Directed Stigmergy-Based Control for Multi-Robot Systems

Fitzgerald Steele, Jr.
The University of Iowa
4225 Seamans Center
Iowa City, IA 52242
+1 319.384.0526
fitzgerald-steele@uiowa.edu

Geb Thomas
The University of Iowa
2404 Seamans Center
Iowa City, IA 52242
+1 319.335.5936
geb-thomas@uiowa.edu

ABSTRACT

Multi-robot systems are particularly useful in tasks that require searching large areas such as planetary science exploration, urban search and rescue, or landmine remediation. In order to overcome the inherent complexity of controlling multiple robots, the user must be able to give high-level, goal driven direction to the robot team. Since human robot interaction is a relatively new discipline, it is helpful to look to existing systems for concepts, analogies, or metaphors that might be utilized in building useful systems. Inspiration from natural decentralized systems guides the development of a computer simulation for stigmergy-based control of multi-robot system, and the interface with which an operator can interact and control mobile robots. In-depth description of the design process includes a description of a basic stigmergy-based control system and an innovative Directed Stigmergy control system that facilitates operator control of the robot team in an interesting and surprisingly effective way.

Categories and Subject Descriptors
I.2.9 [**Artificial Intelligence**]: Robotics – Operator Interfaces, Autonomous Vehicles

General Terms
Algorithms, Documentation, Design, Human Factors.

Keywords
Stigmergy, Robotics, Multi-Robot, Supervisory Control, Swarm, Human Robot Interaction, User Interface

1. INTRODUCTION

The past decade has brought tremendous advances in the use of robotics in search and exploration tasks. The 1996 Mars Pathfinder Mission showcased Sojourner, a shoebox sized mobile robot. Sojourner had few science instruments, and limited range of motion – it could not travel more than a few meters from the Sagan Memorial Lander [12]. The 2004 Mars Exploration Rovers (MER), Spirit and Opportunity, are golf-cart sized mobile robots with expanded science data collection capabilities. The MER robots are not limited to the area near the landing site – they have driven over six kilometers and are still going strong [17].

There have also been advances in the means mission control operators direct the action of the remote robots. Sojourner required a rigid, time stamped sequence of commands. If the robot was not able to meet the schedule for any reason, it retracted itself into a safe mode and awaited new instructions [25]. The Contingent Rover Language, combined with the MarsMap Virtual Reality interface allows mission planners to graphically define a set of possible courses of action from which the rover can pick the most appropriate or opportunistic execution plan [6, 27]. The abstraction toward goal-directed commands continues with the advanced autonomous navigation [16] and science opportunity recognition capabilities [19, see also 8] onboard the MER robots. Mission planners give each robot an end destination and the robot defines its own path to the target, recognizing interesting features and avoid obstacles along the way.

There has also long been a demand for using robots for Urban Search and Rescue, the search for victims in urban settings that have been damaged from earthquakes, terrorist attacks, or other disasters [1]. Mobile robots have shown the ability to traverse spaces that are too small or too dangerous for humans or rescue dogs [4]. There are advantages to using multiple robots in urban search and rescue, including more quickly searching the damaged area, using robots with different shapes in different environments, and coordinating the use of expensive thermal sensors [18]. However, it has been found that even a single robot can be difficult for rescue workers to control in emergency situations [7].

The evolution in robot search missions highlights two key elements for future robotics research. First, increasing the number of robots available extends the area available for exploration and increases the odds of making interesting discoveries. Second, there is a trend toward higher levels of supervisory control over the robots. In order to improve the scientific return of future robotic exploration missions, we must enable teams of robots to cooperatively explore the remote environment, and enable mission controllers to effectively direct and control the multi-robot team through high level, goal driven interaction.

Increasing the number of robots in the system, however, increases the complexity of the system, making it more difficult for the human operator to maintain effective control over the robot team. An operator must switch focus between robots, and quickly review the system to gain situation awareness. Supervisory control of the system would be difficult because the process of monitoring the automation would not scale to large robot teams [23]. Another challenge to managing many robots is that a collection of robots may take on emergent behaviors. Emergent behavior in multi-robot systems occurs when the individual actions combine to form a larger, emergent pattern [9] [15]. Emergent behaviors can be quite useful in that the system automatically adapts to changes in the environment, but emergent behaviors are not easily predicted or understood at the level of the

HRI'07, March 8-11, 2007, Arlington, Virginia, USA.

individual. Human robot interaction techniques are required to control the group behavior of the robot team, while maintaining the advantages of decentralized robot architecture.

Can we devise a way to retain the advantages of decentralized systems while supporting human control? This research question is first addressed by describing a control system for autonomous robots based on *stigmergy*, the manipulation of the environment to form emergent behaviors. We then examine Directed Stigmergy, an augmented control system and user interface for human control of the autonomous multi-robot system. Directed Stigmergy is inspired by examples of emergent behavior in self-organizing, decentralized, systems found in nature: foraging behavior of ant colonies and molecule bonding. Experimental simulations indicate that operators are successful in using Directed Stigmergy to exert influence over the swarm behavior of a group of autonomous mobile robots without directing individual robot behavior. While Directed Stigmergy can be generalized to a range of robot search and exploration tasks such as Urban Search and Rescue or landmine remediation, this paper will use planetary exploration on Mars as an application domain.

