A General, Local Algorithm for Robot Formations

Jakob Fredslund, Maja J Matarić

Abstract—
We study the problem of achieving global behavior in a group of robots using only local sensing and interaction, in the context of formations. The goal is to have N mobile robots establish and maintain some predetermined geometric shape. We report results from extensive simulation experiments, and 40+ experiments with four physical robots, showing the viability of our approach. The key idea is that each robot keeps a single *friend* at a desired angle θ, using some appropriate sensor. By *panning* the sensor θ degrees, the goal becomes simply to center the friend in the sensor's field of view, for all formations. We also present a general analytical measure for evaluating formations and apply it to the position data from both simulated and real-world experiments. We used two surveillance lasers to track the physical robots to obtain real-world position data.

Keywords— **Multiple robot coordination, local sensing, robot formations.**

I. Introduction – What is the Problem?

WHEN a university marching band impresses spectators at a college football game, it is not only because it plays well; it is also because of the amazing coordination it demonstrates. Having a group of almost two hundred individuals, each of whom has very limited knowledge of the others' positions, display a global behavior like forming the letters 'USC' – while performing – is a great feat. Without a conductor with complete knowledge from a position in a small tower overlooking the group, who can give instructions and corrections by megaphone during practice sessions, it would have been an even greater task to produce this behavior. Imagine the conductor being one of the band members: down on the field, with no more knowledge or sensing capabilities than the next band member. A radically different strategy would be needed to avoid any mishaps – like accidentally spelling out 'UCLA'. With no global information available anywhere in the group, how could a global, coordinated behavior be accomplished?

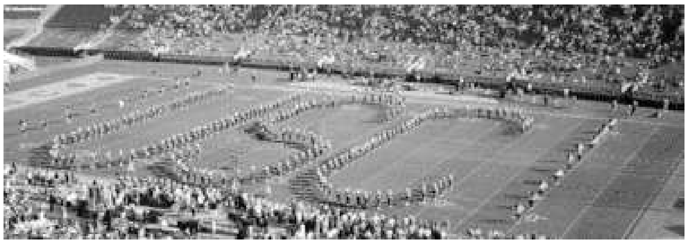

Fig. 1. The football marching band of USC.

In this work we have studied the problem of achieving global, coordinated behavior in a group of robots using only local sensing and interaction. With no centralized world view shared by the group, how can the robots cooperate to solve a task, which is not solvable by a single robot alone?

As an instance of this general problem, we have considered *formations*. We define formations (informally here, more formally in Section IV) to be groups of mobile robots establishing and maintaining some predetermined geometric shape while moving. This task clearly demands (spatial) coordination among robots, and it is an inherently 'global' task: a formation manifests itself at the level of the group as a whole. Robots need to have some notion of their respective roles in the formation, and they can not ignore the others. This is an example of a tightly-coupled coordination problem. We study how to achieve global-level formation coordination *without* providing the robots with global knowledge of other robots' positions or headings.

The rest of this paper is organized as follows. Section II gives the requirements to a solution algorithm for the formation problem. Section III gives an overview of related work. Section IV presents our algorithm, and section V describes an implementation of the algorithm. In section VI, we present and evaluate our experimental results, and finally section VII offers a discussion of our algorithm and some conclusions.

II. What is a Solution?

WHAT is required of a solution to the formation problem? First, a solution algorithm must get the robots into formation from their original positions, and then they must stay in that formation while moving. Consequently, a formal definition of the notion 'in formation' is needed, a measure by which to evaluate the algorithm. Such a measure would also allow for comparisons of different formation algorithms. We present a measure in Section IV.

It does not suffice to show that a group of robots can establish a particular formation and move through an unrealistically noise- and problem-free world. Thus, the second requirement of a solution to the formation problem is that it display the following characteristics:

1. *Generality* in the number of formations it can produce,
2. *Stability* of established formations,
3. *Robustness* to changes in group size,
4. *Dynamic Switching Capability* between formations,
5. *Obstacle Avoidance Capability* of maintaining or re-establishing formations in the presence of obstacles.

Robot behavior can be considered at two levels of abstraction: the algorithmic level, and the embodied system level. At the algorithmic level, only the algorithm itself needs to be evaluated. One can employ mathematical validation tools or make simulations to document the workings

Jakob Fredslund is with Department of Computer Science, Aarhus University, Denmark. Currently he is visiting Maja J Matarić's Interaction Lab at the University of Southern California, Computer Science Department, University Park, Los Angeles, CA 90089, USA.

of the algorithm. At the embodied system level, a physical robot executing the algorithm should be scrutinized; the robustness of the embodied algorithm is thus evaluated. Success on the algorithmic level does not imply success on the embodied system level, but it certainly increases the probability of it.

In summary, an algorithm that addresses all of these issues and performs robustly at the embodied system level, grounding itself in some formalized formation evaluation measure, solves the formation problem. We propose such a solution here, and validate it through extensive simulation and real robot trials.

III. Related Work

A VARIETY of approaches have been proposed to create global behavior in a group of mobile robots. [12] showed how a set of simple behaviors (avoidance, aggregation, and dispersion) based on local sensing only, can be combined so that a global flocking behavior emerges, and demonstrated it on a group of 13 mobile robots. In [18] the point is made that much can be achieved by fully exploiting minimal systems: a robot soccer playing team is demonstrated that has a minimalist, behavior-based control system with only a few basic behaviors. From the interaction of these, two different group formations of the three robots naturally emerge, enabling the team of seemingly 'willed' offensive and defensive team play. In [10], as little as two basic, local behaviors (avoidance and goal seeking) were proven by experiments with five physical robots to be enough to result in a successful collaborative box-pushing behavior. [19] presents work where a group of real robots, equipped with only touch sensors, are able to form a physical chain reaching from a nest to a source of food. Robots not in the chain are able to follow the chain to the source or back to the nest, carrying food. Minimalist communication through the chain is achieved by robots tapping each other's touch sensors. The global chain emerges from a set of basic behaviors that use only local sensing.

[2] presents theoretical work where a large set of robots, represented as points in the plane, congregate at a single position, regardless of their initial distribution. Moving synchronously in discrete time steps, robots iteratively observe the neighboring robots that are within some visibility range V, and follow simple update rules to find their next position.

With formations, however, a more rigid and reliable spatial structure is needed from the group of robots, typically involving fixed inter-robot distance and agreement on overall group heading. Each robot has to somehow determine its position relative to the position of (some of) its peers.

In [16], a group of simulated robots are able to form approximations to circles and simple polygons. The robots are assumed to have a sensor for determining the relative positions of all other robots, but there is no shared global coordinate system. Each robot follows simple geometrical rules, orienting itself in reference to, e.g., the robot that is furthest, and the one that is nearest to it. In [6], complete global knowledge is assumed; i.e., all robots' headings and positions are represented in a shared coordinate system and known by all. Again, each robot references itself to the furthest and nearest other robots. With a small set of simple update rules, the robots over time disperse uniformly and form a circle. By introducing positional and orientational constraints that each robot must satisfy with respect to a pre-selected reference robot, coordination of motion while in formation is also considered. An example of a matrix formation performing a right turn is given.

[20] shows how the shape of a group of mobile robots can be controlled using only local interaction between neighboring robots. Since all robots are connected in the shape, represented as a strongly connected graph, local interaction is distributed out and so eventually influences all robots in the group. The algorithm is demonstrated in simulation on a children's game example, where a chain of 'police robots' gradually fold around a 'thief' and catch him.

In [11], the desired formation is defined by a so-called *virtual structure*, or VS. Each robot is a priori assigned to a fixed point in the VS. The algorithm then iterates between fitting the VS to the current robot positions, displacing the VS in some desired direction, and moving the robots toward their new positions as prescribed by VS. Each robot uses the positions of all other robots in the calculation of its next movement. Both in simulation and with three physical robots, an instantiation of the algorithm is then documented, where the fitting and consequent displacing of the VS cleverly take into account that the robots are nonholonomic; i.e., some translations of the VS will give the robots a better chance of reaching their spots than others.

In [3], three different techniques for achieving and maintaining a given formation within a group of robots are identified: *unit-center-referenced*, *leader-referenced*, and *neighbor-referenced control*. In the first, each robot decides its position relative to the centroid of all robots; in the second, the robot uses the position of a predetermined leader; and in the third, the robot's nearest neighbor is used as a reference point. Each robot determines other robots' positions by dead reckoning, GPS, or by direct perception, and its coordinates in the global coordinate system are broadcast to all robots. This high reliance upon a centralized world view and the need to transmit coordinates between robots might have a negative impact on performance, as the paper states.

Successful simulated results are reported, and experiments with two holonomic Nomad150 robots were also carried out. Each of the two robots kept track of its position through odometry. Finally, a successful demo with four unmanned ground vehicles (UGVs) using GPS is described, highlighting the difference in complexity between controlling holonomic and differentially-driven robots.

The authors point out that not all three control principles will be feasible for all applications. Even though clever, collaborative techniques (like the one introduced in [14]) could be employed to reduce odometry error, if only dead-reckoning is available to a robot for determining its position, *unit-center-referenced* and *leader-referenced* control become problematic. Positional errors will accumulate as a result of wheels slipping or other forms of odometry noise, and hence the calculation of the centroid of all robots would become more and more inaccurate, especially as the number of robots grows. Perhaps a lesser problem in *leader-referenced* control, odometry would probably still have to be relied upon to some extent, since the leader might not be visible to all robots at all times. When possible, direct perception is a much more robust way of obtaining needed information, and therefore it seems that of the three control techniques suggested in [3], only *neighbor-referenced control* is feasible in the absence of accurate positional information. However, there is no reason why this principle could not be generalized to include more reference points than just one neighboring robot. This idea is explored in [7], [8].

In [8], the authors assume that without some global, centralized sensor like GPS, it is unrealistic that all robots know the positions and orientations of all other robots. Therefore, each robot is controlled by one of two reference principles, based on information that is locally available. Either it references itself to one neighboring robot, maintaining a certain angle, ψ, and distance, l, to it (referred to as $l - \psi$ *control*), or it references itself to *two* neighbors, maintaining distances l_1 and l_2 to them ($l - l$ *control*). Thus, the needed information is position and orientation of one or two robots close by and within line of sight. The authors go on to present a nice, theoretical framework in which these relationships between robots are represented in directed graphs. The features of the algorithm, including switching between formations, are demonstrated through several examples, done in simulation; however, [7] states that physical implementation of the approach would be a serious challenge.

