Human-Agent Teams for Extended Control Operations in Space

D. Schreckenghost, P. Bonasso, C Martin, T. Milam, C. Thronesbery

NASA Johnson Space Center, TRACLabs
1012 Hercules, Houston TX, 77058, USA
ghost@ieee.org, r.p.bonasso@jsc.nasa.gov, cmartin@traclabs.com, tmilam@traclabs.com, c.thronesbery@jsc.nasa.gov

Abstract

At Johnson Space Center (JSC) we have developed automated control agents to reduce human workload by automating routine tasks such as system reconfiguration and FDIR. We have evaluated these control agents for extended periods during ground tests of regenerative crew life support systems that represent analogs of crew systems operating in space. In this paper we describe our experience with these ground tests at JSC, including our use of the automated control agents during these tests. We describe how human-agent teams can coordinate their activities when distributed physically. We also describe how human teams can maintain both situation and system performance awareness of the automated control agent when they are distributed physically and communicating asynchronously.

Introduction

The complexity and dynamics of manned space missions has traditionally been handled by providing 24/7 Earth-based tending of operations in space. While such tending is viable for shorter missions like Space Shuttle, it becomes costly for extended stay missions such as International Space Station (ISS) or exploration missions. At JSC we have developed automated control agents for the purpose of reducing human workload by automating routine tasks such as system reconfiguration and FDIR. We have evaluated these control agents for extended periods during ground tests of regenerative crew life support systems that represent analogs of systems operating in space. As a result of these evaluations, we have demonstrated that automated control agents can reduce human workload significantly (Bonasso, et al., 2002; Schreckenghost, et al., 1998).

We also learned that using automated control agents will change the humans' job in operations. The crew is able to spend less time doing control tasks and is freed up to do more science. Flight controllers are able to do tasks in addition to mission control. Yet both are responsible to maintain awareness of control situation, to intervene in control when the unexpected occurs, and to perform control tasks not easily automated. Additionally, for extended operations, they must manage long-term changes in control performance. It is important to support both crew and ground controllers in performing this job ([a]Schreckenghost, et al., 2002). To do this, we have developed human support software called the Distributed Collaboration and Interaction (DCI) environment.

In this paper we describe our experience with crew life support ground tests at JSC, including our use of automated control agents during these tests and our development of the DCI environment as a result of these tests. We describe how human-agent teams can coordinate their activities when distributed physically. We then describe how human teams can maintain both situation and system performance awareness of the these agents when they are distributed physically and communicating asynchronously. We use the DCI system to illustrate some of the capabilities needed for this concept of operations for human-agent teams.

Experience in Extended Space Operations

Crew Life Support Ground Tests at JSC

In 1997, we deployed an automated control agent at JSC for product gas transfer (PGT) during the Lunar/Mars Life Support Technology Project (LMLSTP) Phase III ground test. The PGT agent was responsible to maintain the health of the wheat plants in the plant chamber. To accomplish this it controlled two hardware systems in the plant chamber, the carbon dioxide (CO2) system and the oxygen (O2) system. The CO2 system injected CO2 from the crew into the plant chamber to feed the wheat in the chamber. The O2 system removed O2 produced by this wheat for use by the crew occupying a separate chamber.

The PGT agent also managed the transfer of O2 and CO2 between the plant chamber and the crew chamber. Normally, it transferred CO2 from the crew chamber to the plant chamber, and O2 from the plant chamber to the crew chamber. For a portion of the test, however, the PGT agent redirected O2 flow to a solid waste incinerator every 4th day. This required accumulating O2 for approximately one day prior to the incineration, instead of providing O2 to the crew. The CO2 produced by incineration was used by the plants in place of the CO2 from the crew. Additionally, every 21 days one quarter of the wheat crop was replanted, to ensure that O2 production was consistent over the course of the test.

