

Coordinated Multiagent Teams and Sliding Autonomy for Large-Scale Assembly

A robot team that assembles complex structures can benefit from its human team members' unique skill set to improve reliability and robustness.

By BRENNAN SELLNER, FREDERIK W. HEGER, *Member IEEE*, LAURA M. HIATT, REID SIMMONS, AND SANJIV SINGH

ABSTRACT | Recent research in human-robot interaction has investigated the concept of sliding, or adjustable, autonomy, a mode of operation bridging the gap between explicit teleoperation and complete robot autonomy. This work has largely been in single-agent domains—involving only one human and one robot—and has not examined the issues that arise in multiagent domains. We discuss the issues involved in adapting Sliding Autonomy concepts to coordinated multiagent teams. In our approach, remote human operators have the ability to join, or leave, the team at will to assist the autonomous agents with their tasks (or aspects of their tasks) while not disrupting the team's coordination. Agents model their own and the human operator's performance on subtasks to enable them to determine when to request help from the operator. To validate our approach, we present the results of two experiments. The first evaluates the human/multirobot team's performance under four different collaboration strategies including complete teleoperation, pure autonomy, and two distinct versions of Sliding Autonomy. The second experiment compares a variety of user interface configurations to investigate how quickly a human operator can attain situational awareness when asked to help. The results of these studies support our belief that by incorporating a remote human operator into multiagent teams, the team as a whole becomes more robust and efficient.

Manuscript received June 1, 2005; revised June 1, 2006. This work was supported by the National Aeronautics and Space Agency under Grant NNA04CK90A.

B. Sellner, F. W. Heger, and S. Singh are with the Robotics Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: bsellner@andrew.cmu.edu; fwh@cs.cmu.edu; ssingh@ri.cmu.edu).

L. M. Hiatt is with the Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: lahiatt@cs.cmu.edu).

R. Simmons is with the Robotics Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213 USA. He is also with the Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: reids@cs.cmu.edu).

Digital Object Identifier: 10.1109/JPROC.2006.876966

KEYWORDS | Author, please supply your own keywords or send a blank e-mail to keywords@ieee.org to receive a list of suggested keywords

AQ1

I. INTRODUCTION

As expectations for robotic systems increase, it becomes harder and harder to meet them with the capabilities of a single robot. One approach is to use multiple simple robots that work together to perform tasks that would require a very complex single mechanism. These teams not only bring a much broader spectrum of potential capabilities to a task but also may be more robust in the face of errors and uncertainty.

While it is envisioned that robot teams eventually will be able to autonomously perform complex tasks, such as large-scale assembly in remote environments, the current state-of-the-art falls short of the necessary capabilities. In particular, the number of contingencies that must be considered and provided for to make robots fully autonomous in open and dynamic environments is prohibitively large. On the other hand, pure teleoperated control of such robots is unlikely to be very efficient due to communication delays, the large number of human operators required, and the sensing and visualization problems inherent in any teleoperation domain. Our goal is to develop a framework within which a single human operator can oversee and flexibly intervene in the operation of a team of largely autonomous robots.

One scenario exemplifying this approach is the assembly of large structures in hazardous environments, such as orbital solar power arrays or Mars habitats. In such environments, in-place human labor is either infeasible or scarce and expensive, making less-fragile robots an attractive option. We have been examining how robots and

ground-based humans can work together while assembling large scale orbital structures, such as kilometer-wide solar power arrays. We ultimately envision many teams of robots working independently on different portions of the structure, overseen by a small number of human operators on Earth.

In order to develop the architectures, software capabilities, and models needed for such an endeavor, we have assembled a heterogeneous team of Earth-bound robots. Our current scenario is the construction of a square frame using four beams and four connectors (called “nodes,” see Fig. 1). To (weakly) simulate space conditions, the nodes are free to roll around. Thus, each node must be braced against insertion forces before a beam can be inserted into it. After each side of the square is completed, the next beam must be procured and brought to the worksite. Heterogeneous teams are well suited to such construction scenarios, where many different skills, and correspondingly different hardware, are required. Our team consists of an imprecise heavy-lift robot, a weaker but more precise Mobile Manipulator, and a dedicated sensing robot (Fig. 1).

Human operators bring their own unique set of skills to the team. To make the most efficient use of the flexible problem-solving skills of humans, a robot team should operate mostly autonomously, getting help from a human operator only when a problem arises that it cannot resolve by itself, or when human control provides significant benefits. Sliding, or adjustable, autonomy is an approach that has been developed to address this, yielding more robust and efficient systems. Most work in Sliding Autonomy has been limited to the control of

single robots. In this paper, we extend the concept to heterogeneous multirobot teams and support the sliding of autonomy at a much finer granularity than has been previously reported.

One fundamental difference in applying Sliding Autonomy to multiagent teams is that the human will not always be aware of everything that is happening to all of the agents; whereas, a human who is monitoring only a single agent is able to keep abreast of all relevant developments. This lack of awareness gives rise to three major issues that need to be addressed for multiagent Sliding Autonomy.

Requesting help: Since the human cannot keep track of all the robots in a multiagent team at all times, the ability of an agent to ask for help is critical. By using empirical performance models of both the autonomous system and the human operator, knowledge of typical human learning curves, and information about the team’s state, the system can make reasonably accurate predictions of how long a given task will take under either human or robotic control. This, in turn, allows it to make principled decisions about when to ask the human for help.

Gaining situational awareness: Since the human is not able to monitor all robots at once, or may be called away to attend to another team, it is important that the operator quickly gain situational awareness of a robot’s workspace when help is requested. It is also important that the autonomous system model how long this will take (as discussed previously), in order to make principled decisions about when to ask for help.

Maintaining coordination: During and after human intervention, the components of the system that remain under autonomous control must continue to operate and maintain the interagent coordination necessary for task completion. We address this issue by enabling agents to monitor themselves and other agents as appropriate for the current task, allowing them to remain coordinated even when the operator is in control of one of the agents. This allows the portion of the team still under autonomous control to continue meaningful operation while other agent(s) are under human control.

We conducted two human-subject experiments to validate our approach to Sliding Autonomy in the multiagent domain. The first examined the performance of four different modes of autonomy to determine the efficacy of Sliding Autonomy. The second experiment investigated how situational awareness is regained and how the composition of the user interface affects this acquisition.

The first experiment, which varied the degree of human involvement in the task, consisted of performing our construction task under four different human/robot cooperation strategies: pure autonomy, System-Initiative




Fig. 1. Our three-agent team consists of Roving Eye (left), Mobile Manipulator (right), and Crane (above structure; only black Crane receptacles are visible). Here, they have completed assembly of square. Crane is still bracing final node, while Mobile Manipulator just disengaged from final beam.

Sliding Autonomy (SISA), Mixed-Initiative Sliding Autonomy (MISA), and teleoperation. Pure autonomy does not involve the human, consisting solely of autonomous behaviors and recovery actions. In contrast, during teleoperation the human is in complete control of all aspects of all the robots. Bridging the gap between these two extremes, SISA allows the operator to intervene only when asked to do so by the autonomous system; while in MISA, the human can also intervene at any time of his own volition. SISA is designed to approximate situations where the operator is a scarce resource and must attend to multiple robot teams or other tasks. MISA, on the other hand, captures situations where humans can be more dedicated to observing the team's activities.

The results of this experiment indicate that the autonomous system is consistently faster, but less reliable, than a purely teleoperated approach. In both SISA and MISA, the speed of assembly approaches that of the autonomous system. While operator workload results showed that the preference for either SISA or MISA is very task- and user-dependent, the workload for the human is clearly less than during pure teleoperation. We conclude that our adaptation of Sliding Autonomy improves the multiagent team's reliability without compromising efficiency and that it can be easily reformulated to meet the differing constraints of a variety of domains.

The second experiment studied how human operators can best attain situational awareness in our scenario. The experiment examined the types and amount of information that should be maintained in order to minimize time to achieve situational awareness, given that the operator was attending a different task prior to the request for help. This disattention assumption is particularly valid for multiagent systems, as the human often will be assisting multiple agents and will need to reestablish situational awareness each time he moves his attention from one agent to another. The results show that the interface designer can make clear tradeoffs between the time to achieve situational awareness and the quality of the resulting understanding. In general, accuracy (i.e., quality of understanding) increases as more data is available to the subject. However, the time needed to attain situational awareness is not a monotonic function of the amount, nor type, of information available. Instead, our experiment showed that there is a clear point at which the time taken to absorb additional information outweighs the corresponding decrease in response time. The results from these two experiments bolster our contention that Sliding Autonomy can be an effective approach to robust control of multirobot teams.

II. RELATED WORK

A. Multiagent Robotic Assembly

The most common deployed multiagent assembly systems are in factories, where multiple industrial robots

are involved in the assembly of a product. A system of four industrial robots arranged around a conveyor network for material handling is described in [1]. The setup is typical of industrial applications, with stationary robots and inter-robot interactions limited to scheduling of shared resources (such as the conveyor system or temporary storage areas). Since time is of significant concern in factory settings, such industrial applications typically are managed by a static central controller. In contrast, our system is comprised of multiple mobile and stationary robots with only high-level central control of the team's overall tasks. They must flexibly coordinate their motions in order to complete the assembly task and adapt to a dynamic, uncertain environment; this requires close coordination between various combinations of heterogeneous robots, often involving more than one robot simultaneously manipulating the structure.

