

Planning and scheduling algorithms for the COSMO-SkyMed constellation

Nicola Bianchessi *; Giovanni Righini †

Università degli Studi di Milano
Dipartimento di Tecnologie dell'Informazione
Via Bramante 65, 26013 Crema, Italy

October 2, 2007

Abstract

The COSMO-SkyMed satellite constellation for the observation of the Earth is made of four satellites equipped with radar instruments and is intended for dual use, i.e. for security as well as for environmental monitoring purpose. The planning and scheduling problem for the COSMO-SkyMed constellation consists in selecting and synchronizing the operations needed to acquire the requested images of the Earth surface with the operations needed to transmit the image files to a set of ground stations. The planning and scheduling problem is solved every day on different time horizons and it must comply with a large number of technical and managerial constraints, such as set-up manoeuvres, use of shared resources on board, memory capacity, operational profiles, user quotas, split requests, deadlines and priorities. The objective is to maximize the number of images taken and transmitted. We developed a constructive algorithm with look-ahead and back-tracking capabilities, that produces feasible plans in a very short time, achieving the target of 1800 images per day.

Keywords: planning and scheduling, Earth observation, decision support, optimization.

*Currently at: Università degli Studi di Brescia, Facoltà di Economia, Dipartimento di Metodi Quantitativi.

†Corresponding author: E-mail: righini@dti.unimi.it, Tel: +39 0373 898060, Fax: +39 0373 898010.

1 The COSMO-SkyMed constellation

The COSMO-SkyMed space program for the observation of the Earth stems from an agreement signed in 2001 between France and Italy for the development of a joint space program named ORFEO (Optical and Radar Federated Earth Observation). ORFEO is composed of the French constellation Pleiades, made of two agile satellites equipped with optical instruments, and the Italian constellation COSMO-SkyMed, made of four satellites equipped with SAR (Synthetic Aperture Radar) instruments. The first COSMO-SkyMed satellite has been launched on June 8th, 2007 and the full constellation will be operational by the end of 2008; the forecasted life of the COSMO-SkyMed constellation is about fifteen years.

The four satellites of the COSMO-SkyMed constellation will follow a heliosynchronous low orbit (about 700 km above the Earth surface) around the Earth passing close to the poles; the satellites will be phased at 90 degrees from one another and each of them will run 14.8125 orbits a day, that is 137 orbits every sixteen days. Hence each satellite will repeat the same ground track every sixteen days, but the same ground track will be repeated by a different satellite every four days. Figure 1 shows the ground tracks of the satellites on the Mediterranean area.

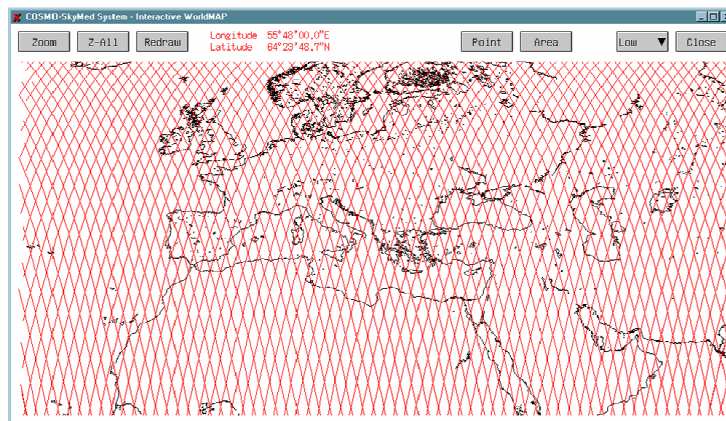


Figure 1: The ground tracks of the satellites cover the whole Earth surface. Each satellite alternates ascending tracks from South-East to North-West and descending tracks from North-East to South-West.

COSMO-SkyMed is a constellation with unique characteristics: its four satellites will allow frequent repeated observations of a same site and the SAR technology guarantees that very high resolution images will be taken in any condition of illumination and weather. The COSMO-SkyMed constellation is intended for dual use, military and civil. Besides being applied to national security purposes such as borders surveillance, the constellation will be also used for cartography and for monitoring the evolution of natural phenomena such as volcanic

eruptions, fires, floods and oil spills in the sea in any part of the world, as well as for providing detailed and timely data on regions hit by an earthquake or a tsunami, where traditional data collection and communication systems are likely to be disrupted.

Many stakeholders, both private and public companies and institutions, are involved in the COSMO-SkyMed space program. ASI, the Italian Space Agency, is the program leader; Alcatel Alenia Space, is responsible for the satellites, while TeleSpazio is responsible for the ground segment that is composed of a main coordination center (CPCM - *Centro Programmazione e Controllo Missione*) and several ground stations communicating with the constellation both for uploading the operational commands to the satellites and for receiving back the image files from them.

We became involved in the COSMO-SkyMed program in May 2003 through a subcontractor of TeleSpazio, namely Space Software Italia SpA, that was in charge for the development of the planning and scheduling algorithms.

2 The planning and scheduling problem

The planning and scheduling problem of the COSMO-SkyMed constellation is made difficult by a lot of particular characteristics of the system, of both technological and managerial nature, related to the selection of images to be taken and the synchronization between acquisitions and transmission to ground stations.

Because the satellites have numerous degrees of freedom, they have several opportunities a day to satisfy each observation request: these are called DTOs (Data Take Opportunities). Figure 2 shows a set of DTOs related to a target in the Adriatic sea.

