

# Continual Coordination through Shared Activities

Bradley J. Clement and Anthony C. Barrett  
Jet Propulsion Laboratory  
California Institute of Technology  
4800 Oak Grove Drive, M/S 126-347  
Pasadena, CA 91109-8099  
{bclement, barrett}@aig.jpl.nasa.gov

## ABSTRACT

Interacting agents that interleave planning and execution must reach consensus on their commitments to each other. In domains where agents have varying degrees of interaction and different constraints on communication and computation, agents will require different coordination protocols in order to efficiently reach consensus in real time. We briefly describe a largely unexplored class of real-time, distributed planning problems (inspired by interacting spacecraft missions), new challenges they pose, and a general approach to solving the problems. These problems involve self-interested agents that have infrequent communication but collaborate on joint activities. We describe a Shared Activity Coordination (SHAC) framework that provides a decentralized algorithm for negotiating the scheduling of shared activities over the lifetimes of separate missions, a soft, real-time approach to reaching consensus during execution with limited communication, and a foundation for customizing protocols for negotiating planner interactions. We apply SHAC to a realistic simulation of interacting Mars missions and illustrate the simplicity of protocol development.

## 1. INTRODUCTION

When interleaving planning and execution, an agent adjusts its planned activities as it gathers information about the environment and encounters unexpected events, and interacting agents coordinate these adjustments to manage commitments with each other. Demand for this kind of autonomous agent technology is growing for space applications. Autonomous spacecraft promise new capabilities and cost improvements in exploring the solar system. Spacecraft (and rovers) that explore other planets have intermittent, delayed communication with Earth, requiring that they be able to manage their resources and operate in for long periods in isolation. The common approach to autonomous decision making is to

**NOTE TO REVIEWERS:** The authors intend to demonstrate simulated spacecraft coordination for the Mars scenario at the conference.

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 2002 ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

place integrated image analysis, planning, and execution systems on-board the spacecraft.

In addition, there is a growing trend toward multi-spacecraft missions. Over forty multi-spacecraft missions have been proposed, including formation flying teams and over 16 planned missions to Mars in the next decade. These spacecraft will coordinate measurements, share images, and route data to and from Earth. Separate missions, such as those to Mars have their own budgets, experiments, and operations teams. As such, the spacecraft represent self-interested entities that benefit from collaborative interactions.

But, even a single spacecraft has multiple science instruments representing different goals of different scientists, and different operations groups will have different areas of expertise over different subsystems for control. These different groups negotiate over mission plans in the same way that different Mars missions must collaborate over spacecraft interactions. Whether this negotiation is done on-board or on Earth, there is a distributed operations planning problem that benefits from automation. Both also have real-time aspects. Onboard systems must plan safely over near- and long-term horizons, and ground systems must also replan based on changing contexts in daily, weekly, and lifelong mission exercises. Ground planning also suffers from communication constraints. Scientists from different universities or opposite sides of the globe will intermittently provide inputs and respond on an irregular basis. A collaboration/negotiation system must be built around communication constraints to meet hard deadlines for coming to consensus on consistent operations plans.

In this work, we will briefly characterize this general problem in terms of activity interaction types and communication constraints and discuss its challenges. The field of multiagent planning has largely focused on fully cooperative planning and execution [4, 5, 13, 8, 2]. Market-based agent systems address near-term resource negotiation but have rarely addressed how near-term decisions affect longer-term goals. Multiagent systems built for Robocup Soccer competitions mainly address collaborative multiagent execution in an adversarial environment and have limited planning capabilities. These approaches do not adequately address real-time planning for self-interested agents.

This paper presents a framework for Shared Activity Coordination (SHAC). SHAC consists of an algorithm for continually coordinating agents and a foundation for rapidly designing and implementing coordination protocols based on a model of shared activities. In the same fashion that a real-time planning system must commit to actions to pass to an execution system, a real-time coordination system must additionally establish consensus on shared activities before they are executed based on communication constraints. Our ultimate goal is to create interacting agents that autonomously adjust their coordination protocols with respect

to unexpected events and changes in communication or computation constraints so that the agents can most efficiently achieve their goals.

