

The RadarSAT-MAMM Automated Mission Planner

Benjamin D. Smith

Barbara E. Engelhardt

Darren H. Mutz

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
{firstname.lastname}@jpl.nasa.gov

Abstract

The RadarSAT Modified Antarctic Mapping Mission (MAMM) ran from September to November 2000. It consisted of over 2400 synthetic aperture radar (SAR) data takes over Antarctica that had to satisfy coverage and other scientific criteria while obeying tight resource and operational constraints. Developing these plans is a time and knowledge intensive effort. It required over a work-year to manually develop a comparable plan for AMM-1, the precursor mission to MAMM. This paper describes 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 mission-design trades. This latter capability informed several critical design decisions and was instrumental in accurately costing the mission.

Introduction

The Modified Antarctic Mapping Mission (MAMM) executed from September through November of 2000 onboard RadarSAT, a Canadian Space Agency (CSA) satellite. This joint NASA/CSA mission is a modified version of the First RadarSAT Antarctic Mapping Mission (AMM-1) executed in 1997 (Jezek *et al.*, 1998). The objective of AMM-1 was to acquire complete coverage of the Antarctic continent, whereas the objective of MAMM is to acquire repeat-pass interferometry to measure ice surface velocity of the outer regions of the continent, north of latitude -80 degrees. The mission objective is to perform synthetic aperture radar (SAR) mapping of the Antarctic over three consecutive 24-day repeat cycles.

Planning SAR mapping missions is a time- and knowledge-intensive process. RadarSAT has a SAR instrument that can be commanded to acquire data in any one of several rectangular swaths parallel to its ground track. The spacecraft can also downlink acquired data to ground receiving stations when its ground track passes over them. The planning problem is to select a set of swaths and downlinks such that the swaths cover the desired region of Antarctica and satisfy science requirements, and the combined acquisition and downlink schedule meets the operational and resource constraints imposed by the RadarSAT Mission Management Office (MMO). The driving operational constraints are the limited on-board tape recorder (OBR) capacity and downlink opportunities,

which together constrain the amount of swath data that can be acquired and saved on the OBR between downlinks.

The AMM-1 mission demonstrated the need for an automated planning capability. The schedule for AMM-1 consisted of 850 acquisitions (swaths) over 18 days, and took over a work-year to develop manually. Despite repeated checking, this plan violated operations constraints that were not detected until the final MMO review. This inability to detect all the operations and resource constraint violations during the planning process required expensive and disruptive last-minute revisions.

This experience led to the development and use of an automated mission planning system for MAMM. The system takes a set of swaths selected by the human mission planner, automatically generates a downlink schedule, then expands the swaths and downlinks into a more detailed plan that it checks for operations constraint violations. With this system MAMM developed its 24-day mission plan containing 818 swaths in a matter of weeks, as compared to the work-year required to develop a comparable mission plan for AMM-1.

In addition to reducing the plan development effort, the MAMM planner also provides resource tracking and other plan details that enable accurate costing and feasibility estimates. The MAMM planner also enables “what-if” studies that were not possible under AMM-1. The planner quickly generated detailed variations of the baseline plan for different ground station availability assumptions. These study plans were instrumental in selecting ground stations and making other decisions about mission alternatives.

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

Mission Planning Problem

The objective of MAMM is to acquire repeat-pass SAR interferometry of Antarctica north of -80 degrees latitude to measure ice surface velocity of the outer regions of the continent.

For mission planning purposes, RadarSAT has two commands: (1) acquire SAR data in one of several rectangular swaths parallel to the spacecraft ground track and either save it to the onboard recorder or downlink it in real time as it is being acquired; and (2) playback and

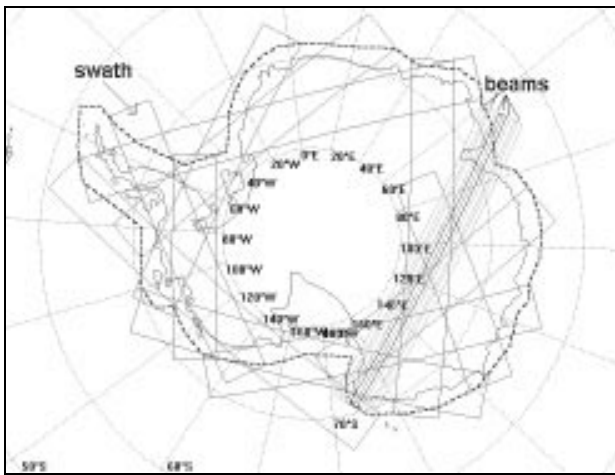


Figure 1: Swath Selection in SPA.

downlink the SAR data on the recorder. The mission-planning problem is to select a subset of the possible swaths and downlink opportunities (real-time and playback) such that the resulting schedule satisfies the scientific requirements and the operating constraints.