The $l - \psi$ *control* control law is adopted in [1] as part of the *Leader-Following Control* mode, one of the modes in a general high-level framework for programming multi-robot systems (behaviors are called *modes*). An experiment with physical robots (though only two) is reported, where the follower robot keeps a pre-set heading and distance to the lead robot. The follower uses a camera and color-blob detection to identify the lead robot and its heading and distance.

In [5], an elegantly simple idea is presented, inspired by the covalent bonds of molecules. All robots have a predetermined set of *attachment sites* spread uniformly around the body at a pre-set distance. When set in motion, the group snaps into a stable geometric shape as robots are 'pulled' towards the nearest of their neighbor's attachment sites through the workings of so-called social potential functions. Thus the formation emerges out of local rules. Depending on the angular offset of the attachment sites, different formations are possible. The approach is validated in simulation; however, since there is no one 'right' spot for each robot due to the symmetrical nature of the attachment sites, multiple configurations with the same attachment sites are possible, while only a specific one may be desired.

A somewhat similar approach is presented in [15], although validated at a more theoretical level where participating robots are modelled as particles in the plane. Such robots could be micro-air vehicles (MAVs) or nano-robots. With simple local rules of interaction, letting particles react to certain, pre-designed forces between neighbors, spatial patterns emerge. Initially, the same multiple-configuration problem arises as in [5], and the overall shape is not perfect. When customized labels are added to all particles and two new force laws are introduced, however, perfect shapes are indeed achieved. Labels are assumed to be detectable by other particles. First, each particle is assigned a binary *spin*, and a force law is added that repels particles of the same spin. Second, a unique, two-value label is given to all particles, and a force law is added that sorts particles column-wise by the first value, and row-wise by the second value. Interestingly, a problem of distributed control is solved by solving a problem of distributed calculation. However, the labelling of the particles seems to be non-trivial and dependant on the desired global shape.

There is thus a spectrum of strategies, ranging from simple, behavior-based, purely local ones out of which global formations emerge, to more involved ones relying to varying extent on global knowledge, typically a global coordinate system and/or complete knowledge of other robots' positions and headings. The former category is characterized by minimalism and robustness but a lack of any guarantees that the desired formation will actually emerge; the latter category by reliability and efficiency but also a need for global knowledge and computation. In [13], Parker defines what exactly 'global knowledge' could mean: knowledge of 1) global goals, and/or 2) of the actions and intentions of other robots. Within this framework, a robot that knows what kind of formation, and with how many robots, it is supposed to participate in, would have Type 1 global knowledge, whereas its knowing the globally required formation heading, or whether another robot is about to evade an obstacle, would be global knowledge of Type 2. She illustrates how the addition of global knowledge can improve system performance through formation simulations with four robots.

[12], [18], [10], [19], [16], [5] all share the feature that a guarantee of the emergence of a desired formation can not be given (in some cases simply because it is not the goal of the work). We have addressed this problem by assigning a 'right spot' in the formation to each robot: each robot only goes one place. Thus, when all robots are in place, the desired formation is established.

In [6], [11], [3], global knowledge is assumed: each robot knows the position of all others. In our approach, robots only use local sensing. Our robots do not share a common coordinate system.

[2], [15] represent robots as points and show simulated experiments with discrete time steps of synchronous robot movements. Each point reacts to all points within some distance of it. In our algorithm, each robot references itself locally to one neighboring robot and keeps a certain bearing and distance to it, like in [20], [1], [7], [8]. In contract, we have demonstrated our algorithm through both simulation

and real-world experiments, and with different numbers of robots doing various formations.

IV. Our Algorithm

Bearing the spectrum of movement coordination strategies we discussed above in mind, we have in our approach sought simplicity yet reliability through locality with minimal communication. Next, we give a brief outline.

The basics of our algorithm are these. Each robot has a unique ID number. One robot is the *conductor*; any robot can take on this role. The conductor defines the heading for the whole formation. With the exception of the conductor, every robot follows one other robot called its *friend*. The distance and relative bearing to this friend, and the friend's ID, are detectable by the robot by means of some appropriate *friend-sensor* with limited visibility. The friend-sensor points forward and is assumed to have a ± 90 degree field of view. Each robot can only have one follower, except for the conductor who can have either one or two followers.

The conductor is in front. Hence, it is followed by either one or two cadres of robots. In the latter case, the two cadres are of the same length. The whole structure is called the *chain of friendships*: it either has the conductor at one end, or in the middle. The ordering of the robots in the chain is by ID.

All robots regularly broadcast heart-beat messages with their ID. The conductor broadcasts the number of the current formation along with its ID.

Thus, each robot senses the relative distance and bearing to its friend, and by radio communication it knows the total number of robots currently participating, and their IDs. It does *not* know the position of any other robot, and it does *not* know the heading of the conductor. There is no shared coordinate system. In section IV-B, our approach is described in more detail.

A. Types of Formations Considered

The types of formations (i.e., their geometric shapes) that our algorithm applies to have the following characteristics. We consider geometric shapes that

- consist of line segments,
- are connected,
- have at most two loose ends,
- are not 'frontally concave'.

Each requirement is explained in term.

In order for the formations to scale with the number of robots, N, the lengths of their constituent lines should be defined in terms of fractions of N rather than absolute values. The formation shapes themselves are defined as sets of such segments and angles between them, and a heading.

The formation shapes have to be connected since we only allow one conductor (connected simply means that a robot's friend has to be within sight). By the same reason, there is only one chain of friendships, and so the formation shape can have no more than two loose ends.

The term 'frontal concavity' refers to shapes that make a backwards curve with respect to the heading. These shapes are not allowed as they would require some robots to look back for their friends, and by the assumption of a forward ± 90 degree field of view of the friend-sensor, this is not possible. Thus, shapes like J, U, and $\sim$, with upwards heading, are not allowed by our algorithm. With this restriction, the possible shapes are essentially those that can be folded from an open bicycle chain, keeping either the middle or the end of the chain in front (see Figure 2).

We assume that all robots are a priori given a comprehensive list of such definitions of shapes, thus *a priori* explicitly embedding the formations that the group will be able to establish.

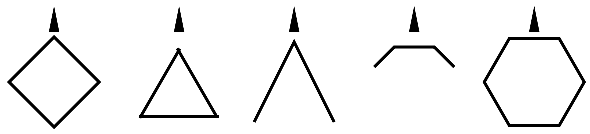

Fig. 2. Examples of formations our algorithm can produce. First from the left is a diamond. From the left, they are: diamond, triangle, arrowhead, wedge, hexagon.

Formations that have the conductor in the middle of the chain of friendships are called *centered formations*. All formations in Figure 2 are centered. In contrast, in *non-centered formations*, the conductor is at one end of the chain of friendships. When switching between two centered formations, or two non-centered formations, all robots keep the same friend. When switching between a centered and a non-centered formation, however, a new robot becomes the conductor, and all robots in between the old and the new conductor must find a new friend; this is similar to a transition between non-isomorphic control graphs in [8]. Such a change is a special case of our algorithm. To illustrate this point, consider the robots switching between the two formations in Figure 3. Robot 2 would have to find a new friend; for centered formations, its friend is 4, for non-centered, it is 1.

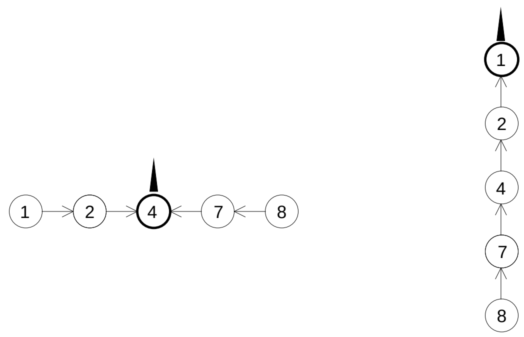

Fig. 3. The chain of friendships; $i \to j$ means i looks for j, a black triangle indicates the formation heading. In the centered formation on the left (line), the robot with the middle ID is the conductor, 4. For the non-centered formation on the right (column), the robot with the lowest ID is the conductor, 1.

B. The Details of our Algorithm

Our key idea is simple: every robot in the group positions itself relative to a designated neighbor robot, its *friend*. This friend references itself to another friend, and so on. To do this, each robot needs to be able to determine the distance and angle to its friend. We assume each robot has some sort of sensor, referred to as the *friend-sensor*, that does this. The friend-sensor is assumed to be pointing forward with a ±90 degree field of view, and have limited visibility.

Each robot has a unique ID, detectable by other robots by the friend-sensor. This ID is broadcast regularly as a heart-beat message. As a result, each robot knows how many robots (N), and which ones, are participating in the formation. The robots are organized by the order of their IDs; a robot with ID i will have a friend whose ID is either the lowest ID greater than i, or the greatest ID lower than i (IDs need not be consecutive). This principle in the *chain of friendships* is never broken.

One robot is the *conductor*, deciding the formation and formation heading, and thus not following any friend. The conductor broadcasts a byte designating which formation to do (f) along with its own ID (thus we have included Type 1 global knowledge in our system, as defined in [13]). The conductor does *not* broadcast its heading.

Any robot can potentially serve the conductor duty; who is conducting is determined by N and f. The conductor is either at the end or right in the middle of the chain of friendships: since the conductor decides the overall formation heading, it should have a clear line of sight and thus be in front. On the other hand, since the conductor defines a kind of center of the formation around which all the other robots move, it should be as much in the center as possible so as to minimize congestion during formation switching.

Therefore, in a formation led by only one robot with only one follower, the lead robot will be the conductor, and it will be the one with the lowest ID. For other formations (the centered formations), the robot with the middle ID of all currently participating is the conductor. For centered formations, robots with IDs *less* than the middle ID find a friend with an ID that is immediately *greater* than their own; robots with IDs *greater* than the middle ID find a friend with an ID that is immediately *less* than their own. For non-centered formations, all robots find a friend with an ID that is *less* than their own. See Figure 3.

The angle a robots needs to keep to its friend depends on its place in the chain of friendships and on the particular formation. By N, f, and its own ID, each robot can determine this angle *locally*. To illustrate, in the column formation to the right of Figure 3, robot 8 knows it should keep its friend, 7, at a 0-degree angle, i.e., straight ahead. For the line formation, on the left of Figure 3, all robots know that the conductor has ID 4 and is in the middle. Therefore, robots with IDs less than 4 find a friend with an ID that is immediately greater than their own, and keep the friend to their right at a 90 degree angle. Robots with IDs greater than 4 keep their friends to the left at 90 degrees.