Coordinated assembly performed by teams of mobile robots is of prime interest to the space community. Stroupe *et al.* [2] use the CAMPOUT architecture to coordinate robots with purely behavior-based strategies to perform very tightly coupled tasks, similar to ours. Two homogeneous robots collaboratively carry a beam and position it with respect to an existing structure with sub-centimeter accuracy. Their agents jointly manipulate the beam, coordinating their motions through force/torque sensors in the manipulators. Throughout the task, the two agents perform symmetric versions of the same task. In contrast, our (very) heterogeneous team manipulates multiple objects to complete a larger, more complex structure, involving multiple interacting tasks with slightly greater docking tolerances. Although no two of our agents simultaneously manipulate the same rigid body, we use only a single sensor (a stereo camera pair on our dedicated sensing agent), which renders this sort of rigid interagent physical coupling infeasible.

B. Human–Robot Interaction

Recently, there has been significant interest in allowing human collaboration with robots in assembly scenarios. 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 guidance control. Typically, the human provides the force input, while the system steers the mechanism into the right place. The roles of the human operator and the system are clear and unvarying and both human and the system must operate simultaneously in order to accomplish the task.

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. While the robot was physically working alongside the astronaut, it was teleoperated by a remote operator in verbal communication with the astronaut. Unlike the complete teleoperation used in ASRO, our system allows

the remote user to take control of parts of the assembly task while leaving the remainder active under robotic control. In addition, the human and robots in our scenario cannot directly interact physically since, unlike [3]–[5], they are not collocated.

A system closely related to our approach to human–robot interaction is presented by Fong *et al.* [6], in which the robot and the user participate in a dialogue. The robot can ask the operator to help with localization or to clarify sensor readings. The operator also can make queries of the robot. However, that framework assumes that the robot will always perform the task: the operator’s role is only to provide the robot with any state information it is lacking. In contrast, our approach makes the operator an equal partner in the team, able to assume control of tasks or subtasks as appropriate.

C. Sliding, or Adjustable, Autonomy

Our use of the term Sliding Autonomy corresponds with the term Adjustable Autonomy as presented by Dorais *et al.* [7].¹ This paper provides several future examples in which Sliding Autonomy is essential for space operations where demands on the operator must be carefully selected and minimized, such as making the most efficient use of astronaut time during a manned mission to Mars.

Using a roving eye and a fixed manipulator, Kortenkamp *et al.* [8] developed and tested a software infrastructure that allows for sliding control of a robot manipulator. The task involved a pick-and-place operation during which Sliding Autonomy allowed the operator to recover from visual servoing errors, participate in high-level planning, and teleoperate the manipulator to complete tasks beyond its autonomous capabilities. Our paper extends this with a more complex assembly task that involves a team of robots and a finer granularity of Sliding Autonomy.

Nielsen *et al.* [9] present a system of three identical robots performing topological map building with the help of a human. The operator’s involvement ranges from teleoperation and landmark selection, to higher level directional commands, to the selection of regions of interest. With increasing robot autonomy, operator workload tended to decrease, but performance was best with maximum human involvement (i.e., pure teleoperation). In contrast to our system, the operator is constantly engaged and in control of the level of autonomy; the autonomous component never explicitly requests assistance.

Scerri has proposed an architecture for Sliding Autonomy applied to a daily scheduler [10]. This autonomous

¹Although 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.

system attempts to resolve timing conflicts (missed meetings, group discussions, personal conflicts, etc.) among some set of team members. Members are able to affect the system’s autonomy adjustment strategies by indicating their intent to attend gatherings or willingness to perform tasks.

Maheswaran *et al.* [11] describe a system of personal assistant agents that can operate under either user-based or agent-based autonomy. The entity in control (either an agent or a human) explicitly reasons about whether and when to transfer decision-making control to another entity (another agent or human). While our system currently does not allow control to be transferred from one robot to another, our approach to Sliding Autonomy allows much finer transfer of control between robot and human, as components of each task (such as monitoring and execution) may be transferred independently.

D. Situational Awareness

Over the years, there has been a large body of research on helping human operators maintain situational awareness. A significant amount of the initial research in this area focused on helping pilots maintain situational awareness while flying [12], [13]. The focus has shifted in more recent research to studying how human operators can maintain situational awareness while teleoperating robots, such as those in the search and rescue domain [12].

Our approach to Sliding Autonomy allows much finer transfer of control between robot and human.

This work is most relevant to systems where the operator is in constant contact with the system. In our domain, however, the operator may have periods where he is not in contact with the system; thus, we are also interested in helping the operator repeatedly attain situational awareness after being out of contact with the system.

Goodrich *et al.* [14] also study situational awareness in situations where the operator may not be in constant contact with the system. Their work examines the effectiveness of an operator when controlling a robot at different levels of autonomy, given increasing inattention to the robot. Their levels of autonomy include full autonomy, provision of goals by the human, provision of waypoints by the human, and safeguarded teleoperation. It attempts to facilitate such awareness by designing a usable interface for the operator. This is similar to our system, where the operator sometimes multitasks between interaction episodes, effectively ignoring the robots for that




Fig. 2. (a) Fully assembled four-beam structure. (b) View of beam being inserted into node. (c) Close-up of node being braced by the Crane.

period of time. They postulate a series of effectiveness/neglect curves characterizing the different levels of autonomy but do not present any experimental results in the cited work.

III. CONTEXT: SCENARIO, HARDWARE, AND ARCHITECTURE

A. Scenario

One area of robotics research is large-scale robotic assembly. Robots are especially useful in areas where humans are not well adapted, particularly hazardous environments, such as space, the Moon, or Mars. Building one robot complex enough to handle the entire construction task on its own is often either very difficult or impossible, especially for larger scale assemblies. Our approach is to use multiple heterogeneous robots that coordinate with one another to complete the task. While this increases complexity due to the necessary coordination [15], it also allows for more flexibility during task execution, as well as simpler and more manageable hardware.

The construction task used in our experiments involves four beams and four planarly compliant nodes that are assembled together into a square structure [Fig. 2(a)]. In a rough attempt at simulating conditions in space, the nodes are supported by casters that roll easily along the floor. A node must be braced before the end of a beam can be inserted into it; otherwise, the insertion forces can cause the node to roll away.

This task decomposes naturally into subtasks that can be completed by heterogeneous agents filling three different roles: an agent that provides information about the state of the world [the Roving Eye, Fig. 3(a)], an agent that braces the nodes during docking [the Crane, Fig. 3(b)], and an agent that does the manipulation and insertion of the beams into the nodes [the Mobile Manipulator, Fig. 3(c)].

To assemble one side of the square, the Crane first braces a node so that the Mobile Manipulator can insert one end of a beam into it. The Crane and Roving Eye then reposition to the other end of the beam, where the Crane braces another node while the Mobile Manipulator completes the node-beam-node subassembly. Once the

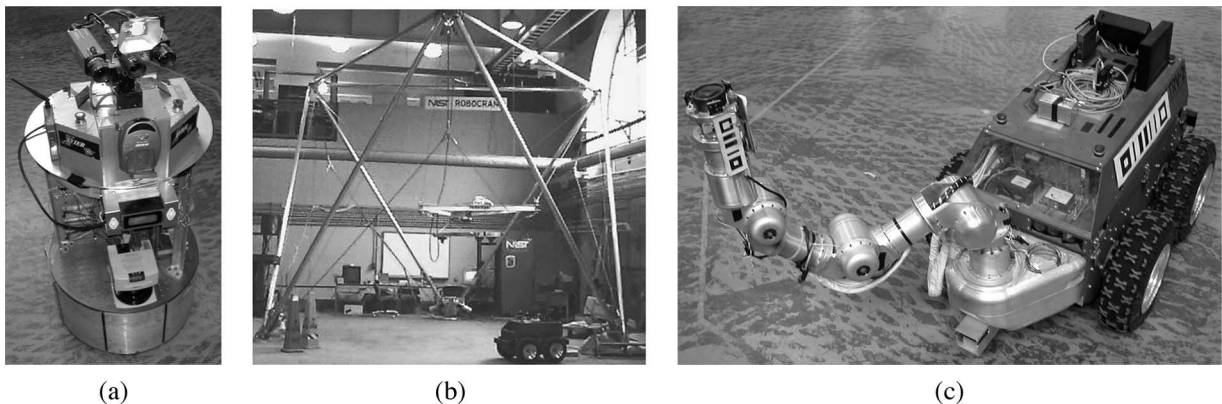


Fig. 3. Three robots used to build our square structure: (a) Roving Eye, (b) Crane, and (c) Mobile Manipulator.

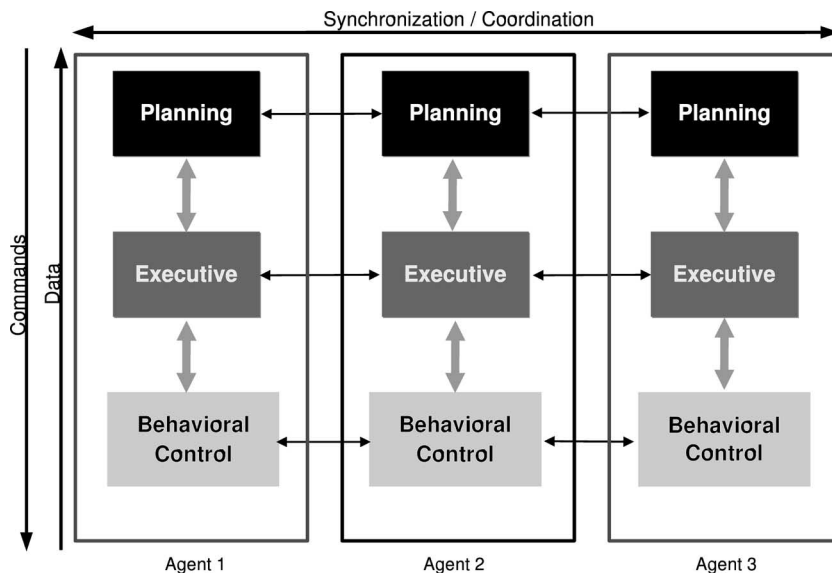


Fig. 4. Syndicate architecture. Note that each layer may communicate directly with all other layers at same level of abstraction; all links are not depicted for reasons of clarity (for instance, Agent 1's behavioral layer may communicate directly with Agent 3's). Our current implementation includes executive and behavioral layers.

beam is securely docked, the Mobile Manipulator releases it and moves to the next side of the square in order to receive the next beam and continue. This process repeats four times until the complete square is assembled. Running autonomously, the three robots can complete each side of the structure in 7.3 min, on average. The success rate per side for completely autonomous assembly is 75%. These statistics are from a set of 24 data points.

