Human-Robot Teams for Large-Scale Assembly

Reid Simmons[†‡], Sanjiv Singh[†],
Frederik W. Heger[†], Laura M. Hiatt[‡], Seth Koterba[†], Nik Melchior[†], Brennan Sellner[†]
[†]*The Robotics Institute,* [‡]*Computer Science Department*
School of Computer Science
Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA 15213

Abstract— **Construction and assembly are complex and arduous tasks, especially when performed in hazardous environments such as in orbit, on the Moon, or on Mars. Effective assembly of structures in such environments, where human labor is expensive and scarce, can be facilitated by the use of heterogeneous robotic teams. Over the past five years, we have developed the architectural framework and tools to coordinate robotic assembly teams, as well as to incorporate the unique skills of remote human operators using an approach that allows authority to "slide" between autonomy and human control at a fine degree of granularity. We have used this approach in several assembly scenarios, and have quantified the gains in reliability and efficiency over both purely autonomous and purely teleoperation approaches.**

I. Introduction

Construction has never been a solo profession. For as long as human society has existed, humans have been banding together to construct edifices that no single person could have built alone. Just as human construction projects require people with differing skill sets, so too construction in space, on the Lunar surface, and on Mars requires a multitude of heterogeneous workers. In such hazardous environments, human labor is expensive and scarce, making less-fragile robots an attractive option. We envision teams of robots working independently on different portions of a large-scale construction project, supervised by a small number of human operators on Earth or in a nearby habitat.

For the past five years, we have been investigating how heterogeneous robotic teams can work together, and with humans, to assemble complex structures. This work has focused on two main aspects: (1) the architecture and tools needed to coordinate the work of different autonomous agents and (2) methods to enable remote human operators and autonomous agents to work together in a flexible manner.

To these ends, we have developed the *Syndicate* architecture, a multi-agent, 3-tiered architecture that allows for the coordination of disparate agents at a variety of levels of abstraction. Using Syndicate, we have developed an approach to *sliding autonomy*. Sliding autonomy is a paradigm that enables control of tasks to "slide" back and forth between a human operator and an autonomous system. Humans bring a unique set of skills to a mixed team, especially their flexible problem-solving abilities. To make the most efficient use of the available expertise, each robot team will operate autonomously, asking for help from an operator only when a problem arises that it cannot solve, or when human control provides significant benefits in reliability or efficiency. Most work to date on sliding autonomy has been limited to the control of single robots. In our work, we have extended the concept to heterogeneous multi-robot teams and support the transfer of autonomy at a much finer granularity than previously reported.

In experiments using several different assembly scenarios, we have demonstrated completely autonomous operation using multiple, heterogeneous robots that are coordinated through the Syndicate architecture. We have also quantified the gains in reliability and efficiency that sliding autonomy brings over pure autonomy and pure teleoperation, respectively.

II. Related Work

In a domain as wide-ranging as multi-robot assembly, the amount of related work is vast. We will briefly survey several related projects; see [1] for an in-depth discussion of related work.

Coordinated assembly performed by teams of mobile robots is of particular interest to the space community. Stroupe et al. [2] use the CAMPOUT architecture to coordinate robots with purely behavior-based strategies to perform coupled tasks, similar to ours. Their team is homogeneous and performs a single, albeit complex, task. In contrast, our agents are heterogeneous, and are able to perform a range of tasks using different objects and manipulators.

Human-robot interaction has also been a prime research topic in recent years. The COBOT project [3] [4] seeks to make manually operated machines more intelligent by providing guidance so that the operator does not have to provide fine control. This approach requires the active participation of both human and robot at all times, rather than using robotics as a multiplier of the human's abilities. NASA's ASRO project [5] developed a mobile robot to assist a space-suited human by carrying tools, helping to manipulate objects, and providing sensor information. This robot was completely teleoperated, in contrast to our mix of autonomy and human control.

Our use of the term "sliding autonomy" corresponds with the term *adjustable autonomy* as presented by Dorais et al. [6],

This work has been supported by NASA under grant numbers NNX06AD23G, NNA04CK90A, NAG2-1294, and NAG9-1226; NSF Fellowship No. DGE-0234630; NSF Grant ECS-0233698; and CNPq-Brazil Grant 200501/86-0.

who present several examples of the utility of this approach. Although Dorais et al. use "adjustable autonomy" in the same sense as we use "sliding autonomy", we feel that the term "adjustable autonomy" carries the connotation that the level of autonomy has been set at a fixed value prior to execution, as opposed to dynamically sliding back and forth during the course of execution. Kortenkamp et al. [7] have developed an infrastructure for sliding control of a robot manipulator similar to our approach in many respects. We have extended their approach with more complex assembly tasks and a finer granularity of task control.

