In the proceedings of the American Nuclear Society Seventh Topical Meeting on Robots and Remote Systems, Augusta, GA. April 1997

Employing MASC to Control Multiple Mobile Robots

Julie A. Adams
Human Factors Laboratory
Eastman Kodak Company
Rochester, NY 14653
adamsj@kodak.com

Abstract

Over the last few decades roboticist have attempted to develop systems which are capable of successfully completing tasks in dynamic environments. Unfortunately, the community has not been as successful with their developments as was once projected. We have previously presented our approach to developing a semi-autonomous system which combines a human supervisor and robotic agents into a team. This paper will briefly review our Mediation Hierarchy theory, the Multiple Agent Supervisory Control (MASC) interface and our robotic test bed. We will also present a description of the experimentation which was performed with this system, review the results and discuss their implications for this system and future semi-autonomous systems.

1.0 Introduction

Over the last few decades many have attempted to create an autonomous robotic system which can execute tasks in a dynamic environment. While great strides have been made in this research area, we are still far short off creating a fully autonomous system. An example of our advances is the Dante II walking robot developed at the Robotics Institute at Carnegie Mellon University [Dante, 1995]. This robot was able to walk 200 meters into a volcanic crater. This is a major accomplishment in the field of robots while it also demonstrates the fragility of such systems. Upon reaching the crater floor, Dante II was unable to ascend autonomously due to "snow melt, mud, and freshly-exposed obstacles" and "was teleoperated during the entire ascent" [Dante, 1995].

One potential solution which in some senses addresses the short comings of autonomous robotic systems were telerobotic systems. Again while the community has achieved many advances in this area [Bejczy, 1990, Drotning, 1992, and Funda, 1992], there are still many restrictions with this type of robotic system. One of the largest restrictions to data has been the requirement that the human operator employ an input device to command every action of the robot. It is true that some researchers have experimented with incorporating autonomy into telerobotics [Hirai, 1989, and Papa, 1992], this two has not been entirely successful. These combined systems tend to incorporate the autonomy as a module. The human operator is typically unable to interact with the robot while it is in the autonomous module. This then again introduces the above referenced issue with autonomous systems.

While neither of these two approaches should be viewed as failures, the community may need to turns its focus to the area of semi-autonomous systems. In such a system the human acts as a supervisor who works to assist the robotic system. The supervisor has overriding control of the system, while the system is capable of working autonomously. The robot(s) and the human supervisor must be able to communicate with one another to fully understand each other's activities, and intentions. This communication must occur with all system levels. High level communications simply do not provide enough feedback. Aside from the creation of deeper communication channels between the human supervisor and the robot(s), we as humans are capable of accommodating dynamic environments. Therefore, the human's ability to interact and have ultimate control over the system is required to accommodate changes in the environment.

In order to achieve semi-autonomy new theories must be developed to model this type of system. Sheridan [Sheridan, 1992] has done much work in the area of supervisory control. This is the first step in the proper direction in order to achieve semi-autonomous systems. Supervisory control must be taken further to create a successful semi-autonomous system. The Mediation Hierarchy theory is one such theory which extends Sheridan's definition. The concept for the theory was to develop the ability for the human supervisor to be able to interact with all levels of a semi-autonomous system. The resulting system should be capable of completing feasible tasks. When such a system encounters a situation in which an autonomous robot would be unable to recover, the supervisor should be able to avert such a problem via the supervisory interactions.

The purpose of this paper is to review the Mediation Hierarchy theory, it's implementation in the Multiple Agent Supervisory Control (MASC) interface and our test bed. We describe the experimentation performed, provide a discussion of the results and their implications not only for this system but for future semi-autonomous robotic systems.

2.0 Mediation Hierarchy

The Mediation Hierarchy theory is composed of four levels: the *task, regulation, processing* and d*ata levels.* This section defines the Mediation Hierarchy. For a complete formulation of this theory please see [Adams, 1995].

In order for a mechanical system to understand what tasks it is to execute, it must be communicated to the system. The *task level* permits the supervisor to specify, either directly or via an planning mechanism, the actions an agent, or a group of agents, are to execute to complete an assigned task.