First we characterize a class real-time, self-interested multiagent planning problem that exists for space. Then we describe the shared activity model, the SHAC algorithm, and its interface to the planner. Then we specify some generic roles and protocols using the SHAC framework that build on prior coordination mechanisms. Then we describe how our current implementation of SHAC is used to coordinate the communication of two rovers and an orbiter in a simulated Mars scenario. We follow with future research needs revealed in this scenario and comparisons to related work.

## 2. CONTINUAL COORDINATION PROBLEM

As mentioned before, agents that interleave planning and execution must commit near-term activities to the execution system while receiving feedback in the form of state updates and activity performance. One such continual planning system, CASPER (Continuous Activity Scheduling Planning Execution and Replanning) identifies the period when the planner commits activities to the execution as a *commit window* [1]. Distributed planning agents must additionally reach consensus on team interactions before execution. As explored in the team plan model given by TEAMCORE [13, 12] formalizations of Shared Plans [8], and coordination interactions of TAEMS [4], these interactions could include

- use and replenishment of shared resources,
- joint actions for achieving a mutually beneficial subgoal,
- choice of methods for achieving a team subgoal,
- participation and role assignments in a joint plan, and
- proposals and commitments of the above.

## 3. SHAC

However, reaching consensus on these interactions is complicated when the agents can only communicate intermittently. Depending on the number of agents involved in a particular interaction, consensus may need to be established far in advance so that negotiations can be propagated far in advance of execution. Thus, for any particular set of interactions, a *consensus window*, within which the agents must limit negotiation and establish agreement, should be defined. For example, if three agents must negotiate a joint action in advance of execution but can only communicate pairwise in disjoint time windows, a consensus window must extend at least to cover windows connecting all three agents. Inside the consensus window, a simple protocol eliminating negotiation (such as all agree or reject) must be employed to guarantee consensus. Interactions beyond the consensus window can be negotiated.

It is an open problem how consensus windows should be defined based on interaction types and communication opportunities. While this paper does not solve this problem, it describes an implemented infrastructure and testbed for developing and evaluating alternative consensus windows and protocols.

Our approach, called Shared Activity Coordination (SHAC), provides a general algorithm for interleaving planning and the exchange of plan information based on shared activities. Agents coordinate their plans by establishing consensus on the parameters of shared activities. Figure 1 illustrates this approach where three agents share one activity and two share another. The constraints denote equality requirements between shared activity parameters in different agents. The left vertical box over each planner's schedule

represents a *commit window* that moves along with the current time. A consensus window is shown to the right of the commit window, within which consensus must be quickly established before committing. Since consensus is hard to maintain when all agents can modify a shared activity's parameters at the same time, agents must participate in different coordination roles that specify which agent has control of the activity. As shown in the figure, SHAC interacts with the planning and execution by propagating changes to the activities, including their parameters and constraints on the values of those parameters.

SHAC continually coordinates by interfacing to a combined planning/execution system that responds to failures and state updates from the execution system. Our implementation interfaces with one such continual planning system, CASPER, mentioned in the previous section. Instead of batch-planning in episodes, CASPER continually adapts near and long-term activities while re-projecting state and resource profiles based on updates from sensors.

### 3.1 Shared Activities

The model of a shared activity is meant to capture the information that agents must share, including control mechanisms for changing that information. A shared activity is a tuple (*parameters, agent roles, protocols, decomposition, constraints*). The parameters are the shared variables and current values over which agents must reach consensus by the time the activity executes. The agent roles determine the local activity of each agent corresponding to the joint action. To provide flexible coordination relationships, the role activities of the shared activity can have different conditions and effects as specified by the local planning model. The shared parameters map to local parameters in the role activity.

