Ad Hoc Teamwork for Leading a Flock

Katie Genter[1], Noa Agmon[2], and Peter Stone[1]

[1]University of Texas at Austin
Austin, TX 78712 USA

[2]Bar Ilan University
Ramat Gan, 52900, Israel

May 10, 2013

UT Austin Learning Agents Research Group

Outline

Ad Hoc Teamwork

Always:

- Only in control of a single agent or subset of agents
- Shared goals
- No pre-coordination

Sometimes:

- Unknown teammates
- No explicit communication

Flocking

- Emergent behavior found in nature
  - Birds, fish, insects
- Animals follow a simple local behavior rule
- Group behavior is cohesive

Example — Leading Teammates in Ad Hoc Settings

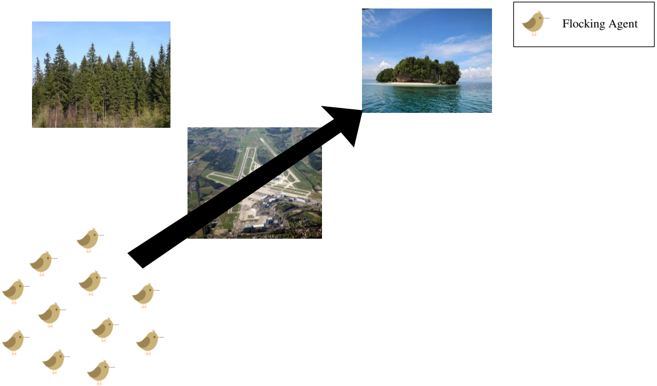

UT Austin Learning Agents Research Group

Example — Leading Teammates in Ad Hoc Settings

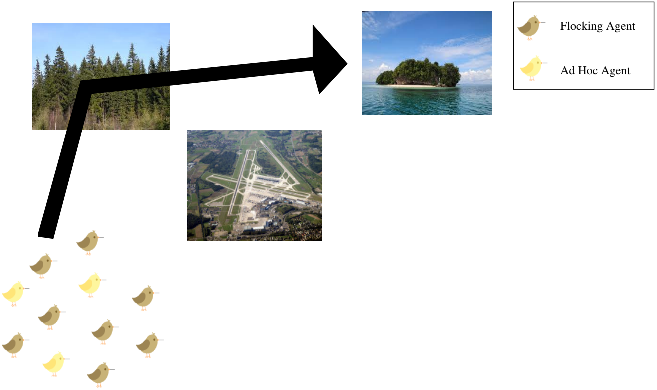

UT Austin Learning Agents Research Group

Example — Leading Teammates in Ad Hoc Settings

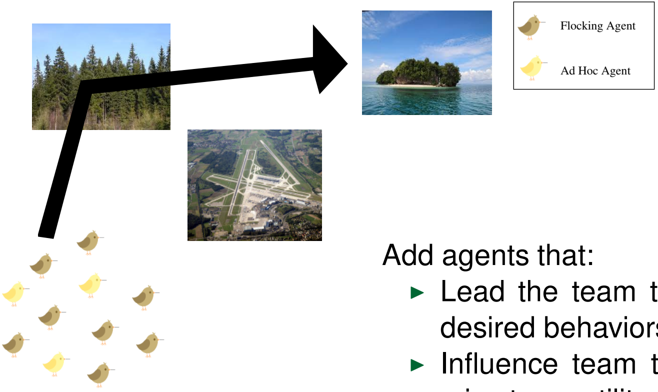

Add agents that:

- Lead the team to adopt desired behaviors
- Influence team to maximize team utility

UT Austin Learning Agents Research Group

Flocking + Ad Hoc Teamwork

Why is this an ad hoc teamwork problem?

- No explicit control of flocking agents
- All agents have shared goals (maximize team utility)
- On-the-fly coordination

Flocking + Ad Hoc Teamwork

In previous work (Jadbabaie et al. 2003, Su et al. 2009), the flock eventually converges to a single controllable agent's heading.

Flocking + Ad Hoc Teamwork

In previous work (Jadbabaie et al. 2003, Su et al. 2009), the flock eventually converges to a single controllable agent's heading.

