

Heuristic Formation Control in Multi-Robot Systems Using Local Communication and Limited Identification

Michael de Denu and John Anderson and Jacky Baltes

Autonomous Agents Laboratory
Dept. of Computer Science, University of Manitoba
Winnipeg, Manitoba, Canada R3T2N2
<http://aalab.cs.umanitoba.ca>
mdedenu, andersj, jacky@cs.umanitoba.ca

Abstract. Groups of individuals often use formations as a means of providing orderly movement while distributing members in a manner that is advantageous to the group's activities. A particular formation may offer a defensive advantage over moving individually, for example, exposing only some of the agents to the proximity of enemies, or might increase group abilities by allowing individuals to limit perceptual focus to one small part of the environment. Formations are used throughout the natural world and in many organized human groups, and are equally valuable to multi-robot systems. Most formation control in multi-robot systems is extremely limited compared to the formations we see in nature: formations are precisely defined, and mechanisms for forming and maintaining formations often require unique labels for individuals and broadcast communication. In this paper, we explore a method for creating heuristic formations - where agents create an overall formation, but forgiveness exists for small variations in form - using only local rules for creating formations and allowing only local communication. Our approach defines relative positions in terms of a probability given the position of one's nearest neighbor, and improves on prior work by assuming that all agents do not begin knowing the unique labels of others in the group. The approach also assumes heterogeneity in sensing, in that agents may not be able to perceive the unique labels of others, and thus may require assistance from those who can. These assumptions make formations robust to the failure of individual agents, and allow previously unknown agents to join an existing formation. An evaluation of this approach is illustrated using Player/Stage, a commonly accepted simulation package for multi-robot systems, for controlled experimentation.

1 Introduction

The ability to move into and maintain formations is an important property in many groups. The potential advantages formations bring are many and diverse, and can be seen in human organizations and in the natural world. In a military unit, for example, a formation can be used as a defensive structure, exposing only a limited number of agents on the periphery to potential enemies (e.g. a square vs. a straight line). In humans this is seen in military situations as well as many sports, but the same advantage can be seen in the much less geometric formation of a school of fish. Formations may

be used to decrease the amount of work individuals must expend (e.g. aerodynamic increases achieved through the formation of a bird flock [1]) or make better use of limited sensory abilities by allowing individuals to focus only on a given area depending on their position in the formation [2]. They may also allow assumptions for increased ability to navigate and decreased decision-making time (e.g. a flight pattern for a group of jets allows one member to predict the likely positions and motion of others, decreasing the number of factors that must be considered when making a change in movement).

While these are most obviously seen in humans in military situations, they are seen throughout the natural world simply because other creatures that form groups experience significantly more vulnerable situations than do modern humans in their daily lives. For the same reason, formations are very useful for groups of robots: like non-human creatures, they are restricted in their decision-making and perceptual abilities compared to humans, and lack the common-sense knowledge to function in highly complex domains. Moreover, many of the applications that we consider robots amenable to - military environments and search-and-rescue settings, for example - present the same dangers that have led human military units to adopt the use of formations.

There are a number of important active areas of research on formations in multi-robot systems: forming various formations based on global or local interactions between robots; changing from one formation to another (again, based on global or local interactions and motivations), and maintaining formations in the face of hazards to navigation, for example. Most of this work is performed using ideal geometric patterns that are similar to those used in military applications (diamonds, squares, etc.). Comparatively little work takes advantage of the fact that formations in nature are rarely perfectly geometric: while schools of fish for example do form ellipses and other approximations to geometric shapes, for example, these are based on flexible local rules that result in a loose aggregate rather than a perfect geometric formation [3]. Similarly, while a flock of birds may form a V-shape, local rules do not necessarily dictate a precise angle, nor whether one side is symmetric with the other. From the standpoint of creating a formation, those in nature are more flexible and less constrained than that with which we are familiar with militarily. Thus, expecting a formation of robots to achieve something close a stated formation but not rigidly precise should similarly allow the benefits of formation-based movement while making the effort and infrastructure required to achieve and maintain formation reasonable.