III. Architecture and Sliding Autonomy

A. The Syndicate Architecture

The Syndicate architecture was designed from its inception to support tight coordination between agents through varying degrees of abstraction. At its core is a three-tiered approach (Fig. 1) similar to that of 3T [8], but extended to support multiple autonomous agents. The three layers (planner, executive, and behavioral) provide differing levels of task granularity and abstraction, easing the design of complex systems by allowing the designer to use the level of abstraction most appropriate to the problem at hand. Within an agent, layers may communicate with the layers "above" and "below" them; that is, the executive may communicate with the planner and behavioral layers, while the behavioral may communicate only with the executive. Commands flow downward from more abstract layers to the more concrete, while data flows upwards. In addition, each layer of an agent may exchange data with the same layer on other agents, allowing coordination to be performed at all levels of the architecture.

The behavioral layer is based on the Skill Manager of [8], and is stateless and reactive. It is the finest-grained of the layers, and acts as the interface between the controllers/hardware and the executive layer. As such, it operates at a very fine temporal granularity and handles the hardware and environmental details that the upper layers abstract away. By doing so, it is able to react quickly to changes in the world.

The executive layer is responsible for managing and sequencing the agents' hierarchical task trees [9], as well as maintaining state relevant to the tasks. The representation of state and sequencing are the primary distinctions between the executive and behavioral layers. A task is loosely defined as an abstraction of one element of the scenario that requires state and/or may be decomposed into atomic behaviors in order to satisfy a goal. The executive's purpose is to order these tasks, and to control the configuration of the behavioral layer in order to accomplish them. The behavioral layer provides processed data to inform the executive's actions, and the executive in turn uses this data to select the appropriate actions. The planning layer of Syndicate has not yet been implemented. As a result, the executive is limited to fixed task orderings, and is able to attempt recovery from failures only at a local level. For instance, the executive may attempt to retry a task, but is unable to replan to select an alternate route or method to accomplish the goal. The implementation of a planning layer for Syndicate is a current focus of our research.

The executive and behavioral layers both support inter-agent communication, allowing coordination to occur at both levels. At the behavioral layer, information may flow between components of multiple agents, enabling distributed servo loops. For instance, a sensing agent may provide a manipulation agent with location data to enable an assembly action. The executive layer may constrain tasks on multiple agents to run sequentially or simultaneously, enabling tightly coordinated activity while avoiding conflicts.

The Syndicate architecture has proved to be a flexible framework on which to build autonomous multi-agent construction teams and explore different aspects of sliding autonomy.

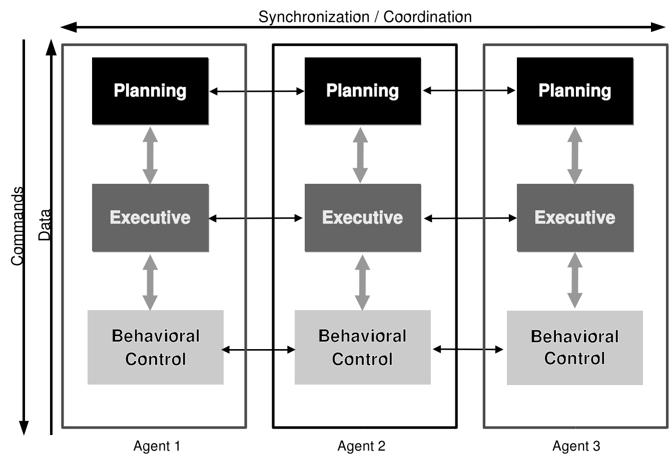

Fig. 1. The Syndicate architecture. Note that each layer may communicate with *all* other layers at the same level. For clarity, the links between Agent 1 and Agent 3 are not depicted.

B. Sliding Autonomy

In complex, multi-agent tasks, such as the assembly of structures, there are inevitably unforeseen errors: it is currently beyond the state-of-the-art to develop an autonomous agent that can adapt to all possible problems that may occur in such dynamic environments. However, agents that can handle many, but not all, of the tasks and possible errors are much easier to produce: the most common cases generally encompass the majority of what occurs during actual operation. What remains is a series of rare failures and odd circumstances. While autonomous systems fare poorly in such corner cases, humans thrive. These sets of complementary skills are what drive sliding autonomy: autonomous agents are excellent at rapidly performing tasks under nominal conditions, while humans excel at troubleshooting and resolving novel problems. Sliding autonomy takes advantage of both by allowing the control of elements of the team to "slide" back and forth between human and autonomous control.

