AI Magazine Volume 23 Number 2 (2002) (© AAAI)

The *RADARSAT*-MAMM Automated Mission Planner

Benjamin D. Smith, Barbara E. Engelhardt, and Darren H. Mutz

■ The Modified Antarctic Mapping Mission (MAMM) was conducted from September to November 2000 onboard *RADARSAT*. The mission plan consisted of more than 2400 synthetic aperture radar data acquisitions of Antarctica that achieved the scientific objectives and obeyed *RADARSAT*'s resource and operational constraints. Mission planning is a time- and knowledge-intensive effort. It required over a workyear to manually develop a comparable plan for AMM-1, the precursor mission to MAMM. This article describes the design and use of the automated mission planning system for MAMM, which dramatically reduced mission-planning costs to just a few workweeks and enabled rapid generation of what-if scenarios for evaluating alternative mission designs.

Planning spacecraft missions is a time- and knowledge-intensive process that can benefit greatly from automated planning and scheduling systems. This article describes the automated mission planning system that was recently developed at the Jet Propulsion Laboratory for the Modified Antarctic Mapping Mission (MAMM). We also compare the automated planning process that was used for MAMM to the manual planning process that was used for its precursor mission, the first *RADARSAT* Antarctic Mapping Mission (AMM-1).

Both AMM-1 and MAMM were joint missions between the National Aeronautics and Space Administration (NASA) and the Canadian Space Agency (CSA) that used *RADARSAT*, a CSA satellite, to acquire synthetic aperture radar (SAR) data of the Antarctic. AMM-1 was conducted in 1997 and acquired the first complete SAR map of Antarctica (Jezek, Sohn, and Loltimeir 1998), as shown in figure 1. MAMM was conducted from September to November of 2000. Its objective was to measure ice surface velocity of the outer regions of the continent, north of latitude –80 degrees. These data provide geoscientists critical information about the geologic forces acting on glaciers, the rate at which ice is pouring into the ocean, and the relation between ice-sheet processes and global climate. Scientists obtain precise velocity data for a given region by comparing SAR images of the region taken from the same viewpoint but several days apart in a process known as radar interferometry. Figure 2 shows a velocity map of the Lambert glacier that was obtained from MAMM data using this process. Scientists also compute surface motion by comparing MAMM data with AMM-1 data from 1997. Figure 3 shows how the Amery ice shelf has been advancing since 1970: the outer line represents the coastline as seen by AMM-1, and the inner line represents the coastline obtained from earlier data.

Planning these missions is a time-intensive and knowledge-intensive process. The mission plans for AMM-1 and MAMM each consist of more than 800 SAR data acquisitions plus supporting spacecraft activities. These plans not only had to meet stringent scientific criteria but also had to obey the satellite's operational and resource constraints. Manually developing these plans takes considerable time and requires intimate knowledge of both science and spacecraft operations. This effort is put in not once but several times: The final plan is produced over several iterations in which a detailed plan is critiqued by other members of the project and then revised, sometimes significantly. The mission plan for AMM-1 consisted of 850 acquisitions in 18 days and took over a workyear to develop manually. Despite repeat-

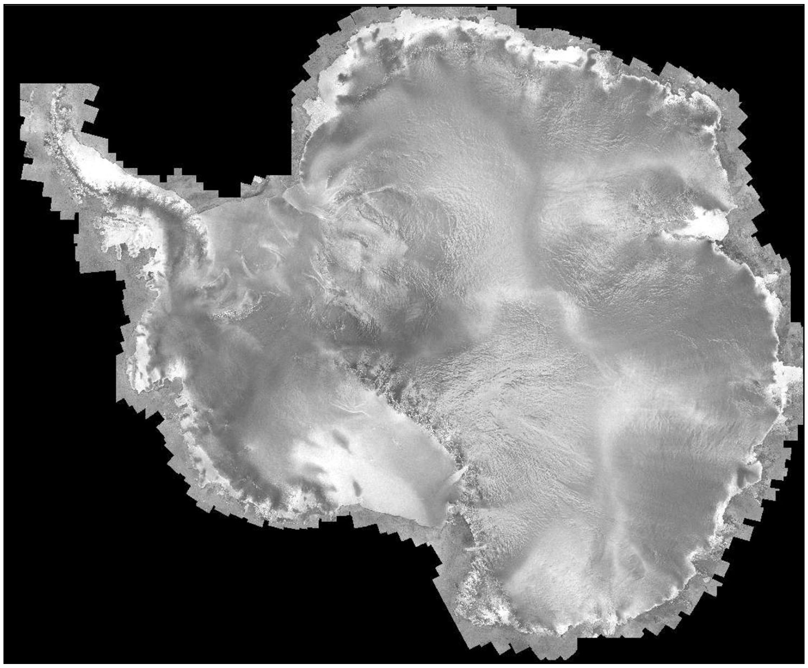

Figure 1. Synthetic Aperture Radar (SAR) Mosaic of Antarctica from the Antarctic Mapping Mission.

Photo courtesy Canadian Space Agency, NASA, Ohio State University, Jet Propulsion Laboratory, Alaska SAR facility.

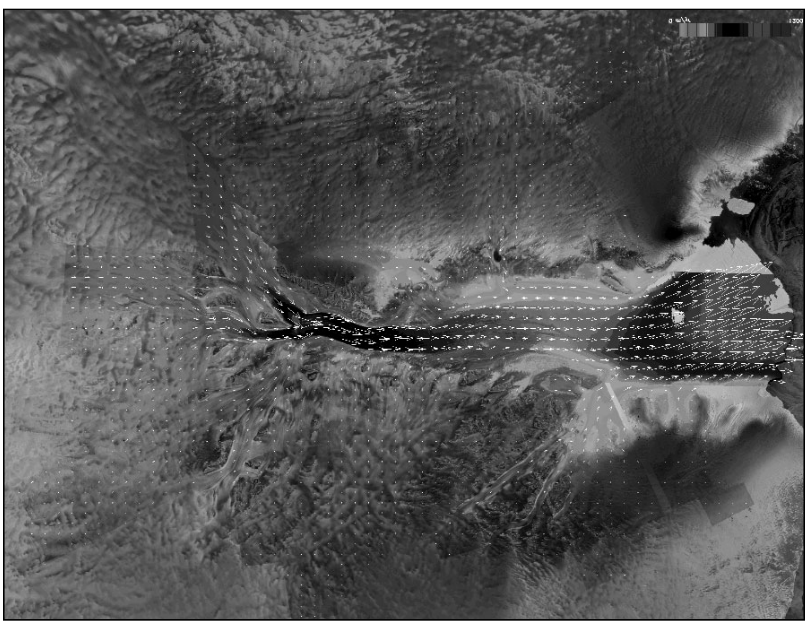

Figure 2. Ice Velocity Map of Lambert Glacier Obtained from the 2000 Antarctic Mapping Mission.

