


INEEL/CON-2000-00430
PREPRINT



A Formation Behavior for Large-Scale Micro-Robot Force Deployment

D. D. Dudenhoeffer
M. P. Jones

December 10, 2000 - December 13, 2000

Winter Simulation Conference

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint should not be cited or reproduced without permission of the author.

This document was prepared as a account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, or any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for any third party's use, or the results of such use, of any information, apparatus, product or process disclosed in this report, or represents that its use by such third party would not infringe privately owned rights. The views expressed in this paper are not necessarily those of the U.S. Government or the sponsoring agency.

A FORMATION BEHAVIOR FOR LARGE-SCALE MICRO-ROBOT FORCE DEPLOYMENT

Donald D. Dudenhoeffer

Michael P. Jones

Center for Human-System Simulation
Idaho National Engineering and Environmental Laboratory
P.O. Box 1625
Idaho Falls, ID 83415-3779

ABSTRACT

Micro-robots will soon be available for deployment by the thousands. Consequently, controlling and coordinating a force this large to accomplish a prescribed task is of great interest. This paper describes a flexible architecture for modeling thousands of autonomous agents simultaneously. The agents' behavior is based on a subsumption architecture in which individual behaviors are prioritized with respect to all others. The primary behavior explored in this work is a group formation behavior based on social potential fields (Reif and Wang 1999). This paper extends the social potential field model by introducing a neutral zone within which other behaviors may exhibit themselves. Previous work with social potential fields has been restricted to models of "perfect" autonomous agents. The paper evaluates the effect of social potential fields in the presence of agent death (failure) and imperfect sensory input.

1 INTRODUCTION

This paper examines multi-agent modeling and simulation for a large-scale number of autonomous agents. Specifically, the paper examines the development of a simulation program to model the interaction and collective behavior of micro-robotic task forces consisting of 1000 or more entities. The objective of the research presented in this paper is twofold. The first objective is to develop a viable simulation tool for studying autonomous agent behavior and collective interaction. The second is to examine behavioral models for group formation and coordinated motion.

Technology will soon make possible the practical deployment of micro-robots on the size order of 2.5cm square. The dispersion of thousands of these micro-robots represents a tremendous capability for application in surveillance and remote sampling. However, many issues, beyond technical feasibility still need to be examined. Notable research papers by Reynolds (1987), Gage (1992), Hodgins and Brogan (1994), Kennedy and Eberhart (1998), Reif and Wang (1999), and Suzuki and Yamashita

(1999), examine one of these issues, namely that of interaction and coordinated motion among autonomous agents. This paper extends and advances previous research efforts by examining the issues of imperfect perception and entity death. The simulation framework developed for this project encapsulates the robots as autonomous entities in which capability is added or reconfigured by the addition of objects (behaviors, sensors, physical characteristics, etc.). The simulation described in this paper was developed using Extend™.

The paper begins by presenting a vision for future micro-robotic deployments. Some of the issues involved in coordinating the group's behavior are also identified. Next the paper discusses the development of the agent model and construction of the simulation. Thirdly, movement coordination is demonstrated using a force function between neighboring entities. Finally, experimental results and future work are discussed.

1.1 Research Motivation

Technological advances in micro-robotics, remote sensors, and artificial intelligence continue to increase the capabilities of micro-robots while decreasing the size of such units. It is easy to imagine producing and deploying thousands of inexpensive, essentially disposable micro-robots in the near future. Although possibly limited in individual capability, deployed in large numbers their cumulative ability represents a tremendous force. Given the proper social behavior set, the agents form a collective; much like a colony of ants or swarm of bees. Importance shifts from the actions of individual agents to the collective behavior. A complex system develops. A complex system is defined as "one whose component parts interact with sufficient intricacy that they cannot be predicted by standard linear equations; so many variables are at work in the system that its overall behavior can only be understood as an emergent consequence of the holistic sum of all the myriad behaviors embedded within." (Levy 1992).

The autonomy, social interaction and system complexity elevate the robots from mere mechanical agents to a point where they approach a semblance of artificial life (AL).

A key element of micro-robot deployment is the construction of behavior sets that facilitate group formation and coordinated motion of micro-robotic forces. This paper examines the utility of Social Potential Forces in maintaining spatial relationships in the face of micro-robot deaths and sensor imperfections. Additionally, this paper introduces an adaptation, which promotes preferential expansion of the collective in a specific direction.

1.2 Application Scenario

To understand the complexities of micro-robot deployment and the need for coordinated control, consider the following scenario. A train with several cars containing hazardous material (HAZMAT) has derailed. A subsequent explosion has scattered debris around the crash site. The extent of the area contamination and the status of the cars carrying the HAZMAT are not known. Due to the dangerous nature of the scattered material, a task force of micro-robots is selected for site evaluation prior to human entry into the area. An Unmanned Aerial Vehicle (UAV) drops a group of 2000 micro-robots equipped with chemical sensors over the debris field and in the vicinity of the overturned cars. As the robotics coordinator, your mission is to utilize the micro-robots to map the contamination levels in the area and assess the hazards before humans enter.

