Multi-Objective Scheduling for the Cluster II Constellation

Mark D. Johnston* and Mark Giuliano**

*Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena CA USA 91109
mark.d.johnston @jpl.nasa.gov
**Space Telescope Science Institute
3700 San Martin Drive, Baltimore MD 21219
giuliano @stsci.edu

Abstract

This paper describes the application of the MUSE multi-objective scheduling framework to the Cluster II WBD scheduling domain. Cluster II is an ESA four-spacecraft constellation designed to study the plasma environment of the Earth and it's magnetosphere. One of the instruments on each of the four spacecraft is the Wide Band Data (WBD) plasma wave experiment. We have applied the MUSE evolutionary algorithm to the scheduling problem represented by this instrument, and the result has been adopted and utilized by the WBD schedulers for nearly a year. This paper describes the WBD scheduling problem, its representation in MUSE, and some of the visualization elements that provide insight into objective value tradeoffs.

1 Introduction

Multi-objective scheduling is an approach to optimized scheduling that offers a number of advantages over the more conventional single-objective approach (Deb 2001; Abraham et al. 2005). By keeping objectives separate instead of combined, more information is explicitly available to the end user or to the scheduling software system for comprehending and deciding on trade-offs among competing objectives. Multi-objective algorithms produce a set of solutions, called a *Pareto surface* (aka trade-off space), where no solution is strictly dominated by another solution for all objectives. Particularly when objectives cannot be cast to commensurate scales, visibility into the Pareto trade-off space can be extremely valuable for the decision maker. Algorithms for solving multi-objective problems have been developed that are effective in building up *populations* of candidate schedules that approximate the Pareto frontier with uniform sampling. However, infusing a multi-objective scheduling approach into an operational setting is faced with some additional challenges. Foremost among these is how the user should best select a specific schedule to execute, give a range of choices as well as visibility into the objective value tradeoffs. This is complicated by the fact that Pareto frontiers in objective spaces of dimensionality greater than three, or even two, can be difficult to visualize and grasp, and is further exacerbated by the discontinuous nature of the Pareto frontier in many real-world problems.

We have applied a multi-objective scheduling approach to several space science missions (Giuliano and Johnston 2008; Johnston and Giuliano 2010) that amply illustrate these challenges: the James Webb Space Telescope (JWST), the Cassini mission at Saturn, and the Cluster II 4-spacecraft Wideband Data (WBD) plasma experiment. In this paper we focus on the Cluster WBD scheduling problem. The scheduling software we have developed for Cluster has been routinely used operationally for nearly a year, and so a significant amount of user feedback has been obtained. In the following we first describe our overall approach and the MUSE (Multi-User Scheduling Environment, (Johnston and Giuliano 2009)) architecture on which it is based. We then describe the Cluster WBD scheduling domain, including the constraints and preferences that apply. This is followed by our adaptation of MUSE for the Cluster domain, and a description of the corresponding graphical user interface. We conclude with a summary and some directions for future research.

2 Approach

We have developed an architecture called MUSE (Multi-User Scheduling Environment) to provide a platform for developing new and integrating existing scheduling components (e.g. scheduling engines and user interfaces) into a multi-objective multi-user scheduling framework. The
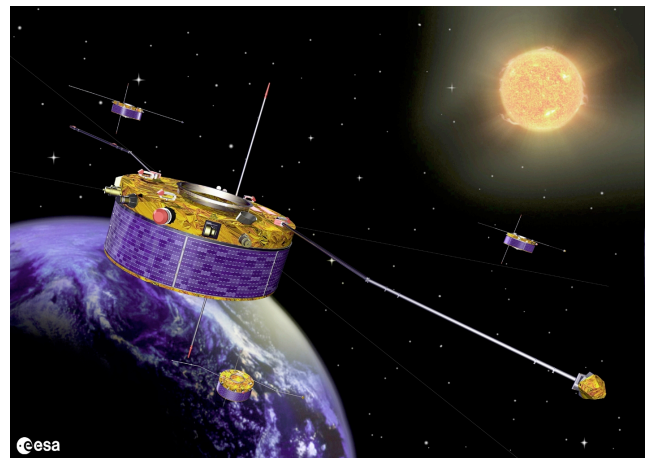

Figure 1. The Cluster II constellation of four spacecraft.

