Challenges in Coordinating Remote Sensing Systems

Robert A. Morris and **Jennifer L. Dungan** and **John L. Bresina**
NASA Ames Research Center
Moffett Field, CA 94035
morris@email.arc.nasa.gov

Abstract

Earth scientists require timely, coordinated access to remote sensing resources, either directly by requesting that the resource be targeted at a specific location, or indirectly through access to data that has been, or will be acquired and stored in data archives. The information infrastructure for effective coordinated observing does not currently exist. This paper describes a set of capabilities for enabling *model-based observing*, the idea of linking scheduling observation resources more directly to science goals. Model-based observing is realized in this paper by an approach based on concepts of distributed planning and scheduling. The problem raises challenging issues related to planning under uncertainty, monitoring and repair of plans, and reasoning about human objectives and preferences.

Introduction

Earth observing systems study different aspects and interacting pieces of the lithosphere, hydrosphere, atmosphere, and biosphere. Scientists are designing increasingly complex, interdisciplinary campaigns to exploit the diverse capabilities of multiple Earth observation assets. Remote sensing platforms are being configured into clusters, trains, sensorwebs, or other distributed organizations in order to improve either the quality or the coverage of observations. Simultaneous advances in the design of science campaigns and in the missions that will provide the sensing resources to support them offer new challenges in the coordination of data and operations that are not addressed by current practice.

This paper shows how Earth science campaigns can be viewed as a distributed problem of planning, scheduling, and execution, requiring capabilities for temporal and spatial reasoning, reasoning about user preferences, and handling uncertainty. We propose a solution based on a notion of *model-based observing*, an approach for image scheduling based on the explicit correlation of assignments of targets to sensors with Earth science campaign goals. We describe a framework for model-based observing and show how such a framework can allow for more efficient use of resources, improve the scientific utility of the data products, and contribute to solving the critical problem of managing increasingly large quantities of remote sensing data. We briefly describe a prototype system that has been developed based on the framework and finally offer a possible list of extensions to the framework based on our experience with the prototype.

Earth Science Campaigns

An Earth science *campaign* is a structured set of remote sensing activities undertaken to meet a particular science goal. Science goals are stated with respect to geophysical or biophysical quantities (often called parameters or variables) that are to be measured. For example, the twenty-four parameters that are the goals of the Earth Observing System set out by NASA (King 1999) include land surface temperature, vegetation, ice and cloud cover, energy flux, and atmospheric concentrations of particles and gasses. Many science goals are directed toward the study of ongoing Earth processes, such as seasonal ice melt, vegetation green-up and senescence, agricultural yield dynamics, the spread of invasive species, and air pollution transport. Toward these goals, campaigns are constructed to assemble regular observations timed to capture snapshots of targeted regions when change is occurring. Other science goals are based on Earth system phenomena that are transient, ephemeral, unpredicted, or simply unknown such as floods, droughts, fire, volcanic activity, algal blooms, storms, insect defoliations, and oil spills.

The phenomena that a campaign aims to observe, which we will refer to as *exogenous events*, have spatial locations and extents and temporal initiations and durations. The structure of a campaign therefore consists of *spatial* constraints, requiring distinct images at specific locations; *temporal* constraints, involving images at different times; and various *sensor capabilities*, requiring images at distinct wavelengths, resolutions, or view geometries. In many cases, the relationships between exogenous events and observations is easily definable but when events are not predicted in advance there can be uncertainty in how to specify observations.

A hypothetical example to illustrate campaign structure is based on the goal to test an emissions model predicting the aerosols released by wildfires. The broad outlines of this campaign were inspired by the Fire Locating and Modeling of Burning Emissions project (Reid *et al.* 2001). Let us say the location of this campaign is in the southern California region of San Diego County. A burn event is anticipated to occur sometime during the "fire season," most likely between August and October in the region. Data on several variables must be gathered in order to accomplish the analysis. In particular, vegetation type or biomass, atmospheric aerosol concentration, and burned area are needed for the region. Fuel moisture content is a variable that also would be useful for the objectives of the science, though not a necessity.