For example, a shared data communication activity can map to a receive role activity for one agent and a send role activity for another. Shared parameters could specify the start time, duration, transfer rate, and data size of the activity. The data size is depleted from the sender's memory resource but added to the receiver's memory. The agents could have separate power usages for transmitting and receiving. In this case the resources are not shared. Another shared activity could be the use of a common transport resource. Although one agent in an active transit role actually changes position, other agents in passive roles have local activities that only reserve the transport resource. Figure 3.1 shows an instance of this shared activity where an orbiter receives communication from a rover.

Protocols are the mechanisms assigned to each agent (or role) that allow the agents to change constraints on the shared activity, the set of agents assigned to the activity, and their roles. In Figure 3.1, both the orbiter and rover use an argumentation protocol to negotiate the scheduling and attributes of the communication. Constraints will be described in the next section, and a variety of protocols will be defined in the Protocols section.

The shared decomposition enables agents to select different team methods for accomplishing a higher level shared goal. Specifically, the decomposition is a set of shared subactivities. The agents can choose the decomposition from a pre-specified set of subactivity lists. For example, a joint observation among orbiters could decompose into either (*measure, process.image, downlink*) or (*measure, downlink*).

### 3.2 Constraints

Constraints are created by agents' protocols to restrict sets of values for parameters (*parameter constraints*) and permissions for manipulating the parameters, changing constraints on the parameters, and scheduling shared activities (*permission constraints*). These

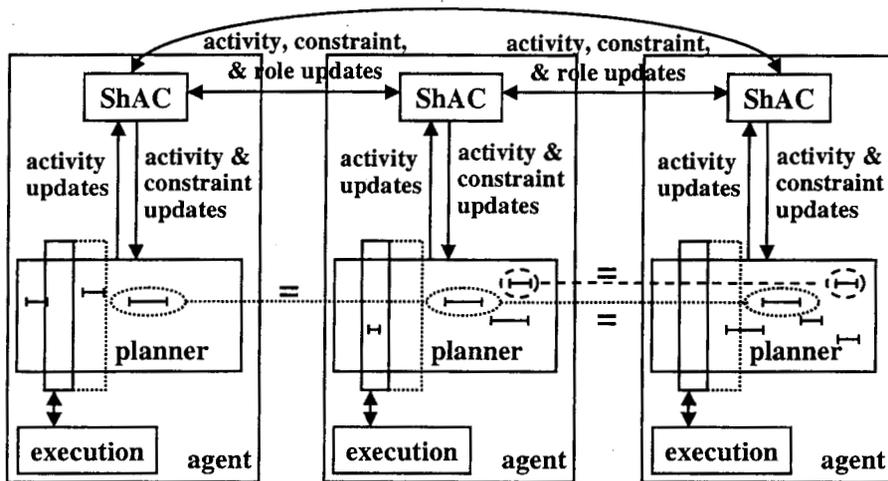


Figure 1: Activities shared among continual planners

```

shared_activity communicate comm_id_12 {
  time start_time = 2004-302:09:30:00; // date
  int duration = 200; // seconds
  int data_size = 25600; // 25.6 Mbits
  real xmit_rate = 128.0; // 5.0 Kbps
  int priority = 1; // critical
  roles =
    receive by orbiter,
    send by rover;
  protocols =
    receive argumentation,
    send argumentation;
  permissions =
    receive (move, delete, xmit_rate),
    send (delete, data_size, priority);
  parameter_constraints =
    rover start_time = ([2004-302:09:30:00, 2004-302:09:38:00],
                       [2004-302:18:30:00, 2004-302:18:38:00]);
}

```

Figure 2: An instance of a shared communication activity between a rover and orbiter

constraints restrict the privileges (or responsibilities) of agents in making coordinated planning decisions. By communicating constraints, protocols can come to agreement on the scheduling of an activity without sharing all details of their local plans.