MUSE architecture integrates both generic and application-specific components. Among the generic components is a means for visualizing objective value spaces for schedule populations, for registering objective limits and acceptable ranges, and for collaborative convergence on mutually acceptable schedules for multiple users. Our approach to visualization includes a variety of techniques to meet the challenges noted above of higher-dimensional objective spaces, including 2- and 3-D projections of the Pareto frontier, histograms and other depictions of values in different dimensions, and attribute exploration techniques that have been successfully used in a number of data visualization applications. We have adapted elements common to mixed-initiative user interfaces that can be applied to our domain.

The MUSE architecture is illustrated generically in Figure 2 and as adapted to a specific mission in Figure 3. Several drivers have led to design decisions as they relate to the architecture:

- MUSE is intended to integrate with existing tools as easily as possible, to leverage existing work in many domains
- The collaborative elements of MUSE require persistent storage of various types of schedule data, hence a server-centric architecture
- Both *online* and *offline* collaboration need to be supported, in consideration of users working across multiple time zones — thus live interaction is available but not required

We distinguish server components (Figure 2 lower half) from those resident on the user's workstation. We also distinguish generic components (left) from those that are generally very domain specific (right). The architecture is designed so that domain specific components can be run as separate processes or can be compiled into the same image as the generic code.

We have adopted the familiar threaded email or news-group interaction model as a metaphor for how MUSE interacts with individual participants. Such interaction can be either on- or offline, in that one can tell upon returning to the interface what has changed since one was last present. This is important in settings where participants may use the system in an infrequent episodic manner.

On the server side, the Multi-Participant Coordinator acts as a central "clearing house" for schedule data, participant's selections, and scheduling runs. The coordinator communicates with the individual participants, providing up to date schedules, schedule status, and other participant's selections of objective value ranges. The Multi-Objective Scheduler is an implementation of an evolutionary algorithm (Deb 2001; Abraham et al. 2005) to evolve a population of candidate schedules towards the Pareto-optimal surface. While various algorithms could be employed here, we are presently using a variant called Generalized Differential Evolution (Kukkonen and Lampinen 2005; Price et al. 2005). More details about this algorithm and how it performs on some relevant domains may be found in (Johnston 2006). The Application Map provides a transformation between decision variable values and domain-specific scheduling decisions as represented and evaluated in the Domain Scheduling Engine components. The Multi-Objective Scheduler supports parallel evaluations of schedules, which can frequently help speed the generation of a Pareto surface for participants.

The Domain Scheduling Engine is the application-specific scheduling software that MUSE uses to evaluate candidate schedules. This evaluation utilizes the decision variable values, and can potentially perform internal conflict resolution or optimization steps on its own before re-
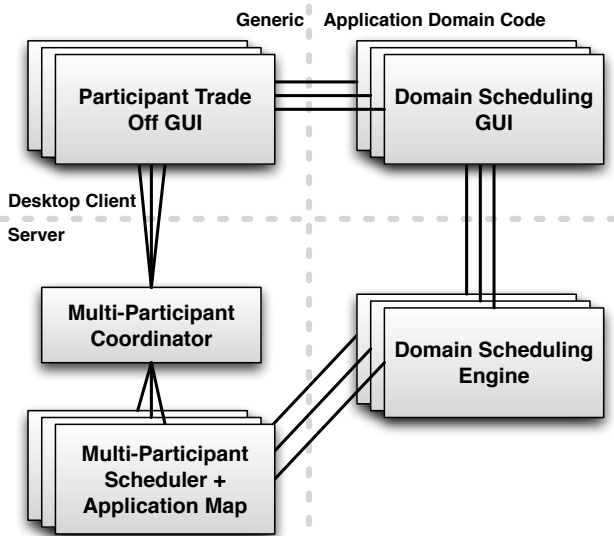
MUSE Architecture Schematic

Figure 2. Architectural overview of the Multi-User Scheduling Environment (MUSE).
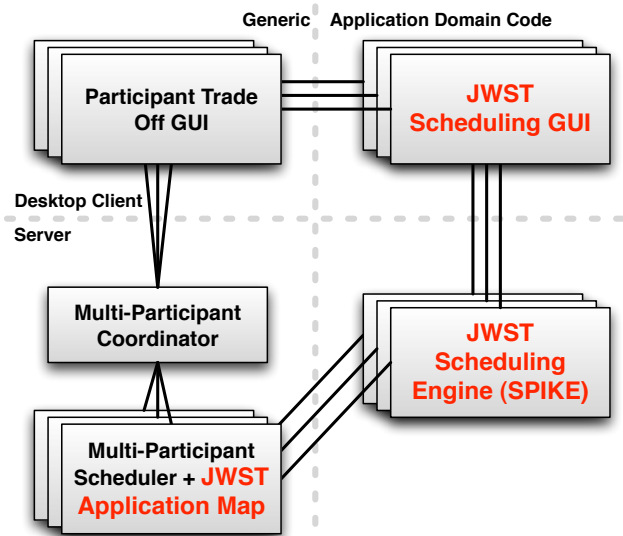
MUSE Architecture — JWST Adaptation

Figure 3. Adaptation of MUSE for a specific domain, here illustrated by James Webb Space Telescope (JWST).

turning a set of objective function values to the Multi-Objective Scheduler. These values are used by the evolutionary algorithm to evolve the candidate population towards a well-sampled Pareto surface.

