Designing Interfaces for Multi-User, Multi-Robot Systems

Adam Rule
University of Washington
acrule@uw.edu

Jodi Forlizzi
Carnegie Mellon University
forlizzi@cs.cmu.edu

ABSTRACT
The use of autonomous robots in organizations is expected to increase steadily over the next few decades. Although some empirical work exists that examines how people collaborate with robots, little is known about how to best design interfaces to support operators in understanding aspects of the task or tasks at hand. This paper presents a design investigation to understand how interfaces should be designed to support multi-user, multi-robot teams. Through contextual inquiry, concept generation, and concept evaluation, we determine what operators should see, and with what salience different types of information should be presented. We present our findings through a series of design questions that development teams can use to help define interaction and design interfaces for these systems.

Categories and Subject Descriptors
A.m. [**Miscellaneous**]: Human Robot Interaction

General Terms
Design, Human Factors.

Keywords
Multi-robot operators, interface design, information presentation, design research.

1. INTRODUCTION
The thought of robots interacting with humans has been around for millennia. One of the oldest written accounts, found in the Daoist Liezi, tells of a West Zhou Dynasty (1066BC-771BC) engineer who created a humanoid so real that it winked at the king's concubines [19]. As the tale goes, the king had to see the machine dismantled before being satisfied that it was not human. The idea of simultaneously controlling groups of robots also has historical roots. As early as the 1920s and 30s Nicola Tesla envisioned teams of 50, even 100 robots being controlled at once by one or more operators [9, 19].

Those dreams are now becoming reality as the use of autonomous robots in organizations continues to increase steadily. For example, autonomous robots are used in search and rescue and field mining missions. Semi-autonomous robots are also in use in hospitals where they assist with straightforward tasks that can reduce the work of staff, for example, collecting linens and cafeteria trays or providing surveillance. These situations are ripe for deploying several robots simultaneously.

In Multi-User multi-Robot System (MURS) scenarios, both the number and complexity of tasks operators face drastically increase over single-user, single-robot cases. A single operator may divide their available cognitive resources over the control of many systems, performing both repetitive and non-standard tasks [28]. In addition, dynamic task allocation, where robots change their behavior in relation to the environment and other robots, drastically increases the complexity of the task at hand [27]. The more robots that are added, the more complexity increases. Multiple operators introduce the problems of team dynamics and fluid coordination.

Researchers have already begun to empirically study how situation awareness changes as the number of robots in a team increases [23] and classify the unique elements of situation awareness in HRI [13]. There is also a mature body of work on how to create a salient display, manage awareness, and command attention [1, 7, 11, 13, 15, 16, 25, 32, 36]. However, what operators want their attention directed towards and how this need changes with respect to operator role and context is unaddressed. Open research questions include: What do operators want to see? Does this change relative to an operator's role on the team? With what importance, or salience, should information be presented? What kind of information can remain hidden unless requested?

To begin answering these questions, this paper presents a design investigation to understand how interfaces should be designed to support multi-robot, multi-operator teams. Through contextual inquiry with expert operators, concept generation, and concept evaluation, we determine what operators should see, and with what salience different types of information should be presented. We present our findings through a series of design questions that development teams can use to help define interaction and design interfaces for these systems.

2. RELATED WORK
We began by conducting a literature review spanning the disciplines of psychology, human factors and robotics, focusing on work related to autonomy of robots, control of operators, situation awareness, and information presentation. This research led to grounding definitions on which to base our design investigations.

2.1 HRI and collaboration
Coordination and communication within HRI systems is more complex than in HCI. HRI systems are more dynamic, exhibit autonomy, operate in dynamic, real-world environments, exhibit different interaction behaviors, and may work with multiple systems simultaneously [15, 35]. For these reasons, solutions from the domains of HCI and software engineering, while applicable, cannot be transferred wholesale into HRI, especially in MURS applications. Take, for example, the case of a search-and-rescue robot that comes across a conscious person trapped in rubble. The interaction of the trapped person with the robot is vastly different from that which the robot operator has with the robot, though both interactions occur simultaneously and in a very real sense the two people are collaborating towards the goal of safely removing the victim from the surrounding rubble. This large difference in operator and victim roles is a situation unique to HRI and

http://doi.org/10.1145/2157689.2157705

highlights the need for unique design. Some research integrating approaches from HCI and HRI offers a taxonomy of design considerations for HRI, including team composition, amount of required interaction, decision support, and space-time location [41].

Additionally, literature in human-human collaboration and team dynamics can be relevant to understanding how MURS interfaces are best designed. This is because types of teams, structures of teams, and numbers of humans and robots involved are important constructs in collaboration. There has been a great deal of research on the basis of differentiation and discrimination between people in one's group and those outside of it. Factors such as intergroup competition, similarity, and status affect the degree of group differentiation; group members can be evaluated more extremely or perceived as threatening [6, 5]. In addition to in-group favoritism, out-group members may be differentiated along a number of social dimensions [33].

