

# Analysis of Decentralized Variable Structure Control for Collective Search by Mobile Robots

Steven Goldsmith, John Feddema and Rush Robinett

Sandia National Laboratories<sup>1</sup>

Albuquerque NM 87185-5800

(505) 845-8926

(sygolds@sandia.gov, jtfedde@sandia.gov, rdrobin@sandia.gov)

RECEIVED

NOV 17 1998

OSTI

**Abstract:** This paper presents an analysis of a decentralized coordination strategy for organizing and controlling a team of mobile robots performing collective search. The alpha-beta coordination strategy is a family of collective search algorithms that allow teams of communicating robots to implicitly coordinate their search activities through a division of labor based on self-selected roles. In an alpha-beta team, alpha agents are motivated to improve their status by exploring new regions of the search space. Beta agents are conservative, and rely on the alpha agents to provide advanced information on favorable regions of the search space. An agent selects its current role dynamically based on its current status value relative to the current status values of the other team members. Status is determined by some function of the agent's sensor readings, and is generally a measurement of source intensity at the agent's current location. Variations on the decision rules determining alpha and beta behavior produce different versions of the algorithm that lead to different global properties. The alpha-beta strategy is based on a simple finite-state machine that implements a form of Variable Structure Control (VSC). The VSC system changes the dynamics of the collective system by abruptly switching at defined states to alternative control laws. In VSC, Lyapunov's direct method is often used to design control surfaces which guide the system to a given goal. We introduce the alpha-beta algorithm and present an analysis of the equilibrium point and the global stability of the alpha-beta algorithm based on Lyapunov's method.

keywords: collective robotics, variable structure control, complexity, complex systems, mobile robotics

<sup>1</sup> Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

## **DISCLAIMER**

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

## **DISCLAIMER**

**Portions of this document may be illegible  
in electronic image products. Images are  
produced from the best available original  
document.**

## I. Introduction

The problem of coordinating the actions of distributed robot teams is an active area of research (Cao, Fukunaga, Kahng, and Meng 1995). Strategies for cooperative action encompass theories from such diverse disciplines as artificial intelligence, game theory/economics, theoretical biology, distributed computing/control, animal ethology, and artificial life. For instance, Reynolds (1987) simulated the formation of flocks, herds, and schools in which multiple autonomous agents were driven away from obstacles and each other by inverse square law repulsive forces. Part of the motivation behind Reynolds' work is the impression of centralized control exhibited by actual bird flocks, animal herds, and fish schools, despite the fact that each agent (bird, animal, or fish) is responding only to its limited-range local perception of the world.

Most of these works do not include a formal development of the system dynamics. Consequently, important system properties such as stability, reachability, observability, and robustness cannot be characterized. Many of the schemes such as the Subsumption Architecture (Brooks 1986; Brooks 1991) approach rely on stable controls at a lower level while providing coordination at a higher level. This coordination is often heuristic and *ad hoc*.

We have shown how to use phase-plane analysis to describe the overall behavior of single and multiple autonomous robotic vehicles with finite state machine rules (Feddema, Robinett, and Dreissen 1998). Each finite state has a trajectory on the phase plane whose initial conditions are determined from the final conditions of the previous finite state. This interpretation of finite state machine behavior lends itself nicely to Variable Structure Control (VSC) design tools. A VSC system changes the dynamics of a system by switching at defined states (in the controls sense of the word) to an alternative function of a set of possible continuous functions of the state (Utkin, 1977; Utkin 1981; DeCarlo, Zak, and Matthews 1988). This switching may be realized with mechanical relays or simple changes of state in a computer program. In VSC, Lyapunov's direct method is often used to design control surfaces that guide the system to a given goal. The importance of this result is that we can design provably asymptotically stable group behaviors from a set of simple control laws and appropriate switching points with decentralized variable structure control techniques. The ability to prove asymptotically stable group behavior is especially important for applications such as locating military targets or land mines, where proof of convergence and an estimate of time to convergence is required.

