Advanced Planning and Scheduling Initiative's XMAS tool: AI for automatic scheduling of XMM-Newton long term plan

Francesco Castellini [(1)], Michèle R. Lavagna [(1)]

[(1)] Politecnico di Milano, Aerospace Engineering Department
Via La Masa 34, 20156 Milano, Italy
castellini@aero.polimi.it, lavagna@aero.polimi.it

Abstract

This paper introduces XMAS (Xmm Mission Apsi Scheduler), a tool for the automatic scheduling of XMM (X-ray Multi-Mirror) – Newton Long Term Planning (LTP), developed under the Advanced Planning and Scheduling Initiative (APSI) study for the European Space Agency. The current work has been built upon the previous experience of the development of a similar tool for the planning of the observations of another ESA mission, the INTErnational Gamma-Ray Astrophysics Laboratory (INTEGRAL). Although tailored for specific XMM's requirements, XMAS is for this reason capable of dealing with modeling features, constraints and objectives common to both missions and may therefore be regarded as a more generic tool, covering many aspects of the scheduling problems for single-instrument space observatories.
After a brief introduction on APSI study, the paper focuses on the peculiarities of XMM-Newton observations features and on the performance indexes, constraints and heuristics that have been necessary to deal with the given requirements. Results will then be presented according to two test sets, representative of typical XMM's LTP problems, and some conclusions and possible improvements will finally be outlined.

1. Introduction

The European Space Agency (ESA) is currently supporting a study named APSI (Advanced Planning and Scheduling Initiative). The study focuses on the Artificial Intelligence techniques possible exploitation to enhance the ESA missions operations management. The development of a general-purpose architecture was a fundamental requirement of the study. To this end the developed framework has been tested on three different planning scenarios, coming from current ESA missions: Mars Express, Integral, and XMM.

The consortium in charge of the study comprehends VEGA Deutschland GmbH as an industrial partner; ISTC-CNR (Rome, Italy), ONERA (Toulouse, France), and Politecnico di Milano (Milan, Italy) as both research centers and academies, each of them responsible for the development of one of the cited applications. ISTC-CNR Was also in charge of the developing the generic planning and scheduling framework, named as APSI hereinafter.
For more details on the general APSI study see [1].
The current paper discusses the design, implementation, and run of the third mission analyzed in the APSI study, that is the long-term planning for XMM observations. The XMM APSI Scheduler (XMAS) greatly leans on the AIMS developed by ONERA for Integral planning solving [2].
The paper is organized as follows: XMM mission requirements are presented in section 2; section 3 describes the modeling approach applied within the APSI framework to successfully cope with XMM observation requests allocations; results are presented in section 4; final remarks are offered in section 5.

2. XMM-Newton mission requirements

The XMM-Newton is an ESA mission launched in 1999 and configured as a space-based X-ray observatory open to the community, with typically 200 Principal Investigators (PIs) selected by an external committee (Observing Time Allocation Committee, OTAC) out of 400 submitters at each Announcement of Opportunity (AO), which covers one year of observation time. The technical and scientific aspects of the mission are covered in [3], [4], [5], [6].

Figure 1 shows the 2-days orbit of XMM-Newton, that constitutes the planning unit for the observations; the entire phase close to the perigee (within 40000 km altitude), that also includes the eclipse phase, is not well suited for science observations due to the spacecraft passage into the radiation belts; that trajectory arc is only exploited for subsystems maintenance activities. This leaves about 36 hours (~ 130 ks) of useful time for observations, which can slightly vary depending on the revolution due to the evolution of the satellite orbit in the solar-terrestrial magnetic system.

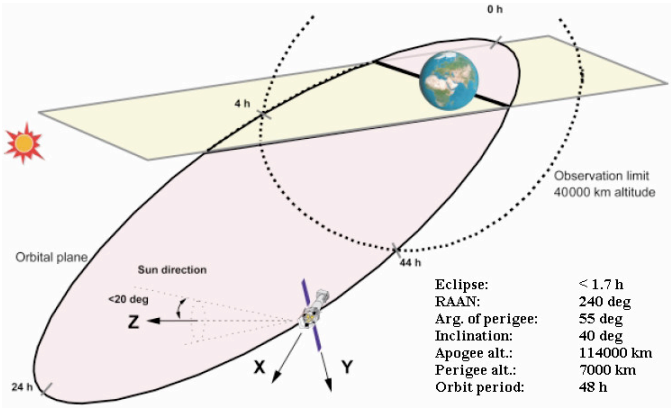

Figure 1: XMM-Newton overview

The typical observations performed by XMM-Newton last from 8 to 132 ks; three priority levels can be attached to the observation requests; constraints on observation allocation in time come from high radiation zone avoidance, as mentioned, and by the so called sky constraints, imposed by the aim of avoiding the Sun, Earth, Moon, planets and major asteroids presence in the current field of view; observation requests may also ask for a *"fixed time"* to occur: in other words, observation is requested to happen at a given Universal Time (UT) or to be phased with some periodic behavior such as for pulsars [7] or other stellar phenomena; some temporal relations according to other observations may be requested, with tolerances ranging from a few hours to months: section 3.5 is going to better clarify those kind of constraints, the so called *"chain observations"* .

