Haptic Shape-Based Management of Robot Teams in Cordon and Patrol

Samuel J. McDonald and Mark B. Colton
Department of Mechanical Engineering
Brigham Young University
Provo, Utah 84602

C. Kristopher Alder and Michael A. Goodrich
Department of Computer Science
Brigham Young University
Provo, Utah 84602

ABSTRACT

There is a growing need to develop effective interaction methods that enable a single operator to manage a team of multiple robots. This paper presents a novel approach that involves treating the team as a moldable volume, in which deformations of the volume correspond to changes in team shape. The team possesses a level of autonomy that allows the team to travel to and surround buildings of interest in a patrol and cordon scenario. During *surround* mode, the operator explores or manipulates the team shape to create desired formations around a building. A spacing interaction method also allows the operator to adjust how robots are spaced within the current shape. Separate haptic feedback is developed for each method to allow the operator to "feel" the shape or spacing manipulation. During *travel* mode, the operator chooses desired travel locations and receives feedback to help identify how and where the team travels. Results from a user study suggest that haptic feedback significantly improves operator performance in a reconnaissance task when task demand is higher, but may slightly increase operator workload. In the context of the experimental setup, these results suggest that haptic feedback may contribute to heads-up control of a team of autonomous robots. There were no significant differences in levels of situation awareness due to haptic feedback in this study.

Keywords

human-swarm interaction; haptic feedback; situation awareness

1. INTRODUCTION

There are many current and future scenarios in which a human must manage a team of robots. Potential scenarios include wilderness search-and-rescue [9], rescue operations in buildings damaged by fire or earthquake [1, 13], searching buildings by law enforcement agencies [10], pollution monitoring and clean-up [11], and military patrol and cordon operations in an urban environment [17]. In each scenario, the robot team serves as an extension of the operator's ability to gather information in complex and often dangerous environments. Enabling the human operator to manage the robot team in an intuitive, effective, and time-efficient manner is therefore critical to the success of operations involving robot teams.

A common approach in controlling autonomous robots (in use by current military, law enforcement, and search-and-rescue agencies) is for a single robot to be controlled and monitored by one or many human operators (see, for example, [20]). This interaction model is clearly not ideal if the objective of employing autonomous robots is to augment the capabilities of humans and maximize the information-gathering capabilities of the team [8, 26]. A preferable interaction model is for a single human operator to control multiple autonomous robots [12, 16, 19]. The effectiveness of such an approach is limited by the operator's ability to command the actions of multiple agents and receive information about the state of the robot team, while accomplishing a primary task, whether it be search-and-rescue, surveillance, etc.

When the robot team possesses appropriate autonomy for the given scenario, the problem becomes one of "team management" rather than "robot control," enabling the operator to focus on task objectives and interpretation of gathered data, rather than on the robots. With the proper management interface, the operator is able to devote more attention and resources to a primary task and remain "heads-up" while managing the team. In this paper, "heads-up" refers to having a sufficient level of competency in a single task to focus visual attention elsewhere on other important tasks. An operator who is "heads-down" tends to focus solely on a single task and may experience difficulty in responding quickly or accurately to additional tasks that require visual attention. In a human-robot interaction scenario, this could occur when an operator is so focused on a graphical user interface that awareness of surroundings or response to incoming commands is degraded. While considering a single-operator-multiple-robot interaction, enabling the operator to achieve heads-up control is challenging.

In general, when multiple tasks demand attention from the same sensory channel, interference can degrade performance or possibly result in task overload [27]. The operator's ability to multi-task can improve with the use of a multimodal interface, using both visual and haptic feedback. Haptic feedback has been shown to reduce collisions when piloting individual robotic vehicles [5, 23], and provide a sense

HRI '17, March 06-09, 2017, Vienna, Austria

DOI: http://dx.doi.org/10.1145/2909824.3020243

of team-level properties when managing multiple robots [21, 22, 24, 25]. There are, however, possible drawbacks to using haptic feedback, as it has been shown to increase operator workload in some studies [14, 15].

The objective of the present work is to investigate the effects of haptic feedback on heads-up control, situation awareness (SA), and workload while managing a small team of robots in reconnaissance and surveillance tasks. A user study was conducted to verify the effectiveness of this approach in completion of a primary task, and quantify its effects on operator workload, situation awareness, and other measures.

2. ROBOT BEHAVIORS

In this work, a simulated team of robots is managed in a military patrol and cordon exercise, in which the robot team must autonomously travel between specified locations (patrol) and surround buildings (cordon), although the interaction approach will be generalizable to other scenarios that require a human operator to control the movement and distribution of autonomous robots. The individual and team behaviors of the robots in this work were developed in [2, 3]; an overview of those behaviors are given in this section. In this scenario, the operator uses the team to search around buildings of interest in an urban environment. Only the exterior regions of the building are considered while searching. The robots are modeled as autonomous, omni-directional agents, and for simplicity, their movement is planar in x and y.

There are two modes available to the team, *travel* and *surround*. Graphs are used to describe the relationships between robots, with each robot corresponding to a node in the graph. An example formation for each mode is shown in Figure 1.

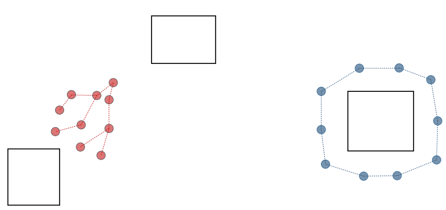

Figure 1: *Travel* (left) and *surround* (right) modes. Individual circles are graph nodes (robots), and colored lines are graph connections. *Travel* mode is represented by a tree formation. *Surround* mode is represented as a spanning ring formation around one building.