2. Stigmergy-based control of multi-robot swarms

2.1 Background

A decentralized system is a system in which the individuals in the population act independently, without receiving instructions from a superior, leader, or central controller. This is different from a *centralized system*, in which a single operator or agent plans, coordinates, and commands the other members in the system. Individuals within decentralized systems interact with each other and their environment to produce complex, coordinated behavior.

Decentralized systems offer some inherent advantages over centralized systems. First, decentralized systems do not have a single point of failure. The loss of a single agent will not cripple the system, as can be the case in single-agent or centrally controlled systems. Second, decentralized systems can achieve complex results with relatively simple system design. The designer need only create simple, low level behaviors, instead of a single, computationally intense control system to govern all possible situations. Finally, decentralized systems are inherently parallel, which allows for extremely scalable systems and faster task completion. These advantages in robustness, ease of design, speed, and scalability make decentralized multi-robot systems an attractive method for searching large areas as required in interplanetary science missions, as well as other domains such as landmine remediation, urban search and rescue, and unmanned aerial vehicles. However, as the multi-robot automation adapts to dynamic changes in the world, emergent group behavior may form that can be difficult for system designers and operators to accurately predict and control. Therefore, it is important to investigate new methods of exerting control over these systems [15].

The food foraging behavior of ant colonies is a classic example of coordinated, emergent behavior in decentralized systems [9]. This behavior demonstrates *stigmergy*, the manipulation of the environment in order to affect the emergent behavior of the group [5, 13, 14]. Ants leave the hive, and randomly wander about the surrounding area in search of food. When an individual ant finds food, it gathers a piece and returns to the hive. The ant leaves a chemical pheromone trail that can be detected by other ants. Other ants will find the trail, follow it to the food source, and return home, strengthening the pheromone trail with its own scent. This positive feedback continues as more and more individuals join the trail, find the food, and bring it home. Finally, the food source is exhausted, and the ants continue to explore the area until a new food source is located and the chain restarts. Thus while each ant independently reacts to its local environment, the emergent behavior of the colony demonstrates rapid search of a wide area, memory of the location of desired targets, efficient allocation of group resources, and dynamic adaptation to changes.

We would like multi-robot systems to exhibit these emergent behaviors. Robots foraging for interesting targets should emulate the ‘random’ wandering behavior in order to explore a wide area, the positive feedback behavior that allows nearby individuals to converge on interesting targets, and the negative feedback to adapt to changes in the local environment [21]. Where ants use chemical pheromones, robots can use radio communication to induce a stigmergy interaction. When a robot finds an interesting target, it can broadcast an “I found something interesting” message. Any robot within range of the signal could then follow the signal toward the broadcasting robot and the target. As it travels, the new robot broadcasts its own message, which may be detected by a third robot.

While this interaction seems intuitive, a stigmergy-based control system must overcome position estimation errors. The Pathfinder Mission to Mars, for example, lost an entire day of science opportunities because the Sojourner rover erroneously stopped short of the desired destination [12]. Designers usually provide robots with some external source of position information such as GPS, or external cameras with stereo capabilities [12]. Another localization method uses a virtual environment of a known area to accurately determine robot position without an external view of the robot [24]. This technique requires an accurate virtual model of the area around the robot, and real-time interaction between the human and the robot. This may not be available to multi-robot systems on Mars because of communication time delays and intermittent communication window between remote robots and mission controllers. While determining absolute positioning is difficult, reliable calculation of relative position between robots has been demonstrated [11][28]. Accurate relative positioning using these methods are assumed.

2.2 Design

We chose to use the *Teambots* robotic simulation platform to implement the stigmergy-based control system. Teambots “is a Java-based collection of application programs and Java packages for multi-agent mobile robotics research. Teambots supports prototyping, simulation and execution of multi-robot control systems” [2]. Although the Teambots software is no longer actively maintained, the full Teambots source code is freely available, which makes it easier to tailor the base software for this particular application (http://www.teambots.org).

The software considered here simulates a team of robots searching a remote field for interesting targets. The team includes between four and eight small rovers, and a central landing craft (from now on, known as the lander) with capabilities similar to Pathfinder’s

Carl Sagan Memorial Lander, including the ability to accurately calculate rover positions via stereo image processing [3], and relay commands from the human team on Earth to the roving vehicles. In order to focus on the human-robot interaction, we make the following assumptions and constraints:

- Robots have a limited, circular communication range. Robots cannot contact the lander, or other robots that are located outside this circle.
- Robots do not have global positioning information (an accurate world map), for themselves or other robots. Robots use information from simulated sonar in order to avoid obstacles
- Robots do have reliable relative positioning information with other robots
- Robots use a gyroscope or magnetic compass for reliable ordinal directions – i.e. due North is the same relative direction for each robot (robots told to travel due north would move in parallel directions, not converge on a single point).
- Robots are able to recognize important or interesting science opportunities. [6, 8, 27].