There are several sensors that provide products at various spatial resolutions relevant to these variables. Landsat's Enhanced Thematic Mapper+ (ETM+) or Thematic Mapper (TM) can be used for mapping vegetation type. Optimal timing for acquiring Landsat data for this purpose in southern California would be the prior June or July in the same year that the fire burned, when forested land can most easily be spectrally distinguished from grassland. For mapping aerosol concentration, images coincident to burning must be obtained. MODIS on the Terra and/or the Aqua satellites would provide data for this variable. MODIS data from either platform could also be used to map the burned area with coarse spatial resolution after (though not too long after) the fire were out. Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) or Landsat TM data would be desirable for mapping the burned area with fine spatial resolution. For mapping vegetation moisture content, hyperspectral data from EO-1 Hyperion instrument are relevant. The most useful data for this purpose would need to be acquired just preceding the fire.

Current Practice and its Limitations

Current practice for developing and executing Earth science campaigns can be summarized by the following salient features:

1. Remote sensing assets are controlled and managed by individual science missions. Mission science planning teams generate command sequences (observation schedules) that are uplinked daily to a spacecraft.
2. There is little coordination among observations and no infrastructure for enabling coordination. Until recently there has been little perceived need for coordination; many satellite missions, such as those within the NASA's Earth Observation System, are global surveyors and image the entire Earth at a fixed temporal frequency and relatively coarse spatial resolution. Under these circumstances, observations are always available for a given geographical region, though not necessarily for a desired time period.
3. In recent years, sensors with fine spatial resolution are proliferating. They cannot cover the entire Earth's surface quickly and so must be tasked selectively.
4. Combining data from multiple sensors may increase their value beyond that of any single sensor.
5. Access to Earth science campaign resources is currently accomplished primarily through a user querying one or more *data archives*, repositories for remote sensing image data. Direct manipulation of the sensing resources, i.e., through submitting *requests* for observations, is much less common.

Although current practice is generally adequate in the present remote sensing environment, there are a number of limitations with this approach. They include:

1. Potential for lack of efficiency in use of resources. For example, sensors with the same or overlapping capabilities might be oversubscribed (i.e., there are more requests for observations than can be serviced), but there is no effective way to enable load sharing.
2. Sub-optimality in quality of science achievable in a campaign. Acquiring data through an archive vs. through direct sensor manipulation is analogous to eating at a buffet vs. a la carte. You may get full but it is not fresh and may not be exactly what you want.
3. Current practice is becoming less tenable as sensor technology becomes more specialized and needs arise to synchronize observations.

Coordination of observations is becoming increasingly important to successful science campaigns. As Earth science and supporting technology mature in parallel, the questions that are being asked involve models that seek to explain the relationships between different processes (e.g., atmospheric, radiative, hydrologic). Data to support these models will continue to involve heterogeneous sensor capabilities and, hence, more coordination of observations.

Approaches For Coordinated Observations

We propose the notion of *coordinated model-based observation* to overcome the limitations of current practice in campaign generation and execution. Coordinated model-based observation is based on the idea of *tying observation strategies more closely to the science goals of the user community* in order to improve the quality of the science obtained and to enable a more efficient use of sensing resources.

Coordinated model-based observation requires an interface for communication between the user community and the observation planning systems for sets of sensor resources. There are a number of different ways this interface can be designed. Figure 1 illustrates current practice (no coordination) and two alternative strategies for coordination: centralized and hierarchically distributed. In the former, a single system provides all the scheduling services to the complete set of observing assets, including generating command sequences. This approach has all the advantages (and disadvantages) inherent to centralized approaches. In particular, an advantage exists in having a single source of control over the scheduling decisions made and complete visibility into the constraints on all the resources. Conversely, a distributed approach may have computational advantages (and disadvantages) that arise from a partitioning of the problem.