As shown in (Feddema et al 1998), each individual finite state does not have to be convergent. However, when the transitions between states are correctly designed, a convergent stable global behavior results. This can be a powerful tool that can be used to overcome many sensing and actuation limitations and external forces. For example, consider the following analogy. Suppose you wanted to make a U-turn with your car, but the turning radius of the car was greater than the width of the street. You would start the turn, but have to stop and back up before continuing with the turn. This action involves two states: forward and backward. The forward state is convergent to the goal, while the backward state is not. You would have continued the forward state except that you were limited by an external force (the curb); you had to go into the non-convergent state to get onto a trajectory that was within the sensing and actuator limitations.

This work is an extension of new results reported by Feddema et al (1998) and Goldsmith & Robinett (1998). Work in discrete-time sliding mode control (DSMC) reported recently (Misawa 1997) is related to this proposal. The alpha-beta algorithm was developed for robotic collectives but may be extended to general adaptive state-space control algorithms for collectives of autonomous control elements. Multiple modes (sliding mode control has two, alpha-beta has three) of control corresponding to states of a finite-state automata has not been addressed in general terms for sampled-data or discrete-time control.

In this paper we focus on the alpha-beta coordination algorithm described in (Goldsmith and Robinett 1998). This algorithm features alpha agents that seek to increase their status by venturing out of local minima, and the beta agents that follow the average location of the alpha agents. This algorithm has the unique characteristic that the beta agents will group together at a local minimum, while alpha agents try to locate a gradient that leads to a global minimum. If an alpha agent finds trajectory to the global minimum, the other members will follow. It appears that the beta agents have only convergent finite states, while the alpha members may contain non-convergent finite

states. In simulation, it has been shown that this algorithm can lead to global minimum. Preliminary results of a similar approach applied to chemical plume tracking have identified important communications issues (Hurtado, Robinett, Dohrmann, and Goldsmith 1998).

## II. Alpha-Beta Coordination

We are concerned with solving the source localization problem by means of a robot team featuring a decentralized coordination strategy we call *alpha-beta coordination*<sup>2</sup>. The *alpha-beta* coordination strategy is a family of collective search algorithms that allow teams of communicating agents to implicitly coordinate their search activities through a division of labor based on self-selected roles and social status. In an alpha-beta team, an agent plays one of two complementary roles. Agents in the *alpha* role are motivated to improve their status by exploring new regions of the search space. Agents in the *beta* role are also motivated to improve, but are conservative and tend to remain aggregated and stationary until the alpha agents have clearly identified and communicated better regions of the search space. An agent selects its role dynamically based on its current status value relative to the current status values of the other team members. Status is determined by a function of the agent's sensor readings, and is generally a measurement of source intensity at the agent's current location. An agent's decision cycle comprises three sequential decision rules: (1) selection of the current role based on the evaluation of the current status data; (2) selection of a specific subset of the current data; and (3) computation of the next heading using the selected data. Variations of these decision rules produce different versions of alpha and beta behaviors that lead to different global properties.

Partitioning the robot team into alpha and beta roles produces a balance between exploration and exploitation. Alpha agents waste energy exploring low-status regions of the search space, but communicate valuable state information to team members that prevents costly reexploration of low-status regions. Alpha agents by nature seek to emulate and ultimately surpass the highest-performing team members and are therefore more sensitive to the effects of transient noise and are more susceptible to the influence of defectors reporting false status values. Beta agents use energy wisely by resisting transient influences and moving in a direct path to high-status regions of the search space identified by alpha agents. Beta agents resist noise and defectors by selective re-sampling and averaging of status data, but must rely on alpha robots to improve their performance. Consequently, beta agents can be misled by noise and defectors under some circumstances through second order effects if many of the the alpha agents are misled.

The success of alpha-beta coordination relies on the following assumptions:

1. Team members have a reliable communications mechanism.
2. The team is positioned in the (noisy) sensate region surrounding a target source.
3. The terminal goal of the team is to converge on the source target.
4. A higher status value implies a higher probability that the source is located near the corresponding coordinates.