In the domain of HRI, recent work on team collaboration with robotic systems has shown the importance of group leaders in developing common understanding and translating team status information across members [24]. Other research conducted with mobile robots in a desert field site showed that different levels of robot autonomy can drastically change the flow of communication and the creation of common ground among team members [39].

Much more needs to be known about collaboration in MURS to better support the interface design of these systems. Systems that are intended for groups may be perceived in many ways, due to the uneven distributions of benefits of the system among members of an organization [21, 31]. Varying levels of autonomy may influence how different people perceive the robot, control, and trust the robot [20]. Social effects such as resistance to use of the technology may result, and workarounds are often not captured in the design [22].

2.2 Factors affecting information presentation

Understanding human-robot collaboration is also important in MURS interface design. Two factors affect how information gets presented in the MURS interface: what operators see and control, and what the robots are controlling.

2.2.1 What operators see and control

Early research defined control as the human monitoring of a complex system with intermittent use of an automated agent [30]. Implicit in this definition is the fact that the supervised system may periodically act autonomously based on sensor data [37]. Subsequent research acknowledged the importance of the human in the interaction loop. Collaborative control was defined as humans and robots engaging in dialog to exchange ideas, resolve differences, perform tasks, and achieve goals [15]. Here, human and robot are seen as equals in collaboration; the human is an imprecise, limited source of information for the robot.

While these are important grounding principles, other research has explored the role of the operator relative to the information he or she will need to see. Essential roles include supervisor, operator, mechanic, peer, bystander, mentor, and information consumer [19, 34, 35]. Collectively, these definitions begin to map out the design space for MURS interfaces.

2.2.2 What the robots are controlling

Because autonomy levels change during dynamic robot control, an open design research question is how those levels may be communicated to human operators. An early model from HCI, which describes ten levels of autonomy, is a good starting point (Table 1) [37]. This model focuses on autonomy in decision making, not specific action, and while the actions of HCI and HRI differ, the process of decision making remains largely unchanged.

Table 1: Ten Levels of Autonomy used in designing interfaces for MURS systems (after [37]).

Level	Description
1	Human does the whole job up to the point of turning it over to the computer to implement.
2	Computer helps by determining the options.
3	Computer helps determine options and suggests one, which human need not follow.
4	Computer selects action and human may or may not do it.
5	Computer selects action and implements it if human approves.
6	Computer selects action, informs human in plenty of time to stop it.
7	Computer does whole job and necessarily tells human what it did.
8	Computer does whole job and tells human what it did only if human explicitly asks.
9	Computer does whole job and tells human what it did and it, the computer, decides they should be told.
10	Computer does whole job if it decides it should be done, and if so tells human, if it decides they should be told.

Alternative control methods include dialogue systems where the robot presents limited information to the operator when seeking guidance at a barrier or decision point [16]; policy systems that create rules for how and when robots are permitted to act autonomously [4]; waypoint systems to assist in navigation or sweeping an area [18]; and playbooks, which give high-level commands to clusters of robots [17].

2.3 Situation Awareness

With dynamic levels and modes of command, situation awareness, which defines how much attention is commanded to the interface at any given time, is of utmost importance. Inefficient attention allocation can lead to operator inefficiencies, significant robot wait time and under-utilization [30]. Thus it is important to consider how operator workload and attention are allocated to maximize system efficiency. Work has already been done to assess how interface design affects both workload and situation awareness in multi-robot control, although for single operators [23].

A straightforward evaluation of perception, comprehension and projection as the stages of situation awareness [14] can be extended for use in HRI. In the domain of HRI, Drury has provided a useful framework of HRI awareness [13], which includes five components (Table 2). All of Drury's HRI awareness components aside from robot-robot awareness can be mediated by a successful interface.

2.4 Salience of information

To design to support situation awareness, we can look to the field of visual design, which is informed by research in many areas: cognitive science, human factors, and semiotics, among others. Research has explored identification and/or search speeds of text

[10], color [8], high-symbolism icons [10], and combined text and high-symbolism pictures [3]. Other research explores how individual design features holistically contribute to a design. Color as a design feature has been shown to result in faster visual search times compared to design features such as size, brightness, and geometric shape [10]. Another study showed that the discriminating effects of color lessen with learning [8]. Other research confirms that text can be processed faster than pictures [10], that concrete icons are more quickly interpreted than abstract ones [29], and that complex icons lead to slower search and identification times (whether the symbols are learned or unlearned by the viewer) [29].

Table 2: Five components of situation awareness in HRI as may affect the design of the interface (after [13]).

