A Method Of Cooperative Control Using Occasional Non-Local Interactions

Brian Shucker, Todd Murphey, and John K. Bennett
College of Engineering, 430 UCB
University of Colorado
Boulder, Colorado 80309
Email: {shucker@cs, murphey@schof, jkb@cs}.colorado.edu

***Abstract*— Current approaches to distributed control involving many robots generally restrict interactions to pairs of robots within a threshold distance. While this allows for provable stability, there are performance costs associated with the lack of long-distance information. We introduce the *acute angle switching algorithm*, which allows a small number of long-range interactions in addition to interactions with nearby neighbors, without sacrificing provable stability. We prove several formal properties of the acute angle switching algorithm, including system-wide connectivity. Further, we show simulation results demonstrating the efficacy and robustness of multi-robot systems based on the acute angle switching algorithm.**

I. Introduction

Teams of cooperating robots may be capable of performing tasks that are difficult or impossible with a single robot. With recent advances in wireless communication and integration, it is becoming feasible to deploy such teams with robots numbering in the hundreds. In such a system, the control and coordination problems become challenging.

We have deleveloped a control scheme designed to address the case of a large number of robots performing a distributed remote sensing mission. Our approach builds upon previous work in distributed control systems, in which each robot implements only local interactions with neighboring robots within a set threshold distance (for example, [1], [2]). Typically, the local interactions are designed to create desirable overall properties in the system. Thus, there is global cooperation without any centralized control, which is essential for systems that must scale to hundreds of robots.

Limiting interactions to neighbors within a threshold distance allows for provable stability, but has several drawbacks. Most significantly, the robots may be partitioned into several disconnected clusters. Causes of partitioning may include environmental disturbances, failure of one or more robots, or temporary communications failure, among others. With a threshold-distance algorithm, there is no guarantee that the disconnected clusters will ever reconnect.

Our work extends the previous approach so that robots interact with selected neighboring robots at larger distances when possible, in addition to interacting with neighbors within a threshold distance. We have developed an *acute angle switching algorithm* that guarantees connectivity at all times when the robots' vision range is sufficient. The acute angle switching algorithm ensures that any disconnected clusters will reconnect if the environment allows.

In this paper, we describe the acute angle switching algorithm itself, and formally show the key properties of the algorithm. We then discuss implementation concerns and show simulation results.

In a companion paper[3] , we formally prove that the acute angle switching algorithm does not destabilize the system.

II. Related Work

There is a significant body of previous work dealing with coordination of small teams of robots, e.g.[4], [5], [6], [7], [8], [9]. More recently, there has been research into behavior-based and virtual-physics based control of large teams of robots[10], [11], [12], [1], [13], [14]. The work most closely related to our own is summarized below.

A. Behavior-based Control

Fully distributed control based upon simple local behaviors has been used in several contexts. Much of this research is based on the intuition gained from observing behaviors such as flocking in animals. In flocking situations, animals seem to draw most of their behavioral cues from the nearby flockmates. Using this observation as a basis, Brooks[10] has investigated behavior-based control extensively; Werger[11] later described the design principles of such systems. Balch and Hybinette[15] suggested the use of "attachment sites" that mimic the geometry of crystals; this is used to create formations with large numbers of robots. A variety of projects have made use of "swarm robotics," e.g., [16] and [17], to carry out simple tasks such as light tracking. Gage[12] investigated the use of robot swarms to provide blanket, barrier, or sweep coverage of an area. Several researches have used models based on the interactions of ants within a colony[18], [17], [19]. These approaches generally seek to define simple local behaviors that lead to large-scale properties that are beneficial in a particular application.

Our work seeks to extend the intuition behind behavior-based control to include small amounts of non-local information. We hypothesize that while animals in a flock mostly follow their local neighbors, they may also make use of some larger-scale observations, especially when there are few neighbors in the immediate vicinity. This inspires us to use a switching function that occasionally allows interaction over longer distances.

B. Virtual Physics

Distributed control based on virtual physics (also called "artificial physics" or "physicomimetics") has also been investigated, although not in the manner described here. Howard, Mataric and Sukhatme[1] model robots as like electric charges in order to cause uniform deployment into an unknown enclosed area. Spears and Gordon[13], [14], [20] use a more sophisticated model analogous to the gravitational force, but make the force repulsive at close range. Both of these models use switching functions based on a threshold distance. McLurkin[21] used a partially-connected interaction graph with a physics model similar to that of compressed springs to produce uniform deployment within a limited indoor environment. These works provide useful heuristic algorithms, but unlike our work, they do not attempt to show any provable properties of the resulting formations.