The alpha-beta strategy is a behavior-based control strategy related to the approach of Mataric (1992;1994;1995). Alpha-beta teams behave in a manner similar to that of simple insect societies (Kube and Zhang 1993). Alpha-beta agents search without centralized leadership or hierarchical coordination. The primary collective mode of an alpha-beta team is to aggregate in a region of high intensity, without any other objectives. Alpha-beta teams are robust to single-point fail-stop failures in team members; agents simply use the latest data transmitted by other team members without regard to the identity of the sender. Alpha-beta coordination requires a minimum of knowledge about the search environment. Agents have no prior assumptions about the nature of the intensity surface, its spatial coherence, gradient field, or any other analytical information. As such, the alpha-beta strategy is intended to be as general-purpose and as assumption-free as possible.

A simple social metaphor provides an intuitively satisfying but imprecise description of the basis for alpha-beta coordination algorithms. The cohesion of an alpha-beta society is based on a common normative goal: each agent is motivated to improve its social status by associating with other agents of higher status. Social status is determined by a scalar function of the shared sensor data communicated by other agents. The only assumption underlying alpha-

<sup>2</sup> The source localization problem entails a single source phenomenon located within a bounded geographic region.

beta algorithms is that the status function orders points in the search space according to the probability that a source is located at the point. On each decision cycle, each agent broadcasts its current social status as a scalar value,  $s_i$ , along with a location vector,  $v_i$ , to all other agents, and receives their status values in return. An agent attempts to improve its standing through emulation by moving to a region occupied by agents reporting superior status. This simple goal pressures agents to: (1) aggregate into groups; and (2) to aggregate in the region of highest known status. To determine its next destination, each agent first computes the common ordered set  $V = \{v_i\}$ . i.e. the set of all location vectors, according to the linear ordering ( $\leq$ ) of agents provided by the status function.<sup>1</sup> The agent then partitions  $V$  to divide its fellow agents into two relative castes. The *alpha caste* is the set  $A_0$  of all agents that have a social standing superior to agent  $a_0$ :  $A_0 = \{v_k | s_k > s_0\}$ . The *beta caste*  $B_0$  is the set of all agents with lower social standing than agent  $a_0$ :  $B_0 = \{v_k | s_k \leq s_0\}$ . The beta set  $B_0$  includes agents of equal status because an agent always seeks to improve its current status. There are a variety of approaches to using the alpha and beta sets to generate the agent's next heading. The vectors in the set  $A_0$  can be used to influence the agent to move toward its members, creating a social pressure to improve called *alpha-pull*. The vectors in the set  $B_0$  can be used to influence the agent to move away from its members, creating a second social pressure to improve called *beta-push*. Either set or  $V$  itself can be used in a variety of ways to provide pressure to aggregate. Alpha-pull and beta push are heuristic in nature and do not necessarily lead to average improvement in arbitrary environments. Designing and testing different decision rules based on the data vectors in  $V$ ,  $A_0$ , and  $B_0$ , or subsets thereof, is the means for investigating the different global behaviors of alpha-beta teams.

A special case of importance is when  $V = A_i = B_i$ . In this case every agent has identical status, corresponding to the zero-information (maximum information entropy) state previously mentioned. When a zero-information state is detected, the team can disperse to broaden the search area by using beta-push (all members are in the beta sets of all other members) to compute a trajectory that leads the agents on the outer edges of the cohort region away from the team's centroid. As the density of the team decreases, more agents are free to move away from the centroid, eventually resulting in a dispersed team. A minimum limit on team density prevents the ultimate loss of team coherence. If the team members cannot find the sensate region, they must resort to an exhaustive collaborative search mode such as Hilbert search (Spires & Goldsmith 1998).