A *travel* is used when the team is traveling between buildings, and represents the patrol action. In *travel* mode, the team formation is governed by the graph connections, where each robot follows another lead robot. This allows the robots to obey forces that repel them from the current building and travel to a new one. The resulting swarm behavior is that of a spanning forest, where the graph connections can be traced to one or two robots that lead the team in the direction of travel. The robots attempt to stay close together until a nearby building is identified, at which point the *surround* mode is enabled.

A *surround* is used when the team is surrounding a building, and represents the cordon action. In *surround* mode, the model determines a surrounding shape and the team forms a spanning ring around a building. The team can successfully surround buildings of convex or concave geometry. Transitions between these two modes are handled by the robot behaviors model. In either mode, robots are attracted to nearby buildings, but repelled by building boundaries to prevent the team from passing through buildings.

3. HUMAN-TEAM INTERACTION

Several interaction methods were developed to allow a single operator to manage a team of robots. Some methods are enabled by default and others are enabled by engaging buttons on a haptic interface. The physics of each interaction method are explained in detail to provide insight into the underlying haptic sensations, thereby giving greater substance to the corresponding human interaction.

3.1 Modeling Clay: A Haptic Metaphor

In a brainstorming breakout session at the 2012 AAAI Fall Symposium on Human Control of Biological Swarms, a breakout group developed the idea of using a deformable medium, such as modeling clay, as a "joystick" to command the distribution of large-scale swarm-like teams of homogeneous vehicles. Diana *et al.* [7] demonstrated a molding scheme in which an operator formed modeling clay into various shapes in the view of an overhead camera and a team of robots replicated the formation commanded by the shaped clay. We modify the modeling clay metaphor so that a human can shape the distribution of robot teams by manipulating a *virtual* deformable volume through stretching, pulling and other operations. The modeling clay metaphor forms the basis for the haptic sensations that the user feels while distributing the robot team. Note that, unlike the work in [7], physical modeling clay is not used in our method; the concept and physics of modeling clay are used to generate the visual and haptic representation of the robot team.

3.2 Force Node Network

In surround mode, virtual modeling clay is modeled discretely by placing potential force field spheres at the location of each robot in the spanning ring and at multiple points between robots. In essence, these potential spheres form the nodes of a "force graph" on which the haptic interaction forces and graphical representation of the deformable volume are based. An example of how the force nodes span between robots is shown in Figure 2. Notice that some force nodes are specifically assigned to a robot location. A virtual haptic cursor is also displayed to the operator that maps to physical haptic device movement.

To provide a dynamic system that the operator can manipulate, the potential spheres also form the mass nodes of a virtual mass-spring-damper network. Virtual frictional forces are also used to model plastic deformation of the volume. Each force node in the network is positioned a distance $\delta_{spacing} = r_n/4$ from neighboring nodes, where r_n is the node radius. In other words, nodes are placed close enough together to cause eight consecutive nodes to overlap. This provides sufficient node density for compelling haptic interaction and responsive team interaction. Parameter values for mass, spring, and damping constants were subjectively chosen to allow each node to maintain sufficient distance relative to neighboring nodes, stabilize the network, and create a distinguishable volume with which to interact. Each

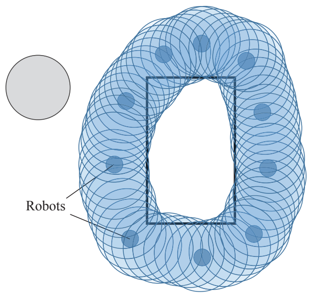

Figure 2: Force nodes (light blue) form a spanning ring shape after the team (dark blue) has surrounded a building. The haptic cursor (gray) is controlled by the operator.

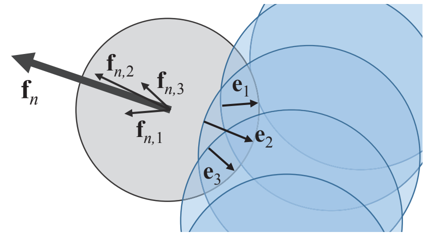

Figure 3: The haptic feedback force, $\mathbf{f}_n$, is produced due to interaction between the haptic cursor and several force nodes.

node has a z-position of zero, such that one hemisphere is above the ground surface, and the other hemisphere is below the surface. The deformable ring naturally maps to the spanning ring in surround mode, and can be mapped to the spanning forest in travel mode; see sections 3.3 and 3.4. The deformable ring forms the basis for computing the haptic feedback force felt by the operator.

3.3 Surround Interaction

In *surround* mode, there are three types of interaction modes with corresponding force feedback algorithms. They are shape exploration, shape manipulation, and spacing manipulation. Each mode is designed to encourage collective, team-level control, rather than micromanagement of each robot.

3.3.1 Shape Exploration

This mode allows the user to explore the shape of the robot team without influencing it. The force nodes are fixed and the operator cannot cause them to move during contact, but force feedback is still provided. The force feedback is computed as a function of penetration distance on each of the nodes with which there is contact. Thus, the nodal force feedback during this mode is given by

$$\mathbf{f}_n = -\sum_{n=1}^{N} k_n \mathbf{e}_n, \tag{1}$$

where k_n is the node stiffness coefficient and $\mathbf{e}_n$ is the penetration vector for the n-th contacted node. Contact between the haptic cursor and virtual ground surface results in a ground feedback force, $\mathbf{f}_g$. The total haptic feedback force in shape exploration is then

$$\mathbf{F}_{sh} = \mathbf{f}_g + \mathbf{f}_n. \tag{2}$$

The nodal force $\mathbf{f}_n$ allows the user to feel the team distribution. The ground-plane force $\mathbf{f}_g$ provides a virtual fixture that enables the user to easily keep the haptic cursor in the same plane as the robot team, making exploration of the team's shape more convenient. A visualization of the feedback produced by coming in contact with a set of fixed force nodes is shown in Figure 3.

