

Extreme Work Teams: Using SWAT Teams As a Model for Coordinating Distributed Robots

Hank Jones

Aerospace Robotics Laboratory
Stanford University
Stanford, CA 94305 USA
+1 650 725 3297
hlj@arl.stanford.edu

Pamela Hinds

Management Science & Engineering
Stanford University
Stanford, CA 94305 USA
+1 650 723 3843
phinds@leland.stanford.edu

ABSTRACT

We present a field study of police SWAT teams for the purpose of enabling grounded design of a system to coordinate distributed field robots. The mission-oriented, spatially distributed SWAT environment provides a rich resource for robotics designers that mirrors field robot deployments in key ways. We highlight the processes with which SWAT team leaders create and maintain common ground among team members and coordinate action in these tightly-coupled, distributed teams. We present a system for coordinating distributed robots that we designed based on our SWAT team observations.

Keywords

Distributed work, distributed teams, leadership, extreme work teams, field robotics, human-robot interaction

INTRODUCTION

This paper describes a qualitative field study of leader behaviors in spatially distributed work teams, namely police Special Weapons and Tactics (SWAT) teams. The primary purpose of this research was to inform the design of a system used by people to coordinate the actions of remote teams of robots. This work also provides insights into the role of leaders in high-pressure, dangerous settings, the role of leaders in distributed work teams, and ways in which work teams compensate for distance when intense collaboration is required.

Extreme Work Teams

Over the last several years, an increasing amount of research in CSCW has focused on geographically distributed work. More is now known about how people coordinate activities with distant colleagues and how technology supports distant work [9, 12, 19]. However, many of these investigations are conducted in a laboratory or in office environments with highly skilled, minimally interdependent workers. Little research has examined

work teams that come together for a single event, are highly interdependent, and whose performance can save or cost lives. We call these intense teams operating under hazardous conditions “extreme work teams.” Although extreme work teams can be collocated or distributed, our interest is in extreme work teams that have the added challenge of coordinating work without a shared physical space.

Spatial distribution in extreme work teams is particularly important because of the challenges introduced for collaboration. Distance typically decreases the likelihood of communication and collaboration [1, 16] in part because distance and reliance on communication technologies makes it more difficult to develop and maintain *common ground* [4, 17, 20]. People cannot easily point to objects of interest, receive feedback on how their message was received, or repair misunderstandings [17]. Maintaining a shared awareness of the work environment and of ongoing activities is also particularly difficult for distributed team members [3, 5]. For example, Cramton [5] reported that team members had difficulty accurately interpreting events at distant sites because they did not have critical contextual information. Lack of contextual information, reduced communication, and less common ground have the potential to severely limit the ability of extreme work teams to coordinate interdependent activities.

One way that distributed teams can successfully overcome the lack of proximity and impoverished communications is to strive for more loosely-coupled systems [20, 21]. Grinter et al [11] observe some benefits of increasing the modularity of the work so that constant interaction and information sharing between distant team members is not required. However, some geographically distributed work does not lend itself to loose coupling [11]. SWAT teams are an example in which there is strong, unavoidable interdependence. Weick’s [24] study of firefighters in the Mann Gulch disaster illustrates the importance of coordination on these teams and the high cost of failure (13 lives lost) when coordination breaks down. Because loose-coupling is not a viable option for

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW’02, November 16–20, 2002, New Orleans, Louisiana, USA.

Copyright 2002 ACM 1-58113-560-2/02/0011...\$5.00.

these teams and the cost of errors is high, they must develop ways to coordinate effectively at a distance.

In complex systems, particularly those susceptible to disaster, constant communication is crucial for coordination and sensemaking [7, 24]. However, because communication is more difficult on distributed teams, the process of sensemaking can be more challenging [5]. Although little research has been conducted on the role of the leader in distributed teams, we suggest that the leader can play a crucial role in ensuring that team members have similar information from which to make sense of their shared reality. This is consistent with Weisband's [25] findings that leaders of more successful distributed teams actively shared "other awareness" information – information about what others in the team are doing. We contend that the leader's role in sharing awareness information will be especially important in extreme work teams because there is not time for all members to communicate and receive feedback from one another and because errors and misinterpretations can, quite literally, be fatal.

Field Robots

Field robots, sometimes called telerobots, are robots that extend a person's sensing and/or manipulating capability to a location remote from that person [23]. Field robots usually have enough autonomous capability to control their own movement and act on high-level commands from their users. These robots are important because they often are called upon to do things that humans cannot, will not, or should not do. Humans cannot yet go to Mars, they will not willingly traverse minefields, and they get bored flying over an area taking photographs for days on end. Consequently, field robots are performing these roles with humans as supervisors from afar.