This scenario represents a practical and seemingly simple application of micro-robots. However, it reveals many of the technical challenges involved in deploying large numbers of micro-robots. These challenges include the following questions.

1. Given a random or batch distribution of the 2000 micro-robots, how do you organize them into a nearly equidistant formation to maximize sensor coverage of the suspect area?
2. Once, the pattern is formed, how do you direct the masses in coordinated motion in the direction of interest?
3. During the formation and sensor sweep, how do you identify and then adjust for the inevitable "deaths" (unit failures) in order to ensure complete coverage?

This example illustrates the nature of some of the collective tasking and coordinated motion issues that still face large-scale micro-robot force deployments. These issues of autonomous behavior and command and control are also faced when using larger robots, but the problem is much more complex for several reasons. First, the diminutive size of the units severely restricts their capability in terms of computational power, memory storage, sensor coverage, etc. Second, the sheer size of the collective, one thousand

or more, makes individual control of the units by an operator unfeasible.

This project draws inspiration from observations of social communities in nature, i.e. bees, ants, birds, fish. The approach taken is to assign each unit a simple set of individual behaviors. Within this framework, each unit senses and reacts to its environment and other units within the group. The units together form a collective behavior. The model presented in this paper encapsulates a behavior set that promotes group formation building and coordinated motion. These capabilities address some of the deployment issues illustrated in the HAZMAT scenario by giving the individuals some autonomous behavior that helps accomplish the prescribed mission.

2 PROJECT DEVELOPMENT

Simulation unfortunately is a necessary evil.
(Arkin 1995)

Simulation, although accepted in many disciplines as an essential tool for gaining insight into system operation, has many skeptics among researchers in the robotics field. The prevailing thought is that the only true test of a new system design is to implement that system on an actual robot and evaluate the robot in a real world environment. No one can deny that this is indeed the ideal method for complete evaluation. However, costs, resources, time and even technology often limit the feasibility of conducting real world testing. In these instances, simulation plays a valuable role in evaluating advanced concepts and designs. Likewise, when resources become available for field-testing, maximum benefit can be achieved by focusing tests on the critical elements identified through prior simulation.

2.1 Simulation of Micro-Robots

Many concerns of the robotics community are the result of prior simulations that have unrealistically represented the challenges of real world robotic system deployment. Many of these concerns are valid and have reduced the usefulness of simulation as a development tool. The following list details some of the specific concerns raised regarding previous robotics modeling and simulation. The list was compiled from the literature and personal discussions with several prominent roboticists.

1. Simulated robots live forever. (immortality)
2. Simulated robots see everything. (perfect sensing)
3. Simulated robots possess unlimited computational ability.
4. The computer code used to drive the simulation does not resemble the same program code used to drive the actual robot. (Arkin 1995)

A simulation should take into account the realistic performance of sensors and the effects of environmental conditions. The level of realism should be related to the focus area and goals of the experiment. Concept exploration may not require a complete world model in terms of exact duplication of the environment, but it should adequately reflect the capabilities of the components in question. Micro-robotic production is not to a point where full-scale deployment is possible, but the simulation described later tries to address the above issues in terms of estimated capabilities. The capabilities of the simulated micro-robots are easily adjusted within the simulation to increase realism as more accurate information becomes available.

2.2 Foundations of Micro-Robotic Agent Modeling

The modeling and simulation of large-scale forces (1000 or more units) of autonomous agents is a relatively new area of research. The roots of such simulation can be traced back to the concept of cellular automata (CA) first conceived by John von Neumann and his colleague Stanislaw Ulam when they were exploring the realm of Artificial Life and self-reproducing automaton. According to Levy (1992), Arthur Burks actually coined the phrase “Cellular Automata” while editing von Neumann’s papers on the subject.

In the late 1960s, a University of Cambridge mathematician, John Horton Conway, took the concept of CA and developed the game of *Life*. This game inspired and influenced generations of researchers in the realm of AL and autonomous agent research. The game of *Life* consists of a two dimensional grid on which entities exist within individual cells. The entities have one of two states, alive or dead. The game traces the generations of entities as they are born, live and die. The state of an entity in the next generation is based solely on the number and states of the neighboring entities in the eight cells adjacent to the entity in question (Levy 1992).

The game of *Life* illustrates some of the key concepts used for autonomous robot modeling and multi-agent simulation. These core concepts include:

1. Identification of and focus on individual entities.
2. A defined rule set governs individual behavior.
3. Individual entities are directly affected by neighboring entities.

