CLUSTER: A CHALLENGING MISSION TO PLAN

Dr. C. R. Haddow[1], Dr. M. Merri[2] & Mr. I. M. Palmer[3]

[1] Logica UK Ltd, ESOC, Robert-Bosch-Str 5, D-64293 Darmstadt, Germany.
Fax: +49 (6151) 902092, E-mail: chaddow@esoc.esa.de
[2] ESA/ESOC, Robert-Bosch-Str 5, D-64293 Darmstadt, Germany.
Fax: +49 (6151) 903010, E-mail: mmerri@esoc.esa.de
[3] Logica UK Ltd, ESOC, Robert-Bosch-Str 5, D-64293 Darmstadt, Germany.
Fax: +49 (6151) 902092, E-mail: ipalmer@esoc.esa.de

ABSTRACT. The Cluster Mission Planning System (CMPS) is an off-line software system that allows the European Space Operations Centre (ESOC) mission operational staff to carry out advanced planning of the 4 spacecraft Cluster mission. Initially, the CMPS supports the ESOC operational staff in the process of schedule construction based on inputs from the Principal Investigators (PIs) via the Joint Science Operations Centre (JSOC) and also the operational staff. This allows resource conflicts to be identified and corrected. In checking for schedule conflicts, the CMPS has to take into account constraints imposed by the Cluster orbits (e.g. eclipses, distance from the Earth), the on-board systems (e.g. data storage, power availability) and the on-ground systems (e.g. ground station visibility, availability). The final output of the CMPS is the generation of machine and man readable schedules to command the four spacecraft and the two ESA ground stations.

1. MISSION DESCRIPTION

Cluster forms part of the Solar-Terrestrial Science Programme and is intended to investigate the small scale and/or transitory phenomena that govern and modulate the transfer of particles and energy between the solar and terrestrial environments. Specific regions of interest include the polar cusp region, the magnetopause, the bow shock and the geomagnetic tail. In order to measure the time-varying, small scale phenomena in these regions, which are crucial to the understanding of the global solar-terrestrial coupling process, the Cluster mission consists of four identical spacecraft which were to be launched into an eccentric polar orbit in June 1996. Cluster's unique four-satellite strategy, employing an adjustable tetrahedral configuration, allows an unambiguous separation of spatial and temporal scales through comparison of the data from instruments carried by the four spacecraft. Each spacecraft carries an identical complement of 11 scientific instruments.

Since the main objective of the mission is to simultaneously capture data from the instruments on all 4 spacecraft, while operating in the regions of scientific interest, it is obvious that planning such a mission raises special problems both from the operational viewpoint, and for the software tools required to support the planning. The planning process for Cluster has to take into account the platform operations, input into the CMPS by means of Operation Request files (OPRQs) generated at ESOC by the operations staff, the observation requests of the science community (co-ordinated by JSOC based at the Rutherford Appleton Laboratories (RAL) in the UK, and submitted in Observation Request files (OBRQs), the availability of on-board resources, the visibility and availability of ground stations and orbital constraints, such as eclipses and distance from the Earth.

Clearly in this case, the challenge to mission planning is the co-ordination of simultaneous operations on the four spacecraft and of the four ground stations used during mission routine phase. There are two ESA stations, Odenwald in Germany and Redu in Belgium and two NASA

Deep Space Network (DSN) stations, Goldstone in the USA and Canberra in Australia. The DSN stations are used for the Wide Band Data (WBD) experiment. The complexity of the scheduling problem is such that it is necessary that the CMPS has full control over determining the contact periods of each of the spacecraft with the appropriate ground stations, at least during the routine operations phase of the mission. In actuality the CMPS, once all resource clashes have been resolved, automatically generates the appropriate commands to put the spacecraft and ESA ground stations into the correct modes by means of “template commands” defined by the operations staff in OPRQs, this is discussed in more detail later. The interface to the DSN stations is not so sophisticated and is handled by means of FAX requests to the DSN scheduling office.