The Phase III test was conducted in the Fall, 1997. The PGT agent operated continuously 24/7 for 74 days. Human shifts supporting the PGT agent required 6-8 hours weekly to support sensor calibration and archival of data logs, with an additional 6 hours for each incineration and 3 hours for each harvest (Schreckenghost, et al., 1998).

Product Gas Transfer Automated Control Agent

The PGT agent was developed using the 3T control architecture (Bonasso, et al, 1997). 3T was developed for autonomous robots, but was first successfully applied to the control of life support systems for extended operations during the Phase III ground test. Since that time it has been used continuously for over a year to control a regenerative Water Recovery System(WRS) at JSC (Bonasso, et al., 2003). The 3T architecture is an example of a layered control architecture. It consists of three tiers of control processing that operate in parallel:

- Deliberative Planner: hierarchical task net planner to manage activities (1) with resources or temporal constraints, or (2) with multi-agent coordination
- Reactive Sequencer: reactive planner to encode operational procedures that can be dynamically constructed based on situational context
- Skill Manager: traditional closed loop control

This approach is designed to handle the uncertainty inherent in complex domains. Control commands flow down through the hierarchy and feedback flows back up through the hierarchy to close the control loop. If a command fails to succeed at any level, it can initiate a repair action (e.g., replanning at the deliberative level, selection of an alternative sequence at the reactive level). Each layer operates at a different time constant, allowing high speed controllers at the lowe level of the architecture to operate in parallel with the slower, deliberative algorithms at the high level.

Support for Team Interaction

During the Phase III test and subsequent WRS tests, we provided a variety of support for human interaction during extended operations with the automated control agents. A key lesson was the need to assist humans in interleaving control support with their "other" job. Because the automated control agent was capable of managing the PGT or the WRS most of the time, the human support task was part–time. For these ground tests, engineers needed to interleave their control support with other tasks on a daily basis by working operations from their office. Alternately, this can manifest itself as shift work, where a person works closely with the system for a period of time and then shifts to on-call work for a period of time. In either case, it becomes apparent that the current approach of dedicated operational support will change. It becomes necessary to provide assistance in detecting and notifying people about operational events affecting their assigned job. Notification should avoid unnecessary interruptions when possible. Such event notification can be complicated by communication latencies (e.g., notice urgency). Situation summarization becomes essential, with easy access to the details of the situation on demand.

As mentioned previously, we expect to see in the future some Earth-based operations supported from offices. There is also the possibility of multiple sites in space for exploration missions in the future (e.g., multiple sites on a planet; an orbiting site and a planetary site). Additionally, as we design missions that include both humans and robots, we will see multiple agents conducting concurrent operations from many locations. Taken together, this concept of distributed team operations will require distributed task management among humans and agents. It will require strategies to adjust autonomy to ensure that concurrent manual and autonomous operations never interfere and, in some cases, are closely coordinated. It will require capabilities for fair allocation of resources and task responsibilities. This will include strategies to ensure that tasks are reassigned when contingencies prevent assigned agents from handling them.

One way to meet these requirements is to develop a mediating agent that assists humans in interacting with automated control agents ([a]Schreckenghost et al, 2002). We have developed the DCI environment that provides such a mediating agent, called an ARIEL agent. ARIEL agents are assigned to each human member of the operational team. The ARIEL agent can interact with both control agents and other ARIEL agents. Each ARIEL agent performs services for its user. The selection of which services are loaded is configurable for a group. The services available in the ARIEL agent include: (1) Task Management, (2) Event Notification, (3) Command Authorization, and (4) Location Tracking. The DCI environment also provides these capabilities: (1) Crew Activity Planner (CAP), (2) Event Detection Assistant (EDA), and (3) Situation Summarization.

We currently are deploying the DCI System in the Advanced Water Lab at JSC for evaluation during a ground test using the Post Processing System (PPS) of the WRS. The PPS is a water "polisher" used to bring relatively clean water to within potable limits. It removes the trace inorganic wastes and ammonium in recycled water using a series of ion exchange beds and removes the trace organic carbons using a series of ultraviolet lamps. The DCI system will provide notification, schedule management, and situation logging for control engineers.