Just as Domain Scheduling Engines can be highly application specific, so are Domain Scheduling GUIs. These GUIs often already exist in many domains and are able to display and manipulate aspects of the scheduling problem that are not common from one domain to another. MUSE is intended to integrate with such GUIs, e.g. to invoke the GUI on one user-selected schedule for detailed examination and assessment.

A key function of the Participant Trade-Off GUI is visualization of the objective space of the problem, in order to comprehend trade-offs and develop a solution acceptable to all participants. For 2- and 3-dimensional objective spaces, there exist commonly used techniques for visualization that can convey the selection possibilities of the candidate schedule population. However, as the dimensionality of the objective space increases, this becomes more and more challenging(Spence 2001; Tidwell 2005). We are investigating a number of techniques in this context for displaying higher dimension objective spaces, including:

- "brushed" histograms or scatter plots that indicate correlations among attributes
- display of neighbors of selected points when projected to 1- or 2-D displays
- use of multi-touch displays for rapid and intuitive manipulations of selections and views
- parallel coordinate plots

We expect that user preferences will play a crucial role in this area, and that a wide range of visualization options should be provided to accommodate the wide range of user preferences. We anticipate defining a "plug-in" mechanism so that it is easy to add additional visualization strategies as they become available.

3 Application to the Cluster II WBD Scheduling Domain

Cluster II is an ESA mission consisting of four identical spacecraft in a tetrahedral formation in polar orbit with apogees of about 20 Earth radii (Figure 1 and 4). The Cluster mission is investigating the Earth's magnetic environment and its interaction with the solar wind in three dimensions. One of the instruments on Cluster is the Wideband Data (WBD) plasma wave experiment (Gurnett et al. 1997; Gurnett et al. 2001). The WBD instrument is on each of the four Cluster spacecraft and operates by providing high-resolution measurements of the electric and magnetic fields in a range of frequency bands. There is no onboard storage for WBD at the high data rates required, so real-time data from the instrument is sent directly to NASA's Deep Space Network (DSN) antennas where it is captured and forwarded to the science team. Mathematical models of the magnetosphere are used to determine scientifically interest-
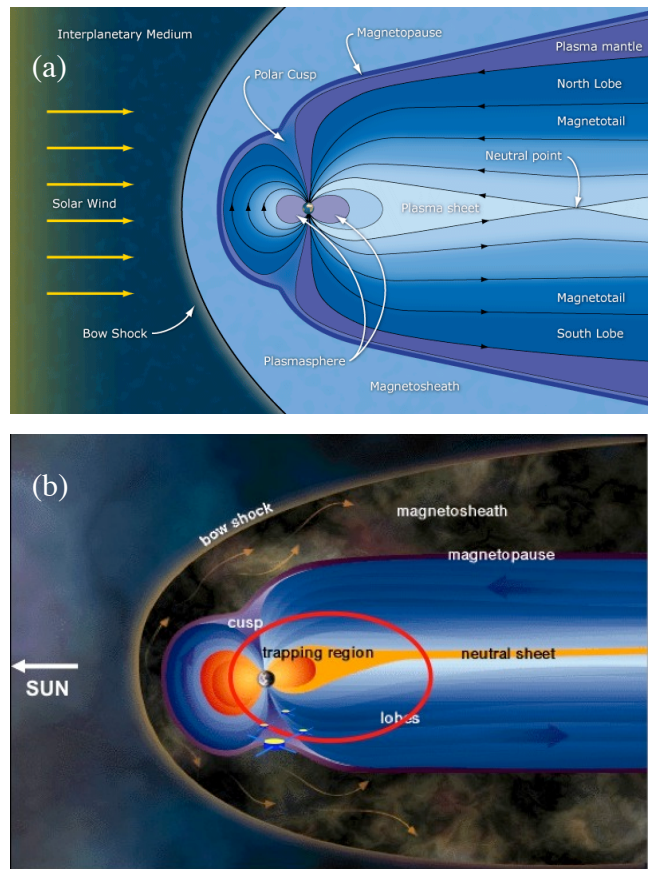