2. THE PLANNING PROBLEM

The basic planning problem addressed by the CMPS is to take the scientific observation requests submitted by the science community, combine these with platform and housekeeping operations requests generated by the operations staff at ESOC and produce a command schedule for each of the 4 spacecraft and 2 ESA ground stations which ensures that the available resources are not exceeded. In order do to this, the CMPS has to take into account the available data storage capacity of the 2 on-board solid-state recorders (SSR) and also the on-board power consumption.

Of the on-board constraints the most significant, in terms of the CMPS, is that imposed by the capacity of the SSRs. The reason for this is that during nominal operations, 2 spacecraft are assigned to each of the ground stations, which are capable of communicating with only one spacecraft at a time. Therefore, although the Cluster orbit does give long visibility periods of up to 28 hours, the contact periods for each spacecraft must be assigned in such a manner as to stay within the on-board storage constraints, to avoid loss of data. The CMPS is responsible for calculating these contact periods, and also the times during contacts when the contents of the SSRs are dumped to the ground. In doing so, the following factors have to be taken into account:

1. Telemetry Data Acquisition (TDA) mode; this basically specifies the rate at which data is generated on the spacecraft and essentially takes one of 3 values, 3972 bps in housekeeping mode, 21845 bps in normal science mode and 131072 bps in burst science mode.
2. Visibility periods of the ground station from a particular spacecraft and, in conjunction with this, the possible unavailability of a ground station due to planned periods of maintenance for example.
3. Resource availability at the ground station; each station has 4 telemetry processors (TMP) which are used to process the telemetry dumped from the spacecraft. Since the processing can take a long time to complete, depending on the quantity of data dumped, and given that there are 2 TMPs allocated to each spacecraft, the CMPS must model the usage of the TMPs when scheduling the contact periods.
4. Antenna switches; since during antenna switches there is a temporary loss of telemetry the CMPS ensures, as far as possible, that telemetry dumps occur at times where they will not be affect by antenna switches (e.g. by performing the switch outside of a contact period if possible).
5. Spacecraft range from ground station; when the separation between the spacecraft and the ground station is greater than 35000 km the link budget is such that a lower downlink bit rate has to be used (131,072 bps as opposed to 262,144 bps) when dumping SSR playback data. In order to avoid the data loss that would occur if the rate was changed part way through a telemetry dump, the CMPS schedules the dumps such that they are carried out at the highest rate possible for the given range and dump duration.

The spacecraft power usage checks performed by the CMPS are, in comparison, much simpler. Each operational and payload command, or command sequence, has a power profile associated with it, describing the power requirements with respect to time. As a result of scheduling the operational and observational requests, the individual power profiles are summed to generate a prediction of the power requirements over the period being planned.

One factor omitted in constructing the model used in the CMPS was the line capacity between the ground stations and ESOC. This was not included because, in the original baseline scenario, where data storage on the spacecraft consisted of tape recorders of 1 Megabit capacity, only one of which was used at a time, the line capacity was not a constraint. Replacement of these by SSRs of 2.25 Megabit capacity, both of which can be used at the same time, mean that the line capacity can at times be a constraint and it may prove useful to enhance the CMPS to include this in the model.

3. SYSTEM CONTEXT AND INTERFACES.

In order to help appreciate the role of the CMPS, Figure 1 below places the system into the context of the complete Cluster ground segment.