Here the focus is on a distributed, *hierarchical* approach. This approach is selected primarily to preserve the current practice of independent management of sensing resources by separate mission scheduling teams. This conservative choice

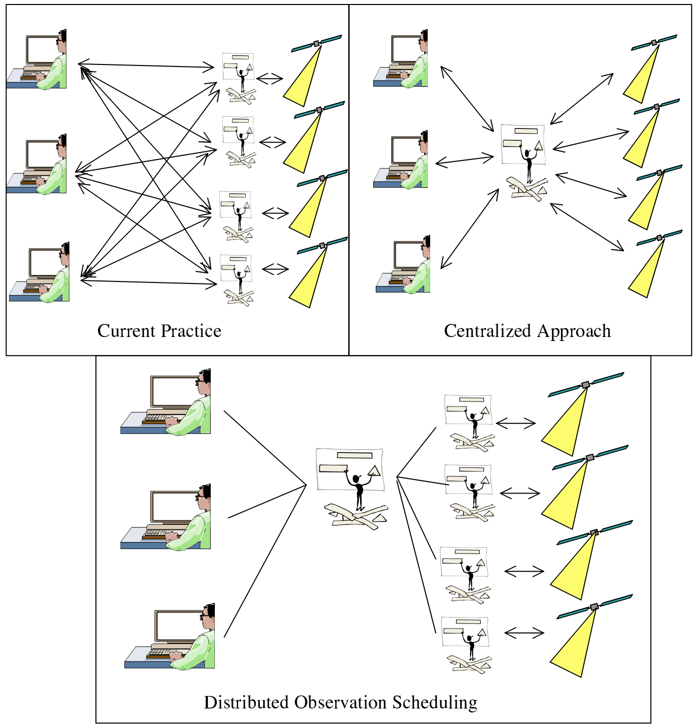

Figure 1: Three types of interface between users of remote sensing systems and those systems themselves. Current practice has no central interface. A purely centralized interface consists of a single system for observation scheduling for a set of assets. A distributed approach distributes the decision making between a central system and a set of schedulers, one for each instrument.

would likely entail smaller incremental costs in an environment where separate mission teams are already in place but may or may not entail lower total costs. There are many possible ways to instantiate distributed, hierarchical observation scheduling based on how the decision-making is distributed. For example, at one extreme the "root" system (henceforth referred to as the *coordinator*) offers nothing more than a mechanism for relaying requests. The disadvantage of this extreme is that it does not exploit the advantages of access and visibility into the sensor network and the end user is completely responsible for building and monitoring a campaign.

If there is to be scheduling decision-making by the coordinator, the question arises of how to partition the decisions between the coordinator and the individual *observation schedulers*. A natural partition is based on the different roles of the coordinator and the observation schedulers. The coordinator is a tool for the science community that produces a *campaign plan* that is to be executed over a long period of time (typically months). By contrast, each of the observation schedulers produces an executable sequence of commands to sensors on a daily basis. This difference in role and in the granularity of the execution horizon induces a natural partition between a set of *campaign constraints* managed by the coordinator, and a set of *resource constraints* handled by the observation schedulers.

In this application domain, it is assumed that the managers of the individual sensors (the *mission planners*) act independently to accomplish mission objectives that may be different from those of the users of the coordinator. This has implications both in the infrastructure for coordination (it is not in the interest of mission planners to consume vast amounts of time in negotiation with the coordinator over the use of a sensor) and in the campaign scheduling itself (there needs to be flexibility in the campaigns generated to allow for uncertainties in the availability of sensors).

The proposed distributed hierarchical approach creates two separate computational problems, the campaign scheduling and the observation scheduling. These are discussed below.

Campaign Scheduling Problem

The core campaign scheduling problem has been formulated as a constraint satisfaction problem (CSP). Each campaign consists of a set of variables representing features associated with requests for a set of *measurements*, as well a set of variables standing for exogenous events (Morris *et al.* 2006). At a minimum, measurement variables are defined for the *sensor* used in the measurement, the *time* of the measurement, and the *location* of the measurement. Domains for these variables are chosen, respectively, from a list of *sensors*, a range of times, and an *area of interest* (for example, a set of latitude/longitudes that make up the corners of a rectangle). An area of interest defines a set of *scenes*, minimal units for defining and querying locations on the Earth corresponding to images transmitted by remote sensors (for example, as defined by the Worldwide Reference System (wrs)).

