satisfactory performance. Although a Leonard-Jones potential form was tried, it was quickly found to be unsuitable for this model, since it prevents spacings of less than r_l with unrealistically large forces.

The behavior of square and triangular configurations of agents was modelled in an environment featuring constraining walls and circular obstacles. Agents were each given a constant force acting in the positive y-direction, and in addition acted under the potential fields between themselves and all other agents in range:

$$\ddot{\mathbf{q}}_i = K_p \mathbf{u}_y - K_s \sum_{j \neq i} \nabla V_j - K_s \sum_{w} \nabla V_w$$

for all j for which $||\mathbf{q}_j - \mathbf{q}_i|| < r_s$, and w are the walls for which $||\mathbf{q}_w - \mathbf{q}_i|| < r_s$, where $\mathbf{q}_w$ is the closest point on the wall to the agent. The potential field is evaluated as a function of distance from $\mathbf{q}_w$. K_p and K_s are positive gains.

5.2.2 Square formation

Achieving a square lattice is considerably more difficult, especially when restricted to a 1D potential function. The method employed was to place minima at spacings commonly found in a square formation (1, $\sqrt{2}$ etc.), while placing maxima at those found in a triangular lattice ($\sqrt{3}$). Clearly even more possible functions of this type exist, even when the range of the potential is less than $2r_l$. In this case, Gaussian curves of various heights and widths were superimposed to get the desired set of maxima and minima, and a steep quadratic section was added for r close to 0, to avoid collisions.

5.3 Successful simulations

5.3.1 Triangular formation

Figure 3 shows a set of frames from a successful simulation of 70 agents travelling in triangular formation through the environment [9]. Collisions are avoided between the walls and obstacles, and between agents. Clearly, more than this is needed if the whole area is to be covered systematically, as would be required in a surveillance exercise. As can be seen, the agents become 'strung out' if the width of the flock is not constricted in some way. Various elaborations on this technique were tried, but with only limited success. The following may be implemented with more success in future work:

- Trapping the flock behind a virtual wall moving at constant speed, or some distance ahead of the centre of mass of the flock.
- Adjusting the spacing between agents to reflect the sparseness of the environment. In wide unconstrained areas, spacing can be increased so that the flock fills the width; in narrow areas, the equilibrium spacing can be reduced.
- Changing the dynamics between pairs of agents in response to managerial commands (for example damping or gain).

5.3.2 Square formation

A square formation, and changes between triangular to square in both directions, were achieved in the static case for a flock of 15 agents [9]. Increasing the number of agents allowed a greater degree of imperfection, especially