B. Hardware

Our Roving Eye [Fig. 3(a)] is Xavier [16], a synchronous robot built on an RWI B24 base and equipped with a stereo camera pair mounted on a pan-tilt unit. The Roving Eye's cameras are the team's only extrinsic sensors—the Crane and the Mobile Manipulator rely on the camera data in order to complete their tasks. By using an independent sensing agent such as this, we avoid conflicts that arise when an agent has multiple duties. For instance, cameras mounted on the Mobile Manipulator might become obscured when it is carrying a beam, forcing the robot to compromise between manipulation and sensing, and inevitably resulting in suboptimal performance.

Because it is not our focus of research, we simplified the vision problem by attaching fiducials (bar codes) to all important objects. The Roving Eye uses the fiducials to identify and locate the objects in the workspace relevant to the task at hand. Sample fiducials can be seen on the side and wrist of the Mobile Manipulator [Fig. 3(c)].

The Crane [Fig. 3(b)] is a NIST-built RoboCrane [17] and consists of a 6-degree-of-freedom (DOF) inverted Stewart platform carrying a simple mechanism that we

designed to allow the bracing of our nodes. The bracing mechanism is a hollow square that can be lowered onto the top of a node, effectively preventing the node from moving [Fig. 2(c)]. We chose the RoboCrane because it is strong enough to immobilize the nodes against sizable insertion forces without impinging upon the workspace of either the Mobile Manipulator or Roving Eye. This vertical bracing lessens the interference problems that often occur when multiple robots are moving in a tight, shared workspace.

The Mobile Manipulator [Fig. 3(c)] consists of a Metrica/TRACLabs 5-DOF anthropomorphic arm [18] mounted on the front of an RWI ATRV-2 skid-steered base. The arm has an electromagnet on its end effector that attaches to a metal plate fastened to the underside of each beam, allowing the Mobile Manipulator to grasp and release the beams.

C. Architecture

We developed the Syndicate architecture to support the closely coupled coordination between the robotic agents required to complete the assembly scenario. One of Syndicate's core features that eases this coordination is its support of a three-layered approach. Each layer is associated with a different task granularity and level of abstraction about the world and may communicate with the layers immediately above and below it. In general, higher and more abstract layers command lower, more reactive, layers, while the lower layers provide the data needed by the higher level decision process (Fig. 4).

Although the Syndicate architecture supports three layers of abstraction, our current implementation only encompasses two. The bottom-most is the behavioral layer, which deals with fine-grained stateless control of the robots. This layer, which is based on the Skill Manager [19], acts as the interface between the controllers/hardware and the executive layer. Thus, it views the world at a small temporal granularity, concerning itself with the hardware and environmental details that higher layers abstract away. By doing so, it is able to react very quickly to changes in the world.

The executive layer is responsible for building and maintaining the agents' hierarchical task trees, thus managing all of the stateful task-level aspects of the system [20]. 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. Many of the tasks rely on data passed up from the behavioral layer to decide when to change state; these changes are then propagated back down to the behavioral layer in order to manipulate the physical world. This relationship illustrates a secondary objective of the executive layer: configuring the behavioral layer based on the current state of the task tree. The top-level planning layer is not yet implemented in our system. Instead, the executive layer contains a fixed task ordering for a given scenario. We are currently investigating different approaches to building a planning layer for Syndicate, which will likely include both planning and scheduling aspects, such as those used in [21].

In addition to communication between the layers of a single agent, Syndicate supports communication between agents at each abstraction layer. This gives agents the ability to directly coordinate multiagent tasks with other agents at all levels of the hierarchy. For example, in the executive layer, we want to ensure that the Roving Eye does not move to the next corner until the Mobile Manipulator is finished with the current beam-node docking operation, since the Mobile Manipulator depends on the raw position data provided by the Roving Eye to complete the docking. To accomplish this, Syndicate enables the executive layer on one agent to constrain its task execution with respect to a task on the other agent, in this case, sequencing two tasks that run on different agents. Coordination at the behavioral level is also required by many tasks. For instance, to dock a beam into a node, the Mobile Manipulator's behavioral layer must communicate with the Roving Eye's in order to receive raw position information, forming a distributed visual servoing loop [22]. Although our planning layer is not yet implemented, an example of coordination at that level can be seen in the FIRE project [21], where communication between peer planning layers is used for auction-based task assignment.

The Syndicate architecture provides a framework upon which we can build to support concepts such as Sliding Autonomy in complex multiagent scenarios.

IV. SLIDING AUTONOMY FOR MULTIAGENT TEAMS

Even with closely coupled coordination, it is nearly impossible to prevent errors from occurring in complex multiagent tasks such as our construction scenario. Because of the complexity and uncertainty involved in such domains, even a highly specific well-programmed robotic team can sometimes fail when operating autonomously. One common way to compensate for this possibility is to use pure teleoperation. In this mode, a human operator controls each of the robots during all of its tasks. Both teleoperation and pure autonomy have been shown to have distinct strengths and weaknesses. In general, teleoperation is slower, but more reliable, while full autonomy is faster, but less robust [23]. Additionally, communication latency, prevalent in space-based applications, make pure teleoperation difficult and tedious. The goal of Sliding Autonomy is to allow human-robot teams to move smoothly along the spectrum from teleoperation to autonomy, making appropriate use of the differing capabilities of team members in order to outperform both pure teleoperation and pure autonomy.

Sliding Autonomy for multirobot teams, however, has complications not found in single-robot systems. In general, we have found that there are three ways in which multiagent Sliding Autonomy is more demanding than the single-agent version: 1) deciding when to ask for help, as the human is not guaranteed to be monitoring any one robot at any given time; 2) providing situational awareness of the requesting robot's workspace when the operator is asked to help; and 3) maintaining coordination of the team as a whole when the human is controlling one of the robotic agents.

A. Our Approach to Sliding Autonomy

Our approach to Sliding Autonomy uses mixed-initiative interactions to orchestrate the collaboration between the human operator and the robotic agents. These interactions consist of either the human or the autonomous system deciding who is in control of different aspects of the task. The goal of these interactions is to optimize metrics such as efficiency and robustness, measured by time elapsed and likelihood of success, respectively.

We introduce the notion of "task switching" to implement these types of interactions. Nearly every task in the executive layer decomposes into separate "monitor" and "action" components, each of which is itself a task that may be controlled by either the human or the autonomous system. The monitoring component of a task is responsible for detecting failures and determining when the task has been completed. The action component interacts with the robot controllers via the behavioral layer to perform the desired actions. While the majority of tasks are split in this fashion, some nonleaf nodes of the task tree that are merely responsible for creating other tasks, and so have no obvious action/monitor split, are not divided.

This methodology yields significant flexibility. For example, the autonomous system is often able to perform a task where its sensing capabilities are not sufficient to reliably determine when the task is completed. In such a case, the human can be assigned the monitoring portion of the task, in order to ensure it is reliably executed. An example is the docking task carried out by the Mobile Manipulator. While the autonomous system is quite proficient at completing the task, the final clearance between beam and node is on the same order of magnitude as the noise in the Roving Eye's sensing system, on occasion making it difficult for the autonomous system to determine when the docking is complete. The human operator can look for additional clues in the video feed and often can make a more reliable determination of when the docking has successfully completed or has become stuck.

On the other hand, the action subtask may be given to the human if the autonomous system detects that it is having problems or believes the operator will be able to perform it more efficiently. In such a case, the human is asked to complete execution, while the system monitors his progress. A prime example of this is the search task carried out by the Roving Eye when it loses sight of a fiducial of interest. Humans are better suited to the action component of this task, as they can decide where to move the camera based on their understanding of the workspace, as opposed to the blind grid-based search pattern the autonomous system uses. However, it is not obvious to the human when the system has detected a fiducial, since the human's vision is so much different than that used by the Roving Eye. Thus, in this instance, the autonomous monitoring component can work side by side with the human to complete the search task.

Regardless of whether a task is switched by the operator intervention or the autonomous system, it interacts as needed with the behavioral layer to switch all necessary low-level components to the appropriate operational mode. This provides fine-grained control over operations, as the human can assume control of very specific portions of the system, while leaving the remainder under autonomous control. For instance, if the user were to take control of the visual search task on the Roving Eye, which looks for fiducials in the environment, the pan-tilt unit would be placed under human control, while control of the Roving Eye's base and responsibility for determining when the search was successful would remain under autonomous control. At the completion of a task, all affected behaviors are returned to the operational mode in which they were prior to the switch.

This approach to Sliding Autonomy still leaves open the questions of when to switch and who decides to do so. In order to investigate how best to structure these interactions, we have experimented with two different approaches, which span the control spectrum, but vary in how tasks may be switched. System-Initiative Sliding

Autonomy (SISA) enables the robot team to ask the human for assistance but does not allow the human operator to interrupt the robots. SISA models situations where the human operator is multitasking and attending to other responsibilities while the robots attempt to operate autonomously, so the human has no ongoing knowledge of what is occurring in the robots' workspace. The operator's attention is needed only when the team decides that the human is better suited for the execution of some task or when an individual robot has encountered an error condition from which it does not believe it can recover on its own. While facilitating multitasking may raise the productivity of the overall system, the time it takes users to regain situational awareness of the robots' workspace when asked to help cannot be disregarded (Section IV-B). We hypothesize that SISA is the best fit when human resources are scarce and the autonomous system is reasonably proficient at error detection.

