

Multi-Agent Simulation for Assessing Massive Sensor Deployment

Captain Sean Hynes, U.S. Marine Corps

U.S. Naval Postgraduate School
833 Dyer Road
Monterey, CA 93943
831-656-3306
sehynes@nps.navy.mil

Neil C. Rowe

U.S. Naval Postgraduate School
833 Dyer Road
Monterey, CA 93943
831-656-2462
ncrowe@nps.navy.mil

Abstract

We present the design and implementation of a multi-agent simulation that models deployment and coverage of sensors performing collaborative target detection. The focus is on sensor networks with enough sensors that humans cannot individually manage each. Experiments evaluated both known and novel deployment algorithms, and considered effects of the sensor type, number of sensors deployed, presence of obstacles, and mobility of the sensors. A particular focus was barrier (traversal) coverage which has many military applications but which has been less studied than other sensor placement problems; experiments showed that good algorithms for it are different than those good for area monitoring. This work provides both useful data for guiding sensor deployment and a valuable testbed for planning of sensor networks.

1. Introduction

We are on the brink of a revolution in robotics and other micro-electro-mechanical systems (MEMS) that has the potential to change much including warfare. Recent advances in wireless communications and digital electronics have enabled development of low-cost low-power multifunctional sensor nodes that are small in size and communicate over short distances [AKY02]. Likewise, recent advances in robotics and distributed computing enable teams of robots numbering in the hundreds to collaboratively perform complex problem solving behavior [HOW02a] [SIN01]. Soon the numbers will be in the thousands, far beyond what is individually manageable by human users. As sensor and robotic technologies continue to advance, modeling and simulation of deployment is becoming increasingly important to conduct low-cost experiments with different configurations, settings, and applications of the technology.

Today, transmitting a bit of data over 10-100 meters by radio frequencies (RF) takes approximately 100 nano-Joules (nJ) of power. Transmitting a kilometer takes 10 to 100 micro-Joules (mJ). Soon, optical communication systems will transmit 10 meters with an energy cost of 10 pico-Joules (pJ) per bit, more than 10,000 times lower than existing radio technology. Computationally, it currently takes around 1nJ per instruction on power-optimized microprocessors, whereas it will soon take around 1pJ per instruction. Batteries provide around 1 joule per cubic millimeter. Solar cells provide approximately 100 microwatt per cubic millimeter (mm³) in full sunlight. The energy cost will be a few nano-Joules for sampling a sensor, performing some relatively simple processing, listening for incoming messages, and transmitting

a simple outgoing message. Putting all this together, a one cubic millimeter battery will provide enough power to sense and communicate once a second for 10 years, and enough energy to transmit 50 billion bits of information [PIS03].

The military is investing heavily in this technology, and it is easy to envision how small sensing devices will aid future warriors. In the future age of network centric warfare, thousands to millions of sensors will pervade the battle space. From chemical agent detection to target acquisition and biological monitoring, disparate sensors will work collaboratively to provide situational awareness and early warnings of imminent danger. Projects such as the *Smart Sensor Web* (SSW) project funded by the Defense Modeling and Simulation Office are aggressively exploring current bounds of this technology, as well as future uses [DMS03]. A Congressional mandate calls for one-third of ground combat vehicles to operate unassisted by 2015.

Current battlefields are far behind this vision of pervasive sensor networks. Ground sensors are usually emplaced by humans, who are exposed to danger in the process. These ground sensors weigh tens to hundreds of pounds, and have limited sensing lives, limited communications distance, and limited maneuverability. Nevertheless new technology is arriving. For example, one of the premier sensor vehicles produced today is quite impressive. Called "Packbot" because it is designed to fit in a pack, Tactical Mobile Robot (TMR) prototype weighs 40 pounds and costs over \$45,000. It can right itself and is waterproof to 3 meters depth. It can climb stairs and survive a 3-meter drop onto concrete. It can reasonably be assumed that in 5-10 years vehicles like Packbot will weigh less than ten pounds, run ten times longer, and cost less than \$1,000. Still this is a big step from such prototypes to extensive employment, including integration with soldiers wearing tiny computers interfaced to battlefield sensors and rear echelon commands.

As sensors and robots continue to get smaller, faster, and more power efficient, it is important to advance the software and algorithms that will take advantage of them. Modeling and Simulation (M&S) is the most cost-effective way to do so, provided we ensure an accurate mapping to and from the real world. One aspect of sensor networks that is difficult to scientifically advance without modeling and simulation is how to deploy sensors efficiently for optimal coverage. It is not feasible to test every conceivable sensor configuration in billions of physical experiments to empirically learn which sensors should be purchased and how they should be employed -- there are too many factors to consider. This work addresses this problem at the application level for sensor networks.

2. Application Area

Consider an infantry company consolidating after a protracted battle against a determined enemy. Suppose they are a part of the main force ordered to establish a defensive position on a recently assaulted objective and prepare to defend it against a counterattack. The company commander surveys the terrain, possibly urban, and orders his platoons to their defensive positions. If this were being conducted today, he would also set up observation posts and launch external patrols to help provide early warning against enemy infiltrations and attacks. A future commander might use a wearable computer to request thousands of dust-size non-mobile sensors be deployed to cover the area forward of his position. The platoon commanders could also release small man-portable autonomous vehicles to deploy more sensors to cover any gaps in coverage. Once the sensors have been deployed, this company might then use wearable computers to monitor the sensor net and visualize targets that may be detected. The wearable computers themselves become part of the sensor network; transparent to this company, the sensor network will also be sending information higher up the command hierarchy. The sensors

could physically deploy, set up ad hoc network communications, manage data, respond to failures, manage power consumption, and much more.

Many such examples exist. Today, the US could use sensor fields to help identify insurgents crossing the Syrian border into Iraq. Battlefield surveillance is just one of many scenarios imaginable for sensor networks. Sensors will also be used for monitoring friendly forces, equipment, and ammunition; reconnaissance of opposing forces and terrain; targeting; battle-damage assessment; and nuclear, biological, and chemical (NBC) attack detection and reconnaissance [AKY02]. With thousands of sensors deployed across the battlefield, human manual positioning and controlling of the sensors can be surpassed. The sensors, and the sensor networks they comprise, will need to be autonomous in every way.

3. Sensor Characteristics

A sensor is a device that implements the physical sensing of environmental phenomena. Sensors typically consist of five components: sensing hardware, memory, battery, embedded processor, and receiver/transmitter [TIL02]. If current trends continue, sensor networks will pervade the battlefield. They include large unmanned flying vehicles high in the sky equipped with hundreds of sensors and providing a high bandwidth communication backbone; medium sized-tanks autonomously searching for lost combatants; small embedded sensors attached to warriors for health monitoring; sensors attached to neurologically guided mice; and tiny artificial "dust motes" light enough to float in the air sensing for chemicals. Millions of sensors will cooperate to form robust fault-tolerant configurable communication networks, collect data, position themselves for good coverage, and answer queries about their locations. Sensor networks -- also referred to as sensor grids, expeditionary sensor grids, sensor fields, expeditionary pervasive sensing, and sensor meshes -- will be expeditionary in that they will be readily deployable, and ubiquitous.

Akyildiz [AKY02] describes important concepts of sensor networks, including a review of the architecture, algorithms, and protocols for communicating in them. That survey makes distinctions between sensor networks and traditional sensors, and between sensor networks and ad hoc networks. Traditional sensors are described as being larger than sensors in a sensor network, and requiring careful, manual positioning and communications topology engineering. The distinction between sensor networks and traditional sensors has more to do with the technical limitations than choice of design, so our work does not make that distinction. Our work assumes that any cohesive group of sensors will be networked together through wireless communications. The differences between sensor networks and other wireless ad hoc networks [PER00] are:

- The number of nodes in a sensor network is high.
- Sensor nodes are densely deployed.
- Sensor nodes are prone to failures.
- The topology of a sensor network changes frequently.
- Sensor nodes mainly use broadcast communications, whereas most ad hoc networks use point-to-point communications.
- Sensor nodes are limited in power, computational capacities, and memory.
- Sensor nodes may not have a global identification (ID) because of the large amount of overhead and large number of sensors.

So sensor network deployment and coverage algorithms must account for sensor failure and rapidly changing communication topologies, and they must be based on minimizing power consumption. Minimizing power consumption means minimizing communications; for mobile sensors, it also means minimizing movement.

Sensors can monitor a wide variety of conditions: temperature, pressure, humidity, soil makeup, vehicular movement, noise levels, lighting conditions, the presence or absence of certain kinds of objects or substances, mechanical stress levels on attached objects, and others [EST99]. Their mechanism may be seismic, magnetic, thermal, visual, infrared, acoustic, or radar [AKY02]. These mechanisms can be grouped into three categories based on how they sense: by a direct line to target (such as visual sensors), proximity to target (such as seismic sensors), and propagated like a wave to the target, possibly bending (such as acoustic sensors).

TODO: Neil, do you want to add more here?

4. Sensor Networks

Virtually every challenge encountered with software occurs in sensor-network design. Akilidiz details several sensor-network design factors, including fault tolerance, scalability, production costs, operating environment, sensor network topology, hardware constraints, transmission media, and power consumption. While these factors are important considerations in this work, only fault tolerance and scalability are addressed directly. Akyildiz gives details on communication protocols and the network stack of wireless sensor networks, but does not go beyond the scope of tiny immobile sensors. That reference mentions various deployment methods (dropping from a plane, delivering in an artillery shell, rocket, or missile, etc.), but does not discuss sensor deployment issues in detail. Sensor deployment strategies, coverage, and mobile robotics are not discussed at all.

Although he did not use the term "sensor networks" to refer to the robotic sensors, Gage [GAG92] [GAG93] [GAG92a] [DIC02] was one of the earliest researchers building mobile sensor systems and has provided many useful ideas. His approach was "to design and implement vehicle behaviors that both (a) can support real-world missions and (b) are realizable with current levels of sensor and processing technology, even in the object-and obstacle- rich environment of ground-based applications, where useful missions generally require high-bandwidth visual perception-based vehicle navigation, guidance, and control beyond the capabilities of current sensor and processing tools" [GAG92]. He envisioned mobile sensor networks of many small and inexpensive vehicles conducting military missions such as minesweeping, mine deployment, surveillance, sentry duty, communications relaying, and combat search-and-rescue (CSAR).