Photo courtesy Canadian Space Agency, NASA, Ohio State University, Jet Propulsion Laboratory, Alaska SAR facility.

ed checking, some operations constraint violations were not detected until just before the mission started, by which time fixing them was expensive and disruptive.

This experience led to the development and use of an automated mission planning system for the MAMM follow-on mission. The human mission planner selects a set of SAR acquisitions that the automated planner expands into a more detailed plan, which it checks for operations constraint violations. Thus, the human mission planner is able to focus on what he/she does best: using scientific knowledge to select SAR acquisitions. With this system, MAMM developed its 24-day mission plan, containing 818 acquisitions, in about 10 workweeks, and the bulk of the time was spent in selecting the SAR acquisitions.

In addition to reducing the plan development effort, the MAMM planner supported trade studies that were not available for AMM-1 by generating what-if plans for various alternatives. These plans could be generated quickly, and the detailed resource use and other information allowed mission managers to make objective comparisons. This capability was instrumental in selecting ground receiving stations, estimating costs, and negotiating spacecraft resource allocations.

The rest of this article describes the design and operation of the automated planning system that was constructed for MAMM based on the ASPEN planning environment (Chien et al. 2000).

Mission Planning Problem

The *RADARSAT* satellite has a handful of parameterized commands that it can perform. For purposes of MAMM and AMM-1, the primary commands are "acquire SAR data and store it on the spacecraft data recorder," "acquire SAR data and downlink it in real time as you get it," and "downlink all the recorded data." A mission plan is a time-ordered list of these activities. The mission planning process is an iterative one. The mission planner develops a detailed draft plan, which the advisory board reviews against scientific, cost, and risk criteria. The planner develops another plan that addresses these concerns, which sometimes requires drastic changes from the previous version. MAMM generated four revisions before arriving at the fifth and final mission plan. The process for generating an individual plan, as shown in figure 4, consists of repeating the following three steps until the plan meets scientific and operations criteria: (1) select SAR swaths, (2) create a downlink schedule, and (2) check the combined schedule for constraint violations. The resulting plan is a time-ordered list of data-acquisition requests and downlink session requests.

The first step is to select SAR swaths that cover the desired target regions in Antarctica and satisfy other scientific requirements. The SAR instrument acquires data in a line that sweeps out a rectangular swath as the spacecraft orbits the earth. It has 16 beam modes, each of which has a different incidence (viewing) angle, footprint, and resolution. Plotting the footprint swept by each beam mode results in thousands of overlapping rectangles that cover Antarctica. Because the instrument cannot be left on continuously, the mission planner must decide what segments of these rectangles should be included in the plan. Any given region is covered by several swaths, although each of these swaths will be on different passes and might use different beam modes. The selected acquisitions must cover the desired regions of Antarctica, fit within the 24-day–plan duration, and use beam modes consistent with the scientific objectives.

The swath selection is partially automated by a tool developed by the CSA called SPA that identifies the available swaths for each beam by plotting the footprint along the spacecraft orbit. The user selects the desired swaths, and SPA generates a swath request file. SPA does not check operations constraints or ensure that the acquired data can be downlinked, so there is no guarantee that the selected swaths make up a valid mission plan.

The second step is to create a downlink schedule. Each SAR data acquisition must either be downlinked to a ground receiving station as it is being acquired or saved to the onboard data recorder for later downlink. The duration and start times of the downlink windows vary from orbit to orbit, and all the recorded data must be downlinked within a single window (the recorder is effectively erased after each downlink). The downlink schedule specifies for each acquisition whether it will be downlinked in real time or stored to the data recorder and specifies the downlink window in which the recorded data will be transmitted. The schedule should also try to maximize objective criteria—certain stations are more reliable or have lower costs than others, and on-board resource costs make real-time acquisitions preferable to recorded ones. Because most of the desired data are of regions where the spacecraft will not be in range of a ground receiving station, most of the acquisitions have to be recorded. It is possible that some of the selected SAR acquisition cannot be downlinked, either in real time or during a later playback. In this case, the mission planner must image these regions with acquisitions from different passes or beams. The automated planning tool generates a downlink schedule for the selected SAR swaths and indicates those that cannot be downlinked. It is then the responsibility of the human mission planner to modify the swath selection to address these problems. The swath selection and downlink schedule are tightly coupled, and one could imagine a planning tool that solves the combined problem. However, MAMM wanted its human mission planners to retain full control