Campaign constraints restrict, or define preferences for, the location, time, or sensors used for measurements. Temporal constraints include constraints on the start time of a measurement, or gaps between measurements, or between a measurement and an exogenous event. Some constraints might restrict the dynamic controllability (Morris, Muscettola, & Vidal 2001) of the campaign; for example, a constraint might specify that a measurement be taken between 1 and 30 days prior to the start of a fire. Other campaign constraints might include restrictions on the amount of cloud cover on the captured image. Finally, it is assumed that a monetary cost is incurred for a measurement; users can define constraints on the overall cost of the campaign.

Users are allowed to specify preferences for sensors and time of a measurement. Users are also allowed to distinguish between measurements that are *required* in a campaign and those that are merely *desired*. Science campaigners also have the choice to trigger the delivery of data directly at the time of an observation request or later through accessing an archive. We assume there are two advantages from the former approach: the user can define and tailor the constraints on the data to directly match the goals of the campaign (and thus enhance their scientific utility) and the data can be acquired more quickly than through an archive, which may be important for some campaigns.

A complete assignment of all the variables that satisfies the campaign constraints defines a *feasible schedule* for a campaign. A feasible schedule will be executed in the distributed setting by packaging the assignments into *observation requests* and submitting them to mission schedulers.

Observation Scheduling Problem

The daily resource scheduling problem can be formulated as an overconstrained CSP: given a set of observation requests, each weighted by a priority, and a set of resource constraints, generate a sequence of observations satisfying a subset of those requests that maximizes the overall priority. Resource constraints restrict the procedures through which data are captured, stored, and downlinked. They include constraints on power, data storage capacity, restrictions on visibility of ground stations for downlink, and on the sensor *duty cycle*, conserving the amount of time the instrument can be on to extend its lifetime.

Scheduling practices differ from mission to mission. For example, daily mission scheduling of Landsat 7 ETM+ images has inputs to daily mission scheduling from a database of requests called the Long Term Acquisition Plan (LTAP) (Potter & Gasch 1998). The LTAP is consulted each day for the purpose of directing the acquisition of images by the Landsat sensor. *Seasonality files* specify which scenes are to be acquired during which periods of time (request period) and the frequency of acquisition during those periods. Other files contain seasonal cloud cover information that are be used as heuristics to bias the scheduler towards cloud-free image requests.

The remainder of this paper focuses on solving the campaign scheduling problem. The following section discusses two core components to a framework for solving the campaign scheduling problem: mixed-initiative campaign generation and repair and campaign execution.

A Framework for Campaign Scheduling

Campaign Generation

Earth scientists formulate campaigns as a set of measurements and exogenous events and their associated constraints. The system determines the consistency of the constraints and manages the execution of the campaign. If the campaign becomes infeasible during execution (i.e., a constraint is violated) then the Earth scientist is consulted again to devise strategies for repair or abandonment of the campaign.

Uncertainty with respect to the availability of sensing resources, as observed earlier, suggests a *flexible* solution to campaign scheduling. We propose a generalization of the flexible temporal planning approach (Dechter 1991), extended to handle temporal preferences and exogenous events (see 2004 and 2005 for more technical details on these extensions).

A *Simple Temporal Problem with Preferences* is a generalization of a Simple Temporal Problem (Dechter 1991). An STPP can be depicted as a pair (V, C) where V is a set of variables representing events or other time-points and $C = \{\langle [a_{ij}, b_{ij}], f_{ij} \rangle\}$ is a set of *soft temporal constraints* defined over V. Informally, a soft temporal constraint consists of an interval that represents a restriction on the distance between arbitrary pairs of distinct events and a user-specified preference function defined on those distances.

To represent temporal uncertainty, we partition V into two groups: the *decision variables* V_d and the *parameters* V_u representing uncontrollable events. We further distinguish between binary *decision constraints* (C_d), those which the agent executing the plan must satisfy, and *uncertainty constraints* (C_u), those which "nature" will satisfy. An uncertainty (temporal) constraint depicts a duration as a continuous random variable. Assuming mutual independence, the constraints in C_u can be expressed in the form $\langle [a_{ij}, b_{ij}], p_{ij} \rangle$, where $p_{ij} : [a_{ij}, b_{ij}] \rightarrow [0, 1]$ is the probability density function over the designated interval. We call the framework $\langle V_d, V_u, C_d, C_u \rangle$, where C_d are soft constraints, a *Simple Temporal Problem with Preferences and Probabilities*, or $STPPP$.