The desired behavior is for robots to randomly explore the area within communication and visual range of the lander looking for science opportunities. When a science opportunity is detected, the robot should approach the target, position its instrument, and perform the science task. While a robot is performs the science task, it broadcasts to other nearby robots that it has found a target. Robots that receive the message should generate a vector toward the broadcasting robot. When the new robot gets near the broadcasting robot, it should circle the broadcasting robot looking for additional nearby science opportunities.

This behavior can be broken down into four distinct behavioral states: Wander, Acquire, Perform Science, and Stigmergy

Wander – Generate a random vector. Travel as close to the direction of the random vector as possible without running into another object.

Acquire – When a science opportunity is found, travel toward the science site. The velocity of travel is inversely proportional to the distance from the science site. Avoid obstacles along the way.

Perform Science – Place instrument on target. Experimental procedure is simulated by a small delay between 1 and 100 seconds. While performing simulated science, broadcast the message "*myRobotNumber* found target."

Stigmergy – If a stigmergy message is detected, calculate a vector to the sending robot. Travel in the direction of the stigmergy vector, avoiding obstacles along the way. When close to the broadcasting robot, circle around it by generating a vector pointing 90° to the left. The combination of this vector and the stigmergy vector allows the new robot to swirl around the broadcasting robot while looking for nearby science opportunities.

2.3 Experiment

The simulation was run on a PC with the Linux operating system. The effectiveness of stigmergy-based foraging is explored in a 3x2 full factorial experiment. The independent variables are the number of robots (2, 4, 6) and the presence or absence of stigmergy interactions between robots. The number of target clusters in the world is fixed at five, and distributed throughout the environment. Clusters of four objects are used instead of single objects to simulate several science tasks on one particular target. This is analogous to the Pathfinder Mission, where the Sojourner robot typically collected several types of data at each science site. The size of the world was fixed to a 12m x 12m area centered on the lander. The communication range for each robot was fixed to 4m. One hundred simulations were run for 1,000,000 timesteps, and sampled at an interval of 300 timesteps. During each trial, the robots count the number of timesteps they are in stigmergy mode, and keep a list of the times at which they find targets. The density of robots in the world (or more precisely, the amount of overlap of communication ranges) was expected to affect the results of the frequency of stigmergy interactions. The more robots there are in a given area, the more chances that a robot with detect a stigmergy-message.

2.4 Results

The simulation data was analyzed in order to determine the affects of stigmergy and the number of robots on the number of targets found in the remote environment. As expected, the number of robots affects the number of targets found ($p < 0.001$), and significantly affects the time robots spend in stigmergy mode ($p < 0.001$). It was surprising to learn, however, that stigmergy-based behavior did not significantly affect the number of targets found ($p > 0.2$). Stigmergy behavior had even less effect on the time in which targets were discovered ($p > 0.7$). It was thought that allowing the robots to cooperate in locating targets would speed the detection process. One reason for the lack of improved performance may be that robots grouped close together in stigmergy mode may hinder each other's ability to close in on a target. In this implementation, the robots have to bounce around and try to force their way past other robots to get to the target. This inefficiency may affect the selected performance metrics. The performance of the stigmergy behavior could be improved by making the robots impatient. Robots that are unable to get to a detected target, or are in stigmergy-mode for a certain time without detecting new targets, should get impatient with the current task, give up and continue on into wander mode.

2.5 Discussion

Stigmergy-based behavior might be used to look for targets in the area near the lander. The first scenario simulated multiple robots searching the area near the lander for targets. Past NASA robot missions and field tests have demonstrated that a single robot tends to perform multiple science tasks at a single site [12]. Stigmergy-based robot interactions may facilitate the identification and analysis of science opportunities at a single site by drawing multiple robots to a target and performing data collection in parallel. This experiment measured the effects of stigmergy-based behavior on the performance of the science foraging task

The use of stigmergy behavior requires that a sufficient number of robots be in the area to invoke the group behavior. If the individual robots are thinly dispersed in an area, there may not be another robot within range of a stigmergy-message, and the stigmergy interactions between robots would be lost. The multi-robot system would behave like a standard foraging system. This is also true of ant colonies – if there are too few ants or they are far apart from each other, the pheromone will dissipate before an effective trail can be established.

An interesting dynamic occurs when two rovers get involved with this stigmergy interaction. The original broadcasting robot will generally finish its science task before the second, which is now broadcasting its own stigmergy message. This invokes the stigmergy behavior in the original robot, which then swirls around second robot looking for new targets. This type of interaction between robots might be particularly effective if science opportunities are clustered together, as might the case with several different science opportunities on a single rock.