Figure 4. (a) scientifically interesting regions of the Earth's interaction with the solar wind, as influenced by its interaction with the Earth's magnetic field, and (b) an illustration of Cluster's orbit (red) in relation to the magnetosphere.

ing opportunities to collect WBD data. These opportunities represent windows during which the spacecraft passes through or in the vicinity of the solar wind and bow shock, the magnetopause, the polar cusps, the magnetotail, and the auroral zones, among others. There may be hundreds of possible opportunities during a week to observe these scientifically interesting regions of the magnetosphere. However, WBD observations with the DSN antennas are limited to about 24 hours total per week, so a relatively small selection must be made from those available.

Cluster WBD Scheduling Inputs

The WBD scheduling opportunity list is a basic input to the Cluster scheduling tool. This list is based on the Cluster orbital positions combined with models of the magnetosphere to determine interesting times for data collection. The opportunity lists includes the following criteria for each opportunity:

- **Priority**: this is one of HIGHEST, HIGH, MEDIUM, or LOW, and specifies the scientific priority of an opportunity. An important aspect of priority in the cluster setting is that it is not strictly an ordering indicator: for

Figure 5. The Cluster II Wide-Band (WBD) plasma wave receiver. One such instrument is on each of the four Cluster spacecraft.

example, the science team would not be happy with a schedule that selected only from the HIGHEST priority set; instead, they desire a mix of priorities so the data in a schedule provides coverage of a range of priority levels.

- **Science type or target**: this identifies the scientific feature to observer, for example FHSK (foreshock), BOWS (bowshock), NOAZ (Northern auroral zone), etc. There are about 20 such categories, each one assigned to one of the priority levels.

- **Start and end times, and DSN complex**: the time range during which observing must take place, and which of the three DSN complexes must be used (Goldstone, California; Madrid, Spain; or Canberra, Australia). These times range from very short (15 minutes or so) to up to about 10 hours; actual observing time for most science types is limited to a few hours, so the longer intervals become windows of flexibility.

- **Spacecraft**: each opportunity can be exploited by from one to four spacecraft, CLU1 to CLU4. Some opportunities are constrained to a minimum number, such that they should only be put into the schedule if at least 3 spacecraft are available. Others can use any number of spacecraft. A complication is that spacecraft can join and leave an opportunity over time. For example, a particular opportunity may start with only two spacecraft, then be joined by a third after a few hours, then by the fourth, after which the first two may drop out. The reason for this is that the four spacecraft are in slightly different orbits, and so spatially they enter and exit the interesting regions with offset times. They also have different visibility from the DSN ground stations, which can also affect simultaneous scheduling. A fur-

ther complication is the spacecraft are not all equally desirable to include in an opportunity so there is a preference ordering on their participation: CLU4, CLU1, CLU2 and CLU3. So, for example, an opportunity in which all four spacecraft *could* participate, but from which one must be *excluded* (say for exceeding the 24h limit), would be evaluated differently depending on which one is omitted, and this in turn depends on the time bounds of the scheduled opportunity.

In addition to the WBD opportunity list, the scheduler needs to know which antennas are available for use by Cluster, including the downtime schedule for those temporarily offline. Along with this is needed the visibility intervals of each spacecraft as seen from each potential antenna. Generally the opportunity windows fall within the antenna visibility periods, but there can be small discrepancies.

One additional important input is the available schedule of other DSN activities for the time range of interest. This is important so that Cluster can be scheduled to avoid activities that it would likely be unable to change during the schedule negotiation process, such as high priority allocations for inflexible events.

The WBD schedule for a particular time frame consists of a selection of opportunities from the WBD list that satisfy key evaluation criteria (see below) and which further specify the spacecraft and time ranges desired. The output of the Cluster scheduling tool is an XML file that can be directly loaded into the DSN scheduling system to represent these observations as scheduling requests.

Cluster WBD Scheduling Criteria

The science team's evaluation of any particular schedule depends on the following preference criteria:

- Priority mix: as noted above, a mix of priorities is required, not selection in order from HIGHEST to LOW priorities. The team has expressed a preferred ratio of priorities as HIGHEST:HIGH:MEDIUM:LOW = 4:3:2:1. In addition, the team has provided guidelines for the mix of science types to be scheduled within each priority category, such as: include "up to two instances of each science type in the HIGHEST category", with similar prescriptions for each priority level.