The official approach - followed by the XMM team to allocate requests in time - foresees two scheduling phases: first, for each AO and thus on a yearly basis, a Long Term Planning (LTP) is developed taking into account only higher priority observations approved by OTAC and filling no more than 50-60% of the total available time; then, a one-month Short Term Planning (STP) period is scheduled to include all other observations. The relatively low percentage of usable time, denoted here as Filling Factor (FF), defined for the LTP is explained by its required robustness: the LTP has to include sufficient free time for the following tasks, to be scheduled within the STP:

re-allocation of observations failed due to possible contingencies: subsystems failures and maintenance, instruments recalibrations necessary after unpredictable strong solar flares, ground segment outages;

addition to the plan of "triggered observations", observations that have to be performed in unpredictable time intervals, as soon as a triggering event has been detected;

addition to the plan of Target Of Opportunities (TOO), unanticipated observations proposed by astronomers during the year covered by the LTP but deemed important enough by the OTAC to temporarily interrupt the XMM-Newton program;

allocation of lower priority observations, that are used as "fillers" during the development of the STP.

For a more thorough description of triggered observations, TOO and other details on the available observations types, the reader should refer to "XMM-Newton Announcement of Opportunities: Policies and Procedures", available at [8].

The goal within APSI study was to develop a new algorithm able to automatically generate an optimized schedule for the LTP of XMM mission, leaving the STP observations out of the loop. In this context, the following simplifying assumptions are applied:

although XMM-Newton experiments make use of three scientific instruments, all these payloads are considered to run in parallel, so that only a high level observation task is considered and the problem reduces to the scheduling of the activities of a single instrument;

the slew time required for targeting is neglected;

no limitations are introduced due to battery or memory resource constraints.

These assumptions need then to be removed for the development of the STP, which is performed manually by XMM team and is not considered within the scope of the APSI study. The following requirements are instead imposed by XMM observations on the LTP process:

only observations with the highest priority level are included in the LTP, therefore all observations have the same priority level;

all observations must be performed in a single block, without being split in sub-observations;

more critical observations, intended as those which have smaller visibility windows with respect to their duration, must be scheduled first;

a maximum filling factor of 50% cannot be exceeded, to be computed in blocks of 4-5 revolutions (simply enforcing such FF in each revolution would immediately exclude the longer observations from being approved);

the final goal of the LTP is to include as many of the LTP observations as possible, and at the same time favor longer observations with respect to shorter ones, that are easier to be rescheduled as fillers in the STP;

improve the plan robustness, all observations should be allocated as early as possible;

since scientific quality of data and instruments safety are affected by the radiation belts, all observations should be scheduled as far as possible from the perigee passage;

"chain observations" constraints: for XMM-Newton, particular observations composed of series of sub-observations are also often foreseen, called, hereinafter,

"chain" observations; for such an observation, the whole series can be moved back and forth in the LTP, but predefined time gaps must be maintained between the sub-observations; these temporal constraints include uneven periodicities and tolerances, as well as different durations of the sub-observations, thus increasing their formalization complexity.

All the requirements stated above have been translated into modeling features of the XMAS tool, which are described in the next section.

3. Modeling XMM mission features

In this section, the models implemented within XMAS to deal with the scheduling problem of single-instrument space observatories are presented; first, a brief description of the basic algorithm developed for the INTEGRAL mission is reported, followed by a detailed analysis of the additional features specifically tailored for XMM-Newton requirements.

3.1. AIMS: APSI Integral Mission Scheduler

As already mentioned, XMAS has been built upon available algorithms developed for the INTEGRAL mission case of APSI study and integrated in a previous tool named AIMS (APSI Integral Mission Scheduler). While a thorough description of the goals, models, algorithms and results achieved with AIMS is available in [2], a brief overview is reported here for better comprehension; note that all features described for AIMS are also included in XMAS, that can be therefore considered as an extension for more generic problems, including also XMM-Newton specific issues.

AIMS uses a local search algorithm that employs stochastic hill-climbing, stochastic heuristics, taboo lists and restarts to generate an optimized schedule over a set of given revolutions and for a set of given observations defined by means of type, duration, priority level, visibility windows and other parameters. In particular, it exploits the scheduling capabilities of the core APSI framework developed in the context of the APSI project [9].