3. Directed Stigmergy Control for Multi-Robot Systems

3.1 Background

The complexity of the planetary exploration requires that these simple algorithms be augmented by the sophisticated human pattern recognition and reasoning ability. Mission scientists must be able to control the direction in which the robot swarm focuses exploration efforts, and specify the breadth of the search. This is addressed with the Directed Stigmergy control system. The Directed Stigmergy control system is similar to the stigmergy foraging behavior previously discussed, but with the additional constraint that each robot must try to stay within the communication range of two other robots. Their circle of communications, then, must stay together like interconnected rings, as in Figure 1. Robot 2 only has freedom to move inside the union of the communication circles of Robots 1 and 4. When robots 1 and 4 are near, and the intersection of their communication circles large, robot two has greater freedom of movement than when they are far apart. Robots 1 and 4, on the other hand, are only bound to stay within range of robot two, and have more freedom of movement.

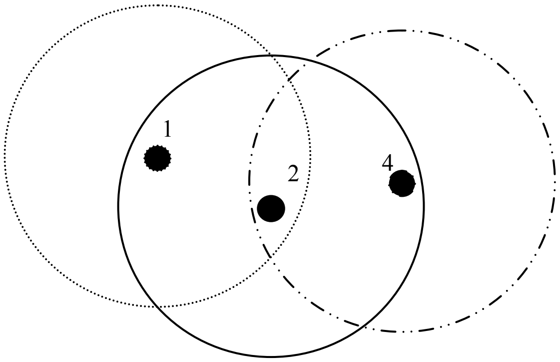

Figure 1. Directed Stigmergy behavior illustrated

The Directed Stigmergy topology has several key advantages over the basic stigmergy behavior. First, the topology supports message passing up and down the members of the chain. The lander, for example, could communicate with a robot linked to the chain but outside its own communication range by sending a message to the closest linked robot, and asking the intermediate robots to forward the message to its intended recipient at the end of the chain. The ability to propagate messages up and down the Directed Stigmergy chain allows the central lander to influence robots outside its radius of communication, and allows distant robots to report results back to the lander more efficiently than returning to the lander's circle of communication. Second, each robot can use the relative position sensors to reliably locate the two robots to which it is linked. By adding these relative vectors together, the lander has a good estimation of the absolute location of all the connected robots. This allows the lander to localize targets outside of its communication range.

To achieve the Directed Stigmergy topology while maintaining the advantages of robustness, scalability, and parallelism inherent in multi-robot systems. inspiration is drawn from the means in which atoms bond to form molecules – the Valence Shell Electron Pair Repulsion Theory (VSEPR) [20]. According to VSEPR, electrons surround an atomic nucleus in various energy shells – with higher energy shells requiring more energy to hold the electron in orbit around the nucleus than lower energy shells. Each shell can hold a certain number of electrons. Atoms seek to minimize its total amount of energy, so higher level energy shells cannot contain electrons until the lower energy shells have been filled. A bond is formed between two atoms when they share electrons in order to fill in their valence shell. For example, sodium (Na) and chlorine (Cl) naturally combine to form sodium chloride (NaCl, table salt) because the NaCl molecule has lower energy than the two individual atoms. The VESPR model is important to the Directed Stigmergy behavior because it provides a model for creating bonds between individuals without an explicit command from a central body or a predefined list of mates to which it is compatible. Instead, each individual atom seeks to satisfy the local need to minimize energy, and naturally forms bonds with other atoms that best satisfy this need.

This valence shell analogy is implemented in the Directed Stigmergy control system architecture. Each robot looks to bond with two other robots (a valence shell n=2). Each robot begins a trial making itself known to other robots by broadcasting its unique ID. When a robot receives a message, it first checks if its valence shell is full, or if it needs to link with more *buddies*. If more buddies are required, it iterates through all the received broadcast messages, and calculates a vector to the robot whose id number is closest to its own. Robots form links with a single buddy each time step. When the required number of buddies have been found (the valence shell is complete) the wander behavior continues, under the Directed Stigmergy constraints that it must stay within communication range of its chosen buddies. All robots continuously broadcast, allowing the system to accommodate additional robots at the end of the chain, or reconfigure the links of the chain in order to assemble robots in ascending numerical order. If, for some reason, a break in the chain occurs, this also helps fix the chain when the two ex-buddies reestablish contact.

Each robot maintains a vector to another robot. This vector is scaled exponentially by the formula $V = e^{r - D_s,}$ where V = the output vector to the target robot, r = the distance to the target robot, and D_s = the safety distance, outside which V becomes large. For this experiment, D_s was set to 80% of the maximum communication range of the sensors. In a real world setting, in which environmental obstacles may reduce signal strength, D_s could be dynamically set based on the signal strength. When two linked robots are near each other, the Directed Stigmergy vector has a small magnitude, and has little effect on the eventual movement of the robot. The magnitude of the random wander vector is larger than the Directed Stigmergy vector, and robots are generally free to move around as though there is no interaction between the two. As the distance approaches the limits of the communication range, however, the Directed Stigmergy vector becomes large very quickly, and exerts a strong influence over the trajectory of the robot to move in the direction of the other robot. This occurs simultaneously in both robots, and the robots quickly moves toward one another, until the Directed Stigmergy vector becomes small, and they are again free to travel randomly.