In this paper, we describe a method for achieving formations in multi-robot systems where formations are formed heuristically, rather than to a precise, pre-defined pattern. This is done by assigning relative positions in the formation a probability, given the position of a nearest neighbor, rather than demanding an exact placement for particular individuals. The technique we employ to achieve formations requires only a simple set of local rules governing the angle and position between any agent and its nearest neighbor, and unlike other approaches, does not require all agents to know a unique identifier for all others in the group, nor to have the ability to broadcast to all members of the group. In our approach, agents begin knowing nothing about the identities of other agents, and some agents will have the ability to perceive the identities of others. Direct inter-agent communication is the only form of communication required to apply this technique to create heuristic formations in a group of robots.

We begin by reviewing related work on formation control in multi-robot systems, and then describe our approach in relation to this work. We then describe an implementation using Player/Stage [4], a well-accepted simulation system for multiple robots, and examine the performance of our approach, comparing the use of limited identification and communication with a baseline group that is given the ability to identify all individuals uniquely. Current work involves a study of the performance of this approach using Citizen Micro-robots in a mixed reality environment, and this and other future work is then discussed.

2 Related Work

In previous work, Yamaguchi [5] describes a method of formation control requiring only local information. Agents establish a link with one or two neighbors. Each agent then updates a formation vector based on the positions of its neighbors. This method is successfully demonstrated in simulation and with real robots. One notable limitation of this method is the fact that it works by maintaining relative distances only. Our work makes use of both distance and angle, allowing for a larger range of potential formations.

Balch and Arkin [2] describe a behavior-based system, in which groups of robots of known size and configuration can move together in formation. Their method relies on knowing the number of other agents and their positions. Our approach calculates the heading and speed of each individual agent as the weighted vector sum of several independent behaviors, as Balch and Arkin's approach does (this is common among behaviour-based agents). The work presented here extends that of Balch and Arkin by not requiring that each agent have a known spot in the formation. Balch and Arkin's work was also important in categorizing formation control approaches by the means with which an agent calculates its appropriate position: *Unit-Center Referenced*, where the center of the formation is determined and positions are taken relative to this; *Leader-Referenced*, where positions are taken relative to a unique leader, and *Neighbor-referenced*, where positions are taken relative to one other predetermined robot. Like many other approaches, our work on formation control is neighbor-referenced, in order to rely more on local information and avoid the bottleneck and failure-recovery problems associated with a unique leader.

Fredslund and Mataric [6] propose a method of formation control similar to that presented here. They assign each agent a unique ID. Each agent passes its ID to a function which determines its desired neighbor and the relative position at which this neighbor should be kept. While this allows for some types of formations that our method cannot currently accomplish, it relies on the ability to locate a unique individual in the group. It also does not have a mechanism for failure recovery. Our proposed method has neither of these limitations.

Howard et al. [7] perform simple formation control with heterogeneous agents. *Lower* agents have sensors to track and follow *Leader* agents. In their approach, sensing differences are strongly tied to specific roles. In our work, all agents share the same role of "formation participant". Agents can perform in this role with different levels

of success, however, depending on their sensing capabilities, and thus agents are still heterogeneous.

Hattenberger et al. [8] describe a method of dynamically adapting a formation to changing environmental conditions. The primary limitation of their technique is the dependence on a lead agent to calculate the relative positions of all other agents. Although it ensures that the intended formation is achieved, it creates a bottleneck that does not exist in our decentralized approach.

Other researchers [9, 10] discuss methods of using only local sensor information combined with simple rules to create formations. We expand on these techniques by allowing the rules to change as our knowledge of the environment grows. This allows us to create formations where different agents obey different formation conditions. The conditions that they obey can also vary dynamically. This is an advantage that neither of these systems offer.