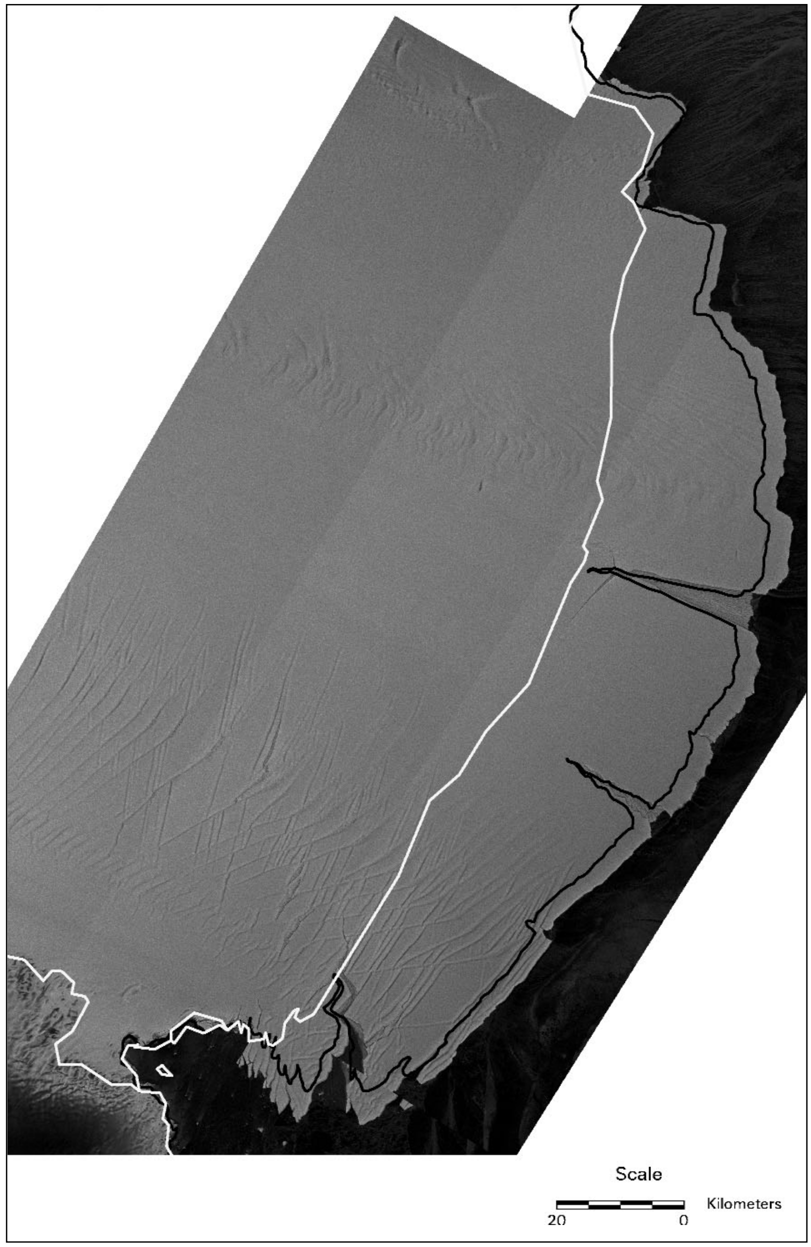

Figure 3. This View of the Amery Ice Shelf Is a Mosaic of Radar Images from the 2000 Antarctic Mapping Mission.

The outer line represents the coastline seen during the 1997 Antarctic Mapping Mission. The inner coastline dates from the mid-1970s. Photo courtesy Canadian Space Agency, NASA, Ohio State University, Jet Propulsion Laboratory, Alaska SAR facility.

Repeat until no constraint violations exist:
1. Select SAR swaths.
2. Create a downlink schedule.
3. Compute resource usage and check for constraint violations.

Figure 4. Observation Planning Process.

over swath selection, which requires human scientific judgment.

The third step is to compute resource use and determine whether the combined acquisition and downlink plan violates operations constraints or resource allocations. If the schedule violates constraints, return to step 1 and modify the selected swaths to correct the problems. Modifications include changing the swath start time, swath duration, or beam or selecting an alternate swath that covers the same target area. The automated planning system expands the acquisition and downlink schedule to include supporting activities, computes the use of each resource over time, and reports any constraint or resource violations.

Replanning during Operations

After the mission plan is developed, it is sent to the Mission Management Office (MMO), which conducts a final verification and uplinks it to the spacecraft for execution. The execution, however, does not always go exactly as planned: SAR acquisitions are sometimes "lost" during operations because of anomalies on the spacecraft or at the ground receiving station. Lost acquisitions are rescheduled for a later opportunity using the same mission planning process but on a smaller scale. The mission operations replanning team selects alternate swaths that cover the missed target regions (step 1), revises the downlink schedule to accommodate them (step 2), and determines whether the resulting schedule is consistent with the operations constraints (step 3). If conflicts are found, the team returns to step 1 and selects different swaths. To minimize schedule disruption, the selected swaths must not overlap acquisitions already in the schedule, and existing acquisitions cannot be moved to make space for the new ones.

Rescheduling several swaths in response to anomalies, as occurred during AMM-1 operations, is a time- and knowledge-intensive task. In addition, mission time pressures demand that new plans be generated very quickly to exploit the next acquisition opportunity, usually within 24 to 36 hours. To generate plans within these time pressures, the AMM-1 mission required a staff of four working from pregenerated contingency plan segments. MAMM wanted to automate this replanning process to improve response time and reduce the burden on the operations team.

Application Description

The MAMM automated planner takes as input a set of desired SAR acquisitions and produces a combined acquisition and downlink schedule annotated with resource profiles and constraint violations. The violations are keyed to the SAR acquisitions that produced them so that the mission planner can quickly locate the offending acquisitions. It also provides summary information, such as total on-board resource use and total downlink time by the ground receiving station, for use in evaluating plan quality.

The human mission planner selects a set of swaths using a swath-selection tool called SPA, which CSA developed for *RADARSAT* missions. SPA generates a swath file that specifies the time, duration, and beam of each swath but leaves the downlink mode (real time or recorded) unspecified. This file is given to the planning system as input. The planning system has two other input files—(1) optimization criteria for the downlink schedule and (2) a mask file that specifies the in-view periods for each ground station—but these are fairly static. These files were modified to generate what-if plans for evaluating mission alternatives and were held fixed while the official mission plan was developed.

The mask and swath files are combined into a single file and passed to the ASPEN planning system, which is described in more detail later. The planner generates a downlink schedule for the swaths and then expands the resulting swath-and-downlink schedule into a more detailed plan that includes support activities such as tape on-off transitions and beam switches and tracks resource use. These details are needed to evaluate the operations constraints. ASPEN checks the plan for constraint violations and finally converts it from ASPEN format to an EXCEL spreadsheet format preferred by the mission planners.

