Influencing a Flock via Ad Hoc Teamwork

Katie Genter and Peter Stone

The University of Texas at Austin
{katie,pstone}@cs.utexas.edu

September 12, 2014

UT Austin Learning Agents Research Group

Background

- Artificial intelligence
  - Multi-agent systems
    - Teamwork
      - Ad hoc teamwork
    - Swarm behavior
      - Flocking

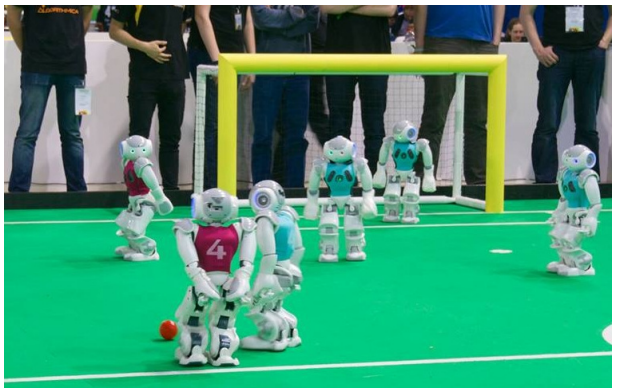

Ad Hoc Teamwork

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

UT Austin Learning Agents Research Group

Flocking

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

UT Austin Learning Agents Research Group