Although currently deployed field robots operate alone, system designers have argued that multiple robots will be required for many future operations including search and rescue [13], space exploration [6], and military logistics [26]. For example, in a search and rescue operation following an earthquake, multiple robots may be working alongside and supporting (human) members of emergency response teams, but operated by users some distance from the site. Current research has focused on the difficult problems of navigation, self-monitoring, and autonomous control, but not on issues that controlling multiple-robot systems will create for their users and for humans working alongside robots in the field. Many of the anticipated problems in coordination derive from the fact that a field robot and its user do not operate 'within sight' of one another. Thus, the human is part of a distant collaboration with a computer acting as the only conduit between the operator, the robots, and other people or objects in the environment.

Prior to developing collaborative systems to coordinate teams of robots, we felt it was important to pay close

attention to the collaborative aspects of the type of work the technology was intended to support. We recognized that future robot operators will be leading teams of robots that are highly interdependent and that operate in hazardous settings, e.g. extreme work teams.

STUDY SETTING

To understand leader behaviors in extreme, distributed work environments, we observed field operations of police Special Weapons and Tactics (SWAT) teams. Such workplace studies are important in CSCW because they increase the likelihood that systems will support actual work practices.

SWAT Teams

Police Special Weapons and Tactics (SWAT) teams are responsible for handling high-risk tactical situations involving barricaded suspects, hostage situations, and the serving of warrants. The primary goal of SWAT, as stated in a SWAT training manual, is the "successful conclusion of high-risk situations through the use of specially equipped and highly trained personnel without injury or loss of life to citizens, suspects, or police officers." The chief components of SWAT teams are assault/entry teams, a Hostage Negotiation Team (HNT), and sniper teams.

We observed field operation exercises of police SWAT teams. We choose to observe field exercises rather than actual SWAT missions for several reasons. First, great pains are taken in the SWAT exercises to replicate the events and activities likely to be encountered in an actual mission. Participants in the exercises are active-duty officers who perform the roles they would in an actual mission and whose lives depend on their readiness for these missions. Second, due to the high level of physical danger in most SWAT missions, it is unlikely that the researchers would have been allowed access, much less the free movement that the exercises provided.

We obtained permission from the leaders of the local SWAT team to observe the field training exercises. We were given vests to signal that we were observers in the exercise and thus were able to move freely through the incident area. This provided an opportunity to observe the evolving state of the exercise as well as view the events from various participants' perspectives. We were introduced as university researchers on-site to observe how SWAT teams functioned and not to analyze the performance of the team.

At an incident site, the command structure consists of an incident commander, the Tactical Commander (TC), the Hostage Negotiator Team leader, and the logistics supervisor. These individuals establish an Incident Command Post (ICP) at some distance from the suspects' location, known as the objective site. The TC was the focus of our research. The TC's role is to coordinate the activities of the tactical teams – assault, emergency response, sniper, and perimeter control – throughout the

exercise. Typically, there is significant physical distance between the TC (at the Incident Command Post) and the tactical teams (at the objective site), and the TC does not have direct sight of the team or even the site itself. The TC relies on radio (voice only) technology to gather information from and convey commands to the tactical teams.

For SWAT field exercises, the primary objective is to execute realistic but controlled scenarios. The nature and location of the incident to be enacted is unknown to most of the team. To avoid causing unnecessary alarm, the exercises are typically held in relatively remote commercial locations on weekend days, although residential areas also are used if sufficient notice can be given. The team we observed had an average of forty members present at the exercises, with around fifteen in the Incident Command Post and the rest deployed to the objective site. The pace of operations is typically measured, although the pace can accelerate greatly.

Our primary method for data collection was the observation of four field exercises. We paid particular attention to the behaviors and communications of the TC, although we also conducted some observations at the objective site. This enabled us to compare the actual situation with the TC's perception of the situation. Throughout the exercise, we used sporadic and informal interviewing as necessary to clarify salient points or fill in background details. We took extensive field notes of the TC's behaviors throughout the exercise, as well as some ancillary notes on the behaviors and conversations between members of the tactical teams, and significant events in the exercise. Where possible, we captured artifacts either by drawing them or photographing them at the scene. We also made audio recordings of the entire exercise from the Incident Command Post.

To familiarize ourselves with police and SWAT procedures before attending the exercises, we participated in two regular patrol 'ride-alongs' with SWAT team members. Through informal discussion, we discovered the functions of the SWAT team, terminology, and other background information. The ride-alongs also provided time to convey the purpose of the study and to build the trust necessary to gain access to the SWAT exercises.

The Field Exercises

We observed four exercises led by three different Tactical Commanders. In the first exercise, the scenario started at a residence and then moved to a hospital clinic. The SWAT team was called to a residential neighborhood where a hostage situation was underway. Upon arrival, an Incident Command Post location was determined, neighbors described floor plans of the residence to the assault team leaders, and uniformed officers cordoned off the area. However, just before the assault was to begin, the administration revealed that the hostage and suspect had left the house before the SWAT team had arrived –

according to the participants, an annoying but authentic scenario indicative of the realism of the exercises. The hostage had been taken to a hospital for treatment, and the suspect had taken additional hostages at the hospital clinic. The SWAT team moved immediately to the hospital site and redeployed the units.