The spreadsheet provides a time-ordered list of acquisition, playback, and downlink commands; identifies the swaths that violate

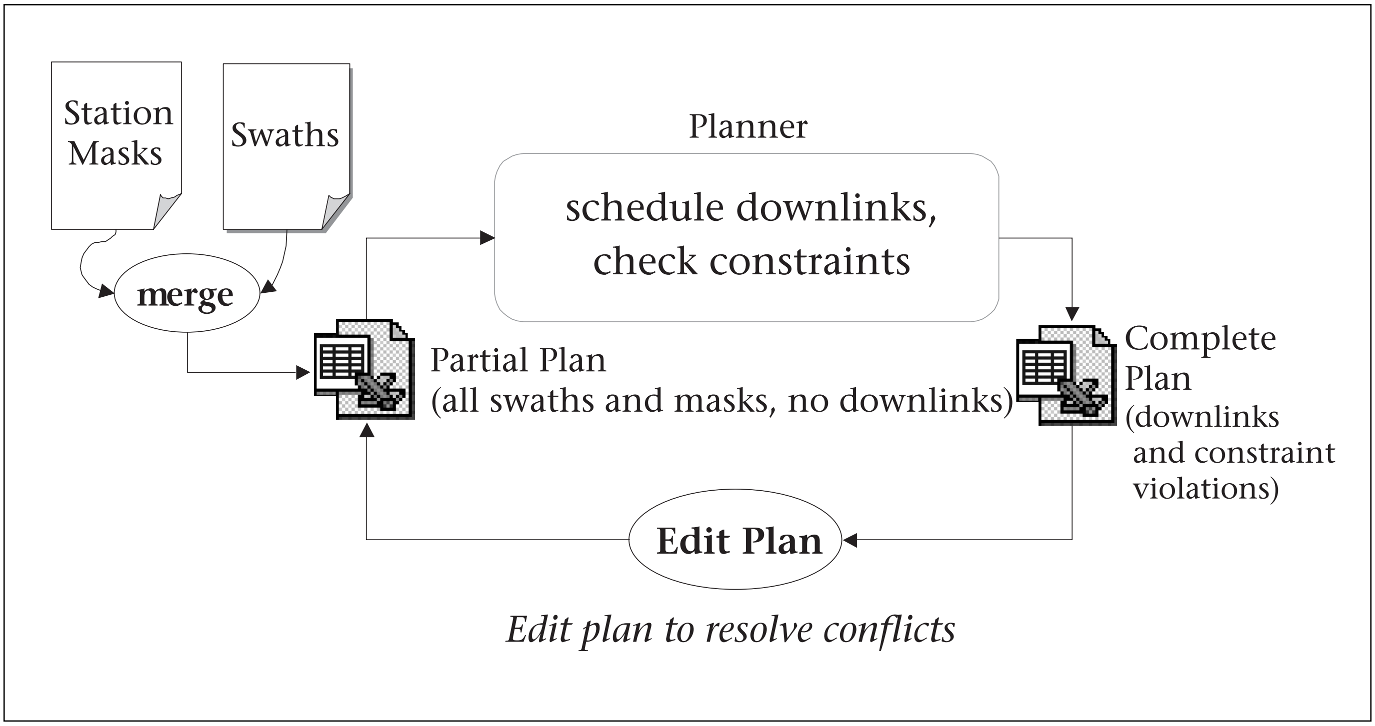

Figure 5. System Data Flow Architecture.

constraints or cannot be downlinked; and provides resource profiles. It also summarizes plan metrics such as resource use totals, ground station connect time (for costing), and the number of real-time and recorded acquisitions. Based on the report files, the human mission planner modifies the selected swaths as needed to resolve the conflicts or improve schedule quality.

Figure 5 summarizes this flow of information graphically. This check-and-edit cycle is repeated until a conflict-free plan is generated. This rapid feedback allows the user to generate a conflict-free plan much more quickly than is possible by hand. Maintaining the human planner in the loop enables the use of human scientific judgment in selecting swaths.

The ASPEN Planner

The core of the MAMM planner is ASPEN (Chien et al. 2000), an automated planning and scheduling system developed at the Jet Propulsion Laboratory (JPL). ASPEN takes high-level goals and produces detailed activity plans for achieving these goals. The ASPEN planning environment consists of a domain-modeling language, an incremental constraint-tracking facility (the plan database), interfaces for planning search algorithms, and a library of planning algorithms that exploit the plan database capabilities by way of these interfaces.

The MAMM planner encodes the *RADARSAT* operations constraints in the ASPEN domain-modeling language. It uses a domain-specific planning algorithm to schedule the downlink activities and expand the swath and downlink requests into a more detailed schedule. The planning algorithm then invokes the plan database's constraint-tracking facility to determine which domain constraints are violated. This structure is shown in figure 6.

ASPEN Domain-Modeling Language

The *RADARSAT* operations constraints are expressed in the ASPEN domain-modeling language (Sherwood et al. 1998). The elements in this language are activities, states, resources, and constraints. An *activity* is an action the spacecraft can perform, such as a data take or beam switch. Activities have a start time and a duration and can overlap each other. A resource represents a physical or logical resource of the spacecraft, such as the on-board recorder tape or instrument on time. A *state* represents a physical or logical state of the spacecraft, such as what the current SAR beam is or whether a given ground station is in view or not in view. Each state and resource is represented as a timeline that shows how it evolves over time.

The activities, states, and resources are related by constraints. These can be temporal constraints among activities (a tape spin down must immediately follow a data take), resource