C. Switched Systems

Jadbabaie and colleagues used algebraic graph theory to show stability for switched networks using nearby-neighbor rules[2], [22], [23]. Hespanha and Morse used dwell-time analysis to show stability in systems with arbitrary switching that is slow on the average[24], [25], [26]. Bullo and colleagues showed stability in a switched system using Voronoi neighbors[27]. These results all differ from our work in that we use a switching function that is designed first to create specific geometric properties.

III. Acute Angle Switching Algorithm

Our system is based on simple spring-like interactions between the robots. Such dynamics are intuitive and easy to implement on a real system. Robots use the acute angle switching algorithm to determine with which other robots to interact.

For a given set of springs, the control law for each robot is

$$\ddot{\mathbf{x}} = \mathbf{u} \tag{1}$$

$$\mathbf{u} = \big[\sum_{i \in S} k_s(l_i - l_0)\hat{\mathbf{v}}_\mathbf{i}\big] - k_d\dot{\mathbf{x}} \tag{2}$$

where $\mathbf{x}$ represents the cartesian coordinates describing the robot's position, $\ddot{\mathbf{x}}$ is the robot's acceleration, $\dot{\mathbf{x}}$ is the robot's velocity, S is the set of springs connected to this robot, l_i is the length of the i'th spring, and $\hat{\mathbf{v}}_\mathbf{i}$ is the unit vector from this robot to the robot on the other end of the i'th spring. Control constants are the natural spring length (l_0), the spring stiffness (k_s), and the damping coefficient (k_d).

At every time step, the current set of springs is determined with an acute angle test. Consider a graph in which the vertices represent robots and the edges represent virtual spring connections. Each vertex has a location equivalent to the estimated location of the robot it represents. By definition, there is an edge between vertices A and B if and only if for all other vertices C, the interior angle $\angle ACB$ is acute. This creates a mesh of acute triangles. Note that the acute-angle test is equivalent to a test for the presence of any vertex C inside the circle with diameter $\overline{AB}$, which is more efficient to compute.

The graph produced by the acute-angle test is equivalent to the Gabriel graph, which was originally described in the context of geographic variation analysis[28][1].

Figure 1 shows two examples of the acute-angle test. In Figure 1(a), an edge exists between A and B, since all interior angles $\angle ACB$ are acute. In Figure 1(b), the edge does not exist because the acute-angle test fails with robot $C4$. The circle with diameter $\overline{AB}$ is also shown; it is equivalent to say that the edge does not exist because $C4$ is inside the circle.

The basic acute angle switching algorithm is fully distributed; robots need only local information in order to determine spring connections. Our simulations show empirically that basic acute-angle meshes are stable. In a companion paper[3], we describe a modified algorithm that uses an additional value called the *energy reserve*, which is propagated through the mesh at low frequency and used to restrict switching in certain cases. We show that the modified acute angle switching algorithm creates provably stable meshes.

IV. Properties of Acute-Angle Meshes

Acute-angle meshes have several desirable properties, including provable connectivity. For completeness, we present proofs of the more relevant properties. Alternative proofs of these properties, as well as additional properties, may be found in [29].

Definition 4.1: $A \bowtie B$ is a relation on robots A and B. $A \bowtie B$ iff $\forall$ robots C distinct from A and B, the interior angle $\angle ACB$ is acute. $A \bowtie B$ indicates that a spring exists between A and B.

Definition 4.2: $A \neg B$ is defined as (**not** $A \bowtie B$).

Definition 4.3: $dist(X,Y)$ represents the distance between robots X and Y.

Lemma 4.4: All robots have a spring connection to the nearest neighboring robot.