3.3.2 Shape Manipulation

During shape manipulation, the network nodes are allowed to move in response to contact from the haptic cursor, as shown in Figure 4. The motion of the nodes is calculated from the governing equations of a mass-spring-damper network. The robot positions are updated from the adjusted robot node positions. This method allows the operator to quickly adjust the shape of the robot team without needing to interact with each robot individually.

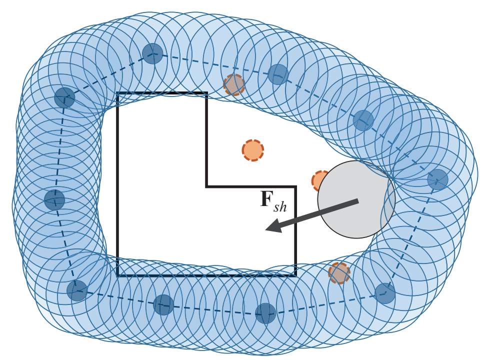

Figure 4: Adjusted team shape, showing previous (orange) and current (dark blue) robot positions. The operator experiences a feedback force, $\mathbf{F}_{sh}$.

The force feedback is governed by the same equations presented in shape exploration. In the case of shape manipulation, equal and opposite forces are exerted on the nodal network, resulting in motion and deformation of the network. The feedback is designed to provide forces similar to what one would feel while plastically deforming physical modeling clay. The operator must use the haptic cursor to interact with the shape and exert forces on the network that exceed the frictional forces, which are included to simulate plastic deformation, in which the force nodes remain in a location even when the haptic cursor is no longer exerting a force on the network.

3.3.3 Spacing Manipulation

The autonomous behavior of the team creates an initially uniform distribution, meaning that each robot maintains an equal distance from neighboring robots throughout the spanning ring shape. Spacing manipulation mode allows the operator to adjust inter-robot distances, moving them closer together in some portions of the shape and farther away from each other in others. The method allows the operator to simply gesture toward a set of robots in the team with the haptic cursor, causing robots to concentrate more (decrease spacing) in the direction of cursor movement relative to the team's center. An example of spacing manipulation is shown in Figure 5.

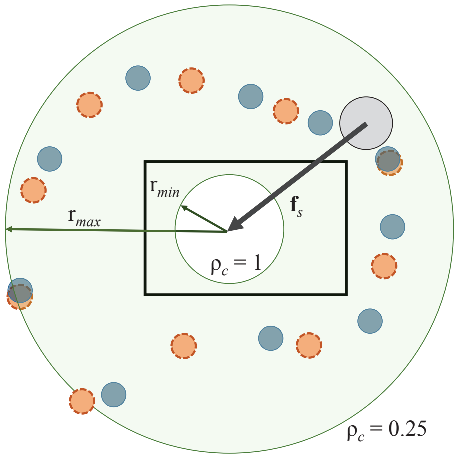

Figure 5: Adjusted team spacing, showing previous (orange) and current (dark blue) robot positions as a result of moving the haptic cursor away from the center of the formation. The robot spacing is a function of the location of the haptic cursor within the manipulation boundary (green).

Since each robot is assigned a node in the network, neighboring nodes act as potential new locations for each robot. The robot spacing within the shape is modeled as a set of equilibrium distances, L_i, that act in the discrete space of the nodal network. Given a team of M robots, the equilibrium distance L_i is assigned to the space between the i-th and the $i+1$-th robots, and the M-th and 1st robots share L_M. These distances determine how the robots are distributed within the shape, irrespective of the dynamics that govern the shape in a more continuous environmental space. Nodal network properties such as total perimeter, P_n, and center of mass, $\mathbf{C}_n$ are used to adjust these distances based on cursor location. The initial equilibrium distance is $L_0 = P_n/M$ and each equilibrium distance is governed by the relationship $L_i = \rho_i L_0$, where ρ_i is an equilibrium distance factor.

As the cursor moves a distance r away from $\mathbf{C}_n$, the robot with closest proximity to the cursor is denoted robot c. The equilibrium distance factor for this robot is set first and is given by

$$\rho_c = \begin{cases} 1 & r < r_{min} \\ 1 - 0.75\frac{r - r_{min}}{r_{max} - r_{min}} & r_{min} \leq r \leq r_{max}, \\ 0.25 & r > r_{max} \end{cases}$$

where r_{max} is the distance from $\mathbf{C}_n$ to robot furthest away from $\mathbf{C}_n$ and $r_{min} = 0.25 r_{max}$. From there, other factors are calculated as

$$\rho_i = \rho_c + \frac{4(1-\rho_c)}{M} n, \tag{3}$$

where n is the minimum number of robot nodes aways from robot c. Equation (3) ensures that $\sum_{i=1}^{M} L_i = P_n$, which is to be expected.