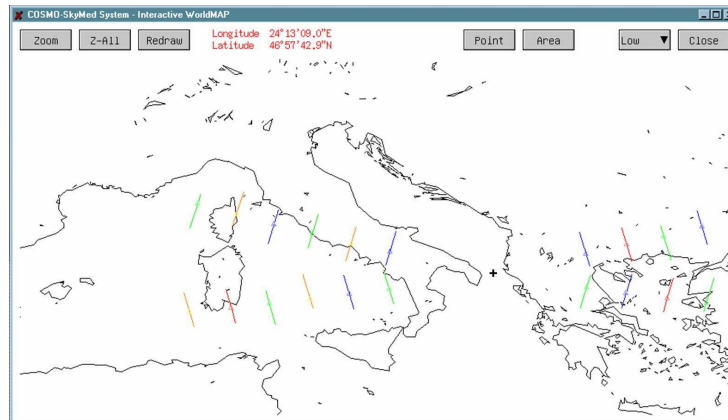


Figure 2: The target indicated by the black cross can be observed by different satellites from different tracks in different time instants. The figure shows twenty opportunities (DTOs) to acquire the image of the target.

For similar reasons every day each satellite has several time windows of variable duration available for transmitting data to ground stations (DLOs, Down-Link Opportunities). Figure 3 shows the DLOs corresponding to a ground station located in Italy. The visibility windows (DLOs) may have different width, depending on the distance between the ground station and the satellite ground track.

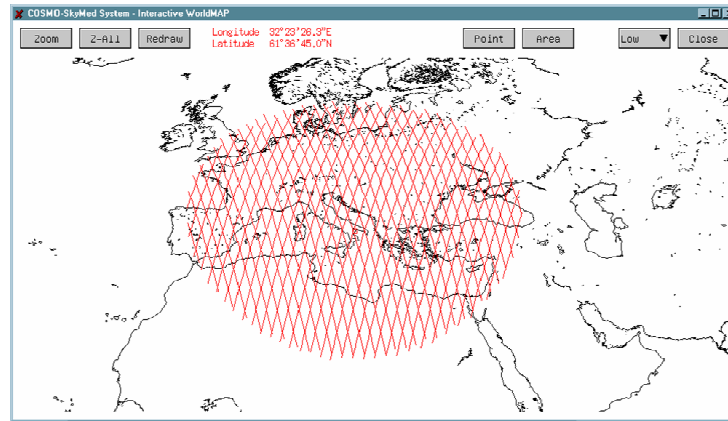


Figure 3: The picture shows many different visibility time windows for a ground station located in Italy. Ground stations located near the poles have larger and more frequent visibility windows (DLOs) because the satellites pass close to the poles after every revolution around the Earth.

Planning horizons. There are three kinds of plans to be computed: a *long-term plan* on a time horizon of sixteen days, that only includes civil requests; a *medium-term plan* on a time horizon of four days, that also includes military requests; finally, a *short-term plan* on a time horizon of one day, that is made by the first day of the most recent medium-term plan, where all time instants for the execution of each operation on board are tuned according to the current orbital parameters of each satellite.

A typical problem instance of the COSMO-SkyMed planning and scheduling problem with a 16 days time horizon may easily involve some millions DTOs and DLOs.

We were asked to provide an algorithm for the long-term and medium-term planning and scheduling problem and we were given some requirements on the computational time available: about fifty minutes for a medium-term plan and 2.5 hours for a long-term plan.

Literature review. The scientific literature on planning and scheduling problems for space is rather large but most contributions consider single satellite problems, short planning horizons and do not take into account transmission of

images to ground stations, not to say many other special problem features we had to consider. Bensana et al. (1999) and Vasquez and Hao (2001, 2003) considered a single-satellite planning and scheduling problem related to an optical satellite. Lemaitre et al. (2002) illustrated several algorithmic techniques for the planning and scheduling problem of the Pleiades constellation. Single-orbit single-satellite problems were also studied by Wolfe and Sorensen (2000) and by Cordeau and Laporte (2003) who solved small instances with up to 50 requests. SAR technology was considered by Harrison et al. (1999), who solved single-satellite instances with up to 50 requests and a time horizon of a few minutes. Globus et al. (2003) presented a genetic algorithm for a problem with up to two satellites on a short planning horizon, with neither memory nor transmission constraints. A more complicated model, more similar to ours, was studied by Frank et al. (2001), who developed a greedy stochastic algorithm without reporting computational results.

3 Problem description

Here we briefly review the main features of the planning and scheduling problem of the COSMO-SkyMed constellation, to put in evidence its data, decision variables, constraints and objectives.

Image acquisition. Each satellite can acquire the image of a *swath*, that is a rectangular portion of the Earth surface at the right or the left of the ground track of the satellite, as shown in Figure 4.

The length of the swath depends on the duration of the acquisition; the width of the swath depends on the SAR instrument *operating mode*, which is also related to the resolution and to the amount of data taken. The operating modes are grouped into two main categories: WIDEFIELD (WF) operating modes for large swaths and NARROWFIELD (NF) operating modes for small targets to be observed with very high resolution. For each request the desired resolution and the corresponding operating mode are specified in input; the starting time and the end time for each DTO are also given, because they depend on the flight parameters of the satellite. Therefore the swath corresponding to each DTO is completely specified in space and time; the algorithm must not solve a *scheduling* problem in this case, in the sense that there are no decision variables representing time instants, but only a *planning* problem, to decide which image must be taken and which of the many possible DTOs must be used to take it. The sets of all DTOs for each image, listed in chronological order for each satellite, are given in input to the planning and scheduling algorithm.