Proof: Suppose $\exists$ at least 2 robots. Pick any robot A. Then some robot B must be closest to A; that is, $\exists$ robot B, $B \neq A$, such that $\forall$ robots C distinct from A and B, $dist(A,B) \leq dist(A,C)$. We want to show $A \bowtie B$. Let b represent interior angle $\angle ABC$ and let c represent interior angle $\angle ACB$. Since $dist(A,B) \leq dist(A,C)$, we know $c \leq b$. Thus c must be acute, since $c + b < 180$ and c is the smaller of the two. This is true for any choice of robot C, which is exactly the condition that defines $A \bowtie B$. ∎

Lemma 4.5: For robots A and B, if $A \neg B$ then there exists a robot C distinct from A and B such that $dist(A,C) < dist(A,B)$ and $dist(B,C) < dist(A,B)$. That is, if A and B are not connected then some C is closer to A and to B than they are to each other.

[1] As the authors do not follow the literature in geographic analysis, we arrived at this algorithm independently. We thank the anonymous reviewer for bringing this reference to our attention.

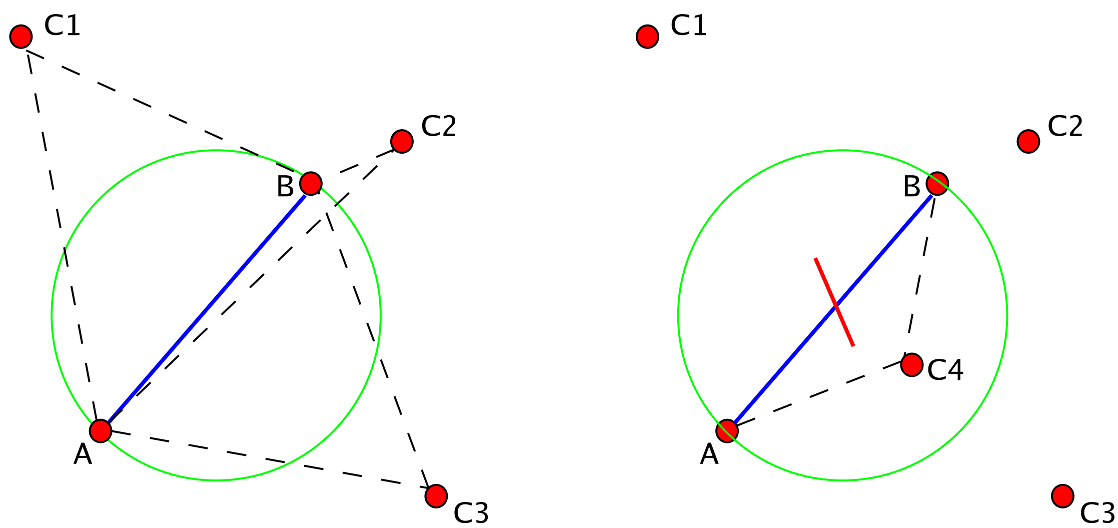

(a) Satisfied acute-angle test (b) Unsatisfied acute-angle test

Fig. 1. Illustration of acute-angle test

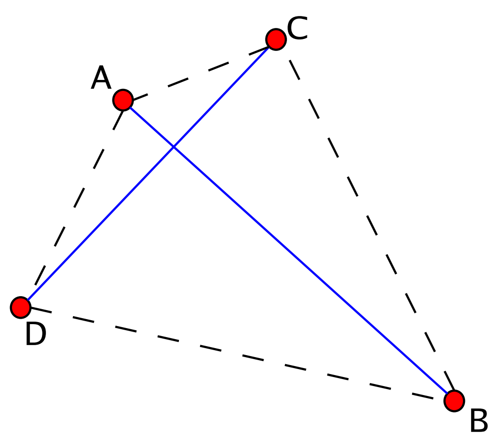

Fig. 2. Illustration of Theorem 4.6.

Proof: By definition, if $A \neg B$ then $\exists$ robot C distinct from A and B such that the interior angle $\angle ACB \geq 90$ (this is the contrapositive of the definition). With this choice of C, segment $\overline{AB}$ is the longest side of triangle ABC, since it is opposite the largest angle. Thus, $dist(A, C) < dist(A, B)$ and $dist(B, C) < dist(A, B)$. ■

Theorem 4.6: Any acute-angle spring mesh is planar.

Proof: By contradiction (see Figure 2): Suppose an intersection exists. Specifically, suppose $A \bowtie B$ and $C \bowtie D$ for distinct robots A, B, C, D and $\overline{AB}$ intersects $\overline{CD}$. Consider quadrilateral $ACBD$. Some angle in any quadrilateral must be at least 90 $\deg$. Without loss of generality, let $\angle DAC \geq 90$. Then by definition, $C \neg D$ since it fails the acute-angle test with A. This contradicts $C \bowtie D$. Since any intersection leads to a contradiction, the mesh must be planar. ■

Theorem 4.7: Any acute-angle spring mesh is connected.