We also investigated Mixed-Initiative Sliding Autonomy (MISA). Here, while the robotic team members still can ask for help, the human operator also has the option of interrupting the system to take control of a task or subtask's action and/or monitoring component. While this has the potential to increase robustness, since the human operator can intervene before a robot makes a fatal error, such gains are realized only if the human actively monitors the progress of the robotic team.

The next three sections describe our approaches to dealing with the three aspects that are endemic to multi-agent Sliding Autonomy: requesting help, gaining situational awareness, and maintaining coordination.

B. Requesting Help: System and User Modeling

In many Sliding Autonomy systems, only the human operator is able to change the control of tasks. In general, this is adequate for single-robot applications, since a human operator generally is more than capable of monitoring the status and progress of a single robotic agent and can take over either if he feels he would do a better job than the robot or if the robot is entering into a dangerous situation. As the number of robots increases, however, it becomes harder and harder for a single operator to keep track of the status of all agents. One way of addressing this problem is to allow the robots to switch the control of tasks to the human by asking for assistance as appropriate. Thus, a robot in trouble can request help from the human, instead of waiting for the operator to realize that there is a problem.

We have developed an approach that enables the autonomous system to make reasoned decisions about when to switch control of tasks based on current conditions and the specific operator available. These decisions are made when tasks are initially launched as well as when a failure is perceived. If the operator has shown he is usually better than the autonomous system, the task will be

If the operator has shown he is usually better than the autonomous system, the task will be assigned to him immediately.

assigned to him immediately. Alternatively, the system may try to perform the task itself but later decide to hand control to the human if it believes it is unable to complete the task. Such decisions need to be made in both SISA and MISA modes, as the system may request help with a task in either case.

While these task allocation decisions could be made via an arbitrary set of heuristics (such as “try twice autonomously, then cede control to the human”), such a strategy is potentially suboptimal, since no heuristic is able to address all exceptional cases. Instead, by building empirical performance models of both the autonomous system and the human operators, we have developed a principled approach to making such decisions that allows the autonomous system to reason about expected task duration and agent reliability to decide who should perform what task.

The decision problem can be phrased as a comparison between the expected time to complete the task if the human makes the next attempt and the expected time if the autonomous system does so. In order to make this comparison, we evaluate two decision trees—one where the human performs the task under consideration and one where the autonomous system controls the task. An example of a tree associated with the autonomous system making the next attempt is diagrammed in Fig. 5.

There are three components to this recursive prediction: the probability of success for a given party’s attempt at accomplishing the task in question, the expected time to complete the attempt given success, and the expected time given failure. In Fig. 5, these correspond to the probability of branch (a), timespan (b), and timespan (c), respectively. To estimate these values, we use prior observations of execution time, conditioned on the failures which have occurred during this instance of the task. The number of preceding failures is roughly equivalent to the current level of the decision tree and is used because a failure is empirically a good predictor of future failures (at least within our scenario). In addition, we condition our expected time calculations on the outcome of the attempt, as failures often take significantly longer than successful attempts. Since our models are updated during task execution and are maintained on a per-operator basis, the system’s decisions will dynamically change in response to the operator’s current performance and will depend on the specific operator available.

By conditioning execution time on previous failures, we turn what was a multimodal distribution into a set of (more or less) unimodal distributions, which greatly eases the calculation of expected time. We estimate how long it will take to complete a task directly from these distributions

$$E(t_s|F_r = i, F_h = j) = \min \left(\begin{matrix} E(t_r|F_r = i, F_h = j) \\ E(t_h|F_h = j, F_r = i) \end{matrix} \right) \quad (1)$$

$$E(t_r|F_r = i, F_h = j) = P(S_r|F_r = i)E(t_r|S_r, F_r = i) + ((1 - P(S_r|F_r = i)) \times (E(t_r|\neg S_r, F_r = i) + E(t_s|F_r = i + 1, F_h = j))) \quad (2)$$

$$P(S_r|F_r = f_r + 1) = 1.0 \quad (3)$$

$$E(t_r|S_r, F_r = f_r + 1) = E(t_r|S_r, F_r = f_r) \quad (4)$$

where

$E(t_s|F_r = i, F_h = j)$ Expected time to complete the task, given i preceding autonomous failures and j preceding human failures. For $i = j = 0$, this corresponds to the expected time of the entire decision tree (timespan (d) in Fig. 5).

$E(t_r|F_r = i, F_h = j)$ Expected time to complete the task if the autonomous system performs the next attempt, given i preceding failures. This is the expected time of one subtree of the decision tree, such as the expanded subtree depicted in Fig. 5.

$P(S_r|F_r = i)$ Probability of the autonomous system successfully completing the task, given i preceding failures. This corresponds to the probability of branch (a) in Fig. 5.

$E(t_r|S_r, F_r = i)$ Expected value of the distribution formed by all data points in which the task was successfully completed by the autonomous system with i preceding failures (timespan (b) in Fig. 5).

f_r Maximum number of preceding failures by the autonomous system that have occurred in practice. $f_r + f_h$ is the maximum depth of the decision tree.

The minimization in (1) represents the decision about whether to assign the next attempt to the human or the autonomous system. Equation (2) weights the time taken to succeed and the time taken to fail plus the remainder of the decision tree by the probability of success or failure, respectively. While the decision tree could, in theory, continue indefinitely, our model contains a finite amount

task queue and include his expected time to service the queue when calculating the expected time to complete the task under consideration if the human performs the next attempt.

While currently we are simply comparing task execution times, we have considered using a full cost model in order to incorporate varying labor costs, the amortized cost of hardware, continuing expenses associated with teleoperation or autonomous control of the robots, the cost of repairs, etc. The specific cost function would be highly dependent on the particular domain in question but would be parameterized at a minimum by $E(t_h)$ and $E(t_r)$ as well as domain-specific price or cost parameters. This is an area for future work.

C. Operator Situational Awareness

Another important issue in Sliding Autonomy is attaining operator situational awareness. This is particularly critical in multiagent domains. In the single-agent domain, the operator needs to monitor only a single robot and workspace. Even if he does lose situational awareness due to attending to other tasks, it is easier for the operator to remember the state of the system and use that to assist him in attaining situational awareness the next time he is asked for help. In the multiagent domain, there are many robots with the ability to ask for help. Not only is there a potentially longer time before an operator assists a robot a second time, but the operator also has more than likely provided assistance to other robots in the interim, further degrading situational awareness.

Our approach is to maintain a buffer of information of the state of the robots' workspace(s), including both synthesized and raw video views, and show these buffers to the operator when he is asked to assist the system. By viewing these buffers, the operator can more quickly gain situational awareness of the pertinent workspace and more efficiently assist the robots in their task than he could from viewing just the current state of the system.

However, what types of information and how much should be shown remain open questions. There is an obvious tradeoff between performance and reliability; having more information and a larger buffer typically leads to a more accurate assessment but increases how long it takes to attain situational awareness. Section VI describes an experiment we performed to help quantify this tradeoff.

D. Maintaining Interagent Coordination

The third issue that arises while adapting Sliding Autonomy to the multiagent domain is how best to maintain interagent coordination of the robot team while one or more of its members is under operator control. We address this by enabling the robots to monitor the progress of their own and other related tasks through monitoring subtasks, updating the execution of their tasks accordingly. For ex-

ample, consider a situation from our scenario in which the operator is asked to brace a node with the Crane. While the operator is performing this task, the Roving Eye continues to monitor the location of the Crane and recognizes when the operator has finished. Once the Roving Eye determines that the operator has performed the bracing, the system automatically continues with the tasks that depended on the completion of the bracing task: the Mobile Manipulator approaches the node with the beam, the Roving Eye changes its focus to the tip of the beam, and the Crane stays where it is in order to continue bracing the node.

This approach allows the team to remain coordinated during a human intervention, as long as the human does not deviate too greatly from the system's plan (e.g., performing tasks in an unexpected order). Since the existing monitoring capabilities primarily monitor for task completion, the human can take whatever course he wishes to complete the task at hand when in control of the action component of a task, and the rest of the agents under autonomous control will remain coordinated. However, if the human decides to accomplish additional tasks first, undo previous work, or make arbitrary modifications to the plan, the current system will lose coordination. Open research questions in this area include the development of action recognition, plan prediction, and cooperative plan generation capabilities for the autonomous system.

E. Summary of Multiagent Extensions to Sliding Autonomy

To summarize, we have extended the traditional single-agent Sliding Autonomy approach to multiagent teams. In the course of our work, we have discovered three primary differences between performing Sliding Autonomy with single agents and multiple agents: 1) Due to the human's inability to simultaneously monitor all of the agents in the team, the autonomous system must reason about when to request help and cannot rely on the human to step in. This increases the importance of detecting that an error has occurred and makes user models a necessity. 2) In addition, the human's inability to monitor all of the agents simultaneously results in a loss of situational awareness between requests for help. Consequently, the human must first attain situational awareness before assisting the robotic team, a step that typically does not need to occur in single-agent Sliding Autonomy. 3) Finally, the multiagent autonomous system must maintain interagent coordination, even when one agent is under the control of the human.

We have conducted an experiment to compare the effects of teleoperation, MISA, SISA, and pure autonomy control strategies on robustness and efficiency and report on the results next. In addition, we carried out an experiment to determine how different interfaces affect the human's ability to attain situational awareness and quantify how long he takes to do so, which can be used to refine the user model.