Set-up and manoeuvres. Each satellite can fly in right-looking as well as in left-looking *orientation* and it can acquire swaths on the right or the left of its ground track accordingly.

Depending on the distance of the target from the ground track, the satellite can also execute roll manoeuvres, choosing among a discrete set of possible *look-*




Figure 4: Satellites observe rectangular swaths. Through suitable roll manoeuvres they can also change the look-angle to acquire targets closer to or farther from their ground track.

angles, that are grouped into three sets, named “Nominal”, “Extended-low” and “Extended-high”.

Each rotation of the satellite from right-looking to left-looking orientation or vice versa and each change of operating mode or look-angle implies suitable *set-up* operations. Each set-up operation has a known duration and during set-up operations it is neither possible to take images nor to transmit them to ground stations.

Requests. The image acquisition requests can be submitted by different users at any time to the CPCM: they are grouped into daily batches and the planning and scheduling algorithm is re-executed daily. While the long-term plan has only provisional value, the medium-term plan is the most important in the sense that the plan for the first of the next four days is translated into a detailed list of commands that is uploaded to the satellites for being executed the day after.

Requests have an associated *deadline*. If a request remains unsatisfied until its associated deadline, it is discarded. This is allowed because the system is designed to be heavily oversubscribed.

Hence at each execution of the planning and scheduling algorithm the input is made of a set of new requests, issued in the last 24 hours, plus a set of accumulated requests, issued in the previous days but not yet satisfied, minus a set of expired requests and the set of the requests satisfied in the last day.

Quotas. Different user classes (civil and military, French and Italian) are defined and each of them has an assigned *quota*, that is defined in order to fairly balance the use of the constellation between the different institutions financing the COSMO-SkyMed program.

Priorities. Each request has a *priority*, that can be “high” or “low”, corresponding to mandatory and non-mandatory requests; high priority (mandatory) requests are typically issued by the military and they do not appear in the long-term plan. Each of them must be inserted into the plan before its deadline expires.

The feasibility of the set of high priority requests is taken for granted, because it is checked in a daily pre-processing phase, also including a quota negotiation between the French and the Italian parts.

Split requests. The requests may concern large areas of the Earth surface, not observable in one shot; in general each request corresponds to a set of images. The partition of a large target area into parts can be done in several ways and choosing one of them is a difficult combinatorial problem in itself. However this was not part of our project: the set of images corresponding to a large target is given as an input datum to the planning and scheduling algorithm.

When a request is composed of several images, it does not make sense to acquire only some of them. In Bianchessi et al. (2007) the authors presented a model, related to Pleiades constellation, in which a convex value function is associated with split requests, so that the above constraint on partial fulfillment is taken into account in a “soft” way. On the contrary in the case of COSMO-SkyMed we manage it as a “hard” constraint, that is we do not consider a split request to be satisfied until the last image associated with it has been taken and transmitted.

Another problem coming from split requests is image compatibility. As explained above, the satellites can take images in different directions (i.e. on ascending or descending tracks with right-looking or left-looking orientation) and with different look-angles. These images must be fused together at the ground stations to reconstruct the whole image of the large target area. To this purpose the images must be compatible, that is they must have been taken in the same direction and with the same look-angle (but not necessarily by the same satellite).

Operational profiles. One of the most original features of the COSMO-SkyMed planning and scheduling problem concerns the operational profiles of the SAR instruments. The operational profiles are intended to represent in a compact way the constraints on the energy consumption of the SAR instruments, that impose some limitations on the operations that can be done in any limited time interval. They are imposed on two different quantities indicating the workload of each SAR instrument in a given time interval: one is the total

acquisition time T in WF operating modes and the other is the number N of images taken in NF operating modes.

A first set of constraints bounds the daily workload: the total acquisition time in WF operating modes in each day and the total number of images taken in NF operating modes in each day are bounded by two given threshold values, T_{day} and N_{day} respectively. The two constraints are independent and they must be satisfied in every time window one day large.

A second set of constraints is used to limit the imbalance of the workload between different orbits. To this purpose the workload is defined as a convex combination of acquisition time T and number of images N ; each image taken in a NF operating mode is converted into an equivalent amount of time, through a given coefficient k , so that the overall workload W can be expressed in seconds. Let us define $W_{orbit} = \frac{T_{day} + kN_{day}}{15}$: this is a lower approximation of the workload that each orbit would have if the maximum daily overall workload would be uniformly distributed over time, because the number of orbits per day is 14.8125. If the overall workload on an orbit exceeds W_{orbit} , the orbit is a *peak orbit* and there cannot be more than one peak orbit in each day. Here by “orbit” and “day” we mean any time window as large as the duration of an orbit (about 98 minutes) or a day respectively. As a consequence of this constraint, the maximum allowed workload in a peak orbit is equal to $2W_{orbit}$: if an orbit would exist with a workload greater than $2W_{orbit}$, then it would be possible to split its workload into two equal parts, belonging to two consecutive orbits; but in such case both of them would have a workload greater than W_{orbit} and hence they would be two consecutive peak orbits, which is not allowed.