These concepts establish the principles for multi-agent model development and they serve as the basis for the design of our micro-robot simulation.

Throughout the remainder of the paper the terms “agent” and “micro-robot” are used interchangeably. The reason for this is twofold. First, we wish to shift focus away from any preconceived mental models and the associated limitations that arise when picturing a robot. The main focus is on the agent’s behavior or psyche. Second,

we want to avoid limiting the simulation principles discussed to just micro-robotic applications. The principles and method discussed in this paper can be adapted easily to other modeling applications. For example, an automated highway system might employ autonomous automobiles that implement similar group spacing behavior.

3 SIMULATION MODEL CONSTRUCTION

The multi-agent simulation was developed using Extend™ by Imagine That, Inc. Extend™ is primarily known as a process simulation language and has not been used to a great degree for multi-agent simulation. However, Extend™ offered a means to develop a prototype system without the overhead of developing a complete simulation environment. The C-like nature of Extend™’s MODL language also supports the project goal of developing simulation code that could be ported rather than re-written when programming real robots for evaluation.

3.1 Environment Model

The simulation is similar to the game of *Life* discussed earlier. The simulation runs in a two dimensional plane represented by a global grid (x, y coordinate system). Additional grids may be layered with the global grid to represent obstacles, terrain, sensory objects, weather, and other environmentally significant elements. The grids may be static or dynamic and reflect changing elements in the environment. This layering of grids allows rapid manipulation of scenarios for elaborate “What if?” analysis. The current work does not include any additional environmental grids.

Agents within the model move within the spaces defined by the global grid. At each simulation step, an agent may move no further than a single block away from its present position. Thus, it may remain in position or move to one of the eight adjacent (horizontal, vertical and diagonal) spaces.

A master status table contains agent position and state information. This table, coupled with information from the global grid, is used to simulate sensing of neighboring entities. The individual sensors modeled on a particular agent can determine the range and direction of other entities within their prescribed sensor range. A more detailed discussion of the sensor model is described in section 3.3.

3.2 Agent Model

In addition to the concepts from the game of *Life*, the principles of complex reactive systems influenced the agent model’s construction. Arkin (1998) defines a reactive system as one that “tightly couples perception to action without the use of intervening abstract representation or time history.” Reactive systems place little emphasis on

planning and utilize agent behavior sets as their core building blocks. Simply speaking, reactive system and behavior based systems sense the world and react.

3.2.1 Agent Model Principles

The agent simulation model is based on the premise that in the near future technology will allow the production and deployment of large-scale masses of micro-robots. The robots will be small. They will likely possess only basic capabilities and mission specific sensors. Direct communication between agents may or may not exist. The maturity of this technology does not yet exist. As a basis for modeling the capabilities of these future agents, the academic experimental robot GrowBot by Parallax, Inc. was used. The GrowBot provides a good test platform in that it was capable, but not too capable in terms of computational power and sensor configuration. Specific assumptions about individual capabilities will be described in the corresponding discussion of those capabilities.

The simulation design is very “object oriented” in its approach to agent construction. Sensors and behaviors are encapsulated when possible. This approach allows individual components to be added and removed from the model as if the corresponding physical component were being added to or removed from a real agent. This modular design permits rapid capability reconfiguration during concept exploration. Additionally, a very conscious effort was made to separate the “simulation artifacts” from the logic code being evaluated. For example, we attempted to account for the real-time parallel nature of individual entity behavior while running in a sequential simulation environment. Furthermore, the boundaries between the true state of the environment and that which can be perceived by an agent are clearly maintained. The goal was to create simulation code that could be ported rather than re-written when programming real robots for evaluation.

3.2.2 Basic Agent Model

The model of an autonomous micro-robot is constructed by building upon a base autonomous agent object. The basic model of the agent can be thought of as simply a physical shell. In abstract programming terms it may also be thought of as an object with general capabilities. The basic agent possesses only locomotion as an innate capability. The agent exists in one of three states: dead, alive or dormant. The only core capability possessed by the agent is motion, which is further restricted by speed and endurance limitations. We make a distinction between the agent’s motion capability and a behavior designed to direct or use that capability of motion.

This basic agent serves as the platform on which additional capabilities (i.e., sensors) and individual behaviors are layered. Sensors are added to the agent model by “plugging in” sensor models. The sensors query the envi-

ronment model to perceive objects or conditions of interest. The agent receives input from these sensors to increase its basic capabilities.

Similarly, new behaviors may be added to take advantage of additional sensory capability. However, it should be noted that sensors and behaviors are not the same thing nor is it necessary to have a one-to-one correspondence between sensors and behaviors. Sensors provide a means for perceiving environmental states or conditions while behaviors are the actions the agent takes based on the perceptions it makes.