Coordinated, Distributed Operations

Distributing operations among humans and automated control agents using the same systems and resources requires that their activities be coordinated. Such coordination necessitates that each member of the team is aware of other team member's activities and their effects on system state, especially when transitioning between manual and automated activities. Such *activity monitoring* also is useful for the human in supervising the automated control agent. For *concurrent activities*, it is important to provide strategies to ensure that ongoing manual and automated activities do not interfere with each other. This can be preplanned for all agents or handled reactively by either adjusting the autonomy level of the agent or by reassigning the agent's role. It is undesirable, however, to achieve this by turning off the automated agent, since this complicates returning to full autonomy later. Beyond avoiding interfering actions, agents should ensure that adequate resources are available to perform activities. In resource limited environments like space, this can require *resource management* planning to ensure that consumable resources are not over-utilized and that regenerable resources are produced in adequate quantities. For extended operations, the configurations of systems, including the automated control agent, will adjust over time due to changes in mission priorities resulting from what is learned in the mission or to workaround the degraded capability of the equipment. It is necessary support *system reconfiguration* activities and to inform team members when they occur to ensure coordinated activity. It also is necessary to provide insight into the configuration of the human-agent team and how it changes over time to ensure coordinated activity.

In this section we describe how human and agent teams can coordinate their activities when distributed physically. We use the DCI system to illustrate some of the capabilities needed for this concept of operations.

Activity Monitoring

Control automation should be able to track human activities and the states they affect for coordination with automated activities. When human activities are computer-mediated, they are observable to the control automation. In some cases, however, human activities are conducted by direct manipulation of physical devices (what Martin, et al., 2004, call *manual* interaction). Such manual activity can be difficult for automation to observe. Strategies for tracking manual activities can include monitoring for direct evidence of manipulation in data (e.g., valve closed), indirect evidence such as the human changing location to where the activity will be performed, or conclusive evidence such as asking for activity acknowledgement. When combined with knowledge of the human's planned activities, this information is an essential part of executing and updating the plan. When there is no knowledge about the human's plans, techniques for plan recognition are required to utilize the evidence collected about human activities. In both cases, a model of the human's activity (a *procedure*) is needed.

Similarly, humans should be able to understand the activities of the control automation and the states they affect for coordination with human activities. This requires the control automation to make available to the user information about its ongoing activities. However, for extended operations, it is important for the automation not to become yet another system that the human must vigilantly monitor. This can be accomplished either by designing the control agent for supervision (Sheridan, 1992) or by providing a separate, mediating agent between the human and the control agent responsible for monitoring the activities of the control agent and drawing the attention of the human when something important occurs ([a]Schreckenghost, et al, 2002).

The ARIEL agent within the DCI system is an example of such a mediating system. It tracks the activities of its user. Each activity can have a unique strategy for activity tracking, such as location tracking or evidence returned from a crew procedure viewer. Because activity tracking is never entirely accurate, we are developing techniques that permit the crew to update task completion assessments made by the agent. We call this *plan reconciliation*. The crew can change the completion assessment to *complete* (not considered for replanning) or *not-complete* (considered for replanning). The task is replanned immediately by the CAP if the agent has not yet marked it complete. The task is replanned at the end of the day if the agent has already marked it complete.

Concurrent Activity

Coordinated activity requires that concurrent activities not interfere. This includes human and automated control activities. It is highly desirable when possible to perform human operations, such as maintenance or repair, without shutting down the control automation, permitting concurrent human and automated operations. This will likely require, however, some reconfiguration of the automation to prevent automation from interfering with the human maintenance or repair activity. For example, a calibration of the O2 gas analyzer involves, among other things, the sensor sending high readings. If these high readings are being processed by the automation during the calibration, O2 concentration will be activated when none is needed, pulling down the O2 concentration in the chamber below acceptable levels. To avoid such errors, the sensor should be taken offline before calibration.