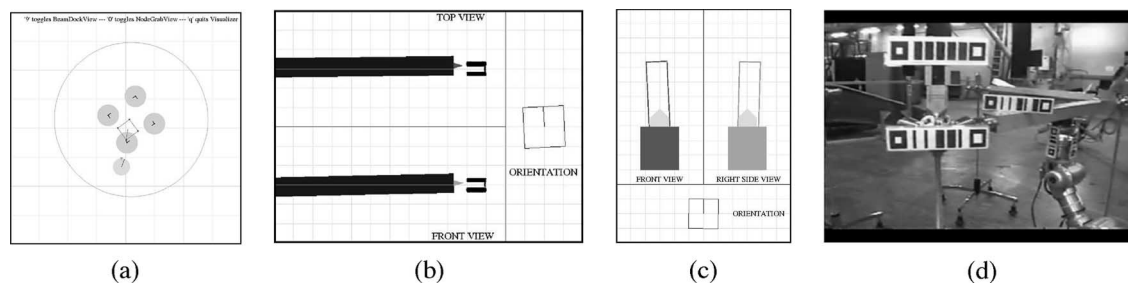


Fig. 6. Screen shots of visualization tools used for controlling robots. From left to right: bird's eye overview of workspace, closeup used during beam docking, another closeup used during node bracing, and raw video feed provided by the Roving Eye.

V. MULTIAGENT SLIDING AUTONOMY EXPERIMENT

This experiment investigated the effects of mixing full system autonomy with teleoperation on the overall efficiency and robustness of our multiagent system.

A. Methodology

For this paper, expert users of the system performed a number of trials for each mode of interaction with the system. We had initially hoped to use operators unfamiliar with the project but found that the time needed for novices to become proficient with the system was excessive: after 14 h of training, the novice users were still slower than experts by a factor of three. Since we needed our subjects to be roughly on par with the performance of the autonomous system in order to obtain meaningful results, we changed the study to employ two project members as subjects.

During an earlier pilot study, we determined that the time taken for the system to complete a node–beam–node subassembly was independent of the side of the structure by using a double-sided t-test with a 0.95 confidence threshold. Therefore, our unit of analysis was a one-beam subassembly, as detailed in Section III-A. Each subject performed this assembly task 12 times in teleoperation mode, 16 times in SISA, and 20 times in MISA. The 48 trials were performed in random order. Together with 24 runs of the autonomous system, that gives a total of 120 datapoints.

In order to create a semirealistic teleoperation experience, the subject sat at a workstation facing away from the robots and the workspace. She was able to see only the raw video output from one of the Roving Eye's cameras and the output of a visualization tool, which displays depth information relevant to the current task (see Fig. 6), as provided by the Roving Eye. Control input was provided through a "Space Mouse" [25], a puck-like device designed for three-dimensional (3-D) CAD work which allows the user to push, pull, and twist a knob similar to a hockey puck to provide full 6-DOF control.

During teleoperation runs, subjects logged which subtask of the overall assembly they were working on in order to enable timing, as autonomous monitoring/timing was not available during pure teleoperation. This selection was used to extract timing information for each user to initialize her user model for the Sliding Autonomy trials. For sliding and full autonomy, the timing information was logged automatically by the system. At the end of each run, the subject completed a NASA-TLX survey [26] to assess her perceived workload while controlling the robots. This survey takes into account factors such as mental, temporal, and physical demand, as well as effort and frustration.

In order to emphasize the difference in initiative between the two Sliding Autonomy modes, the users performed a distractor task while using SISA. The screen of their workstation and the video feed from the Roving Eye were turned off. The only feedback available was audio cues asking for help. During MISA, the users were not presented with a distractor task, instead they actively and constantly followed the system's progress. They were told to exercise their ability to proactively take over control whenever they saw the opportunity to complete a subtask faster than the system.

B. Results

Table 1 summarizes our hypotheses about the effects of introducing Sliding Autonomy into our system on three relevant metrics. With the exception of several cells highlighted in bold in Tables 2 and 3, the results we obtained agreed with our expectations. For instance, teleoperation (729–911 s) took 1.5–2 times as long as any of the modes involving autonomy (437–492 s). The times and TLX results of failed runs were not included in the results.

Both users generally followed the hypothesized trend that the success rate would increase as human involvement increased; pure autonomy was the worst at 75%, followed by SISA, MISA, and finally teleoperation, which had an average 96% success rate. Note that the users did not give up on the failed teleoperation runs; instead, the system reached a point of terminal failure.

Table 1 Expected Results for Sliding Autonomy Performance

	Time to Completion	Success Rate	Human Workload
Teleoperation	worst	best	worst
Autonomous System	best	worst	best
System Initiative	good	good	better
Mixed Initiative	better	better	good

Table 2 Results for User 1. Discrepancies Compared to Expected Results are Highlighted in Bold

User 1	Mean Time to Completion [standard deviation]	Success Rate (# of trials)	TLX Workload [standard deviation]
Teleoperation	729 [139] seconds	100% (12)	42 [10]
Autonomous System	437 [94] seconds	75% (24)	0
System Initiative	462 [63] seconds	75% (20)	16 [16]
Mixed Initiative	492 [140] seconds	81% (16)	17 [11]

Table 3 Results for User 2. Discrepancies Compared to the Expected Results are Highlighted in Bold

User 2	Mean Time to Completion [standard deviation]	Success Rate (# of trials)	TLX Workload [standard deviation]
Teleoperation	911 [193] seconds	92% (12)	71 [9]
Autonomous System	437 [94] seconds	75% (24)	0
System Initiative	458 [106] seconds	85% (20)	50 [14]
Mixed Initiative	445 [72] seconds	88% (16)	34 [11]

Not surprisingly, due to the subjective nature of the NASA-TLX workload survey, the actual values varied strongly between users; although, trends were consistent. Teleoperation had, by far, the highest workload, approximately twice that of the Sliding Autonomy cases (the right-most column of Tables 2 and 3).

Comparing our two users side by side with the autonomous system, we note a completion time comparable with that of the autonomous system for both forms of Sliding Autonomy and a much longer completion time for teleoperation [Fig. 7(a)]. Additionally, the teleoperation time to completion shows differences between the two users. In Fig. 7(b) we see an upwards trend of the success rate proportional to the amount of human involvement

during the assembly task. Finally, Fig. 7(c) shows that trends in the perceived workload measured by the TLX survey are not consistent between the two users—the second user found SISA to be frustrating, while the first reported a low workload. The autonomous system is omitted from this chart since no human operator was involved and hence there is no reported workload.

The histograms of User 1’s performance (Fig. 8) show again that the autonomous modes are much faster than teleoperation. The fastest run time was recorded under autonomous operation at just over 300 s. At the same time, a large portion of the teleoperation runs took over 800 s, much longer than the slowest run during an autonomy trial.




Fig. 7. Comparison of our two subjects’ independent runs and fully autonomous system. Completion time and success rate generally follow our expectations, but perception of workload is very task and user dependent.

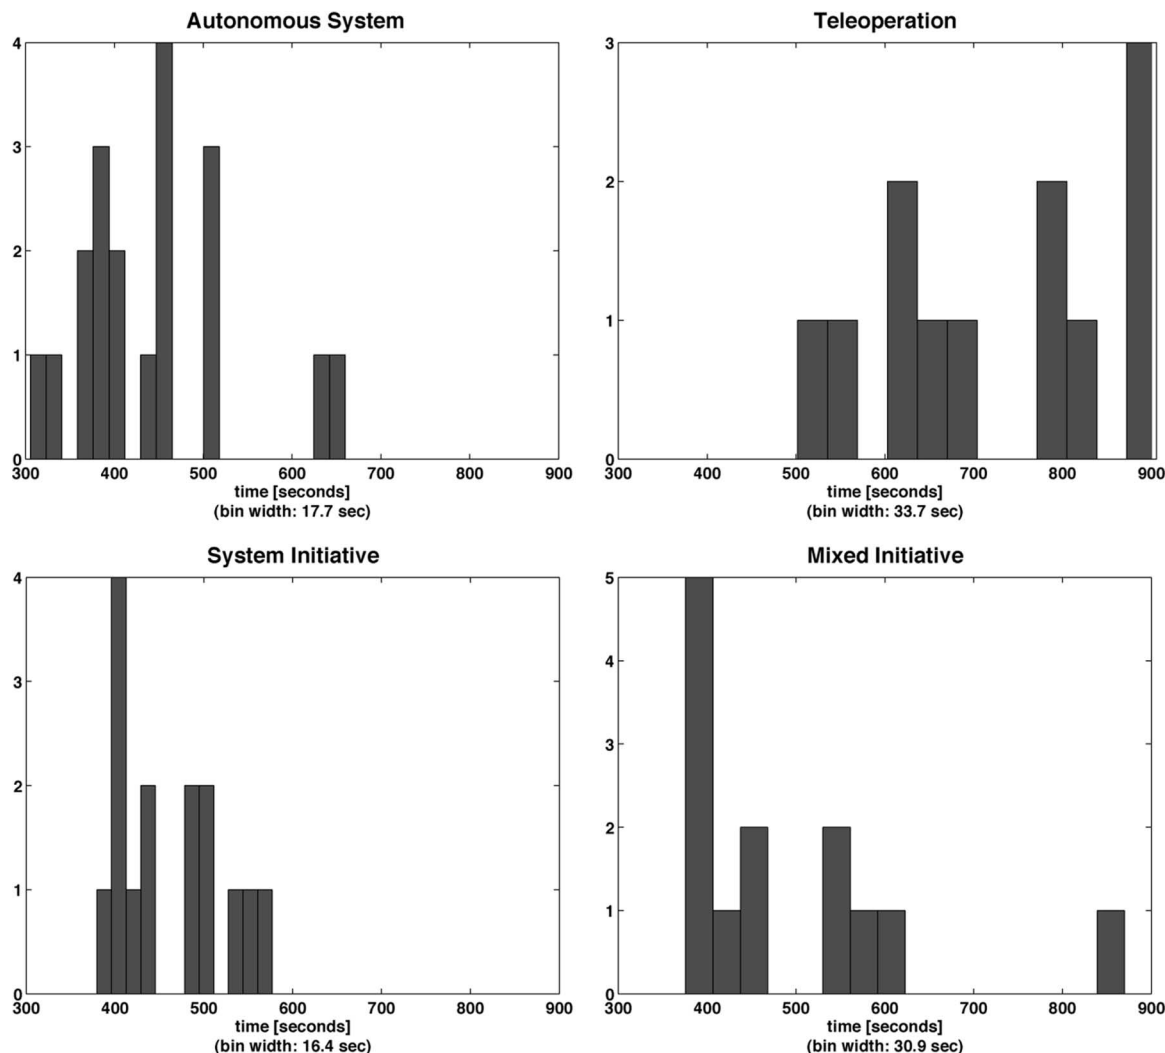


Fig. 8. Histograms comparing User 1's performance when using different Sliding Autonomy modes. Horizontal axis marks completion time in seconds, and height of bars shows number of experimental runs for each time interval.