Proof: Consider a spring mesh $\mathbf{M}$ partitioned into two parts, $\mathbf{M_1}$ and $\mathbf{M_2}$, so that every robot is in either $\mathbf{M_1}$ or $\mathbf{M_2}$ and there is at least one robot each in $\mathbf{M_1}$ and $\mathbf{M_2}$. It is sufficient to show that there exists a spring between some robot in $\mathbf{M_1}$ and some robot in $\mathbf{M_2}$ for any such partitioning.[2]

Pick robots $A \in \mathbf{M_1}$ and $B \in \mathbf{M_2}$ such that for any $A' \in \mathbf{M_1}, B' \in \mathbf{M_2}$, $dist(A, B) \leq dist(A', B')$. That is, pick the robots A and B with the smallest distance between them. Now we show $A \bowtie B$ by contradiction.

Suppose $A \neg B$. Then by Lemma 4.5, there is a robot C such that $dist(A, C) < dist(A, B)$ and $dist(B, C) < dist(A, B)$. Robot C must be in either $\mathbf{M_1}$ or $\mathbf{M_2}$. If $C \in \mathbf{M_1}$ then $dist(A, B) \leq dist(C, B)$ because of how we selected A and B (we are using C as A' and B as B'). However, we know $dist(C, B) = dist(B, C) < dist(A, B)$, which is a contradiction. If $C \in \mathbf{M_2}$ there is a similar contradiction, so $A \bowtie B$.

This proof is valid when there are at least 3 robots. The 2–robot case is covered by Lemma 4.4. The 1–robot case is meaningless. ■

The preceding proofs assume that vision limitations are not significant. In reality, a mesh could of course become disconnected due to great distances or environmental occlusion. However, since the connections between two partitions always include the shortest pairwise distance between robots in each partition, acute-angle meshes tend to be connected whenever the environment allows.

V. Implementation Concerns

The control scheme we have described is straightforward to implement and inherently scalable. The complexity of the

[2]This is true because if some group of robots $\mathbf{M'} \subset \mathbf{M}$ is not connected to the others, then one can set $\mathbf{M_1} = \mathbf{M'}$ and $\mathbf{M_2} = \mathbf{M} \setminus \mathbf{M'}$. This immediately leads to a contradiction, since there must be at least one spring between $\mathbf{M_1}$ and $\mathbf{M_2}$.

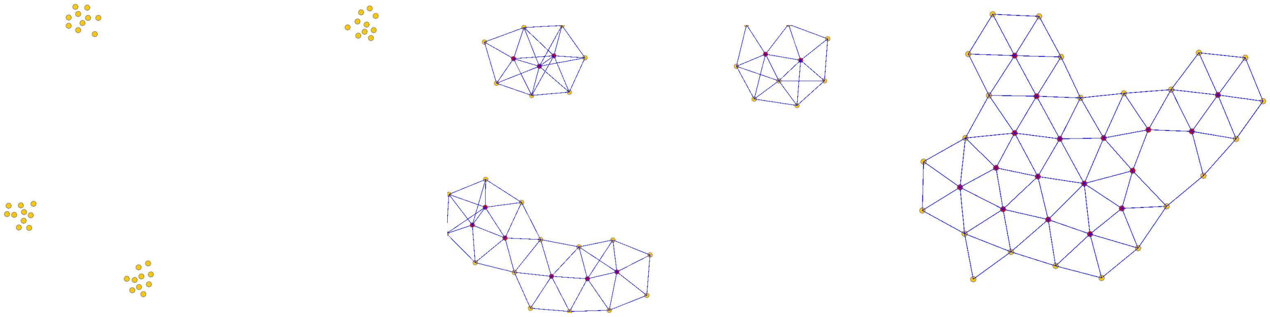

(a) Initial condition

(b) Final condition with threshold distance algorithm

(c) Final configuration with acute angle algorithm

Fig. 3. Multiple clusters of robots deploying into a single mesh

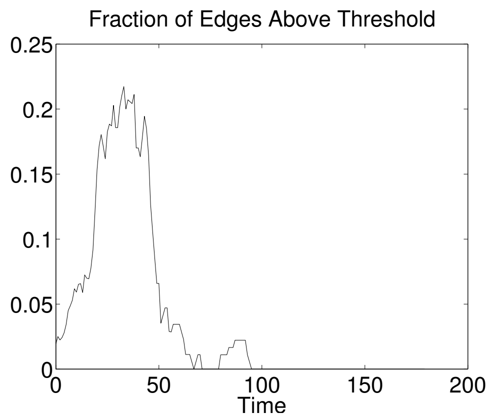

Fig. 4. Fraction of edges that are not within threshold distance, for the situation shown in Figure 3