There many strategies for achieving concurrent manual and automated operations. When activities can be preplanned, a centralized activity planner will ensure that activities do not conflict. In cases where a more reactive approach is needed, distributed planning approaches can be used. Adjustable autonomy techniques enable concurrent manual and automated operations, and cover a range of approaches. They can include the reconfiguration of automation to be less than fully autonomous (Dorais, et al., 1999). This will affect the agent's ability to take action but will not affect its ability to observe (i.e., monitor data or the actions of others). Adjustable autonomy also can include the reassignment of roles among agents (Scerri et al., 2003) in response to changes in situation.

The DCI system has developed an approach for assisting humans in concurrent commanding of systems normally managed by automated control agents (Martin, et al., 2004). The Command and Authorization service authorizes human commanding if the requested procedure does not conflict with ongoing activities. It does this by determining the scope of the effects of the requested procedure on a system and its constituent subsystems, and by ensuring that only one person at a time is commanding within that scope. It also reconfigures the automated control agent to ensure compliance with the authorization granted. If the requested procedure conflicts with an ongoing action, it informs the requestor that authorization is denied and informs the user about what system states are in conflict. It permits the user to override the denial in emergencies. When the user releases authorization, the agent reconfigures the control agent to its former state.

Resource Management

Coordinating human and agent activities requires ensuring that adequate resources are available to perform them. Crew resources include items such as oxygen and water. Vehicle resources include items such as batteries and fuel. Typically, these resources are limited. For short duration operations such as the Space Shuttle, enough resource must be carried to last for the entirety of the mission. For extended operations enough resource must be provided either (1) to last between re-supply missions from Earth (like ISS), which requires managing resource usage, or (2) to regenerate resources (e.g., recycle water), which requires managing both resource production and usage.

Resource availability is usually managed as part of the planning of both human and automated activities. As such, the modeling of these activities should include knowledge of what resources are required to perform the activity. Resources modeled for planning space activities can include tools and equipment as well as consumables such as filters or lubricants. Examples of automated space planning systems that manage resources include Remote Agent Planner (Jonsson, et al., 2000) and Aspen (Chien, et al., 2000). Over extended operation, equipment will wear out or break, and it becomes important to track the degradation or loss of resource as part of the resource model. Similarly the consumption of finite consumables should be tracked. These changes in the availability or capability of resources will affect planning. For regenerative resources, resource usage should be predicted, and resource production activities should be planned. Automation has been used to predict resource usage and schedule production activities, including crop scheduling and the scheduling of waste for incineration.

Humans, robots, and software agents also can be viewed as resources, from an activity planning perspective. Modeling the skills of these agents permits more effective workload balancing, and supports plan adjustment when an agent becomes unavailable or is needed elsewhere.

For example, the DCI system uses centralized planning to build multi-day human plans. We model human activities, including human skills to ensure that tasks are assigned to qualified personnel, and are investigating mapping activities in the plan to procedures. We also use the planner to manage temporal constraints on activities.

For architectures that do not use centralized planning, the allocation of resource budgets can be combined with distributed negotiation about resource utilization.

Reconfiguration over Time

Dynamic environments like space require frequent reconfiguration of both system hardware and software. This occurs for a variety of reasons, such as (1) mission phase changes corresponding to operating configuration changes, (2) reconfiguration of vehicle or crew systems to support a payload or mission objective, and (3) reconfiguration of vehicle or crew systems in response to an anomaly. Such reconfiguration should be coordinated within the human-agent team to ensure awareness of the configuration for activities using the affected systems.