Behaviors may rely on multiple sensory input (stimuli). For example, a robot’s next move may be based on the input it receives from multiple neighbor detection sensors. Similarly an individual stimulus is not necessarily unique to one behavior. Neighbor position information may be used in both a group formation behavior and a collision avoidance behavior. When behaviors conflict or compete for resources, an arbitrating mechanism usually dictates the agent’s reaction. In this way behaviors are layered. Section 3.4.1 describes the behavior arbitration model.

3.3 Sensor Model

Sensors are modeled as encapsulated object classes. The agent uses a fixed set of input and output parameters to communicate with each sensor. Consequently, multiple types and qualities of sensors may be evaluated with complete transparency to the agent model. The agents presented in this model possess two types of sensors, a Nearest Neighbor sensor and an Object Detection sensor.

This project distinguishes itself from much of the past research in this area by the attention dedicated to modeling realistic sensor capabilities. The premise behind agent interaction is that one agent can “see” his neighbor. The ability to detect and identify neighboring agents cannot be taken for granted. Adjacent agents can be identified via two methods. The first method consists of an active broadcast in which agents broadcast position information. Neighbor position may be derived from a relative coordinate system or by strength and direction of the signal. Omni-directional position data is possible.

The second method involves passive detection without open communication between agents. Neighboring agents are detected through passive sensors. Infrared sensors are an example of this type of sensor. Sensor coverage is directly tied to the number and arrangement of sensors. Detection is further dependent on the sensor’s accuracy.

The agent model uses passive detection. Each agent possesses an array of five sensors for detecting neighboring agents. Each sensor has a coverage spread of 45 degrees. Figure 1 illustrates the sensor configuration used in the model. Three additional sensors could have been added for complete 360-degree coverage. This was not done to conserve the resources that would have been consumed by

each additional sensor. These resources include power consumption, physical space, and computing (CPU) time. For this study's purposes, 360-degree coverage was not necessary because an agent does not care who is behind it.

The sensors detect the nearest agent within the sensor's coverage area. The sensor returns the relative bearing, range and type of the neighbor agent detected. Neighbor type is important because neighbor type will determine the agent's reaction to the detection. Figure 1 represents the neighborhood of the agent. The perceived neighborhood, represented by the black dots, consists of only those neighbors correctly detected by the agents. Note that only the nearest neighbor is detected if multiple neighbors exist within the same sector.

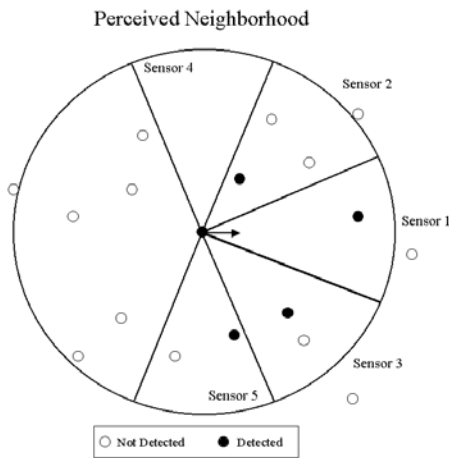


Figure 1: Agent sensor arrangement and neighborhood.

Two types of errors are modeled for each sensor, an inherent offset error and a detection error. The offset error accounts for the imperfect angular alignment of a sensor with respect to the intended relative positioning. This error is constant. The detection error represents the imperfection of the sensor and the degradation of the detection probability as a function of the detection range. The detection function is based on an exponential distribution with a mean detection range of forty inches.

The Object Detection sensor determines whether an obstruction exists along the intended path of the agent. Specifically, an obstruction is detected only if it is immediately in front of the agent. This sensor returns a signal indicating an obstacle was detected. Within the model, no errors are associated with this sensor.

3.4 Behavior Model

The most important element of agent construction is the behavior set. In developing concepts and models for individual agent and collective behavior schemas, biological entities and examples from nature were examined for insight. Important to this research project were the relation-

ship that birds and fish exhibited in flocking and schooling behaviors. Birds and fish have the ability to form and maintain collective patterns. These patterns are formed by the animal's ability to balance the desire to remain close to the flock (or school) and also to avoid collision (Shaw 1975). Within the flock, the bird does not possess universal knowledge (i.e., knowledge of the position of all others in the flock), but it adjusts its position based on the perception on its immediate neighbors. Reynolds used this framework to develop his ground breaking animation work on Boids (Reynolds 1987). These two principles of flocking and local perception provide the basis for the development of the agent's behavior.

3.4.1 Subsumption Architecture

Once a set of individual behaviors has been developed, a framework or architecture must be constructed to initiate behavioral responses and coordinate multiple behaviors. The subsumption architecture (Brooks 1986) provides the basis for behavioral coordination within the micro-robot agent model.