The total distance between robot nodes, d_i, is the sum of the distances between neighboring nodes that connect one robot node to another. Robot nodes transfer to neighboring nodes in order to maintain $|L_i - d_i| \leq 0.25 L_i$. This tolerance allows for a range of distances to acquire equilibrium and prevent desired robot positions from switching unnecessarily between neighboring nodes. Due to the finite amount of nodes in the network, robot node positions simply establish the desired positions of robots. The continuous dynamic motion is handled separately by applying forces to the robots and guiding them directly to these positions.

When this type of manipulation is enabled, a force, $\mathbf{f}_s$, is produced that is directed toward $\mathbf{C}_n$, proportional to the cursor's distance away from $\mathbf{C}_n$, such that

$$\mathbf{f}_s = k_s(\mathbf{C}_n - \mathbf{p}_{cursor}) - b_s \mathbf{v}_{cursor}, \tag{4}$$

where k_s is the spacing force stiffness coefficient, b_s is the spacing damping coefficient, and $\mathbf{p}_{cursor}$ and $\mathbf{v}_{cursor}$ are the position and velocity of the haptic cursor, respectively. This feedback is designed to give the operator a sense of the gesture direction relative to the team by modeling the force as a virtual spring that connects the cursor to the team's center-of-mass. The magnitude also informs the operator of the strength of the spacing adjustment in that direction. As the operator may be quickly gesturing toward various sections of the team shape, a damping force is provided to prevent the device motions from being too abrupt. This damping effect also helps the device feel more like a grounded joystick rather than a free-moving cursor. The ground force feedback is also active in this mode, so the total feedback force felt by the operator will then be

$$\mathbf{F}_s = \mathbf{f}_g + \mathbf{f}_s. \tag{5}$$

3.4 Travel Interaction

Travel interaction allows the operator to move the team from one building to another. The haptic feedback is divided into relative travel and shape exploration forces, which are explained in the following sections.

3.4.1 Relative Travel

The relative travel force occurs while the team is traveling between buildings. It activates once the travel position is being chosen, and deactivates once the team has switched to *surround* mode. The relative travel force is

$$\mathbf{f}_r = k_r(\mathbf{C}_n - \mathbf{p}_{cursor}) - b_r \mathbf{v}_{cursor}, \tag{6}$$

where k_r is the relative travel stiffness coefficient and b_r is the relative travel damping coefficient. Similarly to the feedback force provided during a spacing manipulation, a small level of damping is used to stabilize the cursor motion while choosing a travel position. The total haptic feedback force felt by the operator during relative travel is then

$$\mathbf{F}_r = \mathbf{f}_g + \mathbf{f}_r. \tag{7}$$

This force is designed to help the operator gain a sense for the relative distance between the cursor and the team. The force is strong enough to guide the device toward the team location, but not strong enough to prevent the operator from moving the device as desired. With training, the operator may interpret these forces to gain a better sense for where and how the team is traveling. Figure 6 shows the process of sending a travel command to the team. The user selects a location by pressing a travel button on the haptic interface. The team moves to that location and the

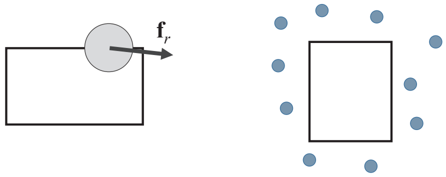

(a) Travel command in process

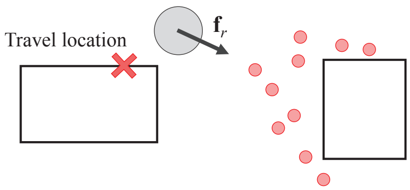

(b) Travel command is applied

Figure 6: This figure shows the travel command process and the direction of the relative force feedback, $\mathbf{f}_r$. In (a), the team is still in *surround* mode and the travel location has been determined, but not applied. In (b), the team begins its motion toward the selected travel location (marked with a red X).

operator experiences the relative travel force from the device as the team travels. This discrete command process is used rather than continuously "dragging" the team via velocity commands so that the operator can use the time on more urgent tasks while relying on the autonomous behavior of the team to complete the travel motion.

3.4.2 Shape Exploration

During *travel* mode, the operator may explore the overall shape of the team by coming in contact with the traveling force nodes. The nodal network is created similar to *surround* mode, except, in *travel* mode, a convex hull [6] is formed around the team and forms an outer boundary that represents the overall team shape.

Since robots are also positioned within the nodal network, a virtual surface is needed to enclose the shape formed by the nodal network. This virtual surface lies tangent with the top of the nodes and within the convex hull. When the cursor comes into contact with this surface, an additional force is produced to simulate an enclosed volume. To help visualize this force, a side-view of the nodal network is shown in Figure 7.

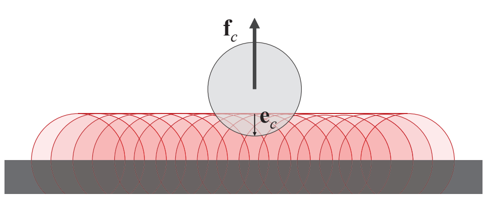

Figure 7: Side-view of interaction with convex hull surface during travel mode.

Similar to the ground force, the convex hull surface is modeled as a flat, stiff surface. The force produced when the haptic cursor comes into contact with this surface is

$$\mathbf{f}_c = -k_c \mathbf{e}_c, \tag{8}$$

where k_c is the convex hull stiffness coefficient and $\mathbf{e}_c$ is the penetration vector into the bounding surface.