The combination of position sensing and message passing provides a means for a human operator to direct the emergent behavior of the multi-robot system. A human command to send a robot due East, for example, is simply translated into a vector (1,0). This operator-defined command is simply added along with the rest of the stigmergy-based wander vector. If given adequate weight, the control vector will override the wandering behavior of the basic foraging system.

A Pendant GUI was developed to direct the robot swarm. The user can drag the mouse over the available space to define a control vector in the desired direction for the robot swarm. The angle theta and the magnitude r of this vector are better thought of as a direction and a weight. The direction of the vector guides the direction the robots will travel. The weight of the vector defines the extent to which operator commands affect the direction of travel. If the control vector has a small weight, the random wander behavior will have proportionally more control over the eventual direction of the robot. The robots stay much closer to the lander, but wander more widely about the specified direction. This is shown in Figure 2. If the control vector has a larger weight, the tendency to travel in the specified direction will be greater than the tendency to move in a random manner. All the robots travel in the same direction, but are constrained to stay in contact with their two buddies. The result is that the robots stretch out in a line, and wander only a limited distance from that line, as shown in Figure 3.

The lander broadcasts the user-defined control vector to all robots within communication range. Receiving robots add the control vector to their Wander vector. The control vector message is also broadcast to robots within its communication range. In this manner, the message is relayed from the lander to all the robots within the chain.

This control system has four basic behavioral states: Wander, Wander Near Buddies, Acquire, and Perform Science. The Wander Near Buddies state which implements the Directed Stigmergy behavior:

Wander – Generate a random vector. Travel as close to the direction of the random vector without running into another object. If any other robots are broadcasting Directed Stigmergy messages, go to *Wander Near Buddies*.

Wander Near Buddies – Generate a random vector. For all detected robots, find the ones with ID numbers nearest own. Generate vectors to those robots. Read the commanded direction message, if any. Sum these vectors. Travel as close to the direction of the resultant vector as possible without running into another object.

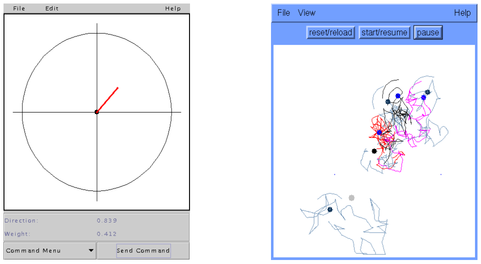

Direction: 0.839
Weight: 0.412

Figure 2. Directed Stigmergy control system with small magnitude command vector from operator

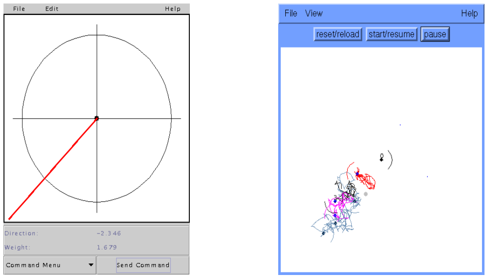

Direction: -2.346
Weight: 1.679

Figure 3. Directed Stigmergy control system with large magnitude command vector from operator

Acquire – When a science opportunity is found, travel toward the science site. The velocity of travel is inversely proportional to the distance from the science site. Avoid obstacles along the way.

Perform Science – Place instrument on target. Experimental procedure is simulated by a small delay between 1 and 100 seconds. While performing simulated science, broadcast the message "*myRobotNumber* found target."

The robots continue to search for science opportunities as they did with the stigmergy-based behavior alone. Robots that discover science opportunities can use communication propagation to send data to the lander in the same way the lander uses communication propagation to control distant robots. When a robot is in the Science state, it broadcasts a Science Location Message indicating it has located an interesting target. This message describes the robot that made the discovery, and the location at which the discovery was made. The robot places its own ID number in the discovering robot field, and initializes the location vector to (0,0). This message is received by the next robot in the chain. Upon receiving that message, the receiving robot adds to the location field the vector pointing to the originating robot. It then rebroadcasts a new, updated Science Location Message, which is received by the next robot in the chain. This third robot adds the vector to the second robot to the location field, and rebroadcasts the message, et cetera, until the message is received by the lander. The lander, then, receives the vector sum of all the links in the chain; the resultant vector is an estimation of the location of the original target.

3.2 Experiment

An operator can use the Pendant GUI to send a direction command from the lander to the mobile robots. One measure of the effectiveness of this type of control is the mean distance from the robots to the lander over time after sending a given command vector. An operator can be said to wield control over the robot system if the robots travel in the direction of the command vector, and if the robots travel farther for a large input vector than they do for a small, or no input vector.