A data acquisition command specifies the start time, duration, downlink mode, and beam. The downlink mode determines whether the data is saved to the onboard recorder (OBR) or downlinked in real-time (RTM). 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 those swaths are often in different orbits and/or different beams.

The playback command plays back and downlinks all the recorded data on the tape, then erases the tape.

Downlink (playback or real-time) may only occur when the spacecraft ground track crosses within range of a ground receiving station (the station is *in-view*). The spacecraft may downlink playback data while also downlinking 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 Mission Management Office (MMO).

In addition to the above, the mission plan must obey operations constraints imposed by the RadarSAT Mission Management Office (MMO), some of which are shown in Table 1. These primarily consist of resource constraints, set-up times between data acquisitions, tape recorder and SAR instrument operating constraints, and downlink policy rules. The resources are onboard recorder capacity, tape transactions (number of times the tape has been started and stopped), and SAR instrument on-time per orbit. The relevant device states referenced by the operations constraints are the tape mode (idle, spinning up, recording,

Data can only be downlinked when a ground station is in view
All recorded data must be downlinked
OBR playback may 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 shall be no shorter than 1.0m
Adjacent data takes shall be at least 5.25s apart when beams are changed
Data takes shall be at least 11s apart when beams are not changed.
OBR takes 10s to spin up, consumes 10s of tape
OBR takes 5.5s to spin down, consumes 5.5s tape
OBR transitions to idle between takes iff OBR data takes are > 30s apart, else continues recording.
There will be <= 6 OBR transactions per orbit
SAR shall be on at most 32.0 minutes per orbit

Table 1: Selected Operations Constraints

spinning-down, playback) and the SAR beam (one of sixteen).

The Planning Process

The mission planning process is an iterative one. The mission planner develops several plan versions before arriving at the final mission plan. Each version is reviewed against scientific, cost, and risk criteria. This analysis informs the approach for developing the next iteration, sometimes drastically. MAMM generated four revisions before arriving at the fifth and final mission plan. The process for generating an individual plan consists of the following four steps. The resulting plan is a time-ordered list of data acquisition requests and downlink session requests.

1. **Select SAR swaths** that cover the desired target regions in Antarctica and satisfy other scientific requirements. The swaths are selected from all the swaths that intersect the target region during one 24-day repeat cycle. This step is partially automated by a tool developed by the Canadian Space Agency called SPA [7] that identifies the available swaths for each beam as shown in Figure 1 by propagating the spacecraft orbit. The user selects the desired (sub-) swaths, and SPA generates a swath request file. SPA does not check operations constraints or ensure that the swaths can be downlinked, so there is no guarantee that the selected swaths comprise a valid mission plan.
2. **Create a downlink schedule.** The downlink schedule specifies which station masks (downlink opportunities) will be used to downlink the data acquisitions, and specifies for each acquisition (swath) whether it will be downlinked in real-time or stored to the data recorder. The schedule must obey resource and operations constraints. In particular, real-time acquisitions must