The force produced by contacting the force nodes, $\mathbf{f}_n$ is given in Equation 1 and the relative travel force is also in effect, so the total haptic feedback force felt by the operator during travel shape exploration is

$$\mathbf{F}_t = \mathbf{F}_r + \mathbf{f}_n + \mathbf{f}_c. \tag{9}$$

4. USER STUDY

Past research has shown interest in military-based scenarios that require a single operator to teleoperate one or more robots while performing other mission-related tasks [4]. A user study was conducted to examine the effects of haptic feedback on an operator's ability to manage a robot team while performing several tasks. The study primarily focuses on how the presence of haptic feedback affects operator performance, whether the operator uses some or all interaction methods available.

4.1 Experimental Apparatus

A simulation system was designed as a means to experimentally determine how well an operator can manage a team of robots while being required to perform additional mission-related tasks and switch attention between multiple displays. Each participant used a dual-monitor, multimodal workstation as shown in Figure 8. The Novint Falcon desktop haptic device was placed to the right of the monitors and is the primary controller used during the simulation. The operator uses a dominant hand to control the Falcon and a non-dominant hand to enter keyboard input.

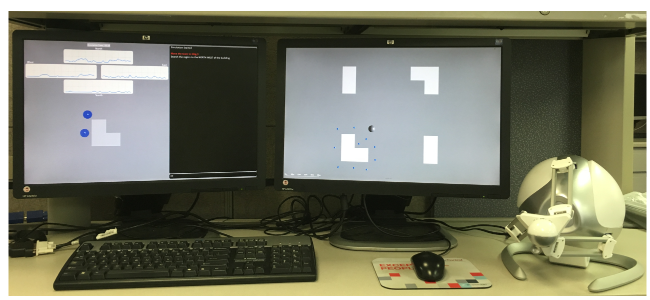

Figure 8: Dual-monitor, multimodal workstation with Novint Falcon desktop haptic device.

4.1.1 Task Description

The operator has two main tasks, which are: (1) Continuously search for and find as many hotspots as possible and (2) Respond to and follow incoming messages. The primary task involves manipulating team shape and spacing in order to maximize the number of hotspots being uncovered. Hotspots are points of interest in the environment, used to simulate a simple reconnaissance task. The operator is only concerned with locating and uncovering hidden hotspots and not in interpreting their meaning. Past research has used hidden points in an environment as a means of measuring team coverage [21]. The operator is provided a sensor reading for each robot. The strength and location of the sensor readings in the environment are used to locate hotspots. Hotspots are hidden in groups of one, two or three, which all need to be detected within a short window of time for the group to be counted as found. This requires the operator to make a specific formation in order to detect and find all the hotspots in each group. For each environment, the percentage of hotspots found is used as the main measure for primary task performance.

Hotspots are detected by means of a sensor attached to each robot. Each sensor reading is binary, in that it has only two possible outcomes, detection or no detection. A true positive detection occurs when a reading is obtained while a hotspot is within sensor range. The true positive rate, p_t, is modeled probabilistically with a Gaussian distribution from the location of the sensor, given by

$$p_t = p_{min} + (p_{max} - p_{min})e^{-d^2/\sigma^2}, \tag{10}$$

where p_{min} is a minimum true positive rate, p_{max} is a maximum true positive rate, d is the euclidean distance from the sensor and σ is the standard deviation of the Gaussian curve. With r_{sensor} being the range of the sensor, the standard deviation was chose to be $\sigma = r_{sensor}/3$. Sensor readings are not always indicative of nearby hotspots. A false positive detection occurs when a reading is obtained while no hotspots are in range. The false positive rate, p_f, is another predetermined constant. Values given to the mentioned constants are $p_{min} = 0.50$, $p_{max} = 0.80$, and $p_f = 0.10$.

The secondary task involves an incoming message feed, which has been preprogrammed to display messages that are timed with the simulation. The operator is asked to respond to incoming messages, which include relocation instructions, search suggestions and probing SA questions to measure the operator's level of SA. The operator responds to relocation instructions and search suggestions by managing the team accordingly. A typed numerical response is required to respond to SA questions.

4.1.2 Multitasking and Heads-up Awareness

Given a single viewpoint of an environment, while it is most convenient and realistic to concentrate all visual information on a single display, our simulator splits information across multiple displays in order to simulate a less than ideal case. For example, robot state information is shown on the main team display, but sensor readings are shown on a second display. Although the user may desire those two pieces of information to be visually displayed on the same screen, they are separated to simulate a situation in which the operator needs to be aware of surroundings rather than being heads-down in the team display. In a practical scenario, unlike this experimental scenario, additional events that may require the operator's attention may include video surveillance, mission planning, mission-related discussions with other people or interruptions due to emergency. Since many of these more realistic scenarios require additional training, a simple division of information was preferred and remains ecologically viable.

It has been posited that the presence of haptic feedback will have a positive effect on an operator's performance while placed in this multitasking scenario. Therefore, experimental comparisons need to be made to test this claim. In situations where task demand is low, it is expected to see little difference between measurements between the presence and absence of haptic feedback. When more task demand is placed on the operator, it is expected that task performance will decrease, SA may decrease and workload may increase as compared to when task demand is lower. If haptic feedback has an influence on the experimental measures, there should be a noticeable difference between the presence and absence of haptic feedback while observing results from environments with high or low task demand.

4.2 Experimental Design

The user study was designed as a full two-factorial experiment with factors being haptic feedback (no haptics/haptics) and world difficulty (easy/hard). Two easy worlds and two hard worlds were developed to accommodate all the combinations of factors and levels. A pilot study was performed to tune the difficulty of the hard worlds until a noticeable difference was seen in task performance between easy and hard worlds. The hard worlds were more difficult in that they required the operator to find more hotspots in the same amount of time. The combinations of factors and levels were balanced by each world type of equal difficulty to prevent one world type from being paired with haptics or no haptics. Results from the pilot study validate this approach.