A workplace hostage incident was the scenario for the second exercise, held at a school supply distribution warehouse. The large, cluttered environment of the warehouse limited the usefulness of hand signals and consequently increased the assault team dependence on the TC as the central node for coordination.

A workplace hostage situation was again the scenario for the third and fourth exercises, held at an abandoned office building.

LEADING EXTREME TEAMS

Across the four exercises, we observed two key roles of the TC – cultivating common ground and coordinating the actions of the distributed team.

Cultivating Common Ground

The leaders of the SWAT teams we observed cultivated common ground that they subsequently relied on to coordinate the distributed work of their teams. Kraut et al [17], Olson & Olson [20], and others have emphasized the importance of common ground for many aspects of distributed work. In every exercise we observed, the SWAT teams cultivated common ground through training, meeting face-to-face before dispersing to the objective site, and re-calibrating during the incident by establishing relative locations of common objects in the environment. The face-to-face meeting at the start of the crisis and the re-calibration process were driven by the TC.

Establishing common ground through training

One way to establish common ground, particularly on distributed teams, is to provide common training [19]. Although we did not observe SWAT training (other than the exercises), it was clear that training in standard vocabulary, communication procedure, and a set of standard procedures for discreet actions (e.g. how to respond when a hostage is released) were common across members of the teams. In the exercises we observed, the TC frequently mobilized action by invoking standard procedures in which members of the teams had been trained.

Calibrating common ground

The SWAT training recognizes that the spectrum of possible scenarios is so broad that the team members can only be given basic building blocks for action, and then they must be adaptable to a wide variety of situations. When a call for the SWAT team goes out, the officers have no idea what aspects of their training will be needed, or what actions will be necessary. In all four exercises, the TC's dealt with this uncertainty by meeting with the

teams face-to-face before dispersing to the objective site. They referred to this meeting as the Situation Report. A photograph of the Situation Report that took place during exercise 3 is captured in Figure 1, with the TC in the center of the group [no headgear, facing camera]. From what we observed, the TC's consistently used these meetings to provide their tactical teams with the known information and reinforced specific elements of their training. The first part of the briefing was always a rundown of the situation specifics – the number and description of suspects, reports of weapons, descriptions of the area and buildings, and any constraints on their operation. The TC's then specified the roles that each officer would perform during the action. In particular, the TC always specified the sniper teams, the scouts, and the Emergency Response Team (ERT). He then dispatched the teams out to the objective site according to the practiced standard procedures for the initial phase of an incident. Any deviations from standard procedures, such as team size or off-limits areas, were explained to each element as it was sent out.

As a means of understanding the situation, the SWAT command structure obtained street maps and building floor plans from city documents and landlord records. The TC's used these documents as exhibits during the Situation Report, consistently referencing them with phrases such as “the door here” and “clear through here.”

In exercise 3, for example, the TC had maps of the building that he lifted high so that the entire team could see. While using his free hand to reference the map, he described the situation and provided the immediate plan:

ERT (Emergency Response Team) is moving into the breach point – the front door. The suspect and hostages are reported to be in the Hercules Room, just off the stairwell on the 2nd floor. Make a rapid clear through the bottom floor but on the top floor, move carefully. We're gonna have to leave people (officers) here (staircase) just in case people (hostages or suspects) come down it.

Across all four exercises, the TC numbered the sides of the building in which the teams would be operating. SWAT teams use numbers to refer to the sides of buildings that are then used for position determination inside and outside of the building. The side with the front door is Side 1 and the numbering proceeds clockwise with Side 2 on the left, Side 3 opposite the front, and Side 4 on the right. In each case, the TC used a map to make the side numbers explicit, saying for instance in exercise 4, “This is going to be the #4 side here,” and so forth. However, in all of the exercises there was some element of the architecture of the building that created confusion.

Re-calibrating common ground

The most challenging component of cultivating common ground took place after the Situation Report ended and

the team moved into their spatially distributed positions. We noticed that team members made attempts to coordinate with each other by referring to common objects. However, because they had different perspectives on these objects after dispersing, there was some confusion. The leader played an important role in reconciling team members' different perspectives.



Figure 1 – Situation Report

The SWAT participants consistently used physical objects and the location of those objects as the foundation for their dialogues, as in the following example from exercise 2, when the TC needed to ascertain the status of doors leading into the warehouse:

TC: B sees two garage doors, and thinks the right one is open but can't be sure because of the car. Can you confirm?