In simplistic terms, the subsumption architecture is based on layering reactive behavior sets on top of each other. These behaviors concurrently react to the perceived environment. A key tenant is that reaction is based on perception and not on planning. Coordination among behaviors involves a hierarchical scheme where higher level behaviors suppress or inhibit lower level behaviors. In this same way, successively more complex behaviors can seamlessly be layered onto the existing behavior set (Arkin 1998).

Figure 2 illustrates the micro-robot agent's behaviors in order of their priority. The priority goes from Collision Avoidance (highest) to Wandering (lowest).

3.4.2 Wandering Behavior

The Wandering Behavior reflects the agent's desire to move about when not under other influences. The wandering may be a random walk or motion in a predetermined direction. In the experimental results presented in Section 4, the Wandering Behavior has a predefined preference to direct the agent toward the east (i.e. right).




Figure 2: Behavior architecture

3.4.3 Group Formation Behavior

The Group Formation Behavior seeks to establish a specific spatial relationship between adjacent neighbors. The work by Reif and Wang' (1999) on Social Potential Fields provides the basis for establishing and maintaining this spatial relationship between agents within the model. In addition to the work by Reif and Wang, this paper introduces the concept of a neutral zone within the social potential field. The neutral zone permits the Wandering Behavior to activate and promotes expansion of the collective in a specific direction.

Social Potential Fields have as an underlying concept that an agent is influenced by his immediate neighbors. A force vector is used to represent the influence exerted by an agent's neighbors. The nature of the force can be attracting or repelling depending on the distance between agents. The sign and magnitude of the force is represented by the force function (Reif and Wang 1999).

Equation (1) is the force function used in the Group Formation Behavior model

$$f(d) = \frac{-c_1}{d^{\alpha_1}} + \frac{c_2}{d^{\alpha_2}} \quad (1)$$

where $c_1, c_2 \geq 0$, $\alpha_1 > \alpha_2 > 0$.

This function creates a repelling force if a neighbor is close and an attracting force if the neighbor is far away. If a neighbor is too close, the agent tends to move away and gain further separation. If the neighbor is too far away, the agent moves toward the neighbor to close the distance between them. The definitions of what is too close and too far away are arbitrary and represent flexibility in configuring the behavior depending on the mission and desired sensor coverage.

In Equation (1), d represents the range between neighboring agents. The constants, c_1 , c_2 , α_1 and α_2 determine the slope and equilibrium point of the force function. The equilibrium point is defined as the distance in which the combined effect of the repelling and attracting forces is zero.

This force function has the following characteristics:

1. Attraction is controlled by the c_2/d^{α_2} term.
2. Repulsion is controlled by the c_1/d^{α_1} term.
3. The equilibrium point where the combined effect is zero is given by $d=(c_1/c_2)^{(1/(\alpha_1 - \alpha_2))}$.

Note that this function represents the force applied by a single neighboring agent. In practice, all perceived neighbors apply forces. The resulting force is the vector summation of all the forces applied by all neighbors. Another model parameter that may be set is to have neighbors of different types that exert forces using different force functions. However, the simulation results presented in the next section use a homogenous set of agent types and hence, a single force function.

To understand the effect of multiple force vectors on a single agent, consider agent A with perceived neighbors N_1, N_2, \dots, N_k with distances d_1, d_2, \dots, d_k . The individual forces applied by the neighbors is given by:

$$f_A(d_i) = \frac{-c_{1i}}{d_i^{\alpha_{1i}}} + \frac{c_{2i}}{d_i^{\alpha_{2i}}} \quad (2)$$

The combined force applied to agent A denoted by $F(A)$ is:

$$F(A) = \sum_{i=1}^k f_A(d_i) \quad (3)$$

Equation (3) represents the force magnitude. It does not represent behavior. Behavior is the reaction to the forces applied and is realized in the agent by either the desire for motion in a certain direction or the desire to remain in place.

As stated earlier, this paper introduces an adaptation to Reif and Wang's presentation of Social Potential Fields. The adaptation is the introduction of a *critical force*. The *critical force* is defined as the magnitude of force below which the agent feels no effect. As an example, a critical force set at 5 implies that a cumulative force, $F(A)$, would require a magnitude greater than 5 to cause a reaction by the agent. By careful selection of the force function $f(x)$ and the *critical force*, a neutral zone between repelling and attracting forces is created. Within this neutral zone, no force effect exists.

This neutral zone accomplishes two purposes. First, it minimizes movement oscillations around the equilibrium point where the sign of the force changes. Second, it provides an opportunity for additional behaviors previously subsumed by the force effect and Group Formation Behavior to have an effect on the agent.

In the agent model, during periods the agent resides within the neutral zone, the Wandering Behavior dictates the desired motion of the agent. In our model, the Wandering Behavior directs the agent to head east. The agent wanders east until the critical force is again reached. At this point, the Group Formation Behavior is activated. This combination of both behaviors working in conjunction not only promotes a uniform spatial relationship between neighbors, but it also causes the entire formation to preferentially expand and move in an easterly direction.