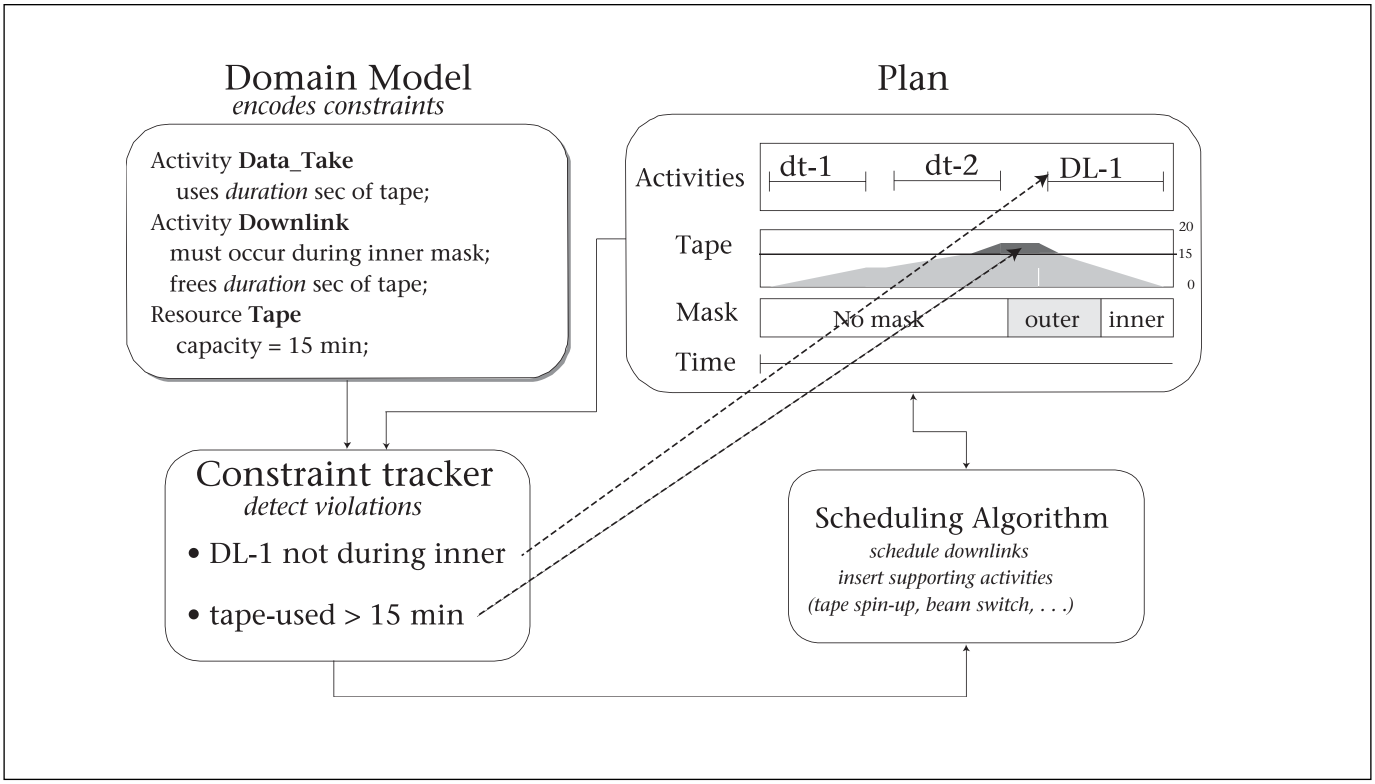

Figure 6. *ASPEN Planning Components.*

constraints (a data take uses *d* seconds of on-board recorder (OBR) tape, where *d* is the duration of the data take), and state constraints (the SAR instrument must be on during a data take). The MAMM operations constraints were encoded in terms of these constraints. Figure 7 shows how some of the MAMM domain knowledge was encoded in ASPEN. Figure 8 shows a sample plan fragment with each of these elements. The full ASPEN domain model has 6 resource timelines, 7 state timelines, and 27 activity types, as summarized in table 1.

Planning Domain Knowledge

For mission planning purposes, *RADARSAT* has two commands: (1) acquire SAR data and either save it to the on-board recorder or downlink it in real time as it is being acquired and (2) play back and downlink the SAR data on the recorder. An *SAR data-acquisition command* specifies the start time, duration, downlink mode, and beam. The *downlink mode* determines whether the data are saved to the OBR or downlinked in real time. The *beam* controls the incidence angle of the SAR instrument and determines which of several swaths parallel to the spacecraft ground track is acquired. The incidence angles of adjacent beams are separated by a few degrees and acquire data in rectangular swaths that partially overlap those of adjacent beams. Several swaths typically cover any given ground region, although these swaths are often in different orbits and beams.

Downlink (playback or real time) can only occur when the spacecraft ground track crosses within range of a ground receiving station (the station is in view). The *playback command* plays back and downlinks all the recorded data on the tape, then erases the tape. The spacecraft can downlink playback data while it also downlinks data being acquired in real time. The station in-view periods are called *masks* and are specified in a mask file provided by the *RADARSAT* MMO.

In addition, the mission plan must obey operations constraints imposed by the *RADARSAT* MMO, some of which are shown in table 2. These constraints primarily consist of resource constraints, setup times between data acquisitions, tape recorder and SAR instrument operating constraints, and downlink policy rules. The resources are on-board recorder capacity, tape transactions (number of times the tape has been started and stopped), and SAR instrument on time for each orbit. The relevant device states referenced by the operations constraints are the tape mode (idle, spinning up, recording, spinning down, playback) and the SAR beam (1 of 16).

Search Algorithm

ASPEN begins with a partial plan that includes the goals and initial states and typically violates several constraints. The search algorithm expands this partial plan into a detailed plan that achieves the goals and does not violate any constraints. ASPEN provides several general-purpose search algorithms and provides support for new application-specific algorithms. ASPEN typically uses an iterative repair (Zweben et al. 1994) search: In each iteration, the search algorithm selects a violated constraint (conflict) and applies an appropriate plan modification operator to resolve the conflict. For example, if two activities were in conflict because the partial plan scheduled them to use the SAR instrument at the same, ASPEN might resolve the conflict by moving one of the activities to a different time. This iteration continues until no conflicts remain, or the search exceeds a user-defined number of iterations. However, the MAMM application did not need iterative repair because the mission planners wanted to resolve all the conflicts themselves. Instead, MAMM used a special-purpose scheduling algorithm to quickly expand the acquisitions into supporting activities and schedule the downlinks.

The initial plan consists of a set of swath request activities and station mask activities. The algorithm first adds the mask activities to the database. The state constraints on these activities set the state timelines for each ground station. The planner then adds the swaths to the database and decides how to downlink them.

The downlink scheduling algorithm is a greedy one. In each iteration, it selects one swath and assigns it to the best downlink opportunity. If no assignment is possible, it backtracks. Because there can be no way to downlink all the selected swaths, it limits its backtracking to a two-orbit window. If no feasible solution can be found in this window, it selects a feasible schedule that downlinks the most data and reports the lost data as a constraint violation.

Once the algorithm has assigned to each swath a downlink mode and downlink opportunity, it reflects these assignments in the plan database. It grounds the downlink mode parameter of each swath to OBR or real time accordingly and creates a downlink activity for each mask that was assigned to one of the swaths.

At this point, the plan consists solely of swath, mask, and downlink activities. The planning algorithm then performs a limited expansion and grounding of the plan. In each iteration, it selects a value for an ungrounded activity parameter or adds an activity to satisfy an open temporal constraint. For example, if activity *A* is in the plan and has an open constraint that it must be before activity *B*, the planner will add an activity instance of type *B* just after activity *A*. At the end of this phase, the plan contains all the activities needed to acquire and downlink the requested swaths. The resource and state timelines have also been computed based on the state and resource constraints made by the activities in the plan.

```
Activity OBR_Data_Take {
    reservations =
      obr_storage use duration,
      obr_state must_be "record";
};
Activity spin_up {
   Duration = 1300;
   Reservations =
         obr_storage use duration, // consumes tape
         obr_state change_to "record";
};
Resource obr_storage {
   Type = depletable;
   Capacity = 91600; // 15.5 minutes = 91600 seconds
 };
State obr_state {
    States = ("idle", "playback", "record");
    Default_state = "idle";
    Transitions = ("idle"-> "playback", "idle"->"record"
                  "playback"->"idle", "record"->"idle");
};
```

Figure 7. ASPEN Domain-Modeling Example.