Example — Leading Teammates in Ad Hoc Settings

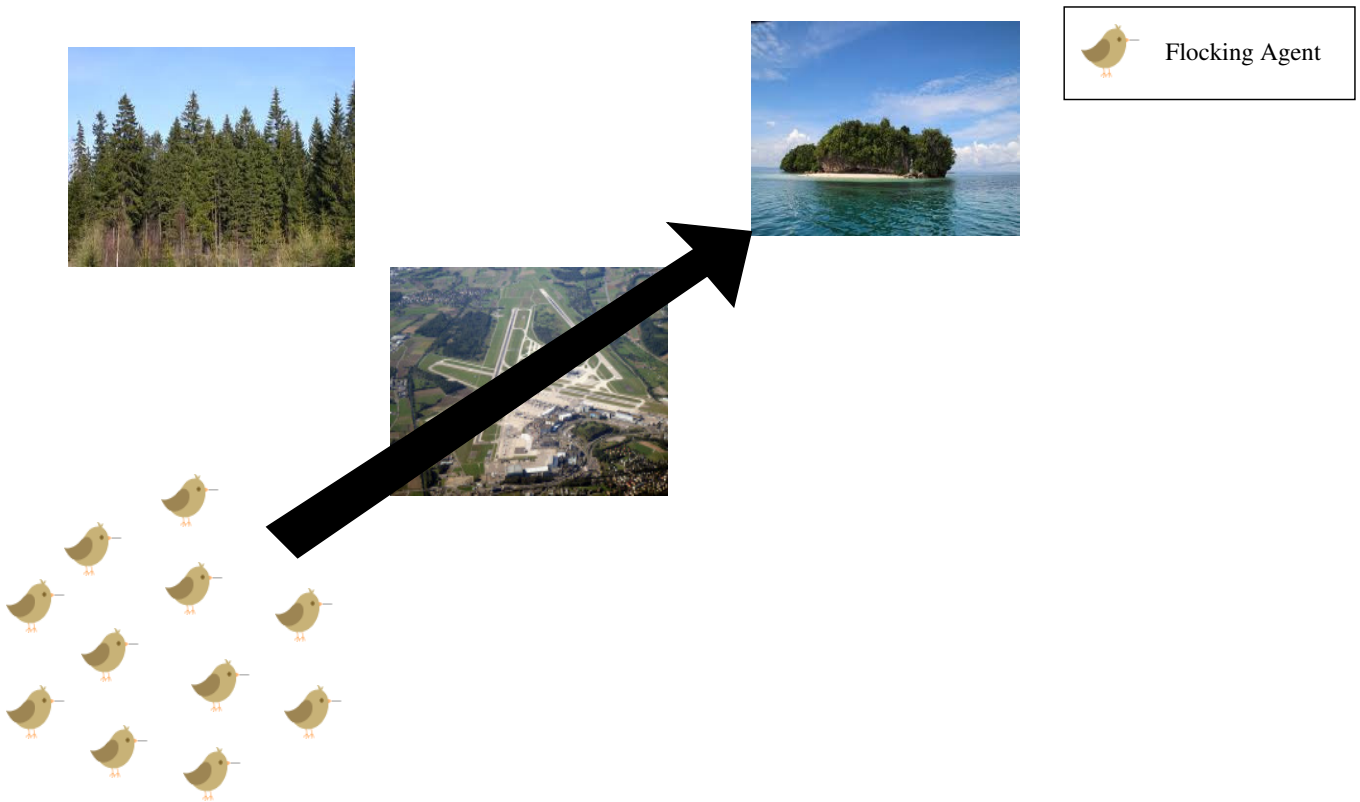

Example — Leading Teammates in Ad Hoc Settings

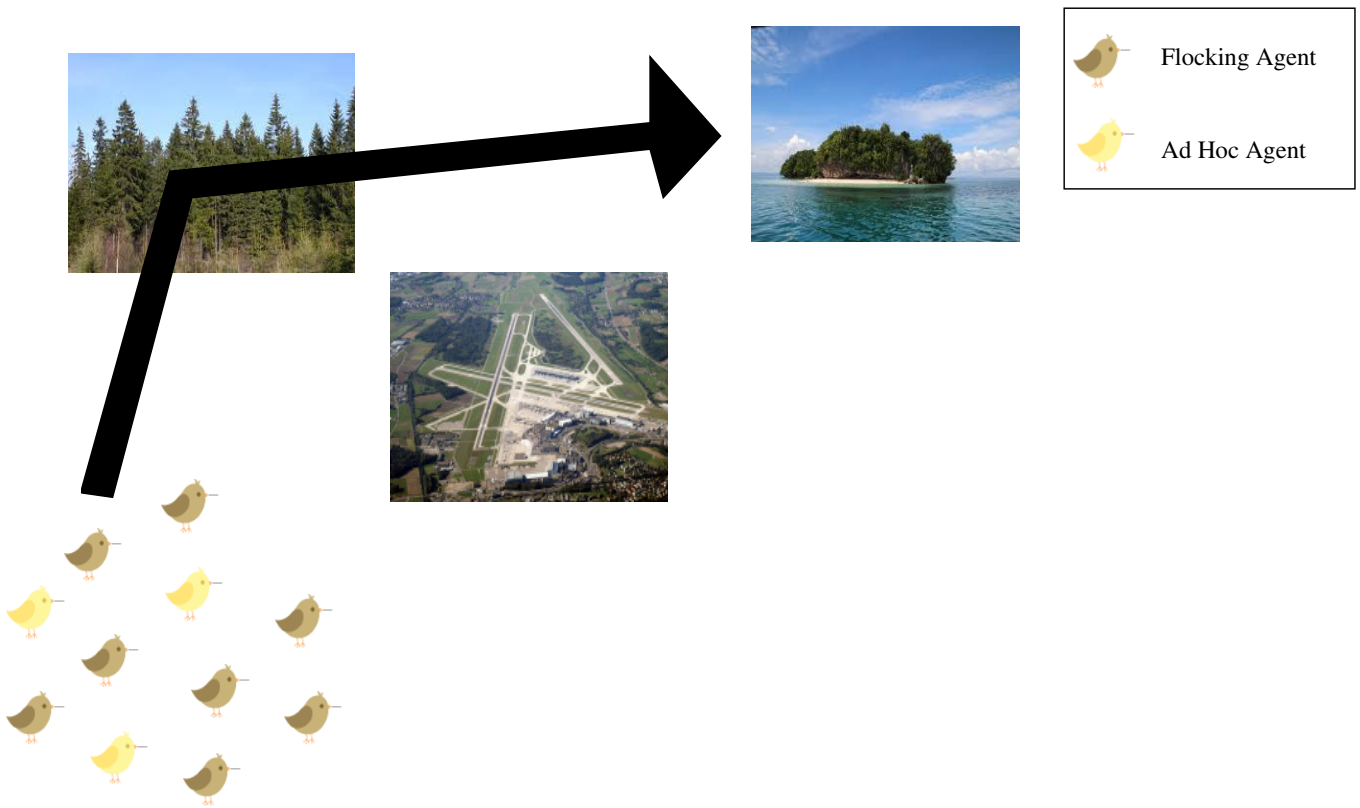

Example — Leading Teammates in Ad Hoc Settings

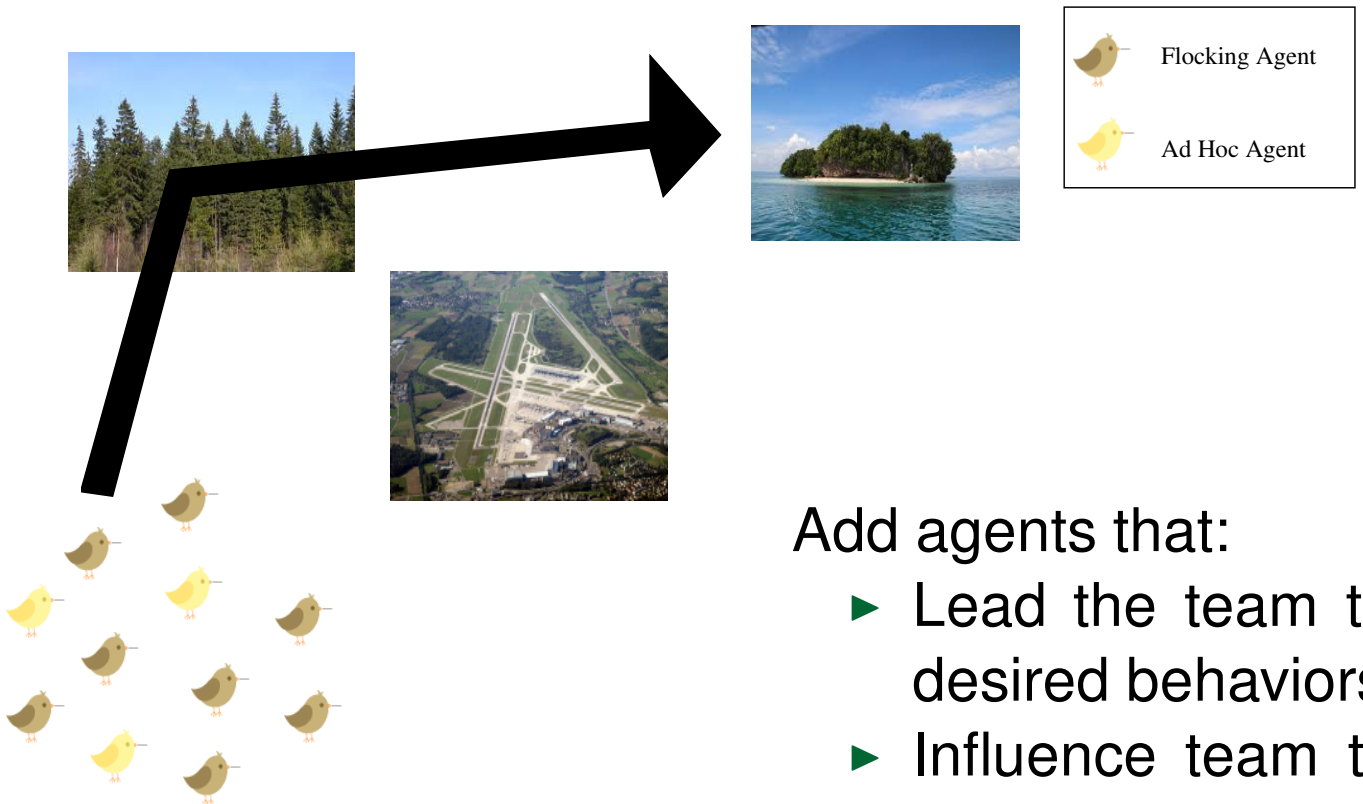

Add agents that:

