Negotiated Formations

David J. Naffin and Gaurav S. Sukhatme
$dnaffin|gaurav@robotics.usc.edu$
Robotic Embedded Systems Laboratory
Department of Computer Science
Center for Robotics and Embedded Systems
University of Southern California
Los Angeles, CA 90089-0781

Abstract

We present a decentralized, behavior-based approach to assembling and maintaining robot formations. Our approach dynamically grows formations from single robots into line segments and ultimately larger and more complex formations. Formation growth occurs by negotiation when robots are in close proximity. The approach has been validated using simulation as well as experiments with physical robots. [1]

1 Introduction

We are interested in studying the problem of assembling and maintaining formations of robots. Why study formations? Besides the challenge of accurately controlling multiple robots, there is the added benefit of having the coordination of multiple sets of sensors. Nature favors animals that have the ability to form formations such as flocks of birds or schools of fish. Animals that can combine their sensing abilities have shown to better avoid predators and efficiently forage for food [10]. Both the Air Force and NASA have identified autonomous formation of spacecraft as key technological milestones for the twenty first century. Applications of space based autonomous vehicles range from ground surveillance to interferometry experiments [9].

Our goal is to have multiple robots organize and maintain simple geometric formations without centralized coordination and using only local sensing. We have designed a bottom-up approach to *growing* formations. Singletons, individual robots that are not part of an existing formation, negotiate with other singletons to form line segments. In turn, these segments continue to negotiate with other sets forming larger and more complicated formations.

Our approach to this problem differs from others in that each robot's role as well as its position in the formation is not determined a priori. All robots begin as singletons and negotiate their roles with others based on their current state. A robot with a positional advantage may change its role from a follower to formation leader. Our approach minimizes the use of global state information.

Formation trials of up to twenty robots were simulated while four Active Media Pioneer 2-DX mobile robots were used as a physical test-bed. Experiments were performed with formations based on simple line segment geometries; lines, columns, wedges and diamonds (Figure 1). The metrics used to study the results of the experiments include the time to complete the formation and the percentage of time in formation.

Our results show that robots equipped with holonomic drives were able to coalesce into formations quicker then those without holonomic drives. With the exception of the line abreast formation, holonomic equipped robots fared better at maintaining their formations. In addition, the time required to coalesce depends on the size of the arena.

2 Related Work

Early work on leader-follower control strategies for formations was reported in [12]. [3] demonstrated distributed control for maintaining robot formations. [1] demonstrated

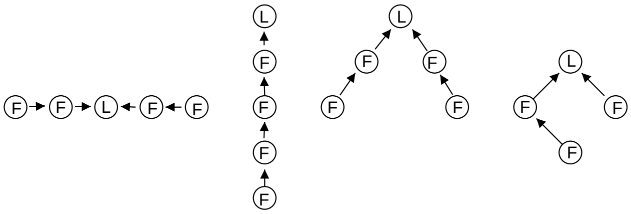

Figure 1: Line, Column, Wedge and Diamond Formations

---

[1] Submitted to the *2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-03)* Las Vegas, NV October 27-31, 2003

behavior-based control on physical robots that could negotiate obstacles. They were the first to adopt a set of metrics for formation evaluation as well as defining the various formation position strategies (i.e. unit-center, leader or neighbor-referenced). More recently, [5] developed a behavior-based approach for formations where each robot is designated a unique friend robot to follow by visual servoing. A conductor robot that is viewed by all followers maintains the overall heading of the formation. Our approach differs from [5] in that there are no preassigned orderings to members of the formation. While there is a robot which performs the role of the formation's leader, as a formation dynamically organizes, the robot that performs this role can (and does) change.

In [2] the authors describe a behavior-based approach to formations based on potential fields. Separate motor schemas compute a vector for moving to the proper formation position, avoiding static obstacles, avoiding other robots and maintaining the current formation. Each robot builds a list of potential attachment sites an generates an attraction vector for the closest. Our approach is similar to [2] in that formations are grown from simpler arrangements. However our approach does not used fixed attachment sites for each robot or predetermine a robot's position in the formation.

