INTECH OPEN ACCESS PUBLISHER

Multiple-Robot Systems for USAR: Key Design Attributes and Deployment Issues

Choon Yue Wong, Gerald Seet and Siang Kok Sim

Robotics Research Centre
School of Mechanical and Aerospace Engineering
Nanyang Technological University, Singapore

Abstract The interaction between humans and robots is undergoing an evolution. Progress in this evolution means that humans are close to robustly deploying multiple robots. Urban search and rescue (USAR) can benefit greatly from such capability. The review shows that with state of the art artificial intelligence, robots can work autonomously but still require human supervision. It also shows that multiple robot deployment (MRD) is more economical, shortens mission durations, adds reliability as well as addresses missions impossible with one robot and payload constraints. By combining robot autonomy and human supervision, the benefits of MRD can be applied to USAR while at the same time minimizing human exposure to danger. This is achieved with a single-human multiple-robot system (SHMRS). However, designers of the SHMRS must consider key attributes such as the size, composition and organizational structure of the robot collective. Variations in these attributes also induce fluctuations in issues within SHMRS deployment such as robot communication and computational load as well as human cognitive workload and situation awareness (SA).Research is essential to determine how the attributes can be manipulated to mitigate these issues while meeting the requirements of the USAR mission.

Keywords Single-Human Multiple-Robot System, Urban Search & Rescue, Communication, Computation, Cognitive Workload, Situation Awareness

1. Introduction

The execution of laborious tasks has undergone various transitions in paradigms especially in the past hundred years. In order to perform tasks on an industrial scale with precision and speed, humans have employed the use of machines. Through this, humans are developing progressively sophisticated methods of working and interacting with machines. Subsequently, it is now the goal for humans to deploy robots that are able to perform in a highly autonomous manner within the complex and unpredictable environments of the real world. Eventually, it is hoped that robots can not only work with only minimal human guidance[1], but also cohesively with other robots by autonomously coordinating their actions amongst themselves when taking on complex undertakings. These steps that humans have made and the milestones that humans hope to reach in terms of using machines and robots suggest that the interaction between humans and machines (and then later robots) is undergoing an evolution. For the evolution described here to be meaningful, it is assumed that the tasks performed are of some difficulty. This could mean that the tasks may be performed in an environment that is complex and unpredictable. Also, available time for completing the tasks may be limited. This evolution is illustrated in Fig. 1.

[1] Whether or not it is desirable for fully autonomous robots to work completely free of human monitoring is still debatable (Seet, Sim, Ong, Wong, & Lau, 2009). At this point in time, humankind has yet to have the luxury of this dilemma.

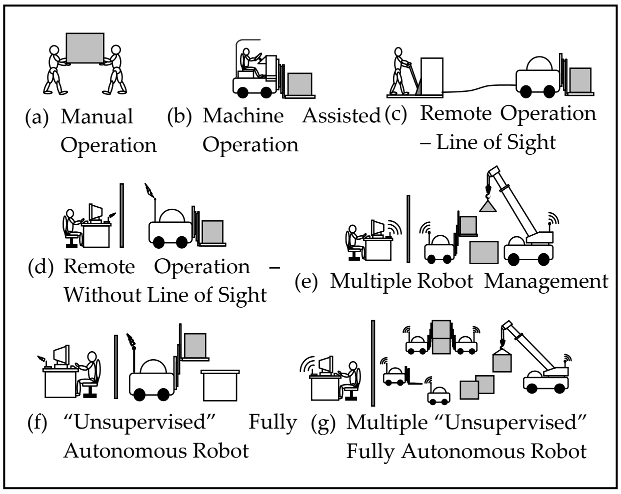

Figure 1. The evolution of interaction between humans and machines (followed by robots) (Adapted from Ong, Seet, & Sim (2003)).

Each stage in this nine-stage evolution is briefly described here. The evolution begins at Fig. 1a with labor performed manually and exclusively by humans. Machines appear at Fig. 1b when they are used to reduce or replace physical labor previously performed by humans. Fig. 1c illustrates the beginning of the use of machines to distance the human and the task performed. Despite the separation, line of sight (LOS) is maintained between human and the immediate task environment.

If the distance between human and machine is very large or if the view that the human originally has of the machine is occluded, then the machine has to be operated remotely without LOS. However, without being able to view the machine and the task directly, control of the machine for task execution is possibly impossible or at best difficult. This is because the lack of LOS means that it is much more difficult for the human to achieve a level of awareness for the robot's situation that is comparable to what the human has for the machine in Fig. 1c. Yet, if the machine is enhanced with the ability to "*extract information from its environment and use knowledge of the world to move safely in a meaningful and purposeful manner*" (Arkin, 1995), then it ceases to be simply a machine but has become an intelligent mobile robot. In that case, the robot may (unlike the machine) use autonomy to perform parts of the task which traditionally required humans to view. For example, the human may command a robot that is not within LOS to move in a set direction, but the robot autonomously performs obstacle avoidance within its immediate surroundings. The remote operation of a robot without LOS is the stage shown in Fig. 1d. It marks in this evolution, the beginning where robot autonomy could be highly essential.

Fig. 1e depicts the expansion on the concept of deploying a robot without LOS by increasing the number of robots deployed. However, the number of humans used for robot supervision is retained at one. Such application of many robots and one human mimics human-only organizations which make use of a single supervisor to monitor and coordinate the actions of multiple subordinates in the hope that the subordinates can efficiently produce effective results. This requires the human subordinates to be skilled at their individual tasks to produce effective outcomes. Just as important is the need for the subordinates' individual efforts to be coordinated so that the desired results can be achieved with minimal costs and in the shortest time possible. In fact, Bond & Gasser (1988) state that the extent of the ability for a team of agents to coordinate is determined by how much the agents can eliminate *"extraneous" activities* when they are working towards a common goal. The need for coordinated subordinate efforts in human-only organizations is mirrored in the human-robot system shown in Fig. 1e. The required coordination of the robots may be performed by the human supervisor, or autonomously by the robots themselves. In addition, Fig 1e illustrates a case of human-robot interaction (HRI) which each robot must definitely be imbued with some degree of autonomy if the robots are to continue operation while the human interacts with individual robots on a turn by turn basis.

If artificial intelligence is able to evolve to human-like levels, then robots would be able to function with human levels of independence. The application of a robot with such intelligence is illustrated in Fig. 1f. Yet, as mentioned earlier, even individuals in human-only organizations require supervision. Therefore, the robot in this case is not functioning entirely independent of supervision from a human. Rather, the robot could be monitored much as a human worker is managed by his supervisor. That is, the robot is left to perform its task independently for long periods. The human checks in on the robot from time to time or when the robot requests for help. Finally, in Fig. 1g, robots perform not only with full individual autonomy, but also with full collective-level autonomy. That means that the robots are able to coordinate their efforts autonomously amongst themselves. Again, despite the high level of autonomy, a human still has to intervene should the robots encounter any problems that they cannot solve by themselves.

Instances of deploying human-supervised robots (Murphy R. R., 2004a; Micire, 2008) indicate that we are at the fourth stage (Fig. 1d) of this evolution. That is, humans are capable of remote interaction with robots even without line of sight between robot and human. However, it is key to note that while the robust and coordinated deployment of multiple robots simultaneously has yet to be accomplished, a number of research efforts suggest that the fifth stage (Fig. 1e) is what many roboticists hope to achieve next.

Having reached the fourth stage of the evolution, it is natural to wonder how humankind can make use of and extend on our current capabilities for interacting with robots.

2. The Applicability of Using Roots for Search and Rescue

With the current state of interaction between humans and robots, researchers have identified a number of applications in which mobile robots using the current state of technology can contribute. Such applications include combat (Yamauchi, 2004; Barnes, Everett, & Rudakevych, 2005), space exploration (Schreckenghost, Fong, & Milam, 2008; Bellingham & Rajan, 2007; Halberstam, et al., 2006), oceanography (Bellingham & Rajan, 2007) as well as search and rescue. In the research described here, the focus will be on robot deployment for search and rescue.

The search and rescue of victims from dangerous and distressing situations is a highly difficult but crucial role that fire fighters and other emergency rescue professionals (also known as first responders) perform. Search and rescue may take several forms such as: mountain rescue, combat search and rescue, air-sea rescue and urban search and rescue. However, complexity/difficulty and limited time to rescue survivors are typical and common aspects of any type of search and rescue effort. For example, in the mountain rescue scenario described by MacInnes (2005), numerous rescuers are usually deployed in a coordinated search using radio, telephone and even road communications. The area in which the missing victim(s) are in would be divided into smaller areas for groups of rescuers to search. The mission environment is likely to consist of rugged and challenging mountainous terrain and could be conducted also in poor weather. Such rescue efforts are more often than not, aided by aircraft such as helicopters. Additionally, the victim(s) must quickly be located in order to minimize exposure to the elements and to increase their chances of survival (MacInnes, 2005). Tadokoro (2009a) pointed out that earthquakes can cause avalanches of soil and debris, fires and even soil liquefactions which are likely to cause ground travel extremely challenging. Added to this difficulty is the limited time with which rescuers must find and extricate victims during urban search and rescue efforts (Tadokoro, 2009a; Casper, Micire, & Murphy, 2000). This is because the likelihood of rescuing survivors decreases significantly 3 days after an earthquake (Fig. 2). As such, rescuers must make optimal use of time to locate and extract as many victims as possible.