These constraints can be graphically represented as in Figure 5, where each orbit corresponds to a point in a two-dimensional space according to its workload in terms of T and N ; the lower-left triangle represents normal orbits, whereas the dark region represents peak orbits.

Data storage. The images acquired by the SAR instruments are temporarily stored as files in memory devices on board of the satellites. Each satellite has a memory device made of two blocks of identical capacity. Owing to the maximum size of an image file, each image is segmented in a given number of files and all its files must be stored in the same memory block. Each image segment file is indivisible and must be transmitted to ground stations completely and without interruption.

The limited amount of memory on board induces capacity constraints that link acquisitions with transmissions.

Ground stations. The COSMO-SkyMed system can exchange data with a set of ground stations, some of them located in Italy and some located close to the poles, where the visibility windows (DLOs) are more frequent and large. DLOs can be used to download files to ground stations: the order in which files are downloaded in a DLO is not important. Different files can be transmitted without interruption, as a unique file, if it is convenient; but if needed, idle time




Figure 5: Each orbit can be represented as a point in a two-dimensional space according to its workload. The point cannot enter the dark region more than once a day and cannot remain in the dark region for more than 98 minutes.

can be planned between two consecutive transmission operations. The only scheduling constraint is on the overall duration of the transmissions which must fall within the time interval of the DLO.

Some images must be transmitted to a specified ground station, some others to a set of ground stations. In this last case we have considered two different options, identified as “AND” and “OR” transmission modes: in the former case the image must be transmitted to all stations associated with it; in the latter case to one of them. When an image has been segmented in more than one file, the corresponding files can be transmitted separately and independently and, if the transmission mode is “OR”, they can be transmitted to different ground stations.

Parallel transmission channels. All satellites are equipped with a transmission device with two antennas corresponding to two transmission channels that can be used in parallel. Some of the ground stations are also equipped with two channels that can be used in parallel. It may happen that the DLOs of two or more stations overlap: in this case the satellite can use its two channels to communicate with two ground stations simultaneously.

Each image has an associated acquisition bit-rate, that is the ratio between the size of the image and the duration of its acquisition. The satellites have a maximum bit-rate, say B_{mem} , available to store information in memory. The two channels on board, identified as Channel 1 and Channel 2, operate at the same given bit-rate, B_{transm} ; however Channel 2 shares a bus with the memory device. Hence, while transmission on Channel 1 is independent on acquisitions, transmission on Channel 2 is not. Owing to the shared bus, when Channel 2 is used for transmission, the available bit-rate to store images in memory is limited to $B_{mem} - B_{transm}$. This implies that only images requiring an acquisition bit-rate not greater than $B_{mem} - B_{transm}$ can be taken in that time interval.

Roll manoeuvres. Another constraint linking set-up and acquisition operations with transmission operations concerns the roll angle of the satellites: transmission cannot take place when the satellite is in “Extended-low” or “Extended-high” position. Therefore the planning and scheduling algorithm must decide the suitable roll manoeuvres needed to allow transmission. For instance it may be possible that two roll manoeuvres from “Extended” to “Nominal” and vice versa are needed between two acquisitions with identical look-angle, for allowing transmission in between.

Pass-through. The satellite can simultaneously acquire and transmit the same image file; this operation is called “pass-through” and it is useful, for instance, to acquire urgent images of target areas close to a ground station. Two different kinds of pass-through can take place: in “low-rate pass-through” the acquisition rate is smaller than the transmission rate; in this case the duration of the transmission in pass-through mode is larger than it would be normally, because the low acquisition rate acts as a bottleneck. In “high-rate pass-through” on the contrary, transmission is the bottleneck and the acquired data are temporarily stored in memory; in this case the memory device is used as a circular buffer which is simultaneously read at one endpoint and written at the other one.

Objective functions. The planning and scheduling algorithm takes into account two different objective functions: the standard objective consists in maximizing the number of satisfied requests per day. To this purpose the images are normalized so that the weight associated with each of them is approximately equal to the corresponding amount of data. This normalized amount of data taken was used as a performance indicator of the planning and scheduling algorithm, which was tested to be able to produce feasible daily plans with at least 1800 normalized images per day.

Under some conditions the COSMO-SkyMed system can be switched into a “very urgent” mode. In this case the objective is to provide some critical images as soon as possible to the required ground stations. Therefore the images corresponding to very urgent requests must be taken at their earliest occurrence

and transmitted at the first available opportunity. They are expected to be in a very small number, so that no conflict avoidance technique is needed.

Commands uplink. Once or twice a day each satellite of the COSMO-SkyMed constellation must receive a list of commands from a specialized ground station. This list of commands contains the detailed plan of the operations to be executed on board in the following 12 or 24 hours. Since the uplink of the commands must be done within a prescribed time window of limited width, the number of commands must be limited too. The width of the uplink time window is sufficient for uploading a plan with as many commands as normally needed for a daily plan with approximately 1800 normalized image acquisitions and the corresponding transmissions; however in certain circumstances it may happen that the flight parameters of the satellite have to be corrected and the part of the program devoted to govern the flight become larger than usual, leaving less time available to uplink the commands related to acquisition and transmission operations. In this case some acquisition and transmission operations must be deleted from the daily plan, still guaranteeing that high priority requests are satisfied.