[4] describes a framework for vision-based control of formations using a single omni-directional camera. Their approach emphasizes the ability to switch between simple decentralized controllers and therefore change their overall formation. A control graph describes the relationships between the leader-followers. In most of the experiments the control graph is static and defined a priori. In a few cases they allow the control graph to vary but within rigid limits. In our approach, the control graph for our formations is implicit and distributed across all the robots involved in the formation. One advantage to the approach in [4] is the ability to prove formation stability over a range of external inputs.

In [13] the authors describe a virtual structure approach to formations. Their approach incorporates formation feedback where by the formation leader receives feedback from its followers. Their approach allows for stability guarantees, and was validated in simulation as well as on physical robots. The emphasis was on maintaining the formation as well as minimizing formation errors. The robot's roles were determined a priori and were static.

In the area of Self-Assembling structures, [8] describes a water-bug model whereby simple capillary forces are used in a open loop control to assemble arbitrary structures. Emphasis was placed on eliminating defects in the resulting structure as well as terminating the construction. Although up to now their work has been purely theoretical, they have proposed using simple robots as a testbed.

3 Approach

Our approach to formation control is to dynamically grow a formation from simpler constituents. *Singletons*, wandering individual robots that are neither followed nor follow others, are the initial state of all robots. As singletons encounter other singletons, they negotiate via a simple broadcast protocol. The outcome of these negotiations is typically a role change for both robots involved. Two singletons leave the negotiating phase as a *leader-follower* pair. The leader-follower pair represents a line segment and is the formation molecule (i.e. the smallest unit of a formation). As leader-follower pairs encounter other leader-follower pairs (or singletons) they negotiate and coalesce into larger formations. For a given formation, there is only one leader. All other singletons and leaders negotiate with the formation's leader to gain membership in the formation.

Followers simply maintain a fixed position relative to a single target robot. Their target is communicated to them by their current formation leader and can be updated at any time. Followers do not participate in any negotiations. A follower's target can be its leader or any other robot in the formation. Often sets of followers will form long line segments. Frequently, while avoiding obstacles and other robots, a follower will lose track of its target. If the target is not re-acquired within a suitable time frame the follower will revert into a singleton and again begin its search for a suitable formation to join.

This bottom up approach has several merits. Although there is a leader robot, any robot can perform this role. In order for a robot to perform any and all roles, each controller must be identical and have access to all behaviors. This decentralized approach adds a element of robustness to the system. The approach requires each robot to sense only its nearest neighbors. No global sensing is required. There are no a priori robot ordering or predetermined organizational strategies. Finally, this approach has been shown to scale to larger numbers of robots.

3.1 Architecture

The controller on each robot consists of a collection of layered behaviors (see Figure 1). At the top of this layering are the *coordinating behaviors*. They are responsible for providing overall coordination, formation bookkeeping, communication and management of all other behaviors. Typically these behaviors are implemented as finite state machines. At any given moment a robot can have only one active coordinating behavior corresponding to its state. Currently we have developed a set of three coordinating behaviors:

Singleton This behavior is active on a robot when it does not belong to any formation. In this state the robot will