At the beginning of the scheduling process, one or more observation activities are initialized with null duration for each visibility window of each observation. Starting from an initial empty schedule, the algorithm then chooses at each step one observation activity to be added, thanks to stochastic heuristics described in Subsection 3.6; if the observation cannot be included due to overlapping with other already allocated activities, the algorithm selects an activity to remove. We call *local move* the execution of such addition (and, if needed, removal) of observation activities, so that one local move is performed at each step of the algorithm, leading to a modification in the cost function exploited to define the optimality of the produced schedule.

If the computation time is high enough, the evolutions observed contain successive peaks, each one corresponding to a new search from an initially empty schedule, that occurs when a restart condition based on the number of iterations without improvements is matched. To allow escaping from local optima, modifications of the current schedule leading to a decrease in the criterion value are also accepted with a given probability, and a taboo list is maintained to prevent from making the same local move several times.

Finally, each observation within AIMS is allowed to be broken down in several *elementary observations*, defined by an atomic duration and allocated in different regions of the time domain; this functionality has been specifically implemented for INTEGRAL mission and is not required for the modeling of XMM-Newton problem, except for the definition of chain observations of class CH in paragraph 3.5.

With reference to the list of XMM-Newton specific issues described in Section 2, the following subsections present the new features added to AIMS algorithms: a criticality index, the filling factor constraint handled in blocks of revolutions, a new multi-objective cost function, the models for chain observations and the resulting heuristics for the local moves selection.

3.2. Criticality index

Since the schedule is generated sequentially by the local search algorithm, the concept of *criticality* of an observation o can be introduced as the ratio between the time of o still not allocated and the total remaining time available in its visibility windows. Such a dynamic index, that varies as the schedule is being filled in, can be used to establish a cardinality in the selection of observations to be added to the schedule: the most critical observations have to be included first being more difficult to find room for them as the schedule crowds.

An additional capability required for XMM-Newton problem is to deal with *"fixed time"* observations, which only have a single visibility window either of the exact duration of the observation or including a given tolerance; the following criteria have been suggested by the XMM team: 1) fixed time observations should in general be allocated first, 2) in a step-by-step planning procedure, lower tolerance observations should come first and 3) in case the same time slot is requested by two fixed time observations, the more critical one should be scheduled for such slot. Note that these criteria are implicitly implemented in the criticality index and can be satisfied by simply tuning the visibility windows: due to the limited time available for fixed time observations in the visibility windows, their criticality index is close to 1, favoring their selection over the other observations; moreover, the tolerance is directly translated in additional available time, thus observations with lower tolerance will also have higher criticality index and will be selected before the others.

3.3. Filling factor constraint

To keep generality as large as possible, two different constraints on the FF have been included in XMAS: first, as available also in AIMS, a maximum FF can be defined for each revolution, and secondly a global maximum FF can be defined for blocks of N revolutions, where N is typically equal to 4 or 5 in XMM-Newton LTP. Therefore it is possible either to impose the two kinds of constraints alone or to combine them, so that a global FF cannot be exceeded in the blocks of revolutions and at the same time the algorithm is led to avoid certain given revolutions (e.g. to bias the observations towards the beginning of each block and thus ensure a higher robustness to the plan).

3.4. Multi-objective cost function

To reflect all the qualitative objectives of the scheduling process for XMM LTP, the multi-objective cost function shown in Equation 1 has been implemented; the following three terms – to be maximized - have been formalized:

$C_{Completion/Quality}$: criterion developed for INTEGRAL and partially re-used for XMM LTP; it considers for each observation o both the completion percentage $q_C(o)$ and the quality of realization $q_{R1}(o)$, that are merged through a coefficient α; the global quality $q(o)$ for each observation of the defined set O is then weighted through the weights $W(P(o))$ according to the priority level $P(o)$. A complete description of this index is available in[2]; however, XMM LTP observations must be performed in a single-block ($q_C(o)=1$ if o has already been scheduled, $q_C(o)=0$ otherwise) and the quality of realization loses significance (thus $\alpha=1$); moreover, all observations have the same priority ($W(P(o))=1$ for any o) and the criterion is thus reduced to the maximization of the number of observations included in the schedule, which is the main goal for XMM LTP;

$C_{TotalTime}$: total scheduled time for observations, aimed at favoring longer than shorter observations allocation; note that for solutions that allow the scheduling of all foreseen observations both criteria $C_{Completion/Quality}$ and $C_{TotalTime}$ assume value *1;*

$C_{EarliestStartTime}$: aimed at the minimization of the gap between the first visibility window for a given observation o and the visibility window where it has actually been scheduled; this has the purpose of increasing the packing of observations in the beginning of the planning period, again to improve robustness; note that $N_{VW}(o)$ is the number of visibility windows defined for o and $i_{VW}(o)$ is the index of the visibility window where o has been allocated, so that $q_{R2}(o)=1$ if o is planned in its earliest possible visibility window and $q_R(o)=0$ if it is allocated in the latest.