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

UT Austin Learning Agents Research Group

Related Work — Ad Hoc Teamwork

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

Related Work — Flocking (1)

- Reynolds 1987, Vicsek et al. 1995
  - Concerned with simulating flock behavior
  - Not concerned with adding controllable agents to the flock
- Turgut et al. 2008
  - Considered the behavioral effects of providing different information to the flock
- Jadbabaie et al. 2003, Su et al. 2009, Celikkanat and Sahin 2010
  - Used controllable agents to influence the flock
  - Only concerned with making the flock converge to some heading eventually

Related Work — Flocking (2)

- Couzin et al. 2005
  - Considered how grouping animals make informed unanimous decisions
- Cucker and Huepe 2008, Ferrante et al. 2010, Yu et al. 2010
  - Used informed agents to influence flock
    - Behave in a fixed way that is predetermined on based on type
- 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

Outline

Problem Definition

Each agent has:

- Constant, non-zero velocity
- 2D Position
- Global heading

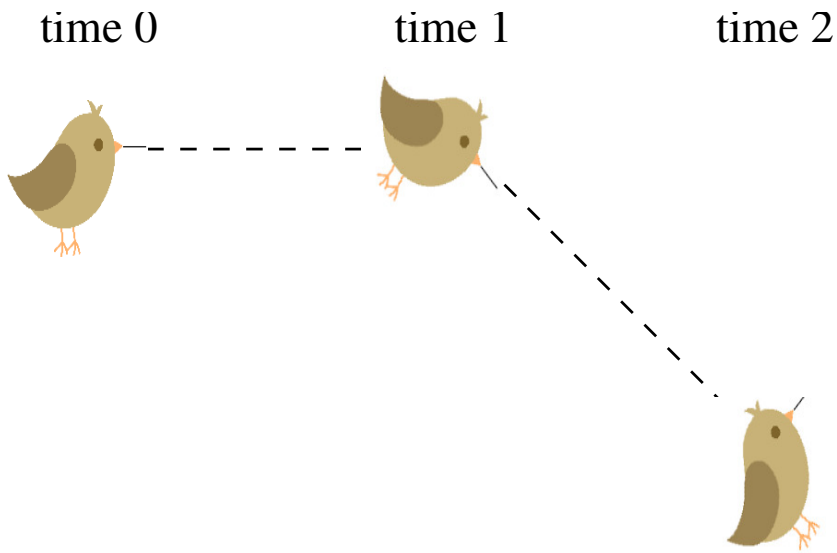

UT Austin Learning Agents Research Group

Problem Definition - Neighborhood

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

- Characterized by a radius in this work