This type of behavior readily supports a scenario in which the agents are batch dropped or are dispensed from a canister and are tasked with establishing a uniform sensor net across a specified area.

3.4.4 Collision Avoidance

Prior to repositioning, an agent will look ahead at the position of his next intended move. If another agent or obstruction is detected, the agent will evaluate a position 90 degrees to the right of the intended position. Again the

agent evaluates this position. If occupied, the agent will turn 90 degrees right and repeat the process. If after turning in a circle, no move is evaluated as “safe,” the agent will remain in place for that simulation step. On the next simulation step, the process begins again.

4 SIMULATION RESULTS

This paper presents the initial development and research into behavior-based command and control for autonomous micro-robots. A primary goal of the project was to develop a simulation framework that permitted us to explore autonomous agent design and the emergence of collective behaviors. A simulation framework was developed using Extend™. At the time of this paper, the simulation is capable of modeling the interactions of over one thousand autonomous agents.

Figure 3 shows the simulation model and the basic screen. The dots displayed represent 350 randomly dis-




Figure 3: Initial dispersion – 350 agents.

persed agents. Animation is not a project focus at this time and the agents are simply represented as an oval icon occupying a pixel on the screen. The simulated world represents each pixel as an inch. Thus the distance represented between consecutive grid lines is 100 inches or just over eight feet. The individual agents are drawn a little larger than they would be if drawn exactly to scale so that they can be seen in the display.

4.1 Model Assumptions

One of the important aspects of this simulation project was to create a credible model of agent performance. In reality, technology has not reached the point of developing micro-robots capable of field deployment. Performance capabilities are therefore based on miniature robot capabilities and on reasonable approximations of expected future performance. The model facilitates capability modification to incorporate new performance data, as it becomes available.

The following assumptions are made about agent and sensor capability:

- Agent Size – one inch in length.
- Agent Speed – one inch per second (maximum).
- Agent Failure – failure equated with agent death, follows an exponential distribution with an average agent life span of 17 minutes.
- Mean range of neighbor detector – 40 inches.
- Maximum offset error of neighbor detector – 2 degrees.

4.2 Force Function and Critical Force Selection

The force function, $f(d)$, used for the examples in this paper is given by Equation 4.

$$f(d) = \frac{-30,000}{d^2} + \frac{1000}{d} \quad (4)$$

Selection of $c_1 = 30000$, implies that the equilibrium point is located at a distance of 30 inches. Figure 4 is a graph of the function. Note the slope of the function. The shape or slope of the function has a large effect on the transition between repulsing and attracting force, particularly in the combination of competing forces exerted by multiple neighbors. The repulsion force (represented by the portion of the graph below the x-axis) is larger than the attraction force. Also, the attraction force flattens out quickly such that its effects do not get too large as the distance between agents increases.

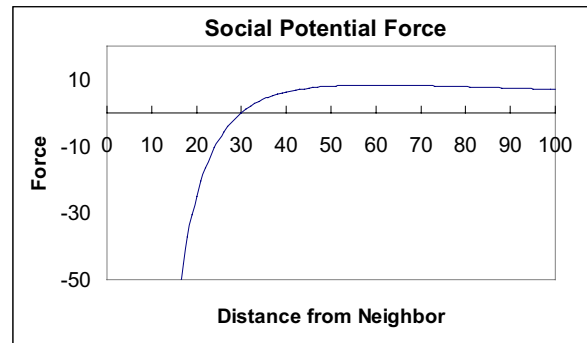


Figure 4: Social Potential Force Function

The critical force used in the simulation run is set at 5. This implies that a force magnitude, $F(A)$, of at least 5 is required to initiate a reaction by the agent. While $F(A)$ is less than the critical force, the Wandering Behavior dominates motion direction (subject to the Collision Avoidance Behavior). When the critical force is exceeded, the Group Formation Behavior takes over. Selection of the force function parameters and the critical force is based on the neighbor detection capability and desired spacing between

agents. In the model, desired average spacing is set at 30 inches with a +/- six-inch tolerance.

4.3 Simulation Runs

The major thrust of this project was to evaluate the robustness of Social Potential Fields in maintaining spatial relationships between agents when confronted with imperfect sensing and agent failure. Additionally, the project introduced the concept of a critical force which, when coupled with a Wandering Behavior, promotes coordinated motion. The following discussions illustrate the project findings. All simulation runs are initiated from the random dispersion illustrated in Figure 3. The simulation time step is set at one second. Each illustration shows the collective after 500 time steps or 8.3 minutes (simulated time).