GPS data. Each satellite of the COSMO-SkyMed constellation also carries GPS instruments, that produce a continuous stream of data which must be segmented and transmitted to the ground stations. Segmentation is done at the beginning of each DLO for each satellite and the corresponding transmission is done within the same DLO time window. The planning and scheduling algorithm must decide when the GPS data transmission is to be scheduled inside each DLO time window.

Interferometry. The COSMO-SkyMed constellation has also been designed for interferometric observations. These are done when two satellites follow the same orbit at a short distance from each other and they observe a same target. This kind of configurations are foreseen to be used during the transient phase, before all four satellites have been launched, and at the end of the constellation life, when less than four operational satellites are left. So we were required that the planning and scheduling algorithm also allow to plan interferometric observations executed by any pair of satellites.

4 The algorithm

For the intractability of the problem and the very large size of the instances we had to give up the idea of looking for provably optimal solutions via exact optimization techniques like constraint programming or mathematical programming algorithms. Then we were left with three main alternatives that were discussed with Space Software Italia SpA in the kick-off meetings of the project: (a) mathematical programming techniques, such as truncated branch-and-bound or

branch-and-price and heuristics based on dynamic programming; (b) constructive algorithms; (c) local search algorithms. During the discussions on the *pros* and *cons* of these alternatives we became aware of some additional important facts that heavily influenced the final choice.

First of all, the planning and scheduling program could not include COTS (Commercial Off The Shelf software) for military software certification reasons. Hence the use of commercial Integer Linear Programming solvers, which is common in our scientific research activity in the academic environment, was out of question. The implementation of mathematical programming algorithms remained an option but constrained to the use of open-source routines.

A second important fact was the lack of details on some data and constraints. We were explicitly told that our algorithm had to be easily modifiable in any moment to take into account new aspects of the problem, not specified in advance.

These two facts convinced us not to follow a mathematical programming approach that requires the definition of a precise model whose structure is then exploited in detail at the largest possible extent. The effectiveness of this kind of approach is likely to be affected even by apparently small changes in the model and this kind of changes had been announced to us (and actually many occurred).

A third fact was that the algorithm had to be executable on different machines with the guarantee of providing the same results, but no hypothesis could be made on the machines themselves and their operating systems. This ruled out the option of non-deterministic algorithms such as GRASP, simulated annealing and many other heuristic and meta-heuristic techniques.

So we were left with options (b) and (c) and with the limitation of deterministic algorithms in any case.

In the literature on planning and scheduling for space a number of techniques are presented: most of them are greedy or local search algorithms. We preferred to design and implement a *constructive algorithm* first, for several reasons:

- the problem has a lot of constraints of different kinds, so that it may be very time-consuming to produce better feasible solutions by simple modifications of previous feasible solutions;
- many constraints are local (for instance the incompatibility constraints between DTOs that overlap in time): this suggests a sequential decision-making process;
- a local search algorithm needs a starting feasible, and possibly good, solution; hence a constructive algorithm is needed anyway.

Therefore we designed and developed a deterministic constructive algorithm. We did not add any local search capability to it, both because of the very strict deadlines and because the results were judged so good that an additional effort to improve them was not fully justified.

The basic structure of the planning and scheduling algorithm is sketched in Algorithm 1.

Algorithm 1 Pseudo-code of the main program.

```
repeat
  SelectActiveSatellite(s)
  if DecisionTime(s) = AcquisitionTime(s) then
    // Let  $DTO(s)$  be the next DTO for satellite  $s$  //
    if Priority( $DTO(s)$ ) > low then
      Take( $DTO(s)$ )
    else
      if Feasible( $DTO(s)$ ) then
        DecisionPolicy( $DTO(s)$ )
      else
        Skip( $DTO(s)$ )
      end if
    end if
  else
    Downlink(s)
  end if
  Update
until EndOfPlanningHorizon
```

Three activities are considered for each satellite: acquisition, transmission on Channel 1 and transmission on Channel 2. For each of these activities a time instant is recorded, that is the earliest time at which the correspondent activity can start. The minimum of the three time instants is the *decision time* for the satellite. The satellite with minimum decision time is the *active satellite*. The algorithm iteratively selects the active satellite, takes a decision about the activity that can occur at its decision time, updates the state of the whole system and recomputes the decision times. The optimization sub-routines of the algorithm are indicated as **DecisionPolicy** and **Downlink** in Algorithm 1, while the **Feasible** and **Update** subroutines take into account all of the constraints. In this way we could clearly separate optimization from constraint satisfaction: new constraints as well as changes to existing constraints were taken into account by suitable modifications of **Feasible** and **Update**, without changing the rest of the code. In parallel we could devise and compare several different optimization policies to be implemented in **DecisionPolicy** and **Downlink** independently of the modifications to the constraints.

Look-ahead and back-tracking. Differently from greedy algorithms, the planning and scheduling algorithm of COSMO-SkyMed has look-ahead and back-tracking capabilities.

Look-ahead is used to identify in advance the presence of DTOs related to a high priority request. In particular every time the transmission of a low priority request on Channel 2 would be in conflict with the acquisition of a high priority request (due to the common resource) the transmission is forbidden. This is

done in the `Update` routine.