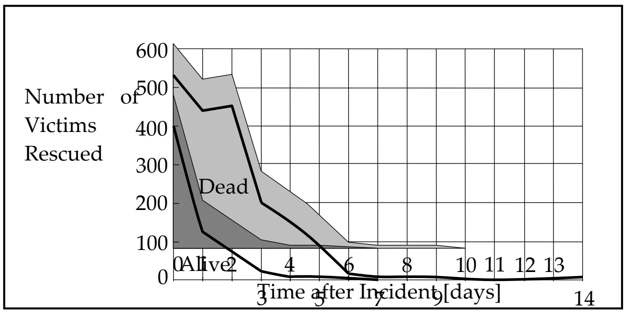

Figure 2. Number of victims rescued versus the number of days elapsed since the Hanshin-Awaji Earthquake (Kobe Earthquake) in 1995 (adapted from Tadokoro (2009a)).

While different types of search and rescue scenarios may share complexity and urgency, it is in particular, urban search and rescue that many researchers claim can benefit from robotics. Nourbakhsh, Sycara, Koes, M, Lewis, & Burion (2005) claim that USAR has become "*the canonical human-robot interaction problem, presenting an obstacle-ridden, unknown environment that can challenge robotic exploration even with the best of human assistance*". Robots can help in a disaster zone by helping to gather information pertaining to victims, the condition of the environment and potential hazards (Tadokoro, 2009a; Casper, Micire, & Murphy, 2000). Additionally, robots can be used to deliver water and medical equipment to victims (Casper, Micire, & Murphy, 2000). In their paper on using augmented autonomy to aid the deployment of human-robot teams, Nevatia, et al. (2008) assert that the exploration of unknown environments (after an urban disaster) is an application that stands to gain from the use of robots. Smaller robots can be effectively employed to examine voids in the rubble of buildings that are too narrow or dangerous for humans and rescue dogs to enter (Burke & Murphy, 2004; Murphy R. R., 2004a; Yanco, Drury, & Scholtz, 2004; Casper, Micire, & Murphy, 2000; Murphy R. R., 2004b). Shiroma, Chiu, Sato, & Matsuno (2005) suggest that robots can be applied to USAR missions such that rescue personnel may operate the robots at a safe distance. The robots can be used to collect information about hazards and victims, allowing human rescuers to skirt highly unstable or dangerous areas and reach victims quickly. Furthermore, robots can be used to carry a multitude of sensors into crevices. Notably, robots have in fact been deployed at ground zero in New York after the September 11 attacks (Murphy R. R., 2004a) as well as in the aftermath of Hurricane Katrina in 2005 (Micire, 2008) to help search for victims.

3. Work Done by Others

The deployments mentioned above have been possible due to research conducted previously. Presently, research into deploying robots for USAR situations is still very active. In addition to the examples of research mentioned in this paper, there are also competitions held to promote studies into using robots for USAR. One of the most prominent is the RoboCup Rescue League (RRL) held annually. The arenas in this competition are simulated disaster zones posing three different levels of difficulty for robots to maneuver and accurately map the disaster area. Competitors are scored according to their ability to locate and identify the status of the victim as well as their ability to generate an accurate map of the simulated disaster zone. Teams are penalized when their robot(s) collides with any victims or parts of the arena. Such competitions allow institutions from all over the world an opportunity not only to test and evaluate their robots but also to examine the solutions offered by other research organizations as well as to share their knowledge.

However, the research into USAR robotics can encompass numerous aspects such as robot mobility and the effects of graphical user interface on HRI. A number of research efforts are also looking into robot autonomy. Additionally, more researchers are paying attention to multiple robot deployment as well. Here, because research efforts in robot autonomy as well as multiple robot deployment are more pertinent to the challenges of deploying multiple mobile robots for USAR, a brief review of both areas of research is provided in the following paragraphs.

3.1 Research into robot autonomy

The term "autonomy", has been described as "*the ability of an agent (in this case, a robot) to act efficiently without any human intervention*" (Braynov & Hexmoor, 2002). Research into robot autonomy is motivated mostly by the desire to alleviate human workload. Yet, research into robot autonomy is also in part driven by concerns for the inability for the robot to receive supervision and guidance from the human when communication fails between human and robot (Bradley, Silver, & Thayer, 2004). In USAR, a robot may be called upon to work in a highly autonomous manner as well as simultaneously execute numerous functions such as navigation, mapping and perception (Casper, Micire, & Murphy, 2000). However, technology in terms of robot mobility, intelligence and perception is still not at a level capable of allowing robots to perform with total independence in dynamic scenarios envisaged in USAR (Shiroma, Chiu, Sato, & Matsuno, 2005). Therefore, researchers still advocate the presence of humans when robots are deployed such that the robots can be supervised (Nevatia, et al., 2008). For example, in their investigation of the benefits of sliding autonomy, Dias, et al. (2008) assert that current robot capabilities in many domains are yet to be sufficient for them to perform robustly and optimally in challenging situations. This is because robots fare rather unfavorably in real-world environments for a number of reasons. Firstly, the high uncertainty of the real world means that it is impossible for developers of autonomous mobile robotic systems to design for or even envision all possible scenarios that may or may not occur during a robot's deployment (Haight & Kecojevic, 2005). Secondly, the high level of complexity expected of such robots means that there is a greater propensity for the robot to fail (Landauer, 1995; Wickens & Hollands, 2000). Thirdly, the current state of technology is yet unable to produce robots that are able to robustly perform tasks that humans or even animals perform easily. Such tasks may include movement (Tadokoro, 2009b), perception (Tadokoro, 2009b; Chadwick, 2005; Fong, Thorpe, & Baur, 2003; Murphy, Kravitz, Stover, & Shoureshi, 2009; Murphy R. R., 2000), learning (Fong, Thorpe, & Baur, 2003) and decision making (Chadwick, 2005; Fong, Thorpe, & Baur, 2003; Murphy R. R., 2000). As such, humans must be called upon to provide properties such as judgment, flexibility, adaptability and experience to aid robotic systems deployed in challenging and unpredictable environments (Haight & Kecojevic, 2005). Such attributes are especially necessary when component failures or difficult situations occur (Wang & Lewis, 2007). With human supervision, robots may also be made simpler, with a lower level of artificial intelligence and fewer sensors.

Before robots can perform with full autonomy in USAR, humans have to make do with the limited autonomy that robots now are capable of and interact with robots in various modes through the spectrum of autonomy. This spectrum ranges from teleoperation to full autonomy. During teleoperation, a human controls the robot to direct its actions (Sheridan, 1992). In other words, the robot exhibits no autonomy. In the case of a mobile robot in USAR, the human may operate the robot by driving it and performing actions such as panning the on-board camera and moving the robot's manipulator. This kind of control occurs via a master-slave scheme such that the robot acts as the slave to the human who directs all of the robot's actions. Not surprisingly, a robot being teleoperated will need constant interaction with the human (Desai & Yanco, 2005; Sheridan, 1992). If a robot has greater autonomy, less interaction and attention is necessary from the human until finally when it has full autonomy, the amount of attention that the human has to provide for the robot reaches a minimum. Therefore, it is reasonable to expect that the human experiences the most cognitive workload when teleoperating the robot. Conversely, the lowest workload is experienced when the robot is operating with full autonomy.

Interactions between the extremes of teleoperation and full autonomy are known as shared control (Desai & Yanco, 2005) whereas operation at teleoperation and full autonomy is known as traded control (Sheridan, 1992). Murphy (2000) claims that a robot deployed using shared control operates with the human supervisor monitoring its actions. When necessary (possibly when something is about to or has gone wrong), the human intercedes and corrects the robot's actions. Murphy also suggests that because "*boring, repetitive control actions*" are allocated for the robot to perform, shared control can help to alleviate cognitive fatigue. The spectrum of autonomy as well as the paradigm for sharing and trading of robot control is illustrated in Figure 3.

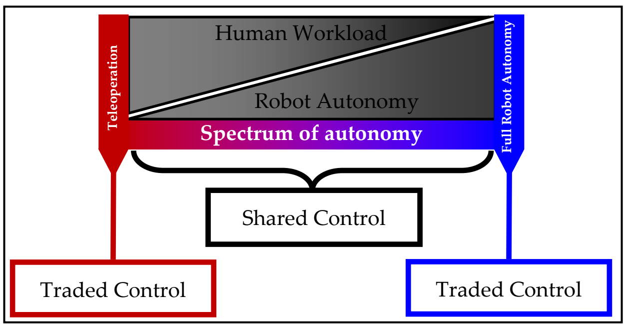

Figure 3. The spectrum of autonomy

Additionally, human and robot may interact at different autonomy levels during the course of a mission (Sheridan, 1992). Such changes in autonomy levels are needed as certain portions of a mission may require more human attention and inputs while other parts of the mission can be more easily accomplished with greater use of robot autonomy. This movement along the spectrum of autonomy has been known as sliding autonomy, adjustable autonomy as well as mixed initiative.

3.2 Research into Multiple Robot Deployment (MRD)

Research into MRD is typically motivated by the desire to improve on existing systems that deploy single robots. A number of benefits can be derived from MRD:

1. Firstly, the deployment of multiple robots means that missions can be completed more quickly since progress towards meeting mission objectives can be made via multiple fronts (Shiroma, Chiu, Sato, & Matsuno, 2005; Dudek, Jenkin, & Milios, 2002; Parker, 2008).
2. Secondly, if multiple copies of the same robot are deployed such that a homogeneous robot collective is formed, then redundancy can be achieved. This is because individual robots can easily replace a failed or lost member of the collective. Therefore, the system as a whole becomes more reliable (Dudek, Jenkin, & Milios, 2002; Parker, 2008; Cao, Fukunaga, Kahng, & Meng, 1997; Balch & Arkin, 1994; Arkin, Balch, & Nitz, 1993; Parker, 2003).
3. The use of multiple robots also helps to address limitations in individual robot payload (Balch & Parker, 2002; Parker, 2008). For instance, a typical USAR task such as determining the state of a trapped victim will require multiple pieces of hardware such as a microphone, cameras as well as a source of illumination. The application of a robotic solution to this task will require multiple robots should there not be any single robot capable of providing all the necessary equipment (Koes, Nourbakhsh, & Sycara, 2005).
4. Rather than integrating all the necessary hardware for mission completion into a single robot, it is possibly more economical and easier to distribute components amongst multiple robots and then enable cooperation between these robots than to integrate all necessary capabilities into a single robot (Balch & Parker, 2002; Parker, 2008; Cao, Fukunaga, Kahng, & Meng, 1997).
5. Multiple robots enable the completion of certain goals[2] that cannot be achieved with the use of a single robot (Dudek, Jenkin, & Milios, 2002; Parker, 2008; Cao, Fukunaga, Kahng, & Meng, 1997).