4.3.1 The Perfect World

Figure 5 represents the dispersion pattern for the case in which sensing is perfect and agents do not fail. Note the nearly uniform spatial relationship between agents. Additionally, note the preferential expansion of the collective in an easterly direction. This “perfect world” represents a performance baseline for comparison. Note that the right-most agents have moved about 500 units (inches). The trailing agents have also progressed though a little more slowly as they are concerned about also providing complete coverage of the area being “swept.”




Figure 5: Model of a “perfect” environment.

4.3.2 Agent Death

In this run (shown in Figure 6) sensors work perfectly, but the agents are subject to death. Death includes hard mechanical failures of the agent and destruction by environmental factors (i.e. falling in a hole or being crushed by an animal). The death follows an exponential distribution with an average agent life span of 1000 time units (sec-




Figure 6: Model agent death (failure).

onds). The agents that have died are no longer displayed in the figure so the density of the collection has decreased.

4.3.3 Imperfect Sensing

This simulation run (Figure 7) represents the case in which a sensor offset error and a sensor detection probability are incorporated into the agent model. Agent failure is not modeled in this execution of the simulation. The sensor offset error is randomly distributed between +/- 2 two degrees from the sensor’s intended main axis alignment. The sensor detection probability is exponential distribution with an average neighbor detection range of 40 inches.




Figure 7: Model using imperfect sensors.

4.3.4 Combined Death and Imperfect Sensing

The combined effects of both agent death and imperfect sensors are illustrated in the simulation run presented in Figure 8.




Figure 8: Model agent failure and imperfect sensors.

4.3.5 Discussion

These simulation runs demonstrate three important results. First, the social potential field method is robust for maintaining spatial relationships when used in the presence of agent death and imperfect sensors. The dispersion patterns are similar for both perfect sensor and imperfect ones. When agents fail, the collective adjusts its shape to fill in the gaps. This behavior produces a fairly uniform coverage of the area being swept by the agents.

Second, the introduction of the critical force permits the collective to wander. That is, without the critical force, the group formation behavior always dominates the wandering behavior and the collective does little more than simply space itself out. When the wandering behavior is allowed to have an effect, the collective can be moved in a predefined direction. This result is particularly apparent in Figures 5 and 6.

The final conclusion is that motion efficiency under social potential force control is highly dependent on accurate neighbor detection. Motion efficiency is defined as the ratio of net distance traveled to total motion. The introduction of imperfect sensing reduces motion efficiency. For example, the average easterly distance traveled in the imperfect sensors case is only 53.22 inches as compared to 392.07 inches achieved with perfect sensing (see Table 1). Intuitively, one would assume that the Wandering Behavior would dominate with imperfect sensor performance. The simulation showed, however, that the agents oscillated between positions significantly more in the case of imperfect sensors. The agents were constantly readjusting their position depending on neighbor detection, lost detection and regained detection.

Table 1 illustrates average agent motion. Easterly motion is the total motion east. X motion is the total movement, east and west, along the x coordinate. The ratio of Easterly motion and X motion is equal to motion efficiency. The lower efficiency of the two simulations with

imperfect sensor configurations indicates movement oscillations back and forth with minimal advancement.

Agent Config.	Agent Motion (Average distances in inches)					
	Easterly Motion	X Dist.	X Motion	Y Dist.	Y Motion	Total Dist.
Perfect	392.07	392.07	465.07	11.94	90.25	392.41
Death	364.95	364.95	448.06	25.36	134.52	366.45
Imperfect Sensors	53.22	53.95	367.27	22.83	310.55	60.69
Death & Imp. Sensors	61.51	61.74	366.35	17.67	306.38	70.46

Table 1: Average Agent Motion

4.3.6 Future Research

This paper represents part of a continuing effort to develop behavior and control concepts for micro-robots, but also for autonomous agent constructs. Social potential fields are shown to provide robust coordinated behavior for dealing with agent death and sensor imperfections in a simulated environment. The next step is to conduct and evaluation on actual robotic platforms.

This paper and Reif's research explore social potential fields in terms of a force function that is uniform and not affected by the relative position of the neighboring agent. Further research is planned to examine the effect of sector dependent force functions which not only depend on the distance from the neighboring agent, but also on the relative angle of the neighbor's position. The use of a sector-based force function may have potential in the formation of intricate patterns with the agents.

The concept of residual forces that decay over time may be a way to address the inefficiencies and reduce motion oscillations. In this manner, the agent possesses a decaying memory of previous forces. The residual force would dampen the oscillation effect by creating some "memory" of a previously detected agent that was not detected during the current detection cycle.

Another interesting concept that must be explored further is the effect of the agents' initial distribution on their subsequent behavior and group formation. This paper examines the case of a batch distribution in which all agents are initially bunched in a small area. Future study involves reviewing the relationship between initial dispersion and the force function in constructing the desired spatial relationship and coordinated motion.