Example 1 *Earth Science Observation Problem.* Inputs: Variables in V_d standing for two controllable events, each consisting of taking an observation, ($Obs1$, $Obs2$), and two uncontrollable events in V_u, the start and end of a fire (FS, FE) (for simplicity, observations are viewed as instantaneous), as shown in Figure 2. There is also an event TR representing the beginning of time. Soft constraints $f_1(t)$, $f_2(t)$ in C_d are associated with the durations between $Obs1$ and FS, and between $Obs2$ and FE, respectively. For example, $f_1(t)$ may express that there is no value for taking $Obs1$ after the start of the fire (FS) and a preference for times that are as close to FS as possible. Similarly, $f_2(t)$ expresses a

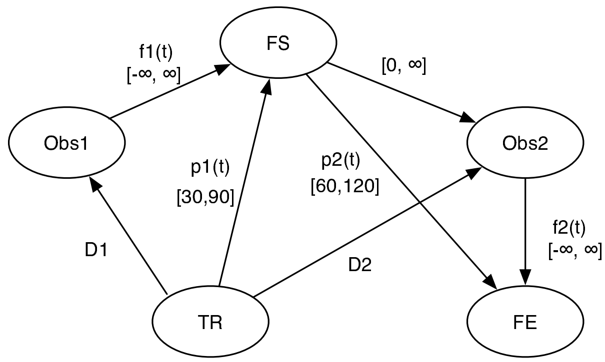

Figure 2: $STPPP$ Representing the Fire Campaign Scenario

preference for $Obs2$ happening before FE as close as possible, with a penalty if the observation is taken after the fire. Uncertainty constraints p_1, p_2 in C_u are associated with random variables representing the start time and the duration of the fire. These constraints are based on Earth science models about fires in the area of interest. For example, p_1 may express a normal distribution over the range of times.

Techniques have been, or are being, developed to formulate and solve flexible planning problems (e.g., generating solutions, determining consistency) and to address issues related to autonomous execution of flexible plans (i.e., determining the degree of controllability). These issues are beyond the scope of this paper.

Campaign Execution

Once a campaign has been generated in the form of a flexible plan, a collection of capabilities are required to ensure its successful execution. A successful (nominal) execution of a campaign is the case in which for all measurements there exist an observation request that was accepted by the corresponding mission. As noted earlier, campaigns may be generated well in advance of their execution and depict time on a coarse scale (most naturally, in terms of days). Furthermore, it is assumed in general that observation schedules may not be generated until hours before (e.g., roughly 8 hours before on Landsat 7) being uplinked to the space craft. Consequently, at campaign generation time, there is uncertainty as to whether, and how, the campaign will be accomplished.

There are four core issues we explore here with respect to ensuring nominal execution: communicating with missions, formulating requests, monitoring a campaign, and campaign repair.

Communication with missions Here we sketch a simple framework for communicating with missions, one that does not include a mechanism for negotiation. As noted earlier, given the independence of missions, it is not necessarily in the best interests of missions to incur the overhead of negotiation, which is why we ignore it here.

The simplest communication between the coordinator and the observation schedulers is *submission-response*, in which a request for a specific time on the resource is requested for a specific area and the response is binary (yes/no). Furthermore, because of the need for coordination, it should be assumed that the coordinator has a time window in which to *retract* a request. For example, if there is a "same day" constraint between two observations and one mission rejects a request, the coordinator should be allowed to retract the request for the other measurement. Similarly, the coordinator should be allowed to submit requests for one measurement to multiple sensor managers and, after one is accepted, retract the others.

Formulation and dispatching observation requests A measurement, the primary elements of a campaign, is defined minimally in terms of a location, time, and sensor. An *observation request* is an instantiation of these attributes formatted to conform to the inputs to the observation scheduler for the resource to which the request is being submitted.