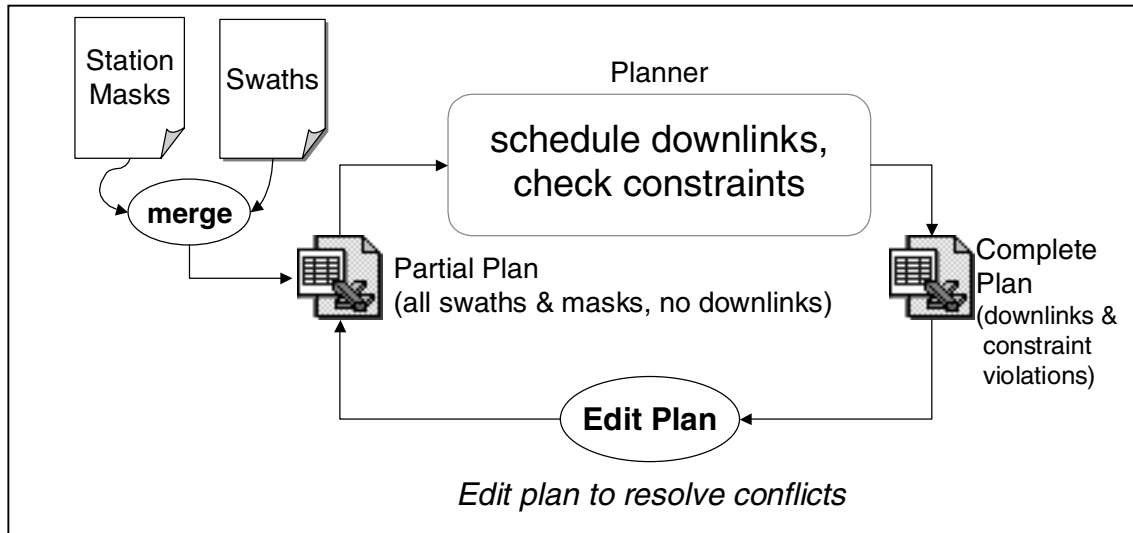


Figure 2: System Data Flow Architecture

occur during a downlink session, and playback sessions must be scheduled during masks that are long enough to downlink all the data on the onboard recorder. The schedule should also try to maximize objective criteria—certain stations are more reliable or have lower costs than others; and resource costs make real-time takes preferable to recorded ones.

3. **Compute resource usage and check for constraint violations.** Determine whether the composite acquisition and downlink plan violates any operations constraints. Checking resource-related constraints requires computing usage profiles for each resource (OBR storage, SAR on-time).
4. **Address violations.** If the schedule violates operations constraints, cannot downlink all of the acquisitions, or is of insufficient quality return to Step 1 and modify the selected swaths to correct the problems. Modifications include changing the swath start-time, swath duration, and/or beam; or selecting an alternate swath that covers the same target area.

Part of what makes the mission-planning problem difficult is the interaction between swath selection (Step 1) and downlink scheduling (Step 2). The ground stations are rarely in view when the satellite is over the desired regions of Antarctica, which means many of the acquisitions have to be recorded and downlinked later. Since playing back data for downlink erases the tape and the in-view periods are shorter than the tape capacity, the swaths selected for a given orbit must fit within the downlink opportunities near that orbit. If the scientifically desired swaths do not fit, an alternate swath must be selected. The alternate may be in a different orbit, which could force swath reselection for that orbit as well.

Replanning During Operations

Planned SAR acquisitions can be lost during operations due to spacecraft and ground station anomalies. Lost acquisitions are rescheduled using the same mission planning process on a smaller scale: select alternate swaths that covers the missed target regions (Step 1), revise the downlink schedule to accommodate them (Step 2), and make sure the resulting schedule is consistent with the operations constraints (Step 3). If conflicts are found, return to Step 1 and select different swaths. In order 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, as can happen with a major anomaly, is a time- and knowledge-intensive task. In addition, mission time-pressures demand that new plans be generated very quickly in order to exploit the next acquisition opportunity, usually within 24 to 36 hours. AMM-1 required a staff of four working from pre-generated contingency plan segments in order to generate plans within these time pressures.

Application Description

The mission planning application automates Steps 2 and 3. The other steps were intentionally not automated since they involve swath selection, which requires human scientific judgment.

The human mission-planner selects a set of swaths (Step 1) using a swath selection and coverage analysis tool called SPA, which CSA developed for RadarSAT missions. The swath input specifies the time, duration, and beam of each swath. These are passed to the planning system along with

downlink priority policy and a mask file, provided by the RadarSAT Mission Management Office (MMO), that specifies the *in-view* periods for each ground station.

The mask and swath files are combined into a single file and passed to the ASPEN planning system, which is described in more detail below. The planner generates a downlink schedule for the swaths (Step 2), 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 usage. This provides the additional details referenced in the operations constraints. ASPEN checks the plan for constraint violations (Step 3), 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 constraints or cannot be downlinked; and provides resource profiles. It also summarizes plan metrics such as resource usage 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 (Step 4).

Figure 2 summarizes this flow of information (Step 1- 4) 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 MAMM planning system is implemented in C++ and runs on a SUN Ultra/60 workstation. The conversion utilities (from SPA to ASPEN, and from ASPEN to Excel) were written in Perl.

