

TECHNICAL REPORT

N.S.F. Research Experience for Undergraduates
in Intelligent and Autonomous Robotics
at the University of Central Florida, Dept. of Electrical and Computer Engineering
May 22nd – July 30th, 2005

Student Researcher: Bess L. Walker, Purdue University
Research Advisor: Keith Garfield, Institute for Simulation and Training, U.C.F.
Program Coordinator: Dr. Zhihua Qu, U.C.F.

Self-Ranked Leadership: A Social Feature Applied to Swarms in Extensions of the Opera Problem

I. Introduction

Research on swarm behavior – where the behavior of a large group of individual agents is determined by their purely local interactions – has been a growing area of interest for years. The list of possible applications for such swarms is long. A swarm is very difficult to disable, since it does not depend on any of its parts. Any individual may be stupid, inexpensive, and easily replaced, but with the proper local interactions the swarm as a whole may exhibit amazingly complex and intelligent behavior.

We have used the MASON simulator to model a swarm of simple mobile autonomous vehicles (MAVs) with three very simple behaviors. Each MAV can seek a target, avoid obstacles, and avoid other MAVs. They are placed in environments which are extensions of the Opera Problem: increasingly complicated systems of hallways and dead ends lie between the swarm and its goal. These simple behaviors suffice for very simple environments. Most notably, MAVs get stuck in dead ends.

We seek to develop robust sets of behaviors for the swarm using social features in addition to the simple behaviors outlined above. A successful MAV gains leadership status, and less successful MAVs follow their local leader. Since MAVs stuck in a dead end are not successful, eventually MAVs are drawn out of dead ends and the swarm as a whole is more successful in reaching its goal.

The goal in developing such a robust behavior is not ensure the swarm achieves its best performance on any particular environment, successful or complex, but rather to ensure is successful on a good-enough level on all environments. Such an approach to behavior design is crucial if swarms are ever to be deployed in the real world, where flexibility trumps all other concerns.

Previous work relating to swarm behavior, leadership, and the opera problem include **2, 3, 4**. Dr. Annie Wu and Keith Garfield at the University of Central Florida are

leading a team investigating a dominance-like social feature applied to the MAVs and opera problem extensions described in this paper [1].

II. Simulation Specifics

II.1. The Opera Problem

The classic opera problem postulates the audience's attempt to escape the opera house through one door if someone suddenly shouts "Fire!" Each person tries to get out as quickly as possible. The panic is such that it is each man for himself rather than an organized process. This is well-modeled by the local decision-making processes used by swarms.

Although the opera problem itself is a toy problem, extensions range from a simple hallway (Figure II.1) to a complex maze. Such extensions quickly become applicable to real problems in the real world. The project seeks to develop simple static behaviors (the agents cannot learn) which allow agents to function successfully in this wide range of environments.



Figure II.1 Agents attempt to get through the hallway in this simple extension of the opera problem.

II.2 The Agents: MAVs

Our agents are simulated MAVs (Mobile Autonomous Vehicles), modeled in the MASON simulation utility.

A MAV has no memory. Its next state is entirely determined by its present state, its perception of its environment, and a small set of rules for movement. Each MAV considers 72 compass directions in both sensing and movement. All MAVs have the same constant speed, but are allowed to halt if necessary. The sensory radius is relatively small (a MAV has radius 4 and sensor radius 30) and the same as the object avoidance radius; a MAV will normally attempt to avoid anything it can sense.

Each MAV is assigned three Goals: seek the goal points, avoid obstacles, and avoid other MAVs. Each Goal generates a 72-element array telling the MAV which directions are positive, negative (could cause collisions), and neutral. A direction marked negative by any Goal is considered forbidden to the MAV. A direction marked positive by any Goal is recommended to the MAV. All remaining directions are neutral.

The MAV chooses the recommended direction which is closest to its direction at the last step off the simulation. If there are no recommended directions, it chooses the neutral direction closest to its last direction. If all 72 directions are forbidden, the MAV halts.

II.3 The Environments

As test environments, we used four simple extensions of the opera problem: a single hallway, a triple hallway, a single hallway with two shallow false hallways (traps), and a single hallway with two deeper traps (Figure II.2). 50 MAVs start at the top and work their way to the bottom, attracted by three goal points.

The leadership social behavior was introduced to allow MAVs to escape from the traps in the environments with the false hallways; it was clear that without any special features, each false hallway would fill with MAVs and none would be able to escape.




Figure II.2 Experimental environments with MAVs in starting positions.