3 Heuristic Formation Control with Limited Knowledge of Others

As mentioned previously, our approach is neighbor-referenced. This means that each agent takes a position in the formation based on that of a particular neighbor, as opposed to via a secondary frame of reference, which in turn allows rules for positioning to be defined locally. This also allows a team of agents to remain in formation when as little as one other agent is within sensor range.

In our approach to formation control, a formation is defined as a set of one or more *formation conditions*. A formation condition describes a particular relationship between two neighboring agents in an overall formation, in terms of distance and angle. A set of formation conditions must describe all types of relationships between two neighbors to describe the structure of a formation. Since the correct position to occupy may not be the same condition that is being used by a neighbor, there will be relationships between formation conditions as well. For example, a V formation consists of three different formation conditions: one describing the angular relationship on one side of the V, the other the inverse forming the other side of the V, and the third the centermost position where the agent is following no one. Agents joining the formation attempt to query their nearest neighbor (which may or may not be possible, depending on whether they know that neighbor's ID) for advice on a space to occupy in the formation. The neighbor responds with a set of probabilities indicating which formation condition(s) best describe the relationship the querying agent should physically assume if it joins the formation following the agent being queried. Representationally, a formation condition thus consists of two components: a vector specifying the desired relative angle and distance to the nearest neighbor, and a list of probabilities (one per formation condition) describing the probability that the respective condition correctly defines how an agent should position itself. In turn, the formation condition adopted then defines the answers that new agent will give to queries from future agents joining the formation.

For example, in a V formation, each agent requires the information shown in Fig. 1. Each row represents one formation condition, and the vector information used will depend on the desired size and spread of the formation, and the size of the robots. The probabilities in each formation condition represent the information that will be im-

parted to a querying agent, advising it as to which formation condition it should likely follow, given the formation condition the encountered neighbor is following. For example, if the encountered neighbor is on the left side of a V formation, the encountering agent must be also, without exception: the probability of following the same formation condition is one, and the others, zero. In a V formation, the only condition that offers an alternative is that of following the central position, since an agent following this could take either side. While this formation is simple, other formations (e.g. a diamond) offer more choice points, since the diamond will branch back in at a given position as well. In addition, there is the possibility of an implicit extra (*null*) formation condition in any formation, that is followed when no information is available as to the condition that should be followed. In a V formation, this can be defined with probabilities 0, 0, and 1 respectively, allowing a new formation to be formed with that agent occupying the center. Thus, each formation can be given a logical starting point. If at any time, an agent has no visible neighbors, it will revert to this state.

Condition	Name	Angle (degrees)	Distance (metres)	P1	P2	P3
1	Right	30	2	1	0	0
2	Left	-30	2	0	1	0
3	Center	-	-	0.5	0.5	0
4	Null	-	-	0	0	1

Fig. 1. Formation conditions for a V formation.

For the current implementation, we consider only fixed values for probabilities and vector components, as opposed to those where these values can be defined as a function of those values in a neighbor. This means that each agent potentially involved in a formation must share the same table of formation conditions.

Unlike some other approaches, ours does not assume that each agent will always be able to uniquely identify and address all others in communication. An agent that cannot uniquely identify others is limited in its ability to participate in the formation, as it cannot direct communications without an ID. Even if it could, it would not be able to precisely determine the physical origin of the response. All messages in the system, however, contain the identity of the sender, so it is possible to reply to others who initiate communication. We make use of this fact by using the *capability* message. Whenever an agent that can determine identifiers encounters a new neighbor, it sends a capability message, asking for the other agent's sensing capabilities. The original sender then uses this knowledge of capabilities to transmit useful information to that agent: in this case, a listing of all other agents that should be visible, and their IDs and positions relative to the receiving agent. The receiving agent can then use this information to its ability to communicate and participate in the formation.