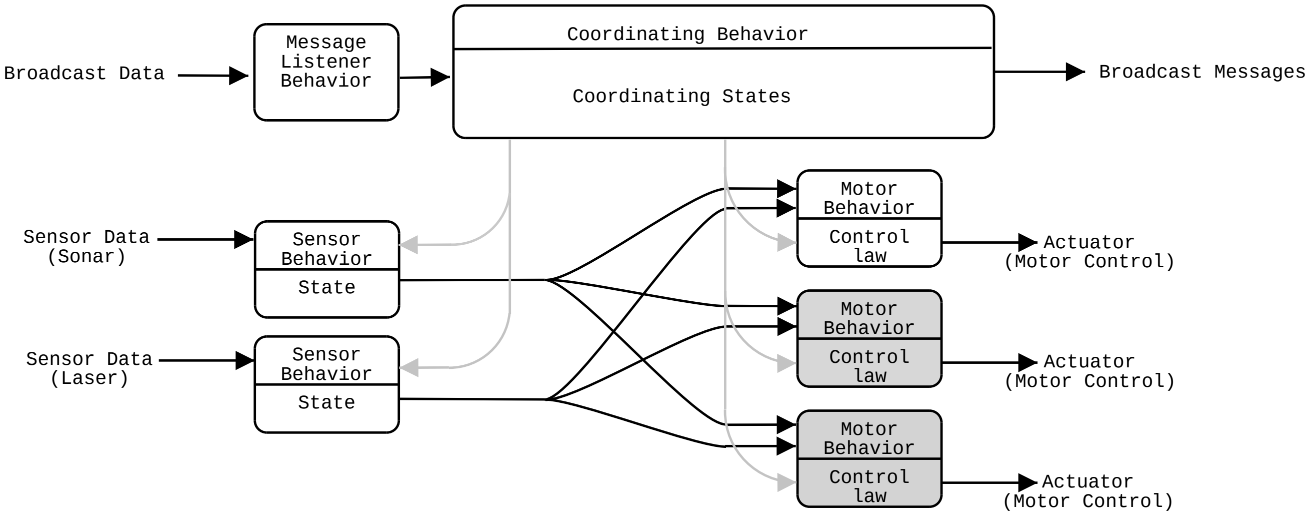

Figure 2: Basic configuration of behaviors on each robot

search for other singletons or join an existing formation. Since all robots start as singletons, when the system is initialized, the coordinating behavior on each robot is the singleton behavior.

Leader This behavior represents the formation leader. A formation may be as few as two or as many as a hundred robots. Any given formation has a single leader. This behavior is the coordinating behavior on the leader robot, and maintains the formation's heading and bookkeeping information (formation geometry, size, angles, etc). It is the overall responsibility of this behavior to gather enough followers to correctly assemble the formation.

Follower This behavior is active on a robot when it is part of a formation yet is not a leader. Its primary function is to maintain a fixed bearing and distance to one (or possibly two) targets. In a few special cases, a follower's target will be its leader. However, most followers will simply target their nearest neighbor.

All lower-level behaviors can be classified into three categories:

Sensor Behaviors These behaviors typically poll a raw sensor such as a camera or laser range finder and update an internal state that can later be queried by other behaviors.

Motor Behaviors These behaviors implement the control policy for the robot's actuators. At any given moment only a single motor behavior will be active. These behaviors implement simple Proportional (P) controllers.

Communication Behavior Similar to the sensor behaviors, this behavior polls a simple broadcast messaging system. It filters messages received but not intended for it. It updates the internal states of other behaviors that subscribe for certain messages.

Bookkeeping Behavior Used by the leader behaviors to maintain the global state of the formation. Each formation geometry has a different bookkeeping behavior. Only a single bookkeeping behavior is needed for the entire formation. These behaviors are responsible for implementing the discovery protocol (see section 3.4) that is periodically communicated between a leader and its followers.

A complete list of behaviors used is shown in Table 1.

3.2 Operation

Initially several singletons are released into an arena and wander about searching for one another. Each is initialized with the final goal, to create a formation of a given size and shape. Upon recognizing another robot's identifier via the `detect-target` behavior, a singleton broadcasts a message to the newly discovered partner asking it to begin negotiations. If the other robot is also a singleton, than the leader of the pair will be chosen as the robot who is at positional advantage to lead. Due to the sensor arrangement, it is not possible for two robots to simultaneously sense each others beacons. When two singletons encounter each other, only one robot will see the others' beacon (discussed in 3.3), and therefore it has a positional advantage to follow the other. Likewise the singleton which can not see its partner is at a positional advantage to lead. At this point the