A parameter constraint is a tuple  $\langle agent, parameter, value\ set \rangle$ . The *agent* denotes who created the constraint. Some protocols differentiate their treatment of constraints based on the agent that created them. For example, the asynchronous weak commitment algorithm prioritizes agents so that lower-priority agents only conform to higher-priority agent constraints [15]. Agents can add to their constraints on a parameter, replace constraints, or cancel them. A string parameter constraint, for example, can restrict a parameter to a specific set of strings. An integer or floating point variable constraint is a set of disjoint ranges of numbers. Scheduling constraints can be represented as constraints on a start time integer parameter. This is shown in Figure 3.1 where the rover restricts the start time of the communication between two eight minute intervals.

Permission constraints determine how an agent's planner is allowed to manipulate shared activities. The following permissions are currently defined for SHAC:

- parameters - change parameter values
- move - set start time
- duration - change duration of task
- delete - remove from plan
- choose decomposition - select shared subactivity of an *or* activity
- add - add to plan<sup>1</sup>
- constrain - send constraints to other agents

In the communication example in Figure 3.1, the receiver is allowed to reschedule (move) the activity, delete it, or change the transmission rate. The sender cannot move the activity, but can delete it and change the requested size and priority of the data.

<sup>1</sup>This permission applies to a class of shared activities (i.e. an agent may be permitted to instantiate a shared activity of a particular class).

### 3.3 Coordination Algorithm

The purpose of the SHAC algorithm is to negotiate the scheduling and parameters of shared activities until consensus is reached. Figure 3 gives a general specification of the algorithm. SHAC is implemented separate from the planner, so steps 1 through 3 are handled by the planner through an interface to SHAC. Step 4 invokes the protocols that potentially make changes to refocus coordination on resolving shared activity conflicts and improving plan utility. SHAC sends modifications of shared activities and constraints to sharing agents in step 5. In step 6, shared activities and constraints are updated based on changes received from other agents.

Ignoring coordination, a continuous planner must determine when it is appropriate to release activities to the execution system. In some cases, an activity involved in a conflict may either be released (requiring the planner to recover from potential failures) or postponed (to allow the planner to recover before a failure occurs). CASPER keeps a *commit window* (an interval between the current time and some point in the near future) within which activities cannot be modified and passes these activities to the execution system.

This interaction with the execution system becomes more complicated when agents share tasks. SHAC must make sure that when a shared activity is released, all agents release it while in consensus on the start time and other parameters of the task. Ideally the agents should establish consensus before the commit window. SHAC avoids changes in the commit window by keeping a *consensus window* that extends from the commit window forward by some period specific for the activity. As time moves forward, the windows extend forward. When a shared activity moves into the consensus window, the agents switch to the simple consensus protocol to try and reach consensus before the activity moves into the commit window.

## 4. PROTOCOLS

In general, protocols determine when to communicate, what to communicate, and how to process received communication. During each iteration of the loop of the coordination algorithm (Figure 3), the protocol determines what to communicate and how to process communication. A protocol is defined by how it implements the following procedures to be called during step 4 of the SHAC coordination algorithm for the shared activity to which it is assigned:

1. modify permissions of the sharing agents
2. modify locally generated parameter constraints
3. add/delete agents sharing the activity
4. change roles of sharing agents

The default protocol, representing a base class from which other protocols inherit, does nothing for these methods. However, even with this passive protocol, the SHAC algorithm still provides several capabilities:

**joint intention** A shared activity by itself represents a joint intention among the agents that share it.

**mutual belief** Parameters or state assertions of shared activities can be updated by sharing agents to establish consensus over shared information.

**resource sharing** Sharing agents can have identical constraints on shared states or resources.

**active/passive roles** Some sharing agents can have active roles with execution primitives while others have passive roles without execution primitives.

**master/slave roles** A master agent can have permission to schedule/modify an activity that a slave (which has no permissions) must plan around.