To alleviate this calculation, each robot maintains a local value called the *lessThanMe*. What matters is really not the robot's ID, but rather its rank in the chain of friendships which is sorted by IDs. The value $lessThanMe$ is precisely this rank, stating how many robots are currently alive with IDs less than this robot's ID. Thus, the $lessThanMe$'s are consecutive (see table I). For centered formations with N robots, the conductor is the one with a $lessThanMe$ of $\lfloor N/2 \rfloor$. For non-centered formations, the conductor is the robot with a $lessThanMe$ of 0.

ID	1	2	4	7	8
$lessThanMe$ (rank)	0	1	2	3	4

TABLE I

RELATIONSHIP BETWEEN ID AND *lessThanMe*.

Consider the leftmost formation in Figure 2, the diamond. The robots form four perpendicular line segments (see Figure 4). The low fourth of the robots (0 and 1 in the same figure), form segment I by keeping their respective friends at a 45-degree angle to the front and left; the second-lowest fourth (2 and 3 in the figure), who are in front of the low fourth, form segment II by keeping their friends 45 degrees to the front and right. The middle robot, 4, is the conductor, and the upper two fourths look left (5 and 6, segment III) and right (7, segment IV) for their friends, respectively.

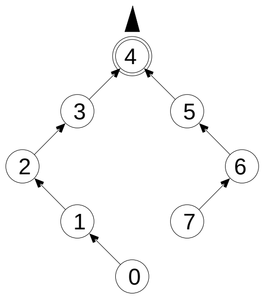


```
n = ⌊N/4⌋
if lessThanMe < n
  camangle = -45
else if lessThanMe <⌊N/2⌋
  camangle = 45
else if lessThanMe > N-n
  camangle = 45
else if lessThanMe >⌊N/2⌋
  camangle = -45
```

Fig. 4. Example: a diamond of 8 robots. The one with a $lessThanMe$ of $\lfloor N/2 \rfloor = 4$ is the conductor. Arrows indicate angle to friend. Filled triangle shows formation heading.

Since formations that are frontally concave are not dealt with, a robot will only have to keep its friend in a range of $\theta \in [-90, 90]$-degree angles with respect to its heading. However, this could still pose a problem; at the extremes, the friend is likely to fall outside the field of view of the friend-sensor. Also, for each angle θ, the algorithm for keeping the friend at θ would essentially be different. For these two reasons, we assume that the friend-sensor can be *panned*, and we propose the following approach. Instead of keeping its friend at angle θ, the robot simply pans its friend-sensor θ degrees and *keeps its friend in the center of view* (see Figure 5). Thus, a buffer zone is introduced at the angular extremes of -90 and 90 degrees, and the friend is more likely to stay within view. Moreover, a very nice structure immediately emerges for the algorithm, described next.

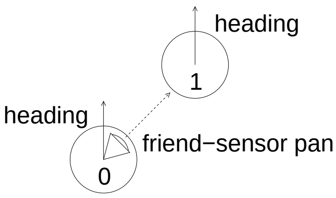

Fig. 5. Example: robot 0 has panned its friend-sensor 45 degrees and keeps its friend robot 1 centered. A filled, black triangle indicates heading.

Based on N, f, and $lessThanMe$, each robot determines the angle θ to keep relative to its friend and pans its friend-sensor to that angle. This represents the highest level of abstraction. At the next level, the core of the algorithm works to center the friend in the view of the friend-sensor. It does so with no regard to the current formation; this has been abstracted away at the level above. Finally, at the lowest level, basic collision avoidance is introduced, if necessary by means of other sensors. This layered structure is depicted in Figure 6.

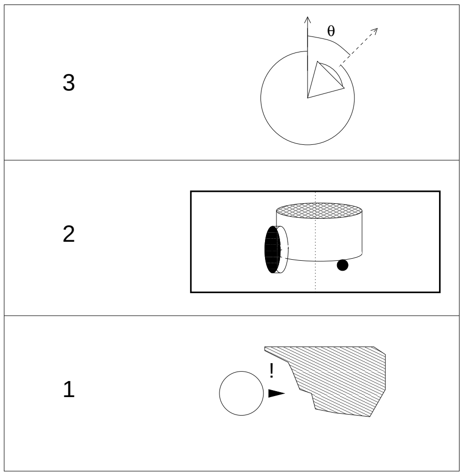

Fig. 6. The three levels of control algorithm abstraction. At the highest level, the appropriate angle to the friend is determined by N, f, and $lessThanMe$. At the intermediate level, the robot simply concentrates on keeping its friend centered in the field view. At the lowest level, the robot handles collision avoidance.

There are several nice implications resulting from this approach. First, once the conductor starts moving, the only way for a robot to keep a stable position relative to its friend is by finding the same heading. Thus, the conductor 'drags' the whole formation into place just by going its own way[1]. Thus, no global heading needs to be agreed upon; it emerges automatically. Since any robot can be the conductor, what seems a centralized element really is not. If the conductor fails, another robot takes over the role.

Second, since the algorithm is basically 'keep your friend in the center', a switch between two centered or two non-centered formations is easily done by gradually panning the friend-sensor into the new appropriate angle; the change in position results automatically. As discussed above, switching between a non-centered and a centered formation is somewhat more involved.

Third, there is no global coordinate system and hence no communication of coordinates and headings. Only minimal heart-beat messages and formation messages are communicated. This information, together with a desired inter-robot distance and a unique ID, is all each robot needs to know to locally determine its place in the global formation.

The following sections present an implementation of this algorithm and address the solution requirements listed in Section II.

V. Implementation

A. The Friend-Sensor

In our implementation, a laser and a color camera are combined as the friend-sensor. The camera identifies the friend, the laser gives the distance to it. However, only the camera can pan.

Friend identification is done using color-blob detection; each robot wears a carefully designed helmet with two fluorescent color stripes. The stripes are clearly visible to other robots, and by their color and order, the corresponding ID can be inferred (see Figure 7). The helmet is placed around the camera, with the lens peering through a 180-degree aperture between the color stripes. As a result, the stripes are at the same height as the sensor detecting them, and furthermore, since the center of the color stripes is also the center of the camera, aligning the camera to the helmet corresponds to aligning robot center to robot center.

Since the experiments reported in this paper only involve four mobile robots, helmets with two colors and two stripes were sufficient for unique identification, but our vision system can be trained to recognize up to 8 colors, so the ID-helmet principle readily scales to a substantial number of robots.

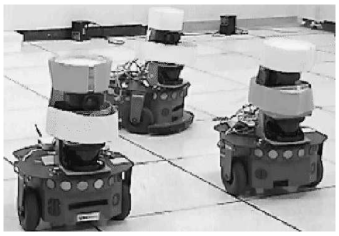

Fig. 7. Three robots with color helmets, each encoding a unique ID by the combination of color stripes.

B. The Controller

Our controller consists of three concurrent behaviors and a module holding state data (Figure 8). Each is described in turn.

[1]Much like when one picks up a mobile that has fallen to the floor.

The state module *WhatDoIKnow***:** Here resides all state information, including:

- this robot's own ID
- the total number of robots, N
- the sorted table of all live robots' IDs
- this robot's rank in the table of IDs, *lessThanMe*
- the current formation, *f*
- the ID of this robot's friend
- the angle the robot should keep to its friend, *camangle* (the angle that the camera will pan)

The behaviors manipulate and make use of this information as follows.

The *channelNListener* **behavior:** This behavior receives the heart-beat messages from the other robots and maintains N and the table of IDs in *WhatDoIKnow*. This information is used to calculate *lessThanMe*. Also the ID of the robot's friend, *myFriend*, is determined, taking *f* into account. If N changes, the robot might be promoted to be the conductor. This happens if its *lessThanMe* value becomes equal to $\lfloor N/2 \rfloor$, in which case it is the middle robot[2]. If so, the *camangle* is set to 0, allowing the robot to look straight ahead. Conversely, it could also be demoted if its *lessThanMe* is no longer $\lfloor N/2 \rfloor$. In this case it looks up a friend in the table of live IDs.

The *channelCListener* **behavior:** This behavior receives formation messages from the conductor and updates the formation variable, *f*, if necessary. If *f* changes, the robot might have to pan its camera to a new angle. The correct angle, *camangle*, is calculated from a simple geometric relationship of *lessThanMe* and N. For any N, the robots' respective *camangles* will result in a formation that is either uniform or incomplete; i.e., attempting a diamond with 5 robots will result in an incomplete 3x3 diamond, not an overcrowded 2x2.

Also, *channelCListener* deals with messages from other robots that have encountered an obstacle and are evading it. If a robot detects an obstacle in its path (see below), it will swerve to evade that obstacle and send out a warning, a *swerve* message, with its ID and a value indicating the direction and sharpness of the turn. Other robots, not necessarily sensing the obstacle, will react to this message by making a swerve of solidarity of equal sharpness and direction, if the sender is swerving their way. Thus, robots share with each other knowledge about upcoming obstacles, categorizing our system as Type 2, following [13].

Combined, *channelCListener* and *channelCListener* represent the highest level of abstraction of Figure 6. Everything pertaining to the actual formation is taken care of in these two behaviors; What friend the robot should stick to, and at what angle. Or if the robot is in fact the conductor. The main behavior, described next, comprises the two lower levels of abstractions of Figure 6.

[2]As explained, for non-centered formations, the robot whose *lessThanMe* is 0 will get the promotion.