A *singleton* observation request is a triple of singleton values consisting of a sensor, time (in days), and *scene* (the minimal unit to represent a region of the Earth). A *flexible* request is one in which either the time or location is a set of values (because we are assuming that an observation scheduler schedules for a single sensor, we assume that sensor values are always unique). A flexible request $\langle s, T, L \rangle$, where T is a set of times and L is a set of locations, is thus the cartesian product $R = s \times T \times L$ of singleton requests. A default interpretation of a request of R is *take exactly one singleton request from* R.

If there are no coordination constraints between distinct measurements, a flexible request format would be an efficient way to increase the chances of a request being serviced. Because of possible coordination constraints, in some cases singleton requests should be submitted to missions. For example, enforcing a "same day constraint" between observations on distinct sensors could not be easily enforced if flexible requests were submitted to each mission.

Monitoring a campaign Monitoring a campaign requires detecting and recording changes to the state of an executing campaign, propagating their effects, and detecting the need for campaign repair.

A campaign is monitored by implementing a *campaign state model*. A campaign state model consists of three state types: plan, measurement, and request. States of the same type are linked by a set of legal transitions. The campaign model defines these legal transitions and also defines changes in state from one state type that trigger changes in another.

Figure 3 shows a state transition diagram for measurements. A measurement starts in a feasible state. It becomes enabled when the temporal preconditions for taking the measurement are met (for example, an exogenous event happens or a dependent measurement has been acquired). It becomes infeasible if the constraints make it impossible for it to be taken; this can happen, for example, if all submissions of requests for the measurement are rejected. Otherwise, a measurement is pending if at least one request for the measurement has been submitted. If a mission accepts the request

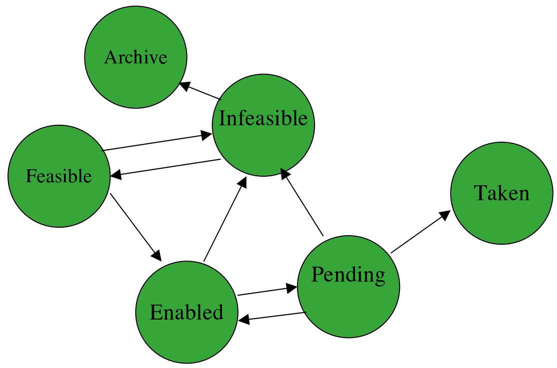

Figure 3: A state transition model for measurements. States and possible transitions between them are depicted.

and the image is acquired, the measurement enters the terminal node *Taken*. The user may decide to acquire the needed data from an archive. If so, requests for measurements are no longer submitted.

Inputs that can cause a transition in state include changes due to the passage of time, the occurrence of an exogenous event, or responses from missions. Such changes can impact the current campaign. In particular, observation requests can be rejected or new constraints are added to the campaign. Exogenous event detection can be facilitated through web agents. For example, the onset of a fire that is crucial to a campaign in the requested area may be discovered on the web by such an agent.

Campaign repair The need for *campaign repair* arises when a constraint on the original campaign is violated, i.e., when the plan becomes *infeasible*. Infeasible plans can be made feasible by *relaxing* constraints. The overall purpose of relaxing constraints is to create a new set of constraints that is consistent; i.e., for which there exists at least one feasible fixed solution. The main repair actions, performed by the human campaigners, are:

- Relax the bounds on the duration constraints between pairs of measurements and/or exogenous events. By increasing the duration bounds, opportunities that previously violated the temporal constraint can be made consistent with the new constraint.
- Create new opportunities that are consistent with the other constraints. Users may expand the set of viewing opportunities by increasing the list of sensors that can be used to take a measurement, by expanding the time window for taking a measurement, or by increasing the area of interest.

A Prototype for Coordinated Observations

A prototype system for the capabilities described in this paper has been implemented. The components of the system, called DESOPS (Distributed Earth Science Observation Planning and Scheduling) is found in Figure 4.