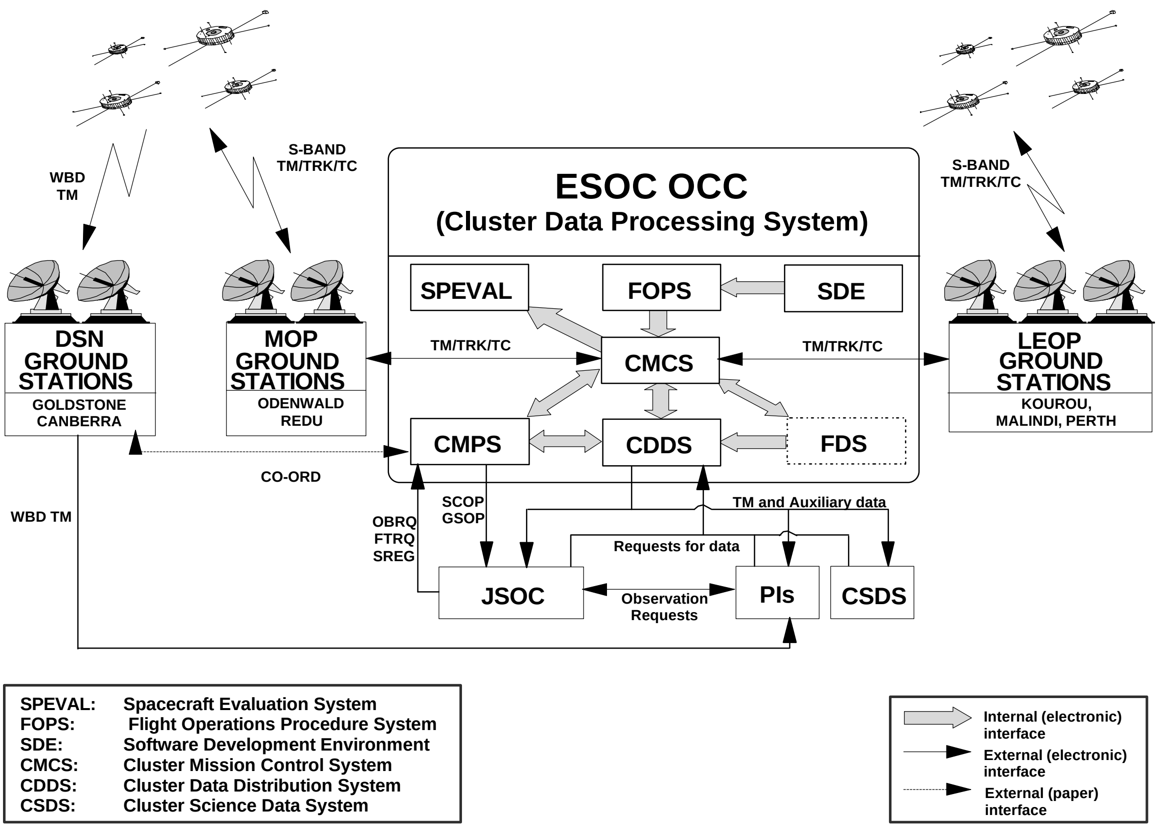

Figure 1. CMPS System Context.

It should be noted that this diagram is in itself something of a simplification. For instance the block labelled CMCS (Cluster Mission Control System) actually comprises 8 separate control systems (a dedicated prime and backup system for each spacecraft), a prime and backup network control system (NCTRS) used to control the network links between the various control systems and ground stations, and a dedicated database system (CDBS) used to maintain the telemetry and telecommand database for all 4 spacecraft. The diagram does however serve to indicate the complexity of the ground segment.

Figure 2 below gives a schematic view of the various data flows to and from the CMPS. In addition to identifying the primary inputs and outputs of the system, the diagram indicates at which stage in the planning cycle the flows occur. There are 4 main stages to the planning process, defined as long, medium, short and operational planning levels. Each data flow is annotated with an identifier signifying at which stage in the cycle the flow occurs and in which order. A more detailed discussion of the planning cycle is contained in a later section.