Given the advantages that the deployment of multiple robots can provide (over single robot deployment), it would appear that many researchers would be keen on studying MRD. In fact, it is indeed so. Furthermore, Casper, Micire, & Murphy (2000) assert that the USAR domain "*lends itself to the use of multiple robots with multiple capabilities*". In the following paragraphs are some examples of relevant research.

One of the earlier works by Goldberg & Matarić (1997) was a study into minimizing instances in which four robots had to avoid one another during a foraging mission. By reducing the number of times robots had to evade one another, system performance can be improved as the mission can be completed more quickly and with less power consumed. In this study, the levels of performance for three different strategies for the foraging mission (which can mimic aspects of a toxic waste cleanup or a victim search and retrieval mission) were compared. The objective of this comparison was to ascertain if physical interference[3] can be used as a tool for designing optimal behavioral schemes for cooperating robots in a collective.

Another effort pertaining to multiple robot cooperation from the 90s was that of ALLIANCE architecture (Parker, 1998). This notable method proposed by Parker suggested that individual robots could be allocated motivations for each of the many tasks necessary to complete a mission. Each robot's motivation to execute a task is influenced by its impatience or acquiescence towards the task. If its impatience to adopt a task reaches a threshold, a robot would then claim the task for its own while at the same time, inhibiting others from doing so. Conversely, if a robot recognizes that it is inept at the adopted task, its level of acquiescence rises until it forsakes the task and frees it for others to perform. Therefore, ALLIANCE is able to achieve simple cooperation between robots without robots having to explicitly[4] negotiate tasks with one another. This architecture has since been extended to become L-ALLIANCE (Parker, 2000) such that robots can use the experiences that they have had with other members of the robot collective to learn about and adapt to capability changes in their teammates.

These examples are only a sample of research pertaining to early work in MRD. A detailed review has been done by Cao, Fukunaga, Kahng, & Meng (1997). Additionally, Parker (2008) offers a survey that is more recent. While earlier research efforts[5] placed emphasis on cooperation techniques between robots, researchers are now as focused towards the interaction between human and robot(s) as they are towards coordination strategies

[2] Such goals include tasks that must be simultaneously performed but are spatially separated or tasks that are too challenging for a single robot.

[3] Goldberg & Matarić describe physical interference as robots getting in the way of one another.

[4] Explicit negotiation or communication between robots can place a burden on available bandwidth.

[5] Such efforts may not have been made to directly address USAR scenarios but rather for simpler missions such as toxic waste cleanup.

between robots. Researchers also investigate the level cognitive workload and SA that the human has for the robots and the USAR mission.

3.3 Implications of work done

The discussion in Section 3.1 has presented the necessity for a robot to be supervised by humans despite its autonomy. Even so, robots may be able to work unsupervised for limited durations. As such, during this time in which the robot is left unattended, a human supervisor can divert attention elsewhere, such as towards the supervision of another robot. In fact, the very possibility a single human supervising multiple robots is attributed to robot autonomy (Goodrich, Quigley, & Cosenzo, 2005). Section 3.2 has presented a number of compelling advantages for MRD. Given the motivation for MRD and the ability for robots to be left temporarily unattended, the simultaneous deployment of multiple robots using fewer human supervisors than there are robots is certainly an appealing possibility.

The use of fewer humans than robots to provide supervision for the robots appears to contradict common sense. If manpower is sufficient, it seems only logical to apply all available human resources to aid in the robots' deployment. More humans for robot supervision would mean closer supervision for each robot, reducing or eliminating the chances of some important piece of information gathered by the robots being missed. More humans will also denote less cognitive workload for each human supervisor. However, the need to limit the number of humans for supervising multiple robots can be explained by considering practical robot deployment in the hazardous environments typical of USAR. In order to minimize human exposure to the dangers of the mission environment, only a small number of human supervisors should be used. Murphy, Blitch, & Casper (2002) reported that during the use of robots to search the rubble of the World Trade Center in 2001, fire chiefs wanted to keep the number of people on a rubble pile or within a structure to a minimum. Hence, the supervision of the deployed robots should be undertaken by just a single human (Kadous, Sheh, & Sammut, 2006). This preference is reflected in the rules of the RoboCup Rescue Physical Agent League Competition which issues penalties to teams with more than one individual in the simulated disaster site. Given the possible advantages of deploying multiple robots and the need to minimize human presence at the mission area, it appears that the SHMRS is ideal for USAR.

A SHMRS consists of two main portions. The first is that of the single human while the second is made up of the robots deployed using the system. Together, the robots form what is termed as a collective.

4. Key Design Attributes and the Issues with SHMRS Deployment

However, while the SHMRS may be promising for the future of USAR, consideration must be given to designing an effective SHMRS. Specifically, key properties or attributes of the SHMRS have to be established prior to deployment. Three key design attributes of the SHMRS are first identified and introduced briefly in the following paragraphs before being discussed further in Section 5 and Section 6.

4.1 Collective Size

A design consideration that is immediately obvious is that of the number of robots or the size of the collective deployed using the SHMRS. The desired collective size may be influenced by the requirements of the mission but limited by the issues (described later) of maintaining a collective of a particular size. For example, a time-limited USAR scenario may call for many robots to be deployed to maximize parallelism and overall collective reliability but greater numbers of robots make the coordination of their actions increasingly difficult (Wagner & Arkin, 2004).

4.2 Collective Composition

Mission requirements may also call for particular types of hardware to be brought by the robots to the mission environment. Depending on whether each robot is able to carry a set of all essential equipment to the work environment, rescuers have to choose either to deploy a homogeneous collective or heterogeneous collective of robots. The discussions regarding collective composition later in Section 5 and Section 6 will demonstrate that employing either type of collective has its benefits and drawbacks. Therefore, consideration must also be provided for the composition of the robot collective.

4.3 Collective Structure

Yet another attribute of the SHMRS that must be established before deploying the SHMRS is the structure of the robot collective. This refers to the organization of the relationships of authority which robots within the collective have with one another. Regardless of whether or not it is the designer's explicit intention, any collection of agents or robots adopts a structure. Horling & Lesser (2005) suggest that consideration for a collective's structure is important because a system's organizational structure can have significant impact on its behavior and performance. It should be noted that structure, unlike the other two key design attributes of size and composition, is not directly driven by the requirements of the collective's mission. Rather, the motivation to select a particular structure over another is derived by its abilities to minimize the cost of autonomously coordinating robot actions while maximizing the ease with which the collective is supervised by the human.

4.4 Issues with SHMRS deployment

However, it is not enough simply to consider the design attributes. In addition, one also has to be aware of the issues with the practical deployment of multiple robots when using only a single human supervisor. These issues may broadly be classified into two categories. The first, is the set of technical issues pertaining to (autonomous) inter-robot coordination (IRC). A major concern within this set is that of communication load imposed on robots when autonomous coordination is performed. If the collective is large, crippling communication overheads can result (Sugihara & Suzuki, 1990; Anderson & Papanikolopoulos, 2008). Yet, communication (which typically occurs wirelessly) is an integral part of autonomous coordination. Unfortunately, wireless communication is known notoriously to be unrealiable. For instance, Tadokoro (2009a) claims that poor quality communication is not unexpected at disaster zones. Additionally, Nevatia, et al. (2008) assert that it is unfeasible to implement methods for inter-robot coordination which require communication between robots to always be ideal. The unreliability of communications has been demonstrated by Casper & Murphy (2002) when wireless ethernet between a robot and the control station was lost on multiple occasions during a field test. The loss of the only robot during rescue efforts at Ground Zero in New York has also been attributed to communications failure (Murphy R. R., 2004b). Another concern within the set of technical issues is the computational demands on robots when performing autonomous coordination. For example, the multiple coordinated exploration study described by Burgard, Moors, Stachniss, & Schneider (2005) reveals that increased computation (resulting from adding robots to the collective) tended to slow the robots. In order to prevent robots from being bogged down with excessive computation, it is essential to keep the amount of computation needed for IRC low. Communication and computational costs have been echoed by Scerri, et al. (2003) as the key weaknesses of existing role allocation algorithms for robots.

The second set of issues pertains to the limits of human cognition. Specifically, these concerns refer to the high cognitive workload that the human supervisor can experience when supervising multiple robots as well as the supervisor's lack of SA for the robots, their individual progress and status of the mission. High levels of cognitive workload will degrade the quality of the supervision provided by the human. Furthermore, the lack of SA can result in poor decision making and an inability to troubleshoot problems when they occur.

Both sets of issues will be also discussed in Section 5 and Section 6 with greater detail. From these sections, it will be possible to appreciate that the identification of the three key design attributes is crucial to SHMRS design due to each attribute's influence on the both sets of deployment issues. Therefore, Sections 5 and 6 will additionally include elaboration on the expected interactions between the key design attributes with both sets of SHMRS deployment issues.

5. The Key Design Attributes and their Influence on Inter-Robot Coordination

This section describes in further detail, the key attributes that have been introduced in Section 4. The anticipated interactions of each attribute with the technical issues of IRC are presented as well. The interactions between size and IRC are first to be presented followed by those of composition and structure.