The following sections describe some subclasses of this abstract protocol, demonstrating capabilities that each protocol method can provide.

### 4.1 Argumentation

Argumentation is a technique for negotiating joint beliefs or intentions [9]. Commonly, one agent makes a proposal to others with justifications. The others evaluate the argument and either accept it or counter-propose with added justifications. This technique has been applied to teamwork negotiation research to form teams, reorganize teams, and resolve conflicts over members' beliefs [14]. It can also be used to establish consensus on shared activities.

A shared activity and associated parameter values are the proposal or counterproposal. Justifications are given as parameter constraints. A proposal is a change to a shared activity that does not violate any parameter constraints. A counterproposal may violate constraints. Protocol method 2 must be implemented to provide the parameter constraint justifications for proposals and counterproposals. In order to avoid race conditions, protocol method 1 regulates permissions.

Argumentation method 1

- if this agent sent the most recent proposal/counterproposal
  - if planner modified shared activity
    - \* remove self's modification permissions
- else
  - give self modification permissions (e.g. move and delete)

Argumentation method 2

- if planner modified shared activity
  - generate parameter constraints describing locally consistent values

As an example, one agent can propose an activity with a particular start time and add justifications in the form of all intervals within which the shared activity can be locally scheduled. Other agents can replan to accommodate the proposal and counter-propose with their own interval restrictions if replanning cannot accommodate others' constraints. If the agents cannot establish consensus before the consensus window, a higher ranking agent can mandate a time that benefits most of the agents. Of course, there are many variations on this example. Agents may be restricted because they are slaves or do not have constraint permissions to counter-propose.

### 4.2 Delegation

Delegation is a mechanism where an agent in a passive delegator role assigns and reassigns activities to different subsets of agents in active subordinate roles. The delegator and subordinate protocols only need

Delegator method 3

- if *agent roles* empty

Given: a *plan* with multiple activities including a set of *shared\_activities* with *constraints* and a *projection of plan* into the future.

1. Revise *projection* using the currently perceived state and any newly added goal activities.
2. Alter *plan* and *projection* while honoring *constraints*.
3. Release relevant near-term activities of *plan* to the real-time execution system.
4. For each shared activity in *shared\_activities*,
  - if outside consensus window,
    - apply each associated protocol to modify the shared activity;
  - else
    - apply simple consensus protocol.
5. Communicate changes in *shared\_activities*.
6. Update *shared\_activities* based on received communications.
7. Go to 1.

Figure 3: Shared activity coordination algorithm

- choose an *agent* to whom to delegate the activity
- add (*agent*, subordinate) to *agent roles*

Subordinate method 3

- if cannot resolve conflicts/threats involving activity
  - remove self from *agent roles*

## 5. APPLICATION TO MARS SCENARIO

Now we describe how SHAC is applied to a simulated three-day scenario involving two Mars Exploration Rovers (MERs), the Mars Odyssey orbiter, Mars Global Surveyor, and the Mars Express orbiter. The delegation protocol described previously was subclassed for the rovers to assign and reassign the routing of images to the orbiters based on how quickly they can deliver the data to Earth. Different master/slave and active/passive roles are defined using permission constraints for the shared activities to implement a basic protocol for coordinating communication to and from Earth. Interactions over communication (once delegated) are between two agents, so the consensus window is defined to cover communication activities spanning two communication opportunities into the future. Once in the consensus window, the rover cannot redelegate activities unless the orbiter cannot resolve conflicts and must decommit. We intend to experiment with other protocols and consensus window definitions in this domain in our future work.

The MERs (MER A and MER B) and the orbiters can communicate with Earth directly, but the MERs can optionally route data through the orbiters, which communicate with Earth at a higher bandwidth. The rovers need daily communication with ground operations to receive new goals. The rovers will often fail to traverse to a new target location and cannot proceed until new instructions come from ground operations. In this scenario both MERs must negotiate with the assigned orbiter to determine how to most quickly get a response from ground after sending an image of the surrounding area.