Man-Machine interfaces must provide some minimal interactions for the human to interact with the system. The *regulation level* combines these interactions, therefore, this level is composed of the *control, request* and *specification interactions.*

When a human interacts with a remote system there must exist a capability which permits the human to control the system. These controls may involve providing commands which instruct the robot where to go, the human must also maintain control over the robot's actions. Therefore, the *control interaction* provides the supervisor with the ability to control the agent's progress while executing a task either for the purpose of deterring or assisting progress.

The human supervisor must also be able to obtain information regarding the system. This permits the supervisor's ability to properly monitor the system and it's progress. The *request interaction* permits the supervisor to request information directly related to the current task.

There exist system actions which require the human supervisor to provide information before the system can begin processing. In these instances, the supervisor must possess a means of providing this information. The *specification interaction* permits the supervisor to specify required information prior the commencement of processing.

The idea which led to the development of the Mediation Hierarchy theory was that of providing the supervisor with the ability to interact with all levels of the robotic system. In order to assist the supervisor in obtaining information from the processes which comprise the system, a communication channel must be established. It is the *processing level* which permits the supervisor to interact and understand the system as well as the system to understand the supervisor. The *processing level* provides the ability to communicate with the various processes which compose the overall mechanical system.

It is a well known fact that mechanical devices can fail. When such a device fails, it would be preferable if the human supervisor could interact with the mechanical system to resolve the issue. For instance, if a sensing mechanism fails, then the preferred resolution is to rely upon another sensing modality. If the mechanical system is unable to determine this on it's own, the human supervisor must be able to instruct the system. The purpose of the *data level* is to permit the human operator to interact with the lowest system levels and ensure that correct information is communicated throughout the system.

3.0 MASC Test bed

The MASC test bed was the multiagents system developed by the University of Pennsylvania's General Robotics and Active Sensory Perception (GRASP) Laboratory. MASC was an integral part of the multiagent's system. The purpose of this system was to investigate the coordination and monitoring of multiagent systems for intelligent material handling [Multi, 1995]. The contributions of this work were an improved understanding of the fundamental problems underlying the control and coordination of multiple agents and well as the development of algorithms for intelligent exploration, organization and coordination of multiple agents.

The multiagents system is composed of four mobile platforms and the human-machine interface. The mobile agent bases are TRC Labmates. A six degree-of-freedom manipulator is mounted on each of the manipulatory agent bases.

The observation agents are equipped with various sensing modalities and a general purpose workstation (SPARC2). These agents utilize WINDATA™ radio ether net communications. These two agents are able to sense the environment and communicate information to the remainder of the system. The SensorBot is equipped with:

- A partial belt of sixteen Polaroid™ ultrasound and infrared sensors.

- A stereo camera pair.
- A light-striping device which employs elementary projective geometry to detect an object intersecting the light planes.

The ultrasound sensors are employed it to detect objects. As this sensing modality may be unreliable, [Mandelbaum, 1995], the infrared sensors are used to verify the ultrasonic readings. A stereo camera algorithm computes a localized correlation and extracts three-dimensional environmental information. The light-striping mechanism is employed to detect objects and then obtain the two-dimensional object information. The odometry and heading readings are utilized to monitor the agent as it moves through out the environment.

The VisionBot, is equipped with a stereo camera pair as well as a pan platform. The stereo pair are employed for visually guided obstacle avoidance, [Kosecka, 1996]. The pan platform is composed of a camera and a turn table and is used to track objects or other agents in the environment.

The manipulator agents are the PumaBot and the ZebraBot. The PumaBot is equipped with a Puma 260 Manipulator and the ZebraBot is equipped with a Zebra-ZERO Manipulator. These agents have no sensors other than force feedback and are therefore “blind”. They are primarily employed for the manipulation and relocalization of objects. The primary algorithms developed with the agents have been for testing the coordination of manipulation and locomotion as well as methods for redundant robots. The coordination of manipulation and locomotion considers the best configuration to carry objects and how to reconfigure the platforms in order to avoid obstacles or to pass through small passage ways, while maintaining the carrying grasp, [Yamamoto, 1994]. The redundant robot research focuses upon “the determination of joint motions for a given end effector displacement in kinematically redundant manipulators”, [Wang, 1995].

4.0 MASC System Overview

MASC is a human-machine interface designed in such a manner that it may be applied to any number or type of robotic agents. The individual robotic agents, their associated manipulators and processes may be controlled by the supervisor through MASC. Our objective was to create a semi-autonomous system which successfully completes assigned tasks.