We have implemented sliding autonomy primarily in the executive layer of Syndicate. Tasks are split into action and monitor subtasks. This allows responsibilities to be better apportioned between human and agent according to their strengths. For instance, the autonomous system may be good at performing an assembly task, but poor at determining when it is completed. In such a case, the action task would be under autonomous control, while a human operator would handle the monitoring task. By providing relatively fine-grained control, we allow the strengths of all elements of the team to be used to their best advantages.

Previous work on sliding and adjustable autonomy focused on the single-agent case, while we have addressed the multi-agent, single-human problem. One fundamental difference in the multi-agent scenario is that a human is unable to simultaneously monitor all agents, and will inevitably be unaware of problems that arise. This lack of awareness yields three issues that must be addressed in multi-agent sliding autonomy:

1) **Requesting help:** Since the human may not be monitoring, a given agent must be able to determine when to request help. In our approach, an agent makes this determination by modeling the past performance of both autonomous and human controllers, and analytically determining whether it is more efficient to make another autonomous attempt or request assistance [10].
2) **Gaining situational awareness:** It is practically impossible for a single operator to simultaneously monitor a team of robotic agents. Any approach to multi-agent sliding autonomy must address the problem of reacquiring situational awareness. Our approach is to maintain buffers of data that may be played back when the human is asked to address a problem. We conducted human-subject experiments to evaluate this approach, and found that a point exists beyond which additional data merely slows the human's response time [11]. The specific value of this point is a function of the scenario.
3) **Maintaining coordination:** After the operator has been in control, the autonomous agent must be able to smoothly resume control from whatever state the human left the system in when he returned control authority. We currently monitor the operator's progress on the task to provide the autonomous controller with accurate initial state. However, this assumes that the human will at least roughly follow the current plan. Introducing additional flexibility into our approach is an area of current research.

Any sliding autonomy system must also answer the question of initiative: who has the authority to change the control of a task, and when? We have conducted a series of human-subject experiments to determine the effects of different approaches to initiative in sliding autonomy, varying whether the human and/or autonomous system could initiate a change of task control [1] [12]. In all sliding autonomy cases, the autonomous system was allowed to request help with a failing task, if its models of itself and the human predicted that the more efficient approach was for the human to make the next attempt. We varied whether the human was constantly monitoring the team (and proactively taking control of tasks) or simply responding to requests for assistance. As baselines, we also performed the task under full teleoperation and complete autonomy.

We found that, in all cases, sliding autonomy approached the efficiency of purely autonomous operation, while retaining much of the reliability of teleoperation. The difference between having the human operator actively monitoring versus only responding to requests for help was relatively slight, except when measuring operator workload: an operator not responsible for constantly monitoring a team experienced much lower subjective workloads, as measured by the NASA-TLX [13] survey.

IV. Scenarios: Past and Current Work

We have demonstrated our approach to large-scale assembly in two past and two current scenarios, as detailed in Table I. The first three scenarios are variations on connecting beams and nodes, and utilize our ground-based 3-robot team of a Roving Eye, Mobile Manipulator, and Crane. The Roving Eye is responsible for reporting the relative positions of objects in the workspace, while the Crane performs heavy lifting operations. In contrast, the Mobile Manipulator handles fine manipulation tasks involving smaller payloads.

In our first scenario (Table I(a)), the Crane coarsely positioned a 3-meter beam above two stanchions. The beam was suspended via a cable, providing a compliant grasp. The Roving Eye then tracked the (swinging) beam and the Mobile Manipulator's end effector. Using the resulting data, the Mobile Manipulator grasped the beam and docked it into the stanchions. This scenario demonstrated the basic skills that we have built upon in our subsequent work: distributed visual servoing and tightly coordinated manipulation between heterogeneous agents.

The second scenario (Table I(b)) built upon the first by requiring more manipulation, movement, and coordination tasks, to assemble a more complicated structure. This structure, seen in Table I(b), consists of four nodes and four beams, yielding a total of eight docking tasks that must be completed. The Roving Eye maintained its sensing role, while the Crane braced the wheeled nodes to keep them from rolling away, and the Mobile Manipulator maneuvered the beams into place. With this scenario, we demonstrated task sequencing, fine manipulation with six degrees of freedom, and our initial approach to sliding autonomy. Scenario 2 was utilized in the human-subject experiments discussed above and detailed in [1], [11], and [12].