When a new agent enters the vicinity of the formation, it first locates its nearest neighbor. It then attempts to communicate with this neighbor. It can only do so if it knows the ID of that agent, since all communication is directed. If communication is successful, that agent will respond with a description of the likely alternatives for the

new agent in the formation, and the agent can select one of these. If an agent cannot communicate with its nearest neighbor, it selects a random formation condition, which may cause a local aberration in the overall formation, but still allows others to build an overall approximation of the intended formation. No prior knowledge of other agents beyond the common knowledge of formations is required, and an agent need not know anything ahead of time about the size of the formation or the other agents involved.

Because an agent can find a neighbor at any given time, new agents can similarly be added in an *ad hoc* manner. Similarly, failure recovery can easily be handled within the agents themselves. When an agent fails to the point where it stops all movement, one of two things will happen: the following agent will collide (or detect a collision with) the agent it was following, or the following agent will stop as well, and no movement on its part will occur. Provided these two conditions are covered in an implementation, failure recovery is assured. An agent must view the lack of movement as not contributing toward its goals, and the stationary former neighbor as an obstacle. This will then allow the agent to begin moving and looking for a new point to join the formation. The ability to deal with failure in this manner means that no member of a formation is ever essential. If any one member fails, it can either be replaced, as described above, or the formation will adapt to its absence. If the center position in a V fails, for example, we will ultimately have two diagonal lines, and the front of each of those will no longer be following any agent, violating their formation conditions. They will both revert to the null formation condition, causing them to recognize themselves as the new center positions. They will then act independently, and will likely merge at some point.

4 Implementation

While the previous section described the overall operation of our approach to formation control, any implementation of this approach requires consideration of agent abilities and the architecture through which agents are designed. This section briefly overviews the decisions we made for the implementation used to examine the performance of the approach (though other forms of implementation are certainly possible). Because Player/Stage was chosen as a platform for evaluating our approach, largely for reasons of control, we employed agents that were easily constructed using available components in Player/Stage [4]. These are simulated Pioneer 2DX robots, using laser scanners with fiducial tracking.

Given that the core ability to form good formations is the ability to communicate directly with others, we want to examine the ability to form adequate formations using groups of agents with the ability to sense the IDs of other agents (i.e. ultimately have the ability to query them) and those that cannot. To achieve this in Player/Stage, we employ on a subset of the population a laser scanner that can read a fiducial marker attached to an agent. These form the agents that can query an ID and communicate. All agents also have sensors that can determine the distance and angle to another agent or obstacle, independent of determining identity.

In order to simulate movement towards a goal, the agents are given the ability to self-localize. This information is used to simplify movement towards a goal position. In a real-world setup, this could be replaced by some sort of goal marker, or a distributed

path planning system. For the purposes of this research, we will simply assume that some way of agreeing on a goal position exists. This localization information is used only to simplify the selection of a common goal position.

4.1 Agent Behaviors

The agents used in our implementation are behaviour-based, which involve a set of interacting weighted behaviours that ultimately determine the control values for the agent. This is also similar to the agents employed in [2]. Our agents employ three behaviors:

- Keep Formation
- Go to Goal
- Avoid Obstacles

The *Keep Formation* behavior is the heart of the approach. It computes the agent's desired relative position (in terms of a distance and angle vector) to a neighbor and compares it to the current relative position of that neighbor. The difference between these two vectors is the output of the behaviour. This generates a vector which always points towards the location in space at which the agent satisfies its chosen formation condition. The magnitude of this vector, and thus the weight it occupies in the agent's decision for movement, depends on the degree to which the agent is out of formation.

The *Go to Goal* behavior is extremely simple. It simply results in a vector, pointing towards the desired final destination of the formation. The magnitude of this vector is constant, reflecting a constant desire to move toward a goal (which is ultimately affected by the blending of the other two behaviours).

The *Avoid Obstacles* behaviour calculates a vector intended to direct the agent away from nearby obstacles, with a magnitude relative to the inverse square of the current distance to that obstacle. The only obstacles considered by this calculation are those that are visible, and within a minimum distance. This is similar to the technique used by [2] for obstacle avoidance, and common in many behaviour-based agents. The primary difference between this work and [2] is that we do not consider other agents separately from environmental obstacles. This is what causes a failed agent to be viewed as an obstacle and avoided (thus allowing an agent to separate itself from a dead neighbor and later re-join a formation).