Gage's concept of a mobile sensor network is a large number of identical elements, each possessing [GAG92]:

- some mobility;
- some sensor capability that allows each element to measure, at least crudely, its position with respect to at least its neighbors;
- some mission-capable sensor;
- optionally, some communications capability; and

- some processing capability that directs the mobility effectors to maintain a specified positional relationship to its neighbors, as measured by its sensors, to accomplish the desired mission objectives.

This paper will address both non-mobile and mobile sensors and will build on the abovementioned works.

Sensor networks are being researched and modeled by academic groups, government groups, and commercial organizations. Most work in wireless sensor networks is not directly related to coverage and deployment of sensors, as it focuses on sensor design, efficient sensor communication, sensor data fusion, and localization [AKY02] [DHI02a]. A significant amount of research has also been conducted for mobile sensors in the field of robotics; however, most of this work has focused on issues such as obstacle avoidance, motion planning, steering, terrain maneuvering, and individual vehicle designs. There has been a surge of interest lately in multiple robot teams working together to accomplish some task or perform some behavior, but usually for less open-ended environments than in the missions encountered by the military.

4. Coverage

Gage [GAG92a] described three useful coverage behaviors: sweep coverage, area coverage, and barrier coverage. The application created for this work is capable of implementing all three.

The objective of sweep coverage is to move a number of elements across a coverage area while maximizing the number of detections per time and minimizing the number of missed detections per area. Examples include minesweeping and Combat Search and Rescue (CSAR). [CHO01] discusses coverage path planning for sweep coverage, focusing on coverage path-planning algorithms for mobile robots in a plane, though the algorithms could be extended to three dimensions. It surveys both heuristic and optimal algorithms. Gage showed randomized search strategies accomplish sweep coverage scenarios well. Randomized searches can be performed by cheaper robots since they do not require localization equipment, robust communications, or advanced computational requirements. Although random search does not guarantee complete coverage [CHO01], it is suitable for tactical scenarios where the target is moving. Additionally, optimal sweep coverage algorithms are not guaranteed to detect even non-mobile targets when the sensors are imperfect. Thus this work uses randomized search where sweep coverage is needed. Lastly, sweep coverage can be accomplished with a moving barrier (see barrier coverage below).

The objective of area coverage [HOW02] is to achieve a static arrangement of sensors that maximizes the detection of targets appearing within the coverage area. We extend this definition to include dynamic arrangements as well. Examples include detecting chemical agent attacks and providing early warning of forest fires. Area coverage is also referred to as blanket coverage, field coverage [GAG92], and grid coverage (although this has special meaning) [DHI02].

The objective of barrier coverage is to achieve a static arrangement of elements that minimizes the probability of undetected penetration from one region to another. We extend this definition to include dynamic arrangements as well. Barrier-coverage experiments in this work are modeled to detect Unauthorized Traversal (UT) problem described in [CLO02]. This problem considers a target traversing a sensor field using some path, and the target is detected if it is recognized at some point on its path. This problem assumes an intelligent target that will then plan to find the path with the worst sensor coverage (with lowest probability of detecting a target that traverses

probability of consensus target detection can be written as: , where η is the value fusion threshold. This

$$D_v(u) = \text{prob} \left[\sum_{i=1}^n \frac{K}{\|u - s_i\|^k} + N_i \geq \eta \right] = \text{prob} \left[\sum_{i=1}^n N_i \geq \eta - \sum_{i=1}^n \frac{K}{\|u - s_i\|^k} \right]$$

work assumes noise has a Gaussian distribution with a mean of zero and a standard deviation of one.

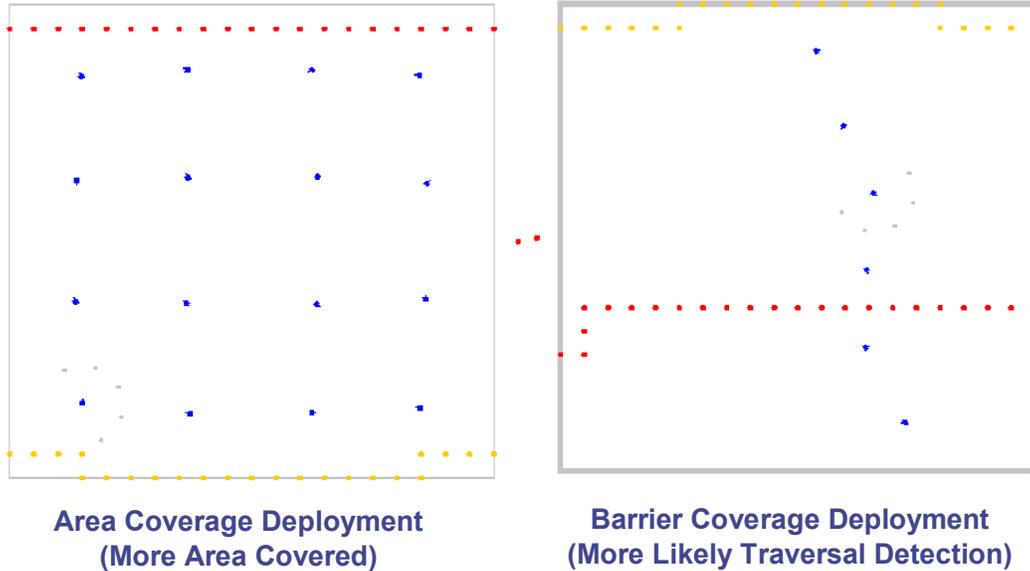


Figure 2. Area vs. Barrier Coverage.

5. Deployment

Sensors may be deployed either manually or autonomously. When deployed manually they either are dispersed to random locations, such as when air-dropped or shot from artillery; or they are placed at specific locations by robots or humans (often in danger). When deployed autonomously, mobile sensors move themselves to sensing locations from an initial arrangement that is easy to realize in a convenient deployment scheme. Possible initial arrangements include (a) from a single source point (e.g., air drop in a canister or off the back of a moving delivery vehicle), (b) in a linear pattern of appropriate density (e.g., sequential deployment from a moving platform), or (c) in a random initial pattern either dense or sparse (e.g., air burst dispersal).

5.1 Random

When deploying sensors randomly, a question is how many sensors should be deployed in an area to ensure a desired detection level is reached. If deploying sensors incrementally, how many should be deployed each step to minimize cost? When the mission is sweep coverage, randomized search techniques have long been recognized as a cheaper than more thorough searches [CHO01].

With the above formulas one can determine the likelihood of detecting a target at any point in the sensor field. In [CLO02], a two-dimensional grid is laid over the sensor field. Value fusion is used to find the detection likelihood at each point on the grid, and this information is used to measure random deployments for barrier coverage. Their measure is termed "exposure" and refers to the detection probability of the path through the grid with the lowest likelihood of detection. The authors develop equations to model overall costs associated with cost for

deployment, the number of sensors for each deployment, the cost for sensors, and the exposure distribution found through simulation. This work uses their exposure measure for barrier coverage deployments.

5.2 Non-Random

Directed sensor deployment algorithms are useful for smaller sensor fields when it is safe for humans to manually place non-mobile sensors or it is possible for mobile sensors to communicate between each other, situate themselves, and maintain internal maps of the area to be covered. A wide range of research has addressed it.

[DHI02] presents simple placement algorithms intended to be fault-tolerant such as putting each subsequent sensor at the least covered point on the grid. Their research shows that directed placement algorithms can significantly reduce the number of sensors needed compared to random placement. Their measure of goodness for deployment of sensors is the minimum detection level of every grid point, so their deployment algorithms address area-coverage missions. They do not address barrier coverage deployments.

[DHI02a] presents a resource-bounded optimization framework for sensor resource management under sufficient coverage of the sensor field. Much of their work is similar to ours from a sensor-modeling standpoint. They represent the sensor field as a grid of points, and model the sensors probabilistically with imprecise detections. They even include obstacles. Their approach is to evaluate the minimum number of sensor nodes required to cover the field. They also experiment with simple ways to transmit or report a minimum amount of sensed data and allow preferential coverage of grid points.

[CHA01] [CHA02] present a coding theory framework for target location in distributed sensor networks. They provide coding-theoretic bounds on the number of sensors needed, and methods for determining their placement in the sensor field using integer linear programming. Although their deployments are strictly for area coverage, they have formulated methods for locating a target based on the location of the sensors that detected the target. This does require a densely packed sensor field, and therefore many more sensors are used than is necessary for detection only. Their approach does not appear scalable due to its computational complexity.

Potential fields have been used to guide robots in a myriad of tasks. [VAU94], for example, uses a potential-field model of flocking behavior to aid the design of flock-control methods. Their work showed a mobile robot that gathers a flock of ducks and maneuvers them to a specified goal. This approach could be used to guide sensors to even-coverage locations.

A problem related to coverage and sensor deployment is that of tracking targets using a network of communicating robots and/or non-mobile sensors. A region based approach to multi-target tracking [JUN02] controls robot deployment at two levels. “A coarse deployment controller distributes robots across regions using a topological map and density estimates, and a target-following controller attempts to maximize the number of tracked targets within a region.” What is similar between this work and ours is a comparison of performance and the degree of occlusion in the environment. Also, they reveal that an optimal ratio of robots to stationary sensors may exist for a given environment with certain occlusion characteristics.

In [GUP03], self-organizing algorithms for area coverage are analyzed for their effect on reduced energy consumption. The authors present techniques to reduce communications which are

different from our quality of coverage metric. Their work deals with selecting sensors that are already placed, whereas our problem deals with the optimal placement of sensors.

The Art Gallery Problem (AGP) studied by numerous researchers considers the placement of cameras in an art gallery such that every key point in the room is visible via at least one camera [BER00] [MAR96]. Some versions find the minimum number of cameras necessary. This problem is similar to complete area coverage with line-of-sight sensors, except that not all sensors are in the visual spectrum. The rooms are usually considered short enough that the cameras' visual quality does not degrade with distance, unlike our probabilistic sensors that have ranges. The art gallery problem has been shown solvable in 2D, but not necessarily in 3D.