Two new scenarios are currently under development. The scenario in Table I(c) extends the previous scenario with different assembly components and a more complex task: assembling a 3x3 grid of nodes and beams. This will involve 9 nodes, 12 beams, and 24 docking operations, in addition to the requisite movement and coordination tasks. The final structure is larger than the workspace of the Crane, necessitating explicit management of the workspace: we can no longer assume there will be sufficient space for agents to reach all portions of the assembly. This structure also is complex enough to make an interesting planning problem; we will be introducing a planning layer, likely based on ASPEN [14],

	(a) Scenario 1: Hanging Beam	(b) Scenario 2: Square
	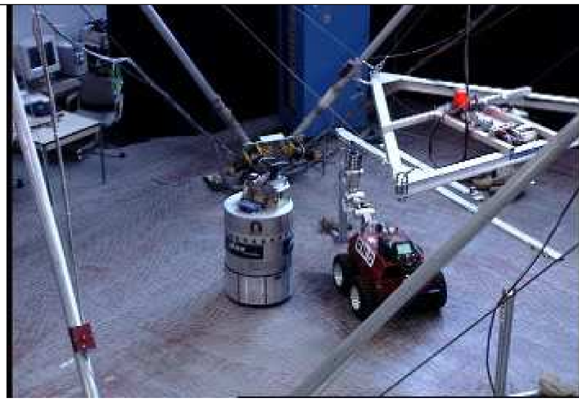	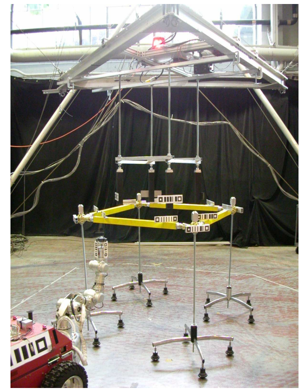
Dates	1999 - 2003	2003 - 2006
Agents	Roving Eye, Mobile Manipulator, Crane	Roving Eye, Mobile Manipulator, Crane
Characteristics	• Physically coupled manipulation along two dimensions • 1 cm tolerances	• Multiple manipulation tasks • Many movement / coordination tasks • 0.5 cm, 5 degree tolerances
New Capabilities	• Visual servoing • Coordinated manipulation	• Task sequencing • Fine manipulation in 6-DOF • Initial sliding autonomy
	(c) Scenario 3: Grid	**(d) Scenario 4: EASE[a]**
	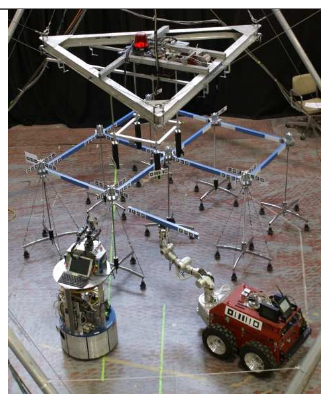	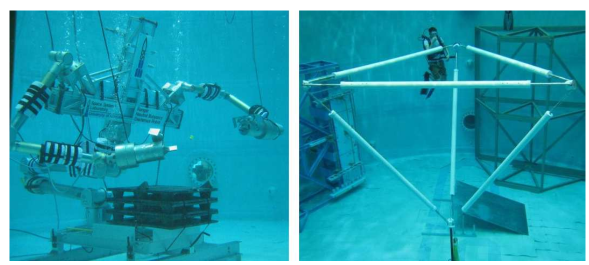
Dates	2006 - present	2006 - present
Agents	Roving Eye, Mobile Manipulator, Crane	4 Ranger arms
Characteristics	• Very large number of tasks • Structure exceeds workspace size • 1-2 mm, 1 degree tolerances	• Space-tested structure • Space-relevant robots • Neutral-buoyancy environment • 1-2 mm, 0.5 degree tolerances
New Capabilities	• Finer-grained sliding autonomy • Multiple manipulators capable of docking • Workspace management • *Assembly planning*	• *Assembly in zero gravity* • *Very fine manipulation* • *Integration with other institutions*

[a] Ranger and EASE photos courtesy Space Systems Laboratory, University of Maryland

TABLE I

MULTI-ROBOT ASSEMBLY SCENARIOS. ITEMS IN ITALICS ARE CURRENTLY IN PROGRESS.

to manage this scenario. In addition, we will be extending our approach to sliding autonomy to enable finer-grained intervention by the operator. Instead of always taking full control of the problematic task, the operator will have the option of giving higher-level guidance, then quickly returning control to the autonomous system. We can currently assemble the new hardware, and are developing the planning system and extensions to sliding autonomy.