If  $V = \emptyset$ , the agent is alone. For the purposes of this research, agents that lose contact with the team remain immobilized. This "hug a tree" philosophy saves energy but may not lead to a reunion with the team and to eventual arrival at the target source. A variety of possible solo behaviors will be investigated later, including random search, gradient search, and using the last known heading to determine the agent's trajectory.

The general form of the alpha-beta update rule uses a linear combination of the vector data in  $V$ :

$$v_i(k+1) = v_i(k) + aT(k)[v_j(k) - v_i(k)] \quad (1)$$

where  $aT(k)$  is a transposed weighting vector derived from the application of some scalar function to the current status measurements corresponding to the vectors in  $V$ . The nature of the function applied to the status measurement vector and the specific subset of vectors selected from  $V$  determine the group behavior exhibited by this version of alpha-beta teaming.

The alpha set  $A$  contains a distinguished subset of elements: the agent or agents with the highest status value. An agent with the highest status in the cohort has no alpha caste;  $A = \emptyset$ . These agents are the  $\emptyset$ -alpha agents and cannot experience alpha-pull. The choice of a decision rule for a  $\emptyset$ -alpha agent is limited two possibilities:

1. Don't move.
2. Move away from the team along a vector derived from the  $B$ -vectors (beta-push).

---

<sup>1</sup> The unordered set of of readings can be used to compute the obvious non-uniform gradient estimates. We have investigated gradient search algorithms and use them as a baseline for comparison of alpha-beta performance. Some forms of alpha-beta algorithms currently under investigation use gradient estimates for alpha decisions.

In the first option, the  $\emptyset$ -alpha<sup>4</sup> identifies the location of highest known status and acts as a stationary beacon for the rest of the team. This is a conservative strategy that saves energy and ensures that the agent remains at the top of the heap, but does not immediately explore the region of highest intensity. The second option uses some form of beta-push to move the  $\emptyset$ -alpha away from the team. This is a risky strategy because the status of the  $\emptyset$ -alpha may decrease, but it provides more information to the team and can possibly shorten convergence time.

The beta set  $B_0$  also contains a distinguished subset of elements: the agent or agents with the lowest status value. These  $\emptyset$ -beta agents represent the social floor of the team, and always use some form of alpha-pull to improve their status.

The remaining members of the cohort have non-empty alpha and beta sets. Such an agent can experience the effects of both alpha-pull and beta-push. There are many possible decision rules for determining the next heading based on the partition  $\{A_0, B_0\}$ . In general, an agent must decide whether to be radical or conservative in its attempt to improve its status. The approach taken here is to provide three classes of behavior. For an agent team with  $N$  agents the update rules are:

1. The  $\emptyset$ -alpha agents use the conservative decision mode and remain immobile:  $v_i(k+1) = v_i(k)$ .
2. The  $m$  agents in  $V$  with the highest status values self-select alpha behavior and use the following update rule:  $v_i(k+1) = v_i(k) + \mu[v^*(k) - v_i(k)]$ , where  $v^*(k)$  is the location of a  $\emptyset$ -alpha agent, selected at random, and  $\mu$  is a factor that provides pressure to move beyond the alpha agent along a line passing through the points  $v^*(k)$  and  $v_i(k)$ .
3. The remaining  $N-m$  agents in  $V$  self-select beta behavior and use the following update rule:  $v_i(k+1) = v_i(k) + a\tau(k)[v_A(k) - v_i(k)]$ , where  $v_A(k)$  are all members of  $A_i$ , and  $a(k)$  is the corresponding vector with elements  $a_j(k) = s_j(k)/S(k)$ , and  $S = \sum s_j(k)$ ,  $j=1, N$ .

Under this regime, self-selected alpha agents attempt to exceed the performance of the stationary  $\emptyset$ -alpha agent by attempting to move to a region near the  $\emptyset$ -alpha along the vector connecting it and the  $\emptyset$ -alpha. Self-selected beta agents compute a weighted average of the alpha vectors based on normalized status values and move towards the resultant. A conservative beta agent seeks to improve its status to the average status of its corresponding alpha set by moving to the point of the center-of-mass of the alpha set. This "safety in numbers" approach provides a tendency to aggregate in the current region of highest known performance, but averages many alpha status positions to reduce noise and the influence of outliers. This behavior provides the beta population with some inertia, but still retains the tendency to improve the status of the population on average.