Officer W: I see there are two doors but they look like one from here. Yes, the right door is open just a little.

and when the TC provided an updated plan to the assault team in exercise 3:

TC: The rest of the team is going to move to that stairwell. OK, and we're gonna go up and take the landing and stage right there on top of it. That's the stairwell on the 3 side. We're gonna go right up there onto the 2nd floor and just clear out from there. We don't know where the suspect is.

The use of objects often went further than just providing the relevant information for the status of one member, becoming a mechanism for determining the spatial relationship between team members, as in this example in exercise 2:

TC: Where are you in the ditch right now? Do you see B near the fence? How far are you from him?

Officer G: I don't see B but I see W near the fence. I am 30 yards from W, in the ditch.

TC: (Knows that B and W are together) OK.

To determine their spatial relationship relative to other team members, members of the tactical teams frequently contacted the TC, who typically had the most complete and comprehensive understanding of the incident. However, we noticed a significant burden that the TC was thus required to shoulder throughout the exercise – he needed to keep track of the perspective and knowledge base of each agent under his command. For example, during exercise 3 a sniper reported:

White male with handgun, 2nd floor, opening 10

which the TC translated without pause into a pertinent description in the frame of reference of the assault team leader:

(Sniper) 01 has got visual on one of the suspects on Side 2 in window above your breach point with at least one gun.

The TC was not only interested in the physical location of the units but also in their capabilities for action. For instance, in this dialogue from exercise 3, the TC worked with the leader of the assault team to get a landline phone to the suspect:

TC: (The suspects' cell) Phones are going dead. Can you spare me two (officers) to come back and deliver the throw phone?

Team Leader: Affirmative.

TC: Alright, send two (officers) to come back and pick up the phone.

A few minutes later, after the officers had returned to their team with the phone, the TC asked the leader:

TC: Are you ready to deliver the phone?

Team Leader: We're ready to deliver the phone.

There was then a delay of several minutes before the TC gave the order to deliver the phone as he waited for agreement from the Hostage Negotiator Team. Through many such dialogues with the distributed teams, the TC was able to add to his mental model of the incident and, in particular, understand the effect that he might be able to have on the situation by commanding action from his teams. The moment-to-moment capabilities of the teams were constrained by the myriad of variables in the environment that were not observable by the TC. Thin walls, open doors, and excited suspects are all examples of small details that effected the options available to a team. By asking the teams directly for reports of their abilities, the TC could be assured that his commands would be valid and relevant for the situation of each team.

The difficulty of maintaining common models of the environment was made evident through one incident during the third exercise when the command team and the assault team were confused about which staircase they were supposed to climb. The building was symmetrical, with stairs on each end. The plan was to use a hostage

exchange to attract all the suspects to one staircase and send the SWAT assault team up the other staircase. However, the diversion almost backfired when the assault team unknowingly went to the wrong staircase. They were surprised to see a hostage in the stairwell just as they began to move. They informed the TC that they had a hostage in custody in the staircase, and in a few seconds the TC realized the source of the confusion and had the team immediately move away from the staircase to avoid the potential conflict.

We observed the SWAT team dispatchers, who are responsible for gathering and disseminating information, making maps of building interiors and exteriors and labeling the unit locations in the first thirty minutes. But, these maps were not updated. Contrary to TC statements in interviews, we observed that the TC's used the maps infrequently and preferred to develop a comprehensive mental picture of the incident. This might be due to the difficulty in representing all the knowledge effectively on paper or reluctance to rely on a dispatcher's interpretation of events. In the field, the assault teams had paper copies of floor plans and street maps, but otherwise they too maintained their model mentally.

Throughout each exercise, the TC moved around the ICP area, monitoring the radio chatter, talking to the Hostage Negotiator Team, getting reports from the various ICP dispatchers as they became available, and looking at the maps and logs. We periodically asked the TC explicitly about the state of the incident and he always had at least a rough but accurate idea of his units' locations.

Using Common Ground to Coordinate Action

Members of the SWAT teams cultivated common ground throughout the exercises. At the same time, we observed that they used this common ground to initialize coordinated action. In this paper, we focus on the role of the TC in coordinating action within the team, but note that coordination of the team's actions with other entities (e.g. negotiators or paramedics) was an extremely important function of the TC and is an equally vital role for robot operators.

In the next sections, we highlight two of the ways in which the TC's relied on common ground to coordinate action.

Referencing standard procedure

Because members of the SWAT team had a base of common ground established through their training, the TC usually assumed that a set of basic SWAT capabilities were available on request, such as 'clear-outs', dynamic or stealth movements, or holding hostages or suspects. For example, in each exercise the snipers asked for rules of engagement as suspects appeared in their sights. The reply from the TC was a simple statement of "standard rules of engagement" which conveyed an extensive set of behaviors and commands that were immediately invoked with seamless precision.