Back-tracking is used to undo previous choices whenever the available resources are not enough to satisfy a high priority or a medium priority request. Medium priority was defined to take into account split requests and interferometric requests: for instance when an image of a split request is acquired, all the other compatible images corresponding to the same request are assigned medium priority. When back-tracking is triggered in routine `Take` because of the lack of resources (memory or operational profiles) that are necessary to take a high or medium priority DTO, the plan is scanned in reverse chronological order and the last selected acquisitions with lower priority are temporarily deleted, until the available resources become sufficient. Then the plan is scanned forward again and the acquisitions temporarily deleted are re-inserted into the plan if and only if this does not conflict with the high or medium priority image that triggered the back-tracking. A similar procedure is also triggered in the routine `Downlink` because of the need of transmitting high priority image files within their deadline. A toy example is reported in Table 1 to illustrate how back-tracking operates. In the leftmost part of the table a list of the most recently planned DTOs is reported. Assume that the residual memory capacity is 10 units and a high priority DTO n.6 of size 30 units must be taken. The lack of memory resource triggers back-tracking: DTOs are examined in reverse order, from 5 to 1, and those with low priority (n.5 and n.2) are temporary deleted until the residual capacity (60 units after the deletions) is sufficient to accommodate DTO n.6. The resulting plan is shown in the central part of the table. After inserting DTO n.6 into the plan, the residual capacity is equal to 30 units. Finally the deleted DTOs are scanned again (from DTO n.2 onwards) and they are re-inserted into the plan if possible. The final plan is shown in the rightmost part of the table.

Req.	Size	Pr.	Req.	Size	Pr.	Req.	Size	Pr.
1	40	L	1	40	L	1	40	L
2	40	L	3	40	H	3	40	H
3	40	H	4	40	H	4	40	H
4	40	H	6	30	H	5	10	L
5	10	L				6	30	H

Table 1: Left: the plan before back-tracking; center: the plan after the deletions and the insertion of DTO n.6; right: the plan after the re-insertions.

Owing to the sequential decision process, the algorithm exploits the local constraints (in particular, all the conflicts between overlapping DTOs), while non-local constraints (due to split requests, high priority requests, interferometric requests, operational profiles) are satisfied through look-ahead and back-tracking. This yields a very fast heuristic algorithm, capable of handling very large instances in a computing time much shorter than the maximum allowed.

Decision policies. When the decision time of the active satellite corresponds to the starting time of a low priority DTO, the algorithm must decide whether to take the corresponding image or not (routine `DecisionPolicy`); when the decision time of the active satellite corresponds to a possible starting time of a transmission on one of the two channels, the algorithm must decide whether a file must be transmitted and which one (routine `Downlink`). We remark that the former is a planning decision, because all time data are given, while the latter is a planning and scheduling decision, because the starting and the ending time of transmissions must be determined.

We employed simple FIFO rules for `Downlink` (with different queues for GPS data and for images with different priority) and we made several experiments with different decision policies in `DecisionPolicy`, including some non-deterministic ones for the sake of comparison with deterministic policies. Here we mention some of them:

1. Always take the next DTO.
2. Take the next DTO with probability increasing with the corresponding amount of data.
3. Take the next DTO with a probability depending on the fraction of data currently stored that must be transmitted to the same ground station. This policy was devised to avoid transmission bottlenecks, by spreading transmissions in a uniform way among all ground stations.
4. Take the next DTO if and only if the number of currently stored files to be transmitted to the same ground station is not maximum, unless they are equal for all ground stations.
5. Take the next DTO if and only if the difference between the maximum and the minimum amount of stored data to be transmitted to different ground stations does not increase or remains under a certain threshold.

5 Computational tests

We defined a reference scenario to compare different decision policies. The number of DTOs for a sixteen days plan was 750,687. Then we considered alternative scenarios generated by changing one of the main parameters of the reference scenario. Finally we tested the algorithm on realistic data generated by Alcatel Alenia Space.

Reference scenario. A 16-days reference scenario was defined as follows:

1. Four satellites at 90 degrees from one another.
2. Three ground stations.
3. Each image must be transmitted to one prescribed ground station.

4. The requests arrive at the rate of 2000 per day.
5. Requests are uniformly generated on the Earth surface with latitude between 60 degrees South and 60 degrees North.
6. Deadlines are set between 1 and 3 days.
7. Uniform mix of all operating modes.

All these seven characteristics were altered, one a time, to generate seven alternative scenarios, as follows:

1. Two satellites at 180 degrees from each other.
2. Two ground stations.
3. Each image must be transmitted to two prescribed ground stations.
4. The requests arrive at the rate of 1000 per day.
5. Requests concentrated in areas of the Earth surface corresponding to continents, with no requests on the oceans.
6. Deadlines set between 4 and 6 days (this modification increased the number of DTOs to 1,681,592).
7. Only WIDEFIELD images.

The scenarios provided by Alcatel Alenia Space had different size and characteristics; the largest ones had 2000 requests per day.

Experiments have been done on a 1.60 GHz Pentium 4 computer. When using non-deterministic policies we ran the algorithm five times and we retained the best solution found. Table 2 summarizes the outcome of our experiments on the reference scenario defined above. Time is indicated in hours:minutes:seconds. We report the computational time of the algorithm and the value of the solutions (the total size of the taken images, measured in Gbits). We also indicate the number and the percentage of taken images, the average delay between request time and acquisition time, the average delay between request time and transmission time, the average delay between acquisition and downlink (aging of information on board), the average percentage level of memory occupation, the average percentage number of files stored in memory, the average percentage of time available for transmission which is actually used for transmission for each station, the average acquisition time per satellite and per orbit and the average setup time per satellite and per orbit.

