

SCHEDULING SPACECRAFT OPERATIONS

Daniel L. Britt Amy L. Geoffroy Philip R. Schaefer
Martin Marietta Space Systems
P.O. Box 179, MS 0427
Denver, Colorado 80201

John R. Gohring
Martin Marietta Data Systems
116 Inverness Drive East
Englewood, Colorado 80112

ABSTRACT

A prototype scheduling system named MAESTRO currently under development is being used to explore possible approaches to the spacecraft operations scheduling problem. Results so far indicate that the appropriate combination of heuristic and other techniques can provide an acceptable solution to the scheduling problem over a wide range of operational scenarios and management approaches. These can include centralized or distributed instrument or systems control, batch or incremental scheduling, scheduling loose resource envelopes or exact profiles, and scheduling with varying degrees of user intervention. Techniques used within MAESTRO to provide this flexibility and power include constraint propagation mechanisms, multiple asynchronous processes, prioritized transaction-based command management, resource opportunity calculation, user-alterable selection and placement mechanisms, and maintenance of multiple schedules and resource profiles. These techniques and the scheduling complexities requiring them will be discussed in this paper.

INTRODUCTION

As the complexity and sophistication of spacecraft and the experiments they carry increase, the cost of operating them increases as well. It is imperative that these spacecraft be operated as efficiently as possible. This will require significant changes in the way spacecraft are managed, including more sophisticated scheduling techniques. Assumptions made in the past to simplify the scheduling problem will no longer be supportable. A system which controls spacecraft must be capable of evolving to meet demands for more payload intelligence and autonomy, more real-time user control, more complexity in the interactions possible between activities aboard one or more spacecraft, etc. In designing a scheduler for spacecraft operations a number of as yet unsolved problems arise as a product of various interactions among experiment and systems requirements, constraints on ground and spacecraft systems capabilities, and so on. The degree to which these problems can

be solved will significantly affect how well spacecraft management is carried out in the 1990's and beyond.

Scheduling, as defined in this paper, consists of fixing the execution times of activities on a timeline, such that all constraints (e.g. resource requirements, environmental conditions, etc.) on these activities are met. This can be contrasted with the definition of planning, in which a set of operations is ordered such that a desired goal state is reached. A scheduler assumes the orderings for operations are fixed, and does not have the knowledge or mechanisms necessary to order them.

It often happens that a partial or completed schedule will prove to be in need of revision due to changes in mission requirements or resource or conditions availabilities. Making the required changes to a schedule, including unscheduling, is also part of the scheduling process.

The scheduling problem is extremely difficult for several reasons. The most critical factor is the computational complexity involved in developing a schedule. The size of the search space, the space of possible schedules, is large along some dimensions and infinite along others. There can be an infinite number of ways to place a single activity, and a large number of choices of crew assignments to activities, for example. Additionally the goal of the scheduling process is ill-specified - the requirement is to produce a "good" schedule, one which meets a number of often conflicting requirements. These requirements can include efficient use of resources, no time or resource constraint violations, and maximum production during a specified time period, for example. There exist many additional factors that make scheduling a difficult problem, e.g. there are interactions between particular activity placements and resource usages that make constraint violations difficult to predict and avoid.

The specific requirements of a scheduler for spacecraft in the Space Station era have not been defined, and are expected to evolve as spacecraft and instruments become more complex. Thus part of the scheduling problem is to create a system which can schedule within a number of possible operational scenarios and management approaches. The next sections discuss some of the solution methods implemented in MAESTRO.

SCHEDULING TECHNIQUES

The approach taken within MAESTRO to scheduling involves a representation of scheduling objects and operations which generate schedules based on the relevant information. (An expanded description is given in [Britt, et al 1986] and [Geoffroy, et al 1987]). Objects of the scheduler include models of the activities to be performed, and models of all relevant constraining resources and conditions. Activities within MAESTRO are modelled as ordered series of subtasks, each

of which requires a set of resources and conditions which does not vary over the duration of the subtask. The duration of each subtask, and any delays between them, can vary. There are several types of constraints which can be considered within MAESTRO. These include resources such as crew time, electrical power and pieces of equipment, consumables such as water and liquid nitrogen, and conditions such as ambient temperature, vibrational stability and spacecraft attitude. The scheduling operations within MAESTRO involve repeatedly executing a selection-placement-update cycle, in which an activity is chosen, the activity is placed on the schedule, and resource availability profiles are updated to reflect that placement.

Selection of an activity to schedule on each cycle is based on heuristics such as relative constrainedness of activities, the priority assigned to each, and the success level, defined as the ratio of performances scheduled to those requested, for each. These criteria are combined using weightings which reflect the relative importance of each of these factors. Selection is thus based on several characteristics of the activities in relation to the current partial schedule. Two of these characteristics, priority and level of success, are calculated in a straightforward manner. Relative constrainedness is a more involved measure.

Relative constrainedness of activities can loosely be defined to be the number of performances of each activity that could be placed given the current partial schedule and resource availabilities. The system first obtains a rough measure of placement opportunities. In order to obtain this rough measure the system maintains knowledge of all possible placement alternatives for every subtask in each activity considered for selection, taking into account resource requirements, subtask temporal specifications, and a number of other factors. This process, called viable intervals calculation, results in a set of time windows for each subtask during which all of the conditions for the operation of that subtask are met. These windows are pruned to take into account temporal constraints between subtasks within an activity as well, but the process achieves only a good approximation to the specification of all and only those subtask time windows which are possible.