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. 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 via those interfaces. The plan database records a partial plan and the constraints that are satisfied and violated by that plan. The plan database supports several plan-modification operators, an operation for incrementally propagating the constraints following modifications, and interfaces for accessing the constraint and plan element information in the database. Search algorithms use these capabilities to determine how to modify the current plan. For a given application one can select one of the general-purpose algorithms in the library or develop a new application-specific algorithm.

The MAMM planner encodes the 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 calls the constraint update operation to determine which domain constraints are violated. This structure is shown in Figure 3.

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 is a time-ordered list of swath, mask, and downlink activities, with one row for each activity. There is one column for each resource. The value of that column for each activity (row) is the value of that resource at the end

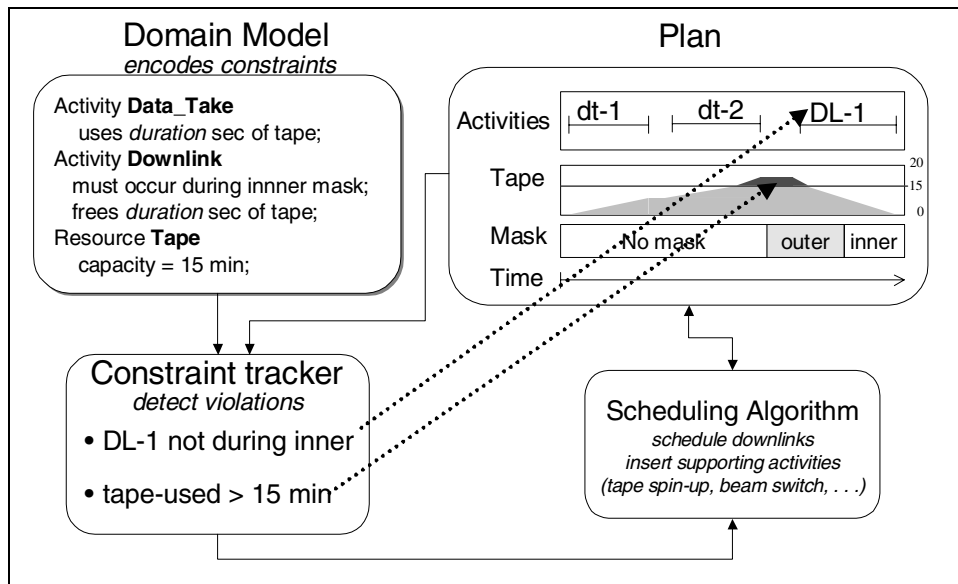


Figure 3: ASPEN planning components

of that activity. The last column holds a list of the operations constraint violations in which that activity is involved. A table maps ASPEN conflicts to corresponding high-level operations constraints, and it is these high-level constraints that are reported in the spreadsheet.