Behavior	Type	Description
`singleton`	Coordinating	Single wandering robot
`follower`	Coordinating	A follower in the formation
`leader`	Coordinating	Formation leader
`message-listener`	Communicating	Receive broadcast messages and update the internal states of other behaviors
`follow-target`	Motor	Maintain relative position to target
`investigate`	Motor	Close gap between robot ans a potential target
`wander`	Motor	Wander about an area while avoiding obstacles
`locate`	Motor	Move toward a potential target location
`spin`	Motor	Rotate in place
`detect-obstacles`	Sensor	Report on obstacle bearings
`detect-target`	Sensor	Report bearings to all potential targets in view
`track-target`	Sensor	Report on the target's range, bearing and orientation
`line`	Bookkeeping	Maintain a line segment formation
`column`	Bookkeeping	Maintain a column segment formation
`wedge`	Bookkeeping	Maintain a wedge shaped formation
`diamond`	Bookkeeping	Maintain a diamond shaped formation

Table 1: Summary of Behaviors

coordinating behaviors on both robots will change accordingly. If the newly discovered robot is not a singleton but a leader of an existing incomplete formation then the singleton is instructed to join the formation and assigned a target to follow as well as a desired range and bearing to keep. If the newly discovered robot is a follower in an existing formation, then it informs the singleton of the broadcast address of its leader to negotiate with.

After a while, there will be fewer and fewer singletons available to join formations. Leaders of incomplete formations will continue to search for new recruits. If a leader discovers another leader, then the two negotiate for overall leadership of the collection. The leader with the positional advantage becomes the new leader. All followers associated with the losing leader will be instructed to join the new formation and instructed as to what target, range and bearing to keep.

In all cases of negotiation, the robots involved slow their forward velocities. This is to aid the newly acquired followers in properly positioning themselves. If a new designated follower is not able to see its proper target, then a message is sent to its new leader asking for a new target. If no proper target is within its range of view, then the robot reverts to a singleton and begins searching again. Only after all newly acquired followers have acknowledged finding their designated targets and properly positioning themselves, will the leader accept them as part of the formation and update its formation bookkeeping. Reverting back to a singleton can also occur for any follower that somehow loses view of its target.

3.3 Choice of Sensors and the Type of Formation Geometries

In order to assemble and maintain a formation, each robot must have the capability to avoid obstacles, distinguish other robots from the environment and determine the bearing, range and heading of other robots relative to itself. The capability of the chosen sensors directly affect the possible formation geometries. For any given position in the formation, there must be a neighbor that is within the robot's view in order to track it. For our purposes, we decided on a laser range finder as our mode of sensing. The SICK LMS-200 laser range-finder can accurately measure the distance to any object within a hemispherical region in front of the unit within a radius of eight meters. The device is also able to distinguish certain highly reflective materials from other surfaces. We constructed beacons from strips of this material attached to cardboard backings. These beacons are attached to robots to distinguish them from other objects in the environment. In addition, by creating binary patterns of these reflective strips, the laser can uniquely identify them as well. By positioning these beacons on the back of each robot, we are able to distinguish another robot from the environment as well as measure its relative distance and bearing. Since the laser range-finder can measure distance within its range to an accuracy of less than a centimeter, we are not only able to measure the distance and bearing to a beacon, but by measuring the distance of both ends of the beacon, we are able to determine the beacon's relative orientation as well. However, accurate beacon identification can only be accomplished within approximately two meters of the beacon.

The type of geometries possible with a forward-looking mode of sensing include those that are connected, are constructed of line segments, and do not require any backward-looking sensing (often referred to as frontally concave [5]). Examples of these shapes are diamonds, triangles, hexagons and wedges. Many of these restrictions are artifacts of the mode of sensing and not our approach. Using multiple laser range-finders or an omni-directional camera would remove the backward-looking sensing restriction. However, in this paper, we concentrate on the forward-looking sensing mode.

Note that the capability to sense other robot's heading is not absolutely necessary. All that it needed is a means for each robot to maintain the same heading as its leader. It is possible for the leader to broadcast its heading changes to its followers. However, this requires that each robot have some means of determining its global orientation (i.e. compass, odometry, a fixed beacon observable by all, etc).