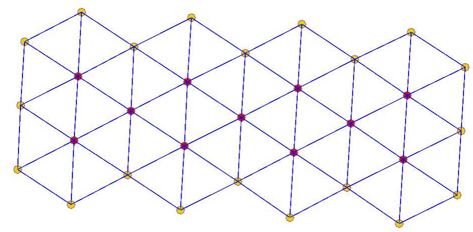

Fig. 5. Initial condition

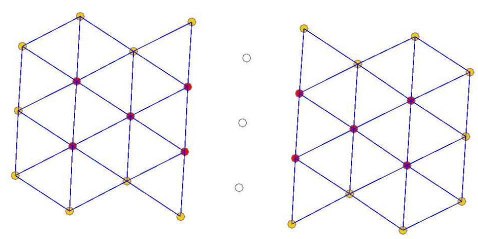

Fig. 6. After multi-robot failure, using threshold distance algorithm

algorithm running on each robot increases with the number of locally visible robots (that is, the number of neighboring robots that are currently in sight), but significantly, *complexity is independent of the total number of robots*. In fact, robots are not explicitly aware of the existence of any other robot that is not locally visible. For this reason, there is no particular limit on the number of robots that may be members of a single cluster.

There is no fundamental communication overhead associated with the control scheme. Since the acute angle switching algorithm and force computations are symmetric, it is not necessary for robots to communicate with each other in order to update the graph topology and execute the virtual physics model. However, the modified switching algorithm described in [3] does require periodic communication with local neighbors.

A. Simulation Results

We have implemented our dynamics model in simulation with both the threshold distance and acute angle switching algorithms. Figure 3 shows four separate clusters of ten robots each deploying into a single mesh. With the threshold distance algorithm, two of the clusters combine, but the other two remain separate. The acute angle switching algorithm successfully combines all four clusters into a single connected mesh. For the latter case, Figure 4 shows the fraction of the edges that are longer than the threshold distance as a function of simulation time. This quantity represents the degree to which the acute angle switching algorithm is behaving differently than the threshold distance algorithm. There are a significant number of long edges during the period when the clusters are combining, but the number of long edges decreases rapidly to zero when the combination is complete. This is not surprising—the dynamics of the system tend to drive all edge lengths to the natural length. Thus, the acute angle mesh becomes equivalent to the threshold distance mesh over time.

Figures 5, 6, and 7 show the performance of both algorithms for a simple test case involving multiple robot failure. Figure 5 shows an initial formation of 31 robots in a stable configuration. In this test case, three robots in the center of the formation fail simultaneoulsy, causing the overall mesh to be partitioned into two smaller meshes.