Knowledge Representation

The RadarSAT operations constraints are expressed in the ASPEN domain modeling language. 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 duration and may overlap each other. A resource represents a physical or logical resource of the spacecraft, such as the onboard recorder tape or instrument on-time. A state represents a physical or logical state of the spacecraft, such as the current SAR beam 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 constraints (a data take uses *d* seconds of 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 4 shows how some of the MAMM domain knowledge was encoded in ASPEN. Figure 5 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 2.

```

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 4: ASPEN Domain Modeling Example.

Activity (27)	Acquire_data
	Acquire_OBR_data
	Acquire_RTM_data
	Downlink
	Downlink_RTM
	Downlink_OBR
	State changers (x 11)
	Mask timeline setter (x 10)
State (7)	Mask (x 5 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 2: MAMM Domain Model Summary

Scheduling Algorithm

The MAMM planner uses a domain-specific planning algorithm to control the plan database. 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 problem is a constrained assignment problem. Each swath must be assigned exactly one 'mode' (real-time or recorded) and exactly one downlink opportunity. That assignment must satisfy the domain constraints—specifically, recorded swaths must not exceed the tape capacity between downlinks and the downlink opportunity must be longer than the amount of recorded data; real-time swaths must be taken while a real-

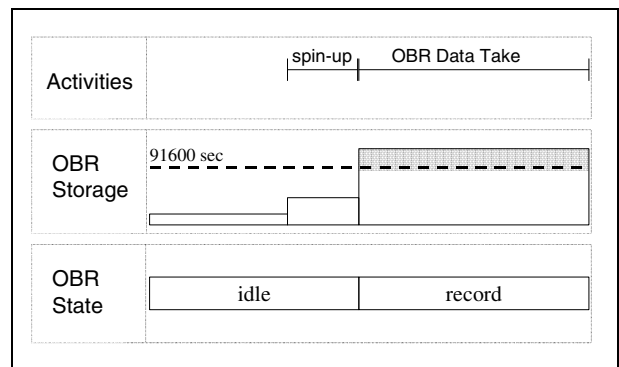


Figure 5: Plan Fragment using activities, states, and resources defined in Figure 2. Each box on the timeline is a *timeline unit* and represents the value of that state or resource over that time period.

time capable station is in-view.

The system employs a greedy algorithm to solve the downlink-scheduling problem. In each iteration it makes the best feasible assignment. If no assignment is possible, it backtracks. Since there may 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 that 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 RTM 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 of 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.

Finally, the algorithm invokes ASPEN’s constraint tracker to identify *conflicts*: violations of constraints in the domain model. These consist of temporal violations (e.g., data take activities are too close together), resource violations (e.g., exceeded tape capacity), and state violations. The algorithm does *not* attempt to fix the constraints, even though that is within ASPEN’s capabilities. The conflict resolution is intentionally left to the human mission planner since it involves swath-selection changes that require human scientific judgment.

Planner Use and Benefits

A development version of MAMM was released to the MAMM mission planners in February of 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 of 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

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.

resource usage information and to 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.

Based on the overwhelming success of this planning system, efforts are now underway to make it available for evaluation at the Alaska SAR Facility (ASF), whose charter includes developing RadarSAT data acquisition plans to satisfy the observation requests of a large scientific community.

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 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% of that time was spent in the initial swath selection, 10% in using the automated planner (setting up runs, learning how to operate it, and getting the results), and 30% revising the swaths to eliminate constraint violations detected by the planner. Constraint violations were removed in 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 work-year, with individual plans taking 3-4 months (12-16 workweeks) to develop. Overall the automated planning system reduced planning effort from over a work-year to 10 workweeks, or a factor of six.

If one includes the development time for the automated system, the automated approach is still 25% less effort than the manual one. The total planning and development effort for MAMM was about 9 work-months (6.75 to develop the planner, and 2.5 to develop the plans) as compared to over 12 work-months for AMM-1. If the system is adopted by ASF those development costs will be amortized over future missions, yielding even greater cost-savings.

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 specific questions it was used to answer during the mission design process are as follows.

1. *Determine the resource requirements for purposes of costing the mission and negotiating spacecraft resource allocations with the CSA.*

This question was addressed with summary statistics that the system generates for each plan. These include total on-board recorder usage, SAR on-time, and total downlink data time broken down by station. This first two were invaluable in negotiating on-board resource allocations. The downlink durations by station were used to estimate ground station costs, forecast usage levels, and to schedule downlink sessions. The detail and early availability of these schedules greatly simplified this process over AMM-1.

2. *How do different downlink scheduling policies impact the mission plan?*

This question was addressed by performing what-if simulations using the ASPEN system. Since downlink station priorities were one of the parameters of the downlink generation algorithm, the plan was expanded and downlink schedules generated using four different possible priority systems. ASPEN supplied the data to reach a decision on the priorities and significantly impact the mission negotiations during the early stages.

3. *What is the impact of not using certain ground stations?*