5 CONCLUSIONS

In this paper, we have described a flexible architecture for modeling thousands of autonomous agents. The agents' behavior is based on a subsumption architecture in which

individual behaviors are prioritized with respect to all others. The architecture used to model individual agents permits specific capabilities to be quickly “plugged in” and tested. Of primary interest in this research was the use of social potential fields as a mechanism for coordinated group behavior. This paper introduced the concept of a neutral zone in the social potential field and demonstrated its effect on the agents’ dispersion. Furthermore, the dispersion patterns illustrate the interaction between the social potential field and a wandering behavior operating within a subsumption architecture.

Simulation has been criticized frequently by members of the robotics community due to the too common use of models that assume perfect performance of agent and sensors. We presented the results of modeling and testing some of the real-world limitations of small-scale micro-robots. The research described here specifically investigated the effects of agent death and imprecise sensors. Initial simulation results suggests that group coordination based on social potential fields is robust to these types of real-world imperfections, but motion efficiency is relies on sensor performance.

Future efforts will explore coordinated behaviors for other mission objectives such as those outlined in the initial HAZMAT scenario. These objectives include forming a perimeter around a region and periodic, operator-induced modifications to the mission.

ACKNOWLEDGEMENTS

A special thanks is extended to the following people who provided valuable insight into robotic system modeling and development: Dr. Lynne Parker of Oakridge National Laboratory, Dr. Douglas Gage of SPAWAR, Mr. Alan Schultz of the Naval Research Laboratory and Mr. Mark McKay of the Idaho National Engineering and Environmental Laboratory.

This work is supported through the INEEL Long-Term Research Initiative Program under DOE Idaho Operations Office Contract DE-AC07-99ID13727.

REFERENCES

- Arkin, Ronald C. 1998. *Behavior-Based Robotics*. Cambridge, Massachusetts: The MIT Press.
- Arkin, Ronald C. 1995. “Just What is a Robot Architecture Anyway? Turing Equivalency versus Organizing Principles.” <ftp://ftp.cc.gatech.edu/pub/people/arkin/webpapers/stanford3.ps.Z>, Georgia Tech Mobile Robot Laboratory.
- Brooks, Rodney. 1986. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, Vol. 2 No. 1, pp.14-23.
- Gage, Douglas W. “Command and Control for Many-Robot Systems”. In *Proceedings of AUVS-92*. Reprinted in *Unmanned Systems*, Fall 1992, Vol. 10, No. 4, pp. 28-34.
- Hodgins, Jessica K. and David Brogan. 1994. Robot herds: Group behaviors for systems with significant dynamics. In *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. Cambridge, Massachusetts: The MIT Press, pp. 319-330.
- Kennedy, James and Russell Eberhart. 1998. Particle swarm optimization. In *Proceedings of the IEEE International Conference on Neural Networks IEEE Service Center, Piscataway, NJ, IV*. pp. 1942-1948.
- Levy, Steven. 1992. *Artificial Life – A report from the Frontier Where Computers Meet Biology*. New York: Vintage Books.
- Reif, John H. and Hohgyan Wang. 1999. Social potential fields: A distributed behavioral control for autonomous robots. *Robotics and Autonomous Systems*, Vol 27., pp. 171-194.
- Reynolds, Craig W. 1987. Flocks, herds, and schools: A distributed behavior model. In *Proceedings of SIGGRAPH’87*. Reprinted in *Computer Graphics*, Vol. 21, No. 4, pp.25-34.
- Shaw, E., 1975. Fish in schools. *Natural History*. Vol. 84, No. 8, pp. 4046.
- Suzuki, Ichiro and Masafumi Yamashita. 1999. Distributed anonymous mobile robots: Formation of geometric patterns. *SIAM Journal of Computing*. Vol. 28, No. 4, pp. 1347-1363.

AUTHOR BIOGRAPHIES

DONALD D. DUDENHOEFFER is research scientist for the Idaho National Engineering and Environmental Laboratory (INEEL) in Idaho Falls, ID. He received his Masters of Science degree in Operations Research from the Naval Postgraduate School in 1994. His research interests include simulation, robotics, human performance modeling, military operations and command and control. Prior to working at the INEEL, he served 12 years on active duty as a submarine officer and currently holds the rank of LCDR in the U.S. Navy Reserves. His email address is <duded@inel.gov>.

MICHAEL P. JONES is an advisory scientist for the Idaho National Engineering and Environmental Laboratory (INEEL) in Idaho Falls, ID. He received a B.S. degree in Computer Science from Brigham Young University in 1990 and M.S. and Ph.D. degrees in Computer Science from the University of Colorado, Boulder in 1992 and 1997 respectively. His current research interests include autonomous agents, intelligent interfaces for command and control, speech and language understanding and machine learning. His email address is <jonemp@inel.gov>.