- Spacecraft number: the more spacecraft participating in an opportunity that permits it, the better. A four spacecraft opportunity is much better than a three spacecraft one, and so on.

- Spacecraft preferences: the order of preference of spacecraft for inclusion in multi-spacecraft opportunities is CLU4, CLU1, CLU2, CLU3.

- Spacing: it is preferred that observations be spaced out more or less evenly in the schedule, so for example a network problem that impacts a day or so of observing does not affect a disproportionate number of Cluster science observations because they were bunched up.

- Collisions with other DSN activities: it is preferable for Cluster to schedule opportunities at times where other

DSN activities do not affect the use of multiple antenna observing. These can be due to other high-priority missions that have limited flexibility, to scheduled or unscheduled maintenance activities, etc.

- Total tracking time: it is expected that Cluster will obtain about 24h of tracking time in a week, but because of changes that occur during the schedule negotiation phase, a somewhat larger amount needs to be in the initial schedule.

Because a high-value Cluster opportunity involving three or four spacecraft will frequently run into contention for use of three or four DSN antennas at one time, avoiding contentious regions of the schedule is an important strategy to avoid later schedule disruption and reduced scientific value

As a consequence of these factors, a number of tradeoffs emerge when generating a Cluster schedule. For example, schedules with many multi-spacecraft observations tend to run into contentious periods of DSN antenna oversubscription by other users, and so they are vulnerable to disruption. Yet these multi-spacecraft observations are often the most scientifically valuable. The ability to explore these and other tradeoffs in the schedule is an ideal application of the MUSE approach.

Scheduling Objectives

Based on the above criteria, we have defined four evaluation criteria that can be quantitatively computed for any given Cluster schedule, no matter how it was derived. For simplicity, each criterion is defined so that larger values mean "better", and, where reasonable, so that the scale is [0,1]. These objective criteria are:

(1) Collision: a measure of how well the scheduled Cluster opportunities avoid other DSN activities with which they are competing for time and antennas. This quantity is calculated as the average collision-free fraction of the needed antenna-hours for the opportunity. For example, if an opportunity for four spacecraft that was 2 hours in duration was in collision on one of the needed antennas for 1 hour with other activities, the collision score for that opportunity would be 0.875. The average for all opportunities is used as the collision score for the entire schedule. A value of 0 means all time is in contention, while 1 means that there are no collisions at all, and is the ideal value.

(2) Spacing: a measure of how evenly spaced in the schedule are the selected opportunities. This quantity is calculated by computing the RMS deviation d^2 from ideal spacing for the same number N of opportunities, then using $\exp(-d^2)$ as the objective value. Perfectly space observations would score an ideal value of 1.

(3) Priority: a measure of how well the actual distribution of priority values meets the science team's ideal proportion as described above. This is calculated from the RMS relative deviation p^2 of the binned opportunities relative to the desired proportions, then using $\exp(-p^2)$ as the objective value. A perfect match to the desired proportions

gives an objective value of 1, and any deviation is scored lower, with a bound of 0.

(4) Spacecraft preferences: this objective combines the preferences for multiple spacecraft along with the individual spacecraft ranking described above. Each spacecraft is assigned a weight w_i=0.80, 0.64, 0.01, 1.0 for CLU1 to CLU4, respectively, then the overall weight of an M spacecraft combination is $M(\sum_i w_i)$. For an entire schedule, the weight of each opportunity is proportional to the summed weighted opportunity minimum duration, which varies depending on the science type and the actual opportunity time range. This objective does not have a natural upperbound, but is scaled by 0.1 so that typical values are in the range [0,10] when times are measured in hours. For example, three 4-spacecraft opportunities of two hours duration (24 hours total tracking time) would score a value of 5.88.

Representation as a Multi-Objective Scheduling Problem

Representing the Cluster scheduling problem could be considered in several possible ways. One would be to define binary decision variables for each opportunity, indicating whether "in" or "out" of the schedule. These would further have to encode which are the participating spacecraft (15 additional choices for each 4-spacecraft opportunity), and the actual time scheduled to identify which of those joining or leaving to include. There are typically hundreds of opportunities per week, many with two or more spacecraft involved. This would lead to an enormous number of decision variables, and all of the preferences would have to be applied to the resulting schedules to discriminate good schedules from unacceptable ones.

Another approach at encoding would be that of an ordering to be used with an incremental schedule construction algorithm, along the lines of a Squeaky Wheel technique (Joslin and Clements 1999).