Research Problem:

Is it possible for one or more agents to lead the team to a desired orientation, and if so - what is the most efficient way of doing so?

Outline

UT Austin Learning Agents Research Group

Problem Definition

Each agent has:

- Constant velocity
- 2D Position
- Global orientation

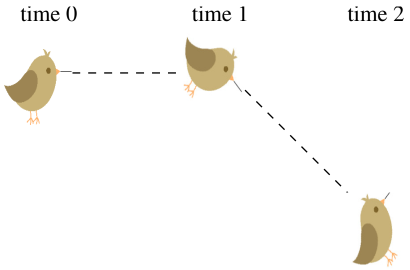

Problem Definition - Neighborhood

Each flocking agent reacts only to agents within a certain *neighborhood* around itself.

- Characterized by a *visibility cone*

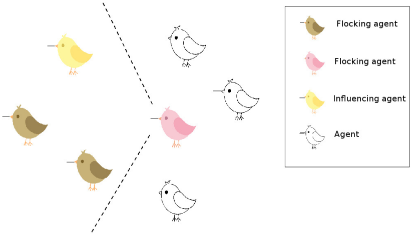

Problem Definition - Orientation Update

A flocking agent's orientation at the next time step is set to be the *average global orientation* of all agents currently within the agent's visibility cone.

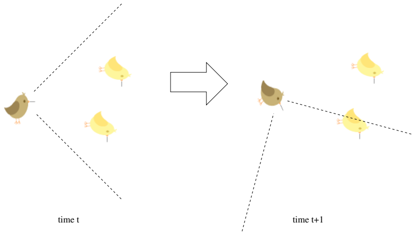

UT Austin Learning Agents Research Group

Problem Definition

(Loading Video...)

UT Austin Learning Agents Research Group

Outline

UT Austin Learning Agents Research Group

Stationary Agents Setup

As the flocking agents are influenced to turn towards a target orientation, different ad hoc agents become available to influence the flocking agents.

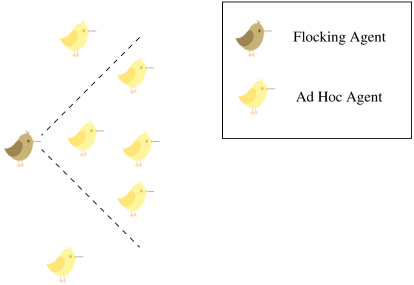

Stationary Agents Theorems

It suffices to consider only algorithms that choose at each time step just one orientation for all of the ad hoc agents to adopt.

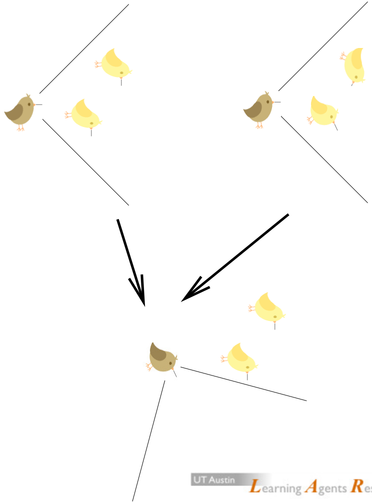

Stationary Agents Theorems

In previous work (Jadbabaie et al. 2003, Su et al. 2009):

- The flock eventually converges to a single controllable agent's heading

Stationary Agents Theorems

In previous work (Jadbabaie et al. 2003, Su et al. 2009):

- The flock eventually converges to a single controllable agent's heading

In this work:

- We prove a tight bound on the number of time steps needed for the ad hoc agents to influence the flocking agents to reach θ^* (when θ^* is reachable).