The average percentage level of memory saturation is always less than 40% for each satellite. As far as the average percentage of time used for transmission is concerned, it is always less than 50% for Station 3 and always less than 24% for the other ground stations. This is a consequence of the place where ground stations are located on the Earth surface (Stations 1 and 2 are close to the North pole). Finally it must be noticed that the average time spent in acquisition per

Table 2: Comparison of different policies on the reference scenario

Decision policy	1	2	3	4	5
Comp. time (sec.)	57	289	276	55	55
Value (Gbit)	61231	61270	60896	60980	60928
Taken im. num	8927	8993	8920	9262	8960
Taken im.(%)	27.90	28.10	27.88	28.94	28.00
Access time	59:41:18	59:36:38	59:56:28	59:16:51	59:54:11
Transm. time	62:45:25	62:41:15	63:02:19	62:04:12	62:59:22
Aging	03:16:11	03:17:20	03:15:29	02:58:16	03:17:11
Mem. occ. (%)	38.47	38.75	37.62	37.08	38.31
Station 1 (%)	24.01	23.77	24.91	26.01	24.16
Station 2 (%)	21.28	21.05	21.11	23.36	20.45
Station 3 (%)	47.89	48.83	45.82	41.94	48.58
Acq. time/orbit	00:03:19	00:03:19	00:03:18	00:03:16	00:03:18
Set-up time/orbit	00:08:42	00:09:11	00:09:13	00:08:47	00:08:51

orbit is about $\frac{1}{3}$ of an orbit duration; this is partially due to the time needed to perform set-up operations among consecutive acquisitions. These results give evidence of how the critical resource of the system (in its default configuration) is neither the memory on board of each satellite, nor the transmission time available for each ground station, but rather the operational profile constraints.

As shown by the reported results, decision policy 4 allowed to plan the largest number of images. Weighing the images according to the given specifications, depending on the operating modes, the algorithm could attain the target value of 1800 images per day. Therefore the experiments on the alternative scenarios have been carried out using this decision policy. Test results on the alternative scenarios are reported in Table 3.

Table 3: Alternative scenarios

Scenario	0	1	2	3	4	5	6	7
Comp. time (sec.)	55	26	57	63	30	59	152	56
Value (Gbit)	60980	29185	39551	43790	38905	34421	48188	33816
Taken im. num	9262	4824	7168	7643	7016	6252	8789	6189
Taken im.(%)	28.94	15.07	22.40	23.88	43.85	19.54	21.47	19.34
Access time	59:16:51	61:31:21	16:53:05	46:46:40	12:30:14	20:40:11	30:24:23	18:10:33
Transm. time	62:04:12	64:18:02	21:28:17	48:12:49	15:51:37	24:13:25	33:46:34	21:43:22
Aging	02:58:16	02:55:31	04:54:55	02:56:44	03:35:06	03:48:53	03:45:15	03:52:53
Mem. occ. (%)	37.08	29.40	86.15	44.80	70.61	68.56	81.11	82.82
Files in mem. (%)	5.98	4.79	12.12	6.43	9.79	9.47	11.33	11.18
Station 1 (%)	26.01	12.99	50.75	35.50	36.84	32.91	39.67	41.30
Station 2 (%)	23.36	11.28	-	32.97	32.13	28.19	35.08	35.93
Station 3 (%)	41.94	17.20	95.57	44.43	73.53	68.67	83.04	80.56
Acq. time/orbit	00:03:16	00:01:38	00:05:04	00:07:14	00:05:47	00:05:12	00:06:55	00:06:25
Set-up time/orbit	00:08:47	00:04:17	00:04:16	00:04:36	00:05:17	00:05:40	00:05:14	00:04:45

The analysis of scenario 1 confirms that the overall acquisition and transmission capacity of the constellation is directly proportional to the number of satellites, whenever it is less than four. In scenario 2 with only two ground stations available, the amount of information obtained is 66% with respect to scenario 0 with three ground stations. Transmission time to Station 3 is more heavily exploited in percentage than transmission time to Station 1. Station 3 tends to be the bottleneck and this is a consequence of the limited time available to communicate with it. The aging of acquired images is about 5 hours, i.e. 2 hours more with respect to the reference scenario, and this leads to an increase of the average percentage level of memory occupation and of the average percentage number of files stored. In scenario 3 the necessity of double transmission for each image results in a loss in number and value of taken images of about 17% and 27%. This suggests that the transmission activity can influence the solution quality even if transmission time is not a critical resource, as shown by the percentages of its exploitation. In scenario 4, the number of requests per day is halved with respect to the scenario 0. With only 1000 requests/day the capabilities of the satellites are not fully exploited: the value and the number of taken images decrease of about 36% and 24%. In scenario 5, where the images requested are more concentrated, so that the number of conflicts is higher, there is a loss of about 44% in value and of about 33% in number of taken images with respect to the reference scenario. In scenario 6, deadlines have been enlarged so that there are approximately twice as many data-take opportunities for each image with respect to scenario 0. Nevertheless, we observe a degradation though smaller than in the previous scenario: about 21% in value and about 5% in number of taken images. In scenario 7, without NARROWFIELD images, the performances of the algorithm degrade since operational profiles are not fully exploited. The analysis of the computing time required by the algorithm on the alternative scenarios confirms the validity of the theoretical estimation of the algorithm complexity: the computing time grows linearly in the number of satellites and the number of DTOs. The dependency on the number of DLOs is not observable because it is dominated by the dependency on the DTOs that are much more numerous.