Using the addressee's reference frame

The other method the TC used for coordinating action was to give commands from the addressee's frame of reference. For the most part, the dialogues were conducted relative to unit position. The TC smoothly alternated between the points of view of the team members under his command. In a typical example, in exercise 4, the TC first told some team members who had a hostage in custody, "Have her (hostage) go out those doors that are right by you" and then immediately told the rest of the team, in a different location, "You are clear to go through the first part (of the floor clear-out procedure) to top of the (next) stairs."

The TC also conveyed the location of the teams relative to one another as well as to other objects. For example, the TC coordinated between the Emergency Response Team (EMT) leader and the assault team by conveying each of their locations. When the EMT leader asked if they could move inside the front door (where presumably they could have a more active role in watching for activity in the building), the TC replied: "I don't have a problem with that. Just remember that the assault team will be moving from your left. Stand by before you move. I'll have assault team stand by in lobby by the double doors."

As with the cultivation of common ground, the TC used objects, referred to in the units' local frames of reference, to coordinate action as shown in this example from exercise 2:

TC: W, do you have a visual on the suspect?

Officer W: No, (there is a) large stack of boxes between me and location (where I hear what) I believe is the suspect.

TC: B, do you see a stack of boxes to your left in the direction of W?

Officer B: Affirmative.

TC: B, do you see a location for W to egress to that remains in cover?

Officer B: Yes, there is a desk with a computer immediately to his left when he comes around the stack that he should be able to get to.

TC: Did you get that W?

Officer W: Affirmative, moving to the desk.

The following example illustrates how the TC could provide wider context for the units when it was needed, but still maintain the use of local referencing.

TC: G, do you see a door in the East wall below the suspect?

Officer G: Affirmative.

TC: We need to get someone in that room to make sure the suspect does not move further into the building.

Officer G: W is in position to move into that door. I can cover the door while he moves in. B can also move to that door if you want him to - he'll just take a second to get into position.

TC: Signal B to move into position to move through the door under the suspect.

The use of relative positions was even used by the unit leaders when talking to the TC, who should not necessarily be expected to be familiar with their frame of reference. In this example, the Assault Team Leader (ATL) requests information from the TC to determine whether some noises being heard were friendly or not:

ATL: We're at the top of the stairs. We hear something over on the right side, are we clear to move through there?

TC: Yes, you're cleared to move through.

These translations by the TC between the various perspectives among the teams and team members were frequent, and relying on the TC as the locus of coordination generally was reliable. However, the reliance on the TC that developed also proved to be a weakness. For example, at one point in exercise 1, the assault team asked for a floor plan for the area in which the suspect and hostages were located. The TC had unintentionally been given the drawings for a different floor than the one desired. An assault plan was devised based on this incorrect floor plan to enter on opposite sides of the building and proceed in two teams around two hallways that passed along the walls of the building. Upon commencing the assault, however, the assault teams reported back that there were no outside hallways, but they could see one another down one central hallway going through the center of the building. This error caused an unexpected, potentially lethal crossfire situation that great pains are usually taken to avoid.

Summary of Findings

In our observations of SWAT field exercises, we found that the cultivation of common ground was an ongoing and critical activity for these extreme teams. One of the most important and challenging activities for the leader of these teams was ensuring that team members re-calibrated when the situation or environment changed and team members no longer shared common ground. This was particularly dangerous when common ground had slipped away unnoticed, leaving team members assuming that they had a shared understanding of the situation and environment. We also observed that the TC coordinated action using common ground as the global frame of reference. While the TC developed awareness of the entire situation by integrating information for all sources, he could decompose that picture to issue commands that used the addressee's frame of reference, and share awareness information about objects and people from the addressee's perspective.

AN INTERFACE TO COORDINATE FIELD ROBOTS

The findings from our observational study of SWAT team leaders has led to a novel interaction design for the operation of multiple robots by a single operator that has been successfully used to command multiple heterogeneous robots [14]. Two of its basic innovations are described here: an object-centered electronic dialogue and the use of distributed local models. From our observations of the SWAT teams, we learned that establishing common ground was a critical foundation for coordinating action. We therefore created a Correspondence Agent that assists the operator by continually merging information from the robots about locations and objects. We also created an electronic dialog that determines the abilities of each robot with respect to the physical objects it senses. When used with the Correspondence Agent, the electronic dialog can be used to coordinate the action of the robots, but without it more errors are made in attempting to accomplish cooperative tasks.

Experimental system

We carried out an experimental implementation of the interaction design using free-floating space robots (FFSRs) and micro autonomous robots (MARs). These robots provided an opportunity to execute the design in a controlled environment that is modeled after the use of robots for Space Station maintenance tasks. The left side of Figure 2 shows the Graphical User Interface (GUI) that is presented to the user. The objects on the screen are manipulated by the user to coordinate robot activity. The large boxy entities are the free-flying robots and the MARs are the small white cylinders. The other entities are mobile or immobile objects. The right side of Figure 2 is a camera snapshot of the environment being modeled and the robots that are ultimately being controlled.