Forward Search Planning Method

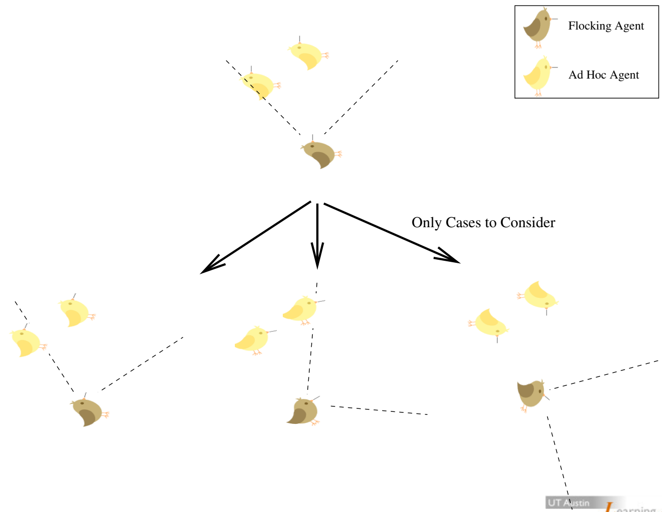

Stationary Agents Video

(Loading Video...)

Outline

Non-stationary Ad Hoc Agents

Heuristic behaviors for ad hoc agents that are *not* within the visibility cone of any flocking agents.

- Towards Flocking Agent
- Towards Visibility Cone

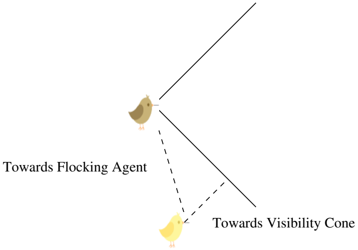

UT Austin Learning Agents Research Group

Non-stationary Agents Video

(Loading Video...)

Experimental Methodology

- Agents randomly placed in a 950 cell by 500 cell environment
- Ad hoc agents have velocity of 50 cells/step
- Each experimental configuration used the same randomization seed

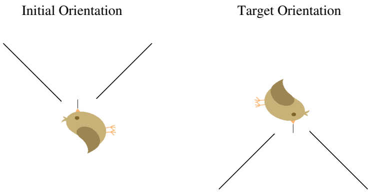

Empirical Results

Is there a significant difference in the number of steps required for the flocking agents to orient to θ^* with each heuristic behavior?

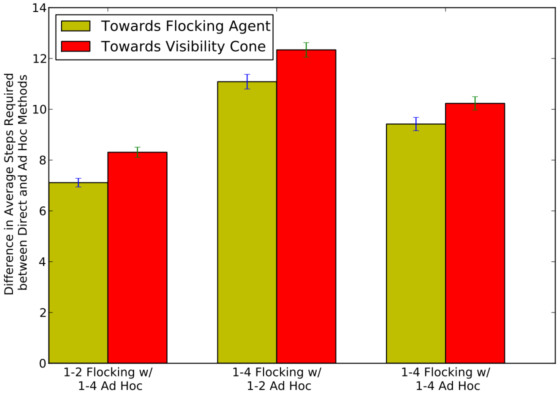

Agents Research Group

Related Work — Ad Hoc Teamwork

- Jones et al. 2006
  - Empirically studied dynamically formed heterogeneous multi-agent teams
  - All agents know they are working as a team
- Agmon and Stone 2012, Stone et al. 2010
  - Leading teammates in ad hoc settings from a game theoretic approach
- Stone et al. 2010
  - Introduced the ad hoc teamwork problem

Related Work — Flocking

- Han et al. 2006
  - Studied how one agent can influence the direction in which a flock of agents is moving
  - Utilized one ad hoc agent with unlimited, non-constant velocity
- Reynolds 1987, Vicsek 1995
  - Concerned with simulating flock behavior
  - Not concerned not with adding controllable agents to the flock
- Jadbabaie et al. 2003, Su et al. 2009
  - Used controllable agents to influence the flock
  - Only concerned with making the flock converge to some orientation eventually

Future Work

- More efficient search (ARMS '13)
- General case of non-stationary agents
- Optimal behavior for non-stationary ad hoc agents

UT Austin Learning Agents Research Group

Summary

Research Problem:

Is it possible for one or more agents to lead the team to a desired orientation, and if so - what is the most efficient way of doing so?

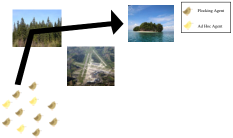

UT Austin Learning Agents Research Group