The experiment was held in the MAGICC Lab on Brigham Young University (BYU) campus and there were a total of 19 participants in the study, of which three were female. The ages of participants ranged from 21 to 30, with an average age of 25. Prior to participation in the study, each participant was given enough time to review and sign a consent form. The purpose of the study was explained in general terms and participants were ensured that they were not required to have any previously knowledge of human-robot interaction to successfully complete the study. Each participant was seated in front of the multimodal workstation shown in Figure 8. The participant took about thirty minutes to complete a tutorial and two practice sessions with and without haptic feedback. Up to that point, the participant could ask any questions about the haptic controller, the simulator or the task description.

The remaining four sessions were each six minutes long and were used to collect experimental data. A psuedorandom ordering of the combinations of factors and levels were used to eliminate the effects of ordering in the data. During each session, the participant managed a team of 10 robots to perform the required tasks. Upon completion of each session, the participant reported NASA TLX workload ratings. Workload weightings were reported after all sessions had been completed. Immediately following the last session, each participant also completed an exit questionnaire to document a preference toward haptic feedback. Each participant was asked to rate their preference between 1 (not preferred) and 5 (preferred) on the use of haptic feedback while adjusting shape, adjusting spacing, traveling and overall preference.

5. RESULTS

This section details key results from the user study. A more comprehensive list of results and analysis can be found in [18]. A mixed models analysis of variance (ANOVA) with blocking on subject was used to determine significant effects on several measures of interest using a significance level of 0.05. Interactions between factors and levels were considered first, and then general effects were considered second if no interactions were significant. Combinations of factors are listed as "NE" for no haptics (No) in an easy world, "YE" for haptics (Yes) in an easy world, "NH" for no haptics in a hard world, and "YH" for haptics in a hard world.

5.1 Task Performance

There was a significant interaction between haptic feedback and world difficulty for percent hotspots found ($F(1, 18) = 4.94$, $p = 0.039$). There was no significant difference in interaction between haptic feedback and no haptic feedback in the easy world, but haptic feedback did improve performance overall (NE-YE mean = -4.805). This suggests that haptic feedback provides a relatively small advantage in finding hotspots when task demand is low.

When comparing haptic feedback with no haptic feedback in the hard world, the presence of haptic feedback significantly improved the operator's ability to find hotspots, as shown in Figure 9. Additional measures also showed that the operator was more effective at finding hotspots with shape manipulation than spacing manipulation. Shape manipulation was also most effective with haptic feedback. Results from the post-experiment questionnaire showed that on average, participants preferred the use of haptic feedback overall (rating of 3.7 on a scale from 1 to 5), which is consistent with statistical results.

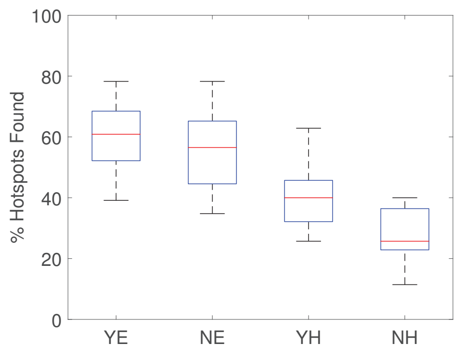

Figure 9: Boxplots for percent hotspots found, showing median values (red) and interquartile ranges (blue).

Measures such as reaction time and response accuracy were used to determine how well the operator responded to incoming messages that defined the secondary task. There were no significant effects due to haptic feedback or world difficulty on secondary task performance. This is not a surprising result because haptic feedback did not directly benefit the operator in responding to incoming messages. It was, however, hoped that there would be some significant difference overall when participants were provided haptic feedback.

5.2 Situation Awareness

There was a significant difference to SA accuracy score due to world difficulty, but not due to haptic feedback. When having the operator give responses to questions of equal difficulty between easy and hard worlds, there was a significantly higher score in the easy world than compared to the hard world. This suggests that the operator had greater task awareness in the easy world because those responses were more accurate, as shown in Figure 10.

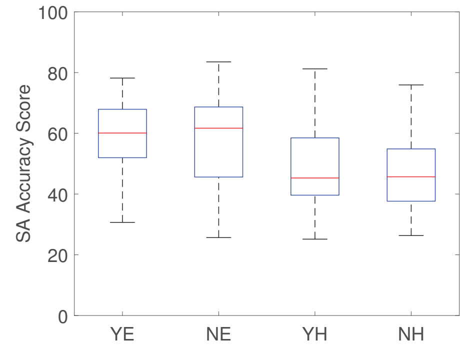

Figure 10: Boxplots for SA accuracy score, showing median values (red) and interquartile ranges (blue).

The SA accuracy score is averaged with the SA response time score to produce the total SA score. Based on these results, the SA accuracy score is useful in measuring key aspects of SA that are not covered by the SA response time measure. There is a slight trend that suggests that haptic feedback may increase SA (N-Y mean = -2.501), but is not significant enough to draw conclusions. These results do suggest, however, that haptic feedback does not have a negative effect on the operator's SA accuracy, which is an advantage of its incorporation with the presented interaction methods.