3.4 Discovery Protocol, Balancing and Inflections

Periodically, a leader's bookkeeping behavior will verify the state of the formation. This is accomplished by a sequence of broadcast messages. The sequence begins by broadcasting a *Find-Target* message to all followers. This message requests that any follower who is currently targeting a specified robot respond back to the sender. The leader broadcasts this message specifying itself as the target. One (or possibly two) followers should respond to this message. The leader then broadcasts one (or two) more *Find-Target* messages using the identifiers of the respondents as the new target(s). This continues until the end points are identified and no followers respond to the *Find-Target* messages. The responses are used to build up the formation's global state, number of followers, followers ordering, and membership list. Figure 4 shows an example wedge formation with a leader (ID 5) and four followers. The resultant responses provide the leader with necessary information needed to build two sets of list of followers, those on the left wing of the formation and those on the right. Each time the discovery protocol begins, all previous global state information is deleted. This way, the current state only depends on the last discovery iteration. Should a leader change roles, it is not necessary that the current state of the system transfer to the new leader. When a new leader is initiated, one of its first tasks is to perform a discovery sequence.

If a formation consists of more than one line segment (for example the wedge or the diamond formations) then often the segments will not grow evenly. This situation can be detected by the leader's bookkeeping behavior and a *balancing* maneuver is initiated. The followers involved are only those physically nearest the leader. This avoids complicated planning of routes between positions within the formation. For any given balancing maneuver, one or possible two followers will have to change their targets, while one follower will only change its bearing to its target (see the example in figure 5). This maneuver is accomplished by issuing a *New-Target* message to those followers involved. This maneuver is performed only once per discovery sequence and will continue until the formation is balanced.

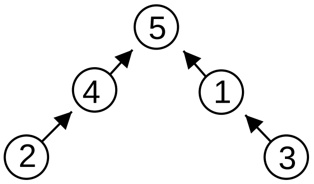

Broadcast #1: Find-Target (5)
Response #1: Follower(4)
Response #2: Follower(1)
Broadcast#2: Find-Target(4)
Broadcast#3: Find-Target(1)
Response #3: Follower(2)
Reposnse #4: Follower(3)
Broadcasr #4: Find-Target(2)
Broadcast #5: Find-Target(3)
No more responses

Figure 3: Example of Discovery Protocol

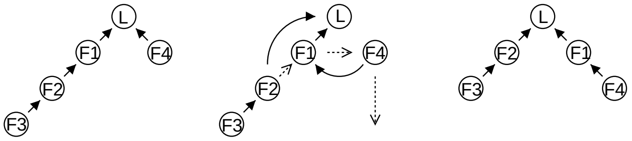

Figure 4: Balance Maneuver

Some formations require multiple line segments where a few of the segments do not have the leader as an endpoint (for example the diamond formation). However, in many these formations, additional segments can be viewed as kinks or inflections in a chain. As the bookkeeping behavior builds its graph of its current formation, if an inflection is needed to maintain the proper geometry, a *New-Bearing* message will be sent to the appropriate followers.

4 Experiments

Experiments were performed in simulation using the Stage simulator [11] as well as on Active Media's Pioneer 2-DX robots using Player [6]. Each robot (simulated or physical) was outfitted with a laser range-finder and a unique laser beacon.

4.1 Experimental Methods

The simulations were performed using five, ten and twenty robots in a thirty-five by twenty-five meter rectangular arena. There were no additional fixed obstacles. Each robot travels at standard velocity of 30 cm/sec and a maximum velocity of 40 cm/sec. Each robot was equipped with a wireless broadcast device, a forward looking laser range-finder, a unique laser beacon and a 360 degree ring of sonar range-finders for use in obstacle avoidance. Each simulation trial lasted sixty minutes.

Similarly, all the physical robot experiments were performed using four Pioneer 2-DX in a ten by five meter arena. Each robot travels at a standard velocity of 20 cm/sec and with a maximum velocity of 40 cm/sec. Each was equipped with wireless Ethernet, a SICK LMS-200 laser range-finder, a unique laser beacon and a 360 degree ring of sonar range-finders. Each trial lasted thirty minutes.