Each MER has a communication state shared with each orbiter that tracks when the image is generated, when it gets to Earth, and when the response from ground operations arrives to the rover.

Shared activities for changing the state are shown for different routing options in Figure 4. The rover's activity for generating an image from its panoramic camera changes the state to *request* to communicate its need to downlink and receive an uplink. Activities for sending the image to Earth (either directly or through Odyssey), change the state to a *wait for uplink* state to indicate that the rover will then be waiting for the uplink. Ground operations needs a period of time to generate new commands for the uplink, so if the uplink is received by Odyssey, the state changes to *received* to indicate that now the rover can get the uplink from Odyssey. Once the rover receives the uplink, the state changes back to the normal *no pending request* state. Rover tasks (such as a traverse) need the uplinked data before executing, so it places a local constraint that shared state be *no pending request* during its scheduled interval. There are no shared resources although communication requests from a MER have effects on many local resources of both the MER and the orbiter. All of the shared activities have active master and passive slave roles. The MERs and orbiters both take the master role for activities labeled for them in Figure 4.

CASPER planners for each of the MERs and orbiters first build their three-day plans separately to optimize the timely delivery of priority weighted science data, resolving any local constraints on memory, power, battery energy, *etc.* The three-day schedules constitute over 900 tasks for each MER and over 1300 for each of the orbiters with 30 state/resource variables for each MER and 22 for the orbiters. Planning is slowed by a factor of 440 to account for differences between a desktop workstation and a radiation-hardened flight processor. Communication for coordination is restricted to times when the orbiters pass overhead. With the exception of Mars Express, the orbiters pass overhead once every eight hours. Because of its irregular orbit, Mars Express sees the rovers only once every 96 hours. Because of this, we actually used no consensus window for communication with Mars Express, putting pressure on the planners to resolve conflicts during image transmission.

When coordination begins, the planners send their communication requests to the other planners while optimizing their plans. Before these updates are received, the initial views of the shared uplink status are shown in Figure 5. The MERs begin with conflicts with their traverse tasks because the uplink has not yet been

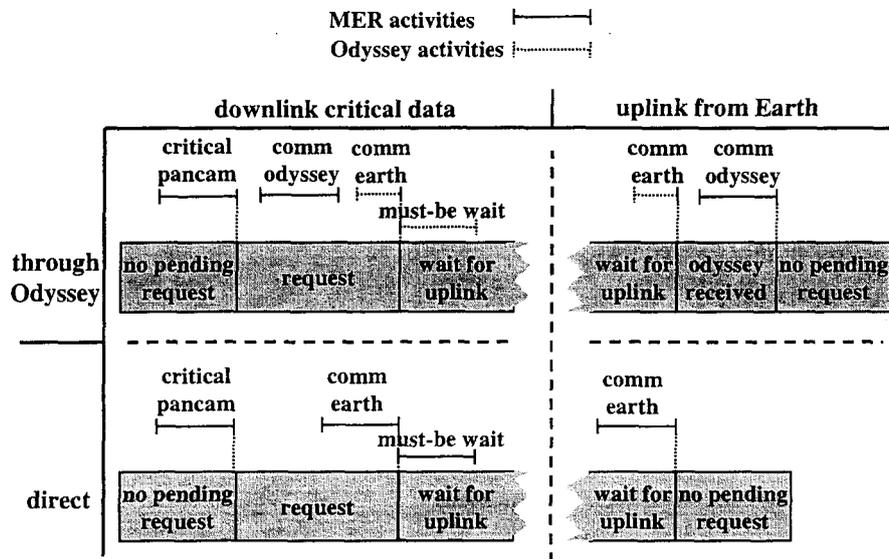


Figure 4: Downlink/uplink states for a rover

received from Earth. The coordination algorithm commands the planners to repetitively process shared task updates, replan to resolve conflicts by recomputing the shared state and modifying scientific measurement operations to adjust for the increased power and memory needs, and send task updates. After a minute and a half, MER A, B, and Odyssey agree on routing the downlink and uplink through Odyssey to get the uplinked commands in time for the traversal on different days. The resulting shared state is shown at the bottom of Figure 5. The planners reach consensus that coordination is complete and sleep while waiting for task updates.