5.3 NASA TLX Workload

The NASA TLX workload measure returned with significant interaction between haptic feedback and world difficulty ($F(1, 18) = 5.19$, $p = 0.035$). Only one significant difference was found between haptics in the easy world and haptics in the hard world, which can be seen in Figure 11. This suggests that workload increases when managing the team in a harder world, which is to be expected due to a higher task demand. When task demand is kept low, as is the case in an easy world, the operator workload is lower and the presence of haptic feedback has little effect.

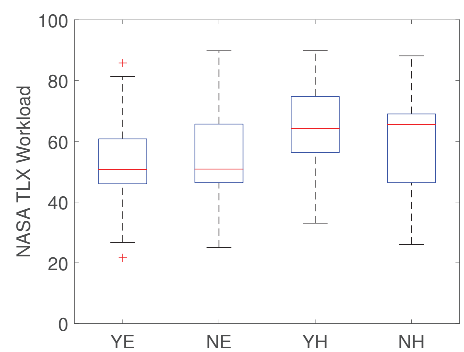

Figure 11: Boxplots for NASA TLX workload, showing median values (red), interquartile ranges (blue) and outliers (red).

There is a slight increase in workload when comparing haptics to no haptics in the hard world (NH-YH mean = -4.757). This difference, however, is not statistically significant and may even be caused by inconsistency of ratings between sessions or a misinterpretation of subscale definitions by participants.

6. DEMONSTRATION MINI-STUDY

Due to the significance of the results from the user study, a short experiment was also conducted to investigate the effects of visual feedback on operator performance with haptic feedback enabled. The aim was to investigate differences in task performance when the operator has obscured vision of the team. The experimental design is identical to the previous experiment, except only the team display is used and sensor readings are shown with team positions. Only 7 individuals were used in this experiment and no statistical analysis was performed.

Results show that with haptic feedback, the operator is able to manage the team and achieve close to the same level of performance without visual feedback of robot positions. As shown in Figure 12, participants performed better with visual feedback enabled. Without visual feedback, haptic feedback still provided enough information for participants to score within 2% in the easy world and 6% in the hard world on average. Workload scores only increased by 1.7% in the easy world and 5% in the hard world without visual feedback. These results complement those of the previous experiment in that haptic feedback does help to improve the operator's heads-up awareness.

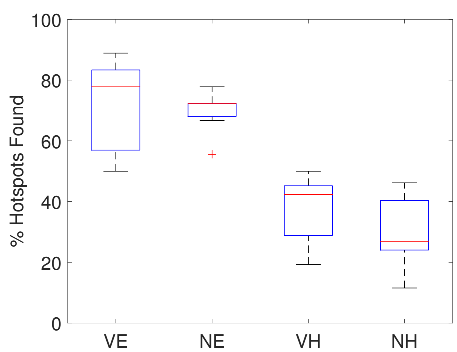

Figure 12: Boxplots for percent hotspots found, showing median values (red), interquartile ranges (blue) and a single outlier (red).

7. CONCLUSIONS

This work presents novel team interaction methods that are based on a modeling clay metaphor. The proposed methods allow a single operator to make intentional, global adjustments to the robot team by treating it as virtual clay. This was made possible by the high level of autonomy that the robots possess. The haptic channel was utilized to improve the virtual clay sensation as perceived by the operator and reduce the amount of information that needed to be visually processed.

A multimodal simulation system was developed as an experimental test bed to investigate the effects of haptic feedback on operator performance, SA and workload. Overall, the results show that haptic feedback significantly improves the operator's ability to find hotspots when task demand is higher. The operator may have used the haptic feedback to more quickly orient the haptic cursor relative to the team in preparation for the next manipulation. Since haptic feedback allows the operator to feel the shape instead of just see it, the operator could look at the sensor readings on the secondary display without even looking at the primary display for periods of time. Observations of the participants' performance throughout the experiment also validate these conclusions. It is clear that haptic feedback did improve performance, and since haptic feedback provides operators with team information without binding them to a visual display, this encourages a more heads-up approach. This improved performance shows that heads-up control of the team can be more easily achieved with haptic feedback than without it.

In the current setup, there were no significant benefits to using haptic feedback during spacing manipulation or travel mode. Haptic feedback likely made no difference in travel mode due to the simplicity of the travel objective. Ultimately, the results from the experiment show that under a higher task demand, haptic feedback made it easier for the operator to maintain heads-up control and achieve a greater level of performance as defined by this specific experimental simulation. Results from the mini-study also suggest that heads-up control can be achieved by means of haptic feedback even without visual feedback. Treating the team as virtual clay shows promise as an interaction model and may be suitable for many other human-swarm interaction scenarios. It is not presumed that this will always be the case, but it is hoped that these results will be useful in the development of future interfaces.

Although there were some improvements in SA score with haptic feedback, they were not significant. The only significant differences to SA score were due to world difficulty level. This does, however, suggest that the chosen SA measurement technique produces SA scores that reflect the expected trend for an operator's SA. The results also show a slight increase in operator workload with haptic feedback, using the NASA TLX measure, but the results were not statistically significant. Given the significant improvement in task performance, however, this trade-off may still be acceptable. With training, haptic feedback can be used to improve operator performance, especially when coupled with an intuitive interaction model such as the proposed virtual clay model.

Future work may include comparing the proposed interaction methods to existing methods to determine whether benefits of haptic feedback can be generalized to all types of interaction. Work could also be done to adjust interaction methods and feedback parameters to optimize user experience and provide more effective feedback. The modeling clay metaphor was used in this work to represent team shape in either surround or travel modes. In developing additional interaction methods, one might consider using this metaphor in other ways such as making topological changes or managing multiple teams at once.