We have adopted an alternative encoding with a goal of building in as many of the scheduling preferences and guidelines as possible directly into the representation, so that they do not have to be eliminated by low objective values. This is best illustrated by an example (Figure 6). Consider the preference mentioned above, for "up to two opportunities of each science type of HIGHEST priority level". This is encoded into a problem-dependent number of decision variables, each representing a combination of priority=HIGHEST and a specific science type, with a range of 0 to 2 selected opportunities out of a total of n possible. With this encoding, one decision variable can represent any choice from the set of: no selection at all (1 choice), a single selection (n choices), and a pair of selections ($\frac{1}{2}n(n-1)$ choices). This can clearly be extended so long as the enumeration of choices does not become too great. For a typical weekly Cluster one-week schedule, there are fewer than 10 decision variables required even when there are several hundred opportunities to schedule. Because we have eliminated schedules that differ too much from the guidelines simply by excluding them from the encoding, the evolutionary algorithm spends much less effort generating op-

Usable opportunities: O1, O2, O3, O4
Science team rule: 0 to 2 must be included in schedule

Usable combinations:
    none: {}
    one: {O1}, {O2}, {O3}, {O4}
    two: {O1, O2}, {O1, O3}, {O1, O4},
           {O2, O3}, {O2, O4}, {O3, O4}
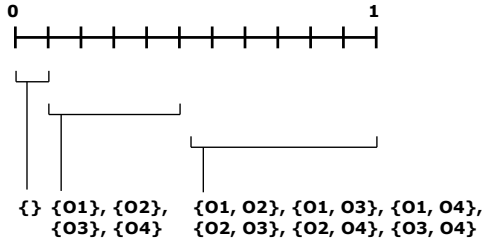
Decision variable encoding to interval [0,1]:

0   1

{} {O1}, {O2},   {O1, O2}, {O1, O3}, {O1, O4},
   {O3}, {O4}   {O2, O3}, {O2, O4}, {O3, O4}

Figure 6. Example of decision variable encoding for 4 opportunities of which 0..2 must be selected.

tions that will be quickly dropped.

Since we are using GDE, which operates on a floating point vector of values, we represent each decision variable as a value in the range [0,1] which maps to one of the choices using the encoding above. The residual value is used to determine which spacecraft combination to select from among all those allowed, while the minimum timing for the opportunity is used for evaluating the schedule. The combination of all choices specified by the vector then fully defines a schedule instance which can be evaluated using the objectives described above.