The Plan Database

The *plan database* records a partial plan and tracks the status of each constraint (satisfied or violated). Whenever the search algorithm modifies the partial plan, the plan database incrementally recomputes the relevant resource timelines and updates the list of conflicts—that is, constraints that are violated by the current partial plan. Conflicts include temporal violations (for example, data-take activities are too close together), resource violations (for example, exceeded tape capacity), and state violations (the tape goes from on to off without first going through the spin-down state). The algorithm does not attempt to fix the constraints, even though that is within

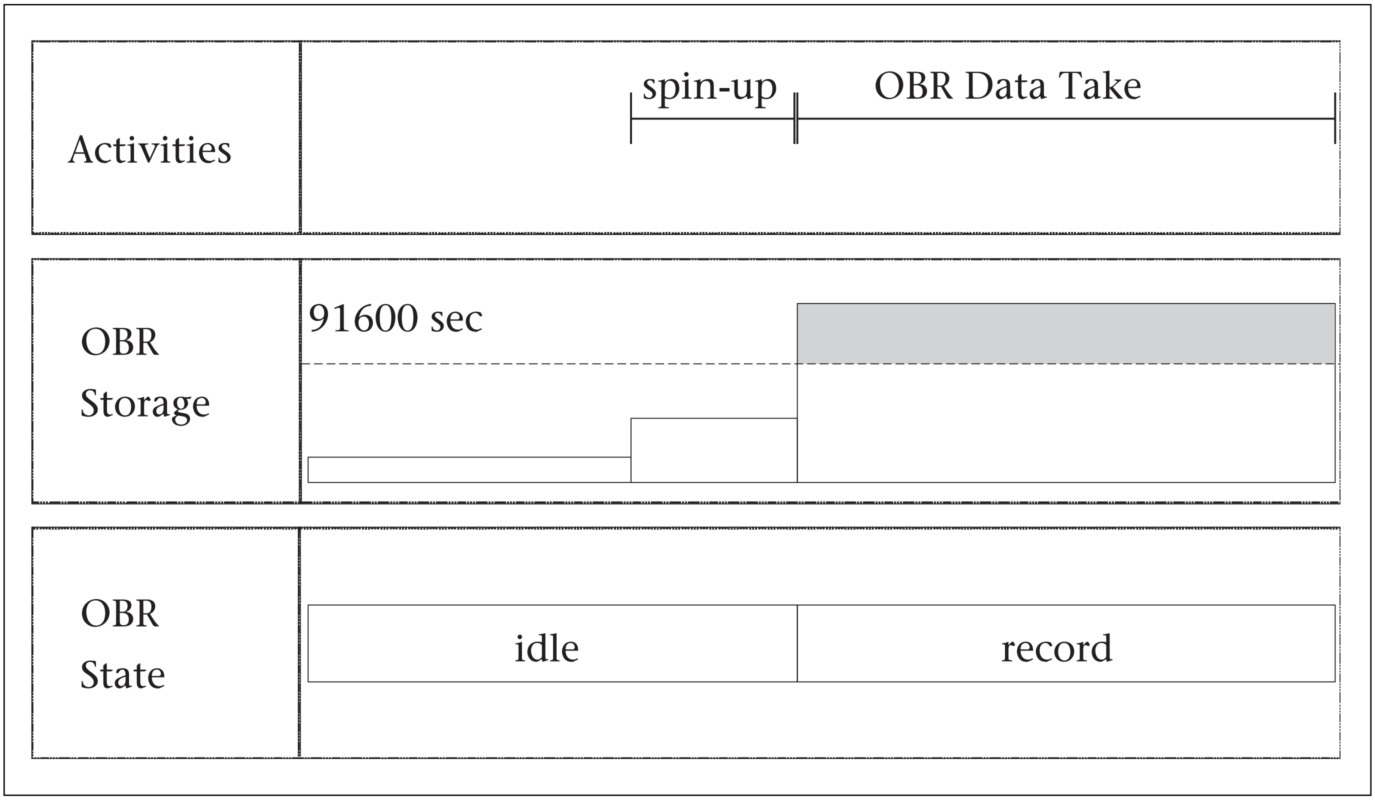

Figure 8. Plan Fragment Using Activities, States, and Resources Defined in Figure 2.

Each box on the timeline is a timeline unit and represents the value of this state or resource over this time period.

ASPEN's capabilities. The conflict resolution is intentionally left to the human mission planner because it involves swath-selection changes that require human scientific judgment.

Final Output

When ASPEN terminates, it saves the plan and constraint-violation information to a file, which is then converted into an EXCEL spreadsheet format preferred by the mission planners. This list is a time-ordered list of swath, mask, and downlink activities, with one row for each activity. The columns indicate activity parameter values and resource values as of the end of the activity. The last column holds a list of the operations constraint violations caused by the activity. A table maps ASPEN conflicts to corresponding high-level operations constraints, and it is these high-level constraints that are reported in the spreadsheet.

Planner Use and Benefits

A development version of MAMM was released to the MAMM mission planners in February 2000 for initial planning and evaluation and was officially deployed in April. The MAMM mission planners used the system from March through July to develop the MAMM mission plan as well as several draft plans and trade-study plans.

The plan development effort for MAMM using the automated system was about one-sixth the manual planning effort for AMM-1. The two missions were comparable: MAMM contained 818 acquisitions over 24 days (repeated three times), and AMM-1 contained 850 acquisitions over 18 days. The MMO review of the final MAMM plan detected no constraint violations, and the plan executed flawlessly on *RADARSAT* from September to December of 2000. In addition to reducing plan development costs, the system's ability to provide detailed resource use information and rapidly generate downlink schedules for different station availabilities and station priority policies were instrumental in evaluating mission alternatives, costing the mission, and negotiating resource quotas.