Among other failed communication attempts, we triggered an anomaly in MER A's plan causing it to cancel its first day's tasks and shift the entire schedule forward a day. Before sending the updated shared tasks, replanning was issued to resolve local constraints to avoid propagating inconsistent state information to Odyssey. All conflicts were resolved in a few seconds except the traverse conflicts with a wait state. Then MER A sends a task update to restart coordination. Coordination completes in less than a minute with data again being routed through Odyssey.

While we have only experimented with simple protocols, this application of SHAC to the Mars scenario shows how planners can coordinate during execution while making minimal concessions to ideal plans and responding to unexpected events. In the next section, we discuss how SHAC builds on related work and discuss new research challenges for decentralized, coordinated planning.

## 6. DISCUSSION AND RELATED WORK

Conflicts among a group of agents can be avoided by reducing or eliminating interactions by localizing plan effects to particular agents [10], and by merging the individual plans of agents by introducing synchronization actions [7]. In fact, planning and merging can be interleaved [6]. Earlier work studied interleaved planning and merging and decomposition in a distributed version of the NOAH planner [3] that focused on distributed problem solving. More recent research builds on these techniques by formalizing and reasoning about the plans of multiple agents at multiple levels of abstraction to localize interactions and prune unfruitful spaces during the search for coordinated global plans [2].

DSIPE [5] employs a centralized plan merging strategy for distributed planners for collaborative problem solving using human decision support. Like our approach, local and global views of planning problem help the planners coordinate the elaboration and repair of their plans. DSIPE provides insight into human involvement in the planning process as well as automatic information filtering for isolating necessary information to share. While our approach relies on the domain modeler to specify up front what information will be shared, SHAC supports a fully decentralized framework and focuses on interleaved coordination and execution.

In many ways this work is following the Generalized Partial Global Planning approach to using a mix of coordination protocols tailored for the domain [4]. SHAC offers an alternative framework for separating implementation of these mechanisms from the planning algorithms employed by specific agents. Unlike GPGP, SHAC provides a modular framework for combining lower-level mechanisms to create higher-level roles and protocols. Our future work will build on GPGP's evaluations of mechanism variations to better understand how agents should coordinate for domains varying in agent interaction, communication constraints, and computation limitations.

Finally, TEAMCORE provides a robust framework for developing and executing team plans [13, 12]. This work also offers a decision-theoretic approach to reducing communication within a collaborative framework. Research is needed to investigate the integration of coordinated planning with robust coordinated execution.

An assumption commonly made in multiagent research is that agents will be able to communicate at all times reliably. In the Mars scenario, the spacecraft communicate with each other in varying time windows and frequencies, and the two MERs can never directly talk to each other. Establishing consensus on beliefs and intentions is impossible without certain communication guarantees [11]. Understanding the communication patterns that make consensus possible and the overhead for establishing consensus is critical for multiagent research.