A **constellation model** allows a campaign planner to reason about the environment within which campaigns are defined and executed. It consists of a database and set of functions for defining the capabilities and dynamics of resources available to the user for observation. Components to a constellation model are:

1. A description of the capabilities of a collection of satellite-borne, Earth-pointing *sensors*.
2. A model of *time*. For the purposes of coordinated campaign, time can be viewed as a finite set of totally ordered values naturally interpreted as the set of days in which some observation can be taken or some other event of interest happens.
3. A *global coordinate system* for data, enabling a user to inquire about satellite imagery over any portion of the world by specifying the location of the data of interest. Examples of coordinate systems are the Worldwide Reference System (WRS) (wrs), or latitude-longitude.
4. A *satellite orbit* function for determining the set of sensor viewing times for a specified region of interest.
5. For each sensor resource, a *mission model* that describes constraints on the process by which tasks on the sensor are scheduled by the mission that manages it.

The constellation model provides a language for specifying the requirements for using a collection of sensing resources and for allowing users to build high quality campaigns. Users of the system define campaigns and can view their execution, monitoring ,and repair over time through a **graphical user interface**.

The **planner** organizes a specification of a campaign into a representation of the flexible plan for accomplishing the campaign goals. The planner generates start times for each sensor in the domain of each measurement from *view paths* over specified regions of interest during specified time windows, using the constellation model. A view path is the intersection of a specified region of interest with the path followed by a satellite over the user-specified time window. View paths are generated by conducting a web search for this data from mission web sites. [1]

The **request manager** contains services for managing the execution of a plan. These services include submitting *observation requests* for each measurement in a campaign to missions, monitoring each submitted request for mission response, monitoring for the occurrence of exogenous events, and triggering repair actions.

Finally, the **plan database** stores the current information about every campaign. The database is used in all phases of campaign planning and execution and contains definitions of all the measurements in the campaign, constraint information, descriptions of the observation requests generated and submitted to missions, and current campaign state information. Different system components query and revise different parts of the plan database at different states of the campaign.

[1] Alternatively, it is possible to generate this data directly through the use of simulators such as STK (Satellite Tool Kit, AGI, Philadelphia, PA).

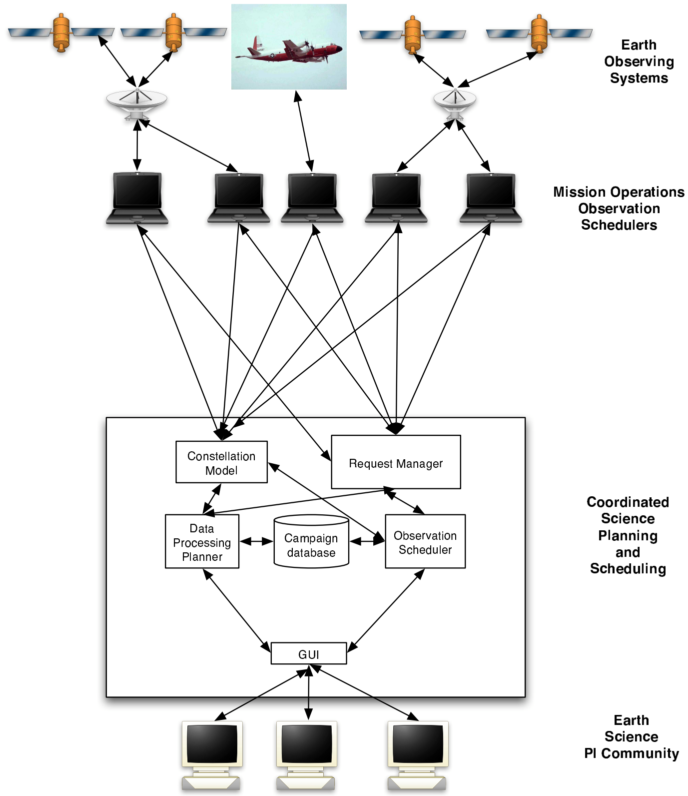

Figure 4: DESOPS architecture for coordinating Earth science campaigns

Extensions

There are numerous ways of extending the list of coordination capabilities described in this paper. Among them are the following.

- *Consideration of the probability of request acceptance*. When the coordinator selects a sensor platform to submit an observation request, it must ensure that the platform is capable of satisfying the constraints of the observation. When more than one platform is applicable, ideally, the coordinator would select the platform that had the highest probability of accommodating the new request. Ideally, the predictive models for the distributed sensors would be based on characteristics of current and future schedule loads and on factors that the scheduler (machine or human) used in constructing schedules.