Extended operations can raise additional reasons for reconfiguration. Systems may be reconfigured to support changes that result from what is learned by conducting the mission. For example, during Phase III we learned that residual iodine in the drinking water was not healthy for the crew. The WRS was reconfigured to add an ion exchange bed to remove iodine from crew drinking and food preparation water. Similarly, for extended operations, it may not be possible or cost-effective to repair or replace degraded or broken equipment. In such cases, systems may be reconfigured to workaround the degraded capability. During the Phase III ground test, we discovered that a flow controller had been installed such that signals were reversed (1 = no flow, 0 = flow). We avoided reinstallation costs by encoding a workaround control procedure that adjusted for the reversed signals.

Typically in space operations reconfiguration has meant changing space systems. But as we distribute teams of humans, robots, and software agents, reconfiguration also can refer to a change in the team configuration – its makeup and the allocation of responsibility among its members. It becomes important to provide insight into the configuration of the group and how it changes over time to assist in team coordination: who the team members are, what responsibilities they hold (i.e., their *roles*), what activities they are performing, and where they are located. Recent work in context-aware computing supports the use of participant relationships as a context for interaction (Tarasewich, 2003). For example, the ARIEL agent provides information about its user's role, location, and activity to other agents for group awareness. The DCI system has modeled user groups for crew, flight controller disciplines, and test engineers. The ARIEL agent provides insight into both the current configuration and a history of configuration changes for team members.

As control agents are deployed, their configuration will need to be adjusted over time as well. Some of this adjustment will result from adaptation to changes required to control changing crew and vehicle systems. Some will result from evolution of the software caused by changes in operations and the discovery of software errors. The reconfiguration of control agents may also be needed for fault tolerance over extended operation, particularly with respect to the degradation of networking and computing capabilities. For example, the DCI system has been implemented with software factories that automate much of the configuration, start up, and shutdown of the DCI system for a particular deployment. These factories would be the basis of reconfiguration for fault tolerance.

Remote, Asynchronous Awareness over Time

The distribution of operations out of a centralized control room to multiple locations (i.e., offices, planetary bases, space vehicles, surface rovers) necessitates providing new approaches for maintaining *situation awareness*. It introduces the possibility of information latencies and asynchronous communication, caused both by physical remoteness as well as human unavailability. When important operational events occur, the responsible people should be notified as soon as possible regardless of location. If the situation worsens, there should be techniques for escalating the urgency and importance of the notices. In cases where primary personnel cannot respond, there should be strategies for automatically notifying backup personnel. Since people may not always be accessible, the important information about the event must be captured and made available for inspection after the fact. Some of the challenges here are providing easily understood summaries of the situation that provide ready access to more detailed information if needed.

In addition to maintaining situation awareness, for extended operations humans must maintain *awareness of system performance*. Performance tends to change slowly due to small but continuous adjustments as the system operates. These adjustments result from parts wearing in and consumables (e.g., filters) wearing out. There also are environmental considerations (e.g., thermal stability). Awareness of performance changes is needed to anticipate hardware degradation before it becomes a problem as well as to adjust control parameters to changes in steady-state behavior.

In this section we describe how human teams can maintain both situation and system performance awareness of the automated control agent over time when they are distributed physically and communicating asynchronously. We use the DCI system to illustrate some of the capabilities needed for this concept of operations.

Situation Awareness

Endsley (2001) states that "situation awareness is the (1) perception of the elements in the environment within a volume of time and space, (2) the comprehension of their meaning, and (3) and the projection of their status in the near future." With this in mind, the elements of the control environment are information about system configuration and status changes, system anomalies, and the impact of these anomalies on operations and crew. To make these elements comprehensible, they must be accessible to a distributed operational team, some of whom are inaccessible at times. Thus mechanisms should be provided for automatically detecting important events, notifying team members about these events, and providing access to additional information about the associated operational situation should it be needed.