The stage simulator can simulate robots with and without holonomic drive capabilities. Robots equipped with holonomic drives have a decisive advantage at maintaining a bearing to a target, especially if the robot needs to maintain a fixed heading as well. Unfortunately we did not have access to physical robots with this capability and therefore all the holonomic robot results are in simulations only.

The wedge and diamond formations were performed using an inter-robot angle of 90 degrees (i.e each robot maintained a 45 degree bearing to its targets). Due to the look-ahead requirement of our mode of sending, the line formation had to be a staggered line, that is, each target robot needed to be a certain minimal distance ahead of its follower to see and correctly identify the target's laser beacon. For the simulations this minimal distance was 20 centimeters. For the physical robots this value was 40 centimeters.

4.2 Metrics

We selected three metrics to measure the performance of our system.

Positional Error: Given a formation, which is defined as the tuple $<\ \mathcal{G}, h, d\ >$ where $\mathcal{G}$ is a connected geometric shape,*h* is a desired heading, and *d* is a desired inter-robot spacing, there exist *K* positions relative to the leader that represent the perfect formation. Given *N* robots attempting to construct this given formation, where $N \leq K$, we define the formation's positional error as $\mathcal{P} = \frac{1}{N}\sum_{i=1}^{N} \mid D(p_i, k_i) \mid$ where p_i is the *ith* robot's position (relative to the leader) and $D(p_i, k_i)$ is the euclidean distance between the two positions. We say that a given set of robots is "in formation" if $\mathcal{P} < \epsilon$, where ϵ is a user-specified tolerance. We used an $\epsilon = 0.33$ meters in simulation and a $\epsilon = 0.5$ meters for the physical robot trials.

Time to convergence: The time to convergence $\mathcal{T}_c(N)$ is defined as the duration of time required for a formation to reach a given size *N* and be in formation for that size. The clock starts when the current leader is elected and ends when the formation has grown to *N* followers that have stabilized into their in formation positions. For the physical robot experiments, these times are wall clock times. However, for the simulations the times are normalized to update cycles of the simulator.

Percentage of time in Formation: Once a formation has reached its current maximum size and settled into the in formation state, the leader's wandering will cause the positional error of the formation to occasionally increase (such as when making turns to avoid obstacles). Often the added positional error will result in the formation being broken. Therefore, the percentage of time in formation is defined as $\mathcal{F} = \frac{t_{in}}{t_{total}}$ where t_{in} is the time in formation since the formation reached its current size and t_{total} is the time elapsed since the formation reached its current size.

In order to measure system performance, the global positions for each robot are needed. For the simulations, this was easily obtained directly from the Stage simulator. For indoor experimentation, we used the mezzanine tool [7]; a system for gathering pose information for a set of robots in a limited size arena. Mezzanine uses a single overhead calibrated camera to detect color fiducial attached to the tops of each robot. The tool can uniquely identify each robot and report its position and global orientation.

4.3 Experimental results

Table 2 summarizes the holomonic simulations. The column formation had the least overall error. Since each follower simply had to maintain a zero degree bearing (i.e. straight ahead) to it target, these results were expected. Diamond and wedge formations had similar results. Only line formations had significantly different results. Table 3 summarizes the non-holonomic simulations. These results are similar to the holonomic results, but reveal that the holonomic drive equipped robots are better able to maintain their positions within the formation. Table 4 summarizes the physical robot trials. These results are similar to the non-holonomic results.

An unexpected result was that non-holonomic drive robots seems better able to maintain the line formation than those equipped with holonomic drives. Since each target in a line formation is just barely within the maximum viewing angle for the laser range-finders, a follower would momentarily lose sight of its target. This situation causes a switch of motor behaviors for the follower. The `locate` behavior for a holonomic drive would perform a lateral movement in an attempt to bring the target back into view, however this maneuver often is not necessary and will move the robot out of formation. Although the behavior is the same for a non-holonomic robots, it is much slower with a non-holonomic drive. Therefore, the target would often be re-acquired before the behavior had a chance to take it out of formation.