The primary task of the human is to “supervise” the agent’s actions during task execution [Sheridan, 1992]. Through MASC, the human supervises the system while observing sensory data and images. The supervisor is permitted to assist the agents when requested and may assume control of an agent when necessary. Each agent is composed of multiple control and processing levels. MASC must permit interaction with these levels for the successful semi-autonomous execution of feasible tasks. This interaction will permit the supervisor to revise incorrect agent decisions and reconfigure the system after partial system failures.

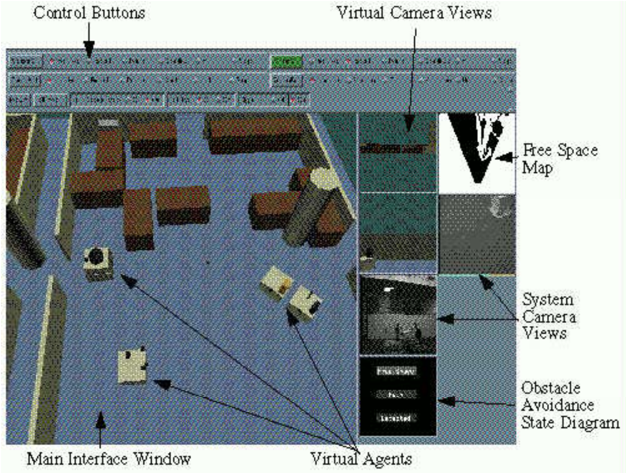

Figure 1: The MASC System Interface.

The MASC interface provides the supervisor with a three-dimensional environmental view. The main working window is the large window in Figure 1. The supervisor may rotate, zoom or translate the view to accommodate his or her current requirements. The right portion of the interface is employed to display two images of the graphical model, as well as images and processed data. The full interface window will permit the display of up to eight such windows. The images of the graphical model are updated from two virtual cameras, located in Figure 1 outside of the door labeled as the “South Doorway” and the second from the camera located on the VisionBot. The third and sixth windows displays are environment images. The fourth and bottom window displays a state diagram. The fifth window displays the free space map created by the visually guided obstacle avoidance process. The interface also provides a set of control buttons displayed on the top of the interface which are marked “Control Buttons” in the figure.

There are four system modes in MACS. The initialization mode permits the human supervisor to specify any agents and their associated processes which may be required for a task. The supervisor is not required to initialize all agents which may be necessary for a task before task execution begins. The Supervisor may edit the graphical model, add objects to the world model from the overlay of system cameras or display various system data on the interface. The purpose of this system mode is to prepare the MASC system for task execution.

The exploration mode is a teleoperation mode. In this mode the human may teleoperate an agent to create locomotion commands. While working within this mode, there exists a "phantom" agent which is similar in principle to Bejczy et al. "phantom" agent [Bejczy, 1990]. The purpose of this agent is to inform the supervisor of the actual position of the real agent. As the supervisor teleoperates the virtual agent to create commands, its position no longer corresponds to the position of the actual agent. Thus, the "phantom" agent is updated with the actual agent's heading and odometry readings.

The navigational mode is an autonomous mode. It permits the supervisor to drive the agents based upon path plans created by one of the two MASC system path planners or the way point specification method. The local R-geodesic path generator, [Multi, 1995 and Wang, 1993], is employed to plan short paths in well-known environments. The set of global path planners are managed as a path plan server at Stanford University. This server is described in [Becker, 1994] and is composed of potential field and cell decomposition planning methods described in [Latombe, 1991]. This planner considers all objects in the environment. The third option is a way point specification ability which requires the operator to choose the desired way points for the agent.

The replay mode permits the supervisor to replay the task execution within the last five minutes. This option does not permit raw image data replay, but does replay all other sensing modalities data displays. It replays a single virtual agent's actions as well as any combination of the active agents. If a particular agent was inactive during the specified time frame the supervisor is notified. If only an inactive agent is chosen for replay, no replay is provided. If other agents are also specified, the human is notified of a particular agent's inactivity and the replay continues with the other specified agents. This particular option is helpful for diagnosing uncertain situations.