5.1 Collective Size and its Influence on Inter-Robot Coordination

As mentioned previously, the requirements of the mission can influence the sizing of a collective. Because many robots can complete a mission quicker, one may assume that it is best deploying a large collective. However, doing so creates a number of problems for IRC. In their work on multi-robot reconnaissance, Wagner & Arkin (2004) assert that greater robot numbers make IRC increasingly challenging. Yet, the benefits afforded by the application of multiple robots such as those described in Section 3.2 cannot be ignored.

The first concern with expanding the size of the collective is communication cost. Robots within larger collectives may have to communicate with more collective members as more robots are added (Fox, 1979). As such, large robot collectives can induce crippling communication overheads (Sugihara & Suzuki, 1990; Anderson & Papanikolopoulos, 2008; Sweeney, Li, Grupen, & Ramamritham, 2003).

The second concern is that of rising computational overheads for individual robots. In addition to having to process communicated information, computational power is sometimes expended on modeling other robots. This in turn leads to delays as the time needed to perform this computation grew as the collective's size increased (Burgard, Moors, Stachniss, & Schneider, 2005).

Further complication can result from expanding the size of the collective due to the increased possibility for inter-robot interference. Such interference may be manifested physically in the form of inter-robot collisions (Goldberg & Matarić, 1997; Rosenfeld, Kaminka, Kraus, & Shehory, 2008) (especially if the mission area is small). Additionally, interference between robots can exist as crosstalk between radio frequency, infra-red and ultrasonic signals (Goldberg & Matarić, 1997).

Also, while the performance of the robot collective can initially improve as robots are added, this improvement cannot be sustained as the collective is further enlarged. This is reflected in a simulation conducted with a robots deployed in a puck-collection mission (Rosenfeld, Kaminka, Kraus, & Shehory, 2008). The number of pucks retrieved grew as the first few robots were added but mission performance declined when even more robots were deployed. The robots spent more time and fuel attempting to avoid other robots than making actual progress towards the mission. Figure 4 depicts the relationship between collective size and collective performance as well as the costs associated with multiple robot coordination.

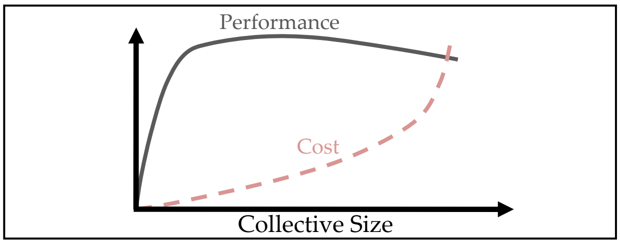

Figure 4. Plot of collective performance and cost of coordination versus collective size.

5.2 Collective Composition and its Influence on Inter-Robot Coordination

A collective is classified as homogeneous if "*the capabilities of the individual robots are identical.*" (Cao, Fukunaga, Kahng, & Meng, 1997). Similarly, a collective is considered heterogeneous if "*members vary in their sensor and effector capabilities*" (Parker, 2008). Homogeneous robot collectives are considered as they offer parallelism and redundancy (Parker, 2003).This means that all members of the collective are able to concurrently contribute towards fulfilling mission objectives since each robot is as capable as the next. Robots are able also to replace one another with greater ease should the need arise.

Heterogeneous robot collectives on the other hand, allow designers of MRS to distribute possibly a limited number of sensors or other hardware components. These hardware items might be too costly, too heavy or too large to be featured on every robot (Parker, 2008). In certain instances, the deployment of dissimilar robots may simply be due to the nature of the mission which makes heterogeneous robot collectives a necessity (Balch & Parker, 2002). Murphy (2004b) asserts that no single robot platform or payload configuration will suffice for a USAR scenario due to the myriad of void sizes and the tasks that rescue robots can be involved in. If multiple robots are to be deployed, designers of such MRS must seriously consider the application of a heterogeneous robot collective.

Unfortunately, heterogeneous robot collectives may require closer coordination. This is due in part to the inability for robots in such collectives to change roles easily, causing cooperative control of robots to be more challenging (Parker, 2003). In a distributed sensing mission for instance, any robot within a homogeneous robot collective may be assigned to perform navigation as well as the necessary detection duties. However, in a collective where some robots are able to navigate well and localize while other robots are better suited for detecting (perhaps signs of life like human voices for instance), closer coordination between surveying robots and navigation robots is necessary such that the surveying roots are deployed at the correct locations. Such a scenario has been described by Parker, Kannan, Fu, & Tang (2003) where two (different) robots for navigation and localization were used to lead two similar robots for acoustic surveillance. The robot responsible for localization relied heavily on the navigation robot to find its way in the environment. In addition to following the navigation robot, it had to teleoperate the two robots responsible for acoustic surveillance. In this scenario, mission objectives would not have been met if simpler coordination methods suitable for a homogeneous robot collective such as robot deployment based simply on 'geographical proximity' or 'areas previously surveyed' had been used. Such close coordination needed when the collective is heterogeneous may demand more communication and computation from robots.

5.3 Collective Structure and its Influence on Inter-Robot Coordination

While there are many structures from the domain of multi-agent systems for organizing agents (such as those in the survey by Horling & Lesser (2005)) that may also be applied to structuring robots into a collective, particular attention is paid in this research to the hierarchical and horizontal structures. This is because it is these structures that have demonstrated their potential and feasibility for use on robots by having been implemented on the MRS domain by other researchers[6]. For example, the hierarchical structure has been used by Elston & Frew (2008) to organize multiple aerial robots. The FIRST (Friendly Interactive Robot for Service Tasks) system (Causse & Pampagnin, 1995) also utilized the hierarchical structure in order to transport heavy loads around a hospital environment[7]. Additionally, the horizontal structure has been applied to robots performing a flocking mission (Matarić, 1992) as well as robots taking part in an urban search and rescue scenario (Wang & Lewis, 2007). Both the hierarchical and horizontal structures are briefly described here.

[6] Although in some cases, the application of these structures was not the focus of work presented.
[7] However in that work, the top level decision making agent was stationary rather than a mobile robot.

The Hierarchical Structure

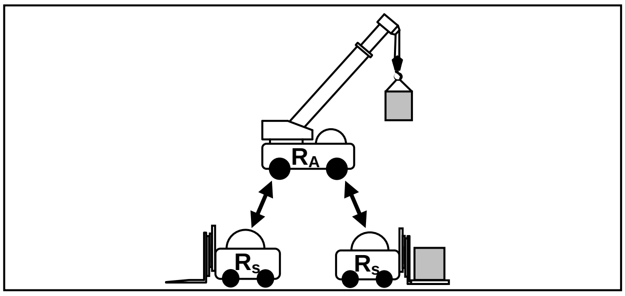

Figure 5. The hierarchically structured robot collective.

The hierarchical structure exhibits a number of characteristics. Firstly, there is a central or apex robot (R_A in Fig. 5) which has a major role in making higher level decisions for the rest of the collective (Fox, 1979). Interactions between members of the hierarchy are limited to between supervisors and subordinates (Dignum & Dignum, 2005). Commands from highly ranked robots flow downwards while feedback gathered by lower ranked subordinate robots will flow upwards (Bond & Gasser, 1988). In addition, individual robots may not have to model many collective members. That is, individual robots (such as subordinate robots (R_S) in Fig. 5) do not have to be aware of aspects pertaining to other members of the collective such as their roles, responsibilities, capabilities, limitations or possibly even their existence. This "ignorance" translates into computational savings for robot individuals.

If certain members of the collective do not have to be aware of one another, then there is no need for such members to communicate. Therefore, the hierarchical structure allows for close coordination between robots while also facilitating reduced communication between robots (Horling & Lesser, 2005). Another possible advantage in using the hierarchical structure stems from the fact that the roles of subordinate robots are not expected to be as complex as higher ranked ones and as such, subordinate robots perform less computation and communicate less than apex robots do.

Unfortunately, MRSs using the hierarchical system are prone to single-point-failures (SPFs) (Horling & Lesser, 2005; Maturana, Shen, & Norrie, 1999) as the failure of certain robots can spell the loss of a large portion of the collective or even the failure of the whole collective (should the apex or another highly ranked robot fail (Parker, 2008)).

The Horizontal Structure

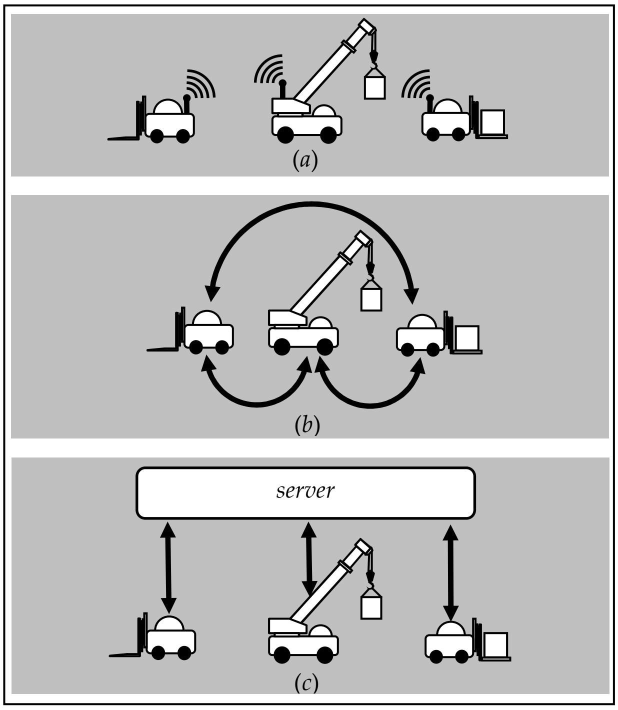

Figure 6. The hierarchically structured robot collective using a) communication via broadcast b) communication between robots on a point-to-point basis and c) communication between robots using a server.