$$
\begin{cases}
C_{Glob} = w_1 \cdot C_{Completion/Quality} + w_2 \cdot C_{TotalTime} + w_3 \cdot C_{EarliestStartTime} \\
with \;\; w_1 + w_2 + w_3 = 1 \\
C_{C_{Completion/Quality}} = \dfrac{\sum_{o \in O} W(P(o)) \cdot q(o)}{\sum_{o \in O} W(P(o))} \\
where \quad q(o) = \alpha \cdot q_C(o) + (1-\alpha) \cdot q_{R1}(o) \\
C_{TotalTime} = \dfrac{\sum_{o \in O} AllocatedTime(o)}{\sum_{o \in O} ObsDuration(o)} \\
C_{EarliestStartTime} = \dfrac{\sum_{o \in O} W(P(o)) \cdot q_{R2}(o)}{\sum_{o \in O} W(P(o))} \\
where \quad q_{R2}(o) = \dfrac{N_{VW} - i_{VW}}{N_{VW} - 1}
\end{cases}
$$

Equation 1: Multi-objective cost function

3.5. Chain observations

As already mentioned in Section 2, *"chain observations"* are a particular class of periodic observations involving unequally spaced sub-observations or sub-observations with different durations. On the contrary, AIMS only includes the modeling of observations composed of two or more sub-observations to be repeated at constant time intervals, defined in the form *(n , t)*, where *n* is the constant periodicity and *t* is the related tolerance, both expressed in terms of integer numbers of revolutions. Uneven spacing or non-uniform duration of the sub-observations cannot be included in this simple model, as well as the definition of periodicities and tolerances in fractions of days instead of integer number of revolutions.

Two new observation classes have then been introduced to overcome these limitations: *CH* (*chain*) class to consider chains of sub-observations with equal duration but uneven periodicities/tolerances and *CS* (*consecutive*) class to consider chains of sub-observations with different durations to be performed according to a given order.

CH class: through *CH* observations any sequence of sub-observations with equal duration can be represented; the first sub-observation can either be constrained to a given time interval or free to be allocated in any of its visibility windows, whereas periodicity constraints with respect to such first sub-observation can be imposed for all the others.

Both the allowed time interval for the first sub-observation and the periodicities/tolerances for the following ones can be expressed either in fractions of days

or integer number of revolutions; the former choice permits to narrow more precisely the allowed time intervals for the sub-observations, while the latter usually results in more relaxed constraints, since the observation activities can be moved along the entire time span of the revolution.

In order to deal with these complex periodicity constraints, an algorithm to determine all feasible start instants for the first sub-observation is run before the beginning of the scheduling process; this leads to the identification of a set of instants out of the first sub-observation's allowed time interval for which the periodicity constraints on all the following sub-observations can be satisfied. One of these time instants is then chosen at random when the observation is selected for addition to the plan, and its visibility windows are then modified to enforce the periodicity constraints on the following sub-observations: given the scheduled start time T_0(o) of the first sub-observation of observation *o*, the allowed time interval for i-th sub-observation becomes:

$$[T_0(o)+p_i(o)-t_i(o);\ T_0(o)+p_i(o)+t_i(o)+eodur(o)]$$

where p_i and t_i are periodicity and tolerance for i-th sub-observation and *eodur(o)* is the duration of an elementary observation of *o*.

***CS class*:** the CS class allows to represent any sequence of observations with different durations to be performed in a given order; such sequences are composed by several observations, for which the following parameters are defined in addition to the duration:

ID: can be equal either to *"FIRST"*, if the observation is the first of a sequence of CS observations, or to the identifier of another previously defined CS observation; if *ID="FIRST"* the observation is always allocated before any other observation of the same sequence, in any of its visibility windows, otherwise the observation is allocated only after the observation with the specified *ID* has already been allocated, with the further constraint specified by parameter *option*;

option = ["A","S","N"]: defines where the observation can be allocated with respect to the observation with identifier *ID*: if *option = "A"*, the observation can be allocated in any of its visibility windows but only after the execution of observation *ID*, if *option = "S"* the observation must be allocated in the same revolution of the observation ID, after its execution, and if *option = "N"* the observation must be performed in the revolution after the one of the observation *ID*.

As for *CH* observations, if a sequence of *CS* observations are to be allocated in the same revolution, as soon as the first observation is scheduled the visibility windows of the successive observations are modified so that the allowed time interval for the i-th observation of the sequence becomes:

$$[Max(T_{end}(i-1),T_{start}(i);\ T_{end}(i)]$$

and observation i will surely be scheduled after observation (i-1), therefore respecting the consequentiality constraint.