One positive aspect of a behaviour-based approach is the ease of extending the capabilities of the agent. For example, if we were to adapt this technique to a team of searching robots, we may want to add a behaviour that would attract agents to any visible targets.

Similar to the techniques described in [2], the above behaviors each generate a vector, indicating the direction that this single behavior would cause us to go. These vectors are then scaled according to pre-defined parameters. For the purposes of this work, we used a scale factor of 2 for the *Go To Goal* behavior, 1 for the *Keep Formation* behavior and 3 for the *Avoid Obstacles* behaviour. These values were determined by trial and error, and are not necessarily optimal. The vector sum of these behaviors is then used to calculate the heading and speed of the agent.

Communication is handled outside of the behaviour-based architecture. In order to more accurately simulate real robots, direct inter-agent communication will be used.

Both types of communication will have a limited range. Yoshida et al. [11] demonstrated the feasibility of using local communication in a formation control domain. They also referenced several results, suggesting that beyond approximately 10 agents, global communication would no longer be feasible.

Because Player/Stage is a simulated environment, it would be easy to have communication go across an unrealistic range. To ensure greater correspondence with the physical world, all messages are passed through a communication server. This server tracks the absolute positions of agents and uses these to filter messages, only delivering those to agents in range. It will also allow us to experiment in future with greater communication faults and difficulties. This server also serves as a convenient place to do tracking of statistics, as it has access to the absolute positions and formation conditions of all agents.

Five types of messages are implemented: *Formation Request*, *State*, *Capability*, *Neighbor*, and *Heartbeat* messages. All messages contain the ID of the sender, to allow for replies even if the recipient cannot identify the sender through perception.

The *Formation Request* message is sent when an agent encounters a new neighbor, and is intended to elicit the formation condition probabilities from that agent, as described in Section 3. If the ID of the neighbor is unknown, this message can not be sent. In response, the neighbor sends a *State Message*, containing the formation condition probabilities described in Section 3.

Upon encountering an agent whose sensing capabilities are unknown, an agent will send a *Capability* message. An agent receiving such a message responds with a description of its sensing capabilities.

If an agent is capable of sensing unique identifiers, it can share this information by sending out *Neighbor* messages. These messages are directed at a target agent, which lacks the ability to sense unique identifiers. A neighbor message contains the polar coordinates of a sensed agent, relative to the target agent.

A *Heartbeat* message is sent from an agent to its neighbor(s) at regular intervals. This is done to help agents to track if they have unsatisfied formation conditions. A lack of a heartbeat for an established period of time indicates that a neighbor is no longer present.

4.2 Evaluation

The approach as described above was implemented using the Player/Stage [4] simulation package. Agents were modeled as Pioneer 2DX robots as described above. Player/Stage's *fakelocalize* package was used to give absolute coordinates of the agents for tracking purposes. We ran a series of trials to examine the performance of the approach in general, and to examine the effect of local communication and the number of agents that could perceive the IDs of others. For a basis of comparison, we considered two measurements of error, Error measurements were taken by a human observer, at the first point in time when every agent in the system, was at the correct relative position and angle to satisfy any formation condition. We define a local error to be an agent following a formation condition that has zero probability, given the actual condition of its neighbor. We define a global error to be a measure of difference from the ideal formation of n agents. This can be determined by finding and counting the largest group

of agents that are in positions consistent with the ideal formation, and subtracting this number from the total size of the formation.

Initial test runs showed that if the goal-seeking behavior was not strong enough, agents could deadlock, by each following one another. Increasing the weight of the goal-seeking behavior corrects this, by moving one agent out of the field of view of the other. A better solution for a future implementation would be some sort of negotiation when an agent chooses a neighbor. This could be skipped in a case where the ID of the neighbor is unknown.