Another interesting aspect of the experiments, was the effect that the boundaries had on the resulting maximum formation sizes. Having a large arena results in longer

Trial No.	Geometry	N	Max	$\mathcal{T}_c$ (sec)	$\mathcal{P}_{min}$ (m)	$\mathcal{P}_{ave}$ (m)	$\mathcal{F}$ %
1	Column	5	5	102	0.142	0.308	81.83
2	Column	10	7	997	0.197	0.318	90.81
3	Column	20	9	1457	0.035	0.309	89.57
4	Line	5	4	125	0.524	1.124	21.22
5	Line	10	5	183	0.499	1.235	5.97
6	Line	20	6	610	0.518	1.347	3.30
7	Wedge	5	5	127	0.232	0.551	63.90
8	Wedge	10	6	147	0.192	0.566	23.02
9	Wedge	20	7	567	0.190	0.631	15.24
10	Diamond	5	4	110	0.116	0.398	71.89
11	Diamond	10	6	211	0.201	0.527	27.82
12	Diamond	20	8	1329	0.216	0.619	15.61

Table 2: Holonomic Simulation Results

Trial No.	Geometry	N	Max	$\mathcal{T}_c$ (sec)	$\mathcal{P}_{min}$ (m)	$\mathcal{P}_{ave}$ (m)	$\mathcal{F}$ %
1	Column	5	5	117	0.179	0.301	88.56
2	Column	10	7	843	0.162	0.330	84.84
3	Column	20	8	1189	0.196	0.342	79.57
4	Line	5	3	141	0.414	0.876	39.11
5	Line	10	5	190	0.443	0.723	8.37
6	Line	20	5	210	0.438	0.747	9.30
7	Wedge	5	5	188	0.289	0.652	69.11
8	Wedge	10	6	301	0.182	0.466	43.02
9	Wedge	20	7	885	0.290	0.481	35.27
10	Diamond	5	4	117	0.136	0.374	68.23
11	Diamond	10	6	288	0.266	0.382	35.53
12	Diamond	20	7	1927	0.295	0.549	21.67

Table 3: Non-Holonomic Simulation Results

Trial No.	Geometry	N	Max	$\mathcal{T}_c$ (sec)	$\mathcal{P}_{min}$ (m)	$\mathcal{P}_{ave}$ (m)	$\mathcal{F}$ %
1	Column	4	4	316	0.14	0.29	92.2
2	Column	4	3	182	0.12	0.27	98.7
3	Column	4	3	202	0.12	0.27	99.1
4	Line	4	2	65	0.31	0.66	39.1
5	Line	4	2	38	0.34	0.62	11.2
6	Line	4	3	241	0.33	0.74	22.2
7	Wedge	4	3	135	0.17	0.48	77.2
8	Wedge	4	3	142	0.15	0.44	74.0
9	Wedge	4	3	137	0.18	0.48	75.1
10	Diamond	4	3	147	0.19	0.39	81.2.
11	Diamond	4	3	137	0.26	0.41	77.3
12	Diamond	4	4	271	0.31	0.42	66.9