On the other hand, there are no relationships of authority between members in a horizontally structured robot collective. Also, robots within the horizontal structure may adopt more than one strategy of sharing information with one another. That is, they may perform a broadcast such that all robots within range can interact with one another (Fig. 6a) or they select to communicate with a particular individual within the collective by addressing it (Fig. 6b) (Stone & Veloso, 2002). Additionally, robots can communicate via a server which acts merely as a storage for information (Fig. 6c) (Stone & Veloso, 2002). Such a server is often termed a "blackboard". Each of these methods has its own strength and weaknesses. For example, IRC via the broadcast paradigm could require all robots to receive and interpret broadcasted messages but this allows for savings in communication bandwidth[8]. Highly coordinated collective behaviours can be achieved if robots communicate by addressing specific individuals. This ensures that only relevant robots receive the necessary information but it may also burden robots with higher bandwidth. Finally, sharing information via the server can help to reduce bandwidth requirements[9] but the failure of the server would be catastrophic to the collective. Of course, a duplicate or server can add

[8] Each robot will not require its own dedicated radio frequency channel but must have the computational capabilities to interpret messages.

[9] Each robot will need to communicate only with the server.

redundancy to the overall system but care has to be given to ensure that data stored within both servers is identical. Also, the robots must seamlessly transition to communicating with the backup server in order for a redundant server to be meaningful.

Another characteristic of the horizontally structured collective is that individuals within the horizontally structured collective may be required to collectively agree on decisions affecting the collective as a whole (Fox, 1979) since a centralized decision-making authority is absent. An instance of multiple robots reaching a consensus has been described by Werger & Mataric (2000), in which physical robots were deployed to observe a number of targets. Each robot would claim tasks for itself by broadcasting its eligibility for a task to the other robots. The robot with the highest eligibility wins the right to perform the contested tasks.

Robots within a horizontally-structured collective taking part in complex missions are also likely to have to model one another (Stone & Veloso, 2002). An example of such modeling is the work presented by Burgard, Moors, Stachniss, & Schneider (2005) pertaining to the simulation[10] of robots assigned to perform an exploration mission. Each robot would compute for itself, target waypoints that do not conflict with those of their peers. Therefore, the intelligence and computational prowess expected of all robots within the horizontally-structured collective are expected to be at a level higher than that of certain robots using a hierarchically-structured collective.

These characteristics of the horizontally structured collective suggest that if close and optimal coordination is expected, then robots should have greater communication and computational competencies than robots using the hierarchical structure. Therefore, it may be presumed that smaller and simpler robots cannot be included in such collectives. After all, the hardware necessary to perform the required computation and communication can only be miniaturized up to a certain limit (Casper, Micire, & Murphy, 2000).

Yet, it is useful to note that horizontally structured collectives (if not communicating via the server) are less prone to SPFs since there is no apex or high-ranking robot at risk of failing (resulting in the loss of subordinates reporting to the apex robot).

6. The Effects of the Key Design Attributes on Multiple Robot Supervision

This section addresses the effects that each of the three key design attributes have on the cognitive issues of multiple robot supervision with a single human. The interactions between collective size and multiple robot supervision are first to be presented followed by those of collective composition. Finally, the effects of altering collective structure on multiple robot supervision are presented last.

6.1 The Effects of Collective Size on Supervising Multiple Robots

To the human supervisor, increasing the size of the collective can result in high cognitive workload since doing so will increase the demand for mental resources. Hart & Wickens (1990) define mental workload as: "*A general term used to describe the mental cost of accomplishing task requirements for the human element of a man-machine (in our case, man-robot) system*".

As a consequence of high cognitive workload, the quality of supervision provided by the human will degrade. For example, the human supervisor may fail to perform lower priority tasks or perhaps even not realize that certain high priority tasks require attention. This degradation of supervision quality in relation to the number of robots that the supervisor must manage is illustrated in Figure 7. Here, as also stated by Wickens & Hollands (2000), it is assumed that workload is high when the supplied mental resources are high.

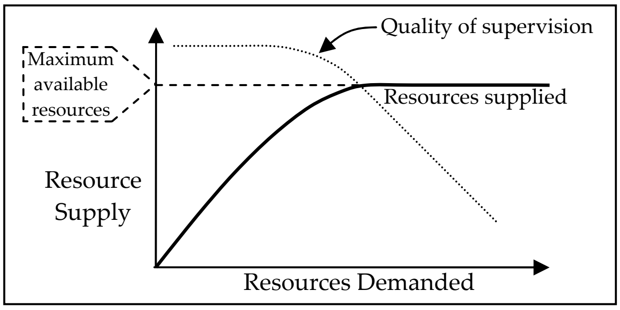

Figure 7. Plot of resource supply vs. resources demanded (adapted from Wickens & Hollands (2000)).

Another component of human cognition crucial to the supervision of multiple robots is situation awareness. Formally, SA is defined as "*the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future*" (Endsley, 1995). Murphy (2004a) asserts that in USAR, data provided by the robots can provide SA about the mission environment and the robot such that robots can navigate safely while providing thorough area coverage. Additionally, she also claims that a search using robots can yield SA about the objects within the rubble and the progress of the search mission. However, Riley & Endsley (2005) claim that SA for the human supervisor can be adversely affected by the multi-tasking and task switching demands of managing multiple robotic vehicles.

As a human is assigned to supervise increasingly large robot collectives, the level of cognitive workload

[10] Comparatively, experiments conducted in simulations are low costs as well as easily set up and repeated. However, the existence of issues such as communication unreliability and other uncertainties of the real world which cannot be fully replicated in simulations mean that experiments and results derived from physical tests have greater validity.

increases in tandem until it reaches a maximum. When this occurs, the human supervisor's level of situation awareness for the robots and the mission as a whole is expected to fall dramatically (Endsley & Jones, 1997). This relationship is depicted in Figure 8. If this occurs, the supervisor's abilities for good decision making (Murphy R. R., 2004a) as well as for troubleshooting robot anomalies or failures will be compromised (Bruemmer & Walton, 2003; Kaber & Endsley, 1997).

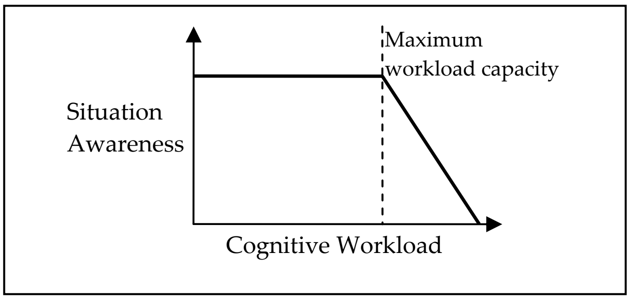

Figure 8. Plot of situation awareness vs. cognitive workload (adapted from Endsley & Jones (1997)).

6.2 The Effects of Collective Composition on Supervising Multiple Robots

During the supervision of multiple robots, the human supervisor is likely to interact with the robots on a turn-by-turn basis although perhaps not always in the same sequence. This requires him or her to switch attention between the robots being supervised. If a homogeneous collective is deployed, then the attention switches occur between similar robots. The converse is true if the human supervises a heterogeneous collective.

One of the interests in this research is to ascertain if the human supervisor will experience difficulty when attention is switched between robots. If so, will attention switching between dissimilar robots result in more cognitive workload for the human supervisor than between similar robots? In their investigation of interface design for systems to interact with multiple robots, Weil, et al. (2006) suggest that the control of multiple remotely-operated vehicles can be made more difficult if it is required of operators to manage vehicles of different types. Yet, Wickens & Hollands (2000) claim that attention switching between highly similar tasks may lead to cognitive interference as the latter task is performed. The views presented by these two works appear to contradict one another but from them, it is possible to infer that both sets of authors subscribe to the belief that the composition of a set of tasks can influence attention switching. Additionally, it is not unreasonable to wonder if it is the switching of attention between dissimilar or similar robots that makes the supervision of multiple robots more challenging or perhaps it is shifting attention between the robots' contexts that adds difficulty to the human supervisor's role. Research into the effect of collective composition on the human supervisor's ability to supervise is still somewhat lacking presently and work has to be done to fill this void in our knowledge towards understanding HRI between humans and multiple robots. Therefore, one of the objectives in this research is to examine how the composition of the supervised collective of robots can influence the human supervisor's switching of attention between robots.

6.3 The Effects of Collective Structure on Supervising Multiple Robots

The discussion on HRI between a single human and multiple robots begins with how organizational structure influences this interaction. As with the earlier study into the effects of collective structure on IRC, the horizontal and hierarchical structures are explored here. However, included too within this investigation are considerations for whether or not the robots are able to share information with one another (and therefore autonomously coordinate their actions) as well as whether or not the human can interact with individual robots. The effects of the robot collective's structure on the ability of the human supervisor to supervise multiple robots is examined by considering five different models of the SHMRS (three using the horizontal structure and two using the hierarchical structure).

Multiple robot supervision using the horizontal structure is first considered. Three variants of this structure applied to the SHMRS are examined in the paragraphs that follow.

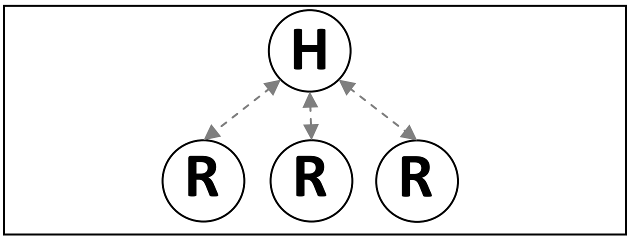

Figure 9. The Horizontal, Manually Coordinated Model.