## 7. CONCLUSION

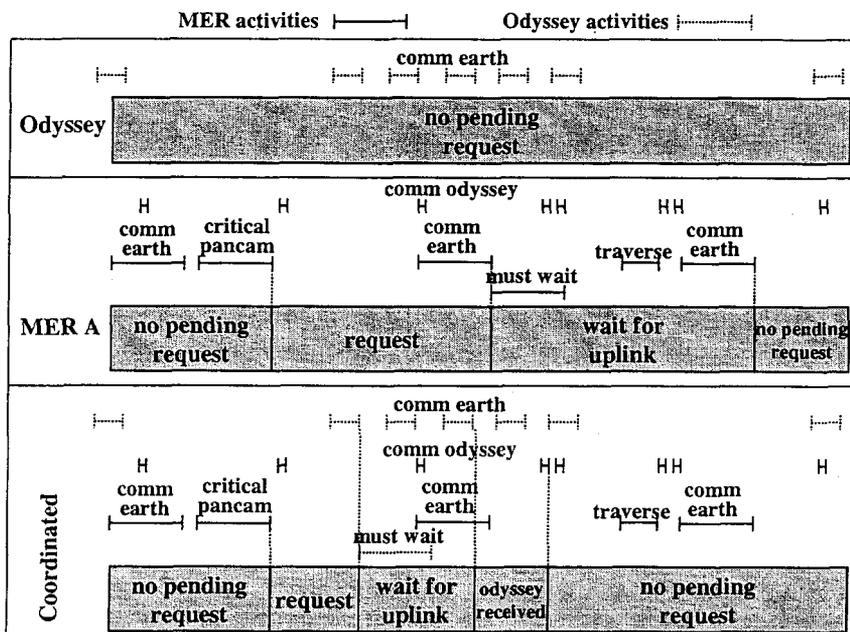


Figure 5: Downlink/uplink shared state for MER A. From top to bottom, Odyssey's initial view, MER A's initial view, and the common view after coordination.

We have introduced shared activity coordination as an approach to designing role-based coordination mechanisms for planning agents. SHAC provides several coordination capabilities upon which we have specified a few higher-level coordination protocols that exercise different aspects of the SHAC model. We have also described an algorithm for continually coordinating planning agents during execution using these protocols. While our future work is aimed at evaluating the benefits of different protocols for different classes of multiagent domains, we validate our approach in coordinating three simulated spacecraft in the presence of an unexpected event.

## 8. REFERENCES

- [1] S. Chien, R. Knight, A. Stechert, R. Sherwood, and G. Rabideau. Using iterative repair to improve the responsiveness of planning and scheduling. In *Proc. ECP*, pages 300–307, 2000.
- [2] B. Clement and E. Durfee. Performance of coordinating concurrent hierarchical planning agents using summary information. In *Proc. ATAL*, pages 213–227, 2000.
- [3] D. Corkill. Hierarchical planning in a distributed environment. In *Proc. IJCAI*, pages 168–175, 1979.
- [4] K. Decker. *Environment centered analysis and design of coordination mechanisms*. PhD thesis, University of Massachusetts, 1995.
- [5] M. desJardins and M. Wolverton. Coordinating a distributed planning system. *AI Magazine*, 20(4):45–53, 1999.
- [6] E. Ephrati and J. Rosenschein. Divide and conquer in multi-agent planning. In *Proc. AAAI*, pages 375–380, July 1994.
- [7] M. P. Georgeff. Communication and interaction in multiagent planning. In *Proc. AAAI*, pages 125–129, 1983.
- [8] B. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86:269–358, 1996.
- [9] S. Kraus, K. Sycara, and A. Evanchik. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence*, 104:1–70, 1998.
- [10] A. Lansky. Localized search for controlling automated reasoning. In *Proc. DARPA Workshop on Innov. Approaches to Planning, Scheduling and Control*, pages 115–125, November 1990.
- [11] S. Mullender. *Distributed Systems*. Addison-Wesley New York, 1995.
- [12] D. Pynadath, M. Tambe, N. Cauvat, and L. Cavedon. Toward team-oriented programming. In *Proc. ATAL*, 1999.
- [13] M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.
- [14] M. Tambe and H. Jung. The benefits of arguing in a team. *AI Magazine*, 20(4), 1999.
- [15] M. Yokoo and K. Hiramama. The distributed constraint satisfaction problem: Formalization and algorithms. *IEEE Trans. on KDE*, 10(5):673–685, 1998.