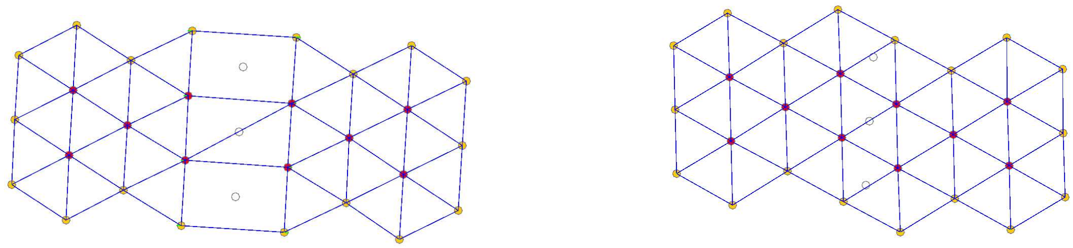

(a) Immediately after failure (b) Final configuration

Fig. 7. After multi-robot failure, using acute angle algorithm

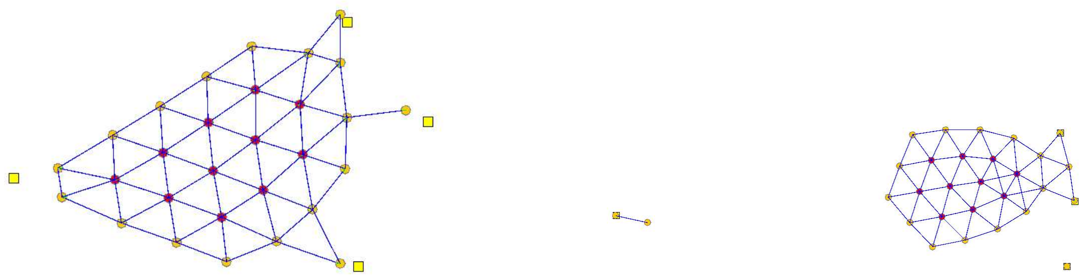

(a) Intermediate state (b) Final configuration

Fig. 9. Target tracking using threshold distance algorithm

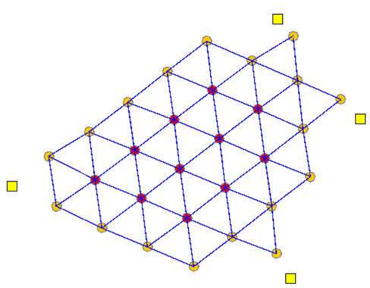

Fig. 8. Initial condition: all targets are moving laterally away from the mesh.

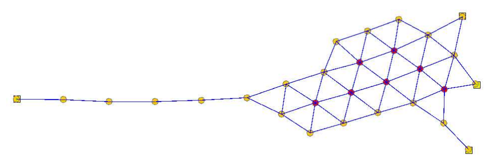

Fig. 10. Final condition using acute angle switching algorithm

As seen in Figure 6, the threshold distance switching algorithm does not cause the two smaller meshes to join, since the robots in each partition are beyond the threshold distance from each other. However, the acute angle switching algorithm will cause the partitions to join, provided that the robots have sufficient vision range (which is true by definition in this case). Figure 7 shows the acute angle mesh configuration immediately after the multi-robot failure, and the final stable configuration.

Figures 8, 9, and 10 show a situation in which the positions of the robots are externally forced by the presence of targets (denoted by squares). All targets are moving away from the mesh, as the robots near the targets move to intercept. The targets "stretch" the entire mesh, but distortion in the mesh is most significant near the targets, since the disturbance caused by the targets is split among more and more springs as it propagates inward. Thus, the longest edges are those connected to the robots that are intercepting the targets.

When using the threshold distance algorithm, this fact means that the robots tracking the targets may split off from the main group, as their connections are stretched beyond the threshold. This is illustrated in Figure 9. Such target-induced splits do not occur with the acute angle switching algorithm, since it is designed to prohibit disconnection. Instead, as seen in Figure 10, a string of robots is pulled out in the direction of each target.

VI. CONCLUSION

We have demonstrated an alternative to the standard distributed control switching algorithms based on a threshold distance. The acute angle switching algorithm creates a switched system that features a provably connected adjacency graph. Without sacrificing stability, this algorithm allows robots to take advantage of long-distance interactions whenever possible, which improves the robustness and performance of the overall system.

REFERENCES

[1] A. Howard, M. J. Mataric, and G. S. Sukhatme, "Mobile sensor network deployment using potential fields: A distributed, scalable solution to the area coverage problem," in *6th International Symposium on Distributed Autonomous Robotic Systems (DARS)*, June 2002.

[2] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Flocking in fixed and switching networks," *IEEE Transactions on Automatic Control (submitted)*, 2005.

[3] B. Shucker, T. Murphey, and J. K. Bennett, "Switching control without nearest neighbor rules," in *American Control Conference*, 2006.

[4] T. Balch and L. E. Parker, *Robot Teams: From Diversity to Polymorphism*. A K Peters Ltd, 2002.

[5] T. Balch and R. C. Arkin, "Behavior-based formation control for multirobot teams," *IEEE Transactions on Robotics and Automation*, vol. 14, pp. 926–939, 1998.

[6] T. Balch and R. C. Arkin, "Motor schema-based formation control for multiagent robot teams," in *First International Conference on Multi-Agen Systems (ICMAS)*, 1995.

