A Short Commentary on "Task Swapping: Making Space in Schedules for Space" by L.A. Kramer and S.F. Smith

Frank J. Diekmann

European Space Operations Centre (ESOC) / European Space Agency (ESA)
Robert-Bosch-Str.5, 64293 Darmstadt, Germany

Abstract
This commentary first summarizes some basic ideas of the paper by (Kramer and Smith, 2004). Some aspects of their article referring to limitations in their Task Swapping method are discussed, also based on the experience gained from the Envisat instrument operations scheduling. A short introduction of this highly decentralized mission planning system demonstrates the difficulty to apply the Task Swapping concept in the current Envisat system. Future developments for improved handling of user scheduling requests might however profit from the methods presented in this paper.

1 The Task-Swapping Procedure by Kramer and Smith

The scheduling of space mission activities is always constraint by the fact that available resources often do not fully satisfy all required tasks or requests from a general user community. (Kramer and Smith, 2004) describe a task swapping procedure in their paper – called MissionSwap – which allows under certain conditions the insertion of new tasks in an already oversubscribed schedule. Their remarkable concept is applicable to multi-capacity and/or multi-resource problems, enabling the shifting of contending tasks to an alternative resource or time. The insertion of a new task does not go at the expense of an already existing one. A number of retraction heuristics are introduced, which demonstrated in several experiments the feasibility to insert a high percentage of initially non-assignable tasks in an oversubscribed schedule. The results of applying the method to different scheduling problems in space domains - like AMC and EOS Fleet scheduling or the SSR resource allocation - demonstrate the high degree of adaptability of Task Swapping and the resulting schedule improvements.

The authors have shown in their outstanding paper that with a suitable task retraction heuristic and a recursive search algorithm, an optimal number of tasks can be allocated in space mission schedules, both in case of schedule generation and in schedule repair.

2 Applicability Limitations

The authors use the term "oversubscribed" in an informal way to mean scheduling problems, where by using typical scheduling algorithms, those algorithms have failed to insert all available tasks into the schedule (personal communication from L. Kramer). The attempt to repair such a schedule by inserting additional and/or conflicting tasks via Task Swapping naturally requires some free capacities, redundancies, or flexibility in task placement in the system in order to be successful. This basic prerequisite is briefly mentioned in the paper for the three main examples. It would perhaps be worthwhile to elaborate these preconditions in more detail. Questions to be asked could be, whether there is a minimum degree of flexibility for a given problem necessary for Task Swapping to work, or what the impact is of the available unused capacity (quantized) on the performance of Task Swapping (measured in number of additional tasks allocated).

The examples given by the authors predominantly refer to 'closed' scheduling systems, where in particular the mission planner is authorized to retract and reschedule specific tasks. A limitation to this system could be imposed through external scheduling constraints of a dynamic nature (in contrary to static constraints like instrument mode transition times, available power or storage capability, etc.). This includes for example the visibility of a target at only a fixed time, availability of ground stations for real time data transmissions, or non-flexible user requests. In particular the latter can hamper any otherwise possible retraction processes completely.

One of the examples given in the paper is the use of Solid State Recorders (SSR) as a possible multi-capacity resource for Task Swapping. From our experience with the two Envisat SSRs, this resource appears only of limited value for accommodating additional tasks. In case of missing data downlink opportunities for a certain period (due to 'blind' orbit cases, that is, no ground station availability) the available SSR storage capacity is simply filled with high priority data takes. In case of regular data downlinks, free SSR memory is not utilized due to instrument operational limitations or fixed user requests.

