Ad Hoc Teamwork for Leading a Flock

(Doctoral Consortium)

Katie Genter
Dept. of Computer Science
University of Texas at Austin
Austin, TX 78712 USA
katie@cs.utexas.edu

ABSTRACT

Designing agents that can cooperate with other agents as a team, without prior coordination or explicit communication, is becoming more desirable as autonomous agents become more prevalent. In my work I examine an aspect of the problem of leading teammates in an ad hoc teamwork setting, where the designed ad hoc agents lead the other teammates to a desired behavior that maximizes team utility. Specifically, I consider the problem of leading a flock of agents to some desired behavior using a subset of the flock that is comprised of ad hoc agents. I consider this problem not only theoretically, but also in a custom-designed simulator FlockSim.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Algorithms, Experimentation

Keywords

Ad Hoc Teamwork, Agent Cooperation, Coordination

1. INTRODUCTION

The growing use of agents in various cooperative domains has emphasized the importance of designing agents capable of reasoning about *ad hoc* teamwork [7]. Such agents can cooperate within a team without using explicit communication or previously coordinating behaviors among teammates.

As an example, consider a team of robots attempting to travel from the parking lot at a local park to a bridge that needs to be repaired. The terrain between the parking lot and the bridge is varied, and the straight line path between the two locations will be difficult for the robots to traverse. Now consider that most of the robots that will be traveling from the parking lot to the bridge have a very simple method of planning a path to travel: they generally head directly towards their goal but will adopt the orientation of any teammates within close proximity. If allowed to traverse from the parking lot to the bridge on their own, these simple robots will be forced to traverse steep inclines and very dense brush. However, consider if a few more sophisticated robots were included in the team that could observe the path planning method used by the simple robots and had some previous knowledge about the terrain of the park. These more sophisticated robots could then influence the simpler robots to alter their path such that the worst of the steep inclines and dense brush would be avoided. This would allow the entire team to reach the bridge in less time and with less damage to the robots.

In this example, the more sophisticated robots influenced — or *led* — the simpler robots to adopt a better path from the parking lot to the bridge. One aspect of ad hoc teamwork involves this idea of *leading* teammates to perform desired actions or adopt particular behaviors. In my work, I consider the problem of using ad hoc agents to lead a flock.

Flocking is an emergent behavior found in various species in nature including flocks of birds, schools of fish, and swarms of insects. In each of these cases, the animals follow a simple local behavior rule that results in a stable, well defined group behavior. Research on flocking behavior can be found in various disciplines such as physics [9], graphics [6], biology [1, 2], and distributed control theory [4, 5, 8]. The main focus of each of these research directions is to characterize the emergent behavior.

In my work I consider the problem of leading a team of flocking agents in ad hoc teamwork settings. The ad hoc teamwork perspective of this problem is highlighted by two facts. Firstly, we are unable to explicitly control the behavior of the flocking agents, thus we can only attempt to influence them implicitly using the behavior of the ad hoc agents (which appear to be identical to the flocking agents, and hence are indistinguishable from other agents in the flock). Secondly, all agents — both flocking and ad hoc — act as one team, and their only desire is to optimize team utility. The ad hoc agents cannot communicate with the flocking agents, but they can communicate and otherwise coordinate their actions among themselves.

2. PROBLEM DEFINITION

In recent work [3], I introduced a flocking model that was initially inspired by Vicsek *et al.* [9]. For clarity, the basics of this flocking model are summarized in this section.

In my flocking model, n visually indistinguishable agents inhabit some environment where each agent a_i moves with some velocity v_i. At each time step t, each agent a_i has a position $p_i(t) = (x_i(t), y_i(t))$ in the environment and an

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

orientation $\theta_i(t)$. Each agent's position $p_i(t)$ at time t is updated after its orientation is updated, such that $x_i(t) = x_i(t-1) + v_i \cos(\theta_i(t))$ and $y_i(t) = y_i(t-1) - v_i \sin(\theta_i(t))$.

We let $N_i(t)$ be the set of $n_i(t) \leq n$ agents (including agent a_i) at time t which are visible to agent a_i. An agent is *visible* to agent a_i if its position is located within a *visibility cone* of angle α centered on orientation $\theta_i(t)$ and extending from agent a_i for an unlimited distance (see Figure 1 for an example). We say that angle α defines the visibility cone for each agent, and that this visibility cone defines each agent's *neighborhood* (i.e., the area in which the agent can see other agents). Under our flocking model, the global orientation of agent a_i at time step $t+1$, $\theta_i(t+1)$, is set to be the average orientation of all agents in $N_i(t)$ (including itself) at time t.

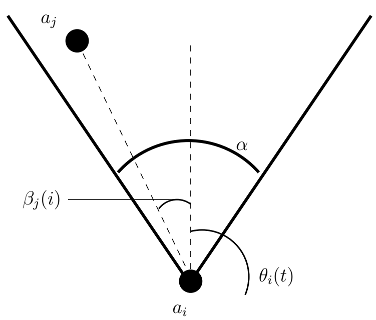

Figure 1: Angle α defines the visibility cone for agent a_i. Agent a_j is visible to agent a_i.