Immunology-derived methods for distributed robotics are shown to be a robust control and coordination method in [SIN01]. [MEG00] addresses the problem of finding paths of lowest and highest coverage, which is similar to our metric for barrier coverage deployments. This work uses Voronoi (nearest-point) diagrams to demonstrate a provably optimal polynomial-time algorithm for coverage. Other than the use of Voronoi to break the environment into discrete segments, their approach for measuring coverage is similar to that used for barrier coverage in this work. They only do random deployments, and do not take into account the effects of obstacles, environmental conditions, and noise.

[HSI03] focuses on area coverage. The objective is to fill a region with robots modeled as primitive finite automata, having only local communication and sensors. This work is similar to the other dispersal papers, but reduces jitter and infinite loops. Their robots follow the leader out of a door and fill an area.

The related paper [HES99] proposes a greedy policy to control a swarm of autonomous agents in the pursuit of one or several evaders. Their algorithm is greedy in that the pursuers are directed to locations that maximize the probability of finding an evader at that particular time instant. Under their assumptions, their algorithm guarantees that an evader is found in finite time. The sensor robots combine exploration (or map-learning) and pursuit in a single problem. While this work does improve upon previous research where deterministic pursuit-evasion games on finite graphs have been studied, it is limited to problems where the evader's motion is random and the evader is assumed to exist [HES99]. If an evader enters the sensor field after the pursuit algorithm has begun, it may still evade the pursuers. Continuously running the pursuit algorithm would result in continuous movement of the sensor vehicles which are power-limited. This type of approach is somewhat useful for barrier coverage and sweep coverage, but not very useful for area coverage. There have been similar approaches to this type of pursuit and evasion problem, such as [LAV97] and [PAR98], but none of them are particularly applicable to our work.

5.3 Autonomous

Autonomous deployment algorithms are useful when communications are limited, the number of sensor vehicles is large, and/or the sensor vehicles do not maintain an internal representation of the area to be covered.

[HOW01] [HOW02] [HOW02a] [HOW02b] describe a number of useful deployment algorithms for robot teams and sensor networks. This work is significant because it has been applied to real robots and appears to be the only published work that demonstrates autonomous sensor deployment. Most of their work centers on incrementally deploying mobile robots into unknown environments, with each robot making use of information gathered by the previous robots. These

deployment schemes are global, limited to the area coverage problem, not very scaleable, detection is assumed to be perfect, and the robots are limited to line-of-sight contact with each other. In [HOW02], however, Howard and colleagues presented a distributed and scalable solution to area coverage using potential fields. This approach was the first of its kind and was repeated as part of this work.

A "virtual force" algorithm is proposed for sensor deployment in [ZOU03]. Sensors are deployed randomly, and then virtual forces direct the sensors to more optimum locations. For a given number of sensors, this attempts to maximize area coverage by using repulsive and attractive forces between sensors. The effects of obstacles are not considered, nor are deployment strategies for barrier coverage. A similar local approach to multi-robot coverage is presented in [BAT02], where the assumption is the absence of any *a priori* global information like an internal map or a Global Positioning System (GPS). This work demonstrated local mutually dispersive behaviors between sensor vehicles and resulted in area coverage within 5-7% of the theoretical optimal solution. Obstacles were included in the environment, but barrier coverage was not.

Our application has several modes of deployment. For random deployment, it finds random points away from obstacles and other sensors. For best-first sensor placement, sensors are deployed one at a time; each sensor is put at the point such that the area around it is the least covered compared to all the other points. Best-path is similar to best-first except the only points considered are on the minimum-exposure path from the previous sensor placement.

For greedy sensor placement, each successive sensor is placed at the point evaluated to improve coverage the most at that time. To compute this, for each point on the grid, a temporary sensor is placed there, the coverage for the entire field is measured, and the temporary sensor is removed. The grid point that resulted in the best coverage becomes the next point to place a sensor. Greedy path is similar to greedy except the only points considered are on the minimum exposure path from the previous sensor placement.

For genetic-algorithm sensor placement, a population of random sensor deployments is generated. For example, if 10 sensors are being placed, 100 sets of 10 random deployments are generated. Until sensor-network performance appears to reach an optimum, the following genetic algorithm is used.

```

sort the population by coverage measured
create a new population {
    keep 10 best from the previous generation
    create 90 new sensor deployments as follows {
        choose 2 sensor deployments, favoring high coverage
        create 2 new sensor deployments {
            find a crossover point
            combine 2 chosen by crossing over
            randomly modify sensor positions
            optionally create new random deployment
        } // repeat until 90 new deployments created
    }
    score each deployment in new generation
    sort new population by coverage
} // repeat until near optimal coverage is achieved

```

Simulated annealing simulates the cooling of molten metal until the crystal structure is created [MIC00]. Although simulated annealing cannot always find a globally optimal solution, it has shown success finding solutions to combinatorial-optimization problems with large search spaces [KIN91] since it is capable of escaping locally optimal solutions. When the algorithm starts, and the “temperature” is high, the search is close to random. As the temperature cools, it becomes less random, focusing in on a global optimum. Our implementation of simulated annealing is below.

```

Tmax = 10;
Tmin = .001;
T = Tmax;

currSensors = a random deployment of sensors
bestCoverage = coverage(currSensors)

while( !haltCondition ) {

    currCoverage = coverage(currSensors)

    termCount = 100;
    while( !terminateCondition ) {

        newSensors = a deployment similar to currSensors
        newCoverage = coverage(newSensors)

        if( currCoverage < newCoverage )
            then currSensors = newSensors
        else if random[0,1] < 1/(1+e(newCoverage-currCoverage)/T)
            then currSensors = newSensors

        if( termCount-- <= 0 )
            then terminateCondition = true

        if( bestCoverage < currCoverage )
            then bestCoverage = currCoverage
    }

    T *= 0.93 // other values could be used

    if( (T<=Tmin) or (bestCoverage>=1.0) ) haltCondition = true
}

```

6. Traversal Detection Model for Barrier Coverage

A grid is created to cover the sensor field. A directed graph object called Digrid represents the points and equal length line segments. The grid is an abstraction used to measure barrier coverage. Higher-order graphs could be inserted into the application in the future. Each time the coverage metric is needed, the Digrid computes the probability for detecting an unauthorized traversal through the sensor field according to the following algorithm [CLO02].

```
Generate the grid points and line segments over the sensor field

For each line segment between adjacent grid points {

    Compute value fusion  $D(u)$  for each endpoint,  $u_1$  &  $u_2$ 

    Assign line segment a weight equal to  $\log(1-D(u_1)) + \log(1-D(u_2))$ 
}

Add a link from virtual start to each grid point on the west
Add a link from virtual end to each grid point on the east

Assign a weight of 0 to all the line segments from the start and end points

Compute the least weight path  $P$  from start to end using Dijkstra's algorithm

Let  $w$  equal the total weight of  $P$ 

Return  $P$  as the least exposure path with an exposure equal to  $10^{-w}$ 
```

Notice that in addition to finding the “least covered” path, this also returns the minimum probability of detecting a traversing target. This is the barrier coverage metric in this work.

7. Autonomous Sensor Deployment

Earlier we described some global algorithms for static sensor placement. For robotic vehicles to use these algorithms, they must know the environment and locations of other vehicles. One vehicle determines where the others will go and tells them. This process reoccurs to adjust to changing conditions and sensor failures. But when the number of sensors is large, readjustment becomes time-consuming and communication becomes excessive.

The goal of autonomous sensor deployment is to minimize communication and power and maximize scalability by having the sensor robots' motion be based on local knowledge. The idea is that by allowing each robot to make individual local decisions, good coverage becomes an emergent behavior of the group. This is similar to how ants perform simple individual behaviors that result in emergent complex behaviors for the colony. The difficulty of this approach is finding the right local behaviors that result in the emergent behavior of good coverage for target detection.

The approach in this work uses potentials and vector fields, common tools for robot motion problems. When the potential force method is used, the forces act upon the sensor vehicles to

guide their movement. [ULR00] identified the limitations of using the potential field method and presented the Vector Field approach as an improvement using real robotic vehicles to demonstrate its effectiveness. Vector fields also allow us to reduce the problem to one dimension. We use motion sensors to gather local data from the environment and determine motion from this vector. Consider the sensor vehicle in Figure 13. The circles spaced along the motion sensors represent the sampling points for “attraction” and “repulsion” forces. Along each direction, the sensor vehicle stores the sum of the forces in a vector. This vector is then used to make local decisions for motion.

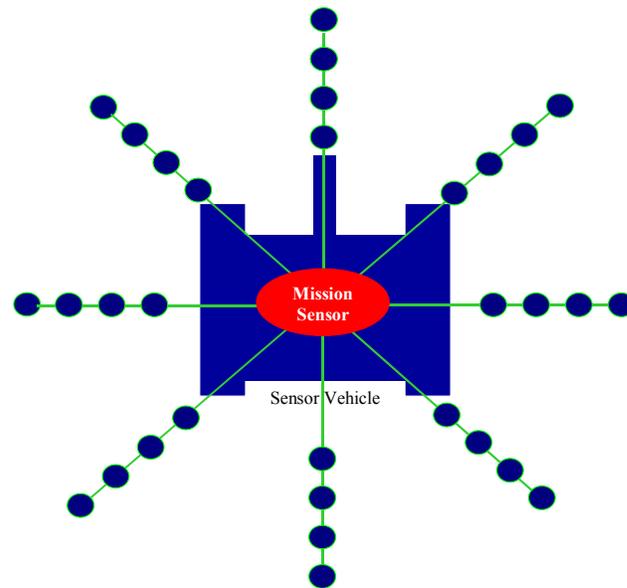


Figure 3. Visualization of sensor vehicle for vector force deployment.

In Figure 3, the sensor vehicle has 8 directions to consider, although we usually use 24. The vehicle agent chooses one of these directions as the desired direction to turn towards. Along each direction, it has evenly spaced sensor points (circles in Figure 13). Each of these sensor points is affected by forces that include repulsive and/or attractive forces from objects, other sensors, local coverage values, sensor field coverage values, etc. A cumulative force vector represents the weighted sum of the forces on each mobile sensor. The forces on sensor points further away from the vehicle are weighted less than those closer to the vehicle.