As a final note, an analysis of the time complexity associated with the increase in the number of CH or CS observations has not been performed; however, the possibility to define completely arbitrary periodicities among different sub-observations may result in a non-polynomial growth of the computational effort required to satisfy such constraints, or may even origin a situation in which constraints satisfaction for all observation chains is unfeasible.

3.6. Local moves selection heuristics

With the implementation of *CH* and *CS* observations, the introduction of more complex heuristics for the selection of the local moves has become necessary to guide more efficiently the local search algorithm.

The selection of an observation activity to be *added* is based on five different criteria combined in a *global priority* level; such criteria are internally tuned so that criterion *j+1* of the ordering reported below is considered only if there exist more than one observation activities with the same value of criterion *j*; then, the observation activity to be added is chosen randomly (if more than one exist) among those with the same value for all the criteria.

1. *Minimum completion percentage*: a non null minimum completion percentage for each observation *o* can be specified within the observation's data; as soon as an observation activity of *o* is allocated without reaching such minimum percentage (i.e. this is of course not possible in case of observations to be performed in one block but can easily occur for CH observations composed of more sub-observations), all other observation activities of *o* get the highest global priority level possible. In this way, at each step if an observation exists that has been only partially allocated, other observation activities of the same observations are added until the minimum required percentage is reached. This constraint has been introduced because for XMM LTP the partial execution of an observation is not considered valuable.
2. *Priority level*: as defined within the observation's data; this criterion is not considered for XMM LTP (all observations have the same priority).
3. *Observation type*: among observations of the same priority level, precedence is given to the observations of type *CH* defined in fractions of days, then to those defined in revolutions and to observations of type CS, finally followed by standard observations (a further classification of these observations, developed for specific issues of INTEGRAL problem, can be found in [2] and can still be used within XMAS tool, but is not useful for XMM problems and is thus not reported here).
4. *Criticality index*: as defined in Section 3.2.
5. *Multi-objective cost function*: as defined in Equation 1.

With this ordering, the algorithm first allocates already started observations with completion percentage lower than the minimum required, then allocates the others according first to the priority and then to the type, criticality and cost function value.

Similar heuristics are implemented also for the selection of the observation activities to be *removed*, with the exception that those being part of *CH* or *CS* observations are prevented from being removed, due to the hard periodicity constraints that make their reallocation very difficult.

4. Scheduling results

In this section, some of the results obtained with XMAS tool for XMM-Newton LTP problem are presented; to test the capability of the developed algorithms to deal with all the peculiar aspects of the mission, the following two test sets of observations provided by XMM team have been studied:

Test set 1: it contains 315 observations defined within the LTP of XMM's AO7 cycle (from 2008-Apr-30 to 2009-Apr-25), all to be performed without splitting and without any periodicity constraint; some observations can be considered "fixed-time" since their available visibility windows are only slightly larger than their duration.

Test set 2: it contains 263 observations still defined within AO7, of which 6 are classified as CH and 5 as CS. Additionally, another 15 observations are given with some kind of periodicity constraint that can however be treated with the standard observation types through the redefinition of the visibility windows.

4.1. Test set 1 – preliminary results

Since all observations of the first test set can also be treated with the tool originally developed for INTEGRAL, a few preliminary runs have been performed with AIMS to obtain some initial results for future comparison. The following considerations can be outlined regarding the resulting scheduling algorithm:

just considering the criterion developed within AIMS ($C_{Completion/Quality}$), with all observations having the same priority and setting α=1, the objective of the optimization stays simply the maximization of the number of scheduled observations;

all observations are of standard type and to be performed in one block; moreover, the criticality index is not included in AIMS, thus the selection of the local moves only depends on the cost function value;

the FF constraint can be imposed within AIMS only on the single revolutions, therefore to reproduce an average FF of 50% over 4 revolutions, the FF has been manually set to 100% over two revolutions and to 0% over the other two of each block; although satisfying the constraint, this approach is not optimal since the FF distribution over the revolutions of a block is arbitrarily enforced.

The results for an optimization run of 1800 s, that allow several restarts to take into account the stochastic aspects of the algorithm, are shown in Figure 2. The best achieved criterion value is *C=0.9428*, which corresponds to the successful scheduling of 298 of 315 observations, with 80.3% of the total observation time actually being allocated. Note that the former index is just given for comparison, because no reference to the total allocated time is included in the cost function.