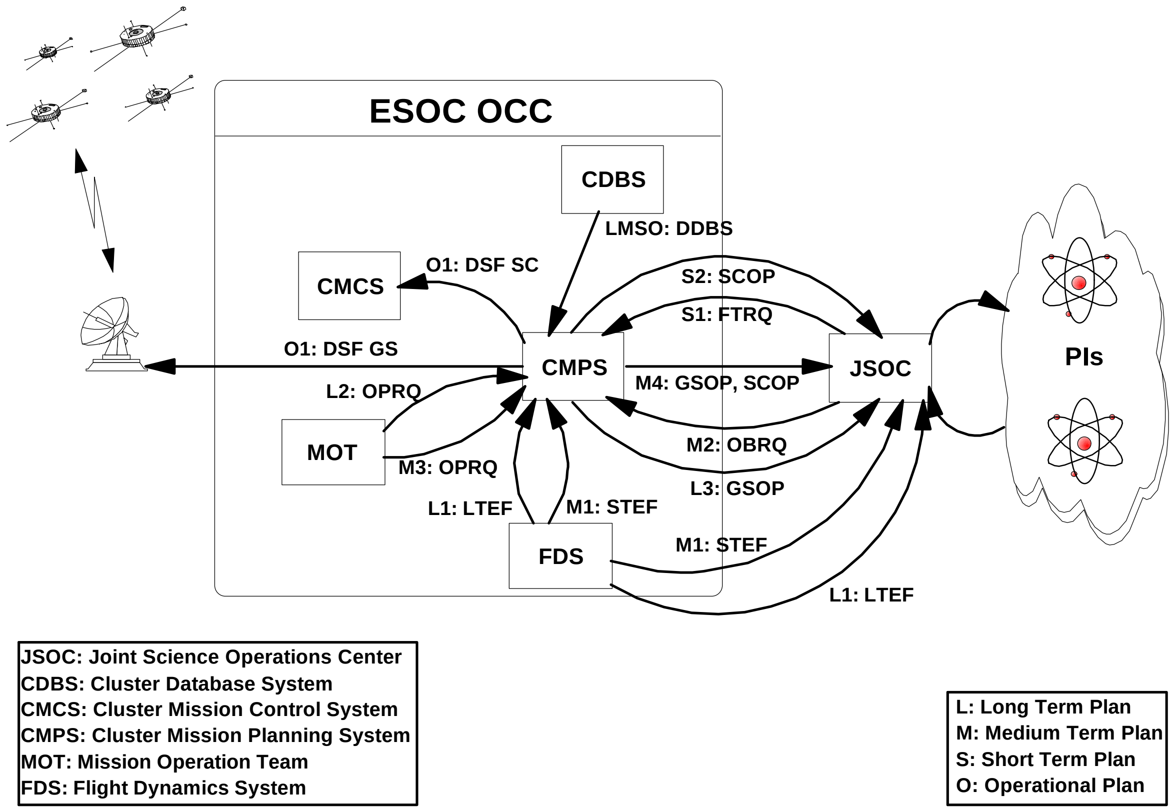

Figure 2. CMPS Data Flows.

The primary inputs to the CMPS are described below:

- STEF; Short Term Events Files are generated by the Orbit and Attitude division at ESOC and contain all orbit events that are relevant to the Cluster spacecraft such as eclipse start/end times, ground station visibility start/end times, antenna switch windows, perigee and apogee crossing, times above 35000 KM etc. These files are generated 2 to 3 times per week and contain reconstituted data (from tracking data) for the previous 10 days and predictions for the next 3.5 months.
- LTEF; Long Term Event Files are similar to the STEFs except that they are generated once for each spacecraft constellation (each spacecraft constellation lasts for approximately 6 months) and contain only a subset of the events in the STEF at a somewhat lesser degree of accuracy. These files are used in the initial phases of the planning.
- OBRQ; Observation Request Files are generated by JSOC based on input from the PIs responsible for the various instruments. These essentially contain the observation modes which the spacecraft should be in (the Cluster spacecraft have basically 3 modes for observation, the first being housekeeping where no science telemetry is acquired, the second is normal science and the third burst science where higher rate telemetry data is generated) at specified times and the commands required to put the payload instruments into the operational condition required during these times.

- FTRQ; Fine Tuning Request Files are generated by JSOC and are used to submit minor changes to a plan once it has nominally been “frozen”. In particular, they allow the insertion, modification and deletion of certain commands or sequences. Those commands/sequences which can be updated are predefined and the main criteria for allowing update of a command/sequence is that it should have no impact on resource usage.
- OPRQ; Operation Request Files are generated by the operations staff at ESOC, and contain either explicit commands to carry out required operations at specified times or “template” commands which are used to put the spacecraft and ground stations in specific modes depending on the schedule requirements. These template OPRQs are generally identified by a naming convention. The syntax of these files is fairly flexible and allows scheduling relative to a large number of events.
- DDBS; Derived Database Files are sent to the CMPS system from the CDBS. The files contain the definition of all commands and sequences including the valid range for parameters which are associated with a particular command/sequence. These are used in the validation of OBRQs received from JSOC, and also to validate commands/sequences in OPRQs generated by MOD. Additionally those command/sequences which affect onboard power usage have a time varying power profile associated with them.