Component	Description
Human-Robot	The understanding that the humans have of the locations, identities, activities, status and surroundings of the robot. Further the understanding of the certainty with which humans know the aforementioned information.
Human-Human	The understanding that the humans have of the locations, identities and actions of their fellow human collaborators.
Robot-Human	The robots' knowledge of the humans' commands needed to direct activities and any human-delineated constraints that may require command noncompliance or a modified course of action.
Robot-Robot	The knowledge that the robots have of the commands given to them, if any, by other robots, the tactical plans of the other robots and the robot-to-robot coordination necessary to dynamically reallocate tasks among the robots if necessary.
Humans' Overall Mission Awareness	The humans' understanding of the overall goals of the joint human-robot activities and the measurement of the moment-by-moment progress obtained against the goals.

A body of design research exists has examined interfaces for safety-critical situations, such as flying and driving. One dual-task study showed that finding items on a secondary display may not significantly hinder a primary task [38]. Other research has focused on icon-function relationships, exploring the relationship between visual design, complexity, and interpretation of particular designs. However, this work does not explore how a display might best convey information for a situation when attention is divided among a number of tasks. This leaves open the exploration of designs and renditions that may use design features like contrast, color, and abstraction to quickly convey information.

From the literature, we can understand that the goal of minimizing complexity on a visual display is important. We also know that bottom-up search happens within 100ms of attending to a visual display. Therefore, designs with high visual salience should employ contrast in size and color as design variables. The details of particular design features relative to one another are important. Collectively, these factors — collaboration, information presentation, situation awareness, and salience of information — must be taken into account when designing interfaces for MURS.

3. RESEARCH THROUGH DESIGN PROCESS

In order to understand the needs of MURS operators and create actionable knowledge to support the design of MURS interfaces, we pursued a research through design process [42], which combined contextual inquiry, affinity diagramming, concept generation, and rapid concept evaluation.

Design research moves from specific to general to specific; observations from a select group of users are used to generate broad design recommendations, which are subsequently used to generate domain-specific design guidelines. In this vein, we chose to study specialized robot operators who not only control robots, but aid in their development. Observations of these operators were used to abstract general principles about MURS control, which were then applied specifically to inform the design of MURS interfaces. The following sections describe the stages of our design process.

3.1 Contextual Inquiry

Contextual inquiry (CI) is the process of interviewing and observing users in context for the purpose of understanding the nature and challenges of their work [2]. CIs allowed us to see first hand the tasks that users complete and the problems they face.

Our team conducted two CI sessions. We interviewed and shadowed participants with the goal of understanding how individuals and groups control multiple robots at one time. Our goal was to elicit information about how to create interfaces that facilitate simultaneous control of more than one robot at a time, allowing for changes in robot autonomy. Our process was to largely observe teams interacting with their robots, asking questions along the way to more fully understand what they were doing. We focused on aspects of interface use, aspects of coordination, and when and how breakdowns occurred and were mitigated.

The first session, lasting about 1.5 hours, was conducted with the primary operator of a social robot. The robot delivers snacks to faculty and researchers in an academic building at Carnegie Mellon University. The robot makes use of head and body movement, an adaptable array of LEDs representing a mouth, and speech synthesis to carry out social interaction. The robot operator must load the robot's tray with snacks, localize the robot's navigation system, drive the robot to its work area using a wired joystick, control robot navigation and dialog using a custom interface, and interact dynamically with the programmer and faculty member responsible for the development of the robot.

The second CI session, lasting about two hours, was conducted with three operators of a search-and-rescue robot. The robot has the form factor of a snake and is constructed of 16 interchangeable modules. The primary goal of the robot is to traverse environments too dangerous for human exploration and identify persons trapped in rubble. The operators make use of a custom interface and standard game controller to drive the robot. Control configurations map components of the game controller to parameters in the control interface, which elicits a response from the robot. The robot has several methods of locomotion, and each method has its own unique set of parameters and control mappings, though many parameters are shared among modes.

In both sessions, operators were asked to describe each element of their current robot control interface and demonstrate normal operation of the robot. For the social robot, our team was able to observe as the robot was used during normal operation, delivering snacks. For the search-and-rescue robot we were unable to observe the robot being operated in a disaster area, though a number of contrived obstacles were used to demonstrate its operation. Along with interface explanation, the CIs included a semi-structured interview with questions related to robot operation, breakdowns, and coordination with other robot operators. Each session was video recorded and notes were taken both during the sessions and upon review of the video recordings.

3.2 Affinity Diagramming

Notes from the CIs were used to create an affinity diagram [26]. In this process, themes are drawn from clusters of notes that had been generated by reviewing the data. Using this process, we uncovered twelve situation awareness questions that MURS operators must answer for robust operation.

To create an affinity diagram, statements or ideas that occurred in the CIs were written onto individual Post-It notes and placed on a common board. Once all team members had exhausted their notes, the team collaboratively organized the notes into categories according to similarity. All team members were able to add and remove category labels, rename categories, and move notes between categories at will. This process ensured consensus on note coding. From this process, our team derived 12 themes in MURS operation: coordination, human roles, possible actions, system status, robot status, directing attention, hardware, recovery and debugging, control modes, customizable and adaptive control, control levels, and robot roles.

