Improving Efficiency of Leading a Flock in Ad Hoc Teamwork Settings

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 7, 2013

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

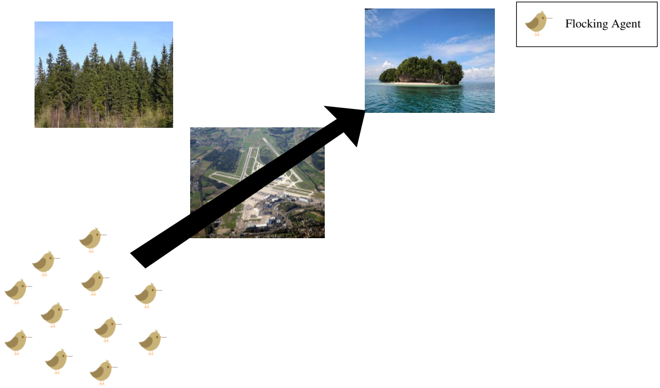

Example — Leading Teammates in Ad Hoc Settings

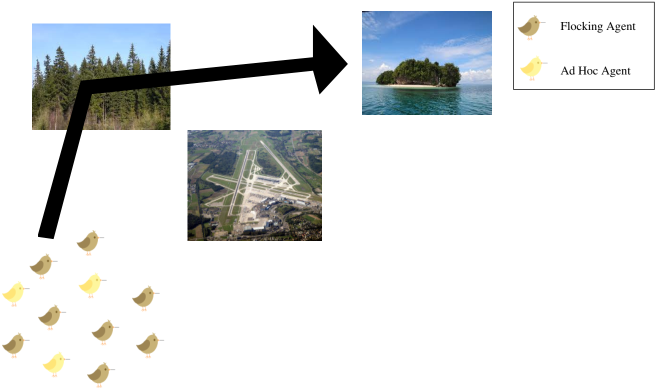

Example — Leading Teammates in Ad Hoc Settings

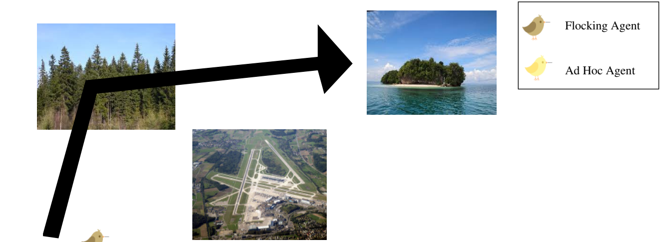

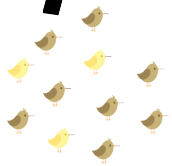

Add agents that:

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

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

Problem Definition

Each agent has:

- Constant velocity
- 2D Position
- Global orientation

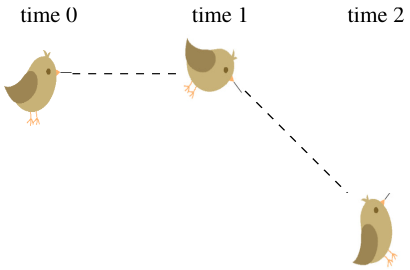

Problem Definition - Neighborhood

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

- Characterized by a *visibility cone*

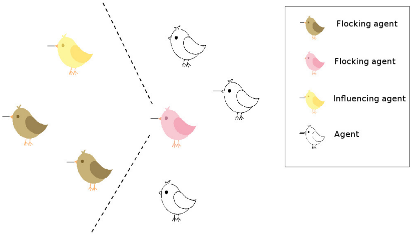

Problem Definition - Orientation Update

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.

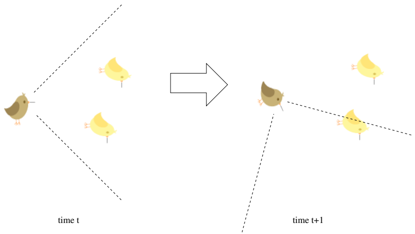

Problem Definition

(Loading Video...)

Outline

Forward Search Planning Method (AAMAS'13)