II.4 The Social Feature: Leadership

Each MAV has a score which functions as a combined approximate measure of success and determination of leadership ability. Each MAV scores itself with no input from others. The leadership score ranges from 0.0 to 1.0 and is modified at each step:

- The leadership score increases when a MAV moves. The closer its direction to its previous direction, the greater the increase in leadership score.
- There is a small constant drain on the leadership score each step. In addition, the leadership score decreases sharply if the MAV does not move.
- Normally a MAV is forbidden to move anywhere near another MAV it senses; the avoidance Goal blocks 180° to prevent any possible collision. However, a MAV is permitted to move *directly* towards the MAV in its vicinity with the highest leadership score.

The practical result is that less successful MAVs exhibit a following behavior. This gives us leadership as a social feature, not just a success score.

This material is based upon work kindly supported by the National Science Foundation under Grant No. 0353918. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

III. Results and Analysis

Leadership outperforms non-leadership in all environments but the triple hallway. Additionally, fewer MAVs are left behind in traps when leadership is used.

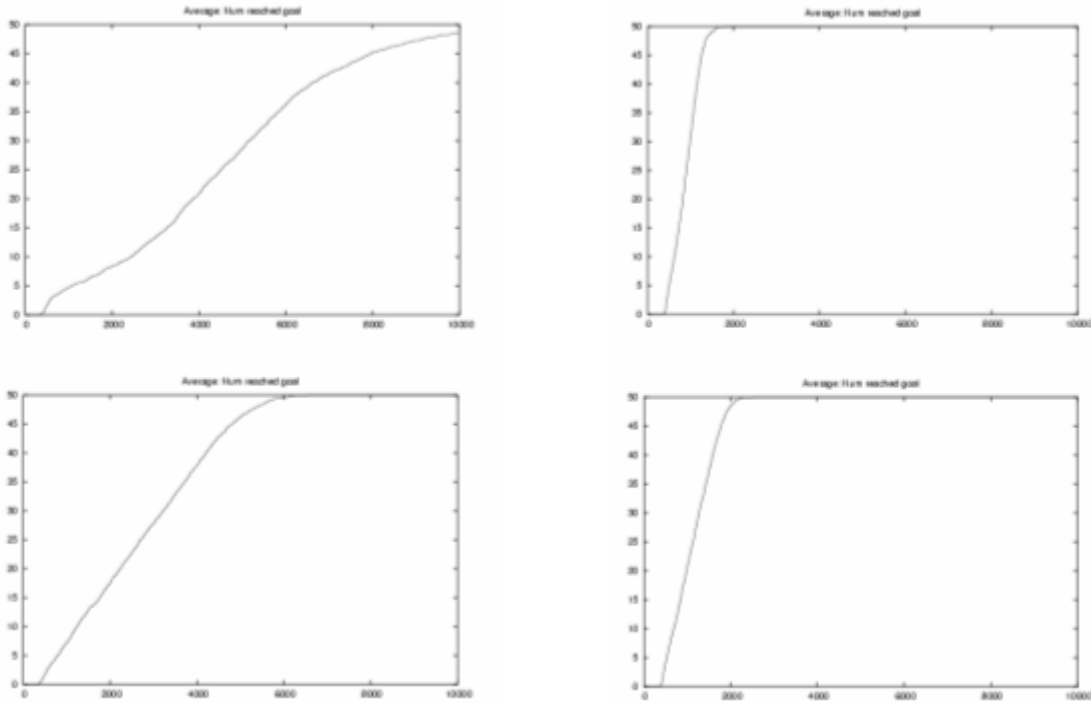


Figure III.1 Untrapped environments, number of MAVs at goal over time, average over 20 runs.

- A. Single hallway, no leadership B. Triple hallway, no leadership
C. Single hallway, with leadership D. Triple hallway, with leadership