The n visually indistinguishable agents that comprise the flock consist of k *ad hoc* agents and m *flocking* agents, where $k + m = n$. The ad hoc agents $\{a_0, \ldots, a_{k-1}\}$ are agents whose behavior we can control, while the flocking agents $\{a_k, \ldots, a_{N-1}\}$ are agents that we cannot directly control. An x-step plan specifies the orientations that each ad hoc agent $\{a_0, a_1, \ldots a_{k-1}\}$ will align to at each time step when given exactly x time steps in which to act.

3. PUBLISHED WORK

I have one published paper on this work [3], to be presented at AAMAS'13. In this work I presented a specification of the flocking problem as a new scenario for studying ad hoc teamwork and contributed an initial theoretical and empirical analysis of the flocking problem.

A portion of my specification of the flocking problem is summarized above in Section 2. Additionally, I defined the Agent Flock Orientation Manipulation Problem to be as follows: Given a target orientation θ^* and a team of n visually indistinguishable agents $\{a_0, \ldots, a_{n-1}\}$, where each flocking agent $\{a_k, \ldots a_{n-1}\}$ adopts the average orientation of all agents that are visible to it, determine whether the ad hoc agents can influence the flocking agents to align to θ^*, and if so, find the plan π that does so with minimum cost $c(\pi)$. Hence, I also presented a definition for the cost of a plan π, $c(\pi)$, in terms of time steps required to orient the flock to θ^*, plan size, and performance error.

Theoretically, I considered the extent of influence that stationary ad hoc agents can have over stationary flocking agents located at a single position. Specifically, I found that $k_i(t)$ ad hoc agents can influence the $m_i(t)$ flocking agents to turn in a particular direction by any amount less than or equal to $\frac{k_i(t)\pi}{m_i(t)+k_i(t)}$ radians in one time step. I also found the extent to which $k_i(t)$ ad hoc agents within the visibility cone of $m_i(t)$ flocking agents can influence the $m_i(t)$ flocking agents to turn in one time step and still have the same $m_i(t)$ flocking agents and $k_i(t)$ ad hoc agents within the flocking agents' visibility cone. Finally, I was able to set a bound on the number of time steps required for $k_i(t)$ ad hoc agents to influence $m_i(t)$ flocking agents to align to θ^*, assuming θ^* is reachable.

I also presented a method for calculating an optimal behavior for the ad hoc agents, such that the flock is led to a particular orientation in minimal time. Empirically, I also began to consider how non-stationary ad hoc agents that are not within the flocking agents' visibility cone at a particular time step should behave if the goal is to orient the flock towards a particular orientation in minimal time. I also showed via an empirical experiment that my ad hoc agent algorithms perform significantly better than naive methods in which the controllable agents orient towards θ^* such that the flock slowly converges to θ^* (e.g. [5, 8]). My ad hoc agent algorithms performed better because I purposely orient the ad hoc agents past θ^* in order to orient the flocking agents exactly to θ^* quickly.

4. ACKNOWLEDGEMENTS

This work has taken place in the Learning Agents Research Group at The University of Texas at Austin. My collaborators in this work are Professor Noa Agmon at Bar Ilan University and Professor Peter Stone at The University of Texas at Austin.

5. REFERENCES

[1] W. Bialeka, A. Cavagnab, I. Giardinab, T. Morad, E. Silvestrib, M. Vialeb, and A. Walczak. Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences*, 109(11), 2012.

[2] H. H. Charlotte K. Hemelrijk. Some causes of the variable shape of flocks of birds. *PLoS ONE*, 6(8), 2011.

[3] K. Genter, N. Agmon, and P. Stone. Ad hoc teamwork for leading a flock. In *AAMAS'13*, May 2013.

[4] J. Han, M. Li, and L. Guo. Soft control on collective behavior of a group of autonomous agents by a shill agent. *Systems Science and Complexity*, 19:54–62, 2006.

[5] A. Jadbabaie, J. Lin, and A. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988 – 1001, June 2003.

[6] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *SIGGRAPH*, 21:25–34, August 1987.

[7] P. Stone, G. A. Kaminka, S. Kraus, and J. S. Rosenschein. Ad hoc autonomous agent teams: Collaboration without pre-coordination. In *AAAI'10*, 2010.

[8] H. Su, X. Wang, and Z. Lin. Flocking of multi-agents with a virtual leader. *IEEE Transactions on Automatic Control*, 54(2):293–307, Feb 2009.

[9] T. Vicsek, A. Czirok, E. Ben-Jacob, I. Cohen, and O. Sochet. Novel type of phase transition in a system of self-driven particles. *PHYS REV LETT.*, 75:1226, 1995.