Seven robots were placed in the world near the lander (their initial positions were not randomized, they were placed in a group near the lander and each other in order to ensure the Directed Stigmergy topology generates properly). Seventy-five simulations were run for 1,000,000 timesteps each. The lander broadcasts a zero magnitude direction vector for the first 50,000 timesteps, at which time the control magnitude was increased to one of three magnitudes: large (weight = 1.5), small (weight = .5), and zero (weight = 0.0). In the language of control theory, we are seeking to measure the step response of the Directed Stigmergy system to the control vectors. At each timestep, the robots record their current distance and direction from the lander into a temporary file. The temporary file is then processed to calculate the average robot distance from the lander at each time step. After each set of data from the large, small, and zero control vectors is collected, a separate script calculates the mean of the average robot distance from the lander across all trials at every timestep.

3.3 Results

The results are summarized in Figure 4, which displays measure of the mean robot distance averaged over 75 trials for three different weights of command vectors. In all cases, the robots stay relatively near the lander until the command vector begins broadcasting at timestep 50,000. The figure clearly indicates that robots range farthest from the lander when a large command vector is broadcast, and stay closer to the lander when a smaller vector is broadcast. There is virtually no change in average robot positions when the input vector remains at zero magnitude. This corresponds with the expected behavior of the system. It is important to note that in this simulation the lander can only directly communicate with robots within four meters. In conventional Mars robot systems, mobile robots must stay within this range of the lander. The Directed Stigmergy control system effectively increases the circle of influence of the lander, allowing scientists more area in which to explore, and creating new opportunities for scientific discovery.

Human control of the multi-robot system is necessary because the system is only as useful as the information it provides or the scientific discovery it makes possible. The clear mapping between the direction of the arrow and the direction to send the linked robots was quickly understood by users that tested the system. Once the users understood that the length of the arrow defined "how much you want to go in that direction" instead of defining a particular point in the space around the center, they were able to effectively maneuver the robot team as they saw fit. The pendant makes it possible to specify a constrained search in a narrow band (Figure 2) or a broad search over a wider area or (Figure 3) without compromising the decentralized nature of the multi-robot system. The robots utilize communication propagation to transfer the request to move in a particular direction from the base of the chain to robots at the periphery. Communication propagation, then, provides a means to affect the behavior of robots outside the direct communication range of the lander.

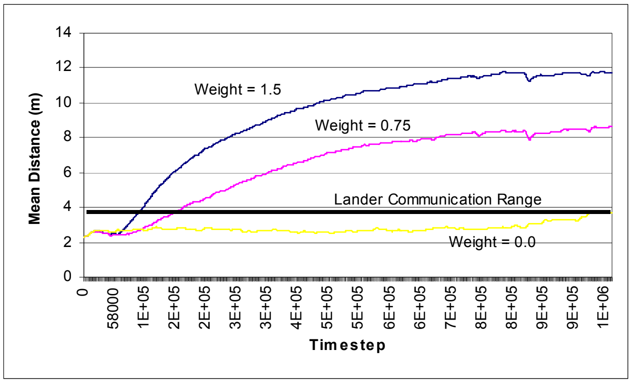

Figure 4. Mean Robot Distance Over 75 Trials

The Directed Stigmergy control system is set up so that each individual robot tries to stay near two other robots within its communication range whose id numbers are nearest its own. Each individual in the system determines which robots to link to dynamically, and without any knowledge of the location or id number of other robots in the area. Instead, each individual recognizes the id numbers of robots in its communication range, and tries to stay close to the robots that have id numbers nearest its own. A robot, therefore, will generate a link to two other robots whether their ID numbers are sequential or not. In one simulated example, six robots were placed in the world. Four of the robots were initialized to be near each other in the center of the world. Robots 2 and six began far away. Robot 1 initially generates a link with robot three, the robot with a higher id and closest to its own. During a random traversal, robot six encountered robot one. Robot 6 had not yet found a buddy, so it immediately formed a link with robot one. Robot 1 however, maintained its bonds with the base and robot three because those ID numbers were closest to its own. Robot 6 stayed near robot one until it found and linked with robot three. In this manner, robot six quickly moved up the chain until it found its place at the end. Similarly, robot two met the chain at the end, near robot five. It traveled down the chain until it assumed its position between robots one and three. Surprisingly, the addition of a new robot in the middle of the chain did not split it in two. The ability of the system to quickly adjust to accommodate new team members indicates the order in which the robots link is not predefined; rather, the emergent behavior of the system is such that the robots set themselves in ascending order. This is consistent with behavior one would expect from the decentralized system. Furthermore, this behavior may assist NASA scientists in locating and reclaiming lost robots. If a robot wanders away from the group, one can point the chained robots in the direction in which the lost robot was last seen.