If possible, the user should be able to choose when to review the information to avoid unnecessary distraction. In some cases, however, information is sufficiently urgent or important that the user's attention should be focused on an event when it occurs. To illustrate, in the DCI system, the ARIEL agent uses knowledge of its user's role to determine whether an incoming event from the control agent is of interest. It determines how to inform its user based on the tolerance for information latency, the need to focus the user's attention on the incoming information, and the user's availability ([b]Schreckenghost, et al, 2002).

Events passed to the human can come directly from the automated control agent. In the ground tests, the control agent sent notices when it reconfigured the system, when an anomaly occurred, and when it took anomaly response actions. One of the challenges for the control agents is providing the user with useful information about the domain system without excessive processing overhead, because software implementation models may not match human models of operation.

To offload some of this processing, events also may be provided to the human by an intermediate process that monitors information from the control agent. An example of such a process is the EDA in the DCI system. The EDA computes Caution & Warning events and failure impacts (e.g., rising CO2 when CO2 removal system is shutdown) using pattern-matching software.

Many operational situations develop over time as a progression of simpler events. For extended operations, it is especially important to capture these events as they occur for inspection after the fact, including comparison to similar events in the future. In the DCI system, we use the Complex Event Recognition Architecture (CERA) (Fitzgerald, et al., 2004) to capture complex events that we call *situation objects*. CERA applies natural language parsing techniques to detect event patterns with complex temporal and hierarchical relationships among them. These situation objects encapsulate the events describing important operational situations in space systems, such as a loss of communication anomaly or a bioreactor slough maintenance activity. In DCI we have developed a XML data structure for life support situations and a generalized viewer for this data structure (Thronesbery and Schreckenghost, 2003). These situations are accessed and logged as notices in the ARIEL agent.

System Trends and Performance

An essential aspect of system awareness over extended operation is knowledge of system performance and trends over time. This knowledge is typically not characterized by specific situations. Instead, it is collective knowledge often described statistically. Control performance can be characterized at many levels:

- component level: changes in the hardware system (e.g., biases, drifts on parameters)
- process level: changes in the configuration of software and hardware (e.g., connectivity, flow)
- product level: changes in the production or consumption of resource (e.g., production rate)

In the DCI system we use the CERA system to capture daily performance objects for the PPS. The information that we capture is a mix of component and product level information, and includes the following:

- % change in bed ion loading (ion exchange beds remove trace inorganic wastes and ammonium)
- average power used to run system

- minimum and maximum total organic carbon (TOC), and time periods in which TOC was hig

We are evaluating this object in the PPS ground test.

Tracking parameter changes over time (called *trending*) has long been an operational technique employed by controllers in manned space operations. Strip charts and real-time plots are used as visual predictors of system behavior. For example, one can see a device going out of limits (e.g., engine overheating) and often can predict when it will reach the limit. Real-time trending, and the ability to interpret trends, is a commonly requested capability for operations. One of the challenges of automating such support, however, is that the interpretation of these trends is very situation-specific. Additionally, when system behaviors are coupled, conclusive interpretation of trends can be difficult.

Conclusions

Introducing automated control agents into manned spaced operations will result in a distributed human-agent team. Automated control agents will operate autonomously in space most of the time. Crew will interact with them for maintenance tasks. Flight controllers will monitor their performance to anticipate problems. If problems occur, both crew and ground may be involved to resolve them.

This distributed human-agent team will require new ways of working together. To address this we first described how humans and automated control agents activities can be coordinated. This includes maintaining awareness of team member activities, providing strategies to avoid interfering activities and to take required actions when tighter coordination is needed, ensuring adequate resources are available, and maintaining awareness of system and team reconfiguration over time. Next we discussed maintaining awareness of both situation and system performance as it changes over time.

We also have proposed that another type of agent, a mediating agent to assist humans in performing distributed operations, is useful. We justify this claim based on our experience in the JSC ground tests. We have developed a prototype of such an agent as part of the DCI system. The ARIEL agent provides services such as event notification, task management, command authorization, and location tracking to assist its user. We are currently evaluating the DCI system during a ground with WRS control engineers.