The primary outputs of the CMPS are described below:

- FVSR; File Validation Status Reports are generated by the CMPS on receipt of a file from JSOC and contain a report of the validity of the files and indicate what syntax errors, if any, were contained in the file.
- SCOP; Spacecraft Operation Files are generated by the CMPS and are sent to JSOC at various stages in the planning process following a scheduling of the observation and operation requests. The contents of the file confirm the inclusion of each observation in the planning period, unless constraint violations have been detected, in which case the reason for the violation is given.
- GSOP; Ground Segment Operations Files, like the SCOP above, are generated by the CMPS and sent to JSOC. They contain schedules regarding ground segment related events, derived as a result of the scheduling process.
- SC DSF; Spacecraft Detailed Schedule Files contain a list of time tagged commands which are sent from the CMPS to the appropriate spacecraft control system for uplink to the spacecraft.
- GS DSF; Ground Station Detailed Schedule Files contain a list of time tagged “jobs” (i.e. commands) which are sent from the CMPS, via the NCTRS, to the appropriate ground station computer in order to ensure the correct station configuration.

An additional point to note about the interfaces external to ESOC is that all data transferred across the interfaces are routed via the Cluster Data Distribution System (CDDS). The reason for this is to add security to the Cluster ground segment. Since the CMPS generates command schedules which are uplinked to the spacecraft it is obviously of paramount importance that the integrity of the system can be guaranteed, and hence there is a “firewall” implemented between the CDDS and the rest of the Cluster ground segment to stop unauthorised access. The syntax checking of files received from JSOC and the “handshaking” carried out by means of the FVSR, SCOP and GSOP files provide an additional degree of security.

Implementation of the CMPS was carried out on a DEC 4000-90A VAX workstation running VMS. The language used was ‘C’, and extensive use of Motif/X-windows was made in the user interface. Graphical and timeline displays were developed using Xrt-Graph, a third-party widget supplied by the KL Group. The use of the Motif and Xrt widget libraries has resulted in a flexible user interface, which provides for easy interpretation of planning information, with facilities such as event filtering and zooming.

Additional software tools used during the development included the public domain software, FLEX and BISON which were used to generate the code that performs syntax and semantic checking of the CMPS files. The powerful syntax of these tools allowed for relatively simple coding to perform these checks, and modifications to file syntax and structure generally led to only minor changes of the FLEX and BISON scripts.

4. PLANNING PROCEDURE

The planning activity for a particular interval goes through 4 separate levels before the command schedules are finally generated. In the nominal case each planning period lasts for 3 orbits, which since the Cluster orbital period is approximately 56 hours, corresponds to closely to 1 week. Under these conditions the interchange of information between ESOC and JSOC has been agreed, to simplify operational use of the planning system, to cover pre-defined intervals. The CMPS itself can plan intervals of arbitrary length (in multiples of 1 orbit), and indeed in case of anomalies may have to. The remainder of this section outlines the procedure that is carried out in the nominal planning cycle and expands further the relative timing of the data flows introduced in Figure 2.

Long term planning; this initialises the planning process. It covers a period of about 6 months, corresponding to a Cluster “Constellation”. This plan includes the initial orbital event predictions generated by Flight Dynamics at ESOC and delivered in the LTEF. Also included are the generic Operation Requests generated by the Operations staff at ESOC. The output from this stage of the planning is the GSOP, which is sent electronically to JSOC, and this contains the preliminary list of ground segment operation.