These themes were further labeled, following from Drury's HRI Awareness model, as human-human, human-robot, robot-human, robot-robot, and overall mission awareness issues [13]. Organizing our notes using this framework allowed us to look for connections between data points, inconsistencies, and areas for further consideration. For example, none of our note categories fit under the category of overall mission awareness, so we re-examined our notes and made sure to consider overall mission awareness in our future work.

3.3 Information Sought by Operators

After several iterations of affinity diagramming, a set of questions was produced that, when answered, lead to situation awareness. These twelve situation awareness questions distill what operators need to know for robust robot operation (Table 3).

Since it is likely that robot-robot awareness is mediated a level away from the operator's interface, the two questions regarding robot-robot awareness were not considered in the remaining design phases. The final ten relevant awareness questions were used to guide the evaluation of MURS interface elements and will henceforth be refereed to as "situation awareness categories".

3.4 Information Salience Framework

With the categories of MURS situation awareness defined, we sought to understand how an interface should display information related to each of these categories. To do so, we modified the bottom seven levels of Sheridan's Ten Levels of Autonomy (Table 1) to describe how much a computer or its human operator bears the load of deciding what information to display (Table 4) [37]. The top three levels were excluded as they involve acting without informing, and information cannot be displayed without informing the user.

Table 3: Essential situation awareness questions that must be communicated in a MURS interface. These were used in our concept generation matrix.

Awareness Type	Awareness Question
Overall Mission	How does any one person know what the overall mission is and the current progress being made towards its completion?
Human-Human	Who is taking what goals and which should I (the operator) cover?
	Who needs assistance or who can help me?
	What robot, area, or interface element is my teammate referring to? (upon asking for assistance)
Human-Robot	What mode, state, and environment is the robot in and how will this affect my commands?
	Which robot needs my attention?
	What is each robot accomplishing?
	What can I accomplish with this robot?
Robot-Human	What should I (the robot) be doing?
	What should I (the robot) do if I hit a barrier?
Robot-Robot	What are other robots doing?
	How can I (the robot) help other robots, or get help?

3.5 Speed Dating

In order to study how MURS interfaces should guide operator attention, we used our situation awareness categories (Table 3) and our information salience framework (Table 4) to guide a rapid concept evaluation (or speed dating) session. Speed dating is a design process used to validate user needs and narrow a design solution space quickly from a large number of design ideas to a few key interventions [12]. Speed dating has two stages: needs validation to focus on broad opportunity areas, followed by a stage of user enactments to understand how contextual elements may aid or hinder a solution's implementation.

Table 4: Final Information Salience Levels used in the concept generation matrix.

Level	Description
1	The computer offers no assistance, the human must call for all information explicitly.
2	The computer offers a complete listing what information may be displayed.
3	The computer narrows the selection down to a few.
4	The computer selects the information to be displayed and the human may or may not display it.
5	The computer displays the selection if the human approves.
6	The computer allows the human a restricted time to veto before automatic display.
7	The computer selects and displays information it deems necessary.

Since we are still at a broad stage of understanding MURS operator needs, we chose to carry out only the first stage of speed dating: needs validation. This stage makes use of a matrix of design issues to cover a broad range of user needs and understand which are most critical to address and in what manner they should be addressed. The matrix is composed when two scales of design intervention are placed orthogonal to each other. Each cell of the matrix is filled with a design solution with attributes corresponding to the cell's position on each of the design scales.

In our case, we placed our ten situation awareness categories orthogonal to three levels chosen from the information salience framework in Table 4. We chose to use only three of the salience levels (the 2nd, 4th and 7th) creating 30, rather than 70, design matrix cells, so as not to fatigue our evaluation participants. We considered these three levels as low, middle, and high information salience. For low salience, the interface simply lists what information may be displayed and may be commanded to display this information. In the medium level, the interface selects which information should be displayed, but does not actually render complete information until commanded to do so. In high salience, the interface both selects and displays information as it deems necessary.

Traditionally, each cell of a design matrix is then populated with a storyboard detailing how a design solution corresponding to the correct values on each design scale could meet user need. We chose instead to populate each cell with an interface wireframe corresponding to the cell's category of awareness and level of salience. In this manner, 30 wireframes were created showing information related to each of the ten situation awareness categories presented at each of the three information salience levels. Figure 1 shows an example wireframe of Overall Mission information shown with high salience.