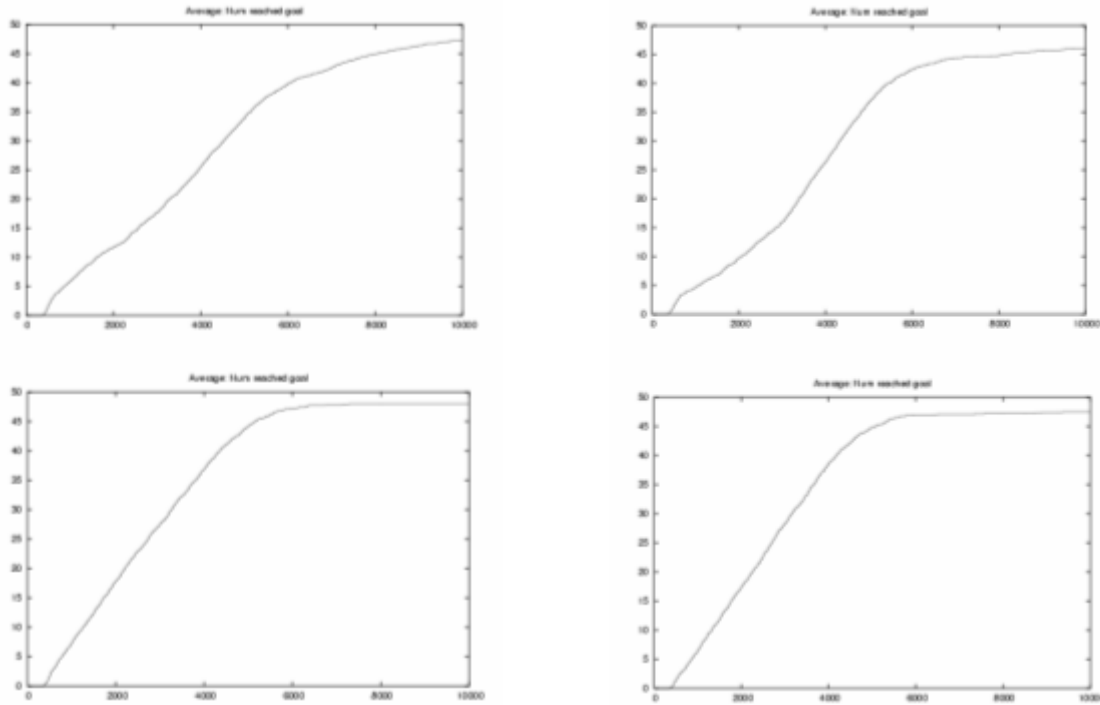


Figure III.1 Trapped environments, number of MAVs at goal over time, average over 20 runs.
 A. Single hallway, shallow traps, no leadership B. Single hallway, deep traps, no leadership
 C. Single hallway, shallow traps, with leadership D. Single hallway, deep traps, with leadership

Results for trapped environments were as expected. Leadership does not prevent *all* MAVs from getting stuck: if there is only one MAV in a trap, it cannot take advantage of the leadership social feature and remains stuck.

It was unexpected that leadership significantly outperformed non-leadership in the case of a single hallway. Surprisingly, when the MAVs are crowded, leadership acts as an overall scattering effect, although locally it is an attractive effect. Leaders develop on the edge of the crowd and draw some MAVs up and away from the congestion, which makes it easier for those MAVs remaining at the hallway entrance to get through. Adding two more hallways has the same effect of reducing congestion, explaining why leadership and non-leadership have roughly the same performance in the triple hallway environment. The lesser performance of leadership on the triple hallway seems to indicate that with another method of reducing congestion, the scattering effects are actually becoming negative.

IV. Directions for Further Study

Currently the sensory radius is equal to the collision radius. As a result, it is difficult for a leader to gather a following of more than one or two other MAVs. We postulate that separating these radii and decreasing the collision radius may allow larger groups to form. This may also allow leadership to be implemented in a way which does not require overriding the collision-avoidance behavior.

This material is based upon work kindly supported by the National Science Foundation under Grant No. 0353918. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Further investigation of the way leadership increases and decreases would also be well-advised. The current values work, according to our results, but we suspect they could be tweaked to encourage better performance. We would also like to continue experiments with increasingly more complex environments; we have shown leadership is a good solution to the problem of simple dead ends, but if it cannot perform in a wide variety of complex environments, the overall goal of robustness has not been reached.

References

1. Garfield, K., A. Wu, M. Onal, Britt, A. Campbell, and R. Shumaker. The Effectiveness of Transferring Multi-Agent Behaviors From a Learning Environment in the Presence of Synthetic Social Features. Submitted to IMEC 2005.
2. Yohei Murakami, Kazuhisa Minami, Tomoyuki Kawasoe and Toru Ishida. Multi-Agent Simulation for Crisis Management. *IEEE International Workshop on Knowledge Media Networking (KMN-02)*, pp. 135-139, 2002.
3. Robinson, A., and L. Spector. 2002. Using Genetic Programming with Multiple Data Types and Automatic Modularization to Evolve Decentralized and Coordinated Navigation in Multi-Agent Systems. In *Late-Breaking Papers of GECCO-2002, the Genetic and Evolutionary Computation Conference*. Published by the International Society for Genetic and Evolutionary Computation.
4. Crespi, V., G. Cybenko, and D. Rus (2001). *Decentralized Control and Agent-Based Systems in the framework of the IRVS*. White paper presented at the PI TASK Meeting held in Santa Fe, NM, on April 2001.
<http://actcomm.thayer.dartmouth.edu/task/crespi/irvs2.pdf>