The important parameters in an alpha-beta regime are  $\mu$ , the search factor that determines the amount by which an alpha will attempt to move beyond a  $\emptyset$ -alpha agent, and the alpha ratio, defined as  $m/N$ , that determines the proportion of alpha agents exploring the search space. This parameter measures the degree of diversity in the population. When  $m/N$  is unity, the population comprises only alpha agents. When  $m/N$  is zero, the population comprises only beta agents.

### III. Collective Dynamics

The current state-space formulation comprises a system of linear, homogenous, time-varying difference equations of order  $N$ , where  $N$  is the instantaneous number of agents:

$$v(k) = F(k)v(k-1) \quad (2)$$

Let  $C$  be the set of all agents,  $\{a_1, a_2, \dots, a_N\}$ , whose states are described by the data pairs  $\{(z_i(k), v_i(k))\}$ , where

<sup>4</sup> Although there may be more than one  $\emptyset$ -alpha, we use the singular hereafter.

$z_i(k)$  is the status measurement vector and  $v_i(k)$  is the position vector of the  $i^{\text{th}}$  agent. Generally,  $z(k) = s(k) + w(k)$ , where  $s(k)$  is the signal and  $w(k)$  is measurement noise and other uncertainties. We will ignore  $w(k)$  in the discussions that follow and focus instead on stability under noiseless operation.

At a given step  $k$ , an agent is in one of three states:  $\emptyset$ -alpha, alpha, or beta. After sensing its local environment and obtaining remote sensor readings from other agents, the agent may transition to another state or remain the same state. Figure 2 shows the state transition map of an alpha-beta agent. Transition excitations are decision functions that determine the switching points along an agent's trajectory in state space.

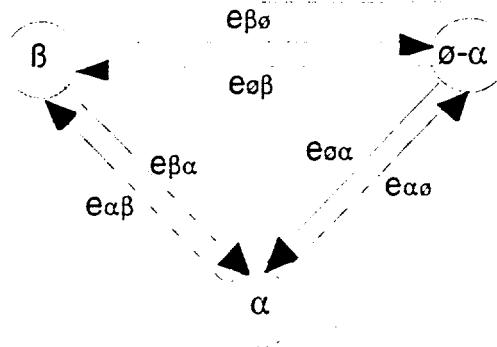


Figure 2: Alpha-Beta State Transition Diagram

The excitation equations for one form of alpha-beta coordination are:

**$\emptyset$ -alpha rule:**

Excitation  $e_{\alpha\emptyset}$ ,  $e_{\beta\emptyset}$ :  
 IF  $z_i(k) = z^*(k)$  THEN ISA( $a_i$ ,  $\emptyset$ -alpha)  
 where  $z^*(k)$  is max  $\{z_i(k)\}$

System Equation:  $v_{\emptyset}(k+1) = v_{\emptyset}(k) : v_{\emptyset}(k) = \max_{i|z_i(k)} \{[z_i(k), v_i(k)]\}$

**alpha rule:**

Excitation  $e_{\beta\alpha}$ ,  $e_{\beta\beta}$ :  
 IF  $z_{\beta}(k) < z_i(k) < z^*(k)$  THEN ISA( $a_i$ , alpha)  
 where  $z_{\beta}(k)$  is the maximum beta measurement

System Equation:  $v_i(k+1) = v_i(k) + \mu(v_{\emptyset}(k) - v_i(k))$

**beta- rule:**

Excitation  $e_{ab}$ ,  $e_{\beta b}$ :  
 IF  $z_i(k) \leq z_{\beta}$  THEN ISA( $a_i$ , beta)