Speed dating sessions were carried out with the CI participants. In order to provide some context beyond personal experience, participants were first shown two basic MURS scenarios, one involving a set of five nurses controlling five robots to dispense medication and a second involving two operators controlling four heterogeneous terrain surveying robots with various abilities. Participants were then shown interface elements in sets of three, grouped by the situation awareness categories they addressed. Participants were asked to explain how important information related to that category was and identify which of the three information salience levels was the most appropriate. Since the appropriate salience for a specific situation awareness category may change based on any number of factors, participants were asked to answer in detail and explain intervening factors accordingly. At the end of the session, participants were asked to rank the three most important and three least important situation awareness categories. Video recordings were taken of the speed dating sessions and notes were taken upon further review of the recordings.

In the next section, we share the results of the concept generation and speed dating sessions. We focus in particular on situation awareness and varying levels of information salience.

4. Results of Concept Generation

The main goal of our speed dating sessions was to understand how interfaces should support MURS. Specifically, we wanted to understand what operators should see, what kind of information salience should be assigned, and what kind of information can be kept hidden unless requested.

4.1 No one-to-one mapping

We wanted to understand if our participants, acting as operators, wanted a specific level of salience for information relating to human-human awareness and a different level for human-robot awareness, and so on with the other HRI awareness components. What we found was that there was no one to one mapping of these

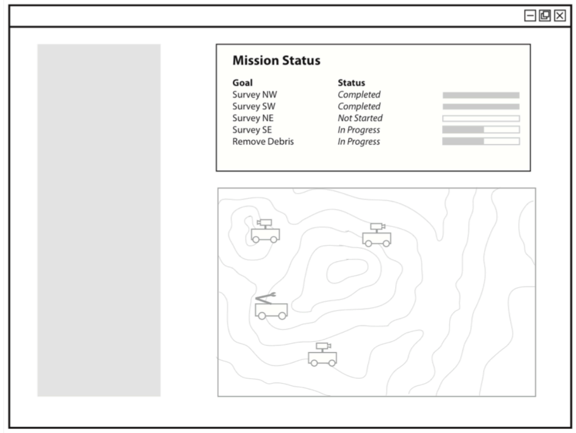

Figure 1: Example interface element used in the speed dating session.

types of information. Instead, robot operators often spoke of wanting different levels of salience to display the same information in different situations. Additionally, they desired a holistic view of the task allocation

4.2 Important vs Urgent

Important information did not always warrant the highest level of information salience. For example, robot operators felt that it was very important that they understand which goals are delegated to them for completion, but they wanted this information displayed with low salience. They felt that a listing of goals would be referenced at the start of operation to develop a clear model of what needed to be accomplished, and only periodically throughout operation, and thus should not take up screen space for most of the operating session.

Additionally, operators wanted urgent information displayed with the highest salience. Cases where inaction would lead to human or robot harm were deemed urgent. For example, operators wanted a high level of display salience in the case of a robot encountering an unresponsive patient in an assistive care setting, but low salience when the robot was experiencing low battery. In the first case, action must be immediate to ensure patient safety, but in the second case, there is a period of time in which action may be taken without consequence. A general hierarchy of information, from most urgent to least urgent, was human safety first, followed by robot failures, robot status, individual operator goals, and finally, overall team goals.

4.3 Operator Roles

There were also some interesting findings related to operator roles. Our participants felt that operators who were in charge of the mission should have more access to information, and subsequently more control over its modification. This is interesting, because it naturally follows that they will have less overall situation awareness of a mission at any given time.

In the next sections, we describe specific findings from each of the situation awareness questions used in the concept generation matrix.

4.4 Overall Mission

Operators thought the middle and high levels of display salience were appropriate for displaying overall mission goals. If goals were to be dynamic and reassigned often, operators wanted the interface to inform them of this, but if goals were more static, they felt it was appropriate to have the middle level of salience where they could decide to check on goals periodically. The operators felt that an understanding of the overall mission scope was more important for team leaders to understand rather than each individual operator.

4.5 Goal Delegation

Operators felt an understanding of which goals were delegated to them was important and desired for a low or medium level of salience. They stated that they would likely look at their set of goals only periodically, or when significant tasks had been completed. However, they felt that information about robot status was more important and would want to be able to collapse or navigate away from goal information so their screen space could be devoted to robot information.

4.6 Assistance

Operators wanted the middle level of display salience for notifications of other operators needing assistance. They did not want to be interrupted from a current task they were completing, but also did not feel they would have time to be continually monitoring the status of other operators. Some kind of alert was desired that would notify them of another operator needing help. A higher level of display salience was desired for larger operator networks where a central leader, or system AI, would assign tasks, similar to how job bids are delegated to service workers.

4.7 Area Reference

Operators desired for a low level of display salience in conveying information about which robot, interface element, or part of the environment another operator may be referring to when asking for assistance. They felt a shared interface element that could be manually edited by each operator, such as an overview map that could be tagged, was most appropriate.

4.8 State Awareness

Operators wanted a medium level of display salience for robot state awareness. They wanted to be able to select which aspects of robot information to view at any one time. They also wanted the ability to “zoom-out” to basic overview information when they were doing general control, and then click into a more specific view when trouble shooting or manually controlling a single robot. The ability to resize, and turn on or off the display of robot sensor data was essential to the operators.