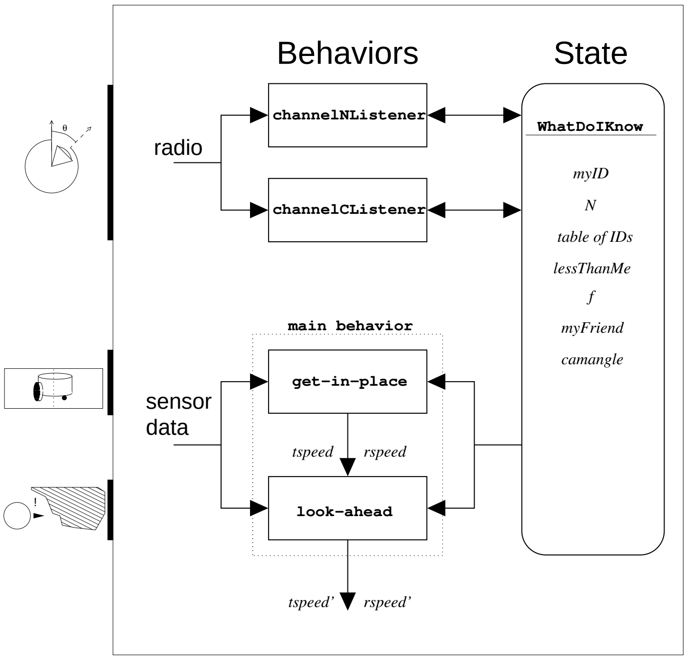

Fig. 8. The controller consists of a state module (*WhatDoIKnow*) and three concurrent behaviors. The first two, *channelNListener* and *channelCListener*, receive messages from other robots. They represent the highest level of abstraction, marked by the icon on the left. The main behavior consists of two serialized sub-behaviors that control the robot based on local sensor data; they represent the two lower levels of abstraction. *get-in-place* keeps the friend in the center of view, *look-ahead* does clever collision avoidance.

The main behavior: A robot R moves by setting two parameters, translational and rotational speed (*tspeed* and *rspeed*). In the main behavior, R cycles through a control loop that reads the sensors, sends out its heart-beat message, sets *tspeed* and *rspeed* to their default values (20 mm/sec and 0 deg/sec, respectively), and then passes them to the *get-in-place* sub-behavior.

The *get-in-place* **sub-behavior:** If R is the conductor, it modifies neither *tspeed* nor *rspeed* — unless it has just circumnavigated an obstacle or made a swerve of solidarity. In this case it will gradually modify *rspeed* so as to return to the heading it had before the interruption. Previous heading is stored in *WhatDoIKnow*.

If R is not the conductor, it will first locate its friend. It looks up its friend's ID in *WhatDoIKnow*, and deduces the color helmet to look for. If other robots are swerving its way to evade obstacles, R makes a swerve of solidarity. Otherwise, R will first get close to its friend: it does so by panning its camera straight ahead (yielding a straight catch-up path towards the friend) and speeding up. Once appropriately close, R will start panning the camera towards *camangle*, the angle ordered by *channelCListener*, all the while keeping the friend in the image. Continuously, R makes small corrections to *rspeed* and *tspeed* so as to center its friend in the image. These suggested values of *rspeed* and *tspeed* are then given to the $look - ahead$ sub-behavior.

The *look-ahead* **sub-behavior:** A central element of *look-ahead* is the *aheadbuffer*. From *tspeed* and *rspeed*, a bounding box for the resulting movement is calculated and a buffer is added: the width of the robots, `robot-size`, on the sides, and $aheadbuffer$ in the front. Then, it is checked if any obstacles are found within this bounding box. If so, a correction is made to *tspeed* and *rspeed* proportional to the proximity of the sensed obstacle.

Thus, $aheadbuffer$ induces immediate collision avoidance. In addition, if it is set to a high value, it allows R to look far ahead for obstacles, resulting in an elegant, smooth avoidance behavior (see Figure 9). However, R cannot keep $aheadbuffer$ high if it is not yet in its place in the formation, as it may have to get close to other robots. Consequently, $aheadbuffer$ is set *low* (to `robot-size`) if R has not yet been in place in the formation, R has been out of place for a long time, or R has other robots in front of it when moving in the formation (derived from N, *lessThanMe* and *f*). Otherwise, $aheadbuffer$ is set to a higher value proportional to `robot-size` and *N*; a large formation needs more space to negotiate obstacles than a small one. A higher $aheadbuffer$ gives more time to react, and more time yields more space. If R is the conductor, $aheadbuffer$ is always set high. As a safety measure, another sensor (in our case sonar) can be used for lowest-level collision avoidance.

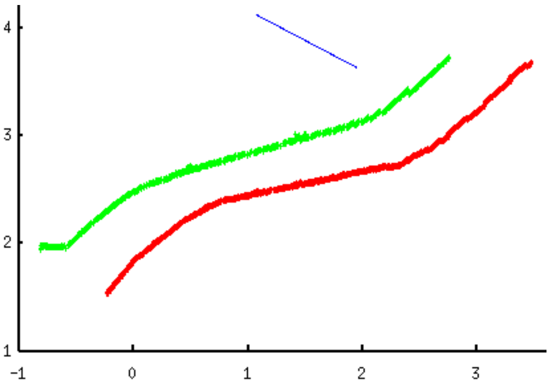

Fig. 9. Example: two robots in an incomplete diamond smoothly avoid a wall, taking advantage of a high $aheadbuffer$. Real robot data, units are meters.

Finally, the *look-ahead* behavior sends the revised (*tspeed*, *rspeed*) command to the wheels, and R makes the corresponding movement. Thus, through a series of simple update rules that modify the translational and rotational speed, based only on local sensor information, a small set of state variables, and the heart-beat information, each robot eventually falls into place relative to its friend. In turn, a global formation emerges.

VI. Experimental Evaluation

Our assumption is that robots start out in the right order with respect to the chain of friendships, but not necessarily with their respective friends in the field of view. This is reasonable, given that the problem of aggregating robots into such a configuration has already been empirically demonstrated [12]. Furthermore, following ideas from [19], it is possible through only local interaction to have N robots form a chain.

To validate our approach, we have performed extensive experiments, both in simulation and with physical robots. Our simulation experiments were done using *Player* and *Stage*. *Player* is a server and protocol that connects robots, sensors, and control programs across a network; *Stage* simulates a set of Player devices [9], [17].

For the real robot experiments, we used the ActivMedia Inc. Pioneer2 DX robot with the SICK LMS200 laser, sonars, and the Sony PTZ camera, and with the *Player* software [9]. To enable the detection of IDs, we then constructed the above-mentioned brightly-colored helmets and placed them around the camera, atop the laser. A picture of an instrumented robot can be seen in Figure 10.

To obtain ground truth data for evaluating system performance, we set up a laser-based metrology system. Two lasers were placed at the boundaries of the experimental workspace, for tracking the robots. Tracking was made possible by placing a reflective marker on top of the helmet. This system allowed us to track the individual robots as they moved, and record their positions for analysis.[3]

Fig. 10. A robot with laser and panned camera, the lens peeping out through the hole in the color helmet.

A Pioneer robot is approximately 50 cm across. Due to the limited space available in our lab, an approximately 4x6 meter arena (see Figure 11), we had the robots move very slowly in order to allow time to get in formation, and for trailing robots to catch up, given the confined workspace (the default *tspeed* was 20 mm/sec).

A. Experimental Design

Summarizing the solution requirements of Section II, we designed our experiments so as to validate the following properties and capabilities: generality, stability, robust-

[3] The tracking program was written by Andrew Howard.

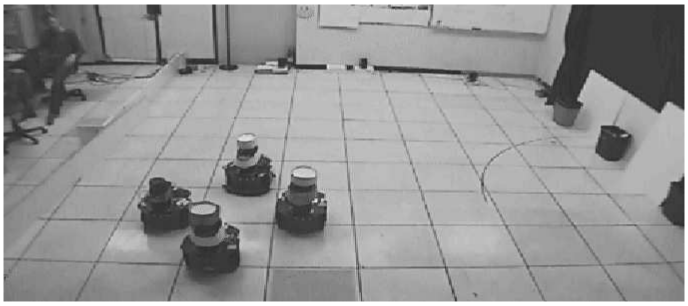

Fig. 11. One author and four robots in our lab arena.

ness, switching between formations, and obstacle avoidance.

As outlined in Section IV-A, our algorithm applies to a very general set of formations. In the real robot experiments, we focused on line, column, wedge, and diamond, following [3], but we also verified that an arrowhead (i.e., approximately a down-pointing $\surd$), can be formed and maintained by four robots. In simulation, we performed experiments with line, column, wedge, diamond, arrowhead, and circle. However, to keep consistency between simulation and real world experiments, here we report only results for line, column, wedge, and diamond. Thus, below, any mention of formations will mean one of these four, unless otherwise stated.

To display stability, the group of robots must be able to *establish* any formation from some arbitrary initial configuration (within the boundaries of the above assumption), and to *maintain* the formation over time, travelling through an environment with no obstacles.

To display robustness, the group of robots must be able to adapt any formation in case of one or more robots failing.

To display formation switching, the group of robots must be able to perform a switch between any two formations in an environment with no obstacles.

To display obstacle avoidance, the group of robots must be able to negotiate different types of obstacles, in any formation. Either by maintaining the formation while avoiding the obstacle, or by re-forming the formation after the obstacle.

Before we report our results from experiments in each of these scenarios, we present the formal measure we have developed and applied to evaluate if our robots are in formation.

B. Evaluation Criteria

We suggest the following *formation evaluation criteria* as a means of judging quantitatively the notion of being in formation:

Definition1 Given the positions of N mobile robots, an inter-robot distance $d_{desired}$, a desired heading h, and a connected geometric shape $\mathcal{G}$ completely characterizable by a finite set of line segments and the angles between them, the robots are considered to be in formation $\mathcal{G}$ iff:

(1) *uniform dispersion*: $\exists d$, such that $\forall$ pairs of immediate neighbors (R_{i_1}, R_{i_2}) with distance $dist(R_{i_1}, R_{i_2})$, $|d - dist(R_{i_1}, R_{i_2})| < \epsilon_{d_1}$, and $|d - d_{desired}| < \epsilon_{d_1}$,

(2) *shape*: $\exists$ a 'stretching function' f with $f(\mathcal{G}) = \tilde{\mathcal{G}}$, such that $\forall$ angles $\theta \in \mathcal{G}$, $|f(\theta) - \theta| < \epsilon_a$, and such that $\forall$ robots R_i, with distance $dist(R_i, \tilde{\mathcal{G}})$ to $\tilde{\mathcal{G}}$, $dist(R_i, \tilde{\mathcal{G}}) < \epsilon_{d_2}$,

(3) *orientation*: $|f(h) - h| < \epsilon_a$; for small $\epsilon_{d_1}, \epsilon_{d_2}, \epsilon_a > 0$.

Criterion 1 states that the same distance should be kept between all neighboring robots. Criterion 2 states that it should be possible to lay out the desired shape over the positional data and perhaps adjust the angles a little, so that all robots are close to this tweaked shape. I.e., no angle in the original shape must be stretched more than ϵ_a to make the data points fit. Allowing the heading of the shape to be one of its defining angles, criterion 3 states that the stretching from criterion 2 must not skew the heading more than ϵ_a.