The horizontal, manually coordinated model (Figure 9) for supervising multiple robots is essentially, the simplest of the three presented here. Incidentally, this model is the same as that classified as "*one human, multiple robots*" by Yanco & Drury (2002). Within this model, the human interacts with each robot via direct channels of communication. That is, each robot acts as if it is the only deployed robot within the environment and is incapable of sharing information with other collective members. As such, any requirements for coordinating the actions of the collective must be met entirely by the human. That being the case, the human supervisor is likely to experience high levels of cognitive workload. SA is expected to be poor as well. An example of a horizontal, manually coordinated collective supervised by a single human is that of the (simulated) UAV collective presented by Dixon, Wickens, & Chang (2005).

Fortunately, robot collectives may be designed such that autonomous coordination of robot actions can be performed. To facilitate autonomous coordination, robots must have the capability to share information with one another without enlisting the help of the human supervisor. As illustrated by Fig. 6, information can be communicated in a number of ways between robots within the horizontal structure. However, for brevity, only the communication paradigm shown in Fig. 6b is used to illustrate the supervision of multiple autonomously coordinating robots in Figure 10.

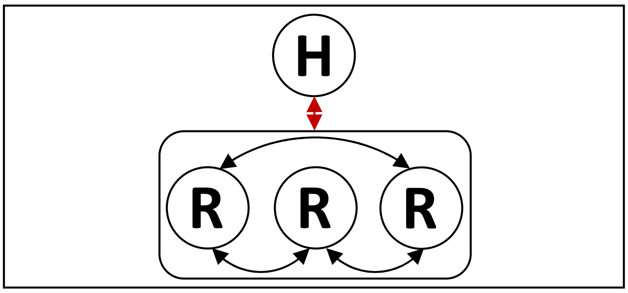

Figure 10. The Horizontal, Autonomously Coordinated Model.

With autonomous coordination, cognitive workload for the human supervisor is decreased. Here, the human instead of communicating with individual robots, sends high-level commands to the collective. The robots amongst themselves deliberate how each of them is to contribute to meeting command requirements. The model depicted in Figure 10 falls into the category of "*one human, robot team*" described by Yanco & Drury (2002)[11].

Unfortunately, autonomous coordination is not expected to come cheaply in terms of computational and communication costs. Much of each robot's available bandwidth may be dedicated to sharing information amongst themselves and as such, communication regularity between robots and human is expected to be low. A large portion of on-board computational power can also be vested in performing autonomous coordination and therefore any information communicated to the human may be highly abstracted. Therefore, without adequate feedback to the human, the collective effectively masks from the human, details for the robots performance during mission execution. This leads to the notion that the human supervisor's level of SA from supervising via the horizontal, autonomously coordinated model is most likely woefully short of the level necessary for missions in dynamic environments expected of applications such as USAR. The possibility for a human facing such a predicament may be inferred from studies in aviation systems. For instance, Sarter, Woods, & Billings (1997) reported that pilots did not always comprehend what the automated flight management system was doing or the reasons behind the system's behaviour. Sadly, confusion and the lack of SA due to automation-induced system opacity have led to a number of disasters involving jetliners such as the 1995 loss of a Boeing 757 near Cali Columbia[12] (Sheridan, 2002). In addition, Endsley & Rodgers (1996) suggested that because of the complexity inherent in automated systems, anomalies with the system could be difficult to detect. Even when problems are found, diagnosis of the problem might still prove challenging.

For the human to take advantage of autonomous coordination while still maintaining an adequate level of situation awareness, the direct channels of communication between human and robots from the horizontal, manually coordinated model (Figure 9) can be integrated with the horizontal, autonomously coordinated model (Figure 10) such that the horizontal, autonomously coordinated model *with* direct communication channels is produced (Figure 11).

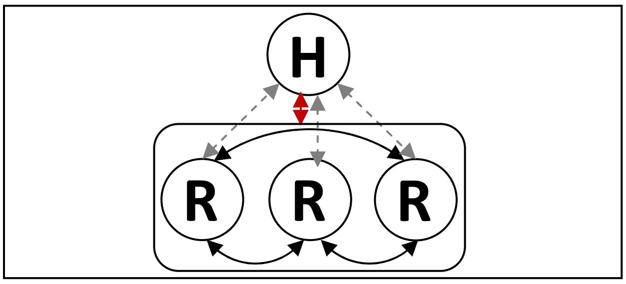

Figure 11. The Horizontal, Autonomously Coordinated Model *with* Direct Communication Channels.

Using this model, some of the coordination that is perhaps too challenging for the robots to handle by themselves can be off-loaded to the human via the direct channels of communication (dashed arrows). This is much like the suggestion put forth by Sheridan (2002) that states that robot control can be shared and traded so that parts of a task can be allocated either to the human or machine. Doing so decreases the amount of computation that the robots must handle by themselves as well as reduces the communication that must occur to produce coordination. It also helps to derive a synergy between properties provided by the human such as perception and decision making with the autonomous coordination performed by the robots (Wang & Lewis, 2007). The autonomous coordination helps to alleviate workload demands on the human while interaction between the human and individual robots through the direct channels of communication helps to deliver feedback to the human

[11] However, it should be noted that in that article, Yanco and Drury did not make any distinctions between the structures employed by the robot collectives. Rather, they classed a human-robot system within this category as long as the system consisted of one human and a collective of autonomously coordinating robots. In effect, this puts the rest of the models which follows as within the "*one human, robot team*" category as well.

[12] Unknown to the crew who were flying at night, the flight management system had directed the aircraft towards a mountain. The flight crew was confused as to why the aircraft was not flying to the desired waypoint.

with greater regularity and richness. This will help to facilitate in improving human situation awareness for the robots.
Autonomous coordination between robots is possible too if an individual within the collective is given the capability to make decisions for others. By doing this, the supervision of multiple robots using the hierarchical structure is derived. Two variants of the hierarchical structure applied to the SHMRS are to be examined next. The models below are examples of the hierarchical structure being applied to the SHMRS. Figure 12 depicts multiple-robot supervision via the basic hierarchical model.

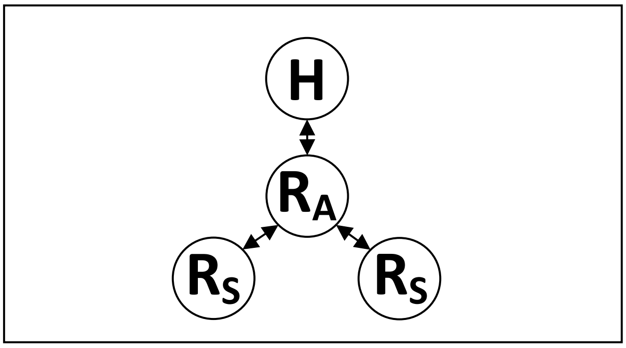

Figure 12. The Basic Hierarchical Model.

With this model, the responsibility for coordinating the actions of all the robots within the collective belongs to the apex robot. As with autonomous coordination attained using the horizontally-structured models illustrated by Figure 10 as well as Figure 11, cognitive workload for the human supervisor is reduced as a result. The supervision of multiple robots using the hierarchical model also means that autonomous coordination can be achieved with less computational power and lower communication overheads for the collective (than the horizontally structured collectives).
The advantages of using the basic hierarchical model can be extended by implementing direction channels of communication between the human and individual robots. This gives the supervision of multiple robots via the hierarchical model *with* direct communication channels (Figure 13).

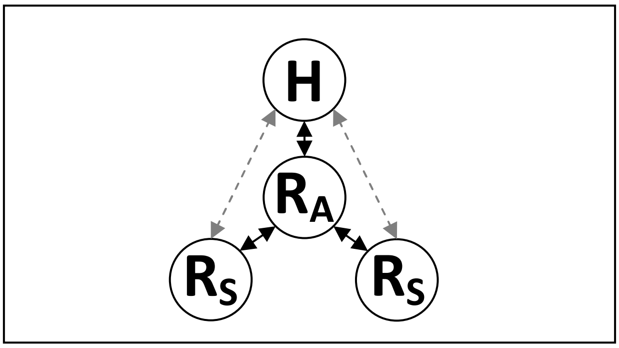

Figure 13. The Hierarchical Model *with* Direct Communication Channels.

This model would retain the benefits of the previous model and therefore is able to help reduce human supervisor cognitive workload through autonomous coordination which incidentally is also performed with less computational power and lower communication overheads. Additionally, direct channels of communication between human and individual robots offers the human an opportunity to perform part of the coordination, further freeing up the computational and communication resources of individual robots for interaction with the human. Through this, situation awareness for the robots and the mission is expected to improve.
However, despite the expected benefits of applying the hierarchical model *with* direct channels of communication to the SHMRS, it is worth pondering the scalability of the proposed solution to optimizing MRD using a single human supervisor. As the collective's size is scaled upwards, the expected benefits of deploying such a SHMRS are not likely to remain.

7. Summary and Conclusion

In this paper, the evolution of interaction between humans and machines (followed by robots) has been described. In addition, the potential for the application of robots to USAR has also been explained. After which, an overview of the relevant work done by others has been provided through brief surveys pertaining to robot autonomy as well as MRD. The survey on robot autonomy has revealed the reasons behind the need for human supervision despite a robot's ability to function with autonomy. Within the survey for MRD, a number of key advantages for using multiple robots have been explained. The implications of robot autonomy and of the motivation for MRD suggest that it is highly advantageous for multiple robots to be deployed with fewer human supervisors than there are robots. Furthermore, safety considerations for multiple robot deployment in USAR indicate that the use of only a single human to supervise the deployed collective of robots is preferred.