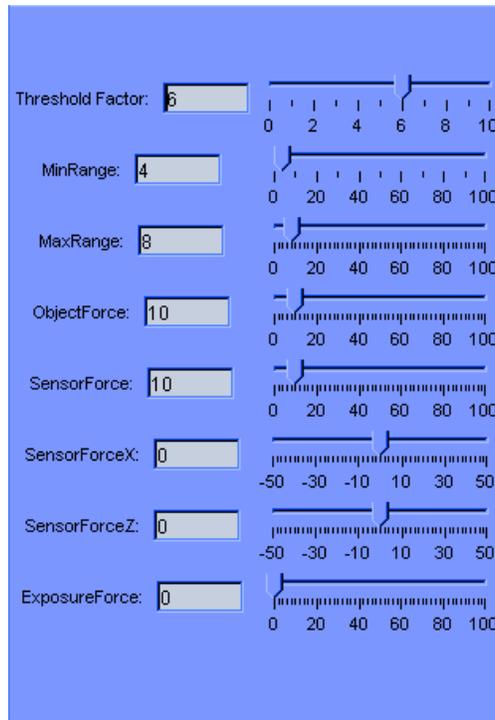


Figure 4. Force panel.

Figure 4 shows the force control panel provided by the application that allows the user to adjust attractive and repulsive force strengths. In practice, these values would be determined in advance but for simulation purposes, it is convenient to be able to modify force parameters to view their effect on the model. In the figure, SensorForce controls the force applied directly from the direction of the entity causing the force. SensorForceX and SensorForceZ allow the user to control the force applied from an entity in the x and z directions, respectively. This is useful because our research shows that if both of these forces are positive, good area coverage results. However, if SensorForceX is negative and SensorForceZ is positive good barrier coverage results. The effective value of each force is added to the direction vector according to this equation:

$$EffectiveForce = \frac{Force}{(Distance)^k}$$

by, *Distance* is the distance from the entity to the sensor point, and *k* is a coefficient that controls how quickly the effect of the force decays with distance. A positive force is repulsive, causing the vehicle to move away.

One way to travel is to select the direction with the least amount of cumulative forces acting upon it and head that way. The vehicle would move this direction to move away from objects and other sensors. The problem with this is that it causes the vehicles to oscillate between directions too much. One improvement is to chose a threshold and only consider directions with force values below the threshold [ULR00]. Instead of picking the lowest valued direction, pick the direction that is closest to the sensor's current direction. Another approach is to add the directions to the left of the sensor's current direction, add the directions to the right of the sensor's current direction, and turn towards the left or the right depending on which sum is lower. There are many ways to combine the various approaches.

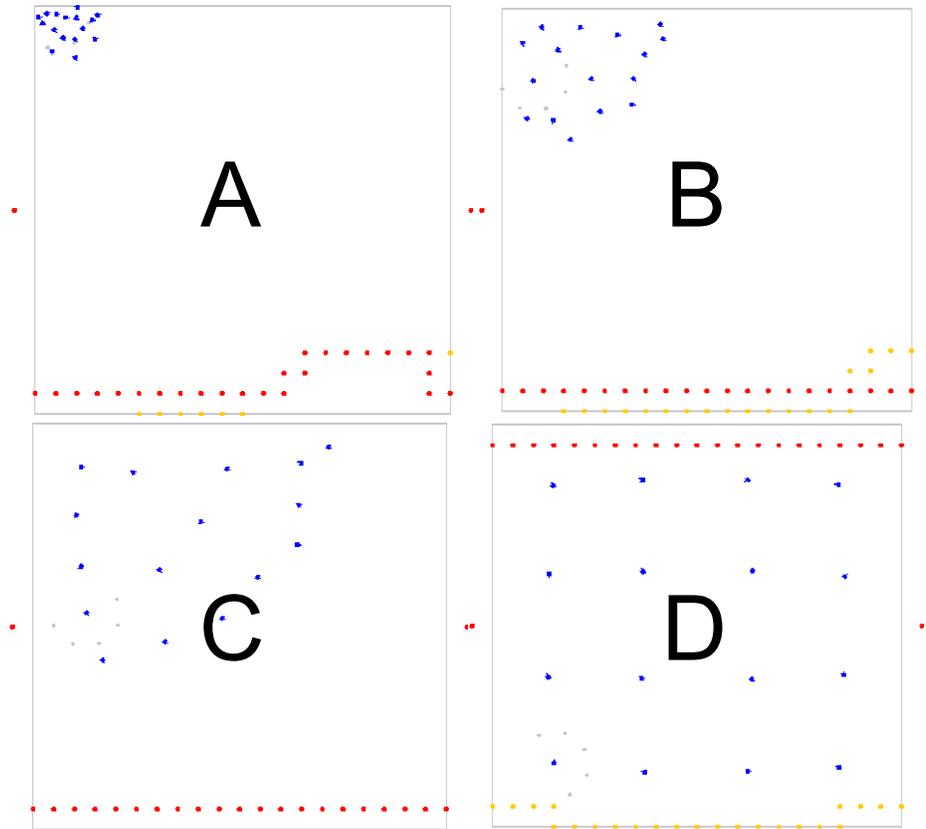


Figure 5. Autonomous deployment for area coverage without obstacles.

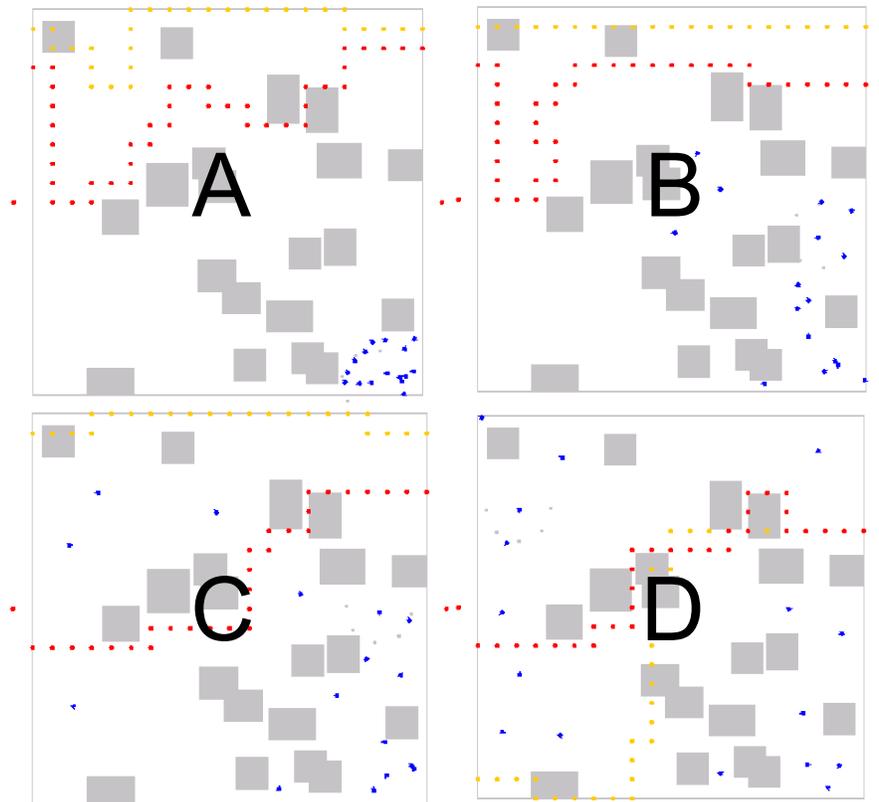


Figure 6. Autonomous deployment for area coverage with obstacles.

Figures 5 and 6 show the progression as sensors deploy from an initial starting position. The first shows a deployment without obstacles present. The second shows deployment with obstacles. These autonomous deployments were conducted with equal forces being applied in each direction, resulting in good area coverage but poor barrier coverage.

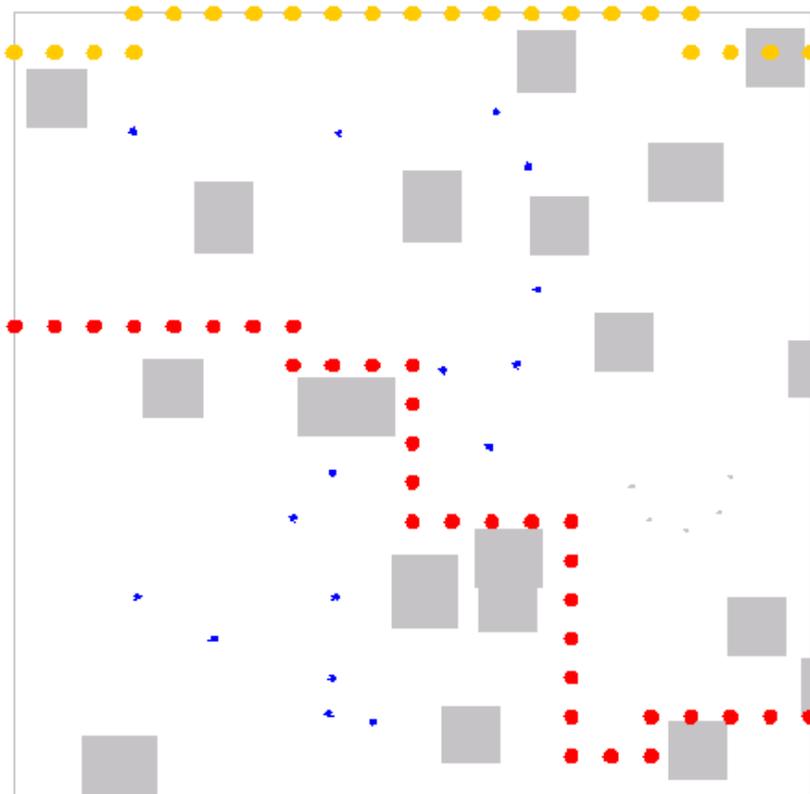


Figure 7. Autonomous deployment for barrier coverage with obstacles.

The deployment sequence in Figure 6 resulted from obstacles and other sensors causing repulsive forces in all directions. When the polarity of the force in the x direction is inverted, the Figure 7 deployment resulted. In this case, sensors repel each other in the y direction, but attract each other in the x direction, resulting in good spacing along a line running north and south. This has the effect of forming a gauntlet that provides good barrier coverage. Again, the red circles show the least covered path along the overlying grid.