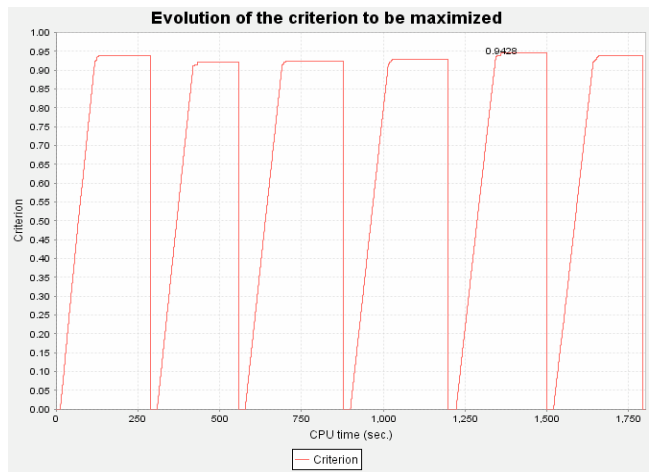

Figure 2: criterion evolution for a maximum FF=50% manually imposed in blocks of 4 revolutions.

4.2. Test set 1 – improvements

The addition of the modeling features described in subsections 3.2, 3.3, 3.4, 3.5 and 3.6 allows to remove the former limitations described above, by including the observations criticalities in the local moves selection, the maximum FF constraint handling in blocks and the complete multi-objective cost function. The results obtained with XMAS are presented sequentially and summarized in Table 1, showing the improvement in the criteria that each feature has brought to the scheduling process. Starting from the top of the table, the results presented in the previous subsection are reported for comparison, followed by those obtained first adding only the criticality index for the local moves selection, then the criticality and the maximum FF constraint in blocks of revolutions and finally the multi-objective cost function with different values of the weight parameters w_1, w_2 and w_3; note that for the first three lines the only criterion considered for the optimization is $C_{Completion/Quality}$ (corresponding, as for AIMS results, to the number of scheduled observations), with $w_1=1$ and $w_2= w_3=0$. Four columns report the average values obtained in 10 restarts for the three criteria alone and for the combined cost function, while the column reporting the successful runs percentage shows the number of times the algorithm has

reached the complete scheduling of all 315 observations over these 10 restarts.

Note that the introduction of the criticality index alone causes a strong improvement of the average criterion value, obtaining a best schedule with 311 allocated observations over 315, against 298 for the original AIMS code. Employing the FF constraint handling in blocks then allows to obtain for the first time a schedule that includes all 315 observations, even though the percentage of success is only of 20%.

Note also that all results reported in the table are referred to a maximum FF=50% imposed in blocks of 4 revolutions, as required for XMM-Newton mission. However, the success percentage for the optimization runs is affected by the number of revolutions per each block. By manually varying this parameter, it has been verified that an increase in its value brings to an increase in the success percentage, because more freedom is left to the algorithm to distribute the FFs over the revolutions with more efficient packing. For example, using a single block including the entire planning period of 181 revolutions, all performed runs have been successful in scheduling 100% of the observations; however, as shown in,*Figure 3*, the FF distribution over blocks of 4 revolutions is in this case irregular, and not acceptable in the context of XMM mission due to its insufficient robustness to possible contingencies.

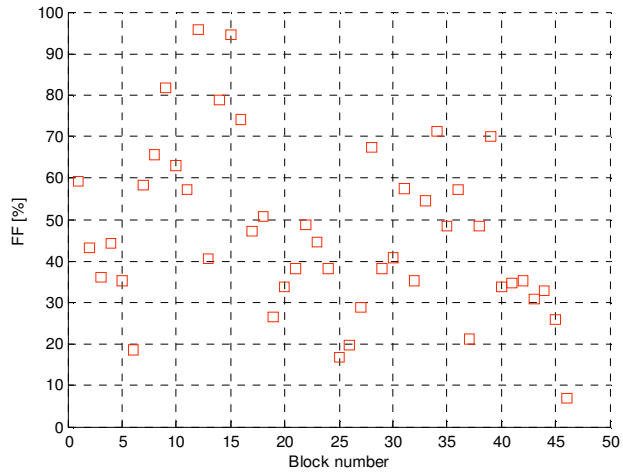

Figure 3: FFs for the 46 blocks of 4 revolutions obtained with a maximum FF=50% imposed on the entire planning period

Algorithm	w_1	w_2	w_3	Average C_{Glob}	Average $C_{Completion/Qualit}$	Successful runs %	Average $C_{TotalTime}$	Average $C_{EarliestStartTime}$
AIMS	1.00	0.00	0.00	0.9238	0.9238	0 %	0.7956	-
XMAS (criticality only)	1.00	0.00	0.00	0.9696	0.9696	0 %	0.9422	-
XMAS (criticality + FF in blocks	1.00	0.00	0.00	0.9883	0.9883	20 %	0.9897	-
XMAS (complete)	1.00	0.00	0.00	0.9980	0.9979	40 %	0.9986	0.4744
XMAS (complete)	0.50	0.50	0.00	0.9914	0.9915	0 %	0.9913	0.4607
XMAS (complete)	**0.70**	**0.30**	**0.00**	1.0000	1.0000	**100 %**	1.0000	0.4914
XMAS (complete)	**0.80**	**0.20**	**0.00**	1.0000	1.0000	**100 %**	1.0000	0.4846
XMAS (complete)	0.90	0.10	0.00	0.9980	0.9979	40 %	0.9986	0.4744
XMAS (complete)	0.50	0.40	0.10	0.9395	0.9926	0 %	0.9902	0.5125
XMAS (complete)	0.70	0.20	0.10	0.9528	0.9958	20 %	0.9903	**0.5217**
XMAS (complete)	0.70	0.25	0.05	0.9669	0.9947	0 %	0.9880	0.5012
XMAS (complete)	0.85	0.10	0.05	0.9684	0.9947	0 %	0.9920	0.4923