5.0 Experimentation

The motivation for the Mediation Hierarchy's development was to develop the ability for a human supervisor and semi-autonomous system to complete feasible tasks. Therefore, proof of the Mediation Hierarchy theory entails executing various tasks until the agents require supervisory assistance, then demonstrating the supervisor's ability to assist and correct the problem through the MASC system interface, followed by the agent's ability to continue with the task execution to completion. During the experimental development, we found the multiagents system as a whole was not sufficiently sophisticated to fully test the Mediation Hierarchy theory. Also, it was determined that difficult experiments would require a vast amount of training for a novice user. Thus we developed the experiment to employ only a portion of the MASC system capabilities. The experiment was designed to provide data encompassing the MASC system's usability and preliminary feedback on the Mediation Hierarchy. The purpose of this Section is to provide the experimental design methodology.

5.1 Tasks

The participants were required to carry out three tasks: the single agent, the two agent and the four agent tasks. These tasks were executed twice, sequentially during two sessions.

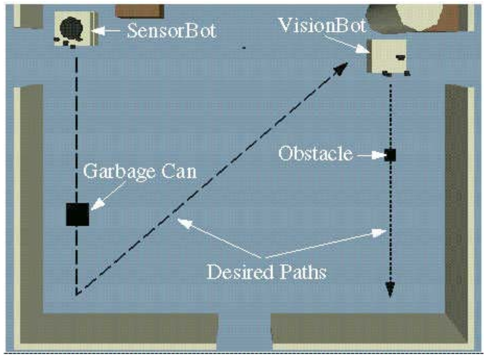

Figure 2. The single and two agent tasks.

The single agent task required the participant to drive the SensorBot parallel to the southwest wall into the corner. Once the agent obtained the position in the corner, the participant was required to turn the agent and drive it diagonally across the room (the desired path is the dashed line in Figure 2). A tall garbage can was placed approximately two thirds of the distance between the agent's initial position and the corner (see Figure 2). The participants were required to drive the agent around the obstacle and into the corner. This requirement stems from the fact that the SensorBot is not equipped with an on-board obstacle avoidance procedure.

The two agent task required the participant to drive the SensorBot as described above while simultaneously driving the VisionBot parallel to the Southeast wall. The desired path for the SensorBot is the dashed line while the VisionBot's desired path is the dotted line in Figure 2. There was an obstacle placed in the VisionBot's path. The

VisionBot's obstacle avoidance process was to be employed to automatically avoid the obstacle.

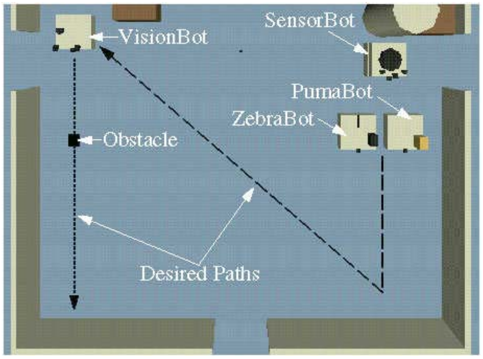

Figure 3: The Four Agent Task.

The four agent task required the participants to simultaneously control all four agents. The VisionBot was positioned as in Figure 3. The participants were instructed to drive the agent along the Southwest wall into the corner, (the dotted line in the Figure). There was an obstacle placed in front of the agent which the obstacle avoidance process was to avoid. The other three agents were positioned as in Figure 3. The two manipulatory agents were in a side-by-side configuration ahead of the SensorBot. The SensorBot's purpose was to observe the manipulatory agent's actions. The participants were instructed to control the manipulatory agents in a combined control method[1] for as much of the task execution as was feasible. All three agents were to be driven along the Southeast wall into the corner, (the dashed line in the Figure). When this position was obtained, the agents were to turn and move diagonally across the room to the goal position marked in Figure 3.

5.2 Participants

The participant group was composed of thirteen computer literate members of the University of Pennsylvania community. The Participants were novice users with mobile robots and most had minimal training with a graphical user interface. The participant's ages ranged between seventeen and thirty-three years and their educational backgrounds ranged from some high school to doctoral level education. There were three female participants and ten male.

All participants received identical training which included a system description. The participants were paid a predetermined amount for the entire experiment. Payment was not contingent upon completion of the experiments or the amount of time required.

5.3 Apparatus

The MASC system version employed for these experiments was pared down from the complete system. This version permitted the participants to use all four agents and their respective sensing modalities. The locomotion command generation method was teleoperation and the autonomous locomotion methods were not employed. The robotic agents employed are those described in Section Three.