[7] M. Roth, D. Vail, and M. Veloso, "A world model for multi-robot teams with communication," in *IROS-2003*, 2003. (under submission).

[8] W. Burgard, D. Fox, M. Moors, R. Simmons, and S. Thrun, "Collaborative multi-robot exploration," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2000.

[9] B. P. Gerkey, R. T. Vaughan, K. Stoy, A. Howard, G. S. Sukhatme, and M. J. Mataric, "Most valuable player: A robot device server for distributed control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2001.

[10] R. A. Brooks, "Integrated systems based on behaviors," *SIGART Bull.*, vol. 2, no. 4, pp. 46–50, 1991.

[11] B. B. Werger, "Cooperation without deliberation: A minimal behavior-based approach to multi-robot teams," *Artificial Intelligence*, vol. 110, pp. 293–320, 1999.

[12] D. W. Gage, "Command control for many–robot systems," in *Proceedings of Nineteenth Annual AUVS Technical Symposium*, June 1992.

[13] W. M. Spears and D. F. Gordon, "Using artificial physics to control agents," in *Proceedings of IEEE International Conference on Information, Intelligence, and Systems*, 1999.

[14] D. F. Gordon, W. M. Spears, O. Sokolsky, and I. Lee, "Distributed spatial control, global monitoring and steering of mobile agents," in *Proceedings of IEEE International Conference on Information, Intelligence, and Systems*, 1999.

[15] T. Balch and M. Hybinette, "Behavior-based coordination of large-scale robot formations," in *Proceedings of the Fourth International Conference on Multiagent Systems (ICMAS*, pp. 363–364, July 2000.

[16] G. Baldassarre, S. Nolfi, and D. Parisi, "Evolving mobile robots able to display collective behaviors," in *Proceedings of the International Workshop on Self-Organization and Evolution of Social Behaviors*, pp. 11–22, September 2002.

[17] E. Şahin and N. Franks, "Measurement of space: From ants to robots," in *Proceedings of WGW 2002: EPSRC/BBSRC International Workshop Biologically-Inspired Robotics: The Legacy of W. Grey Walter*, (Bristol, UK), pp. 241–247, Aug. 14-16, 2002.

[18] S. Koenig and Y. Liu, "Terrain coverage with ant robots: A simulation study," in *Proceedings of the International Conference on Autonomous Agents*, pp. 600–607, 2001.

[19] B. B. Werger and M. J. Mataric, "From insect to internet: Situated control for networked robot teams," in *Annals of Mathematics and Aritificial Intelligence*, pp. 173–197, 2001.

[20] D. F. Gordon-Spears and W. M. Spears, "Analysis of a phase transition in a physics-based multiagent system," in *Proceedings of NASA–Goddard/IEEE Workshop on Formal Approaches to Agent-Based Systems*, 2002.

[21] J. McLurkin and J. Smith, "Distributed algorithms for dispersion in indoor environments using a swarm of autonomous mobile robots," in *7th International Symposium on Distributed Autonomous Robotic Systems (DARS)*, June 2004.

[22] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Stable flocking of mobile agents, part i: Fixed topology," in *IEEE Conf. on Decision and Control*, 2003.

[23] H. G. Tanner, A. Jadbabaie, and G. J. Pappas, "Stable flocking of mobile agents, part ii: Dynamic topology," in *IEEE Conf. on Decision and Control*, 2003.

[24] J. P. Hespanha, "Extending lasalle's invariance principle to switched linear systems," in *IEEE Conf. on Decision and Control*, 2001.

[25] J. P. Hespanha and A. S. Morse, "Stability of switched systems with average dwell-time," in *IEEE Conf. on Decision and Control*, 1999.

[26] J. P. Hespanha, D. Liberzon, A. S. Morse, B. D. O. Anderson, T. S. Brinsmead, and F. D. Bruyne, "Multiple model adaptive control, part 2: Switching," *International Journal of Robust and Nonlinear Control*, 2001.

[27] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, 2004.

[28] K. R. Gabriel and R. R. Sokal, "A new statistical approach to geographic variation analysis," *Systematic Zoology*, vol. 18, pp. 259–278, 1969.

[29] D. W. Matula and R. R. Sokal, "Properties of gabriel graphs relevant to geographic variation research and the clustering of points in the plane," *Geographical Analysis*, vol. 12, pp. 205–222, 1980.