The final scenario represents a different aspect of the domain. We are collaborating with the Space Systems Laboratory of the University of Maryland to use their four-arm Ranger robot [15] to assemble the EASE structure (Table I(d), right). EASE is an inverted tetrahedron, consisting of 4 nodes and 6 beams, that has been assembled several times in the Space Shuttle's bay while in orbit. It was used as a test of astronauts' ability to assemble and maintain structures. While the structure involves fewer components than our other current scenario, it is three-dimensional, rather than planar, will be assembled in a neutral-buoyancy tank, and requires much tighter assembly tolerances than previous structures. This work is still at an early stage. We have demonstrated basic integration of our software with the Space System Laboratory's by performing grasping operations with the Ranger hardware. We will be moving on to assembly operations shortly, and plan eventually to demonstrate aspects of sliding autonomy.

V. Future Work

We plan to continue work on our two primary research focii: autonomous multi-agent assembly and sliding autonomy.

A. Autonomous Multi-Agent Assembly

In the area of assembly, three areas stand out for further work: coordinated manipulation, overlapping agent capabilities, and planning. In our first scenario, the Mobile Manipulator and Crane simultaneously manipulated the same rigid beam, although the Crane had a very compliant grasp. We plan to bring such coordinated manipulation back into our current scenarios, which will require tighter inter-agent coordination and expose interesting research questions.

We also plan to introduce overlapping capabilities into our agents. Currently, our three land-based agents have non-intersecting capabilities. By expanding their skill sets, and possibly adding additional agents, we will extend the solution set for a given problem and increase the team's flexibility.

Finally, we will add planning capabilities to the system that will take advantage of the flexibility inherent in multi-agent teams, as well as providing more comprehensive recovery strategies in the face of failure. The planning/scheduling system we envision will enable the team to predict when tasks will take longer than expected and act accordingly to minimize task completion time or maximize the number of tasks that are completed. One approach to this is transferring agents back and forth between teams as execution warrants.

B. Sliding Autonomy

There are a broad range of interesting research questions in the sliding autonomy arena; we plan to focus on three of them: finer-grained control, in-situ humans, and flexible human-autonomous handoffs. Our current system allows the control of individual tasks (and the associated hardware) to pass back and forth between the remote human operator and the autonomous control software. However, the only way for an operator to intervene is to take complete control of the robot in question. We plan to extend Syndicate to support a greater range of interventions, allowing the operator to reparameterize tasks (e.g. "use Node B instead of Node A") or reorder the plan (e.g. "build the left side first, instead of the right side"). This will further reduce the load on the operator, allowing him to service more robotic teams.

To date, our work on sliding autonomy has focused on a remote human operator supervising an on-site robotic team. We are interested in investigating how best to incorporate an in-situ human into the team. Interesting questions include how the human will communicate with the team, what types of help he can render, how to determine the human's current task, and whether robotic agents should proactively offer help to the human. We have performed some initial work on gestural commands and frames of reference [16], but have merely scratched the surface.

Finally, an important aspect of sliding autonomy is the way in which the autonomous system determines how to proceed when control is returned to it by the operator. Our current system assumes the human completes the task he was asked to help with, and immediately returns control to the robot. However, this is by no means a safe assumption: the human may have given up, completed additional tasks beyond the assigned one, or performed the task in such a way as to violate some of the postconditions the automated system relies on. For sliding autonomy to be truly robust, the autonomous system must monitor the operator's actions, determine his intent, and ascertain his progress in order to smoothly resume control.

VI. Conclusion

Our research to date has resulted in a wealth of experience and knowledge in multi-robot coordination, human-robot interaction, and assembly. Our work with large-scale assembly domains, and our attempts to answer the challenges that arise from such tasks prompted us to develop the Syndicate architecture, which addresses the coordination of disparate robots and humans in these types of complex problems. Syndicate has been employed in a number of past and ongoing scenarios, demonstrating our approach to multi-agent coordination and sliding autonomy. This approach has shown that it is possible to leverage the skills of human operators by giving them supervisory roles, in which they are asked for help by the autonomous robots when problems arise. Making more efficient use of human time in this way allows a small group of humans to supervise a much larger number of robots, thus multiplying the effectiveness of the humans and magnifying the scale of the tasks that can be undertaken. In addition to our land-based demonstration scenarios, we are currently applying our approach to space-relevant robots and hardware, specifically the multi-arm Ranger robot and the EASE structure. We feel that multi-robot teams have a great future in complex space scenarios and that, by working together with remote humans, effective and reliable assembly can be realized.