Schedule Visualization and Decision

We have developed a graphical user interface (the Cluster Schedule Explorer, or CSX) tailored to the Cluster scheduling domain, and in particular to the exploration of alternatives and tradeoffs as described above. Users specify which opportunity file to load, along with a existing DSN schedule file representing other missions already allocated to
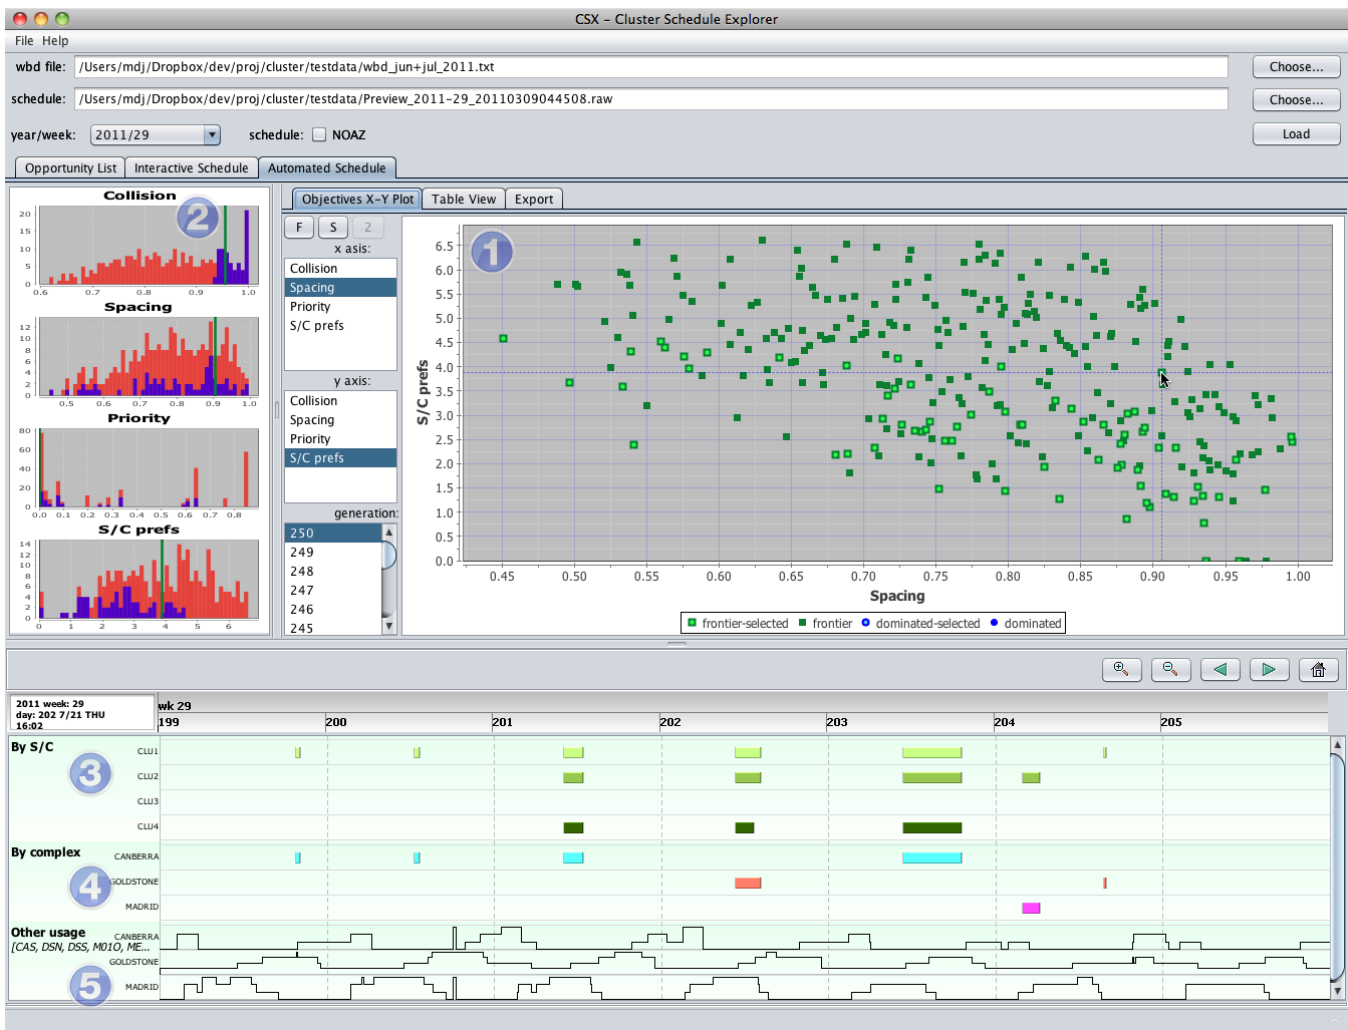

Figure 7. The Cluster scheduler graphical user interface, illustrating functions designed to assist the user in viewing tradeoffs and selecting among schedules for execution. Numbered items are described in the text.

antennas that could be in conflict with Cluster – these are used for the collision metrics calculations. Loading these two datasets for a particular week will run the evolutionary algorithm for a selected week, and present the results in a graphical view.

Figure 7 shows an example population of schedules generated with the prototype Cluster scheduler.

❶ The X-Y plot shows the evolved population as points graphing any pair of objectives – here showing spacecraft preferences vs. spacing. Better schedules are to the upper right, and the family of points clearly shows the tradeoffs. Because this is a projection 4D→2D, the interior points are also visible.

❷ A set of histograms of all of the the objective values are provide on the left side of the interface, to show the overall distribution (in red). We use a "brushed histogram" technique to show correlations among the objectives: the user can select a range in one or more histograms, and the corresponding values in the other histograms are highlighted (in blue). In addition, the selected points are highlighted in the X-Y plot. In Figure 7 the highest values of the collision objective have been selected, and are shown as open squares on the X-Y plot. It is easy to see the tradeoff of good values of collision avoidance with poor values for spacecraft preferences.

❸ From the X-Y plot, users can click on any schedule population member to get a Gantt chart view of all of the activities. This shows the timeline of which spacecraft are scheduled at what times, and resting the mouse over any item shows details of the timing, science type, and priority, as well as a reference to the WBD input table. The Table View tab provides a textual summary of the selected schedule in tabular form for more detailed assessment if needed.

❹ Underneath the spacecraft activity view is a view of Cluster activities at each of the three DSN complexes.