Dialogue content

Interaction between the TC and the SWAT team members was necessarily conducted via voice dialogues, and we sought to determine whether an interaction that utilized a dialogue was possible for field robot operation. Dialogue with a robot is an old topic in artificial intelligence, going back to Nilsson's classic SHAKEY robot in the 1960's. Much more recently, robotic researchers have developed many systems for *collocated* human-robot interaction using natural language interfaces, sometimes supported by a GUI [e.g., 10]. There have been two projects to use some form of dialogue with a field robot – Fong [8] used dialogues to collaborate on object classification, and Zelek [27] directed a robot on navigational tasks. However, dialogues have not been used with systems of multiple robots, and no interactions have integrated findings from observations in non-robotic environments of how humans actually use dialogue to coordinate actions.

Based on our observations of the SWAT TC's dialogues with the tactical teams, we designed the dialogues in our human-robot interaction to be centered on physical objects in the world. The goal of our dialogue is to create valid and relevant commands of the form Subject-Verb-Direct Object, such as "Robot A grab Component 1." This is an extension of a well-known method for the operation of single robots based on a conceptual command of "Put that there." The limitation on possible actions in the "Put that there" method – one can only ask the robot to move an object – is both positive and negative, since it means that the operator has little confusion over what actions are possible by the robot, yet is given no choice in the matter. Our method provides the possibility of many more verbs than simply 'to put' and replaces the location focus with a physical object focus, creating a new paradigm for operation that is best described as "Do what to whom?"




Figure 2 – Graphical user interface and robot hardware


The electronic dialogue starts when the robot generates a list of the types and positions of objects it senses. The operator then selects, by clicking, a particular object with which to work. The interface requests a list of tasks that the robot can enact on that object, and this list, when returned, is displayed in a dialog box on the interface. The operator selects a task and the command – valid and relevant – is sent. In effect, the operator is assured of a consistent set of affordances for solving mission problems and accomplishing tasks. An example of this interaction is shown in Figure 3, with a screen shot of the interface on the left and a representation of the electronic dialogue on the right.

Local model maintenance

In our observations of SWAT teams, we learned that the TC's constructed a mental model of the current situation and environment from the initial information known and the subsequent updates from the teams, providing many different perspectives that had to be integrated to form an overall world view. In these SWAT situations, such a process is the only way to construct such a global model, as no external sensor, such as a camera or a position reporting system, are available.

In contrast, existing multiple robot systems currently rely on access to a global model of the world to allow the

robots to interact, cooperate, avoid conflicts, and otherwise function harmoniously. To accomplish this, a remote sensor is used to determine the position of the robots, objects, obstacles, and the structure of the environment [e.g. 2, 22]. However, for field robot operation, such a sensor is not likely to be available. In addition, our SWAT observations and our previous robot development [15] convinced us that we needed a system that integrated different frames of reference but maintained the perspective of each.



Robot: I see three objects - purple, red, and black.
 User: What can you do to the purple object?
 Robot: I can escort, watch, or locate it.
 User: Please escort it.

Figure 3 – Example dialogue interaction

For example, the GUI keeps track separately of each robot’s sensory perception of nearby objects, and combines them for display. This provides a more realistic system, but not without cost. As shown in Figure 4, two robots will typically sense their environment slightly differently. (To provide a clear example, the FFSR and the MAR are each shown only with the cone each senses and not the other robot.)

This presents a problem for conducting a dialogue involving both robots and the cone, since there are, conceptually, two cones in the world – one sensed by the FFSR and one sensed by the MAR. We observed the same problem during our SWAT observations, when members of the distributed team referred to the same objects (e.g. the staircase in exercise 3) from different perspectives. The TC resolved this issue by maintaining separate models that represented the perspective of each team member.

We addressed this problem for the robots by introducing an autonomous software agent we call the Correspondence Agent (CA). The CA listens to status reports from the robots about objects that are sensed and makes determinations of which objects correspond to which others – basically forming ties between the worlds based on the objects being sensed, just like the SWAT TC. This agent was automated because of the likely scale of the task, which for even a small system of eight robots and five objects can require the analysis of over 390,000 (# objects to the power of the # of robots) possible correspondence pairs. Nonetheless, there are instances where this automation links objects in error, or does not

link objects that the operator recognizes as representing the same object. We consequently included a mechanism in the CA to manually determine that objects be linked or unlinked. A more comprehensive description of the CA and its design can be found in Jones and Rock [14].




Figure 4 – Different perspectives

The effect of the Correspondence Agent is shown in Figure 5. On the left, with the CA turned off, two cones are presented to the operator. When the CA is enabled, the interface is informed that those two objects are actually the same, and the interface presents only one. The operator can now select the two robots and give them commands for cooperative tasks on the cone. Most importantly, however, the interface sends out commands in each robot’s own frame of reference using the cone that it senses, not the unified cone shown to the operator. Again, this is the reason for using distributed models and a Correspondence Agent – to be able to provide a straightforward interface to the operator, but maintain an infrastructure that reflects the realities of the robotic system.