It was also noticed, as one might expect, that when the desired formation distance was close to the obstacle avoidance distance, formation members would drift out of formation. This often led to changes in the nearest neighbor over time.

Our first trial consisted of a line formation with five members. Varying the number of agents able to sense IDs had no noticeable impact on the time needed to establish the line formation. Times ranged between 58 and 64 seconds, with no discernable pattern. This makes sense, as the only piece of information communicated by agents is their current formation condition, and there is only one condition in this formation. It also goes without saying that there were no errors, as the line formation has only a single condition.




Fig. 2. An example of a mostly correct V formation achieved by our system

Next, we examined ten agents in a V formation. This formation is more interesting, as it has three formation conditions. each condition describing the arms of the V is valid if the neighbor shares that condition (as per Fig. 1). In the ideal V formation, there is a single agent in the central position that is a neighbor to two agents. Qualitatively, agents who can sense IDs do generate more straight line formations, where the group is reduced to one half of the V. This is likely due to the fact that agents are more likely to communicate with agents other than the one in front. The results suggest that the num-

ber of local errors increases as the number of agents who can sense IDs decreases. This makes sense, as agents who cannot sense IDs have limited communication abilities. The results also suggest that there is a critical number of agents who cannot sense IDs, beyond which the number of relative errors is roughly constant. The results of these trials are displayed in Tables 1 and 2.

We should also note that the formations established with no ID sensing were more prone to sudden change, since agents can't uniquely identify one-another without the help of an observer. This inability to identify one-another makes them more-likely to re-calculate their desired position in the formation. These re-calculations tend to result in formations where members shift around. The ability to uniquely identify others reduces this occurrence in the other two sets of trials.

Number of agents who can sense Ids	Trial 1	Trial 2	Trial 3	Trial 4
10	1	4	4	4
5	4	4	2	2
0	3	4	4	3

Table 1. Number of global errors in the formation with ten agents

Number of agents who can sense Ids	Trial 1	Trial 2	Trial 3	Trial 4
10	0	0	0	0
5	2	0	1	2
0	2	1	1	2

Table 2. Number of local errors in the formation with ten agents

5 Discussion

We are currently working on an evaluation of this approach using a larger team of physical robots. In order to support a large team in a small area, we are using twenty Citizen Micro-robots (Fig. 3, left), each approximately an inch in size. These operate in a mixed reality environment: the robots form their own physical reality, along with any other objects that are introduced, and a virtual reality is provided by running the robots on a large horizontally-mounted monitor (Fig. 3, right), allowing a global vision system to perceive both physical and virtual elements, and precisely track the movements of robots. In prior work [12], we have shown that this approach allows better control and repeatability of experiments while allowing large numbers of small robots to operate in a physical environment. Here, this approach will allow us to substitute human judgement with computer-vision based tracking to examine the accuracy of the formations

and their adaptivity over time, as well as generate random obstacles and automatically track collisions between these and robots moving in formation. The micro-robots contain no laser scanners and currently have no local vision. Limited local vision, differing between types of agents, can be provided by restricting the viewpoint of the agents in our global vision system.




Fig. 3. Left, Citizen Micro-robots; Right, a Mixed Reality application.

One limitation of our current approach is that we consider only formation conditions where the neighbor's position is a fixed value. An interesting extension would be to allow the neighbor's position to be based on a function instead. For example, we could vary the desired angle as distance between agents changes, and create curved formations.

Some positions in a formation are more important than others: in particular, there are situations where agents may be required to share a common neighbor. One potential mechanism to deal with this is the use of *mandatory* formation conditions. In a mandatory formation condition, specific neighbors are tagged as mandatory, and if there no neighbor(s) satisfy the condition, an agent will send a message to all neighbors it can identify, requesting that these conditions be filled. Upon receiving this message, an agent can decide to fill the request, or pass it along to its neighbors. An agent will only choose to fill the request if it cannot be passed along further.