A second process, temporal constraint propagation, based on a technique developed for scene understanding by Waltz [1975], further refines the specification of placement opportunities providing an exact measure of all possible start and end times for each of the subtasks. This function handles a variety of constraints on the start and end time points, including minimum and maximum durations of all subtasks making up an activity, delays between subtasks, duration of each performance of an activity, delays between performances of an activity, starting and/or ending time windows for activities or subtasks imposed by mission requirements, and the set of ordering relations between activities enumerated by Allen [1983] such as precedes or follows. The result of these two processes specifies all and

only those points on the timeline which are candidate start and end times for each of an activity's subtasks. These results can be used to measure constrainedness - how hard it will be to find a place on the timeline where the activity can be scheduled meeting all of its constraints.

Once an activity has been selected for scheduling there are typically a large number of times each subtask could start or end. This necessitates making use of placement heuristics appropriate to each activity, determining the placement of the activity in relation to the overall scheduling time period, maximizing or minimizing subtask durations, minimizing or maximizing delays between subtasks or between performances of the activity, and placing the activity in relation to other activities already scheduled. In determining exact placements, these placement heuristics are used in conjunction with the Waltz function described above to prune possible placements down to a unique specification of each subtask's start and end times.

Unschedulering may be required for a number of reasons - e.g. a new high priority item may need to bump some previously scheduled activities, or there may be a downward revision in projected resource availabilities. In these cases, a number of factors must be considered when deciding which performances of which activities must be taken off the schedule. Heuristics for unscheduling when constraints are violated include goodness of fit between activity resource use and magnitude of resource overbooking, base priorities of activities, dependencies between activities, other opportunities to place each activity, the ratio of performances scheduled to requested for each activity, interruptibility and restartability of each, and so on. As with selection for scheduling, these factors are combined with weightings and compared to determine which performances to unschedule.

These and other automated decision-making functions are complemented in MAESTRO with a highly interactive user interface, allowing the user to choose the level of interaction or intervention in the scheduling process that he desires.

DESIGN ISSUES

Consider the contrast between two operational scenarios - one for control of unmanned orbiting platforms with numbers of simple instruments, the other for control of experiments in a Space Station core module. In the first scenario, it is likely that control will be geographically distributed, schedule development will occur relatively close to actual schedule implementation (especially for those experiments determined by targets-of-opportunity, or recent atmospheric or political occurrences), significant on-going schedule revision will be required, and resource availabilities and requirements may be somewhat unpredictable. In the Space Station scenario, experiment planning and scheduling will be much more fixed, the

environment will typically be more predictable, and scheduling may tend to be more centralized. For these two scenarios scheduling philosophies may differ radically - e.g. resource envelopes may be used which exactly specify the resources which will be used for an activity, or may reflect a loose operational envelope in which the activity must fit; the system may host a single user or a variety of user types, each with different requirements and different levels of authority for scheduling decisions; and scheduling may occur either as a batch or an incremental process. For these different scenarios, the core scheduling problem remains the same - what differs is the implementation of the interfaces surrounding the core scheduling system.

Because these and future interfaces may differ, the MAESTRO system has been developed such that the core scheduling functions are independent of the transactions that interact with the scheduler. The scheduling core does not differentiate between interactions with a user on the host processor, a transaction log on a file, or a user utilizing a workstation in a different location. MAESTRO and its interactive displays may be implemented on a single processor or may function as the scheduling node in a larger network of computers and/or systems. Further, processes external to MAESTRO may be used to directly or indirectly enforce the appropriate philosophy. The scheduler has different selectable options for scheduling in batch or incremental mode. External processes can determine which users are allowed to perform which operations on the scheduling system. Decisions regarding how loosely or tightly resources will be assigned can be determined by the way in which activity resource requirements are modeled, and how closely the resource profiles provided to the system reflect the actual resource availabilities. This design permits the implementation of interfaces appropriate to various scheduling philosophies and viewpoints while maintaining the core capabilities of the scheduler.

CONCLUSION

Scheduling is a difficult problem. The complexity of the scheduling problem can be overcome by heuristic decision-making, temporal constraint propagation, maintenance of multiple schedules and resource availability profiles, and other techniques. The problems introduced by the considerable variability in possible operational scenarios can be vitiated by separation of the scheduler from its interfaces, use of multiple asynchronous processes, prioritized command management, and intelligent preprocessing of scheduling requests. These and other techniques are implemented in the prototype scheduling system MAESTRO. Further work is necessary to refine these techniques and make them execute more efficiently, but a solid base has been laid for scheduling in the Space Station era. MAESTRO appears to be a suitable vehicle both for future research and as a starting point for production software.

REFERENCES

1. Allen, J.F., "Maintaining knowledge about temporal intervals", Communications of the ACM, v.26 no. 11 pp.832-843, November 1983.
2. Britt, D.L., Geoffroy, A.L., and Gohring, J.R., "A Scheduling and Resource Management System for Space Applications", Proceedings of the Conference on Artificial Intelligence for Space Applications, Huntsville AL, November 1986.
3. Geoffroy, A.L., Britt, D.L., Bailey, E.A., and Gohring, J.R., "Power and Resource Management Scheduling for Scientific Space Platform Applications", Proceeding of the 22nd Intersociety Energy Conversion Engineering Conference, Philadelphia PA, August 1987.
4. Waltz, D., "Understanding line drawings of scenes with shadows", in P. Winston (Ed.), The Psychology of Computer Vision, 1975.