Mission Plan Development

The MAMM mission designers used the automated planner to develop a series of four draft plans and the final mission plan. Each draft was reviewed against scientific, cost, and risk

Activity (27)	Acquire_data
	Acquire_OBR_data
	Acquire_RTM_data
	Downlink
	Downlink_RTM
	Downlink_OBR
	State Changers (x11)
	Mask Timeline Setter (x10)
State (7)	Mask (x5 stations)
	Beam
	OBR Mode
Atomic Resources (2)	SAR in use
	OBR in use
Depletable Resources (4)	SAR on time
	OBR storage
	Tape transactions
	Data_not_downlinked

Table 1. Modified Antarctic Mapping Mission Domain Model Summary.

OBR = on-board recorder. RTM = real time. SAR = synthetic aperture radar.

criteria, and the results determined the swath-selection strategy for the next version. The average development time for each plan was about two workweeks. Roughly 60 percent of this time was spent in selecting the initial swath, 10 percent in using the automated planner (setting up runs, learning how to operate it, and getting the results), and 30 percent in revising the swaths to eliminate constraint violations detected by the planner. Removing the constraint violations required between one and four check-and-edit iterations. Table 3 summarizes the development times for each of the plan revisions.

The total development time for all MAMM plans was 10 workweeks. By comparison, mission planning for AMM-1 required over a workyear, with individual plans taking 3 to 4 months (12 to 16 workweeks) to develop. Overall, the automated planning system reduced planning effort from over 1 workyear to 10 workweeks, or a factor of 6.

If one includes the development time for the automated system, the automated approach is still 25 percent less effort than the manual one. The total planning and development effort for MAMM was about 9 workmonths (6.75 to develop the planner and 2.5 to develop the plans) compared to more than 12 workmonths for AMM-1.

Costing and Trade Studies

In addition to reducing development costs, the automated system provided valuable information for the plan-evaluation phases. For each plan, it provided detailed resource and summary information that informed the cost and risk assessments. It also automatically generated what-if variations of draft plans for evaluating mission alternatives. The mission designers and project managers perceived both of these capabilities as highly beneficial, and the information was directly used to estimate ground-station costs and negotiate *RADARSAT* resource quotas.

Some of the questions that were answered by the what-if capability during the mission design process are described here.

What are the resource requirements for these observations? This information is needed early in the planning process to negotiate the spacecraft resource allocations. Because *RADARSAT* is shared by other scientific and industrial customers, MAMM had to operate within negotiated resource limits rather than have full access to all the spacecraft resources. This question was addressed with summary statistics that the system generates for each plan. These statistics include total OBR use, SAR on time, and total downlink data time bro-

Data can only be downlinked when a ground station is in view.
All recorded data must be downlinked.
OBR playback can only occur during downlink.
SAR acquisitions cannot overlap.
Cannot transmit RTM data when recorder is in record, spin-up, or spin-down modes.
Data takes will be no shorter than 1.0 meters.
Adjacent data takes will be at least 5.25 seconds apart when beams are changed.
Data takes will be at least 11 seconds apart when beams are not changed.
OBR takes 10 seconds to spin up, consumes 10 seconds of tape.
OBR takes 5.5 seconds to spin down, consumes 5.5 seconds of tape.
OBR transitions to idle between takes iff OBR data takes are >30 seconds apart, else continues recording.
There will be ⇐ 6 OBR transactions an orbit.
SAR will be on at most 32.0 minutes an orbit.

Table 2. Selected Operations Constraints.

OBR = on-board recorder. SAR = synthetic aperture radar. RTM = real time.

ken down by station. These first two were invaluable in negotiating on-board resource allocations. The downlink durations by station were used to estimate ground station costs, forecast use levels, and schedule downlink sessions. The detail and early availability of these schedules greatly simplified this process over AMM-1.

How do different downlink scheduling policies impact the mission plan? The downlink policy is implemented as an objective function that the downlink schedule tries to maximize. The mission manager had several policy variations that he wanted to evaluate. His question was addressed by generating what-if downlink plans for each objective function.

What is the impact of not using certain ground stations? During the early stages of mission planning, MAMM wanted to know whether an additional, and expensive, ground station was needed to receive data acquired over a region that had limited coverage by the other ground stations. This question was addressed by generating what-if plans in which this station was not available and plans in which it was available. The mission planners reviewed these plans to understand the impact of ground-station availability on the science plan, spacecraft resources, and other downlink stations. These data enabled the mission planners to determine that the station could safely be eliminated.

Anomaly Replanning during Mission Operations

As discussed earlier, scheduled data acquisitions can be lost during operations because of spacecraft or ground station anomalies. The operations team reschedules these acquisitions for a later opportunity. However, the replanning staff must submit the rescheduled swaths at least 36 hours before they are executed to provide the MMO enough time to process and uplink the requests. In most cases, the replanning staff has to submit a new acquisition plan within 24 to 36 hours of the anomaly. To manually turn around plans within these time constraints, AMM-1 required a team of four people working from pregenerated contingency plan segments. The missed observations were placed into gaps in the original plan to minimize coverage holes. More extensive changes, such as altering the remaining (unexecuted) planned swaths, were avoided to minimize the planning effort and the chance of introducing errors into the plan. Unfortunately, it was sometimes impossible to find a way to reschedule all the missed observations within this time frame using these manual procedures. These observations were simply dropped from the schedule.

For MAMM, the automated planner was available during operations for identifying operations conflicts in manually generated replan schedules. The system took as input the replanned schedule and provided a list of conflicts within minutes. This capability enabled the replanning team to quickly identify and correct any constraint violations before submitting it to the MMO for a final (and more costly) check. If this system had been available on AMM-1, which had 10 spacecraft anomalies and lost a primary ground receiving station early in the mission, the AMM-1 project manager estimated that the replanning staff could have been reduced from four to one. The replanning system was part of the mission operations software and performed well during operations rehearsals. Unfortunately for the system—although fortunately for the mission—there were no anomalies during the mission that required its use.

Development and Deployment

The automated planning system was developed using the ASPEN planning environment. ASPEN provided the domain-modeling language and constraint-checking facilities. The development process was fairly typical: Acquire the specifications and domain knowledge (operations constraints), encode the knowledge, develop the infrastructure, and then test it. The work force breakdown is shown in table 4.