Medium term planning; This is the level at which most of the planning process takes place. To start with a Long term plan is used to generate a medium term plan. The length of the plan generated is normally 6 orbits in duration (corresponding to 2 planning periods). It should be noted that, since each of the 4 spacecraft is in a slightly different orbit, the start time of the orbit is defined to be the ascending node crossing of an arbitrarily chosen “reference” spacecraft. At this stage the refined Flight Dynamics event predictions are included from the appropriate STEF files, as specialised platform and payload operation requests and the OBRQs from JSOC are included into the plan. Each OBRQ from JSOC normally contains observation requests covering 3 orbits. Thus a “nominal” medium term plan will contain 2 OBRQs, the first covering the 3 orbits of the plan, and second the last 3. The CMPS issues warnings to the operator if this is not the case. Activities between ESOC and JSOC start 6 weeks before the start of the period being planned. There then follows various iterations of OBRQs and SCOP/GSOPs up until one week before the planning period starts when a final OBRQ covering the initial 3 orbits of the planning period is sent to ESOC from JSOC and included. Once this OBRQ has been included into the plan, and successfully scheduled, the plan is then “frozen”.

Short term planning; once a medium term plan has been frozen it is used to generate a short term plan. This typically covers the first 3 orbits of the medium term plan from which it was generated. At this level no further scheduling is carried out. Minor changes can be carried out to the plan, known as fine tuning, which have no effect on resource usage, by the submission of FTRQs by JSOC. Any FTRQs for a particular planning period must be received at ESOC at least 72 hours before the start of that planning period. Once the deadline for receipt of the FTRQs has passed, any previously received FTRQ is included into the plan and a SCOP and GSOP generated and sent back to JSOC. On completion of short term planning activities an operational plan is generated.

Operational planning; An operational level plan normally covers the same period as the short-term plan from which it was generated. This plan is completely frozen, no changes at all being possible.

It is used for generation of the detailed schedule files for the spacecraft and ground stations. Once these files have been generated the spacecraft DSFs are automatically sent to the appropriate computers of the CMCS where they are expanded into the Cluster commands that can then be uplinked to the on-board time-tagged queue. The ground station DSFs after generation are automatically transferred to the NCTRS from where they are sent to the appropriate ground station and then executed by the station computer. Due to size limitations of the on-board time tagged queues the spacecraft DSFs are generated on a daily basis, while the ground station DSFs are generated for the complete duration of the operational plan.

A considerably more detailed of the Cluster planning procedure is given in Ref. 1.

5. SCHEDULING ALGORITHM

The CMPS scheduling algorithm, which forms the core of the CMPS, is deterministic and is driven by a number of configurable rules and constraints, the values of which can be altered to produce different planning results based on the same observation and operational request inputs. An outline of the algorithm is given below and a more expanded description can be found in Ref. 2.

The first point to note about the algorithm is that for each ground station and spacecraft combination, it loops many times over the duration of the plan in order to determine the TDA modes to be in effect for the spacecraft. The starting point for the algorithm is the plan start time, or later if a ground station contact is in effect as a result of the initialisation of this planning period from an existing plan. The algorithm then searches for possible downlink windows in which it can assign contact between a spacecraft (SC) and the ground station (GS) assigned to it. It attempts to do this in such a manner as to avoid the oversubscription of the on-board recorders which are modelled and updated continuously throughout the processing loop.

The first step in the processing is to determine which SC requires downlink the most. This is based on a number of calculations concerning the forecast usage of the recorders for each SC, from which the latest possible time that each SC can dump the recorded data to the ground station without loss of data (i.e. both recorders spilling over) is calculated. Once a SC has been identified as a candidate for having downlink with the ground station, the algorithm determines if the SC recorder usage exceeds the minimum dump threshold. If so, a recorder dump is scheduled as soon as possible, allowing for user-configurable margins. The start time and duration of the recorder dump is calculated, and the TDA mode to facilitate the recorder dump is inserted into TDA model. The provisional downlink window end-time is also evaluated, and is the latest time that another SC assigned to the GS requires a recorder dump. The next step is then to calculate an exact time for the end of the downlink window, based on a number of factors such as the usage of recorders on other spacecraft and the time of the next ground station visibility period.

When the exact downlink window end time has been calculated, the TDA modes are selected in the TDA model for the downlink SC. In addition, the Acquisition Of Signal (AOS) and Loss Of Signal (LOS) events of the chosen downlink SC are inserted into the plan. These are required to allow preparatory operations at the GS and on-board the SC to be scheduled using the generic operation request templates. After successfully selecting the real-time TDA modes for the downlink SC, the TDA modes for all other spacecraft assigned to the GS are selected, based on the fact that they must be recording for the duration of the calculated downlink window.

