Placing Influencing Agents in a Flock

Katie Genter and Peter Stone

Department of Computer Science, University of Texas at Austin
Austin, TX 78712 USA
{katie,pstone}@cs.utexas.edu

Abstract

Flocking is a emergent behavior exhibited by many different animal species, including birds and fish. In our work we consider adding a small set of influencing agents, that are under our control, into a flock. Following ad hoc teamwork methodology, we assume that we are given knowledge of, but no direct control over, the rest of the flock. In our ongoing work highlighted in this abstract, we are specifically considering the problem of where to initially place influencing agents that we add to such a flock. We use these influencing agents to influence the flock to behave in a particular way - for example, to fly in a particular orientation or fly in a particular pattern such as to avoid an obstacle.

Introduction

Flocking is an dynamic, emergent swarm behavior found in various species in nature. Each animal in a swarm follows a simple local behavior rule that results in a group behavior that often appears well organized and stable. In this abstract, we briefly introduce some of our ongoing work on the problem of *leading* a team of flocking agents in an ad hoc teamwork setting. An ad hoc teamwork setting is one in which one or more teammates — which we call *influencing agents* — must determine how to best achieve a team goal when given a set of possibly suboptimal teammates that we are unable to directly control or alter (Stone et al. 2010).

As a motivating example, imagine that a flock of migrating birds is flying directly towards a dangerous area, such as a wind farm or an airport. Our goal is to encourage the birds to avoid the dangerous area without significantly disturbing them. Since there is no way to directly control the flight path of the birds, we must instead alter the environment so as to induce the flock to alter their flight path as desired.

In our work, we assume that each bird in the flock dynamically adjusts its heading based on that of its immediate neighbors. We assume that we control one or more influencing agents — perhaps in the form of robotic birds[1] or ultralight aircraft[2] — that are perceived by the rest of the flock as one of their own. It is through these influencing agents that

we alter the birds' environment so as to induce them to alter their path.

Our previous work has considered how randomly placed influencing agents should behave so as to influence the flock to face a particular direction or maneuver along a path so as to avoid an obstacle (Genter and Stone 2014). In our current research summarized in this abstract, we wish to examine how the influencing agents should be placed in the flock. Specifically, we are considering: *where should influencing agents be initially located within a flock to maximize their influence on the flock?*

Problem Definition

We use a simplified version of Reynolds' Boid algorithm for flocking (Reynolds 1987) in which each agent calculates its orientation for the next time step to be the average heading of its *neighbors* at the current time step. An agent's *neighbors* are the agents located within some set radius of the agent. In order to calculate its orientation for the next time step, each agent computes the vector sum of the velocity vectors of each of its neighbors and adopts a scaled version of the resulting vector as its new orientation. At each time step, each agent moves one step in the direction of its current vector and then calculates its new heading based on those of its neighbors. All agents maintain constant, identical speeds. Our domain is not a torus, so once agents leave the visible domain we consider them lost forever and irretrievable by influencing agents.

The n agents that comprise the flock consist of k *influencing* agents and m *flocking* agents, where $k + m = n$. The influencing agents attempt to influence the flock to travel in a pre-defined direction — we refer to this direction as θ^*. We conclude that the flock has converged to θ^* when every agent that is not lost forever and is not an influencing agent is facing within 0.1 radians of θ^*.

Control of Influencing Agents

In recent work we considered how to influence a large, non-stationary flock to (1) quickly *orient* towards a target orientation and (2) *maneuver* through turns quickly but with minimal agents becoming lost as a result of these turns (Genter and Stone 2014). We introduced a 1-step lookahead algorithm for determining the individual behavior of each influencing agent. This 1-step lookahead algorithm considered

---

[1]www.mybionicbird.com

[2]www.operationmigration.org

all of the influences on the neighbors of an influencing agent and allowed the influencing agent to determine the best orientation to adopt (where best is defined as the behavior that exerts the most influence on the next step).

Maximizing Influence

In our current work, we consider where to place k influencing agents in a flock initially — in other words, at the point the influencing agents begin influencing the flock. In the following subsections we shortly discuss our base case, as well as the various approaches we are currently considering for initially placing the influencing agents.

Random Approach

Our past research has randomly placed the influencing agents within the dimensions of the flock (Genter and Stone 2014). Hence, we use random placement as the base case for evaluating our other placement approaches.

Grid Approach

Grid placement places k influencing agents at predefined, well-spaced, gridded positions throughout flock. The placement of the influencing agents is dependent on the space covered by the flocking agents, and not on the positions of flocking agents and the actual grid size is dependent on k. In any particular grid, the influencing agents are spread out among the possible positions as much as possible.

Border Approach

Our border approach works by placing k influencing agents as evenly as possible around the space covered by the flocking agents. As in the grid approach, the placement of the influencing agents is not dependent on the positions of flocking agents. Instead, we place influencing agents on each side of the flock such that at most $\frac{k}{4}$ influencing agents are positioned on any particular side of the flock. If more than one influencing agent is placed on a particular side of the flock, the influencing agents spread out as much as possible on that side of the flock.

Graph Approach

The graph approach places the influencing agents such that the number of flocking agents not connected to any influencing agent directly or indirectly is minimized. If multiple placements of influencing agents lead to minimal unconnected flocking agents, then the graph approach attempts to maximize the number of direct or indirect connections between flocking agents and influencing agents. This graph approach can be thought of as a graph connectivity problem, where each agent is a node and each set of neighbors are connected by an edge.

Preliminary Results

We utilize the MASON simulator (Luke et al. 2005) for our experiments. In our experiments, we use each of the four methods discussed above to initially place the influencing agents and then the influencing agents behave according to the 1-step lookahead algorithm. Due to the limited length of this abstract, we do not report the details of our experimental setup and we only present one set of preliminary experimental results.

Figure 1 shows the average number of flocking agents lost by the different approaches when $n = 10$. Note that the graph approach loses fewer flocking agents than the other approaches for each of the five k's considered. This is because the flocking agents are sparsely populated in the environment, and hence tend to have fewer neighbors. This is important because flocking agents with few neighbors are less likely to be influenced by distant influencing agents. Additionally, the graph approach performs better than the other approaches when the percentage of influencing agents in the flock is higher. This is because the graph approach focuses on minimizing the number of unconnected flocking agents, so as a higher percentage of the flock is composed of influencing agents, the number of unconnected flocking agents decreases quicker for the graph approach than for the other approaches. Finally, note that no flocking agents are lost by the graph approach when 50% of the flock is composed of influencing agents. This is because the graph approach places the influencing agents such that each flocking agent is influenced by at least one influencing agent.
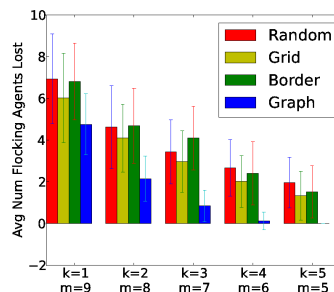

Figure 1: The average number of flocking agents lost across 100 trials.

Discussion

Our ongoing work considers how influencing agents can best be placed into a flock. We give a high-level overview of four approaches that could be used to place the influencing agents into the flock, and provide one set of preliminary results on a small flock of ten agents.

References

Genter, K., and Stone, P. 2014. Influencing a flock via ad hoc teamwork. In *ANTS'14*.

Luke, S.; Cioffi-Revilla, C.; Panait, L.; Sullivan, K.; and Balan, G. 2005. Mason: A multi-agent simulation environment. *Simulation: Transactions of the Society for Modeling and Simulation International* 81(7):517–527.

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

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