8. Results

The simulation can be run as a Java application or a Java applet. Once started, the objects and the internal component structures are created. Based on the parameters specified in an XML initialization file, a default scenario starts. Scenarios run to completion or until the user interrupts with input. Several scenarios can be started by clicking on an appropriate button, and others can be loaded from XML scenario files. In most cases, starting a scenario also causes parameters to be loaded from XML files as well. When a scenario provides active 2D or 3D output, the user has several options for modifying the view, such as zooming in/out and choosing which objects in the environment are visible. The program consists of over 50 Java classes and several XML data files. The source files occupy approximately 700kB and consist of over 30,000 lines of code.

8.1 Experiments

8.1.1 Random Deployment for Traversal Detection

For random deployment, we generate random locations in the terrain and assign the sensors there.

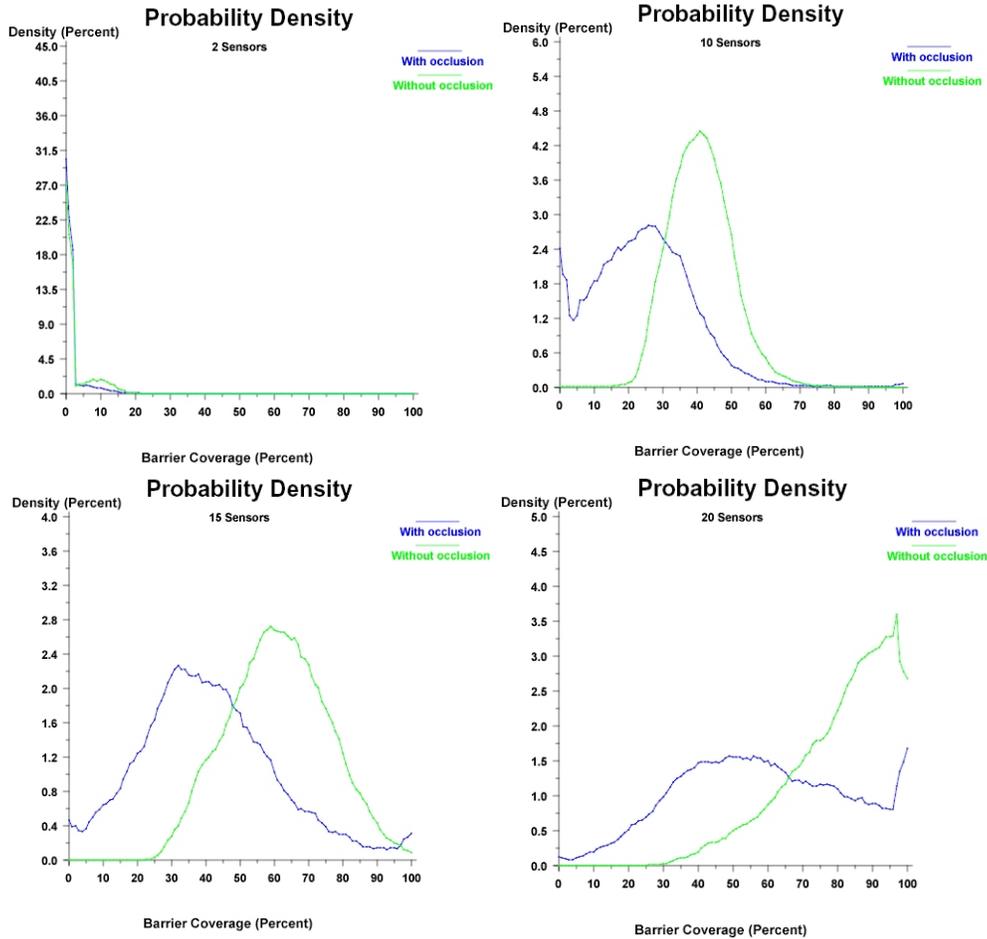


Figure 8. Probability Density Function of Coverage for a Set Number of Sensors

We estimated the probability of traversal detection as a function of the number of sensors by repeated experiments. Each time the barrier coverage probability was calculated and saved. The saved probabilities are used to compute a probability density function. Figure 8 shows some results for 2, 10, 15, and 20 sensors for occluded and unoccluded measurements. Occluded measurements take the effects of obstacles into account, while unoccluded measurements do not.

Probability density function for the distribution of traversal detection coverage measured for deployments of 2, 10, 15, and 20 sensors.

For 2 sensors, the highest density is near zero for unoccluded sensor detections and even lower for occluded. For 10, 15, and 20 sensors, the peak coverages get progressively higher, as expected. This also shows that 20 sensors provide ample coverage for traversal detection.

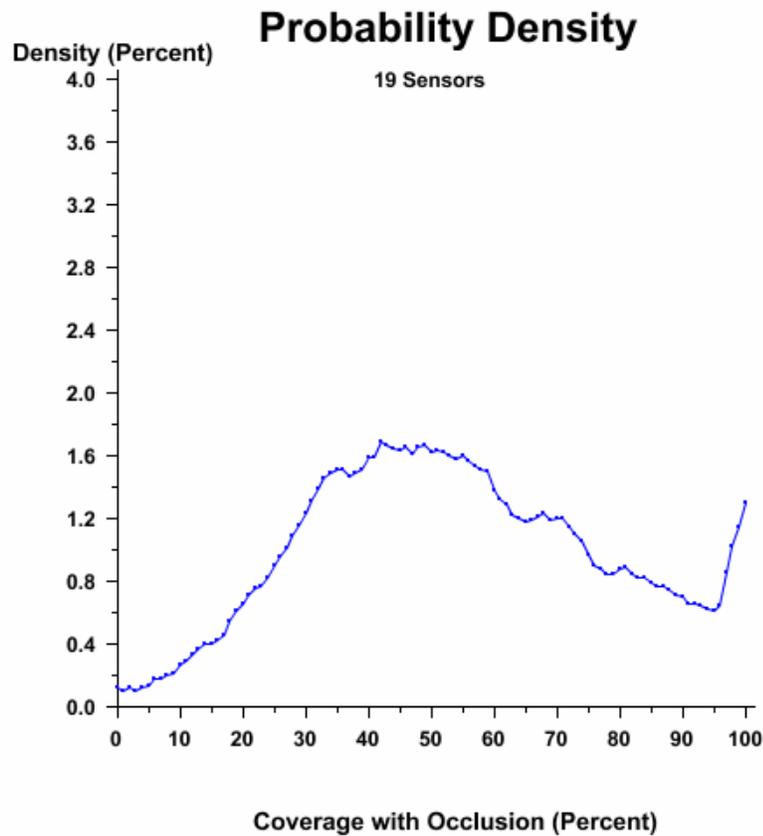


Figure 9. Probability density function for 19 sensors.

Now consider Figure 9, which shows the coverage for deployments of 19 sensors. Notice how the density fluctuates around 1.2 percent density between 30 and 100 percent coverage, varying no more than 1.0 percent. Later we show the number of sensors causing an even spread of density values to be the number of sensors that result in minimum cost when deploying in steps.

8.1.2 Confidence Level for Variable Number of Sensors

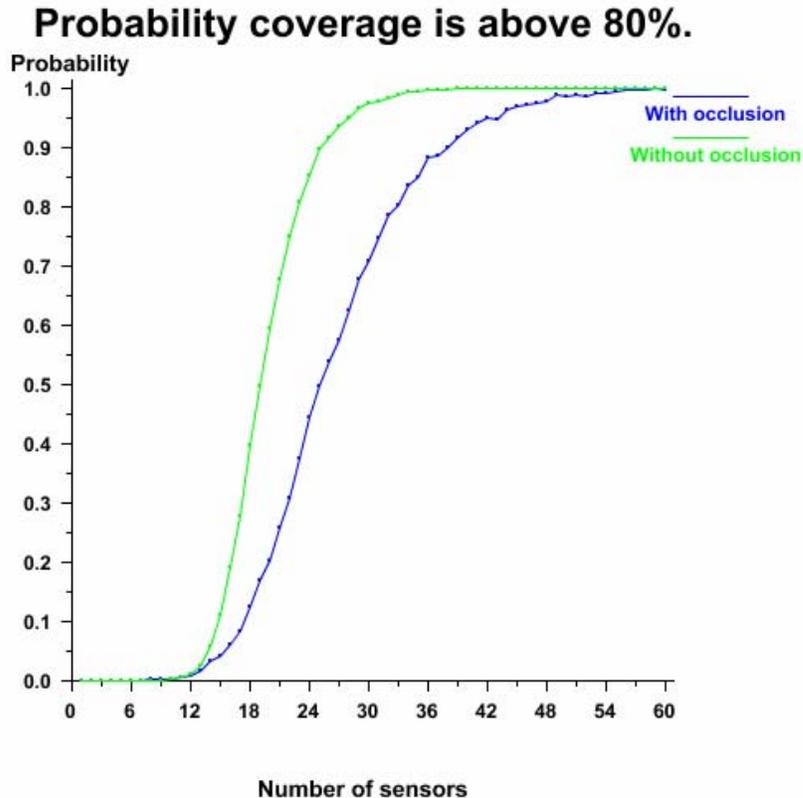


Figure 10. Probability that coverage is above 80% for varying number of sensors.

Obtaining the probability of reaching a specified level of coverage, or the confidence level, is accomplished by randomly deploying sensors repeatedly and computing the percent of time the deployment achieves the specified level. Figure 10 shows a graph of the confidence level reaching 80% coverage for 0 to 60 sensors. The confidence level increases with the number of sensors deployed.

8.1.3 Cost to Achieve Coverage

When deploying sensors randomly, such as when dropping from aircraft or launching from artillery, the question is how many to deploy. A simulation can deploy a set of sensors and measure the coverage. If the desired level of coverage is not reached, deploy another set until the desired level of coverage is reached. Obviously, if there are more sensors available than is necessary to completely cover the sensor field, deploying all sensors at once would be wasteful. If there are 1,000 sensors and only 100 are required, then less than 1,000 should be deployed. We could deploy sensors one at a time, which would minimize the number of sensors deployed at each step; however, if each deployment includes a fixed cost, this would be expensive.