4.9 Robots Needing Attention

Operators saw a need for both medium and high levels of display salience in being notified of which robots needed their attention next. For critical status updates such as a robot mechanical or software failure, or cases requiring immediate attention, such as a patient being unresponsive, operators wanted a high level of display salience and their attention to be directed to the problem area immediately. For less vital issues, such as a robot going idle, operators wanted a medium level of salience, being informed of the problem but able to choose whether or not to respond.

4.10 What Robot is Accomplishing

Operators wanted low or medium levels of display salience for understanding what a robot was accomplishing currently. As distinct from critical issues discussed in the section above, current progress is information that operators want to be able to access, but do not want the system to draw their attention to unless they decide.

4.11 Robot Capability

Operators wanted a low level of display salience for understanding the capabilities and commands that could be given to a robot. They felt that the understanding of what a robot can do should be developed during some period of training and reference materials should be on hand, possibly in the form of a wiki, detailing robot capabilities. A medium level of display salience was considered appropriate for MURS where operators would not have access to training beforehand or where very few actions could be carried out by the robots.

4.12 Robot Control

For actually commanding a robot to complete a task, operators desired a medium level of display salience. Tabs or buttons that are pervasively displayed and can be used to switch between robots or levels of control were mentioned. Operators did not want for the display of these elements to be dynamic or brought to their attention except in critical situations as detailed in Section 4.4.8.

4.13 Robot Idle Control

In order to increase robot productivity, some operators let robots decide what to do if they become idle or finish an assigned task. Communicating to the robot the order in which tasks should be completed, or “safe” tasks that they can accomplish while waiting for further input is sometimes beneficial. Operators felt that a low level of display salience, such as a secondary window where task priorities could be assigned pre-mission would be appropriate for predictable tasks, but that more dynamic missions and environments would require a medium level of display salience. A robot becoming idle was not deemed a critical alert and operators felt a high level of display salience that could direct attention to an idle robot was inappropriate.

4.14 Summary

While operators wanted a holistic sense of their mission, and wanted a clear communication of goal state, the most important information for the robot operators to know was the delegation of individual goals, the state of their robots and the environment around them, and which robot needs immediate attention. The least important information to know was what a robot’s capability,

Table 5: Summary of levels of display salience for components of the MURS interface.

Level	Description	Level	Description
H	Convey errors related to human safety.	M	Convey state awareness.
H	Communicate robot failures.	M	Convey robot idle control.
H/M	Convey that robot needs attention.	M/L	Convey team goals.
H/M	Convey a sense of the mission.	M/L	Convey need for robot assistance.
M	Convey robot status.	M/L	Convey goal delegation.
M	Convey operator goals.	M/L	Convey what robot is accomplishing.

which other operators need assistance, and telling a robot what to do when it is idle and waiting for attention. Operators stated that increased robot autonomy, increasing numbers of robots under an operator's control, and increasing task critically would necessitate higher levels of display salience, directing their attention more to critical decisions.

5. DISCUSSION

Our investigation into interfaces to support MURS teams led to a number of interesting findings regarding how information should be presented to support team leaders and collaborators. To implement these ideas, the growing literature on interface design patterns in HCI can be used as reference material [40]. Modular tabs, collapsible panels, and multiple windows, all lend themselves well to low levels of display salience. Accordions, list inlays, and alert lists similarly work well for medium levels of display salience. Responsive disclosure, movable panels, and pop-up windows all work well for high levels of display salience.

While the needs of operators, and therefore, interface elements change over time in a MURS system, we can put forth some general suggestions for whether or not information should be presented with high, medium, or low salience. These are presented in Table 5. In addition, based on the results of our design investigation, we can offer several design suggestions for those who are designing and implementing MURS systems. These design recommendations follow from Table 5, and include the following:

First, high salience should be reserved for communicating information such as errors related to human safety, errors related to robot safety, items needing immediate attention, and information about the dynamics of the mission.

Second, medium or low salience can be used to present information on the dynamics of team context. This includes information like robot status, operator goals, operator state of awareness, and team goals.

Third, low salience on should be used for operators controlling the mission. These people will want high control of the information that they are interacting with.

Overall, operators desire a holistic view of the system at any time, and it is important that they can exert control over all types of information if needed. For example, override features should be presented so operators can either intervene with or ignore high salience warnings if needed. In the end, operators must be trusted to override information that the system may think is urgent, but actually isn't.