Note that by using the term "immediate neighbor", Definition 1 does not demand completeness of formations; this means that 6 robots can actually form a diamond. Note also that the measure is *global* in the sense that N robots are not considered to be in line even if the angular offset between neighboring robots is small, if in fact overall they form, say, an arc. In other words, even if the behavior of a robot is locally meaningful, it may not be so when considered globally. This point was also made in [2]. Definition 1 is easily extended to shapes that do not consist of line segments, such as circles; we omit these special cases to enhance clarity.

In the following sections, we report our experimental results. In *all* experiments, both simulated and real-world, we used Definition 1 to determine whether the robots were in formation. To establish basic notion: The distance travelled by the conductor up to the point where the robots establish a **f**ormation **t**he first **t**ime is called *ft2*.

C. Stability Results

In simulation, the 'establish-formation' and 'maintain-formation' components were done simultaneously. The robots started out in the right order next to each other, but with inter-robot distance 70 cm, and with random headings (except for the conductor whose heading was correct). With four robots, we ran 10 trials for each formation. In every trial, the robots had to establish and maintain the formation, travelling 19.0 meters.

We then measured the % of time-steps the robots were in formation after $ft2$. From the %'s of the 10 trials, we report the weighted average %, and the standard deviation, σ. Using Definition 1 (here as well as in all other simulation experiments reported in the following sections), we set ϵ_{d_1}=ϵ_{d_2}=10% of the desired inter-robot distance ($d_{desired}$=60 cm), and ϵ_a= 3.6 degrees. Results are shown in table II.

The large standard deviation for the wedge in table II is due to one run where the % of time in formation after *ft2*

$N = 4$ 19.0 meters	$ft2$ (meters)	av.% time in formation after $ft2$ (σ)
line	0.7	88.4 (4.9)
column	3.0	100.0 (0.0)
diamond	2.6	100.0 (0.0)
wedge	1.8	92.0 (12.7)

TABLE II

Simulation: Stability, weighted averages over 10 runs.

was only 54.6. If this possible outlier is excluded, the % and standard deviation are 95.3 and 3.7, respectively.

Next, we evaluated the stability of the approach with real robots. As mentioned, our lab space is quite small, and in fact there was not enough room for the robots to first establish a formation, and then maintain it for a significant amount of time. Therefore, we had to separate the 'establish-formation' and 'maintain-formation' portions to show stability.

With 'establish-formation', the same was done as in simulation: in the initial configuration, the robots were placed in the right order, next to each other but with random headings, and they then established the formation within the boundaries of our lab arena. We verified this for all formations several times. Figure 12 shows four trials.

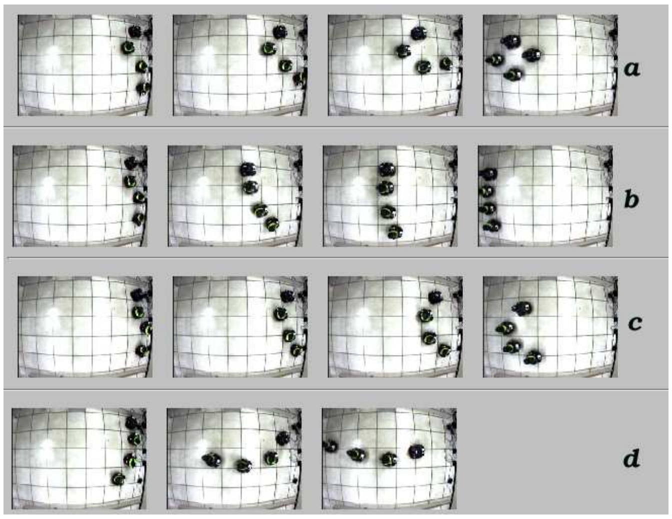

Fig. 12. Overhead camera images, order is left to right; robots move from right to left in each image. Establishing a diamond (a), a line (b), a wedge (c), and a column (d).

For 'maintain-formation', the robots were first placed close to the desired formation, thus saving most of the time required for 'establish-formation'. Then, we again recorded the % of time they were in formation after actually establishing it; i.e., the % of time in formation after $ft2$. As trials were very time consuming, only 5 trials with each formation were done. Instead of standard deviation, we simply show the %'s of all 5 trials, together with their weighted average. Results are shown in Table III.

In all real robot experiments (here and in following sections), we used Definition 1 with $d_{desired} = 80$ cm,

N=4	T1	T2	T3	T4	T5	Av.
line	94.1	46.8	70.4	34.7	64.9	60.9
colm	96.1	100.0	98.9	100.0	100.0	99.2
diam	75.0	99.7	88.6	99.8	100.0	91.9
wedg	100.0	99.9	99.6	50.3*	49.8	78.2

TABLE III

Real robot data showing stability for each formation, with four robots in 5 trials, T1-T5. Values are % of time in formation after $ft2$. Final column is the weighted average of the five trials.

ϵ_{d_1}=$(0.20 * d_{desired})$, ϵ_{d_2}=$(0.08 * d_{desired})$, and ϵ_a=$(0.08 * 2\pi)$. In other words, the dispersion of robots was set fairly loosely, allowing the average inter-robot distance to differ up to 20% from $d_{desired}$. Criterion 2 was applied more strictly: all robots had to be at most 8% of $d_{desired}$ from the line segments they belonged to[4]. Angles were allowed to deviate up to 28.8 degrees. We used a simple line-fitting algorithm to fit data points to straight lines. Examples of shapes that did and did not qualify, using our measures with these parameters, are shown in Figure 13.

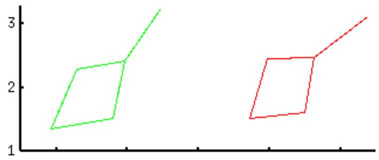

Fig. 13. Example: two shapes formed by four real robots with the conductor showing its heading. Only the shape on the right was judged to be a diamond by our measure.

Tables II and III demonstrate that the column and the diamond are indeed very stable. The wedge seems to be either very stable or rather unstable, and the line is unstable. These facts stem from two related problems.

The first problem is that the wedge and the line share the challenging feature that a subset of robots must pan their cameras ± 90 degrees to look for their friend. This orientation proved to be the most difficult to maintain. With a ± 90 degree camera pan, it is very hard for a robot R to ascertain whether it is ahead of or behind its friend F, if their headings are slightly different. If R thinks it is ahead, it will speed up, possibly resulting in an unstable, oscillatory course. An example of this is shown in Figure 14 which shows the T4 trial, indicated with a * in Table III.

The second problem is that the SICK laser range is only ±90 degrees: if R is right next to F and its camera is panned ±90 degrees, F is at the boundary of R's laser range. If F drops beyond that boundary, R can no longer judge the distance to F accurately, and so again an oscillatory course can result.

[4]For the line, we set $\epsilon_{d_2} = (0.30 * d_{desired})$. Otherwise, at least one robot would often find itself too far from the best-fit line.

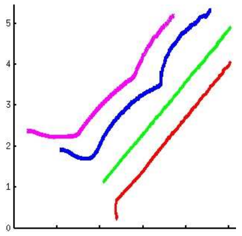

Fig. 14. Example: a poor wedge trial from Table III; the second robot from the top oscillates, and consequently, so does its friend (to its left).

Fig. 15. A trial of 6 robots in a circle (a hexagon); one is terminated, and the rest adapt to a 5-circle, a pentagon.

Thus having addressed stability, we posit that for robustness, formation switching, and obstacle avoidance, it suffices to show that upon initially establishing the formation, the robots do get back in formation after the occurrence of failure, switch, or obstacle, respectively. In other words, in these three categories of experiments, a trial is a success if the robots establish the formation, overcome some incident, and re-establish an adapted formation within some pre-set distance.

D. Robustness Results

The structure of the simulation experiments addressing robustness is as follows. The robots establish the formation, one or two robots are terminated when the conductor has travelled 6.0 meters, and the robots then establish an adapted formation. We tested this by showing that 1) the initial formation is established (at *ft2(1)*) before the conductor has travelled 6.0 m, and 2) an adapted formation is achieved (at *ft2(2)*) before the conductor has travelled 19.0 m. Thus (*ft2(1)*) should be less than 6.0, and (*ft2(2)*) should be greater than 6.0 and less than 19.0. Table IV shows results from two experiments with 6 robots; 10 trials were done of each.

robot drop-outs (start: $N = 6$)	*ft2 (1)* (std.dev.)	*ft2 (2)* (std.dev.)
circle, 1 out	3.2 (0.077)	8.9 (0.130)
diamond, 2 out	4.0 (0.205)	8.3 (0.235)

TABLE IV
SIMULATION: ROBUSTNESS, AVERAGED OVER 10 RUNS FOR EACH.

Once the conductor reached the 6.0 meter mark, 1 or 2 robots were eliminated. For the circle (i.e., hexagon), the conductor itself was eliminated. The remaining 5 robots themselves realized this (having stopped receiving the heart-beat from the conductor), promoted a new conductor, and eventually reached an adapted circle (a pentagon) after an average of 8.9 m. Figure 15 shows a similar trial where it was not the conductor that was eliminated.

In the diamond case, the 6 robots established the diamond the first time at 4.0 m on average – an incomplete '6-diamond' like the one shown on the left in Figure 16. At 6.0 m, the two extreme robots were terminated (0 and 5 in Figure 16, leaving 3 the conductor). With only 4 robots remaining, a complete 4-diamond could be formed through a restructuring of the remains of the incomplete 6-diamond; on average, this adapted formation was established by 8.3 m. (In Figure 16, robot 1 slides in underneath 2 and ends up behind 3 as the back corner.)

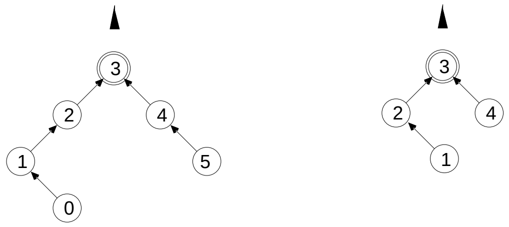

Fig. 16. Example: switching from an incomplete 6-diamond to a complete 4-diamond after robots 0 and 5 are terminated.

We performed real robot robustness experiments for all formations, both terminating the conductor and other robots. However, because robots tended to occlude each other while negotiating terminated colleagues, it was impossible to get tracking data. Instead, Figures 17 and 18 show video frames from actual trials. Figure 17 from a column trial where the conductor was terminated, Figure 18 from a wedge trial where a follower robot was terminated and later revived.