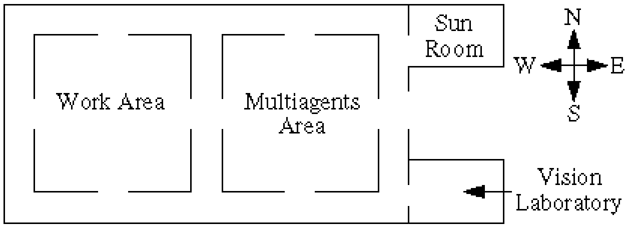

Figure 4: The GRASP Laboratory layout.

Figure 4 displays the GRASP Laboratory floor plan. The curtains surrounding the laboratory's eastern portion (Multiagents Area), in which the experiments would occur, were closed. Also, doorways into this section of the laboratory were blocked off. The closed curtains restricted the participant's ability to view the area between trails. The MASC system was run on an Silicon Graphics Indigo[2] with 96 megabits of memory in the Vision laboratory. This is the room to the east of the main laboratory in figure 4. The participants were physically unable to view the multiagent work area during the experiments.

6.0 Discussion

This section discusses some of the experimental results and their implications for this particular system and semi-autonomous robotic systems in general.

Task	Completion		Incomplete		Accident		System Issues	
Session	1	2	1	2	1	2	1	2
Single Agent	25	25	0	0	0	0	1	1

---

[1] Combined control permits the supervisor to create a single command to be executed simultaneously by both manipulator agents.

Two Agent	20	25	2	0	2	1	2	0
Four Agent	20	19	0	0	2	3	4	4
Total	65	69	2	0	4	4	7	5

Table 1. The number of completions per task.

The results related to the task completions are displayed in Table 1. This information is presented by task per session. As the table illustrates, there were 134 successful completions out of 156 trials. This is an 85.9% completion rate. The system issues column relates to instances when issues arose with the actual robots. A recurring problem which was unresolvable was the lose of communications between one or both of the manipulatory agents. When instances of this nature occurred and the participant was working with more than one robot, they were instructed to continue the task with the remaining robot(s). If the ten “completions” with system issues are included, the completion rate raises to 92.3%. The completion rates were much higher than anticipated and displayed a good result.

Session	One	Two
Single Agent	368.8	272.6
Two Agent	411.6	343.1
Four Agent	651.7	458.4

Table 2. Mean completion times in seconds.

The average task completion time is displayed in Table 2 by task and session. As anticipated, the completion times increased as did the number of active robots. The other interesting aspect was that the completion times decreased, on average, between sessions. Statistical analysis of the results showed that the relationship between the single and two agent tasks was significant with a ($f(1,50) = 6.04$, $P = 0.018$) while the relationships between the single and the four agent tasks as well as the two and four agent tasks were insignificant. The analysis also showed that the relationship between sessions for the single agent task was significant, ($f(1,24) = 12.4$, $P = 0.002$) while the two and four agent tasks between sessions were insignificant. The data appears to be significant when viewing the average completion times but during testing it was apparent that the participants each had a personalized manner in which to command the robots. This lead to large variations in the completions times for each task. The insignificance of the results is likely the cause of this large variation between completion times.

The number and types of errors that the participants made with the interface are presented in Table 3. As this interface combines both button menus and direct manipulation and does not require the user to type commands, it was anticipated that errors should be reduced as opposed to other interface metaphors which require the user to employ the keyboard. In all, there was a total of 37 errors. The most frequent error was attempting to work with an agent different from the one specified. This error primarily occurred during the four agent task. The method for choosing an agent to work with was found to be too time consuming and convoluted. The participants would have preferred a more direct manner of choosing a robot.

Error	Occurrence
Working with un-chosen agent	20
Re-initializing an agent	7
No Camera Server Process	3
No ZebraBot Robo Process	3
No SensorBot Robo Process	2
No Path Following Process	1
No Ultrasound Process	1
Total	37

Table 3. The interaction errors by frequency of occurrence and category.

The remaining errors are all related to the participants initialization process. The method of displaying menus as a level of buttons may have hindered the participants ability to mentally recall what processes they had initialized and which they had not. It should be noted that the largest number of errors occurred in the four agent task with a total of 23 errors while the single agent and the two agent tasks each had 7. As well, the number of errors fell between sessions. The first session had a total of 27 errors while the second session had 10.