The examination of these results was not only useful to choose the decision policy to be implemented in the final version of the planning and scheduling algorithm, but also to identify system bottlenecks and to perform sensitivity analysis on certain parameters. For instance it was possible to have a quantitative estimate of the criticality of each ground station, depending on its latitude. Failures of the memory blocks as well as deterioration of the bit-rate of the transmission channels were also simulated and their consequences on the overall system performance were evaluated.

Not only technical but also managerial parameters were studied: for instance if the constraints are imposed that each image must be transmitted to two different prescribed ground stations or to at least one of them, the overall system performance is affected: this provided useful information to COSMO-SkyMed managers at the ground segment CPCM.

Another remarkable effect of this kind of analysis was the identification of which of the many constraints are really binding the overall system performance.

Remarkably the computing time required by the algorithm was never larger than five minutes and usually less than one minute. This allows the repeated execution of the planning and scheduling algorithm in the allotted time, which is useful for negotiation purposes between different user classes.

6 Conclusions

We have solved the COSMO-SkyMed planning and scheduling problem by means of a deterministic constructive algorithm with look-ahead and back-tracking capabilities. This heuristic was tested to be capable of computing plans including even more than the prescribed number of normalized images per day in much less than the maximum allowed computing time.

This work leaves a number of open opportunities for further developments.

The determination of which constraints are really binding in realistic situations can be seen as a first step to the development of mathematical programming algorithms for the COSMO-SkyMed planning and scheduling problem, in which only the “really binding” constraints are initially taken into account while the others are relaxed. This may allow to produce almost-feasible solutions at a reasonable computational cost. This in turn would provide both upper bounds on system performance and a starting point for a heuristic post-processing procedure to produce better feasible plans. Preliminary results have been presented by Bianchessi (2006). However the solutions provided by the constructive algorithm achieve the theoretical resource saturation level defined by the constellation designers. Hence it is unlikely that more sophisticated approaches can increase the number of normalized images taken by large amounts.

On the other side, the high costs of the space program (more than a thousand million Euros) and the availability of under-exploited computing time suggest to exploit the system capabilities of generating revenues at the largest possible extent. For this purpose a promising research direction is to develop optimization algorithms based on local search or constraint programming or mathematical programming to maximize the overall value of the images taken instead of their (normalized) quantity. In particular dynamic pricing of the images required by civil users can be the subject of a specific revenue management project.

Acknowledgements. We acknowledge the kind collaboration of Vincenzo Piuri and Francesco Di Tolle for bringing us into the project, Manuel Roveri, Marco Gamassi and Daniele Sana for their help in the software code development and Giovanni Laneve, Antonio Zigrino and Alessandro Pacaccio from Space Software Italia SpA for their continuous and fruitful collaboration. The observations of two anonymous referees were very useful at improving the paper.

References

- [1] Bensana E., Lemaitre M., Verfaillie G., “Earth observation satellite management”, *Constraints* 4 (1999) 293-299.
- [2] Bianchessi N., “Planning and scheduling problems for Earth observation satellites: models and algorithms”, Doctoral dissertation, Università degli Studi di Milano, Italy, 2006.
- [3] Bianchessi N., Cordeau J.-F., Desrosiers J., Laporte G., Raymond V., “A heuristic for the multi-satellite, multi-orbit and multi-user management of Earth observation satellites”, *European Journal of Operational Research* 177 (2007) 750-762.
- [4] Cordeau J.-F., Laporte G., “Maximizing the value of an Earth observation satellite orbit”, *Journal of the Operational Research Society* 56 (2005) 962-968.
- [5] Frank J., Jonsson A., Morris R., Smith D.E., “Planning and scheduling for fleets of Earth observing satellites” In Proc. 6th International Symposium on A.I., Robotics and Automation for Space, 2001.
- [6] Globus A., Crawford J, Lohn J., Pryor A., “Scheduling Earth observing satellites with evolutionary algorithms”, *International Conference on Space Mission Challenges for Information Technology (SMC-IT)*, 2003.
- [7] Harrison S.A., Price M.E., Philpott M.S., “Task scheduling for satellite based imagery”, *PLANSIG99, 18th Workshop of the U.K. Planning and Scheduling Special Interest Group*, 1999.
- [8] Lemaitre M., Verfaillie G., Jouhaud F., Lachiver J.F., Bataille N., “Selecting and scheduling observations of agile satellites”, *Aerospace Science and Technology* 6 (2002) 367-381.
- [9] Vasquez M., Hao J.-K., “A logic-constrained knapsack formulation and a tabu algorithm for the daily photograph scheduling of an Earth observation satellite”, *Journal of Computational Optimization and Applications* (2001) 137-157.
- [10] Vasquez M., Hao J.-K., “Upper bounds for the SPOT5 daily photograph scheduling problem”, *Journal of Combinatorial Optimization* 7 (2003) 87-103.
- [11] Wolfe W.J., Sorensen S.E., “Three scheduling algorithms applied to the Earth observing systems domain”, *Management Science* 46 (2000) 148-168.