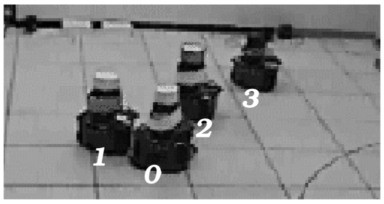

Fig. 17. Example: originally, the four robots were in a column, and then the conductor (robot 0) was terminated. Robot 1 was then promoted to be the new conductor, and it is shown circumnavigating the terminated robot 0. The others follow it.

Our algorithm proved to be very robust. Due partly to the extremely high reliability of the (minimal) radio communication, and partly to the high recognizability of the color helmets, robots were always able to determine a new (correct) friend and locate it, or to assume the conductor duty. By means of the table of IDs of other live robots that each robot maintains in its $WhatDoIKnow$ state module (Section V), it is simply a matter of a table look-up to find the right new friend to look for, if the current friend stops sending heart-beat messages. Hence, we believe it to

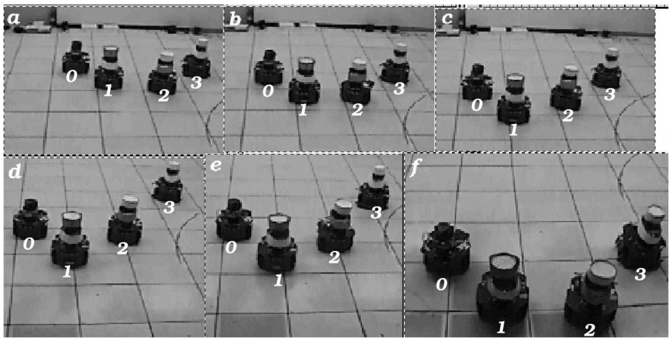

Fig. 18. Time sequence of frames from a video recording. In image (a), four robots have established a wedge, but then robot 3 is terminated. In (b), robot 2 is waiting for 1 to pass it. This has happened in (c), and in (d), the three remaining robots (0, 1, and 2) have established an adapted wedge. Then robot 3 is revived. In (e), 2 has sped up to get back on the side of 1, while 3 is struggling to catch up. And finally in (f), the original wedge is re-established.

be satisfactory to demonstrate robustness by only showing results from a few sample experiments in simulation and reality.

E. Switching Results

Ideally, many trials of switching between each and every pair of formations should be conducted to fully document the algorithm. This leads to an intractable number of trials, however, so instead we consider grouped formation switches: a switch between centered formations (call it Type A), and a switch between a non-centered and a centered formation (call it Type B). Since the only non-centered formation considered is the column, there is no need to consider switches between non-centered formations.

Type A switches are simple, since all robots keep the same friend, and the conductor stays the same. All that is changed is a robot's *camangle* to its friend. As this happens gradually, the change in position also happens gradually, and the robot has no difficulty in keeping its friend in the image.

Realizing this, there is no essential difference between switching from, say, the diamond to the line, and switching from the diamond to the wedge, or vice versa. Basically, it is a question of 'folding' or 'unfolding' the chain of friendships, with the degree of folding depending on the formations involved in the switch.

In our simulation experiments, therefore, we have demonstrated two examples of Type A switches: diamond→wedge and line→diamond. The former represents an unfolding of the chain from the closed diamond shape to the open wedge shape; the latter represents a folding from the flat line to the closed diamond.

We performed 10 (9) trials of each; results are shown in table V. In all trials, the robots established the first formation for the first time at *ft2(1)*. Then, after 6.0 meters, they were ordered to make the switch, and before 19.0 m, the new formation was established (*ft2(2)*). The results for the line→diamond switch are averaged over 9 runs only. In one of the 10 runs we performed, the robots failed to form the diamond within the 19.0 meter mark, and thus for this run the *ft2 (2)* is void. In Figure 19, we show another example switch.

$N = 4$	*ft2 (1)* (std.dev.)	*ft2 (2)* (std.dev.)
(A) diamond→wedge	3.6 (0.144)	8.3 (0.089)
(A) line→diamond	1.1 (1.373)	9.7 (0.193)
(B) line→column	4.0 (1.747)	9.2 (0.372)

TABLE V

Simulation: Switching, averaged over 10 (9) runs for each.

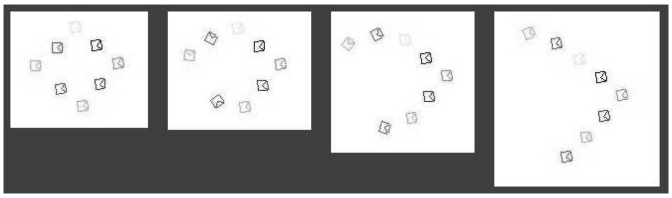

Fig. 19. Example: 8 robots switching from a diamond to an arrowhead.

When switching from the line to the diamond, the line folds back and the two extreme robots end up next to each other (their cameras panned in opposite directions). This is an example of a situation where letting the middle robot be the conductor minimizes congestion. If the leftmost robot in the line had been the conductor, the rightmost robot would have had to make a much longer correctional movement.

Type B switches are more complex, as previously explained. When switching from diamond to line, for example, the new conductor finds itself *behind* the other three robots when the order to switch is received. We therefore made sure that any switch to column went through the line first, so that the conductor for the column (the leftmost robot) would have a clear path when actually becoming the conductor. I.e., any folded shape was forced to unfold before allowing a switch to the column. Once this was done, a switch to the column was trivial. Table V shows trials of the switch line→column.

Switching from column to a centered formation is tricky. Here, the robot in the middle of the column becomes the conductor, and again its path is blocked. Moreover, all robots in front of the new conductor now have to locate a new friend *behind* them. Treating this situation as a special case is necessary. Currently, our algorithm causes the front half of the robots to make a swerve to the left, allowing the new conductor to pass them. The further from the middle a robot is, the bigger a swerve it will make. Ideally, this means that it will end up on the outside of the robot behind it, which makes a slightly smaller swerve. Figure 20 explains this further.

Since the maneuver in its current form does not scale with the number of robots, we do not report simulation results for this kind of Type B switch. However, we have verified in simulation that the idea works for up to 4 robots.

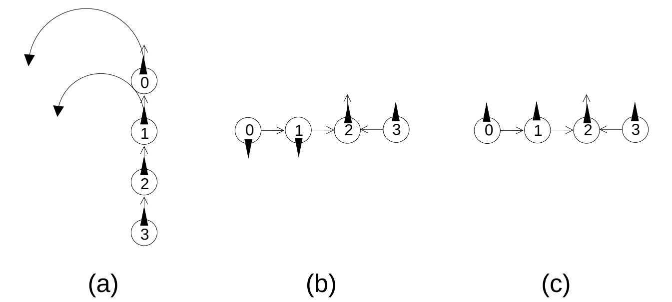

Fig. 20. Switching from column to line. Filled triangles indicate robot headings, thin arrows show the pan of the camera. First, the front half of the column makes left swerves. The further from the middle, the bigger the swerve (a). They then pan their cameras left to look for their new friend (b). After seeing the new friend, they gradually pan their cameras from left to right, causing them to turn left while keeping their friend centered, so that they in the end have the same heading as the conductor (c).

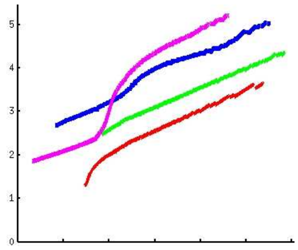

Fig. 21. Real robot data, units are meters. A switch from diamond to wedge.

With real robot experiments, we again encountered the problem of occlusions. It was not possible to reliably track the robots during a switch, so we only show one graph of position data, demonstrating a switch from diamond to wedge (Figure 21). However, again we verified that all Type A switches were performed reliably. Figure 22 shows overhead camera images of a switch from diamond to line.

Fig. 22. Sequence of overhead camera images of a switch from diamond to line, order is left to right. Robots move from right to left in each image.

For Type B switches, we verified that line→column was performed reliably (essentially, the same thing happens as shown in Figure 12d). For the switch column→line, we have validated the above-mentioned maneuver for three robots only. A sequence of video images from this trial is shown in Figure 23.

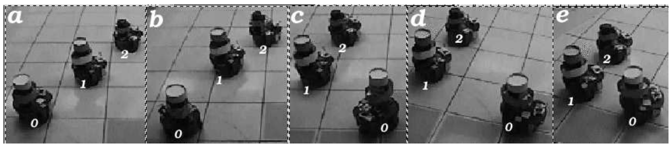

Fig. 23. Video sequence of a Type B switch: column→wedge. In image (a), the three robots have established a column. In (b), a switch has been called, and so robot 0 makes a swerve to the left. Robot 2 starts to move into position behind and to the right of robot 1, the new conductor. In image (c), robot 0's swerve is complete, and it sees its new friend, robot 1. In (d), 0 has turned toward 1, and finally in (e), the wedge is complete.

F. Obstacle Avoidance Results

When a robot formation encounters an obstacle, it should preferrably move around it as rigid body, preserving the formation as much as possible. With more than one obstacle, or with very large or very complex shapes, the robots should at least re-establish the original formation on the other side.

We have tested our algorithm in simulation with various scenarios. All four formations were tested with 10 trials each, in a scenario where a wall (called wall_1), oriented orthogonally relative to the formation heading, was in the path of the formation. In addition, two more scenarios were tested with 10 trials each: a diamond formation crossing an obstacle field with 7 robot-sized obstacles, and a line crossing a long, flat wall (called wall_2), parallel to the formation heading. Results are shown in Table VI. The evaluation goal was set for the robots to establish the formation before 6.0 m, and then re-establish the formation within 19.0 m.

$N = 4$	*ft2 (1)* (std.dev.)	*ft2 (2)* (std.dev.)
line, wall_1	1.2 (0.900)	10.5 (1.081)
column, wall_1	3.8 (0.018)	13.1 (1.464)
diamond, wall_1	2.8 (0.070)	10.7 (1.274)
wedge, wall_1	2.1 (0.737)	11.8 (1.614)
diamond,obstacles	2.7 (0.143)	11.5 (0.925)
line, wall_2	0.7 (0.936)	14.7 (0.361)

TABLE VI

SIMULATION: OBSTACLE AVOIDANCE, AV. OVER 10 (9) RUNS.