One should note that this interaction does not require a 3-D GUI like the one we have developed – robot operation based on the same electronic dialogue could be conducted strictly via a text-based cell phone browser or through a speech-based interface.




Figure 5 -- Demonstration of Correspondence Agent

DISCUSSION

The purpose of our study was to inform the design of a system to coordinate teams of remote robots expected to perform in extreme settings. We took new observations from existing human-human interaction and applied them to future human-robot interaction. Although few work settings currently boast robot supported cooperative work

with autonomous robots, this eventuality is not far away. Autonomous robots are currently deployed in hospitals (<http://www.pyxis.com/products/newhelpmate.asp>) and are being developed for a wide range of other applications (<http://www.service-robots.org>). We therefore believe that this is an important future direction for research in CSCW.

This study also adds to the minority of CSCW papers that derive implications for design from field studies and impact the architecture of the system, not just its interface or user-level functionality. Such work is particularly important when designing robots that will work with and be operated by humans because of the time and expense involved in designing and building robots. In addition, there are currently few situated users. To increase the likelihood of building robots that support work practices and are effectively operated by humans, we argue for the importance of observing likely environments in which robots will be implemented.

Our observations of leaders of SWAT teams also provide more general knowledge about leading distributed work teams. In this study, we found that strong leadership in a face-to-face meeting at the outset of the activity was crucial for establishing common ground. In the meetings we observed, leaders focused on making sure that team members had a shared understanding of the situation and understood others' roles. This is consistent with Armstrong and Cole's [3] finding that distributed teams were more effective if they met face-to-face to discuss how they planned to approach the task and the roles team members would play. Along with Weisband [25] we found that leaders are particularly important in establishing common ground during these early phases of a project. Thus, we argue that leaders of distributed teams may be more effective if they assemble new teams early on (using the richest possible medium) and build a foundation of common ground that can be leveraged later in the project.

We also observed that a vital role of leaders on extreme teams was to share awareness information about the activities and positions of other team members. It was inefficient for team members to share this information because they did not know who would need it and they did not know how to convey the information from the perspective of the recipient. Leaders who are in contact with all team members and have a global picture of the team's status are in the unique position to provide this translation function and ensure that team members understand their colleagues' situations. Weisband [25] observed that distributed teams were more effective when their leaders shared awareness information about others on the team. Our results support and extend this work, suggesting that leaders of distributed teams will be more effective if they convey information from the perspective of the recipient. We observed that sharing information

based on the recipient's perspective aided significantly in teams being able to make sense of the situation and their role in it.

Limitations & Future Work

The work we presented here has a number of limitations and there are numerous opportunities to replicate and extend this work. We highlight a few. First, to provide an environment appropriate to the eventual role of field robots, we chose to study SWAT teams. SWAT teams are extreme teams and therefore unlike many other teams in several ways – their work is dangerous, they are extremely and unavoidably interdependent, and they are temporary. Although these factors make SWAT teams ideal for understanding how we might coordinate field robots and informs our understanding of teams in extreme settings, these factors limit the generalizability of these results for more traditional teams.

Second, we used the results of our qualitative study of SWAT teams to generate principles for a new system intended to coordinate teams of field robots. Future research is needed in the form of usability studies to assess the impact of our design and determine the extent to which it is an improvement over existing single or multiple robot systems. Such studies would be a substantial contribution because there are currently few, if any, usability studies performed with field robots.

ACKNOWLEDGMENTS

The authors sincerely thank the Palo Alto – Mountain View (California) Special Weapons and Tactics team, particularly Lt. Dennis Burns of the Palo Alto Police Department and Agent Russell Barcelona of the Mountain View Police Department, for their cooperation and assistance. This research was supported by funding provided by NASA grant NCC2-1154 to the first author and by NSF grant #ITR/PE-0121426 to the second author.

REFERENCES

1. Allen, T. 1977. *Managing the Flow of Technology*. MIT Press, Cambridge MA, 1977.
2. Arkin, R., & Balch, T. AuRA: Principles and Practice in Review. *Journal of Experimental and Theoretical Artificial Intelligence*, 9 (August 1997), 175-189.
3. Armstrong, D.J., & Cole, P. Managing distances and differences in geographically distributed work teams, in S. E. Jackson & M. N. Ruderman (eds.), *Diversity in work teams: Research paradigms for a changing workplace*. American Psychological Association, Washington DC, 1995, 187-216.
4. Clark, H.H., and Wilkes-Gibbs, D. Referring as a collaborative process. *Cognition*, 22 (1986), 1-39.
5. Cramton, C.D. The mutual knowledge problem and its consequences for dispersed collaboration. *Organization Science*, 12, 2001, 346-371.