Table 1: Summary of results obtained for test set 1 with different versions of the scheduling algorithm and different weights tuning.

The introduction of the multi-objective cost function (thus including all XMAS capabilities) brings an enhancement of the success percentage; as expected, the results are in this case strongly dependant on the values of the weight parameters; tuning such parameters, the best behavior has been obtained for $w_1=0.7\text{-}0.8$ and $w_3=0$, which means without considering the earliest start time criterion. This can be explained by the fact that with $w_3 \neq 0$ the algorithm favors the allocation of observations that still have available time in one of the first VWs, regardless their duration, therefore preventing the efficient packing of the observations.

However, the behaviour of the algorithm with the earliest start time criterion has been verified to be correct through some runs performed with a reduced set of 200 observations, for which the algorithm is always capable of allocating all observations and the results are thus more clear; in this case, as shown Figure 4, using the earliest start time criterion alone allows to allocate most of the observations closer to the beginning of the planning period with respect to using the other criteria alone, with only few observations allocated in the last 80 revolutions.

The use of the earliest start time criterion alone is therefore in general not suggested, but it can prove to be useful in combination with the other criteria in order to increase the schedule's robustness in case of under-constrained problems.

As a final consideration, note that the computational time required for XMAS execution is comparable with that of AIMS. Figure 5 shows an example of successful scheduling run with the complete XMAS algorithm; the time required to get to convergence is slightly higher than for AIMS (~190 s against ~130 s) due to the additional operations introduced in the algorithm, but this is completely acceptable in front of the much better results achieved.

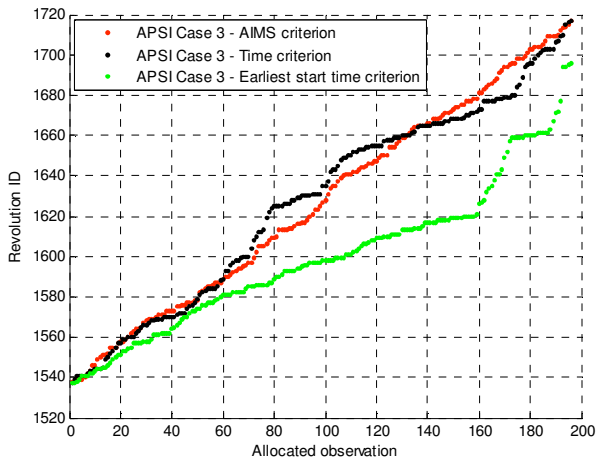

Figure 4: number of scheduled observations in the revolutions of the planning period for the three separately considered criteria

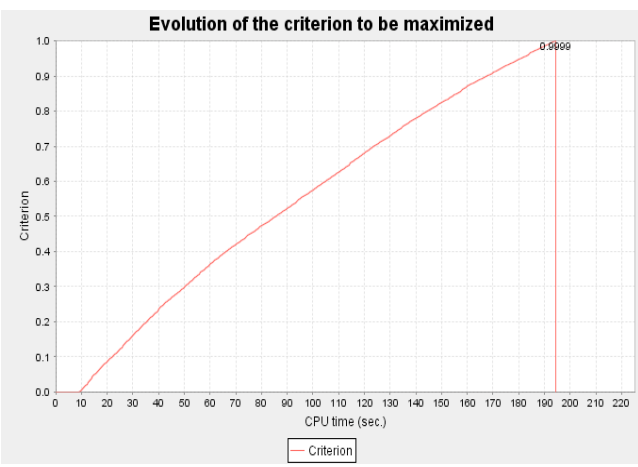

Figure 5: example of criterion evolution for a successful run with XMAS tool, comparable with that of AIMS (Figure 2).

4.3. Test set 2 – chain observations

For test set 2, a preliminary phase has consisted in the manual modification of the visibility windows for those observations that can be treated as standard provided that a reduction of the "effective" visibility windows is implemented, according to specific periodicity constraints. All runs have then been performed with a maximum *FF=50%* in blocks of 4 revolutions, $\alpha=1$, $w_1=0.8$, $w_2=0.2$ and $w_3=0$ in light of the good results obtained with these settings for test set 1.