3 The ENVISAT "Regional Mission" Scheduling Problem

Envisat is currently the largest and most powerful remote sensing satellite. It provides both global and precise regional / local observations of the Earth's surface and its atmosphere. The majority of all measurements of the 9 instruments on-board contribute to the so-called Global Mission (GM) by continuous monitoring of the Earth, providing for most sensors a complete coverage of the globe within one to three days. Two instruments (the Advanced Synthetic Aperture Radar (ASAR) and Medium Resolution Imaging Spectrometer (MERIS), allow in addition with their high data rate measurement modes regionally and temporally limited observations of surface scenes. Such specific observations for these instruments forming the Regional Mission (RM) can be requested directly from the user community. In the absence of any user requests, ASAR and MERIS perform standard measurements within the Background Regional Mission (BRM). On top of these routine missions, calibration and validation activities for all instruments regularly require dedicated instrument operations. The planning of these rather different missions and ensuring a continuous and error free instrument operations is the main challenge of the Envisat mission planning system. This task is distributed over several functional entities and locations. All user related and basic background measurement tasks are handled by the Mission Control Facility (MCF) of the Payload Data Segment at ESRIN (Italy). The ultimate merging of all instrument and onboard data handling activities takes place at the Mission Planning System (MPS), which is part of the Flight Operations Segment (FOS) at ESOC (Diekmann et al., 2004).

While the low data rate instrument (GM and BRM) scheduling is a straight forward activity (all instrument activities can usually be planned with the resources available), the Regional Mission tasks are – similar as the EOS example presented by the authors – discretely constraint by the observability of the target scenes. In addition, the ASAR and MERIS scheduling is an 'open' system, in that all nominal instrument measurement tasks are requested by external users. The scheduling of these user requests works via a strict priority scheme and a 'first come – first serve' basis within the same user category and priority class until 14 days before on-board data acquisition (personal communication from S. Vazzana and M. Sansone, ESRIN, Frascati). Between 14 days and two days before data take, all scheduled user requests are protected via an increased priority, such that late requests cannot overrule already planned tasks. This is only possible through emergency requests issued under the International Charter on Space and Major Disaster. This type of request will always bump out already scheduled data takes. The same happens to lower priority tasks in case of conflict with a higher priority one. In this case, the MCF operator contacts the affected user and – if possible – alternatives are negotiated. Users are always encouraged to provide their requests with acquisition margins (moving windows), but this option is almost never considered, such that there is no scheduling flexibility left for the operator. He can only reject requests in case of instrument constraint violations (invalid modes, maximum operations time exceeded per moving window, etc.). The tasks finally scheduled usually fit to the available SSR memory (variable due to commandable partition sizes) or are transmitted in real time to ground. This scheduling set-up does therefore not allow the application of a retraction heuristic and a task swapping algorithm as presented by Kramer and Smith, in spite of potentially unused resources in the system.

4 Possible Application of Task-Swapping for Re-Scheduling of ENVISAT User Requests

In case a lower priority task is bumped out the schedule by a higher priority one, the user is contacted who gets the possibility to reschedule his task. An automatic re-scheduling algorithm could in the future significantly improve this process. A possible retraction heuristic could be based on the weighted sum of the number of re-scheduling alternatives. These include not only shortening or slight movements in time of the user requested data take, but also the potential use of a different ASAR viewing swath in the next orbit (not always possible) or a delay until the next scene overflight (normally after three days). Rescheduling of a retracted task should however avoid bumping out and rescheduling of another task, which would imply contacting yet another user. The result of this re-scheduling algorithm could then be presented to the affected user without negotiating the case with the MCF operator.

This is just one thinkable option to improve the current ASAR and MERIS scheduling scheme based on the concepts presented in the paper by Kramer and Smith. Further applications appear possible and will certainly be discussed in the future.

5 References

F.J. Diekmann, A. Rudolph and P. Bearman. ENVISAT Payload Operations Planning – the end-to-end view and experience. In *Proceedings 16th IFAC Symposium on Automatic Control in Aerospace, to appear,* St. Petersburg, Russia, 2004.

L.A. Kramer and S.F. Smith. Task Swapping: Making Space in Schedules for Space. To be presented at *4th International Workshop on Planning and Scheduling for Space (IWPSS'),* Darmstadt, Germany, 2004.