At this point the processing is now complete for the downlink window. The above processing is then repeated, taking the end of the previous downlink window as the window start time for the next iteration. This process continues until the entire planning period has been covered.

6. CONCLUSIONS

A planning system has been developed capable of supporting the planning activities required by the Cluster mission, arguably the most complex ESA mission to date, involving as it does the co-ordination and simultaneous control of 4 spacecraft and associated scheduling and control of ground stations.

Our experience in implementing the Cluster mission planning system reinforces what has been found before (Refs. 3 and 4), that while the basic requirements of a spacecraft control system (e.g. telecommanding, telemetry reception etc.) are readily definable in considerable detail near the beginning of a project the same is not true of mission planning. Here it is most definitely the case that a very steep learning curve must be climbed before the user requirements of the planning system can be accurately defined. Unfortunately it is usually the case that the manpower available is insufficient to allow enough time to be spent on mission planning at the user requirements phase of a mission. Problems then arise at a later stage when the software developed is found not to be completely adequate. Also adding to these difficulties are the “Gottchas”, aspects of the mission whose impacts on the planning process only become obvious late in the day. Further care must be taken in applying lessons learned from previous missions. Even missions which appear to be broadly similar turn out to have planning requirement which are substantially, if not completely, different (cf. Ref. 5). For instance in the case of the CMPS some of the initial concepts were derived from ERS, and turned out later not to be completely appropriate.

In order to overcome these difficulties it is desirable that someone on the operations staff (the end users) can spend most (ideally all) of their time working on aspects of mission planning, and that they and the software developers form a close working relationship, where ideas can be freely, and informally, exchanged. It is also desirable that the software be designed from the start to be as flexible as possible (subject to budget constraints). Experience has shown that mission planning requirements are never engraved in stone and continue to evolve and change up to, and after, launch. This has certainly been the case with Cluster, Ref. 3 also supports this for EURECA.

With respect to future developments in the mission planning area, it is interesting to note how the market in scheduling tools is developing. The scheduling algorithm developed for the CMPS is completely bespoke, since when development started, in 1993, there were few if any tools available which could realistically be used to provide a scheduling engine capable of handling all the constraints which apply to the Cluster mission. In the intervening period various products have surfaced, such as ILOG Solver, a constraints based scheduling ‘engine’, which could possibly be integrated with bespoke software to form the core of a mission planning system. It is suggested that for future planning systems some time be spent at an early stage evaluating these products to see if it is realistic to build a planning system around them. This could potentially reduce the budget required to implement mission planning systems, and also provide a higher degree of flexibility, something to be greatly desired in view of how late mission planning requirements usually stabilise.

7. ACKNOWLEDGEMENTS

The authors would like to express their thanks to the Cluster spacecraft operations manager P. Ferri, and in particular his mission planning team of B. Charrat and M. Beckers, for a fruitful, if at times lively, partnership in the development of the CMPS.

8. REFERENCES

1. The Cluster Flight Operations Plan, April 1996, Issue 4, Volume 4 (CL-ESC-PL-0001).

2. Cluster MPS Architectural Design Document (CL-ESC-RS-0311), Appendix D - Scheduling Algorithm.
3. M. De Deus Silva, D. J. Evans, "Turning Plans into Reality - The EURECA Mission Planning Experience "; EURECA Symposium, April 1994, ESTEC, Noordwijk, The Netherlands.
4. C. R. Haddow, M. Jones, E. M. Sørensen, T. Wolff, "A Review of Mission Planning Systems"; Second International Symposium on Ground Data Systems for Spacecraft Control, Paper D.1, 1992, Pasadena, U.S.A.
5. C. R. Haddow, K. O. Roeska, E. M. Sørensen, "Mission Planning Systems for Scientific and Remote Sensing Satellites: A Comparative Analysis"; First International Symposium on Ground Data Systems for Spacecraft Control, 1990, Darmstadt, Germany.