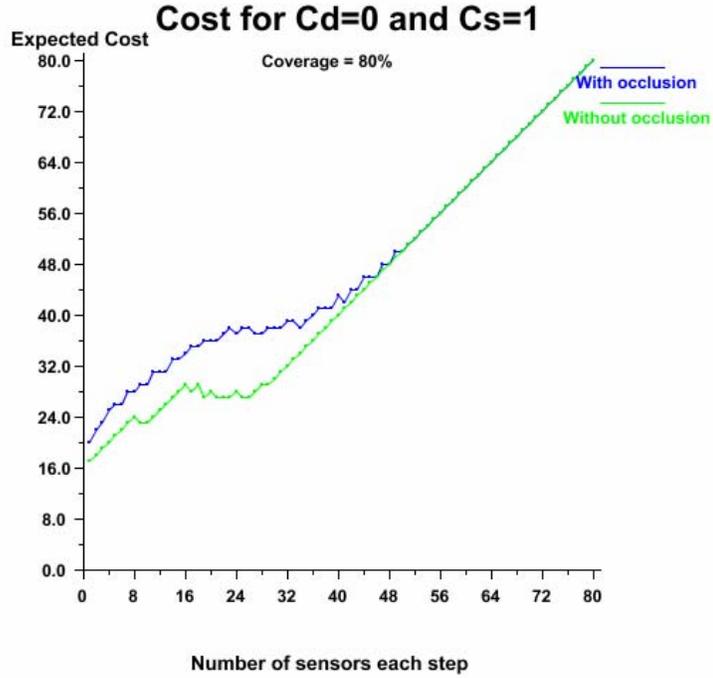


Figure 11. Cost of achieving 80% coverage as a function of the number of sensors with $C_d = 0$ and $C_s = 1$.

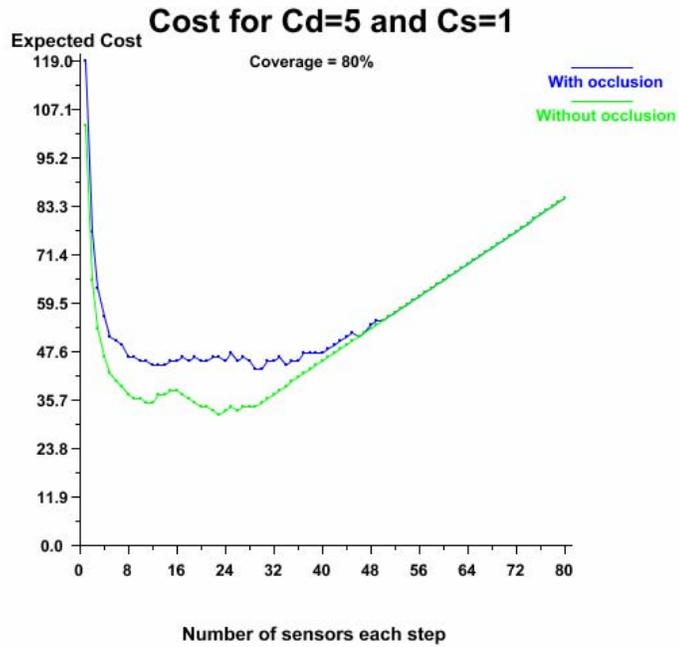


Figure 12. Cost of achieving 80% coverage as a function of the number of sensors with $C_d = 5$ and $C_s = 1$.

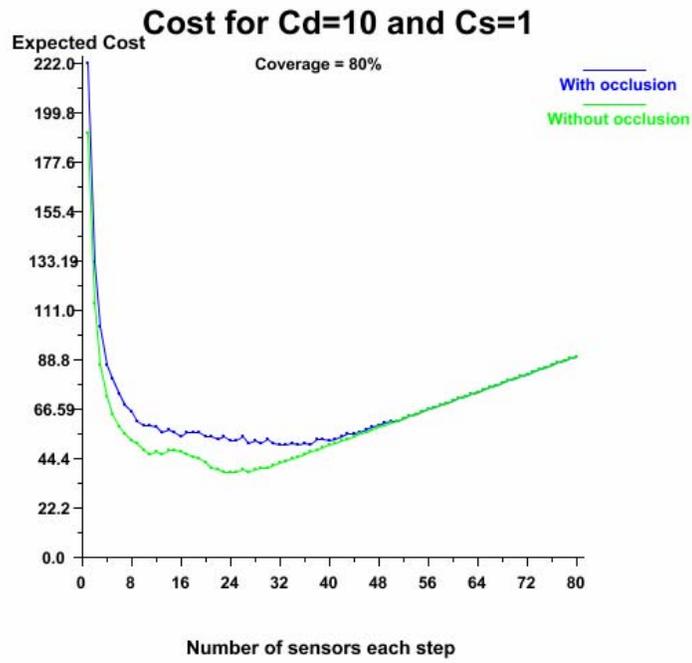


Figure 13. Cost of achieving 80% coverage as a function of the number of sensors with Cd = 10 and Cs = 1.

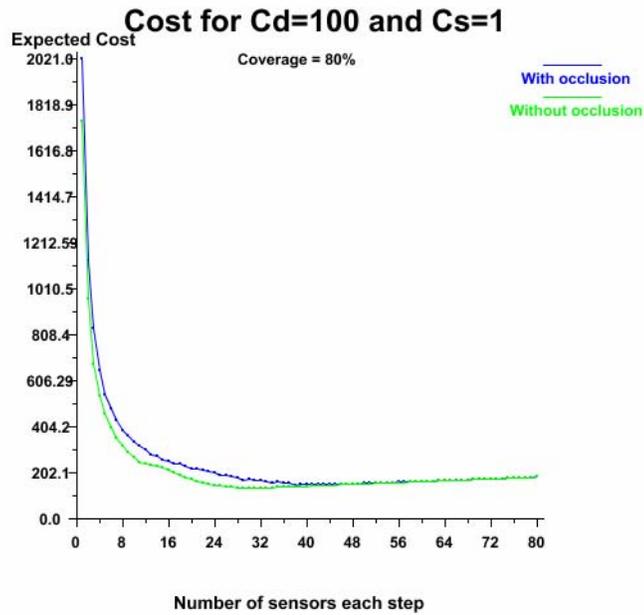


Figure 14. Cost of achieving 80% coverage as a function of the number of sensors with Cd = 100 and Cs = 1.

To find the optimum number of sensors to deploy at a time we first assign a cost to deploy a set of sensors (C_d) and a cost to each sensor (C_s). Figures 11 through 14 are graphs of how much it costs to deploy a set of sensors at a time to achieve the specified coverage.

The first figure shows that if deployment is free then it is most cost-effective to deploy one sensor at a time until achieving the desired coverage. The last shows that if deploying costs 100 times as much as each sensor then it is better to deploy many sensors each step. The other figures show that deploying approximately 23 sensors each step minimizes cost. Cost is achieved at nearly the same number of sensors for occluded and unoccluded sensor measurements so finding the minimum for one should be sufficient in practice.

Note that the local minimums of 23 for cost correspond to Figure 9, where coverage fluctuates between 30 and 100 percent with 19 sensors. [CLO02] developed an analytical explanation for the correspondence between the local minimum and the number of sensors with a flat probability density. They used the following equation to describe the expected cost

$$E\{C\} = (C_d + n \cdot C_s) \left[\sum_{i=1}^S i \cdot \left(\prod_{j=1}^i F_{j,n}(v) \right) (1 - F_{i,n}(v)) \right]$$

where n is the number of sensors deployed each step for a total of S steps, and $1 - F_n(v)$ is the confidence of obtaining coverage of amount v with n sensors. As described in [CLO02], when the probability density varies widely the weight associated with the first term of the sum increases rapidly while the weights associated with the higher number of sensors decrease. The cost increases again after this local minimum as the increase in n compensates for the decrease in weights.

8.1.4 Deployment Algorithm Performance Comparison

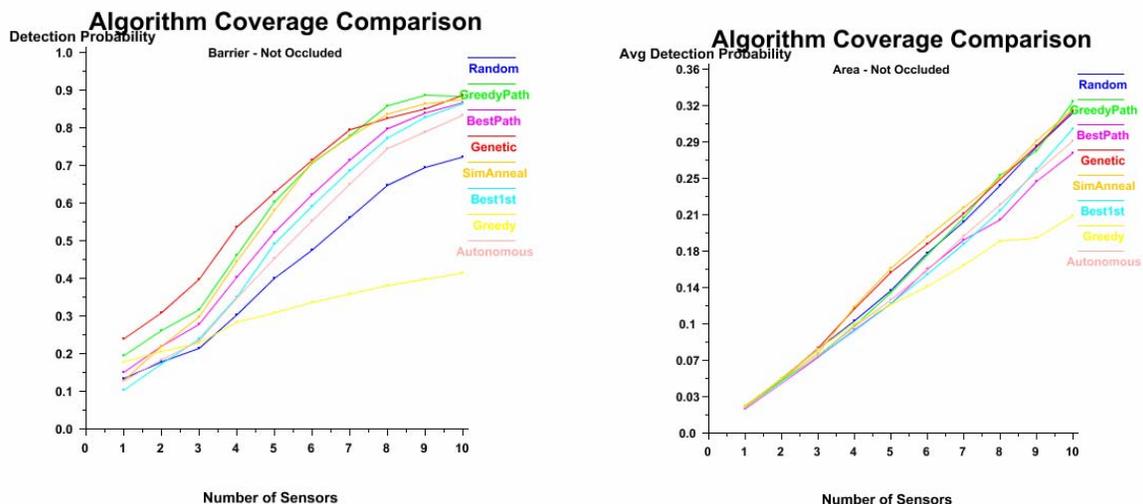


Figure 15. Deployment algorithm comparison, not occluded.

Figure 15 compares barrier coverage for each of the placement algorithms in a 400m² sensor field. Occluded means that objects in the environment block the energy from target to sensor. Measurements were obtained for area coverage and barrier coverage, occluded and not occluded, for each of the deployment algorithms described in the previous chapter.

The graphs in Figure 15 show the coverage differences between algorithms when obstacles present barriers to traversal but not to sensing. Our analysis shows that the Greedy algorithm is the worst coverage algorithm; however, our GreedyPath algorithm performed well. Note Random placement performed fairly well, which is significant because it will usually be cheap to deploy randomly.