System Equation:  $v_i(k+1) = v_i(k) + \sum P_{ij}(k)(v_j(k) - v_i(k))$   
 $v_j(k) \in A_i$

where  $P_{ij}(k) = z_i(k)/\sum z_{\beta}(k)$ : normalized sensor readings

and  $A_i = \{v_j(k), i \neq j, z_j(k) > z_i(k)\}$ :  $i^{\text{th}}$  alpha-cohort

The rows of the system matrix at step  $k$  reflect the independent decisions of the  $N$  agents. At step  $k$  the system state in terms of the transition matrix and initial state is

$$v(k) = \Psi(k, k_0) v(k_0) \quad (3)$$

The behavior of  $\Psi(k, k_0)$  as  $k \rightarrow \infty$  determines the stability of the system. The rows of  $\Psi(k, k_0)$  represent the transition function of the  $i^{\text{th}}$  agent at step  $k$ . An agent's behavior will be stable if its corresponding transition function is bounded for all time. The  $\emptyset$ -alpha and beta subsystems are asymptotically stable. The transition function of an  $\emptyset$ -alpha agent at step  $k$  is given by:

$$\psi_{\emptyset}(k, k-1) = 1, \text{ for initial state condition } v_{\emptyset}(k-1) \quad (4)$$

Thus  $x_{\emptyset}(k-1)$  is a stable node of the system at step  $k$ . The transition function a  $\beta$ -alpha agent is given by:

$$\psi_{\beta}(k, k-1) = [\rho_j^k \mid v_j \in A_{\beta}] \quad \text{for initial state condition } v_j(k-1) \quad (5)$$

Equation 5 represents a stable system iff all  $\rho_j < 1$ , the case that holds in a conventional alpha-beta regime. The alpha subsystem that performs exploration may or may not be stable, depending on the value of the convergence parameter  $\mu$ . The transition function of an alpha agent is given by:

$$\psi_{\alpha}(k, k-1) = [0 \ 0 \ \dots \ (1 - (1 - \mu)k) \ \dots \ (1 - \mu)^k \ \dots \ 0 \ 0] \quad (6)$$

For exploration beyond the  $\emptyset$ -alpha node at  $v_{\emptyset}(k-1)$ ,  $\mu > 1.0$  must hold. An asymptotically stable transition function occurs when  $1.0 < \mu < 2.0$ , and  $v_{\emptyset}(k-1)$  is the equilibrium point of the subsystem. The alpha subsystem is stable in the sense of Lyapunov when  $1.0 < \mu \leq 2.0$ , with an undamped oscillation about  $v_{\emptyset}(k-1)$  occurring for  $\mu = 2.0$ . Widespread exploration to escape a local extremum or plateaus may require that  $\mu > 2.0$ , resulting in a temporarily unstable alpha subsystem. We now give a heuristic proof of alpha subsystem stability based on Lyapunov stability. The sufficient condition for Lyapunov stability of the time-varying alpha subsystem is:

$$|\Psi(k, k_0)| < M \text{ for } k \geq k_0 \quad (7)$$

where  $|\Psi(k, k_0)|$  is the norm (largest eigenvalue) of the transition matrix and  $M$  is a finite constant. Eigenvalues with exponential growth are associated with alpha and  $\emptyset$ -alpha points in the state space as described by (6). If  $\mu > 2.0$  and (6) holds for a specific alpha agent for all  $k$ , then the alpha trajectory will be oscillatory and will increase without bound along the line  $L$  determined by  $L = v_{\emptyset}(k-1) - v_{\alpha}(k-1)$ . Assuming the intensity  $Z(v, t)$  is a Lyapunov function, i.e.,  $Z' \leq 0$  (the total derivative of  $Z$  is negative semidefinite) holds in the region of interest, and if  $v_{\emptyset}(k-1)$  is not the maximum intensity point along  $L$ , the alpha agent will eventually become a  $\emptyset$ -alpha agent upon discovering a point of higher intensity than  $v_{\emptyset}(k-1)$  along  $L$ . The agent will then be governed by (4) and will remain a stable node until surpassed by a higher intensity value.