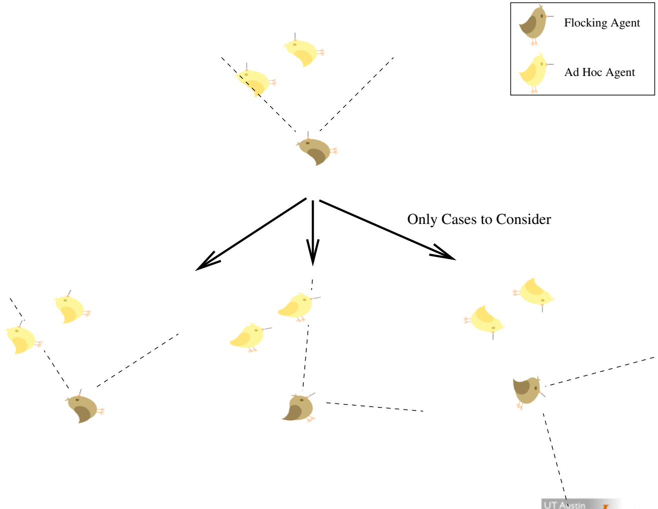

Backward Search Planning Method

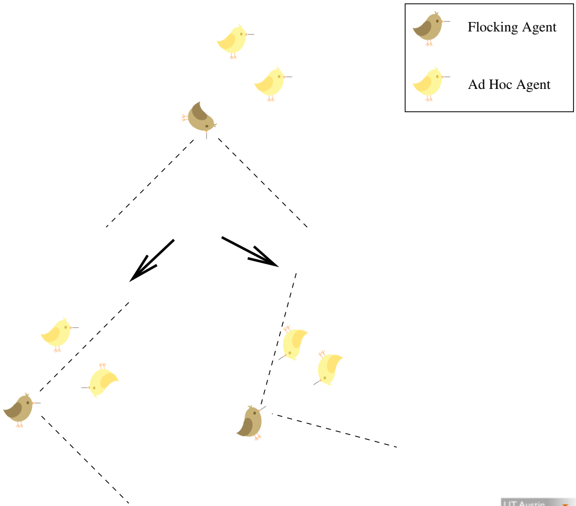

Comparison of Forward and Backward Search Methods

- Forward Search
  - Planning for moving ad hoc agents is easier and more intuitive
  - Less efficient ($2^{\text{numAdHoc}} * \text{numAdHoc} + 1 * \text{maxSteps}$)
- Backward Search
  - Planning for moving ad hoc agents is more difficult
  - More efficient ($\text{maxSteps} * 2\text{numAdHoc}^2$) due to better pruning

Outline

Motion Can Be Helpful

Non-stationary ad hoc agents can influence the flocking agents to reach θ^* faster than stationary ad hoc agents.

(Loading Video...)

Motion Can Be Harmful

Non-stationary ad hoc agents can influence the flocking agents to reach θ^* slower than stationary ad hoc agents.

(Loading Video...)

Outline

Overview

- Altering Ad Hoc Agent Behavior
- Replanning Ad Hoc Agent Behavior
  - Move Inside Visibility Cone
  - Move Border Closer to θ^*

Altering Ad Hoc Agent Behavior

- Keeps the same desired sequence of orientations for the flocking agents
- Recalculates ad hoc orientations
- May not be possible in some situations

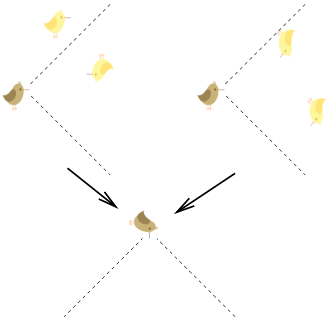

Replanning Ad Hoc Agent Behavior

Move Inside Visibility Cone

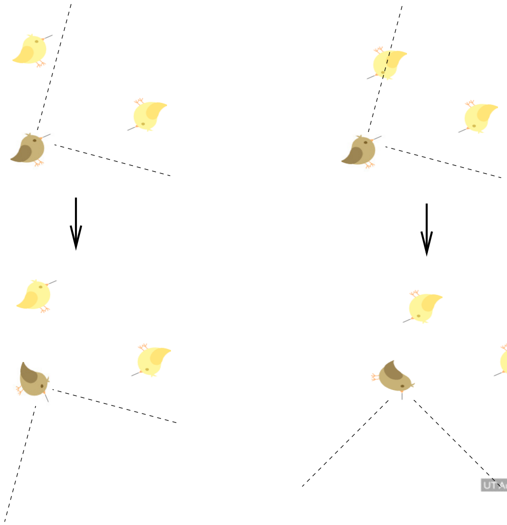

Replanning Ad Hoc Agent Behavior

Move Border Closer to θ^*

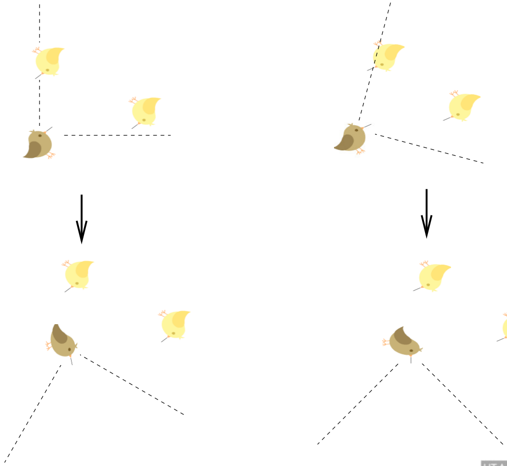

Conjecture

Running the plan repair methods on all the minimal size plans returned by the search will obtain an optimal plan for moving ad hoc agents.

- Must all minimal size plans be repaired?
- Can just one minimal size plan be repaired?
- Or must all plans be repaired?

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

- Optimal behavior for non-stationary ad hoc agents
  - Repair one minimal plan?
  - Repair all minimal plans?
  - Repair all plans?
- General case of non-stationary agents

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?