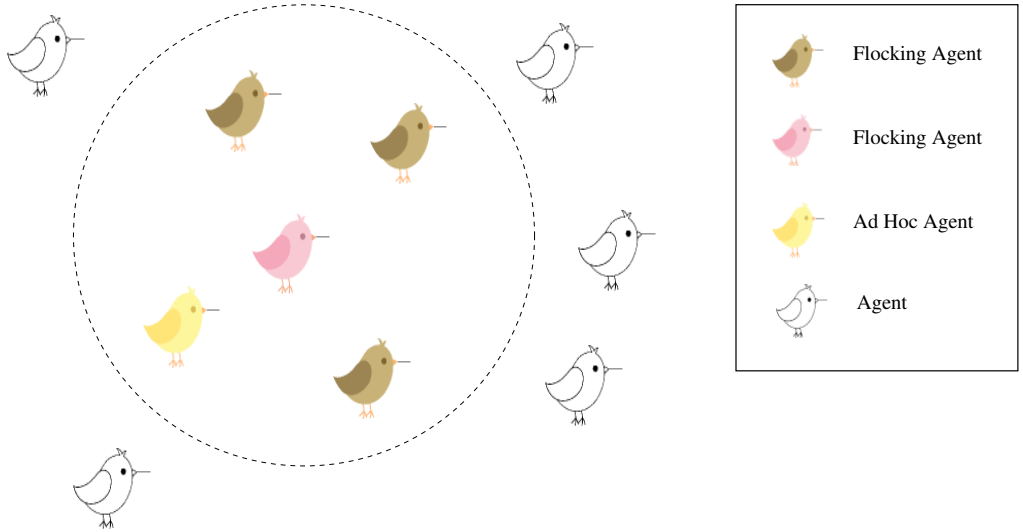

UT Austin Learning Agents Research Group

Problem Definition - Heading Update

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

- We only consider the Alignment aspect of Reynolds' model
- Agent can turn any amount instantaneously (not fully realistic)

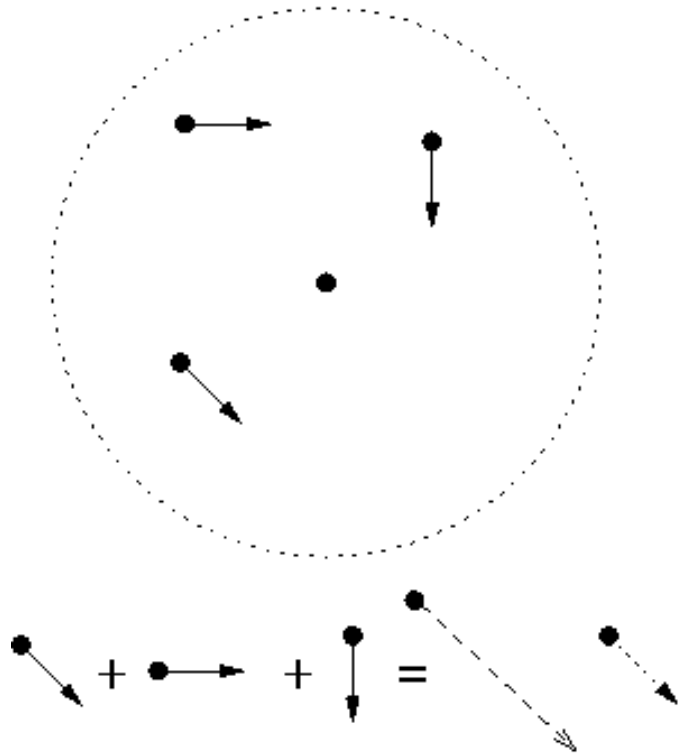

Research Questions

Research Problem:

How should ad hoc agents behave so as to:

- *orient* the rest of the flock towards a target heading as quickly as possible
- *herd* the rest of the flock through turns quickly but without compromising the composition of the flock

Outline

1-Step Lookahead Behavior

- Each ad hoc agent determines the best heading to adopt at each time step
  - 'Best' is the behavior that will exert the most influence on the next time step
- Considers all influences on each neighbor of the ad hoc agent

1-Step Lookahead Behavior

- For each potential ad hoc agent heading, consider how each of the neighbors of the ad hoc agent will be influenced
  - Consider each neighbor of each neighbor of the ad hoc agent
- Pick the heading that results in the least difference between the goal heading and the neighbors' new headings

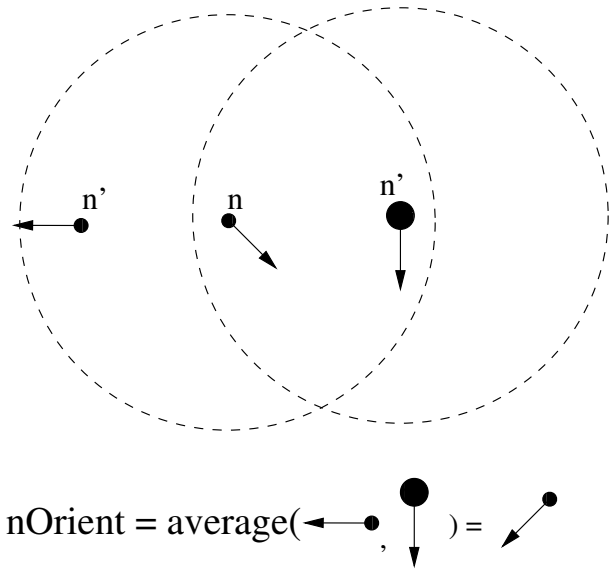

Outline

Research Questions

How should ad hoc agents behave so as to:

- *orient* the rest of the flock towards a target heading as quickly as possible
- *herd* the rest of the flock through turns quickly but without compromising the composition of the flock

Experimental Setup

Orient, Before

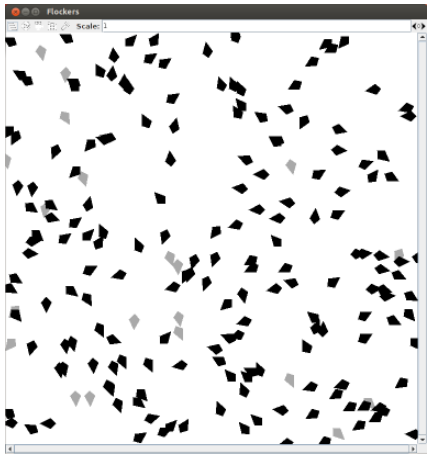

Orient, After

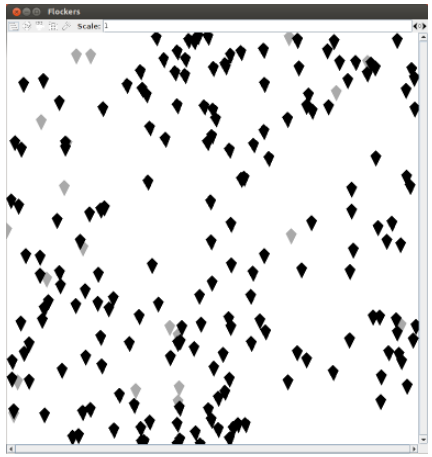

Herd, Before

Herd, After

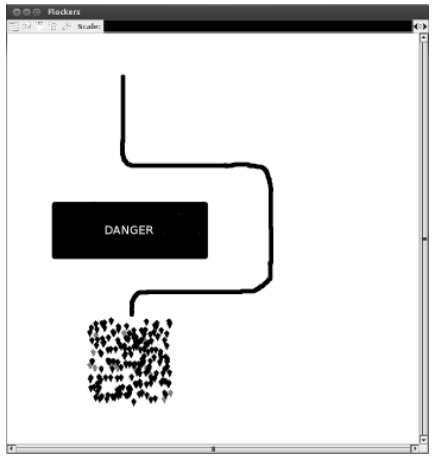

UT Austin Learning Agents Research Group

Baseline Behavior

Face Desired Orientation Behavior

- Ad hoc agents always orient towards the desired orientation vector
- Inspired by Jadbabaie, Lin, and Morse (2003)