4. DISCUSSION

The Directed Stigmergy is a unique way to give the human operator control over the emergent behavior of the complex, distributed multi-robot system. Human operators interact with the Directed Stigmergy system as an Open Loop controller, rather than a closed loop controller found in teleoperation of single or multi-agent systems. The system is therefore less sensitive to the effects of time delay between the human controller and the remote system. The pendant GUI proved to be an effective means to control the multi-robot system and extend the area available for scientific exploration.

This work demonstrates that control systems for multiple robots can consist of simple reactions to low-level sensor input. The multi-robot control system relies on the sum of the simple actions of the individuals to complete the given task. Fast and accurate relative position measurements between robots, combined with communication propagation provides a means to localize both robots and targets outside the effective range of the lander. The Directed Stigmergy position estimation still has error, which will accumulate based on the number of robots in the chain, and the estimated error of the positioning sensors. Without Directed Stigmergy, however, it would not be possible to estimate the position of targets outside the range of the lander at all – this estimation is better than no estimate at all.

The system also demonstrates the transmission of information from distant robots to the lander via transmission propagation. The interconnected communication rings permits localizing interesting targets outside the visual range of the base via communication propagation. It also allows robots to more efficiently send data to the lander by sending message down the chain instead of having to travel back within communication range of the lander. In this simulation, the transmitted Science Location Message simply contains the ID of the robot that found the target, and the resultant of the vectors that define its location. In real life, the contents of this message could also contain an image buffer, spectrometer data, or any scientific data set [10]. Finally, Directed Stigmergy provides operators a means to search for, capture and control lost robots.

5. CONCLUSIONS

This paper demonstrates a novel paradigm for human control over a group of robots. Mobile robots are called into action to search and explore areas that are inaccessible or dangerous to humans – areas such as the surface of Mars, damaged urban structures, or unexploded minefields. In order to effectively and efficiently use the robot team to explore the remote environment, the human operator must be able to direct the emergent behavior of the group of robots, instead of having to send detailed instructions to each individual. Directed Stigmergy provides a means for the operator to direct the high level goals of the robot team, and maintains the advantages of decentralized, stigmergy-based exploration.

This research was inspired by both natural and artificial multi-agent systems. Natural organizations like birds, termites, and ants help us appreciate the elegant simplicity and utility of decentralized systems. These natural systems also contribute the concept of stigmergy, which is used to organize and coordinate behavior of simple, reactive agents. Stigmergy-based multi-agent control systems such as the Directed Stigmergy control system contribute to planetary exploration by not only speeding the scientific data collection process, but also by extending the area in which robotic systems can safely explore. The Directed Stigmergy system also maintains realistic assumptions and constraints. Arming laboratory and field mobile robots with these control systems would be an excellent means to test and refine the viability of the simulation.

The future is bright for this line of research. One could start by measuring the effects of dynamic human control over internal reactive behavior parameters, especially the weight of noise parameters, on the performance of the foraging behavior. If the random vector is given low weight, the system would probably act as though it were directly teleoperated. A heavily weighted random vector, on the other hand, would make human control ineffectual. Does the weight of the random vector have a point of maximum performance? How does this point vary with specific tasks? Can a human operator reliably set this point, and if not, what algorithm could be used to automate the process? It would also be interesting to utilize Directed Stigmergy in other search and foraging tasks such as urban search and rescue. This could be done by competing in the Robocup Urban Search and Rescue competition [22, 26].

6. ACKNOWLEDGMENTS

This work is partially supported by NASA's Applied Information Systems Research Program under NASA grant NNG05GA51G and the Iowa College of Engineering GAANN grant P200A030132. Thanks to John Lee for his insights and suggestions, as well as the comments from the anonymous reviewers.

7. REFERENCES

[1] Kobayashi, A. & Nakamura, K. (1983). Rescue Robot for Fire Hazards. *Proceedings of the 1983 International Conference on Advanced Robotics,* 91-98.

[2] Balch, T. (1998). Behavioral Diversity in Learning Robot Teams. *Computer Science*. (pp. 200). Atlanta, GA: Georgia Institute of Technology.

[3] Blackmon, T. T., Ngyuen, L., Neveu, C. F., Rasmussen, D., Zbinden, E., Maimone, M. W., Matthies, H., Thayer, S., Broz, V., Teza, J., Osborn, J., Herbert, M., Thomas, G., & Steele, J. (1999). *Photorealistic virtual reality mapping system for Chornobyl accident site assessment*. Paper presented at the Human Vision and Electronic Imaging IV, San Jose, CA.

[4] Blitch, J. & Mauer, R. (2000). Knobsar: A Knowledge Based System Prototype for Robot Assisted Urban Search and Rescue. *Simulation*, *66(6)*(6), 375-391.

[5] Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems (Santa Fe Institute Studies on the Sciences of Complexity).* Oxford University Press, USA.

[6] Bresina, J., Golden, K., Smith, D., & Washington, R. Increased Flexibility and Robustness of Mars Rovers*. The 5th International Symposium on AI, Robots, and Automation in Space*.