Autonomous deployment performed well for unoccluded (Figure 15) and occluded (Figure 16) measurements. The sensor vehicles were able to deploy using only local information and no central planning. For scenarios where many sensors have no internal map of the environment, limited communication capabilities, and limited power, this will be essential.

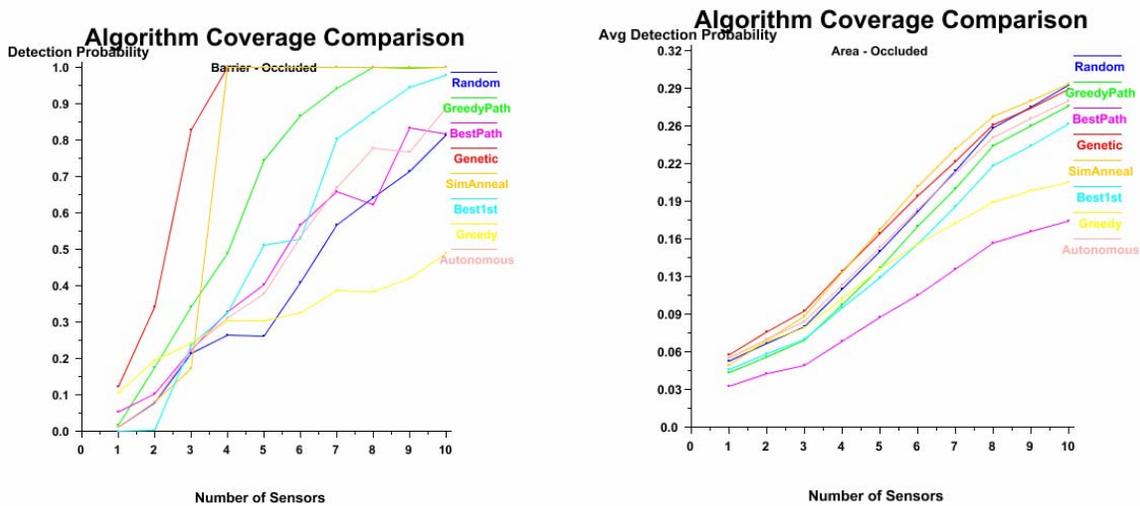


Figure 16. Deployment algorithm comparison, occluded.

When the effect of obstacles is taken into account, the graphs in Figure 16 show the choice of algorithm can have a significant effect on coverage. Greedy again performs the worst. The Genetic and Simulated Annealing algorithms perform the best by far, achieving 100% barrier coverage with far fewer sensors than the other algorithms. Figure 17 gives a time comparison of the algorithms. GreedyPath performed quite well considering it is also one of the fastest of the algorithms. Random did not perform as well, but did perform adequately considering it is the fastest. Best1st and Greedy are the slowest. Genetic and Simulated Annealing are the slowest with few sensors, but with more sensors find optimal solutions quickly.

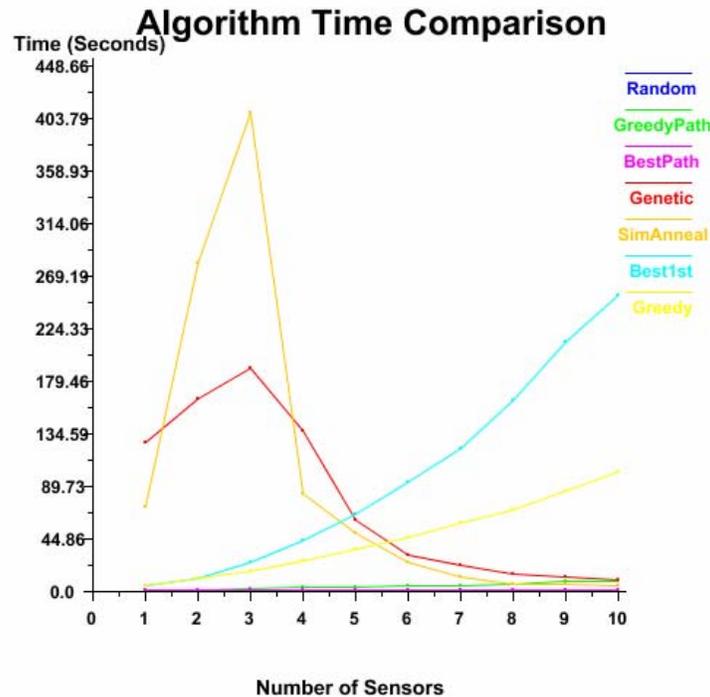


Figure 17. Deployment algorithm time comparison.

8.2 Future Work

This work developed a simulation environment to enable varied testing of coverage and deployment issues in a wide variety of sensor networks. Many types of sensors exist and with different characteristics (mobility, power, sensing ability, etc.). The energy measured by the sensors varies. Sensors exist that detect sound, light, movement, vibration, radiation, chemical substances, and so on. The environment may have varying occlusion characteristics as well as areas where the desire to detect a target is higher. Different coverage measurement schemes, different data fusion schemes, and different deployment algorithms will need to be tested. The number of possible combinations that are testable with the program implemented for this work is large.

As for future work, the directed grid used to measure barrier coverage can have different levels of granularity. With higher granularity, there are more grid squares and the coverage metric should be more accurate. Future work should examine the effect of granularity on accuracy and time. [MEG00] describes the use of Voronoi diagrams for computing coverage similar to our use of directed grids. Future work could evaluate the effective of granularity and compare and contrast Voronoi diagrams with directed grids.

In [PRO00] the authors present a vigilance model which allows birds to control the area a flock occupies as well as their vigilance rate. An optimal strategy is found for the birds under a variety of conditions. Performing the two contradictory tasks presented, feeding and avoiding becoming food, is analogous to two sensor tasks, sensing phenomena and conserving energy. Future work

could examine the models presented in [PRO00] and extend the theories to optimize the tradeoff between frequent sensing and power conservation, for example.

Future enhancements can be made to the size and scope of the environment. For simplicity, this work focused on a limited number of homogenous sensors in a fixed area. The application supports testing for heterogeneous sensors, as well as different sensing models. Also, the area to be covered does not need to remain static. Sensors covering the flank side of a mobile force would need to move with the force and continuously update their positions.

The size of the environment and the number of entities can be increased in a few ways. One approach would be to have a larger sensor field aggregate the information of several smaller sensor fields for coverage and deployment. However, a limiting factor is that all the sensors run on one machine unlike in the independence of the real world. Future work could simulate separate computational entities and network communication, which should enable the size to increase greatly.

9. Conclusions

This work has explored coverage and deployment issues for mobile and non-mobile sensors. A simulation of a multi-agent expeditionary sensor network was created to formulate and test search, coverage, and deployment algorithms. In the course of this research we evaluated cost and performance of deploying multiple homogenous sensors with a wide variety of constraints. The comparison of several deployment algorithms showed significant insights, and novel autonomous deployment schemes were presented. The application and algorithms developed enable future modeling and simulation efforts.

While the components implemented in this work do not target any existing framework, the concepts are applicable to a wide range of sensor coverage and deployment problems. Many possible scenarios can be built, tested, and compared. This allows for extensive simulation of varied platforms to guide future acquisition and minimize cost. The design of the application allows for future extensibility with minimal modification to existing code. The algorithms described can be implemented in any programming language or sensor network.

10. References

- [AKY02] Akyildiz, I.F., Su, W., Sankarasubramaniam, Y., Cayirci, E., “Wireless sensor networks: a survey,” *Computer Networks*, 38, 393-422, 2002.
- [AXE84] Axelrod, R., *The Evolution of Cooperation*, Basic Books, 1984.
- [AXE97] Axelrod, R., *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*, Princeton University Press, 1997.
- [BAT02] Batalin, M.A., “Spreading Out: A Local Approach to Multi-Robot Coverage,” Proceedings of the 6th International Symposium on Distributed Autonomous Robotics Systems, Fukuoka, JA, 2002.
- [BAT03] Batalin, M.A., “Efficient Exploration Without Localization,” Computer Science Dept., USC, 2003.

- [BER00] Berg, M., Kreveld, M., Overmans, M., and Schwarzkopf, O. *Computational Geometry: Algorithms and Applications*, Springer-Verlag, 2000.
- [BIG98] Bigus, J.P., and Bigus, J., *Constructing Intelligent Agents with Java™*, John Wiley & Sons, Inc, 1998.
- [BOR90] Borenstein, J., Koren, Y., “Real-Time Obstacle Avoidance for Fast Mobile Robots in Cluttered Environments,” Proceedings of the 1990 IEEE International Conference on Robotics and Automation, Cincinnati, OH, 572-577, 1990.
- [BOR96] Borenstein, J., Everett, H.R., Feng, L., Wehe, D., “Mobile Robot Positioning – Sensors and Techniques,” *Journal of Robotic Systems, Special Issue on Mobile Robots*, Vol. 14, 231-249, 1996.
- [BUC02] Buckland, Mat, *AI Techniques for Game Programming*, Premier Press, 2002.
- [BUG00] Bugajska, M.D., Schultz, A.C., “Co-Evolution of Form and Function in the Design of Autonomous Agents: Micro Air Vehicle Project,” Navy Center for Applied Research in Artificial Intelligence, 2000.
- [BUL01] Bulusu, N., Heidemann, J., Estrin, D., “Adaptive Beacon Placement,” UCLA, Los Angeles, CA, 2001.
- [CHA01] Chakrabarty, K., Iyengar, S.S., Qi, H., Cho, E., “Coding Theory Framework for Target Location in Distributed Sensor Networks,” Dept. of Computer and Electrical Engineering, Duke University, 2000.
- [CHA02] Chakrabarty, K., Iyengar, S.S., Qi, H., Cho, E., “Grid coverage for surveillance and target location in distributed sensor networks,” *IEEE Transactions on Computers*, vol. 51, 2002.
- [CHO01] Choset, H., “Coverage for robotics – A survey of recent results,” *Annals of Mathematics and Artificial Intelligence*, 31, 113-126, 2001.
- [CLI96] Cliff, D., Miller, G.F., “Co-evolution of pursuit and evasion II: Simulation Methods and Results,” *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior (SAB96)*, MIT Press Bradford Books, 1996.
- [CLO01] Clouqueur, T., Ramanathan, P., Saluja, K.K., Wang, K., “Value-Fusion Versus Decision-Fusion for Fault-Tolerance in Collaborative Target Detection in Sensor Networks,” Fusion 2001 Conference, 2001.
- [CLO02] Clouqueur, T., Phipatanasuphorn, V., Ramanathan, P., Saluja, K.K., “Sensor Deployment Strategy for Target Detection,” WSNA’02, Atlanta, GA, 2002.
- [DHI02] Dhillon, S.S., Chakrabarty, K., “A Fault-Tolerant Approach to Sensor Deployment in Distributed Sensor Networks,” Dept. of EE & CE, Duke University, Durham, NC, 2002.