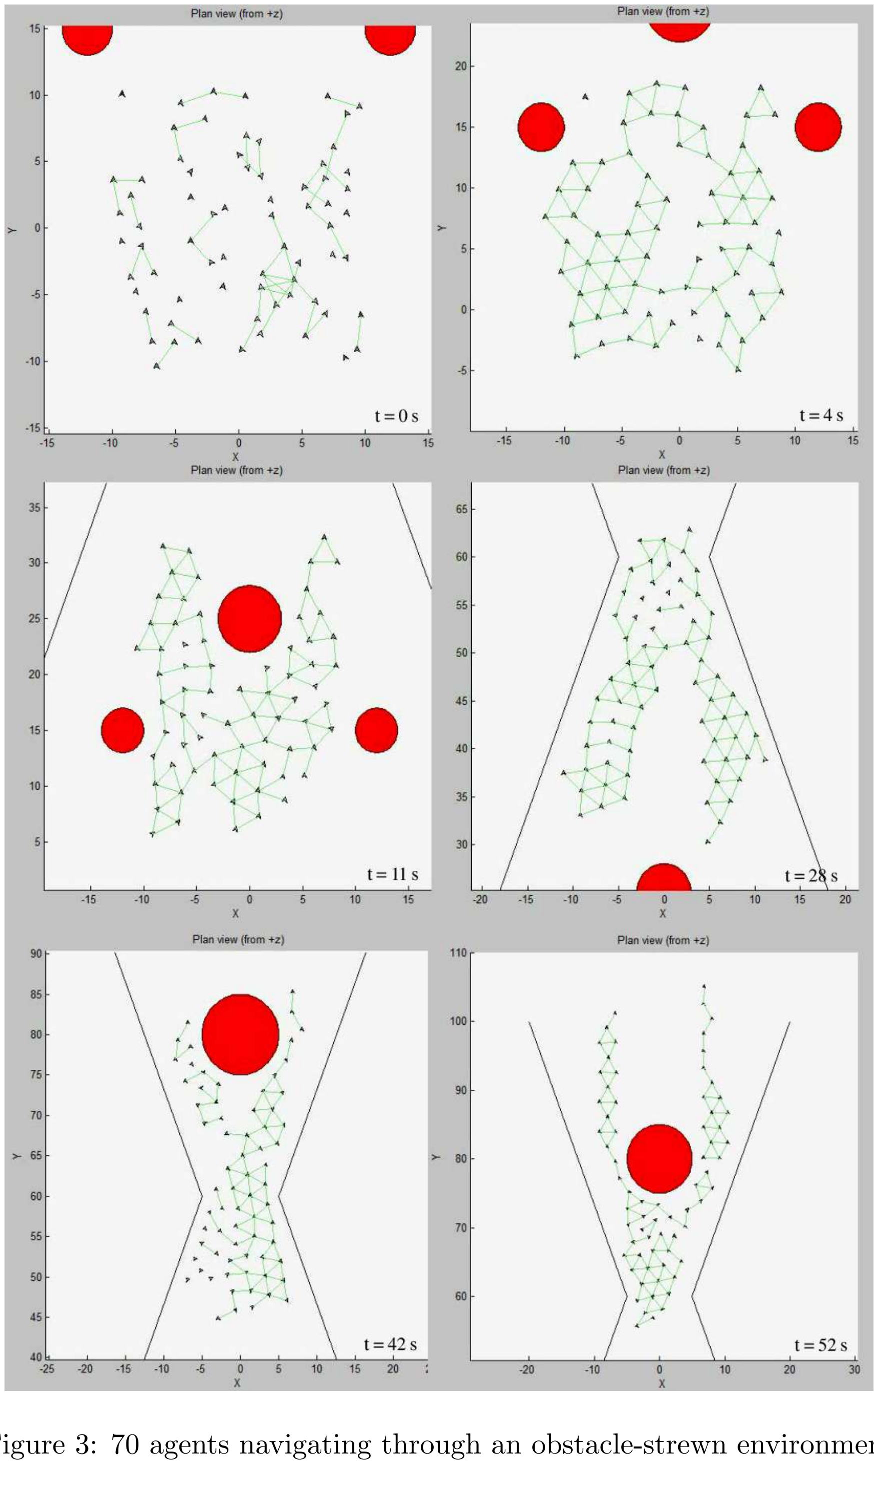

Figure 3: 70 agents navigating through an obstacle-strewn environment

in the square lattice. It proved much more difficult, however, to sustain a square formation in the dynamic case, i.e. within a moving flock. This is something which can be worked on in future.

6 References

[1] Biologically-inspired Reactive Collision Avoidance; Ross, Marsden, Shadden, Sarohia, C.I.T.
[2] A Direct Solution for Fuel-Optimal Reactive Collision Avoidance for Collaborating Spacecraft; Scharf, Aikmese, Ploen, Hadaegh, Proceedings of the 2006 American Control Conference June 2006.
[3] Flocks, Herds, and Schools: A Distributed Behavioral Model; Reynolds, in Computer Graphics, SIGGRAPH Conference Proceedings, 1987.
[4] Exact robot navigation using artificial potential fields; E. Rimon and D. E. Koditschek, IEEE Transactions on Robotics & Automation, October 1992.
[5] Gyroscopic Forces and Collision Avoidance with Convex Obstacles, Chang and Marsden 2003.
[6] Collision Avoidance for Multiple Agent Systems; Chang, Shadden, Marsden, Olfati-Saber, Proceedings of the 42^{nd} IEEE Conference on Decision and Control 2003.
[7] See *http://www.cds.caltech.edu/~nair/academics.php*
[8] See *http://www.cds.caltech.edu/~pdutoit*
[9] For movies, see *http://www.cds.caltech.edu/~jcw*

7 Acknowledgments

Many thanks to Sujit Nair for his guidance throughout the project. Further thanks to Konda Reddy Chevva, who provided guidance throughout the second part of the project. Thanks to Philip Du Toit for his guidance on using potential-based methods to achieve formations. Thanks to Shane Ross for supplying code from his own collision avoidance simulations for this project. Thanks to Dr. Marsden for overall direction of the project.