C. Discussion

When comparing the two extremes of the autonomy spectrum (pure autonomy and complete teleoperation), it is clear that there is an inherent tradeoff of speed against robustness. If we were willing to allow an increase of 50%–100% in the time needed to complete the structure, we could achieve near perfect reliability by simply allowing the human to teleoperate everything. However, our workload data indicate that in addition to the significantly increased time to complete the assembly, operators would swiftly become mentally overloaded in addition to being unable to multitask during assembly.

Our experimental results suggest that this dilemma can be resolved by employing some form of Sliding Autonomy. As is shown in Fig. 7, adding any amount of autonomy reduces the completion time to a level comparable with the purely autonomous approach, while

adding any amount of human involvement increases the reliability of the overall system. Sliding Autonomy sits in the intersection of these two trends. It combines the low completion time advantage of autonomous operation with the increased success rate due to human involvement. The increase in system reliability when Sliding Autonomy is used can be attributed to the operator's intuition and ability to quickly understand problematic conditions and then initiate recovery measures to help the system avoid failure conditions. In addition, there is a clear benefit to introducing Sliding Autonomy from a workload perspective. Since the system still performs the tasks it is skilled at autonomously, the perceived workload reported by our users was significantly lower than during pure teleoperation.

The subjective nature of the TLX scores does not allow a direct user-to-user comparison, but we can make

statements about trends in the data. We expected SISA to have the lowest workload because the user has to work with the system only when asked for help. However, there are at least two factors that are not captured well within the TLX framework: boredom and the frequency of interruption by the system. Both of these are directly linked to the specific user's ability to perform the task at hand. For example, a relatively strong user will likely be asked for help with many tasks, to the point where she is called back to help with the next task almost as soon as she returns to her distractor task from assisting the team. On the other hand, a relatively weak user will not be asked for help very much; she may be called to help only a few times during a 15-min experimental run. Depending on the individual user, this could be interpreted either as a relaxing situation with low workload or as a frustrating situation where the individual has to be on call for the system but is never asked to do anything. On the other hand, an experienced user may be frustrated due to being constantly interrupted by the system asking for help. This frustration is likely responsible for high workload results such as the data shown for User 2 in Fig. 7.

The multimodal grouping of the completion time results shown in the histograms in Fig. 8 corresponds to assembly attempts with varying degrees of success. The left-most group represents smooth runs without any failures, and subsequent groups indicate increasing numbers of autonomous or operator errors and/or near-failure conditions.

There are three discrepancies between our hypotheses and our actual results: User 2's workload, User 1's mean time to completion in the MISA and SISA cases, and User 1's success rate under SISA (the bolded entries in Tables 2 and 3). The high workload of User 2 can be explained by the subjective nature of the TLX survey, as discussed. User 1 took longer on average to complete the task under MISA than SISA (although there is not a statistically significant difference between the means). This may be due to an overeager operator intervening in cases where the autonomous system is, in fact, more efficient, resulting in an overall decrease in efficiency. Finally, User 1's success rate under SISA is not superior to the autonomous system's. Success rate improvement between pure autonomy and SISA is dependent upon the autonomous system's ability to discern failures with which it requires assistance; if few discernible failures occur, little gain will be realized by applying SISA.

Our experiment shows that adding a human agent to a multirobot team via a Sliding Autonomy framework combines the advantages of autonomous robot operation with the reliability of teleoperation at a mental workload level tolerable by the operator. The amount of attention the operator pays to our task (i.e., SISA versus MISA) has no measurable effect on the time to task completion compared to fully autonomous operation, but it does manifest itself in overall system robustness: the greater the op-

erator's involvement in the team's operation, the higher the success rate.

The choice between forms of Sliding Autonomy (SISA or MISA) depends heavily on the system as a whole. If the operators are comparable in skill level with the autonomous system, the system is able to perform significant portions of the task on its own, and the system is able to detect the existence of an error reliably, then the humans can productively multitask when operating under SISA. In a situation involving many robot teams, SISA may be preferable to MISA, as its slightly lower success rate is outweighed by the ability of a few operators to simultaneously oversee a number of teams. The additional time required to attain situational awareness can be minimized by selecting the most effective user interface (see Section VI).

For comparatively weak human operators, a similar multitasking argument holds, but a higher degree of autonomy is required of the system since the operator's ability to help can be rather limited. For very skilled humans, however, their abilities often lead to them being continuously asked for help. Switching between different teams performing different tasks very often and rapidly is more confusing and stressful to the operator than helpful to any team. Instead, for strong operators, a MISA setting is often preferable to SISA, despite the same requests for help, since it allows the human more control over her workload. Additionally, if the autonomous system were unable to detect most failures, MISA would be the preferred method in all cases in order to compensate for the autonomous system's lack of reliability.

VI. SITUATIONAL AWARENESS EXPERIMENT

In addition to experimentally evaluating the differences between our different approaches to multiagent Sliding Autonomy, we also conducted an experiment to test how well users attain situational awareness in our robotic system. In order to help the operator gain situational awareness more quickly, we maintain a buffer containing the state of the system over the last n seconds, which can be replayed to the operator when the system asks for help. Two natural questions arise: "What kinds of information should be included in the buffer?" and "How much data should it save?". To help answer these questions, we tested how quickly users attained situational awareness of our system when using various combinations of displays and data buffer lengths. In addition to guiding the design of operator interfaces, this experiment provides the information needed to account for the time necessary to gain situational awareness in the system's model of the human (Section IV-B).

A. Methodology

The experiment tested four different combinations of information streams. The first was simply a video feed

from one of the Roving Eye’s cameras (Fig. 9.2). The second was the Roving Eye video along with the synthesized “technical drawing”-style visualizer (Figs. 6 and 9.1), showing the relative positions of the beam and node from above and the front of the beam. The third was the Roving Eye video along with two other video feeds—one from a fisheye camera placed on top of the Crane, looking down (Fig. 9.4), and one from an external camera placed outside the robot workspace, looking towards the structure (Fig. 9.3). Finally, the fourth combination included all four of the previous elements.

We also varied the length of the data feed that was presented to the subjects. The four possible lengths were 0 (still shot), 5, 10, and 20 s. This, in combination with the four different displays, yielded a total of 16 different test conditions.

During each trial, each of the 32 subjects was shown the data buffer from an attempted docking and was asked to identify through a dialog (Fig. 9.5) why the robot requested help. The experimental procedure was a combination of training and testing. The subject’s training began with reading a written overview of the task and hardware at hand, with the experimenter answering any questions. The subject was then shown one example of each of the seven types of errors via the graphical interface (Fig. 9), using the maximal data and 20-s playback condition.

After training, the subjects began the actual experiment. They were instructed to play a video game in between trials to simulate multitasking. Each user was tested on four of the 16 conditions, performing six trials per condition. We applied Latin squares to both the data feed and length effects to account for ordering and practice effects. A Latin square is a statistical technique that allows

experimenters to test effects while controlling for two other known sources of variation (here, intersubject variability and ordering effects).

B. Results

During this experiment, we recorded the time that elapsed between the beginning of playback and when the user chose which error she believed had occurred. We also recorded her response to allow analysis of user accuracy. Users were not allowed to choose a response until they had watched the entire playback clip. Also, although users were able to change their answer, we only considered the data collected from their final responses. We used a univariate ANOVA test to analyze this data.

The data show that in our scenario, with respect to response time, the best display and data feed length combination is the Roving Eye video plus visualizer, viewed for between 5–10 s [Fig. 10(a)], as this combination had the shortest average response time. Considering solely data feed length, significant differences were found between 0- and 5-, 0- and 10-, 5- and 20-, and 10- and 20-s playbacks. If instead we consider the composition of the display, significance was found between the Roving Eye and the Roving Eye plus other videos displays, the Roving Eye and all displays, the Roving Eye plus videos and Roving Eye plus visualizer displays, and the Roving Eye plus visualizer and all displays.