However, before a SHMRS can be deployed, its designers must consider the key design attributes of size, composition and structure. Furthermore, designers must be aware of the issues inherent in the use of the SHMRS. These issues can be categorized into two sets: the technical set and the cognitive set. Firstly, IRC is difficult to achieve due to the technical issues of communication and computational demands on robots. The second set of cognitive issues consists of cognitive workload and reduced SA. This is because the supervision of the robots with a single human can induce high levels of cognitive workload and diminish SA. The quality of supervision provided by the human will degrade if cognitive workload is excessive. Consequently, good decision making and troubleshooting by the human becomes difficult.

It is expected that by altering each key design attribute, variations will also be seen in the issues with SHMRS deployment. Therefore, the projected effects that each attribute may have with the issues of IRC were first presented. The potential and feasibility of the hierarchical and horizontal structures for implementation on physical collectives of robots warrant their consideration in this paper. Robots using the hierarchical structure may not have to communicate with and model as many robots than those within a horizontal structure. However, the hierarchically structured collective is more prone to SPFs. In terms of collective size, it is felt that robots can be placed under high computational and communication loads as a collective's size increases. When collectives are heterogeneous, closer coordination between robots is necessary. Greater computational and communication burdens can be incurred when dissimilar robots are deployed.

The influence of each SHMRS design attribute on the quality of supervision provided by the human has been examined as well. With regards to collective size, it is felt that increasing the number of robots can result in high cognitive workload and poor SA. Therefore, the quality of supervision provided will be compromised. Additionally, while the consensus is that collective composition can affect attention switching between robots, the ease of supervising a homogeneous collective compared to that of supervising a heterogeneous collective is still debatable. Hence, research pertaining to collective composition and its influence on multiple-robot supervision is necessary. The effects of structure on the multiple robot supervision were presented by considering both the hierarchical and horizontal structures. Also, consideration is provided for whether or not robots are able to communicate between themselves as well as whether or not the human can communicate with robots individually. The investigation indicates that the allowing the human to communicate with individual robots when the collective is hierarchically structured can contribute to reducing communication and computational load for the robots. In addition, such a model of the SHMRS can prevent the human supervisor from being overly taxed in terms of cognitive workload while at the same time, providing adequate SA.

This paper has painted a picture of the considerations necessary for designing a SHMRS prior to its deployment. In particular, it has indicated that the structure of the collective within the SHMRS can be manipulated such that the technical and cognitive issues involved with practical deployment of the SHMRS can be mitigated. This also reduces the need for the size and composition of the robot collective to be compromised due to SHMRS deployment issues. However, despite the anticipated interactions between the key design attributes with both sets of issues, the SHMRS has to be physically implemented and tested such that empirical data may be used to validate the expected gains derived from manipulating the collective's structure.

8. References

Anderson, M., & Papanikolopoulos, N. (2008). Implicit cooperation strategies for multi-robot search of unknown areas. *Journal of Intelligent and Robotic Systems: Theory and Applications, Volume 53, Issue 4* , 381-397.

Arkin, R. C. (1995). Intelligent Robotic Systems. *Intelligent Systems , 10* (2), 6-8.

Arkin, R. C., Balch, T., & Nitz, E. (1993). Communication of Behavioral State in Multi-Agent Retrrieval Tasks. *IEEE International Conference on Robotics and Automation* (pp. 588-594). Los Alamitos, CA, USA: IEEE Computer Society Press.

Balch, T., & Arkin, R. C. (1994). Communication in Reactive Multiagent Robotic Systems. *Autonomous Robots, Volume 1, Issue 1* , 27-52.

Balch, T., & Parker, L. E. (2002). *Robot Teams: From Diversity to Polymorphism.* Natik, Massachusetts: A K Peters.

Barnes, M., Everett, H. R., & Rudakevych, P. (2005). ThrowBot: Design Considerations for a Man-Portable Throwable Robot. *International Society for Optical Engineering 5804*, (pp. 511 - 520).

Bellingham, J. G., & Rajan, K. (2007). Robotics in Remote and Hostile Environments. *Science , 318* (5853), 1098-1102.

Bond, A. H., & Gasser, L. (1988). An Analysis of Problems and Research in DAI. In A. H. Bond, & L. Gasser (Eds.), *Readings in Distributed Artificial Intelligence* (pp. 3-35). San Francisco, USA: Morgan Kaufmann.

Bond, A. H., & Gasser, L. (1988). *Readings in Distributed Artificial Intelligence.* Morgan Kaufmann Publishers.

Bradley, D. M., Silver, D., & Thayer, S. M. (2004). A Regional Point Descriptor for Global Topological Localizationin Flooded Subterranean Environments. *IEEE Conference on Robotics, Automation and Mechatronics*, (pp. 673 - 678).

Braynov, S., & Hexmoor, H. (2002). Quantifying Relative Autonomy in Multiagent Interaction. In H. Hexmoor, C. Cristiano, & R. Falcone, *Agent Autonomy* (pp. 43 - 56). Kluwer.

Bruemmer, D. J., & Walton, M. C. (2003). Collaborative Tools for Mixed Teams of Humans and Robots. In A. C. Schultz, L. E. Parker, & F. E. Schneider, *Multi-Robot SystemsL From Swarms to Intelligent Automata, Volume II.* Springer.

Burgard, W., Moors, M., Stachniss, C., & Schneider, F. E. (2005). Coordinated Multi-Robot Exploration. *IEEE Transactions on Robotics. Vol. 21, No. 3* , 376-386.

Burke, J. L., & Murphy, R. R. (2004). Human-Robot Interaction in USAR Technical Search: Two Heads Are Better Than One. *IEEE International Workshop on Robot and Human Interactive Communication*, (pp. 307-312). Kurashiki, Okayama Japan.

Cao, Y. U., Fukunaga, A. S., Kahng, A. B., & Meng, F. (1997). Cooperative Mobile Robotics: Antecedents and Directions. *Autonomous Robots 4(1)* , 7-27.

Casper, J. L., & Murphy, R. R. (2002). Workflow Study on Human-Robot Interaction in USAR. *IEEE International Conference on Robotics and Automation 2*, (pp. 1997-2003).

Casper, J., Micire, M., & Murphy, R. R. (2000). Issues in Intelligent Robots for Search and Rescue. *SPIE - The International Society for Optical Engineering*, (pp. 292 - 302).

Causse, O., & Pampagnin, L. H. (1995). Management of a multi-robot system in a public environment. *IEEE/RSJ International Conference on Robitc Systems*, (pp. 246-252). Pittsburgh, Pennsylvania, USA.

Chadwick, R. A. (2005). The impacts of multiple robots and display views: An urban search and rescue simulation. *Human Factors and Ergonomics Society*, (pp. 387-391).

Desai, M., & Yanco, H. A. (2005). Blending Human and Robot Inputs for Sliding Scale Autonomy. *International Workshop on Robots and Human Interactive Communication*, (pp. 537-542).

Dias, M. B., Kannan, B., Browning, B., Gil Jones, E., Argall, B., Freddie Dias, M., et al. (2008). *Sliding Autonomy for Peer-To-Peer Human-Robot Teams.* Pittsburgh, Pennsylvania: Carnegie Mellon University.

Dignum, V., & Dignum, F. (2005). Structures for Agent Organizations. *International Conference on Integration of Knowledge Intensive Multi-Agent Systems, KIMAS'05: Modeling, Exploration, and Engineering*, (pp. 215-220). Waltham, Massachesetts.

Dixon, S. R., Wickens, C. D., & Chang, D. (2005). Mission Control of Multiple Unmanned Aerial Vehicles: A Workload Analysis. *Human Factors, Vol. 47, No. 3* , 479-487.

Dudek, G., Jenkin, M., & Milios, E. (2002). A Taxonomy of Multirobot Systems. In T. Balch, & L. E. Parker, *Robot Teams: From Diversity to Polymorphism* (pp. 3 - 22). Natik, Massachusetts: A K Peters.

Elston, J., & Frew, E. W. (2008). Hierarchical Distributed Control for Search and Tracking by Heterogeneous Aerial Robot Networks. *IEEE International Conference on Robotics and Automation*, (pp. 170-175). Pasadena, CA, USA.

Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors: The Journal of the Human Factors and Ergonomics Society, Volume 37, Number 1* , 32-64.

Endsley, M. R., & Jones, W. M. (1997). *Situation Awareness Information Dominance & Information Warfare.* United States Air Force Armstrong Laboratory.

Endsley, M. R., & Rodgers, M. D. (1996). Attention Distribution and Situation Awareness in Air Traffic Control. *Proceedings of the 40th Annual Meeting of the Human Factors and Ergonomics Society*, (pp. 82-85). Santa Monica, California, USA.

Fong, T., Thorpe, C., & Baur, C. (2003). Multi-Robot Remote Driving With Collaborative Control. *IEEE Transactions on Industrial Electronics , 50* (4), 699-704.

Fox, M. S. (1979). *Organization Structuring: Designing Large Complex Software.*

Goldberg, D., & Matarić, M. J. (1997). Interference as a tool for designing and evaluating multi-robot controllers. *The National Conference on Artificial Intelligence*, (pp. 637-642).

Goodrich, M. A., Quigley, M., & Cosenzo, K. (2005). Task Switching and Multi-Robot Teams. *Multi-Robot Systems. From Swarms to Intelligent Automata , 3*, 185-195.

Haight, J. M., & Kecojevic, V. (2005). Automation vs. Human Intervention: What Is the Best Fit for the Best Performance? *Process Safety Progress 24* , pp. 45 - 51.