The development process was repeated over two iterations. The first iteration (R1) produced an operational system that had the most critical capabilities and operational constraints. This system was used to develop a draft plan for use in making cost and feasibility estimates. This development process also provided feedback on ease of operability and needed and unnecessary capabilities and uncovered some minor refinements to the operations constraints. Development of R2, the second and final version, was informed by the feedback from R1. The total work effort was just under seven workmonths.

Related Work

There are a number of related systems for automating space mission planning. APGEN (Maldague et al. 1997) is a mainstream planning system used by several flight projects at JPL. APGEN has a procedural scripting language for generating plans but does not have a declarative, model-based planning and scheduling engine. The SPIKE scheduler (Johnston and Miller 1994) was initially developed for the Hubble Space Telescope (HST) and has since been used for other applications. HST is a "general observatory" class mission on which many astronomers request observing time. The focus of SPIKE is to optimize observation schedules to satisfy as many requests as possible. The focus of the MAMM planner, however, is to reduce the effort of constructing a single feasible plan. The ASTER scheduler (Muraoka et al. 1998) constructs near-term observation schedules for the advanced spaceborne thermal emission and reflection radiometer instrument on the *Terra* satellite. It does not separate the scheduling engine from the mission domain knowledge but combines them within a single mission-specific scheduling algorithm.

Version	Date	Iterations	Workweeks
1	3/6	3	2
2	4/12	2	2
3	4/27	2	2
4	6/8	4	3
Final	6/19	1	1
TOTAL		12	10

Table 3. MAMM Plan Development Effort.

Task	R1	R2	Total
Knowledge Acquisition	1.0	0.5	1.5
Knowledge Engineering	6.0	2.0	8
Scheduling and Downlink Algorithm	2.0	1.0	3
Infrastructure	6.0	2.0	8
Testing	1.0	6.0	6
Total	16.0	11.5	27.5

Table 4. Application Development Effort in Workweeks.

Conclusions

Mission planning is a time- and knowledge-intensive task. It required over a workyear to manually develop the mission plan for AMM-1. We developed an automated planning system that reduced the mission planning time for MAMM, the follow-on mission to AMM-1, to just 10 workweeks. In addition to reducing mission-planning effort, it also enabled rapid generation of what-if plans for evaluating mission alternatives and provided resource use information that was used for costing the mission and negotiating spacecraft resource allocations. These what-if analyses contributed to the quality and success of the mission, and the mission planners considered this capability an invaluable tool. Automated planning was overwhelmingly successful for MAMM, and we would expect similar successes for future *RADARSAT* missions.

Acknowledgments

This article describes work performed at the Jet Propulsion Laboratory (JPL), California Institute of Technology, under contract from the National Aeronautics and Space Administration (NASA). The Antarctic Mapping Mission is a joint project between NASA and the Canadian Space Agency (CSA). The project is led by Ohio State University at Columbus in partnership with the Alaska Synthetic Aperture Radar Facility at the University of Alaska at Fairbanks, the JPL, California Institute of Technology, and the Vexcel Corporation. The authors want to thank John Crawford for championing this technology; Richard Austin, the mission planner, for using the system and providing design feedback; the CSA for providing the *RADARSAT* operations constraints; and Kim Partington for his support and sponsorship during his tenure as manager of NASA's Polar Science Program.

References

Chien, S. A; Rabideau, G. R.; Knight, R. L.; Sherwood, R. L.; Engelhardt, B. E.; Mutz, D. H.; Estlin, T. A.; Smith, B. D.; Fisher, F. W; Barrett, A. C.; Stebbins, G.; and Tran, D. Q. 2000. ASPEN—Automating Space Mission Operations Using Automated Planning and Scheduling. Paper presented at the Sixth International Symposium on Technical Interchange for Space Mission Operations and Ground Data Systems (SpaceOps 2000), 19–23 June, Toulouse, France.

Jezek, K. C.; Sohn, H. G.; and Noltimeir, K. F. 1998. The *RADARSAT* Antarctic Mapping Project. Paper presented at the International Geoscience and Remote Sensing Symposium (IGARSS '98), 6–10 July, Seattle, Washington.

Johnston, M., and Miller, G. 1994. SPIKE: Intelligent Scheduling of Hubble Space Telescope Observations. In *Intelligent Scheduling,* eds. M. Zweben and M. S. Fox, 391–422. San Francisco, Calif.: Morgan Kaufmann.

Maldague, P.; Ko, A.; Page, D.; and Starbird, T. 1997. APGEN: A Multimission Semiautomated Planning Tool. Paper presented at the First International Workshop on Planning and Scheduling for Space, 28–30 October, Oxnard, California.

Muraoka, H; Cohen, R. H.; Ohno, T.; and Doi, N. 1998. ASTER Observation Scheduling Algorithm. Paper presented at the Fifth International Symposium on Technical Interchange for Space Mission Operations and Ground Data Systems (SpaceOps '98), 1–5 June, Tokyo, Japan.

Sherwood, R. L; Govindjee, A.; Yan, D.; Rabideau, G. R.; Chien, S. A.; and Fukunaga, A. S. 1998. Using ASPEN to Automate EO-1 Activity Planning. In CD-ROM Proceedings of the IEEE Aerospace Conference. Washington, D.C.: IEEE Computer Society.

Zweben, M; Daun, B.; Davis, E; and Deale, M. 1994. Scheduling and Rescheduling with Iterative Repair. In *Intelligent Scheduling,* eds. M. Zweben and M. S. Fox, 241–256. San Francisco, Calif.: Morgan Kaufmann.

Benjamin Smith is the manager of the Autonomy Technology Program at the Jet Propulsion Laboratory and was the task lead for the MAMM automated planner. His personal research interests are in planning and scheduling and machine learning. He received his Ph.D. in computer science from the University of Southern California in 1995.

Barbara E. Engelhardt is a part-time member of the technical staff in the Artificial Intelligence Group at the Jet Propulsion Laboratory and a Ph.D. student at the University of California at Berkeley in computer science. Engelhardt was a member of the Replanning Team during MAMM mission operations. Her research interests are in probabilistic reasoning with graphic models, machine learning, and data mining. She received her B.S. and M.S. from Stanford University in 1998.

Darren H. Mutz is a doctoral student and a member of the Reliable Software Group in the Computer Science Department at the University of California at Santa Barbara. He was a member of the Artificial Intelligence Group at the Jet Propulsion Laboratory from 1995 to 2001. His research interests are in the area of computer and network security. He holds a B.S. in computer science from the University of California, Santa Barbara.