With respect to accuracy, the best data feed condition is that with the longest (20 s) data feed length [Fig. 10(b)]. Similarly, the display that had the highest accuracy was the Roving Eye video plus other videos [Fig. 10(b)], although it is not significantly different from the all displays condition. Considering data feed length, significance was

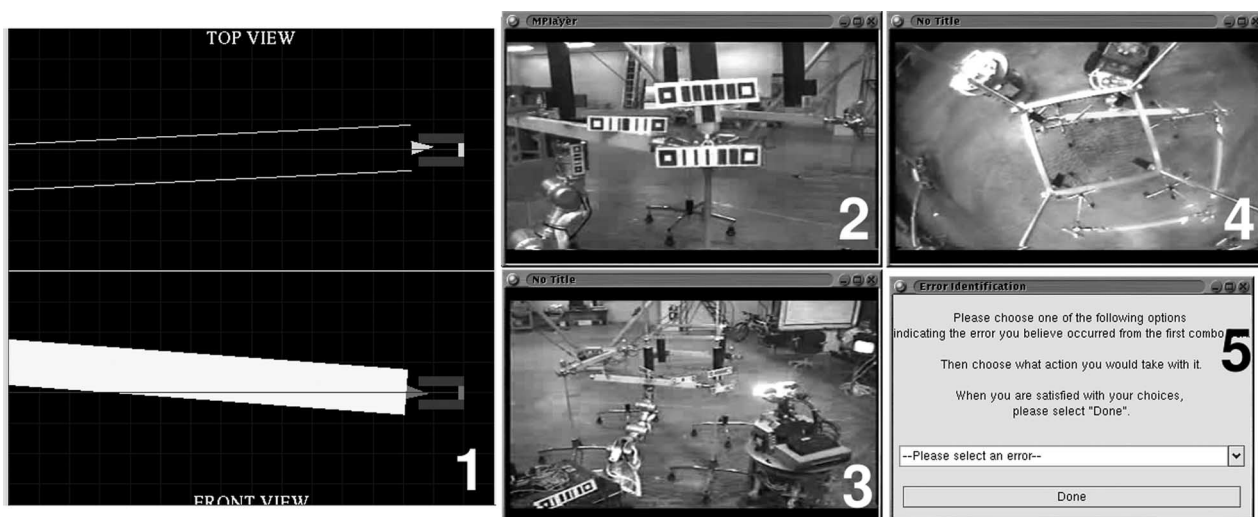


Fig. 9. Subject interface for situational awareness experiment, including three video streams [Roving Eye’s cameras (2), external camera (3), and Crane-mounted camera (4)], synthesized view of beam and node (1), and error categorization input dialog (5).

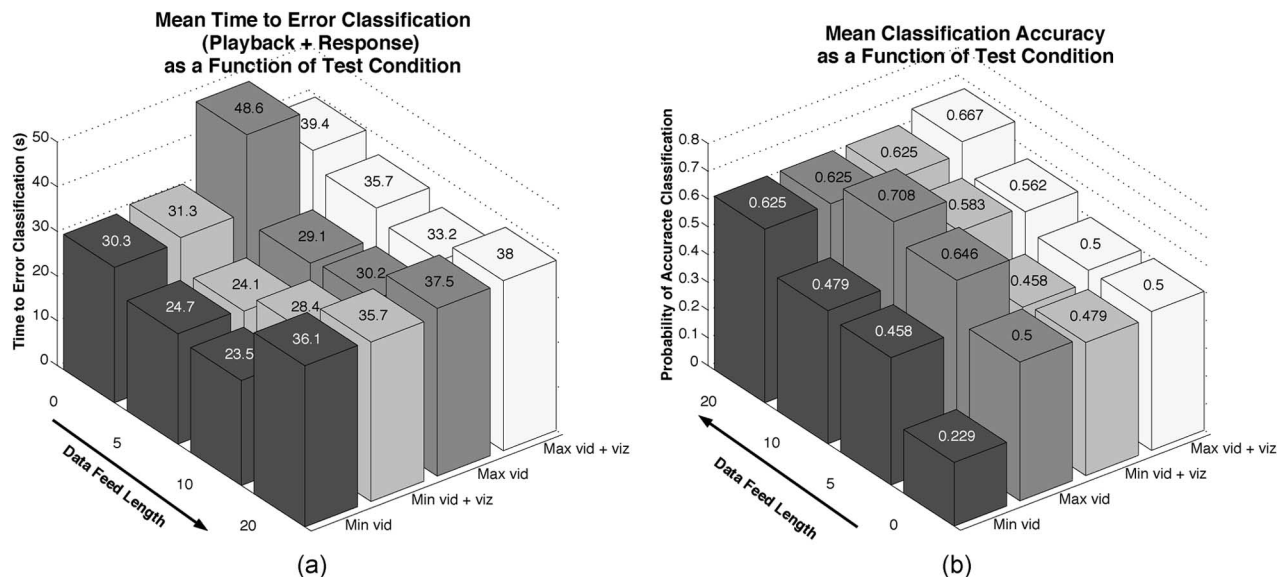


Fig. 10. Effects of available information and data feed length on (a) how long subjects took to choose an error and (b) how accurate their decisions were. Note that ordering differs between two figures, although shading is consistent.

found between 0- and 10-, 0- and 20-, and 5- and 20-s playbacks. Significance in display composition was found between the Roving Eye and Roving Eye plus videos displays, as well as the Roving Eye and all displays.

C. Discussion

If response time is the only metric under consideration, the results are clear—the best conditions were the Roving Eye video alone or the Roving Eye video plus visualizer, with 5–10 s of playback. We believe that this is because as more raw videos are added, the mental overhead required to process, interpret, and merge the available information increases, resulting in slower response times. The inclusion of the more abstract visualizer view does not significantly increase user response time. We believe this to be due to it requiring less mental overhead but have no direct proof of this. Similarly, we believe that a 5–10-s playback was the best range because with longer data feed lengths the amount of meaningful information the user can process and remember plateaus, making the extra information relatively less useful (when only considering response time) in the decision-making process.

The accuracy measures, however, suggest a slightly different story. When considering this metric, the longer the data feed playback, the more accurate the users' answers become. This is most likely because users can make a more informed decision, even though they take longer to respond in order to process the additional information. The same argument applies to the display condition. Although the extra videos require more time to process, they allow users to make more informed decisions, leading to higher accuracy.

From this experiment, we can make some recommendations about attaining situational awareness in systems such as ours that involve a remote human as part of a multiagent team. There is clearly a tradeoff between accuracy and response time; in general, the most accurate conditions were those that had the longest response time. This effect, however, may decrease as users become more expert. The information they derive from the data feed may plateau earlier. Thus, a follow-up to this experiment should use expert users as subjects and examine whether the accuracy trends still hold or if they begin to correlate with the response time results. Either way, a choice needs to be made between efficiency and accuracy. If timing is critical, then it might be worthwhile to accept a lower accuracy in order to encourage the operator to respond more quickly; similarly, if accuracy is more important, then users should be given as much information as possible, and sufficient time to process and merge the information, in order to allow them to make more informed decisions. It is important to keep in mind, however, that our measures of time did not include actual task execution, instead only measuring the time taken to determine what should be done. Since a misdiagnosis may lead to wasted recovery effort, it is quite possible that classification errors can affect the speed of the system to a greater extent than these results indicate, increasing the importance of accurate classification.

VII. FUTURE WORK

Determining how to proceed when the human relinquishes control of a task remains an open research question. If the

human returns control after completing the assigned task (and only the assigned task), it is straightforward: our current system is able to monitor the task's completion and continue with the tasks that were dependent upon it. However, the human may instead perform additional or different tasks than those that the system expected. In addition to deducing the human's goals and tracking the arbitrary effects her actions may have on the structure, the system needs to be able to replan to accomplish the scenario's goals from whatever state the human leaves the system in when she relinquishes control.

Our current scenario encompasses a relatively complex task, but contains minimal coordinated manipulation of a single object by multiple manipulator robots. Our earlier work involved a much simpler task but required extensive coordinated manipulation [27]. We are now moving to a new scenario that combines a significantly more complex assembly with a need for this type of coordinated manipulation. Addressing both these issues while using Sliding Autonomy in a multiagent setting should uncover many new issues, for example, deciding how best to coordinate tightly coupled manipulation between a human and robot, instead of between two robots. Additionally, we are looking into adding additional robots with overlapping capabilities to the team. With these additional robots, the planning problem will become much more interesting, as a wider variety of solutions to the scenario will be possible. Their overlapping capabilities will also allow the robots to ask each other for help in addition to, or instead of, asking the human. This additional step will allow the human to be treated as simply another agent in the system, just one with a slightly different skill set.

VIII. CONCLUSION

We have presented various issues involved with extending Sliding Autonomy into a multiagent domain. One of the main challenges is allowing the robots to ask for help, since the human is not always guaranteed to be paying attention to each robot at any given time. We make this possible by incorporating user models into our system that allow the robotic agents to make informed and reasonable decisions about when to request the operator's assistance. Experi-

mental data showed that this way of incorporating a human into a multirobot team combines the advantages of autonomous robot operation with the reliability of teleoperation, resulting in a more efficient and robust system as a whole.

By allowing the system to ask the operator for help, we also introduce the question of how to best enable operators to quickly gain situational awareness of the robot's workspace and state so that they can provide assistance more quickly and effectively. To this end, we conducted an experiment testing different operator displays and playback times to find out which combination resulted in the most efficient interaction. Based on the results, we feel that a human operator in our system will be able to quickly gain situational awareness and assist any robotic team member that asks for help, whether or not the operator had previously been paying attention to the situation.

In order to allow interagent coordination even when a human is controlling some task, we allow the robotic team members to monitor their and other team members' progress as related to their current task. This allows the robots to continue with their respective tasks even when their team members are being controlled by an operator or are waiting for an operator response. In other words, a team member can use its knowledge of the other agents' actions to ensure that its tasks remain coordinated. Our experiments have shown that this is an effective way of maintaining task coordination while using Sliding Autonomy.

Overall, we have shown some shortcomings of controlling large-scale construction systems through both full autonomy and complete teleoperation. We have demonstrated that by allowing humans to work as team members via Sliding Autonomy in a multiagent system, a blend of complementary skills and abilities emerges that increases the system's overall robustness and efficiency. ■

Acknowledgment

The authors would like to thank the many individuals who contributed to the DIRA and Trestle projects over the years: R. Ambrose, D. Apfelbaum, J. Brookshire, R. Burrige, B. Hamner, D. Hershberger, M. Hwangbo, D. Kortenkamp, S. Mehalek, Metrica/TRACLabs, J. Ramos, T. Smith, P. Staritz, and P. Tompkins.