[7] Casper, J. & Murphy, R. (2003). Human-Robot Interactions During the Robot-Assisted Urban Search and Rescue Response At the World Trade Center. *Systems, Man and Cybernetics, Part B, IEEE Transactions on*, *33(3)*(3), 367-385.

[8] Chien, S., Sherwood, R., Tran, D., Cichy, B., Rabideau, G., Castano, R., Davies, A., Mandl, D., Frye, S., & Trout, B. (2005). Using Autonomy Flight Software to Improve Science Return on Earth Observing One. *Journal of Aerospace Computing, Information, and Communication*, *2*, 196-216.

[9] Drogoul, A. (1995). When ants play chess - (Or can strategies emerge from tactical behaviors?). *From Reaction to Cognition*. (pp. 13-27). Berlin, Germany: Springer-Verlag.

[10] Estlin, T., Gaines, D., Fisher, F., & Castano, R. (2005). *Coordinating Multiple Rovers with Interdependent Science Objectives*. Paper presented at the Autonomous Agents and Multi-Agent Systems Conference, Utrecht, Netherlands.

[11] Fox, D., Burgard, W., Kruppa, H., & Thrun, S. (2000). A Probabilistic Approach to Collaborative Multi-Robot Localization. *Autonomous Robots*, *8(3)*(3), 325-344.

[12] Golombek, M. P. & Anderson, R. C. (1999). Overview of the Mars Pathfinder Mission: Launch Through Landing, Surface Operations, Data Sets, and Science Results. *Journal of Geophysical Research*, *104(14)*(14), 8523-8553.

[13] Grassé, P. P. (1959). La Reconstruction Du Nid Et Les Coordinations Inter-Individuelles Chez Billicsitermes Natalensis Et Cubitermes Sp. La Théorie De La Stigmergie: Essai D'interprétation Du Comportement Des Termites Constructeurs. *Insect Sociology*, *6*, 41-80.

[14] Grassé, P. P. (1984). Termitologia, Tome II. *Fondation des Sociétés. Construction.* Paris: Masson.

[15] Lee, J. D. Emerging Challenges in Cognitive Ergonomics: Managing Swarms of Self-Organizing Agent-Based Automation. *Theoretical Issues in Ergonomics Science*, *2(3)*(3), 238-250.

[16] Maimone, M. W., & Biesiadecki, J. J. (2006). *The Mars Exploration Rover Surface Mobility Flight Software: Driving Ambition*. Paper presented at the IEEE Aerospace conference, Big Sky, Montana.

[17] (2006). Mars Exploration Mission Status Page. http://marsrovers.jpl.nasa.gov/mission/status.html

[18] Murphy, R., Casper, J., Micire, M., & Hyams, J. (2000). Mixed-Initiative Control of Multiple Heterogeneous Robots for Urban Search and Rescue. *Univ. Central Florida, Orlando, CRASAR-TR2000-1*.

[19] NASA Mars Team Teaches Old Rovers New Tricks to Kick Off Year Four. Retrieved January 4, 2007, from http://marsrovers.jpl.nasa.gov/newsroom/pressreleases/20061228a.html

[20] Oxtoby, D. N., & Freeman, N. (1994). *Chemistry: Science of Change.* Chicago, IL: Saunders College.

[21] Resnick, M. (1994). *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds (Complex Adaptive Systems).* Mit Pr.

[22] Scrapper, C., Balakirsky, S., & Messina, E. (2006). Moast and Usarsim: A Combined Framework for the Development and Testing of Autonomous Systems. *Proceedings of SPIE*, *6230*, 62301T.

[23] Sheridan, T. B. (2002). *Human and Automation: System Design and Research Issues. Part II: Design of Human-Automation Systems.* Santa Monica, CA: Human Factors and Ergonomic Society and John Wiley & Sons, Inc.

[24] Steele, J., Fitzgerald, Thomas, G., & Blackmon, T. T. (1998). *An Operator Interface for a Robot-mounted, 3D Camera System: Project Pioneer*. Paper presented at the IEEE Virtual Reality, Houston, TX.

[25] Stoker, C. R., Zbinden, E., Blackmon, T. T., & Kanefsky, B. (1999). Analyzing Pathfinder Data Using Virtual Reality and Superresolved Imaging. *Journal of Geophysical Research*, *104(E4)*(E4), 8889-8906.

[26] Wang, J., Lewis, M., Hughes, S., Koes, M., & Carpin, S. (2005). Validating Usarsim for Use in Hri Research. *Proceedings of the 49th Annual Meeting of the Human Factors and Ergonomics Society, September*.

[27] Washington, R., Golden, K., Bresina, J., Smith, D. E., Anderson, C., & Smith, T. *Autonomous Rovers for Mars Exploration*. Paper presented at the IEEE Aerospace Conference, Aspen, CO.

[28] Wu, J., Chan, W., & Thomas, G. (2001). Rapid and Accurate Inter-Robot Position Determination in Robot Teams. *IEEE Transactions On Instrumentation And Measurement*, *50(1)*(1), 163-168.