This question was addressed using what-if scenarios in which ASPEN was not allowed to downlink data to certain stations. This was accomplished by simply excluding the masks for those stations from the input file—the station was never in-view, and therefore never available for downlink. This enabled a closer examination of the impact of removing a ground station on the other stations and on the science collection in general. Using this information, the mission identified an unnecessary ground station early in the mission planning phase, and saved a significant amount of funding that would otherwise have been needed to support that station during operations.

Anomaly Replanning During Mission Operations

Spacecraft or ground station anomalies during operations can cause scheduled data acquisitions to be lost. These acquisitions can be rescheduled.

The operations re-planning 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 this means the replanning staff has to submit a new acquisition plan within 48 to 72 hours of the anomaly. To manually turn around plans within these time constraints on AMM-1 required a team of four people working from pre-generated 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 that 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.

Use of the system for anomaly re-planning was part of the operations procedures, was available during operations, and successfully replanned simulated anomalies during the operations rehearsals. However, it was never needed during the mission. Few anomalies occurred in the first cycle, and they only impacted acquisitions that could be manually rescheduled trivially and confidently.

Nevertheless, this capability is expected to be useful on future missions. If it 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 re-planning staff could have been reduced from four to one.

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 was used to develop a draft plan for use in making costing and feasibility estimates. That development process also provided feedback on ease of operability, 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 7 work months.

TASK	R1	R2	Total
Knowledge Acquisition	1.0	0.5	1.5
Knowledge Engineering	6.0	2.0	8
Scheduling & 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.

Difficulties

The primary difficulty was in the size of the plans. A typical 24-day MAMM input plan has over 800 swaths and 1,000 downlink masks, and expands into a plan with over 8,000 activities and 16,000 timeline units. ASPEN typically generates plans about a tenth this size in a few minutes, but these large plans require about an hour to generate. The performance degradation was a result of constraint propagation costs and memory swapping.

To reduce propagation costs we redesigned the scheduling algorithm to eliminate unnecessary “downstream” propagation. When an activity is added to the schedule and imposes a resource reservation, it forces all of the resource timeline units downstream of the activity to be recomputed. Placing activities in increasing time order, where possible, minimizes the number of downstream activities. The algorithm uses heuristics ensure the most computation-efficient ordering.

We further improved performance by re-engineering the domain model to minimize the size of the expanded plan. This reduced the expanded plan for an 818-swath input from about 12,000 activities and 20,000 timeline units to 8,000 activities and 16,000 timeline units, or about 25%. This reduced the plan size below the memory limit where swapping drove the performance to unacceptable levels.

Without these improvements a typical 800-swath plan required over 10 hours to run. With the modifications, the expanded plan was about 25% smaller and only required about an hour to process.

Lessons. Very large planning problems encounter performance issues that do not arise for more moderate problem sizes. The impact of performance tuning on development and maintenance need to be considered in projecting costs and selecting planning systems.

Maintenance

Maintenance has not yet been an issue. The RadarSAT operations constraints have been static for several years and are expected to remain so. Should maintenance be needed, the update mechanism is to modify the domain model and, if necessary, update the expansion-ordering heuristics. End-users should be able to make simple modifications to the APEN model themselves. The language is designed for non-AI experts, and such personnel have successfully

developed detailed ASPEN models (Willis, Rabideau, and Wilklow 1999; Sherwood *et al.* 1998). However, major changes would probably require additional performance tuning, which would require an experienced developer.

Conclusions

Mission planning is a time- and knowledge-intensive task. It required over a work-year 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 two work-months. In addition to reducing mission planning effort it also enabled rapid generation of “what-if” plans for evaluating mission alternatives, and provided resource usage information that was used for costing the mission and negotiating spacecraft resource allocations.

These 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.

Acknowledgements

This paper describes work performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract from the National Aeronautics and Space Administration. The authors also wish to thank John Crawford, the MAMM project manager, for championing this technology; Richard Austin, the mission planner, for using the system and providing design feedback; the Canadian Space Agency for providing the RadarSAT operations constraints, and Dr. Kim Partington for his support and sponsorship during his tenure as manager of NASA’s Polar Science Program.

References

- Chien, S.; Rabideau, G.; Knight, R.; Sherwood, R.; Engelhardt, B.; Mutz, D.; Estlin, T.; Smith, B.; Fisher, F.; Barrett, T.; Stebbins, G.; and Tran, D. 2000. ASPEN—Automating Space Mission Operations using Automated Planning and Scheduling. In *SpaceOps 2000*. Toulouse, France.
- Jezeq, K.C., H.G. Sohn, and K.F. Noltmeir. 1998. The RadarSAT Antarctic Mapping Project. In *Proceedings of IGARSS '98*. p. 2462-2464.
- Sherwood, R.; Govindjee, A.; Yan, D.; Rabideau, G.; Chien, S.; and Fukunaga, A. 1998. Using ASPEN to Automate EO-1 Activity Planning. In *Proceedings of the IEEE Aerospace Conference*. Aspen, CO. IEEE Press.
- Willis, J.; Rabideau, G.; and Wilklow, C. 1999. The Citizen Explorer Scheduling System. In *Proceedings of the IEEE Aerospace Conference*. Aspen, CO. IEEE Press.