❺ The bottom view on the Gantt chart shows a histogram of usage by other DSN users that are used in the collision avoidance objective calculation. For example, there are three antennas at Canberra that can be used by Cluster, and the time when any of these are available can be seen at a glance.

Once the user has assessed the population, a specific schedule can be selected for export, from the Export tab. This generates an XML file for import directly into the DSN scheduling system.

Following an initial period of domain information gathering, the initial Cluster Schedule Explorer version was demonstrated to the Cluster WBD schedulers in early February 2010. Following a number of iterations on some of the constraints and what would be the most appropriate objective values, a second demonstration was conducted in early March, which led to a decision to use the tool in "shadow mode" for a trial in April. As that got underway in late April 2010, the team decided that the schedules being generated by the tool were already of sufficient quality to replace their pre-existing manual process, and they have been used directly since that time.

In summary, the Cluster Schedule Explorer has come into use routinely by the Cluster scheduling team to generate nearly a year of operational schedule inputs. The use of the MUSE scheduling software significantly shortens the time required to generate each schedule, and more importantly provides the scheduler with confidence that a good schedule has been created, by providing visibility into tradeoffs and quantitative insight into how well a schedule meets the various evaluation criteria.

4   Conclusions

We have described the MUSE Multi-User Scheduling Environment as an architecture for multi-user multi-objective scheduling, and in particular its adaptation to the Cluster WBD science scheduling domain where it is now in routine use. Future plans include the adaptation of MUSE to additional missions for the purpose of further validating our overall approach, and to evolve the framework for broader use. The combination of improved schedule comprehension and visibility, along with collaborative schedule development, offers the potential for a significant advance in scheduling support for future missions.

Among the challenges for the future are several that we are actively working on now:

- schedule revision: the most important remaining challenge for the Cluster domain is that of providing a mechanism for optimally *revising* an existing schedule as external factors impact the original choices. The DSN schedule is negotiated, and Cluster's inputs are among many that may have to be revised in order to resolve contention in the final schedule. Several strategies are being explored for re-importing an existing DSN schedule and re-scheduling around contentious time ranges.

- visualization: another area of active work is that of other visualization approaches that can be used for higher dimension objective spaces. Cluster makes use of four objectives, but more would not be uncommon.

References

Abraham, A., L. Jain and R. Goldberg (2005). *Evolutionary Multiobjective Optimization*. Berlin, Springer.

Deb, K. (2001). *Multi-Objective Optimization Using Evolutionary Algorithms*. New York, John Wiley & Sons.

Giuliano, M. and M. D. Johnston (2008). *Multi-Objective Evolutionary Algorithms for Scheduling the James Webb Space Telescope*. International Conference on Automated Planning and Scheduling (ICAPS), Sydney, Australia.

Gurnett, D., R. Huff, J. Pickett, A. Persoon, R. Mutel, I. Christopher, C. Kletzing, U. Inan, W. Martin and J. L. Bougeret (2001). "First results from the Cluster wideband plasma wave investigation." *Annales Geophysicae* **19**(10/12): 1259-1272.

Gurnett, D. A., R. L. Huff and D. L. Kirchner (1997). "The Wide-Band Plasma Wave Investigation." *Space Science Reviews* **79**(1): 195-208.

Johnston, M. D. (2006). *Multi-Objective Scheduling for NASA's Deep Space Network Array*. International Workshop on Planning and Scheduling for Space (IWPSS-06), Baltimore, MD, Space Telescope Science Institute.

Johnston, M. D. and M. Giuliano (2009). *MUSE: The Multi-User Scheduling Environment for Multi-Objective Scheduling of Space Science Missions*. IJCAI Workshop on Space Applications of AI, Pasadena, CA.

Johnston, M. D. and M. Giuliano (2010). *Multi-User Multi-Objective Scheduling for Space Science Missions*. SpaceOps 2010, Huntsville, AL.

Joslin, D. E. and D. P. Clements (1999). "Squeaky Wheel Optimization." *Journal of AI Research* **10**: 353-373.

Kukkonen, S. and J. Lampinen (2005). *GDE3: The Third Evolution Step of Generalized Differential Evolution*. The 2005 Congress on Evolutionary Computation.

Price, K., R. Storn and J. Lampinen (2005). *Differential Evolution: A Practical Approach to Global Optimization*. Berlin, Springer.

Spence, R. (2001). *Information Visualization*, ACM Press, Addison-Wesley.

Tidwell, J. (2005). *Designing Interfaces*, O'Reilly.