The question that remains is: what does this imply about the Mediation Hierarchy and the participants ability to interact with the agents via the MASC interface? First and foremost, it shows that when supplied with a usable interface, novice users are able to control multiple robots in an environment in which they must rely upon the robot’s sensing capabilities in order to complete assigned tasks. While it is true that the interface version employed for these experiments was a reduced version, we were still able to obtain pertinent information regarding the participants ability to complete tasks.

The results which relate to the participant’s ability to employ the MASC system for such a task are positive. Their reactions during training were that of disbelief that they would be capable of controlling these robots to complete tasks. Upon completion of their first experimental session, most participants felt comfortable with their ability to control the robots and found the tasks fun. Upon completion of the second testing session, the usability results were favorable, based upon the subjective questionnaires. In fact, many participants had desired different more challenging tasks during their second session. This was interesting from the perspective that entering the experiments we were unsure how many of the participants would even be able to successfully complete the simplest tasks.

The experiments raised some essential usability issues with MASC. Many participants felt they became lost while working within the various system levels of the button menus. As well, they desired the ability to reorganize the interface layout to represent their preferences. Some participants did not like to have to look to their right to view the smaller images and a preferred layout would have been to place them along the top of the window. As well, within the area in which images were display, the participants wished to dynamically re-arrange the images employing a drag and drop method. The general consensus throughout was that MASC was a fun, and easier to use than anticipated.

Provided that the system version employed for the experiments was a reduced version, we were unable to fully test the Mediation Hierarchy theory. We were provided information pertinent to further the theory development and implementation in the MASC system. Essentially, the participants activities with the interface substantiated the *task* and r*egulation levels*. The participants were capable of specifying the tasks for the agents, commanding the agents, requesting information and pre-specifying information. The participant's ability to properly instruct the agents on the task execution based upon the MASC interface and the sensory information provided is an indication that they were able to work with the system at the *task level*. Their ability to understand the creation of locomotion commands and other commands such as stop and pause indicate their ability to succeed with the r*egulation level's command interaction.* The participants successfully requested the sensory information they felt pertinent to the current task, this is the definition of the *request interaction*. As well, the participants were able to initialize the system prior to beginning the tasks. While this is not exactly what we had in mind for the *specification interaction*, it does resemble the tasks required of the user for this interaction type.

While these experiments did not formally test the Mediation Hierarchy's *processing* and *data levels*, we were able to draw conclusions as to the participants desire and need to interact at these levels with the agents via the interface. A suggestion from those participants with backgrounds in image processing was the desire to be able to modify the image quality or the frame rate of the images they received. They also expressed a bit of frustration with their inability to fully understand what the VisionBot was doing when it encountered a corner. They wanted to be able to question the process, to determine what it thought it was seeing and why it chose a particular resolution to the problem. They also wanted to be able to override this behavior. This is exactly the type of interaction for which we have defined the *processing level*. While this does not in any manner indicate that the participants would have been able to understand the agent's behavior, it does uphold the concept of providing such a capability.

Another frustration for the participants was their inability to obtain a true situation awareness of the agents positions during execution. This is primarily due to the type of wheels these mobile platforms which have a high incidence of slippage in their readings. Due to this error, it was sometimes the case that the participants were not completely certain of the robots location. They could rely upon the actual images they received but they desired the ability to localize the agent relative to the environment and then reset it's reading accordingly. This frustration was found in almost every participant. This represents the type of interactions for which the Mediation Hierarchy's *data level* was designed to address.

So what does this imply about the use of the Mediation Hierarchy, and more importantly, man-machine interaction in the realm of robotics. Twenty five years ago the idea was that robots would be able to work autonomously to assist and relieve humans of various duties from very simply tasks to those which are unsafe. While the robotics community has made many in roads, we still have not obtained a robot which can autonomously solve a problem successfully, in a dynamic environment. We do know that humans are very capable of dealing with dynamic situations, therefore, it is even more important that we combine the practices of Human Factors Engineering with Robotics to develop systems which can further man kind.

7.0 Conclusions

As was discussed, while there have been many wonderful achievements in the field of robotics, the goal of a fully autonomous robot which is able to survive in a dynamic environment still eludes us. It is now time for the robotics community to embrace the concept of semi-autonomous robotic systems. We should not abandon research in the areas of autonomous or teleoperated systems but should expand our field to the development of semi-autonomous systems. In order to follow endeavors with these types of systems and be successful in their development, we must combine the research areas of Human Factors engineering with the Robotics discipline. This will ensure the development of systems which permit the human supervisor to efficiently interact and assist the robotic system.