Orient Experimental Results

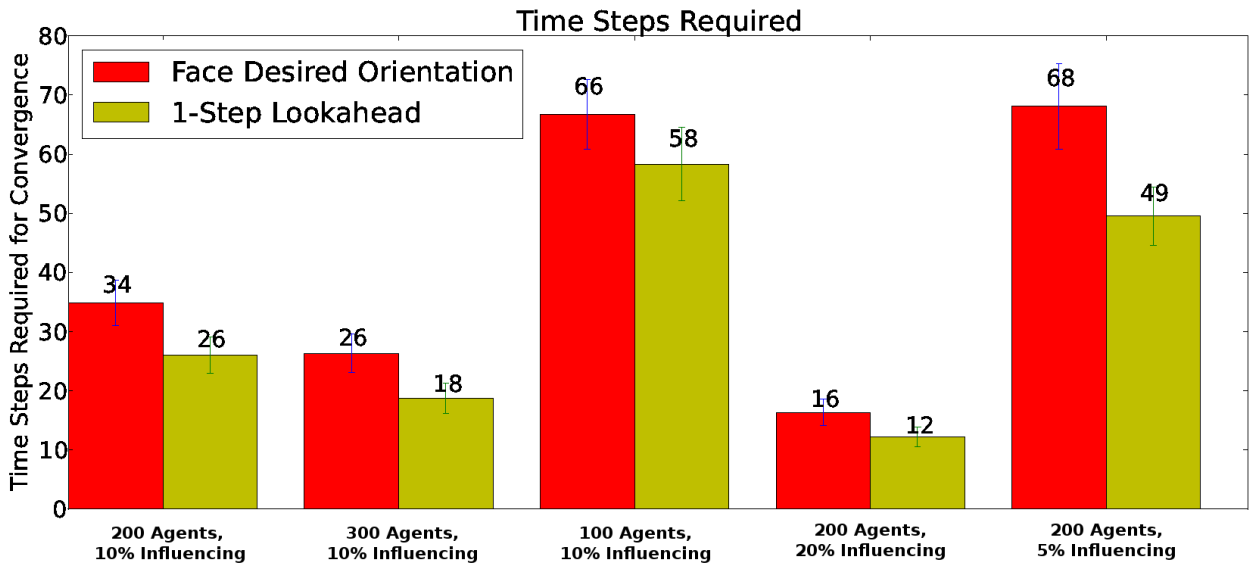

The results shown in this figure are averaged over 50 trials and the error bars represent the 95% confidence interval.

Orient Video - Flock Size 200, Ad Hoc Percent 10%

(Loading Video...)

UT Austin Learning Agents Research Group

Orient Video - Flock Size 200, Ad Hoc Percent 5%

(Loading Video...)

Orient Video - Flock Size 200, Ad Hoc Percent 20%

(Loading Video...)

Orient Video - Flock Size 100, Ad Hoc Percent 10%

(Loading Video...)

UT Austin Learning Agents Research Group

Orient Video - Flock Size 300, Ad Hoc Percent 10%

(Loading Video...)

UT Austin Learning Agents Research Group

Herd Experimental Results

	Steps-Converge	Steps-Optimal	Diff
10 Steps to Turn - Baseline	1243.0 (4.6)	1205	38.0
50 Steps to Turn - Baseline	1245.8 (2.2)	1225	20.8
100 Steps to Turn - Baseline	1261.0 (1.6)	1250	11.0
200 Steps to Turn - Baseline	1301.9 (1.0)	1300	1.9
10 Steps to Turn - 1-Step Lookahead	1237.0 (5.4)	1205	32.0
50 Steps to Turn - 1-Step Lookahead	1238.6 (3.0)	1225	13.6
100 Steps to Turn - 1-Step Lookahead	1254.5 (1.3)	1250	4.5
200 Steps to Turn - 1-Step Lookahead	1300.6 (0.6)	1300	0.6

These results are averaged over 100 trials. The numbers in parentheses show the 95% confidence interval.

Herd Video - 10 Steps to Turn

(Loading Video...)

Herd Video - 50 Steps to Turn

(Loading Video...)

Herd Video - 200 Steps to Turn

(Loading Video...)

Ongoing Research

- Other types of algorithms for ad hoc agents
  - Deeper lookahead searches
  - Coordination between ad hoc agents
- Extend to other interaction models
  - Consider flock separation and cohesion when calculating the next heading

Summary

Research Problem:

How should ad hoc agents behave so as to:

- *orient* the rest of the flock towards a target heading as quickly as possible
- *herd* the rest of the flock through turns quickly but without compromising the composition of the flock

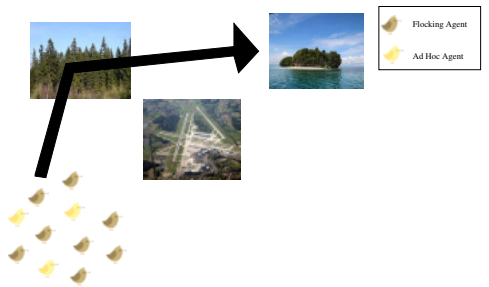

UT Austin Learning Agents Research Group