Acknowledgements

We acknowledge the support of Dr Michael Shafto, manager of Human-Centered Computing in NASA's Intelligent Systems Program, who sponsored this work.

References

Bonasso, P., Firby, J., Gat, E., Kortenkamp, D., Miller, D, and Slack, M. 1997a. Experiences with an Architecture for Intelligent, Reactive Agents. *Journal of Experimental Theory of Artificial Intelligence*. 9: 237-256

Dorais, G., Bonasso, R. P., Kortenkamp, D., Pell, B., and Schreckenghost, D. Adjustable autonomy for human-centered autonomous systems. Proceedings of Adjustable Autonomy Workshop. *International Joint Conference on Artificial Intelligence (IJCAI)* Aug 1999.

Bonasso, R. P., Kortenkamp, D., and Thronesbery, C. 2003. Intelligent Control of A Water Recovery System: Three years in the Trenches. *AI Magazine* 24 (1): 19-44.

Chien, S., G. Rabideau, R. Knight, R. Sherwood, B. Engelhardt, D. Mutz, T. Estlin, B. Smith, F. Fisher, T. Barrett, G. Stebbins, D. Tran , "ASPEN - Automating Space Mission Operations using Automated Planning and Scheduling," *SpaceOps* 2000,.

Endsley, M. Designing for Situation Awareness in Complex Systems. *Proceedings of Second International Workshop on Symiosis of Humans, Artifacts, & Environment*. Japan. 2001.

Fitzgerald, W., R. J. Firby, A. Phillips, J. Kairys. Complex Event Pattern Recognition for Long-Term System Monitoring. *AAAI Spring Symposium 2004 on Interaction between Humans & Autonomous Systems over Extended Operation* Stanford Univ., Palo, Alto, CA. Mar 2004.

Jonsson, Morris, Muscettola, Rajan, and Smith. Planning in Interplanetary Space: Theory and Practice. Proceedings of the *Fifth International Conference on Artificial Intelligence Planning Systems*, Breckenridge, CO. April, 2000.

Martin, C., Schreckenghost, D., Bonasso, P. Augmenting Automated Control Software to Interact with Multiple Humans. *AAAI Spring Symposium 2004 on Interaction between Humans & Autonomous Systems over Extended Operation* Stanford Univ., Palo, Alto, CA. Mar 2004.

Scerri, P., L. Johnson, D. Pynadath, P. Rosenbloom, N. Schurr, M. Si and M. Tambe. Getting Robots, Agents and People to Cooperate: An Initial Report. *AAAI Spring Symposium on Human Interaction with Autonomous Systems in Complex Environments*. March 2003.

Schreckenghost D.; Ryan, D.; Thronesbery, C.; Bonasso, P.; and Poirot, D. Intelligent control of life support systems for space habitats. *AAAI IAAI* 1998. Madison, WI. July 1998.

[a]Schreckenghost, D, Thronesbery, C., Bonasso, P., Kortenkamp, D., and Martin, C. 2002. Intelligent Control of Life Support for Space Missions. *IEEE Intelligent Systems* 17 (5): 24-31.

[b]Schreckenghost, D.,C. Martin, and C. Thronesbery. Specifying Organizational Policies and Individual Preferences for Human Software Interaction. Proceeding *AAAI 2002 Fall Symposium on Etiquette for Human-Computer Work*. Nov 2002.

Sheridan, T. *Telerobotics, automation, and human supervisory control*. MIT Press. Cambridge, MA, USA . 1992.

Thronesbery, C., and Schreckenghost, D., Situation Views: Getting Started Handling Anomalies. *IEEE International Conf on Systems, Man, & Cybernetics*. Wash, D. C. 2003

Tarasewich, P. Towards a Comprehensive Model of Context for Mobile and Wireless Computing. *Proceedings of the AMCIS 2003 Conference*, 114-124.