  Organizations that control the sensor platforms typically share little or no information about the scheduling loads and methods. Without such information, one could still try to build a predictive model for each of the distributed sensor resources by collecting statistics based on the history of request-response pairs. One could estimate, for example, the probability of acceptance for different types of observation requests, where the requests were characterized in terms of temporal windows and spatial areas. Of course, a more effective predictive model could be constructed if the proprietary information in the schedules were abstracted away and the coordinator had access to the aspects that influence inclusion in a schedule for each of the distributed sensors.

- *Automatically incorporating process models and results of analysis*. A campaign arises from the need to build or validate process models about the earth's eco-systems. Data from observations are processed and analyzed, and new observation goals generated. The capabilities described here depend on the human user to generate observation goals. However, this process could be automated. *Goal generation* is the process of generating specifications of new measurements from data and models. For example, the combination of a fire model plus climate forecasts might combine to predict that a certain area of the western US is highly likely to experience brush fires during a certain period. This prediction can be transformed into a specification of a measurement. Similarly, the results of analysis may indicate gaps or uncertainty in the knowledge obtained, which can generate new observation goals. Thus, this extension allows for more domain knowledge to be integrated into the specification of the campaign.

- *Coordinating multiple campaigns*. More effective campaign management can result from *joint requests*, i.e., requests for the same measurement by multiple campaigners. The impetus for joint requests is cost-sharing, which results in lower overall costs for each campaign.

- *Load balancing*. As noted earlier, different sensors with overlapping capabilities can have different demands, and more efficiency can be gained by load balancing. In this extension, missions with oversubscribed resources could submit requests for observations that did not fit into their schedule. The coordinator would manage the submission of the requests to other missions.
- *Coordinating mixed observation platforms*. Future missions will combine observations obtained with remote sensing resources with sub-orbital (aircraft) and ground sensing (for example, the Intex mission (Singh 2004)). Unlike remote sensing assets, coordinating sub-orbital in situ sensors will require the generation of *flight plans* for the aircraft on which they reside.

Conclusion

This paper has described a complex problem in distributed scheduling for Earth science campaigns and a solution based on the principle of model-based observing. Challenging issues arise in the need to formulate and execute campaign submitted by users. The benefits of added coordination are the standard ones, viz., more efficient use of resources and better quality in the plans generated.

References

Dechter, R. 1991. Temporal constraint networks. *Artificial Intelligence* 49:61–95.

King, M., ed. 1999. *EOS Science Plan: The State of Science in the EOS Program*. Greenbelt, Maryland: NASA Goddard Space Flight Center.

Morris, R.; Morris, P.; Khatib, L.; and Smith, N. Y. 2005. Temporal planning with preferences and probabilities. In *Proceedings of the ICAPS Workshop on Constraint Programming for Planning and Scheduling*.

Morris, R.; Dungan, J.; Edgington, W.; Wood, T.; Williams, J.; and Yorke-Smith, N. 2006. Coordinated observation scheduling. *Proceedings of i-SAIRAS 2006*.

Morris, R.; Morris, P.; and Khatib, L. 2004. Strategies for global optimization of temporal preferences. In *Proceedings of CP 2004*.

Morris, P.; Muscettola, N.; and Vidal, T. 2001. Dynamic control of plans with temporal uncertainty. In *Proceedings of the 17th International Joint Conference on A.I. (IJCAI-01)*. San Francisco, CA: Morgan Kaufmann.

Potter, W., and Gasch, J. 1998. A photo album of earth: Scheduling daily Landsat 7 activities. In *Proceedings of Space Ops 98*.

Reid, J. S.; Prins, E. M.; Westphal, D. L.; Christopher, S.; Schmidt, C.; Richardson, K. A.; Theisen, M.; Reid, E. A.; and Eck, T. 2001. Flambe: The fire locating and modeling of burning emissions project. In *Proceedings of the Fall Meeting of the American Geophysical Union*.

Singh, H. 2004. Intex-na: Intercontinental chemical transport experiment - north america. *http://geo.arc.nasa.gov/sgg/singh/*.

The worldwide reference system. http://landsat.gsfc.nasa.gov/documentation/wrs.html.