Figure 24 shows a simulated trial of a column evading the wall_1 obstacle. In all 10 trials of the diamond in the wall_1 scenario, the % of time in formation after $ft2$ was over 60. In other words, after first establishing the diamond, before the obstacle, the robots managed to maintain the diamond while negotiating the wall in more than 60% of the time. In four trials, this number was 90% or more. For both sets of line trials, results are for 9 runs only; in one trial in each scenario, the robots did not re-form the line in time.

With real robot experiments, space was a problem. By the proportionality of $aheadbuffer$ to the number of robots, four robots simply demanded too long a clear line of sight to fit into our lab space. Reducing the $aheadbuffer$ means a deterioration in performance: robots can get closer to obstacles before they initiate their avoidance behavior. With a long $aheadbuffer$, a rigid-body avoidance of the whole formation is possible; with a small $aheadbuffer$, robots tend to get tangled up in each other – and moreover, the obstacle behavior one can show, collapses into individual robots avoiding obstacles and each other. And this

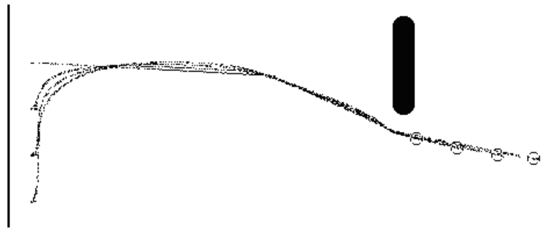

Fig. 24. Example: the conductor sights the obstacle from a distance and makes an early swerve. Once past the wall, the conductor gradually returns to the original heading.

local obstacle avoidance behavior works very reliably; our robots never hit each other. So, with a small $aheadbuffer$, 'formation obstacle avoidance' loses its meaning.

Therefore, we instead conducted experiments using only two robots. That meant a comparatively smaller $aheadbuffer$, and hence it was possible to demonstrate the desired rigid-body avoidance in two scenarios. The robots performed an incomplete '2-diamond' – meaning simply that robot 0 kept behind and to the left of robot 1 at an angle of 45 degrees (exactly as in Figure 5). In what follows, we refer to the conductor as C, and to its follower as R.

N=2	T1	T2	T3	T4	T5	Av.
O_1	100.0	97.9	99.8	97.8	93.9	98.0
O_2	81.8	82.0	89.4	21.0	26.7†	61.6

TABLE VII

REAL ROBOT DATA: OBSTACLE AVOIDANCE, TWO ROBOTS. 5 TRIALS, T1-T5, FOR EACH SCENARIO. NUMBERS ARE % OF TIME IN FORMATION AFTER $ft2$. FINAL COLUMN IS WEIGHTED AVERAGE OF THE FIVE TRIALS.

In the O_1 experiment of Table VII, the robots had to negotiate a wall that was only in R's path. By the combined workings of the long $aheadbuffer$ and the *swerve messages* sent out by R, C made swerves of solidarity in time to let R keep its position. Hence the robots maintained the formation while still evading the wall, as seen by the high %'es of experiment O_1 in Table VII. The data shown in Figure 9 are from an O_1 run. A picture of the scenario is shown in Figure 25.

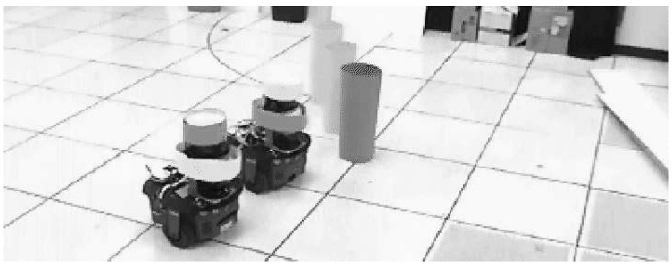

Fig. 25. The scenario in obstacle avoidance experiment O_1.

In O_2, the wall was in both of the robots' paths, but this time the wall had a passage in the middle, not wide enough to allow the robots to pass through it while keeping in formation. As seen by the three high %'es in trials T1–3 in Table VII, only for a short time did R lose its position, going through the wall, before regaining it on the other side. The T5 trial (marked with †) is shown in Figure 26; here R did not follow C through the passage in the wall, but instead went around it to the left, completely losing sight of C. Another very similar trial is shown in Figure 27. Generally, in this situation, a robot attempts to get back to the heading it had when it was last in position; thus R made a right curve tightly following the perimeter of the wall. Eventually it spotted C in the distance and managed to catch up and get into position for the last 100 time steps. The T4 trial (with a % of only 21.0) also demonstrated a recovery where R re-established the formation before the last 100 time steps; i.e., in both T4 and T5 did the robots actually get back into formation.

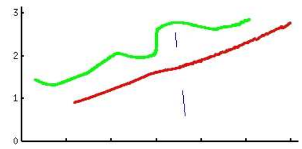

Fig. 26. Real robot data: Two robots negotiating a wall with a hole in the middle. The conductor goes through the wall, the follower gets confused and goes all around the wall, losing sight of the conductor. Upon seeing it again, the follower catches up.

No experiments with column or line are reported with two robots. The column trivially worked since this formation does not take up any lateral space beyond that of the conductor. The line resembles the incomplete diamond in having the robots next to each other. Usually it took more time for the robots to establish the line than the diamond, and since the $aheadbuffer$ is set low in the follower robot until it has found and maintained its right spot over some time, the obstacles had to be moved back further. But then they turned out to be overly close to the back wall of the workspace, and so it proved impossible to validate obstacle avoidance for the line. However, we are confident that a behavior comparable to that of the diamond would have been observed.

With some scenarios and formations, deadlocks between two robots could occur. One of the features of our current implementation of the algorithm is, as mentioned above, that if a robot loses track of its friend, it will try to get back to the heading it has when it was last in the right position with respect to its friend. Under unfortunate circumstances, two robots both in this recovery proces might be following the same obstacle perimeter in opposite directions. If they encounter each other in this situation, each with a heading pointing *between* the obstacle and the other robot, they get stuck; a turn either way would result in a collision, and robots are not allowed to back up. More thought needs to be given to the recovery behavior so that a robot can evade another up-coming robot in either direc-

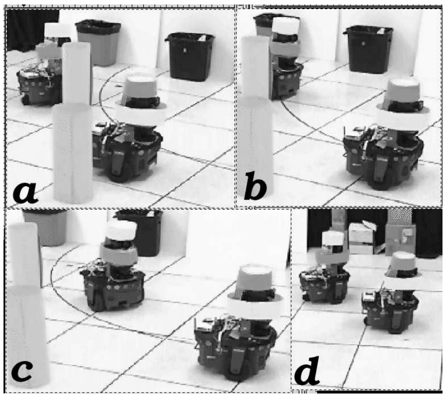

Fig. 27. The conductor, C, goes through the hole between the two pillars in image (a), while its follower, R, goes around them, momentarily losing sight of C. In (b), R sees C again, and in (c), it catches up following a straight path, its camera panned straight ahead. In (d), the incomplete diamond is re-established.

tion.

VII. Discussion and Conclusions

Our robots use only local sensing and use simple communication to communicate the global goal: formation f with N robots. Also, robots inform each other of upcoming obstacles. Hence, our system is of Type 2, according to [13]. Our key concept is that each robot follows a designated 'friend' robot at the appropriate angle and distance, using a panning sensor that can give this angle and distance to the friend. By panning the sensor an appropriate number of degrees, the algorithm simply keeps the friend centered in the sensor's view. This also enables easy switching between formations. Unique IDs and a protocol for minimalist radio communication provide robustness to drop-outs and help negotiate obstacles. A conductor that leads the way solves the problem of determining the friend's heading; by the nature of the algorithm, the only stable configuration is when all robots eventually have the same heading as the conductor. Any robot can assume the conductor role.

We implementated the algorithm equipping each robot with a laser and a panning camera to measure distance and angle to its friend; we then validated the implementation through extensive simulation trials and 40+ experiments with physical robots[5]. Our general, quantitive evaluation criteria proved useful in analyzing positional data; a formal measure like the one we propose lends itself to an easier way of comparing different trials or even different strategies.

Our method proved highly succesful for certain formations (column, diamond, and, to some extent, wedge), while the line proved more challenging. The reason for this is that keeping an angle of ±90 degrees to your friend proved to be the most difficult, while an angle of 0 degrees was the most easy. This suggests a way of measuring the difficulty of formations: the more robots in the formation that have to keep high angles to their friends, the more difficult the formation. Or, put formally:

Definition2

Given a connected geometric shape $\mathcal{G}$ to be formed by N robots R_i, with θ_i being the angle that R_i has to keep to its friend (letting the conductor's angle be 0), the *difficulty* $D(\mathcal{G})$ of $\mathcal{G}$ is defined as:

$$D(\mathcal{G}) = \sum_i^N |\theta_i|$$

Thus with four robots, ignoring the conductor, we have:

$$D(column) = 0 + 0 + 0 = 0$$
$$D(diamond) = 45 + |-45| + 45 = 135$$
$$D(wedge) = |-45| + |-90| + 45 = 180$$
$$D(line) = |-90| + |-90| + 90 = 270$$

Given the ±90 degree field of view of the friend-sensor, 270 is the maximum difficulty with four robots. Of course, one angle of 90 degrees is not necessarily twice as difficult as one of 45, but Definition 2 does provide an ordering on the formations that matches our empirical results. In Figure 28, we show the weighted averages from the stability results (section VI-C), plotted against difficulty of formation.

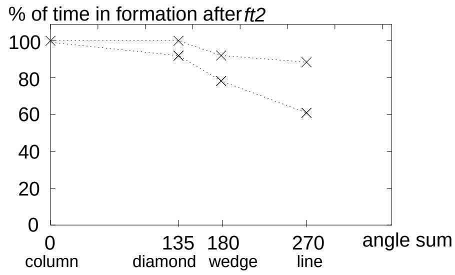

Fig. 28. Four robots; weighted average of % of time in formation after $ft2$ (over the stability trials), plotted against the difficulty of the formation. Top curve is from the simulation trials (table II), bottom curve is from the real-world trials (III).

qqq The graphs of Figure 28 capture the intuition given in the previous sections: 0-degree angles are trivial, and 90-degree angles are much harder than 45-degree angles, in the current implementation of our algorithm. There is a clear drop in performance from the diamond to the wedge, whose only 'angular' difference is that an angle of 45 degrees has been swapped with one of 90.