Equation (6) will govern the agent's behavior as  $k$  approaches infinity if  $v_{\emptyset}(k-1)$  represents the point of maximum intensity (global maximum) in the region of interest. Thus instability is possible only when the global maximum has been captured by at least one agent. Adaptive gain-scheduling techniques have been evaluated that reduce  $\mu$  at convergence to stabilize the collective, but are beyond the scope of this paper.

#### IV. Conclusions

We have analyzed the stability of the alpha-beta coordination regime and we identified the conditions under which the system is unstable. Specifically, the alpha-beta regime has an unstable mode when  $\mu > 2.0$  and the global equilibrium point has been discovered by at least one agent. The analysis has lead to a better understanding of the limitations of alpha-beta with respect to the intensity surface being searched and the parameters of the search components (alpha agents). Generalization of the alpha-beta algorithm to other forms of variable structure control based on finite-state automata (VSC/FA) is currently being investigated.

#### V. References

Brooks, R. 1986. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation* RA-2.

- Brooks, R. 1991. Intelligence without reason. In *Proc. IJCAI-91*.
- Brooks, Rodney A. and Flynn, Anita M. 1989, Fast, Cheap and Out of Control: A Robot Invasion of the Solar System, *Journal of the British Interplanetary Society*, Vol. 42, pp. 478-485.
- Cao, U., Fukunaga, A., Kahng, A., and Meng, F. 1995. Cooperative mobile robotics: Antecedents and directions. In *Proc. of IEEE/RSJ IROS*.
- DeCarlo, R.A., Zak, S.H. Matthews, G.P. 1988, Variable Structure Control of Nonlinear Multivariable Systems, *Proceedings of the IEEE*, Vol. 76, No. 3, March 1988.
- Feddema, J.T., Robinett, R.D., Driessen, B.J. 1998. Designing stable finite state machine behaviors using phase plane analysis and variable structure control. In *Proceedings of 1998 IEEE International Conference on Robotics and Automation*, Belgium, May 16-21, 1998.
- Goldsmith, S.Y., Robinett, R.D. 1998, Collective Search by Mobile Robots using Alpha-Beta Coordination. In *Proceedings of the Workshop on Collective Robotics (CRW-98)*, Agent World 98, Paris France.
- Hurtado, J., Robinett, R., Dohrmann, C., Goldsmith, S. 1998. Distributed sensing and cooperating control for swarms of robotic vehicles. In *Proceedings IASTED International Conference on Control and Applications*, Honolulu, Hawaii.
- Kube, C. and Zhang, H. 1993. Collective robotics: From social insects to robots. *Adaptive Behavior*, 2(2).
- Mataric, M. 1992. Behavior-based systems: Key properties and implications. In *IEEE International Conference on Robotics and Automation, Workshop on Architectures for Intelligent Control Systems*.
- Mataric, M. 1994. Interaction and intelligent behavior. *MIT AI Lab Tech report AITR-1495*.
- Mataric, M. 1995. Designing and understanding adaptive group behavior. *Adaptive Behavior* (4:1).
- Misawa, E. A. 1997. Discrete-Time Sliding Mode Control: The Linear Case. *Journal of Dynamic Systems, Measurement, and Control*, Vol. 119.
- Reynolds, Craig W. 1987. Flocks, Herds, and Schools. *Computer Graphics*, Volume 21, No. 4, pp. 25-34.
- Spires, S., and Goldsmith, S. 1998. Exhaustive Geographic Search with Mobile Robots Along Space-Filling Curves. In *Proceedings of the Collective Robotics Workshop*, Agent World '98, Paris France.
- Utkin, V.I. 1977. Variable Structure Systems with Sliding Modes, *IEEE Transactions on Automatic Control*, Vol. AC-22, No. 2, April 1977.
- Utkin, V.I. 1981 Sliding Modes in Control Optimization, Springer-Verlag.