6. LIMITATIONS

While we present some general conclusions about MURS interfaces, this study has a number of limitations. First, the small sample population used for the CIs and speed dating mean that more extensive study is needed to provide greater certainty in the results presented here. Design is a general set of skills which applies in many situations, and we feel confident that our findings are readily generalizable. Second, the CIs were conducted with single robot operators who collaborate with other operators rather than MURS operators. This deficiency is partly a result of there being few MURS currently in operation, and those in development are still at the stage of individual robot development rather than team development. As MURS become more widespread, operators will be able to provide richer context knowledge for understanding the design of MURS interfaces. Despite these limitations, we believe the study we completed shows the unique nature of MURS, provides several useful frameworks for future research, and lays a foundation for future MURS interface development.

7. CONCLUSION

Robotic systems have matured substantially, but little is known about how to design the interaction for multi-robot, multi-user systems. In this paper, we present a design investigation leading to several design recommendations for interfaces for MURS systems. While operators desire a holistic view of interaction with the system, our findings show that low, medium and high salience can successfully be used to convey different aspects of the interaction. We hope that our findings will help to advance the field of multi-robot interaction and operator collaboration.

8. ACKNOWLEDGMENTS

This material is based upon work supported by ONR and the National Science Foundation under Cooperative Agreement EEC-0540865. We thank the QOLT REU program for their assistance. We also thank the participants who assisted with our research. Special thanks to Steffi Paepcke, Vikram Chatterji, and Eva Shon for their assistance in the research process.

9. REFERENCES

[1] Baker, M., Casey, R., Keyes, B., and Yanco, H.A. (2004). Improved Interfaces for Human-Robot Interaction in Urban Search and Rescue. Proceedings of 2004 IEEE Transactions on Systems, Man, and Cybernetics. New York, NY: IEEE Press, 2960-2965.

[2] Beyer, H., Holtzblatt, K. (1998) Contextual Design: defining customer-centered systems. Morgan Kaufmann.

[3] Booher, H.R. (1975) Relative comprehensibility of pictorial information and printed words in proceduralized instructions. Human Factors, 17, 266-277.

[4] Bradshaw, J., Feltovich, P., Jung, H., Kulkarni, S., Taysom, W., and Uszok, A. (2004) Dimensions of adjustable autonomy and mixed-initiative interaction. in Agents and Computational Autonomy: Potential, Risks, and Solutions, (M. Nickles, M. Rovatos, and G. Weiss, eds.) Springer, Berlin/Heidelberg, 17-39

[5] Branscombe, N.R., Wann, D.L. and Noel, J.G. (1993). In-Group or Out-Group Extremity: Importance of the Threatened Social Identity. PSPB, 19/4, 381-388.

[6] Brewer, M. B. (1979). In-Group Bias in the Minimal Intergroup Situation: A Cognitive-Motivational Analysis. Psychological Bulletin, 86/2, 307-324.

[7] Bruemmer, D., Frew, D., Nielsen, C., and Walton, M. (2007). Mixed-initiative control for collaborative countermine operations. IEEE Transactions on Robotics. IEEE Transactions on Systems, Man, and Cybernetics, 35/4, 438-449.

[8] Carney, C., Campbell, J.L. and Mitchell, E.A. (1998) In-vehicle display icons and other information elements: Literature review, Federal Highway Administration.

[9] Cheney, M. (1981) Tesla: Man out of Time. Touchstone, New York.

[10] Christ, R.E. (1984) Research for evaluating visual display codes: An emphasis on colour coding. In Easterby, R. and Zwaga, H. eds. Information design: The design and evaluation of signs and printed materials, John Wiley & Sons, 209-228.

[11] Crandall, J.W., Goodrich, M.A., Olsen, D.R., and Nielsen, C.W. (2005). Validating Human-Robot Interaction Schemes in Multi-Tasking Environments.

[12] Davidoff, S., Lee, M., Dey, A., Zimmerman, J. (2007) Rapidly exploring application design through speed dating. In Proceedings of UbiComp '07, 429-446.

[13] Drury, J., Scholtz, J., Yanco, H. (2003) Awareness in Human-Robot Interactions. In Proceedings of IEEE Conference on Systems, Man and Cybernetics, Washington, DC.

[14] Endsly, M. (1995) Towards a Theory of Situation Awareness in Dynamic Systems. Human Factors, 37 (1), 32-64.

[15] Fong, T., Thorpe, C., Baur, C. (2001) Collaboration, Dialogue, and Human-Robot Interaction. In 10th International Symposium of Robotics Research, Lorne, Victoria, Australia.

[16] Fong, T., Thorpe, C., Baur, C. (2003) Multi-Robot Remote Driving with Collaborative Control. IEEE Transactions on Industrial Electronics 50/4, 699-704.

[17] Goodrich, M., McLain, T., Anderson. J., Sun, J. Crandall, J. (2007) Managing Autonomy in Robot Teams: Observations from Four Experiments. In proceedings of Human Robot Interaction '07, Arlingon, VA, 25-32.

[18] Goodrich, M., Olsen, D., Crandall, J. Palmer, T. (2001) Experiments in Adjustable Autonomy. In Proceedings of IJCAI Workshop on Autonomy, Delegation and Control: Interacting with Intelligent Agents.