8. ACKNOWLEDGMENTS

This work was funded by the United States Army Research Laboratory through contract W911NF-14-1-0633.

9. REFERENCES

[1] L. Alboul, J. Saez-Pons, and J. Penders. Mixed human-robot team navigation in the GUARDIANS project. In *Proceedings of the International Conference on Safety, Security, and Rescue Robotics*, Sendei, Japan, October 2008.

[2] C. K. Alder. A distributed control algorithm for small swarms in cordon and patrol. Master's thesis, Brigham Young University, 2016.

[3] C. K. Alder, S. J. McDonald, M. B. Colton, and M. A. Goodrich. Toward haptic-based management of small swarms in cordon and patrol. *2015 Swarm/Human Blended Intelligence Workshop (SHBI)*, pages 1–8, 2015.

[4] M. Barnes, F. Jentsch, J. Y. C. Chen, E. Haas, and K. Cosenzo. Soldier-robot teams: Six years of research. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 54(19):1493–1497, sep 2010.

[5] A. M. Brandt and M. B. Colton. Haptic collision avoidance for a remotely operated quadrotor UAV in indoor environments. *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, pages 2724–2731, 2010.

[6] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer, 2000.

[7] M. Diana, J. P. De La Croix, and M. Egerstedt. Deformable-medium affordances for interacting with multi-robot systems. *IEEE International Conference on Intelligent Robots and Systems*, pages 5252–5257, 2013.

[8] M. A. Goodrich. On maximizing fan-out: Towards controlling multiple unmanned vehicles. In M. Barnes and F. Jentsch, editors, *Human-Robot Interactions in Future Military Operations*. Ashgate Publishing, Surrey, England, 2010.

[9] M. A. Goodrich, B. S. Morse, D. Gerhardt, J. L. Cooper, M. Quigley, J. A. Adams, and C. Humphrey. Supporting wilderness search and rescue using a camera-equipped mini UAV. *Journal of Field Robotics*, 25(1-2):89–110, 2008.

[10] H. Jones and P. Hinds. Extreme work teams: Using SWAT teams as a model for coordinating distributed robots. In *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work*, CSCW '02, pages 372–381, New York, NY, USA, 2002. ACM.

[11] S. Jung and M. A. Goodrich. Shaping couzin-like torus swarms through coordinated mediation. In *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 1834–1839. IEEE, 2013.

[12] A. Kolling, P. Walker, S. Member, N. Chakraborty, K. Sycara, and M. Lewis. Human interaction with robot swarms : A survey. *IEEE Transactions on Human-Machine Systems*, 46(1):1–18, 2015.

[13] V. Kumar, D. Rus, and S. Singh. Robot and sensor networks for first responders. *Pervasive Computing, IEEE*, 3(4):24–33, Oct 2004.

[14] T. Lam, H. Boschloo, M. Mulder, M. van Paassen, and F. van der Helm. Effect of haptic feedback in a trajectory following task with an unmanned aerial vehicle. In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, volume 3, pages 2500–2506. IEEE, 2004.

[15] T. M. Lam, M. Mulder, and M. M. van Paassen. Haptic interface in UAV tele-operation using force-stiffness feedback. In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics (SMC)*, pages 835–840, 2009.

[16] M. Lewis. Human interaction with multiple remote robots. *Reviews of Human Factors and Ergonomics*, 9:131–174, 2013.

[17] P. Maxwell, A. A. Maciejewski, H. J. Siegel, J. Potter, and J. E. Smith. A mathematical model of robust military village searches for decision making purposes. In *IKE*, pages 311–316, 2009.

[18] S. J. McDonald. Haptic shape-based management of robot teams in cordon and patrol. Master's thesis, Brigham Young University, 2016.

[19] M. J. Micire. *Multi-Touch Interaction for Robot Command and Control*. PhD thesis, University of Massachusetts Lowell, 2010.

[20] R. R. Murphy and J. L. Burke. The safe human-robot ratio. In M. Barnes and F. Jentsch, editors, *Human-Robot Interaction in Future Military Operations*, chapter 3, pages 31–49. Ashgate Publishing, 2010.

[21] S. Nunnally, P. Walker, N. Chakraborty, M. Lewis, and K. Sycara. Using coverage for measuring the effect of haptic feedback in human robotic swarm interaction. *IEEE International Conference on Systems, Man, and Cybernetics*, pages 516–521, 2013.

[22] S. Nunnally, P. Walker, M. Lewis, N. Chakraborty, and K. Sycara. Using haptic feedback in human robotic swarms interaction. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 57(1):1047–1051, 2013.

[23] R. M. Philbrick and M. B. Colton. Effects of haptic and 3D audio feedback on operator performance and workload for quadrotor UAVs in indoor environments. *Journal of Robotics and Mechatronics*, 26(5):580–591, 2014.

[24] T. Setter, H. Kawashima, and M. Egerstedt. Team-level properties for haptic human-swarm interactions. *American Control Conference*, 2015.

[25] D. Sieber, S. Musić, and S. Hirche. Multi-robot manipulation controlled by a human with haptic feedback. *IEEE International Conference on Intelligent Robots and Systems*, 2015-December:2440–2446, 2015.

[26] J. M. Whetten, M. A. Goodrich, and Y. Guo. Beyond robot fan-out: Towards multi-operator supervisory control. In *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics*, Istanbul, 2010.

[27] C. D. Wickens. Multiple resources and performance prediction. *Theoretical Issues in Ergonomics Science*, 3(2):159–177, 2002.