Figure 6 shows the criterion evolution for a long run of 3600 s; the criterion value for the solutions obtained with the different restarts range between 0.9679 and 0.9898 and by going through the details of the obtained schedules it has been verified that the constraints for all the CH and CS observations are satisfied. It is however impossible to schedule 100% of the observations in this case due to the high number of long observations (i.e. ~130 ks) that alone occupy an entire revolution.

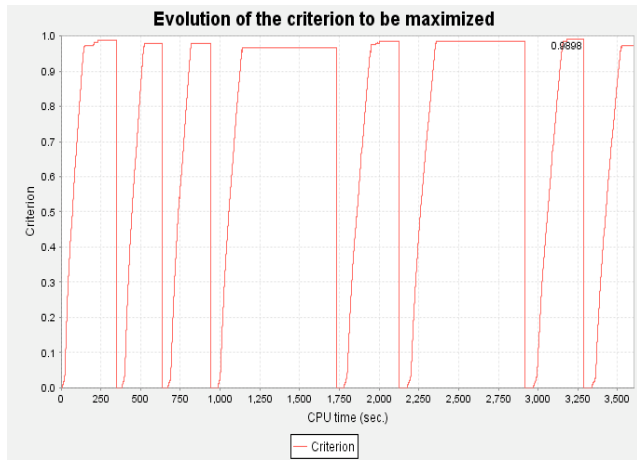

Figure 6: criterion evolution for the test set 2, FF=50% in blocks of 4 revolutions, $w_1=0.8$, $w_2=0.2$, $w_3=0$.

5. Conclusions

The tailoring of the APSI framework supported by the AIMS development, to deal with the observation requests allocations for XMM mission, gave successful results. All sets of observations turned out to be allocated thanks to the multi-objective approach and to the observation chain modeling. Moreover, XMAS software generality has been maintained by allowing the user to exploit all the functionalities of the original AIMS tool.

Possible improvements of the code refer to the further generalization of the filling factor constraint, that could be considered in blocks of variable numbers of revolutions across the planning period and the development of an additional more powerful class of observations merging the characteristics of both CH and CS classes.

More constraints, coming from the available on board resource profitles, may be included in the optimization process.

Finally, obtained results need a further validation through users' feedback on XMAS. At the time being the tool has been distributed to the XMM team to allow them time for playing with and checking for possible further improvements.

Acknowledgements

The authors thank Alessandro Donati and Nicola Policella from ESOC for their support in the APSI study and all the members of the APSI team for fruitful exchanges. Many thanks to Pedro Rodriguez and the XMM team from ESAC who supported the authors in understanding and modeling the XMM long-term planning problem supporting the APSI project application to XMM with enthusiasm.

References

1. Steel R., et al. 2009. Advanced Planning and Scheduling Initiative: AI techniques in a space domain. Submitted for acceptance at the 6th International Workshop on Planning and Scheduling for Space, IWPSS-09, July 19th-July 21st 2009, Pasadena, CA.

2. Pralet C. and Verfaillie G. 2009. AIMS: A Tool for Long-term Planning of the ESA INTEGRAL Mission, submitted for acceptance at the 6th International Workshop on Planning and Scheduling for Space, IWPSS-09, July 19th-July 21st 2009, Pasadena, CA

3. Jansen F.A. 1999. XMM: Advancing Science with the High-Throughput X-Ray Spectroscopy Mission. *ESA Bulletin 100*, December 1999.

4. Barré H., Nye Y., Janin G. 1999. An Overview of the XMM Observatory System. *ESA Bulletin 100*, December 1999.

5. Bagnasco G., Eggel H., Giannini F. and Thürey S. 1999. The Scientific Instruments On-board XMM - Technical Highlights. *ESA Bulletin 100*, December 1999.

6. De Chambure D., Lainé R., Van Katwijk K. and Kletzkine P. 1999. XMM's X-Ray Telescopes. *ESA Bulletin 100*, December 1999.

7. Koh D.T, Bildsten L., Chakrabarty D. et al. 1997. Rapid Spin-Up Episodes In The Wind-Fed Accreting Pulsar Gx 301. *The Astrophysical Journal*, 497: 933-947, April 1997.

8. http://xmm.esac.esa.int/external/xmm_user_support/documentation/ao7_policy/ann_ao7_02.html

9. Cesta, A., Fratini, S., Oddi, A., Cortellessa, G.,. 2009, MrSpock: A Tool for Long-term Planning of the ESA Mars Express Mission, submitted for acceptance at the 6th International Workshop on Planning and Scheduling for Space, IWPSS-09, July 19th-July 21st 2009, Pasadena, CA.