Another possible addition would be to make an agent's state depend not only on the state of its nearest neighbor, but also on any known agents already following it. This would potentially resolve some of the issues associated with branching formations. For example, if the leading agent in a V formation has a left follower, but no right follower, it should be considered to be in a different state than if it has a left follower, but no right follower. This extension, combined with the mandatory conditions described above could help to overcome some of the difficulties that this system has in reducing the number of global errors.

In this paper, we have described an approach to heuristic formation control in groups of agents with different types of sensors, and have described an implementation using ten agents. Since the approach does not rely on each agent knowing the identity of all others, the approach is robust to agent failure and adapts to adding new agents as well.

Though our current evaluation has not examined scalability, we argue that this approach should scale very well, because there are no communications bottlenecks, and the per-agent processing is not strongly related to the total number of agents in the system. In fact, we expect that results with very large numbers of agents will yield better results, as a larger number of agents increases the likelihood that all of the different conditions of the formation will be met. The adaptability and scalability should make this approach one that is useful for large numbers of agents, and environments such as search-and-rescue, where failure and agent replacement is not only possible, but expected.

References

1. Dimock, G.A., Selig, M.S.: The aerodynamic benefits of self-organization in bird flocks. In: Proceedings of the 41st AIAA Aerospace Sciences Meeting, Reno, NV (January 2003)
2. Balch, T., Arkin, R.: Behavior-based formation control for multirobot teams. *IEEE Transactions on Robotics and Automation* **14** (December 1998) 926–939
3. Sugawara, K., Tanigawa, H., Kosuge, K., Hayakawa, Y., Mizuguchi, T., Sano, M.: Collective motion and formation in simple interacting robots. In: Proceedings of the 2006 IEEE International Conference on Robots and Systems (IROS), Beijing, China (October 2006) 1062–1067
4. Gerkey, B., Vaughan, R., Stoy, K., Howard, A., Sukhatme, G., Mataric, M.: Most valuable player: a robot device server for distributed control. In: Proceedings of the 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Volume 3. (November 2001) 1226–1231
5. Yamaguchi, H.: Adaptive formation control for distributed autonomous mobile robot groups. In: Proceedings of the 1997 IEEE International Conference on Robotics and Automation. Volume 3. (April 1997) 2300–2305
6. Fredslund, J., Mataric, M.: A general algorithm for robot formations using local sensing and minimal communication. *Robotics and Automation, IEEE Transactions on* **18**(5) (2002) 837–846
7. Howard, A., Parker, L.E., Sukhatme, G.S.: Experiments with a large heterogeneous mobile robot team: Exploration, mapping, deployment and detection. *The International Journal of Robotics Research* **1**(5–6) (June 2006) 431–447
8. Hattenberger, G., Lacroix, S., Alami, R.: Formation flight: Evaluation of autonomous configuration control algorithms. In: Proceedings of the 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), San Jose, USA (October 2007) 2628–2633
9. Spears, W., Gordon, D.: Using artificial physics to control agents. In: Proceedings of the 1999 International Conference on Information Intelligence and Systems, Bethesda, MD (October 1999) 281–288
10. Lee, G., Chong, N.Y.: Adaptive self-configurable robot swarms based on local interactions. In: Proceedings of the 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), San Jose, USA (October 2007) 4182–4187
11. Yoshida, E., Arai, T., Ota, J., Miki, T.: Effect of grouping in local communication system of multiple mobilerobots. In: Proceedings of the 1994 IEEE/RSJ/GI International Conference on Intelligent Robots and Systems (IROS). Volume 2., Munich, Germany (September 1994) 805–815
12. Anderson, J., Baltes, J., Tu, K.Y.: Improving robotics competitions for real-world evaluation of AI. In: Proceedings of the AAAI Spring Symposium on Experimental Design for Real-World Systems. AAAI Spring Symposium Series, Stanford, CA (March 2009)