Halberstam, E., Navarro-Serment, L., Conescu, R., Mau, S., Podnar, G., Guisewite, A. D., et al. (2006). A Robot Supervision Architecture for Safe and Efficient Space Exploration and Operation. *10th Biennial International Conference on Engineering, Construction, and Operations in Challenging Environments*, (pp. 92-102).

Hart, S. G., & Wickens, C. D. (1990). Workload assessment and prediction. In H. R. Booher, *Manprint, an emerging technology, advanced concepts for integrating peorple, machine and organizations* (pp. 257-296). New York, USA.

Horling, B., & Lesser, V. (2005). A survey of multi-agent organizational paradigms. *The Knowledge Engineering Review Vol.19:4* , 281-386.

Kaber, D. B., & Endsley, M. R. (1997). Out-of-the-Loop Performance Problems and the Use of Intermediate Levels of Automation for Improved Control System Functioning and Safety. *Process Safety Progress 16 (3)* , 126-131.

Kadous, M. W., Sheh, R. K., & Sammut, C. (2006). Controlling Heterogeneous Semi-autonomous Rescue Robot Teams. *International Conference on Systems, Man, and Cybernetics*, (pp. 3204-3209). Taipei, Taiwan.

Koes, M., Nourbakhsh, I., & Sycara, K. (2005). Heterogeneous Multirobot Coordination with Spatial and Temporal Constraits . *National Conference on Artificial Intelligence 3*, (pp. 1292-1297).

Landauer, T. (1995). *The trouble with computers.* Cambridge, Massachusettes: MIT Press.

MacInnes, H. (2005). *International Mountain Rescue Handbook* (Fourth Edition ed.). London, United Kingdom: Frances Lincoln Ltd.

Matarić, M. J. (1992). Minimizing Complexity in Controlling a Mobile Robot Population. *IEEE International Conference on Robotics and Automation*, (pp. 830-835). Nice, France.

Maturana, F., Shen, W., & Norrie, D. (1999). MetaMorph: An adaptive agent-based architecture for intelligent manufacturing. *International Journal of Production Research, Volume 37, Issue 10* , 2159-2174.

Micire, M. J. (2008). Evolution and Field Performance of a Rescue Robot. *Journal of Field Robotics , 25* (17), 17-30.

Murphy, R. R. (2004a). Human-Robot Interactions in Rescue Robotics. *IEEE Transactions on Systems, Man, and Cybernetics - Part C: Applications and Reviews, Vol 34, No. 2* , 138-153.

Murphy, R. R. (2000). *Introduction to AI Robotics.* Cambridge, Massachusetts, USA: The MIT Press.

Murphy, R. R. (2004b). Trial by Fire: Activities of the Rescue Robots at the World Trade Center from 11 - 21 September 2001. *IEEE Robotics and Automation Magazine , 11* (3), pp. 50 - 61.

Murphy, R. R., Kravitz, J., Stover, S. L., & Shoureshi, R. (2009, June). Mobile Robots in Mine Rescue and Recovery. *IEEE Robotics and Autonomation Magazine , 16* (2), pp. 91 - 103.

Murphy, R., Blitch, J., & Casper, J. (2002, Spring). AAAI/RoboCup-2001 Urban Search and Rescue Events: Reality and Competition. *AI Magazine , 1*, pp. 37 - 42.

Nevatia, Y., Stoyanov, T., Rathnam, R., Pfingsthorn, M., Markov, S., Ambrus, R., et al. (2008). Augmented autonomy: Improving human-robot team performance in urban search and rescue. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, (pp. 2103 - 2108). Nice, France.

Nourbakhsh, I. R., Sycara, K., Koes, M., M, Y., Lewis, M., & Burion, S. (2005). Human-Robot teaming for Search and Rescue. *IEEE Pervasive Computing , 4* (1), 72-77.

Ong, K. W., Seet, G., & Sim, S. K. (2003). A Hierarchical Mixed Multi-Agent System. *The 4th IEEE International Confrence on Contol and Automation*, (pp. 644 - 648). Montreal, Canada.

Parker, L. E. (1998). ALLIANCE: An Architecture for Fault-Tolerant Multi-Robot Cooperation. *IEEE Transactions on Robotics and Automation 14 (2)*, (pp. 220 - 240).

Parker, L. E. (2000). Lifelong Adaptation in Heterogeneous Teams: Response to Continual Variation in Individual Robot Performance. *Autonomous Robots , 8* (3), 239-267.

Parker, L. E. (2008). Multiple Mobile Robot Systems. In B. Siciliano, & O. Khatib, *Springer Handbook of Robotics* (pp. 921-941). Springer.

Parker, L. E. (2003). The Effect of Heterogeneity in Teams of 100+ Mobile Robots. In A. Schultz, L. E. Parker, & F. Schneider, *Multi-Robot Sytems: From Swarms to Intelligent Automata* (pp. 205-215). Kluwer.

Parker, L. E., Kannan, B., Fu, X. Q., & Tang, Y. F. (2003). Heterogeneous Mobile Sensor Net Deployment Using Robot Herding and Line-of-Sight Formationgs. *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems*, (pp. 2488-2493). Las Vegas, Nevada, USA.

Riley, J. M., & Endsley, M. R. (2005). Situation awareness in HRI with collaborating remotely piloted vehicles. *Human Factors and Ergonomics Society 49th Annual Meeting*, (pp. 407-411).

Rosenfeld, A., Kaminka, G. A., Kraus, S., & Shehory, O. (2008). A study of mechanisms for improving robotic group performance. *Artificial Intelligence 172* , 633-655.

Sarter, N. B., Woods, D. D., & Billings, C. E. (1997). Automation surprises. In G. Salvendy, *Handbook of Human Factors and Ergonomics.*

Scerri, P., Pynadath, D., Johnson, L., Rosenbloom, P., Si, M., Schurr, N., et al. (2003). A Prototype Infrastructure for Distributed Robot-Agent-Person Teams. *International Conference on Autonomous Agents 2*, (pp. 433-440). Melbourne, Australia.

Schreckenghost, D., Fong, T. W., & Milam, T. (2008). Human Supervision of Robotic Site Survey. 6th Conference on Human/Robotic Technology and the Vision for Space Exploration.

Seet, G., Sim, S. K., Ong, K. W., Wong, C. Y., & Lau, W. S. (2009). Mixed Multi Agent Collaboration. *International Conference on Man-Machine Systems (ICoMMS).* Penang, Malaysia.

Sheridan, T. B. (2002). *Humans and Automation: Sysem Design and Research Issues.* Santan Monica, California, USA: John Wiley & Sons.

Sheridan, T. B. (1992). *Telerobotics, Automation, and Human Supervisory Control.*

Shiroma, N., Chiu, Y. H., Sato, N., & Matsuno, F. (2005). Cooperative task execution of a search and rescue mission by a multi-robot team. *Advanced Robotics , 19* (3), 311-329.

Stone, P., & Veloso, M. (2002). A Survey of Multiagent and Multirobot Systems. In T. Balch, & L. Parker, *Robot Teams: From Diversity to Polymorphism* (pp. 37-92). Natick, Massachusetts, USA: A.K. Peters.

Sugihara, M., & Suzuki, I. (1990). Distributed motion coordination of multiple mobile robots. *IEEE International Symposium on Intelligent Control*, (pp. 138-143). Philadelphia, PA, USA.

Sweeney, J. D., Li, H., Grupen, R. A., & Ramamritham, K. (2003). Scalability and schedulability in large, coordinated distributed robot systems. *IEEE International Conference on Robotics and Automation 3*, (pp. 4074-4079).

Tadokoro, S. (2009a). Earthquake Disaster and Expectation for Robotics. In S. Tadokoro, *Rescue Robotics: DDT Project on Robots and Systems for Urban Search and Rescue* (pp. 1-16). London, United Kingdom: Springer.

Tadokoro, S. (2009b). Summary of DDT Project, Unsolved Problems, and Future Roadmap. In S. Tadokoro, *Rescue Robotics: DDT Project on Robots and Systems for Urban Search and Rescue* (pp. 175-189). London, United Kingdom: Springer.

Wagner, A., & Arkin, R. (2004). Multi-robot communication-sensative reconnaisannce. *IEEE International Conference on robotics and Automation*, (pp. 4674-4681).

Wang, J., & Lewis, M. (2007). Human Control for Cooperating Robot Teams. *The ACM/IEEE Conference on Human-Robot Interaction -- Robot as a Team Member*, (pp. 9 -16).

Weil, S. A., Freeman, J., MacMillan, J., Jackson, C. D., Mauer, E., Patterson, M., et al. (2006). Design of a multi-vehicle control system: System design and user interaction. In N. J. Cooke, H. L. Pringle, H. K. Pedersen, & O. Connor, *Advances in Human Performance and Cognitive Engineering Research, Volume 7: Human Factors of Remotely Operated Vehicles* (pp. 223-236).

Werger, B. B., & Mataric, M. J. (2000). Broadcast of Local Eligibility: Behavior-Based Control for Strongly Cooperative Robot Teams. *Interantional Conference on Autonomous Agents*, (pp. 21-22). Barcelona, Spain.

Wickens, C. D., & Hollands, G. J. (2000). *Engineering Psychology and Human Performance 3rd Edition*. New Jersey: Prentice Hall.

Yamauchi, B. (2004). PackBot: A Versatile Platform for Military Robotics. *The International Society for Optical Engineering 5422*, (pp. 228 - 237).

Yanco, H. A., & Drury, J. L. (2002). A Taxonomy for Human-Robot Interaction. *AAAI Fall Symposium on Human-Robot Interaction*, (pp. 111-119).

Yanco, H. A., Drury, J. L., & Scholtz, J. (2004). Beyond Usability Evaluation: Analysis of Human-Robot Interaction at a Major Robotics Competition. *Human Computer Interaction* , *19*, 117-149.