[19] Goodrich, M., Schultz, A. (2007) Human-Robot Interaction: A Survey. Foundations and Trends in Human Computer Interaction, 1(3), 203-275.

[20] Groom, V., Takayama, L., Ochi, P., Nass, C. (2009) I am my robot: the impact of robot building and robot form on operators. In proceedings of HRI '09, La Jolla, CA, 31-36.

[21] Grudin, J. (1988) Why CSCW applications fail: Problems in the design and evaluation of organizational interfaces. In Proc. of CSCW'88, 85-93.

[22] Grudin, J. (1994) Groupware and Social Dynamics: Eight Challenges for Developers. Comm. of the ACM, 37(1): 92-105.

[23] Humphrey, C., Henk, C., Sewell, G., Willians, B., Adams. J. (2007) Assessing the scalability of a multi-robot interface. In proceedings of HRI 2007, New York. 239-246.

[24] Jones, H., Hinds, P. (2002) Extreme work teams: using SWAT teams as a model for coordinating distributed robots. In proceedings of CSCW 2002, New Orleans, 372-381

[25] Kaupp, T. and Makarenko, A. (2008). Decision-Theoretic Human-Robot Communication. Proceedings of HRI08. New York, NY: ACM Press, 89-96.

[26] Kawakita, Jiro. (1986) The KJ Method: Seeking order out of chaos. Choukouron-sha, Tokyo (In Japaneese).

[27] Lerman, C., Jones, C., Galstyan, A., and Mataric, M.J. (2006). Analysis of Dynamic Task Allocation in Multi-Robot Systems. International Journal of Robotics Research, 25:225, 225-241.

[28] Lewis, M., Wang, H., Chien, S.-Y., Velagapudi, P., Scerri,, P. and Sycara, K. (2011). Process and Performance in Human-Robot Teams. Journal of Cognitive Engineering and Decision Making, 5/2, 186-208.

[29] McDougall, S., de Bruijn, O. and Curry, M. (2000) Exploring the effects of icon characteristics: Concreteness, complexity and distinctiveness. The J. of Experimental Psychology: Applied, 6(4), 291-306.

[30] Mitchell, P. J., Cummings, M. L. (2005) Management of Multiple Dynamic Human Supervisory Control Tasks. In Proceedings of the 10th International Command and Control Research and Technology Symposium.

[31] Mutlu, B., Forlizzi, J. (2008) Robots in organizations: the role of workflow, social, and environmental factors in human-robot interaction. In proceedings of HRI 08. New York, 287-294.

[32] Park, B., and Rothbart, M. (1982). Perception of Out-Group Homogeneity and Levels of Social Categorization: Memory for the Subordinate Attributes of In-Group and Out-Group Members. Journal of Personality and Social Psychology, 42/6, 1051-1068.

[33] Parasuraman, R., Galster, S., Squire, P., Furukawa, H., and Miller, C. (2005). A Flexible Delegation-Type Interface Enhances System Performance in Human Supervision of Multiple Robots: Empirical Studies with Robo-Flag. IEEE Transactions on Systems, Man, and Cybernetics, 35/4, 481-493.

[34] Scholtz, J. (2002) Creating Synergistic CyberForces. in Alan C. Schultz and Lynne E. Parker (eds.), Multi-Robot Systems: From Swarms to Intelligent Automata. Kluwer.

[35] Scholtz, J. (2002) Theory and Evaluation of Human-Robot Interactions. In Proceedings of the 38th Hawaii International Conference on System Sciences.

[36] Scholtz, J., Antonishek, B., and Young, J. (2005). Implementation of a situation awareness assessment tool for evaluation of human-robot interfaces. IEEE Transactions on Systems, Man, and Cybernetics, 35/4, 450-459.

[37] Sheridan, T., Verplank, W. (1978) Human and computer control of undersea teleoperators. Technical report, Man-Machine Laboratory, Massachusetts Institute of Technology, Cambridge, MA.

[38] Somervell, J., McCrickard, D.S., North, C. and Shukla, M. (2002) An evaluation of information visualization in attention-limited environments. In Proc. VISSYM 2002.

[39] Stubbs, K., Hinds, P.J., and Wettergreen, DD. (2007). Autonomy and Common Ground in Human-Robot Interaction. IEEE Intelligent Systems, 22/2, March/April, 42-50.

[40] Tidwell, J. (2011) Designing Interfaces. O'Reilly, Beijing.

[41] Yanco, H.A. and Drury, J.L. (2002). A Taxonomy for Human-Robot Interaction. AAAI Technical Report FS-02-03.

[42] Zimmerman, J., Forlizzi, J., and Evenson, S. (2007). Research through design as a method for interaction design research in HCI. In proceedings of CHI 07. New York, NY: ACM Press, 493-502.