ACKNOWLEDGMENTS

The authors would like to thank the many individuals who have contributed to the project over the years, including Jonathan Brookshire, Robert Burridge, David Hershberger, David Kortenkamp, Josue Ramos, and Trey Smith.

REFERENCES

[1] B. Sellner, F. Heger, L. M. Hiatt, R. Simmons, and S. Singh, "Coordinated multi-agent teams and sliding autonomy for large-scale assembly," *Proceedings of the IEEE*, vol. 94, no. 7, July 2006.

[2] A. Stroupe, T. Huntsberger, A. Okon, and H. Aghazarian, "Precision Manipulation with Cooperative Robots," in *Multi-Robot Systems: From Swarms to Intelligent Automata*, L. Parker, F. Schneider, and A. Schultz, Eds. Springer, 2005.

[3] R. B. Gillespie, J. E. Colgate, and M. Peshkin, "A General Framework for Cobot Control," in *International Conference on Robotics and Automation*, Detroit, MI, 1999.

[4] W. Wannasuphoprasit, P. Akella, M. Peshkin, and J. E. Colgate, "Cobots: A Novel Material Handling Technology," International Mechanical Engineering Congress and Exposition, Anaheim, ASME 98-WA/MH-2, 1998.

[5] R. R. Burridge, J. Graham, K. Shillcutt, R. Hirsh, and D. Kortenkamp, "Experiments with an EVA Assistant Robot," in *Proceedings of the 7th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*, Nara, Japan, 2003.

[6] G. Dorais, R. Banasso, D. Kortenkamp, P. Pell, and D. Schreckenghost, "Adjustable Autonomy for Human-Centered Autonomous Systems on Mars," Presented at the Mars Society Conference, 1998.

[7] D. Kortenkamp, R. Burridge, P. Bonasso, D. Schrenkenghost, and M. Hudson, "An Intelligent Software Architecture for Semiautonomous Robot Control," in *Autonomy Control Software Workshop, Autonomous Agents 99*, 1999.

[8] R. Bonasso, R. Firby, E. Gat, D. Kortenkamp, D. Miller, and M. Slack, "Experiences with an Architecture for Intelligent, Reactive Agents," *Journal of Experimental and Theoretical Artificial Intelligence*, vol. 9, no. 2-3, pp. 237–256, 1997.

[9] R. Simmons and D. Apfelbaum, "A Task Description Language for Robot Control," in *Proceedings of the Conference on Intelligent Robots and Systems (IROS)*, Victoria, Canada, 1998.

[10] B. Sellner, R. Simmons, and S. Singh, "User Modelling for Principled Sliding Autonomy in Human-Robot Teams," in *Multi-Robot Systems: From Swarms to Intelligent Automata*, L. Parker, F. Schneider, and A. Schultz, Eds. Springer, 2005.

[11] B. Sellner, L. M. Hiatt, R. Simmons, and S. Singh, "Attaining Situational Awareness for Sliding Autonomy," in *Proceedings of HRI 2006*, Salt Lake City, UT.

[12] F. W. Heger, L. M. Hiatt, B. Sellner, R. Simmons, and S. Singh, "Results in Sliding Autonomy for Multi-Robot Spatial Assembly," in *8th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*, Munich, Germany, September 5-8, 2005.

[13] S. G. Hart and L. E. Staveland, *Human Mental Workload*. Amsterdam: North-Holland, 1988, ch. Development of the NASA-TLX (task load index): Results of Empirical and Theoretical Research, pp. 139–183.

[14] S. Chien, *et al.*, "Aspen – automated planning and scheduling for space mission operations," in *Space Ops*, Toulouse, June 2000.

[15] S. Roderick, B. Roberts, E. Atkins, and D. Akin, "The ranger robotic satellite servicer and its autonomous software-based safety system," *IEEE Intelligent Systems*, Sep/Oct 2004.

[16] L. M. Hiatt and R. Simmons, "Coordinate frames in robotic teleoperation," in *International Conference on Intelligent Robots and Systems (IROS)*, Beijing, China, 2006.