The Mediation Hierarchy is a theory which expands traditional supervisory control in an attempt to create a successful semi-autonomous robotic system. The MASC interface is also the first step at combining the disciplines of Human Factors and Robotics.

We have reviewed the Mediation Hierarchy, the MASC interface and our multiple robotic test bed. While the Mediation Hierarchy may not provide the ultimate solution to this problem it is a step in the proper direction.

We have also presented a description and discussion of experiments employing our system. While these experiments did not fully test our theory, we were able to develop conclusions as to the concept of such a system. Admittedly, there are problems associated with this work, and there still remains more integration and testing in order to fully substantiate this theory, but the results obtained to date appear to reflect a positive direction for the field of robotics.

8.0 Bibliography

[Adams, 1995] Adams, J.A. *Human Management of a Hierarchical System for the Control of Multiple Mobile Robots*. Ph.D. thesis, University of Pennsylvania, 1995.

[Becker, 1994] Becker, C. *Internet Path Planner Server Protocol Manual.* Stanford University, Jan., 1994.

[Bejczy, 1990] Bejzy, A., Kim, W., and Venema, S. The phantom robot: Predictive displays for teleoperation with time delay. In *Proceedings of the IEEE International Conference on Robotics and Automation,* pages 546 - 551. IEEE, 1990.

[Dante, 1995] Wettergreen, D., Pangels, H., and Bares, J. Behavior-based Gait for the Dante II Walking Robot. In *IEEE/RJS International Conference on Intelligent Robots and Systems*. IEEE/RJS, Aug., 1995.

[Drotning, 1992] Drotning, W., Christensen, B., and Thunborg, S. Graphical Model Based Control of Intelligent Robot Systems. In *IEEE Control Systems Magazine*, v. 12, num. 4, Pages 13 - 18, Aug., 1992.

[Funda, 1992] Funda, J. Lindsey, T. and Paul, R. Teleprogramming: Toward Delay-Invariant Remote Manipulation. In *Presence: Teleoperators and Virtual Environments*. v. 1, num. 1, pages 29 - 44, Winter, 1992.

[Hirai, 1989] Hirai, S. and Sato, T. Multi-Level Manual and Autonomous control Superposition for Intelligent Telerobot. In the *NASA Conference on Space Telerobotics*, pages 131 - 140, 1989.

[Koseka, 1996] Koseka, J.. *Discrete Event Systems for Autonomous Mobile Agents*. Ph.D. thesis, University of Pennsylvania, 1996.

[Latombe, 1991] Latombe, J-C. *Robot Motion Planning*. Kluwer Academic Publishers, 1991.

[Mandelbaum, 1995] Mandelbaum, R. *Sensor Fusion for Mobile Robot Localization, Exploration, and Navigation*. Ph.D. thesis, University of Pennsylvania, 1995.

[Multi, 1995] Adams, J.A., Bajcsy, R., Kosecka, J., Kumar, V., Mandelbaum, R., Mintz, M., Paul, R., Wang, C., Yamamoto, Y., and Yun, X. Cooperative material handling by human and robotic agents: Module development and system synthesis. In *IEEE/RJS International Conference on Intelligent Robots and Systems*. IEEE/RJS, Aug., 1995.

[Papa, 1992] Papanikolopoulos, N.P., and Khosla, P.K. Shared and Traded Telerobotic Visual Control. In *IEEE International Conference on Robotics and Automation*, pages 878 - 885. IEEE, May, 1992.

[Sheridan, 1992] Sheridan, T. *Telerobotics, Automation, and Human Supervisory Control*. MIT Press, Cambridge, Mass., 1992.

[Wang, 1993] Wang, C. Local Path Planner. University of Pennsylvania. Personal Communications.

[Wang, 1995] Wang, C. *The performance of Repeatable Control Schemes for Redundant Manipulators*. Ph.D. thesis, University of Pennsylvania, 1995.

[Yamamoto, 1994] Yamamoto, Y. *Control and Coordination of Locomotion and Manipulation for a Wheeled Mobile Robot*. Ph.D. thesis, University of Pennsylvania, 1994.