Table 4: Pioneer 2-DX Results

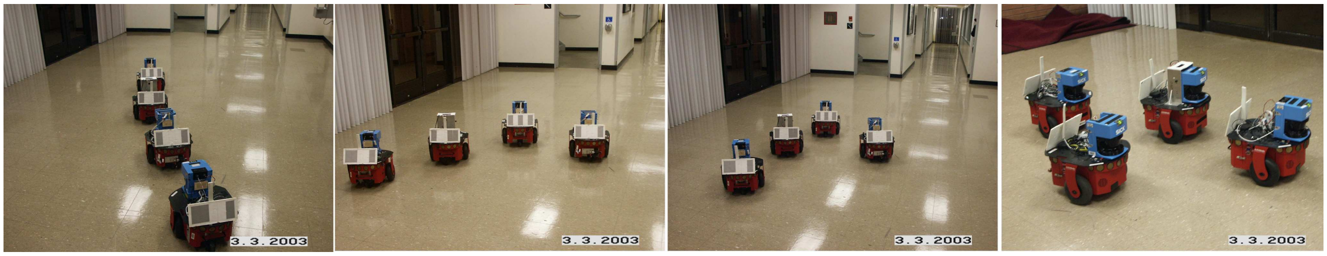

Figure 5: Pioneers in Formation: Column, Line, Wedge, Diamond

periods of time traversing the open spaces, therefore allowing the followers more opportunity to settle into their proper positions before being disturbed by a leader's heading change to avoid the boundary. However, the boundaries also act as a catalyst that forces singletons and smaller formations to interact. Choosing the ideal size for the arena is an important consideration. An excessively large arena means that the time required to coalesce into a formation may be too prohibitive. Choose a small arena, and there may be too much interference to allow larger formations to assemble.

5 Conclusion and Future Work

Our approach to robot formation control is based on growing formations from simpler configurations (i.e single robots, line segments, etc.) into more complicated formations. By using simple local broadcast communications, we are able to dynamically reconfigure each robot's role in formation as the formation grows. Our approach is decentralized and requires only a single local sensor for each robot.

In the future we plan on investigating dynamically changing formations as well as more complicated formations (webs, lattices, etc.). We plan on deriving models that describe the behavior of the formations. Using these models we expect to predict how the formations would coalesce and behave under various situations (i.e. static obstacles, interacting with other formations, etc) which can be verified by our approach.

6 Acknowledgments

The authors would like to thank Boyoon Jung, Maxim Batalin, Andrew Howard, and Brian Gerkey for all their assistance. This work is supported in part by the DARPA MARS program under grants DABT63-99-1-0015 and 5-39509-A (via UPenn), and by NSF CAREER program under grant IIS-0133947.

References

[1] T. Balch and R. Arkin. Behavior-based formation control for multi-robot teams. *IEEE Transactions on Robotics and Automation*, 14(6):926–939, December 1998.

[2] T. Balch and M. Hybinette. Social potentials for scalable multi-robot formations. In *IEEE International Conference on Robotics and Automation (ICRA '00)*, pages 73 – 80, April 2000.

[3] Q. Chen and J. Luh. Coordination and control of a group of small mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2315–2320, 1994.

[4] A. K. Das, R. Fierro, V.Kumar, J. Ostrowski, J. Spletzer, and C. Taylor. A framework for vision based formation control. *IEEE Transactions on Robotics and Automation*, 18(5):813–825, October 2002.

[5] J. Fredslund and M. J. Matarić. A general algorithm for robot formations using local sensing and minimal communications. *IEEE Transactions on Robotics and Automation*, 18(5):837–846, October 2002.

[6] B. Gerkey, R. T. Vaughan, K. Støy, A. Howard, G. S. Sukhatme, and M. J. Matarić. Most valuable player: A robot device server for distributed control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'01)*, pages 1226–1231, October 2001.

[7] A. Howard. Mezzanine user manual; version 0.00. Technical Report IRIS-01-416, Institute for Robotic and Intelligent Systems, Univrsity of Southern California, 2002.

[8] E. Klavins. Toward the control of self-assembling systems. In *Workshop on Control Problems in Robotics and Automation*, pages 153–168, December 2002.

[9] J. Lawton R. W. Beard and F. Hadaegh. A coordination architecture for spacecraft formation control. *IEEE Transactions on Control Systems Technology*, 9(6):777–790, November 2001.

[10] Craig Reynolds. Flocks, herds and schools: A distributive behavioral model. *Computer Graphics*, 21(4):25–34, 1987.

[11] R. T. Vaughan. Stage: A multiple robot simulator. Technical Report IRIS-00-393, Institute for Robotic and Intelligent Systems, University of Southern California, 2000.

[12] P. K. C. Wang. Navigation strategies for multiple autonomous robots moving in formation. *Journal of Robotic Systems*, 8(2):177–195, 1991.

[13] B. J. Young, R. W. Beard, and J. M. Kelsey. A control scheme for improving multi-vehicle formation maneuvers. In *American Control Conference*, pages 704–709, June 2001.