REFERENCES

- [1] P. Valckenaers, H. V. Brussel, L. Bongaerts, and F. Bonneville, "Programming, scheduling, and control of flexible assembly systems," *Comput. Industry*, vol. 26, no. 3, pp. 209–218, Aug. 1995.
- [2] A. Stroupe, T. Huntsberger, A. Okon, and H. Aghazarian, "Precision manipulation with cooperative robots," in *Multirobot Systems: From Swarms to Intelligent Automata*, L. Parker, F. Schneider, and A. Schultz, Eds. New York: Springer, 2005.
- [3] R. B. Gillespie, J. E. Colgate, and M. Peshkin, "A general framework for cobot control," in *Proc. Int. Conf. Robotics Automation*, Detroit, MI, 1999.
- [4] W. Wannasuphprasit, P. Akella, M. Peshkin, and J. E. Colgate, "Cobots: A novel material handling technology," in *Proc. Int. Mechanical Eng. Congr. Exposition*, Anaheim, CA, 1998, ASME 98-WA/MH-2.
- [5] R. R. Burrige, J. Graham, K. Shillcutt, R. Hirsh, and D. Kortenkamp, "Experiments with an EVA assistant robot," in *Proc. 7th Int. Symp. Artificial Intelligence, Robotics Automation in Space (iSAIRAS)*, Nara, Japan, 2003.
- [6] T. Fong, C. Thorpe, and C. Baur, "Robot, Asker of questions," *Robotics Autonomous Syst.*, vol. 42, 2003.
- [7] G. Dorais, R. Banasso, D. Kortenkamp, P. Pell, and D. Schreckenghost, "Adjustable Autonomy for human-centered autonomous systems on mars," in *Proc. Mars Soc. Conf.*, 1998.
- [8] D. Kortenkamp, R. Burrige, P. Bonasso, D. Schrenkenghost, and M. Hudson, "An intelligent software architecture for semiautonomous robot control," in *Proc. Autonomy Control Software Workshop, Autonomous Agents*, 1999.

- [9] C. N. Nielsen, M. A. Goodrich, and J. W. Crandall, "Experiments in human-robot teams," in *MultiRobot Teams: From Swarms to Intelligent Automata*, A. C. Schultz and L. E. Parker, Eds. Boston, MA: Kluwer, 2003, vol. 2.
- [10] P. Scerri, D. Pynadath, and M. Tambe, "Towards Adjustable Autonomy for the real world," *J. Artificial Intelligence Res.*, vol. 17, pp. 171–228, 2002.
- [11] R. T. Maheswaran, M. Tambe, P. Varakantham, and K. Myers, "Adjustable Autonomy challenges in personal assistant agents: A position paper," in *Lecture Notes in Computer Science*. London, U.K.: Springer-Verlag, 2004.
- [12] H. A. Yanco and J. Drury, "Where am I? Acquiring situation awareness using a remote robot platform," in *Proc. IEEE Conf. Systems, Man Cybernetics*, Oct. 2004.
- [13] M. R. Endsley, S. J. Selcon, T. D. Hardiman, and D. G. Croft, "A comparative analysis of SAGAT and SART for evaluations of situation awareness," in *Proc. 42nd Annu. Meeting Human Factors Ergonomics Soc.*, Chicago, IL, Oct. 1998.
- [14] M. A. Goodrich, D. R. Olsen, J. W. Crandall, and T. J. Palmer, "Experiments in Adjustable Autonomy," in *Proc. IJCAI Workshop Autonomy, Delegation and Control: Interacting With Intelligent Agents*, 2001.
- [15] D. Hershberger, R. Simmons, S. Singh, J. Ramos, and T. Smith, "Coordination of heterogeneous robots for large-scale assembly," in *Robot Teams: From Diversity to Polymorphism*. London, U.K.: A.K. Peters, 2002.
- [16] R. Simmons, J. Fernandez, R. Goodwin, S. Koenig, and J. O'Sullivan, "Lessons learned from Xavier," *IEEE Robot. Automat. Mag.*, vol. 7, no. 2, pp. 33–39, Jun. 2000.
- [17] R. Bostelman, J. Albus, N. Dagalakis, and A. Jacoff, "RoboCrane project: An advanced concept for large scale manufacturing," in *Proc. AUVSI Conf.*, Orlando, FL, Jul. 1996.
- [18] Metrica/TRACLabs Arm Website. [Online]. Available: <http://www.traclabs.com/robotarm.htm>
- [19] R. Bonasso, R. Firby, E. Gat, D. Kortenkamp, D. Miller, and M. Slack, "Experiences with an architecture for intelligent, reactive agents," *J. Experimental Theoretical Artificial Intelligence*, vol. 9, no. 2–3, pp. 237–256, 1997.
- [20] R. Simmons and D. Apfelbaum, "A task description language for robot control," in *Proc. Conf. Intelligent Robots Systems (IROS)*, Victoria, Canada, 1998.
- [21] D. Goldberg, V. Ciciello, M. B. Dias, R. Simmons, S. Smith, and A. Stentz, "Market-based multirobot planning in a distributed layered architecture," in *Multirobot Systems: From Swarms to Intelligent Automata: Proc. 2003 Int. Workshop Multirobot Systems*, vol. 2. Boston, MA: Kluwer, 2003, pp. 27–38.
- [22] D. Hershberger, R. Burrige, D. Kortenkamp, and R. Simmons, "Distributed visual servoing with a Roving Eye," in *Proc. Conf. Intelligent Robots and Systems (IROS)*, Takamatsu, Japan, Oct. 2000.
- [23] F. W. Heger, L. M. Hiatt, B. Sellner, R. Simmons, and S. Singh, "Results in Sliding Autonomy for multirobot spatial assembly," in *Proc. 8th Int. Symp. Artificial Intelligence, Robotics Automation Space (ISAIRAS)*, Munich, Germany, Sep. 5–8, 2005.
- [24] B. Sellner, R. Simmons, and S. Singh, "User modelling for principled Sliding Autonomy in human-robot teams," in *Multirobot Systems: From Swarms to Intelligent Automata*, L. Parker, F. Schneider, and A. Schultz, Eds. New York: Springer, 2005.
- [25] SpaceMouse Manual. [Online]. Available: <http://www.3dconnexion.com/spacemouseplus.htm>
- [26] S. G. Hart and L. E. Staveland, "Development of the NASA-TLX (task load index): Results of empirical and theoretical research," in *Human Mental Workload*. Amsterdam, The Netherlands: North-Holland, 1988, pp. 139–183.
- [27] J. Brookshire, S. Singh, and R. Simmons, "Preliminary results in Sliding Autonomy for assembly by coordinated teams," in *Proc. Conf. Intelligent Robots Systems (IROS)*, Sendai, Japan, Sep. 28–Oct. 2, 2004.

ABOUT THE AUTHORS

Brennan Sellner received the B.S. degree in computer science and human-computer interaction, in 2001, and the M.Sc. degree in robotics, in 2004, both from Carnegie Mellon University, Pittsburgh, PA. He is pursuing the Ph.D. degree at the Robotics Institute from the same university.

His research interests include planning and execution systems for flexible, coordinated, and cooperative robot-human teams, especially the prediction of conflicts and execution-time task adaptation, Sliding Autonomy, and human-robot interaction.



Laura M. Hiatt received the B.S. degree in symbolic systems from Stanford University, Stanford, CA, in 2004. She is pursuing the Ph.D. degree from the School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.

Her current research interests include execution-time evaluation and modification of multi-agent plans.



Reid Simmons received the Ph.D. degree from the Massachusetts Institute of Technology, Cambridge, in 1988, in the field of artificial intelligence. His thesis work analyzed the relationships between different aspects of expertise and developed a domain-independent theory of debugging faulty plans.

He is a Research Professor in the School of Computer Science, Carnegie Mellon University, Pittsburgh, PA. Since coming to Carnegie Mellon in 1988, his research has focused on developing self-reliant robots that can autonomously operate over extended periods of time in unknown, unstructured environments. This work involves issues of robot control architectures that combine deliberative and reactive control, probabilistic planning and reasoning, monitoring and fault detection, and robust indoor and outdoor navigation. More recently, his research interests have included the areas of coordination of multiple heterogeneous robots, human-robot social interaction, and formal verification of autonomous systems. He has been involved in the development of over a dozen autonomous robots.

Dr. Simmons was part of the Remote Agent team that won the 1999 NASA Software of the Year Award.



Frederik W. Heger (Member, IEEE) received the B.S.E. and M.S.E. degrees in mechanical engineering from the University of Pennsylvania, in 2003 and 2004, respectively. He is pursuing the Ph.D. degree from the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.

His current research interests include introducing a Sliding Autonomy component into multi-agent planning.



Sanjiv Singh received the Ph.D. degree in robotics from Carnegie Mellon University, Pittsburgh, PA, in 1995.

He is currently an Associate Research Professor at the Robotics Institute, Carnegie Mellon. His recent work has two main themes: perception in natural environments and multiagent coordination. A recent project was developing techniques that compute egomotion using passive vision and inertial measurements. His research interests also include providing accurate tracking and situational awareness in dynamic environments such as encountered in search and rescue, using radio signals to compute location. This research seeks to assist emergency response personnel in their operations. A related project seeks to collect and act on environmental data from a sensor web in precision agriculture. The second theme is developing a multirobot architecture to perform complex tasks such as assembly that cannot be performed by a single machine, where tasks are distributed between robots and human operators such that the strengths of each can be exploited and control can “slide” effectively between humans and robots.



Dr. Singh is the Editor-in-Chief of the *Journal of Field Robotics*.

AUTHOR QUERY

AUTHOR PLEASE ANSWER QUERY

AQ1 = Author, please provide keywords.

END OF AUTHOR QUERY