- [DHI02a] Dhillon, S.S., Chakarabarty, K., "Sensor Placement for Grid Coverage under Imprecise Detections," Dept. of EE & CE, Duke University, Durham, NC, 2002.
- [DIC02] Dickie, Alistair, *Modeling Robot Swarms using Agent-Based Simulation*, Master's Work, Naval Postgraduate School, Monterey, CA, June 2002.
- [DMS03] Defense Modeling and Simulation Office (DMSO). "Smart Sensor Web: Experiment Demos Network Centric Warfare," https://www.dmsomil/public/pao/stories/y_2002/m_04/7-1-8, May, 2003.
- [EST99] Estrin, D., Govindan, R., Heidemann, J., Kumar, S., "Next century challenges: scalable coordination in sensor networks," ACM MobiCom'99, 263-270, 1999.
- [EST02] Estrin, D., Culler, D., Pister, K., Sukhatme, G., "Connecting the Physical World with Pervasive Networks," Pervasive Computing, IEEE, 59-69, 2002.
- [GAG92] Gage, D.W. "Command Control for Many-Robot Systems," Proceedings of AUVS-92, Huntsville, AL, 22-24 June 1992.
- [GAG92a] Gage, D.W. "Sensor Abstractions to Support Many-Robot Systems," RDT&E Division, Naval Command Control and Ocean Surveillance Center, 1992.
- [GAG93] Gage, D.W. "Randomized search strategies with imperfect sensors," *Proceedings of SPIE Mobile Robots VIII*, Boston, 9-10 September 1993, pp 270-279.
- [GAG98] Gage, D.W. "Randomized Search Strategies," <http://www.spawar.navy.mil/robots/research/manyrobo/randomize.html>, June, 2003.
- [GOL01] Goldberg, D., Mataric, M.J., "Robust Behavior-Based Control for Distributed Multi-Robot Collection Tasks," USC, Los Angeles, CA, 2001.
- [GOR00] Gordon, D.F., Spears, W.M., Sokolsky, O., Lee, I., "Distributed Spatial Control, Global Monitoring and Steering of Mobile Agents," AI Center, NRL, Wash., DC, 2000.
- [GUP03] Gupta, H., Das, W.R., Gu, Q., "Connected Sensor Cover: Self-Organization of Sensor Networks for Efficient Query Execution," Mobihoc'03, Annapolis, MA, 2003.
- [HES99] Hespanha, J.P., Kim, H.J., Sastry, S., "Multiple-Agent Probabilistic Pursuit-Evasion Games," Proceedings of the 38th Conference on Decision and Control, 1999.
- [HOW01] Howard, A., Mataric, M.J., "Cover Me! A Self-Deployment Algorithm for Mobile Sensor Networks," Robotics Research Lab, USC, 2001.
- [HOW02] Howard, A., Mataric, M.J., Sukhatme, G.S., "Mobile Sensor Network Deployment using Potential Fields: A Distributed, Scalable Solution to the Area Coverage Problem," Proceedings of the 6th International Symposium on Distributed Autonomous Robotics Systems (DARS02), Fukuoka, JA, 2002.

- [HOW02a] Howard, A., Mataric, M.J., Sukhatme, G.S., "An Incremental Deployment Algorithm for Mobile Robot Teams," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2002), EPFL, SW, 2002.*
- [HOW02b] Howard, A., Mataric, M.J., Sukhatme, G.S., "An Incremental Self-Deployment Algorithm for Mobile Sensor Networks," *Autonomous Robots, Special Issue on Intelligent Systems, 2002.*
- [HSI03] Hsiang, T., Arkin, E.M., Bender, M.A., Fekete, S.P., Mitchell, J.S.B., "Algorithms for Rapidly Dispersing Robot Swarms in Unknown Environments," Stony Brook University, NY, 2003.
- [JUN02] Jung, B., Sukhatme, G.S., "A Region-Based Approach for Cooperative Multi-Target Tracking in a Structured Environment," 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2002.
- [KOH00] Kohout, B. "Challenges in Real-Time Obstacle Avoidance," Veridian/Pacific-Sierra Research, Rosslyn, VA, 2000
- [KOR91] Koren, Y., Borenstein, J., "Potential Field Methods and Their Inherent Limitations for Mobile Robot Navigation," Proceedings of the IEEE Conference on Robotics and Automation, 1398-1404, 1991.
- [KIN91] Kindl, M.R., Shing, M., Rowe, N.C., "A Stochastic Approach to the Weighted-Region Problem: Design and Testing of a Path Annealing Algorithm," *Naval Center for Applied Research in Artificial Intelligence, 1991.*
- [LAV97] LaValle, S.M., Lin, D., Guibas, L.J., Latombe, J., Motwani, R., "Finding an Unpredictable Target in a Workspace with Obstacles," Proceedings of the 1997 IEEE International Conference on Robotics and Automation, 1997.
- [LUN01] Lund, H.H, Pagliarini, L., "Edutainment Robotics: Applying Modern AI Techniques," Proceedings of International Conference on Autonomous Minirobots for Research and Edutainment (AMIRE-2001), 2001.
- [MAR96] Marengoni, M., Draper, B., Hanson, A., and Sitaraman, R. "System to place observers on a polyhedral terrain in polynomial time." *Image and Visual Computing, 1996.*
- [MEG00] Meguerdichian, S., Koushanfar, F., Potkonjak, M., Srivastava, M.B., "Coverage Problems in Wireless Ad-hoc Sensor Networks," CS Dept and EE Dept, UCLA, 2000.
- [MEG01] Meguerdichian, S., Koushanfar, F., Qu, G., Potkonjak, M., "Exposure in Wireless Ad-Hoc Sensor Networks," ACM SIGMOBILE, Rome, Italy, 2001.
- [MIC00] Michalewicz, Z., Fogel, D.B., *How to Solve It: Modern Heuristics*, Springer-Verlag Berlin Heidelberg, NY, 2000.

- [PAR00] Park, S., Savvides, A., Srivastava, M.B., "SensorSim: A Simulation Framework for Sensor Networks," EE Dept., UCLA, 2000.
- [PAR98] Parker, L.E., "Cooperative Robotics for Multi-Target Observation," Center for Engineering Systems Advanced Research (CESAR), Oak Ridge, TN, 1998.
- [PER00] Perkins, C. *Ad Hoc Networks*, Addison-Wesley, 2000.
- [PIS03] Pister, Kris. "Smart Dust. Autonomous Sensing and Communicating in a Cubic Millimeter," <http://robotics.eecs.berkeley.edu/~pister//SmartDust/>, August, 2003.
- [PRO00] Proctor, C.J., Broom, M., "A Spatial Model of Anti-Predator Vigilance," Centre for Statistics and Stochastic Modeling, University of Sussex, 2000.
- [REY99] Reynolds, C.W., "Steering Behaviors for Autonomous Characters," Sony Computer Entertainment America, Foster City, CA, 1999.
- [ROM02] Romer, K., Kasten, O., Mattern, F., "Middleware Challenges for Wireless Sensor Networks," *Mobile Computing and Communications Review*, Vol. 16, 2, 2002.
- [ROU98] Roumeliotis, S.I., Pirjanian, P., Mataric, M.J., "Ant-Inspired Navigation in Unknown Environments," Computer Science Dept, USC, 2002.
- [RUS95] Russell, S. and Norvig, P. *Artificial Intelligence: A Modern Approach*. Prentice-Hall, Inc, 1999.
- [SIN96] Singh, S., Bertsekas, D. "Reinforcement Learning for Dynamic Channel Allocation in Cellular Telephone Systems," *NIPS96, Section: Applications*, 1996.
- [SIN01] Singh, S.N.P., Thayer, S.M., "Immunology Directed Methods for Distributed Robotics: A Novel, Immunity-Based Architecture for Robust Control & Coordination," Robotics Institute, CMU, 2001.
- [SPE99] Spears, W.M., Gordon, D.F., "Using Artificial Physics to Control Agents," AI Center, Naval Research Laboratory, Washington, DC, 1999.
- [STE95] Stentz, A., "The Focussed D* Algorithm for Real-Time Replanning," Robotics Institute, CMU, 1995.
- [STO96] Stork, K.A., *Sensors in Object Oriented Discrete Event Simulation*, Master's Work, Naval Postgraduate School, Monterey, CA, September 1996.
- [TIL02] Tilak, S., Abu-Ghazaleh, N.B., Heinzelman, W., "A Taxonomy of Wireless Micro-Sensor Network Models," *Mobile Computing and Communications Review*, Vol. 1, Num. 2, 2002.
- [ULR00] Ulrich, I., Borenstein, J., "VFH*: Local Obstacle Avoidance with Look-Ahead Verification," 2000 IEEE International Conference on Robotics and Automation, San Francisco, CA, 2505-2511, 2000.

- [VAU94] Vaughan, R., Sumpter, N., Henderson, J., Frost, A., Cameron, S., “Robot Control of Animal Flocks,” BioEngineering Division, Silsoe Research Institute, Silsoe, Bedford, UK, 1995.
- [WAT01] Watt, A., Policarpo, F. *3D Games: Real-time Rendering and Software Technology*, Addison-Wesley, 2001.
- [WEI99] Weiss, Gerhard. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. The MIT Press, 1999.
- [ZOU03] Zou, Y., Chakrabarty, K., “Sensor Deployment and Target Localization Based on Virtual Forces,” Dept. of EE & CE, Duke University, Durham, NC, 2003.