As explained earlier, the problem with keeping an angle of 90 degrees to a friend is two-fold: the friend may fall off the edge of the laser's range, thus forcing the distance calculation to rely on the size of a color-blob. Since this is obviously a very unreliable measure, varying with angle and from robot to robot, the resulting distance measurement will hardly be correct.

The other problem is inherent in our algorithm. If a robot R is on the side of its friend F, close to its right spot, it becomes very hard to distinguish whether R is ahead of F or behind it. Or, analogously, if R is headed away from R or towards it. Figure 29 demonstrates the ambiguity.

[5]Videos and additional imates of the formation performance are found at http://robotics.usc.edu/~agents/projects/formations.html

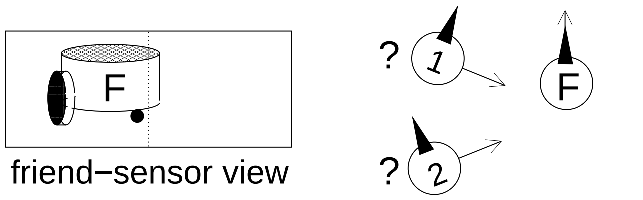

Fig. 29. Robot R has panned its friend-sensor 90 degrees to the right. The sensor view is shown on the left in the figure: R's friend F is a little off center to the left. Unfortunately, R has no way of telling whether it is headed towards F (situation 1), or away from F (situation 2). Filled triangles show heading, thin arrows show sensor pan.

In case 1, the right thing to do would be to turn away from F and slow down; in case 2, the exact opposite is needed. Unfortunately, once the wrong move is made, the same ambiguity tends to arise again. If R erroneously gets ahead of and in the way of R, congestion and confusion result. If the friend-sensor pan is less than ± 90 degrees, R will be very unlikely to ever overtake F, since in this situation, its right spot is *behind* F, not abreast with it.

With a purely reactive system based on local sensing only, with no way of telling other robot's headings, this ambiguity would have been unsolvable; the sensor view in Figure 29 alone simply does not convey enough information. (In [14], a clever way of displaying a robot's heading to its peers is devised.) With global knowledge, however, a purely reactive system could easily deal with this and related dilemmas, since R would know the position and heading of F. So what can be done in the absence of global knowledge?

We have remedied this problem somewhat by including some history of distance and bearing to the friend in the state module $WhatDoIKnow$. In case 1, for instance, if the wrong move is made, the distance to F will decrease rather than increase. Also, a 'better-safe-than-sorry' principle is adopted so that a robot will hardly ever speed up when in doubt. With a wider distance sensor range than the lasers had in our implementation, the former problem with the ± 90 degree orientation angles could be eliminated, and the exact magnitude of the latter problem could be assessed.

As long as the occasional oscillatory behavior is observed, it would be unjust to claim that our algorithm scales to any number of robots. However, the problem only arises with the infamous ± 90 angles, and so column and diamond and other less difficult formations would be reliable done by a substantially larger number of robots. Figure 30 shows a simulation trial with 8 robots establishing a diamond. We have done trials of circles and diamonds with more than 10 robots.

By the same 'better-safe-than-sorry' principle, a robot that is close to its relative position to its friend will only ever make very small rotational corrections. Therefore, small oscillations vanished, in the sense that they were not passed on to follower robots. If the ± 90 degree angle problem could be effectively reduced, and if a more general solution to the problem of switching from a non-centered to a centered formation was found, we would feel confident in claiming that our algorithm would scale to large numbers of robots. Neither problem seems insurmountable.

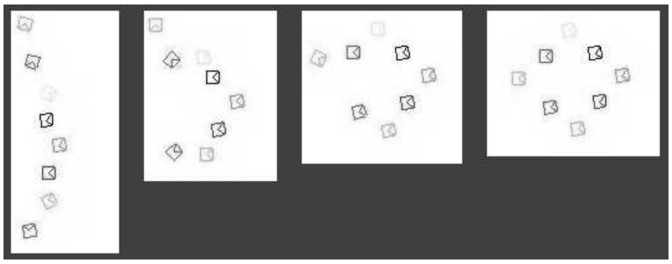

Fig. 30. A trial of 8 robots establishing a 3x3 diamond.

Many real-world issues affected the performance of our implementation of the algorithm, as is showing in Figure 28, where the simulation curve is higher than the real-robot curve. For example, since the physical shape of a robot is not completely symmetrical, when a robot perceived a distance of n millimeters to its friend, its base center was usually not at distance n from its friend's base center. Moreover, the extent of this error varied with the angle between the robots. In addition, however well our colored helmets worked, they were not perfect: depending on ambient light and perception angle, the center of their color stripes was not always aligned with the camera, thus introducing a small error in the angular positioning of a follower robot to its friend.

Alternatively, we could have used color markers that were perfectly cylindrical and placed exactly on the base center of each robot. With this solution though, a robot would have to compensate for the distance of its camera from the base center, no longer keeping the friend in the center of view but rather at some small offset from the center. For three reasons, this seems unfeasible. First, this perceived offset would vary with the angle to the friend, and so the nice structure of three layers of different levels of abstraction (Figure 6) would have to be abandoned. Second, the buffer zone of the camera field of view to each side of the friend would be smaller; especially so when keeping a 90 degree angle to the friend. And third, the marker would not be visible from all angles. With the color helmets placed *around* the camera, all three problems vanish. This can be considered another example of fully exploting a minimal system ([18]).

Upgrading our algorithm from using only Type 1 global knowledge to using Type 2 as well ([13]), immediately proved beneficial. With the swerve messages sent out by robots encountering obstacles, other robots were able to make swerves of solidarity. This, in turn, would leave place for the originator of the swerve message and thus increase the probability that a rigid-body avoidance behavior of the formation as a whole could be performed. All in all, we found that our scheme of minimal communication (heart-beat, swerve, and formation messages, all one byte) provided for a leap in performance. Since the formation problem offers no implicit environmental communication between robots, this performance leap matches the findings of [4].

We believe that our work shows that no global positioning system is needed. With a global map and global knowledge of all robots' positions, more reliable performance is possible, but this paper demonstrates that even with only local control, a group of robots can display global, coordinated behavior in the form of stable, robust, switchable formations.

Acknowledgments

The authors are grateful to Richard Vaughan, Andrew Howard, Brian Gerkey, Esben Østergaard, Boyoon Jung, and Gaurav Sukhatme for valuable help.

References

[1] R. Alur, A. Das, J. Esposito, R. Fierro, G. Grudic, Y. Hur, V. Kumar, I. Lee, J. Ostrowski, G. Pappas, B. Southall, J. Spletzer, C. J. Taylor, *A Framework and Architecture for Multirobot Coordination*, International Symposium on Experimental Robotics (ISER 2000), Hawaii, December 10-13, 2000

[2] Hideki Ando, Ichiro Suzuki, Masafumi Yamashita, *Formation and Agreement Problems for Synchronous Mobile Robots with Limited Visibility*, Proceedings of the IEEE International Symposium on Intelligent Control, Monterey, CA, August 1995, pp. 453-460.

[3] Tucker Balch, Ronald C. Arkin, *Behavior-based Formation Control for Multi-robot Teams*, IEEE Transactions on Robotics and Automation, Vol. 14, No. 6, December 1998, pp. 926–939.

[4] Tucker Balch, Ronald C. Arkin, *Communication in Reactive Multiagent Robotic Systems* Autonomous Robots, 1(1), 1995, pp. 27–52.

[5] Tucker Balch, Maria Hybinette, *Social Potentials for Scalable Multirobot Formations*, IEEE International Conference on Robotics and Automation (ICRA-2000), San Francisco, 2000.

[6] Qin Chen, J. Y. S. Luh, *Coordination and Control of a Group of Small Mobile Robots*, IEEE International Conference on Robotics and Automation (ICRA-1994), San Diego, 1994.

[7] Jaydev P. Desai, Vijay Kumar, James P. Ostrowski, *Control of Changes in Formation for a Team of Mobile Robots*, IEEE International Conference on Robotics and Automation, Detroit, May 1999.

[8] Jaydev P. Desai, James P. Ostrowski, Vijay Kumar, *A Theoretical Framework for Modeling and Controlling Formations of Mobile Robots*, IEEE Transactions on Robotics and Automations, (under review).

[9] Brian Gerkey, Kasper Støy, Richard T. Vaughan, *Player Robot Server*, Institute for Robotics and Intelligent Systems Technical Report IRIS-00-391, University of Southern California, 2000.

[10] C. Ronald Kube, Hong Zhang, *Collective Robotic Intelligence*, Second International Conference on Simulation of Adaptive Behavior, pages 460-468, December 7-11, 1992.

[11] M. Anthony Lewis, Kar-Han Tan, *High Precision Formation Control of Mobile Robots Using Virtual Structures* Autonomous Robots 4, pp. 387–403 (1997).

[12] Maja J Matarić, *Designing and Understanding Adaptive Group Behavior*, Adaptive Behavior, vol. 4, 1:50–81, MIT December 1995.

[13] Lynne E. Parker, *Designing Control Laws for Cooperative Agent Teams*, IEEE International Conference on Robotics and Automation (ICRA-1993): 582-587.

[14] Ioannis M. Rekleitis, Gregory Dudek, Evangelos E. Milios, *On Multiagent Exploration*, Proc. Vision Interface, Vancouver, BC, 1998, pp. 455-461.

[15] William M. Spears, Diana F. Gordon, *Using Artificial Physics to Control Agents*, Proceedings of the 1999 IEEE International Conference on Information Intelligence and Systems.

[16] Kazuo Sugihara, Ichiro Suzuki, "Distributed algorithms for formation of geometric patterns with many mobile robots," Journal of Robotic Systems 13(3), 1996, pp. 127–139.

[17] Richard T. Vaughan. *Stage: a multiple robot simulator*, Institute for Robotics and Intelligent Systems Technical Report IRIS-00-393, University of Southern California, 2000.

[18] Barry Brian Werger, *Cooperation Without Deliberation: A Minimal Behavior-based Approach to Multi-robot Teams*, Artificial Intelligence 110(1999) 293–320.

[19] Barry Brian Werger, Maja J Matarić, *Robotic 'Food' Chains: Externalization of State and Program for Minimal-Agent Foraging*, From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, eds. Maes, Matarić, Meyer, Pollack, and Wilson, 625–634, MIT 1996.

[20] Hiroaki Yamaguchi, Tamio Arai, *Distributed and Autonomous Control Method for Generating Shape of Multiple Mobile Robot Group* IEEE/RSJ IROS, 800–807, 1994.