6. Doggett, W. Robotic Assembly of Truss Structures for Space Systems and Future Research Plans. *2002 IEEE Aerospace Conference Proceedings* (Big Sky MT, March 2002), IEEE Press.
7. Eisenhardt, K. High reliability organizations meet high velocity environments: Common dilemmas in nuclear power plants, aircraft carriers, and microcomputer firms, in Karlene H. Roberts (ed.) *New Challenges to Understanding Organizations*. Macmillan, New York NY, 1993, 117-135.
8. Fong, C., Thorpe, C., & Baur, C. Collaboration, Dialogue, and Human-Robot Interaction, in *Proceedings of the International Symposium of Robotics Research 2001* (Victoria Australia, November 2001), Springer-Verlag.
9. Fussell, S, Kraut, R., & Siegel, J. Coordination of communication: Effects of shared visual context on collaborative work, in *Proceedings of the Conference on Computer-Supported Cooperative Work* (Philadelphia PA, December 2000), ACM Press, 21-30.
10. Green, A., Huttenrauch, H., Norman, M., Oestreicher, L., Eklundh, K. User Centered Design for Intelligent Service Robots, in *Proceedings of the IEEE International Workshop on Robot and Human Interactive Communication* (Osaka Japan, September 2000), IEEE Press, 161-166.
11. Grinter, R.E., Herbsleb, J.D., & Perry, D.E. The geography of coordination: Dealing with distance in R&D work, in *Proceedings of SIGGROUP* (Phoenix AZ, November 1999), ACM Press, 306-315.
12. Herbsleb, J., Mockus, A, Finholt, T, & Grinter, R. 2000. Distance, dependencies, and delay in a global collaboration, in *Proceedings of the Conference on Computer-Supported Cooperative Work* (Philadelphia PA, December 2000), ACM Press, 319-328.
13. Hornbuckle, J. Unmanned Aerial Vehicles in the U.S. Coast Guard. *Proceedings of AUVS '93* (Washington DC, June 1993), AUVSI.
14. Jones, H., and Rock, S. Dialogue-Based Human-Robot Interaction for Space Construction Teams. *2002 IEEE Aerospace Conference Proceedings* (Big Sky MT, March 2002), IEEE Press.
15. Jones, H., and Rock, S. Human-Robot Interaction for Field Operation of an Autonomous Helicopter. *SPIE Conference on Mobile Robots XIII* (Boston MA, November 1998), SPIE 3525, 244-251.
16. Kraut, R., Egido, C, & Galegher, J. Patterns of contact and communication in scientific research collaborations, in J. Galegher, R. Kraut, and C. Egido (eds.) *Intellectual Teamwork*. Lawrence Erlbaum, Hillsdale NJ, 1990, 149-172.
17. Kraut, R.E., Fussell, S. R., Brennan, S. E., & Siegel, J. Understanding the effects of proximity on collaboration: Implications for technologies to support remote collaborative work, in P. J. Hinds & S. Kiesler (eds.), *Distributed work*. MIT Press, Cambridge MA, in press.
18. Liang, D., Moreland, R., & Argote, L. Group versus individual training and performance: The mediating role of transactive memory. *Personality and Social Psychology Bulletin*, 21, 1995, 384-393.
19. Mark, G. Merging multiple perspectives in groupware use: Intra- and intergroup conventions in *Proceedings of SIGGROUP* (Phoenix AZ, November 1997), ACM Press, 19-28.
20. Olson, G., & Olson, J. Distance matters. *Human Computer Interaction*, 15, 2000, 139-179.
21. Olson, J.S., & Teasley, S. 1996. Groupware in the wild: Lessons learned from a year of virtual collocation, in *Proceedings of the Conference on Computer-Supported Cooperative Work* (Boston MA, November 1996), ACM Press, 419-427.
22. Parker, L. Multi-Robot Team Design for Real-World Applications, in T. Fukuda, T. Arai, & I. Endo (eds.), *Distributed Autonomous Robotic Systems 2*. Springer-Verlag, Tokyo, 1996, 91-102.
23. Sheridan, T. *Telerobotics, Automation, and Human Supervisory Control*. MIT Press, Cambridge MA, 1992, 4.
24. Weick, K. The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly*, 39, 1993, 628-652.
25. Weisband, S. Maintaining awareness in distributed team collaboration: Implications for leadership and performance, in P.J. Hinds & S. Kiesler (eds.), *Distributed work*. MIT Press, Cambridge MA, in press.
26. Wilson, J. Unmanned helicopters begin to deliver. *Aerospace America* (June 1999), 38-42.
27. Zelek, J. Human-Robot Interaction with a Minimal Spanning Natural Language Template for Autonomous and Teleoperated Control, in *Proceedings of the International Conference on Intelligent Robots and Systems* (Grenoble France, September 1997), IEEE Press, 299-305.