

Science Goal Driven Observing: A Step towards Maximizing Science Returns and Spacecraft Autonomy

Anuradha Koratkar^a, Sandy Grosvenor^b, Jeremy Jones^c, Nargess Memarsadeghi^c, Karl Wolf^d
^aSpace Telescope Science Institute; ^bBooz Allen Hamilton; ^cNASA/Goddard Space Flight Center;
^dAquilent

ABSTRACT

In the coming decade, the drive to increase the scientific returns on capital investment and to reduce costs will force automation to be implemented in many of the scientific tasks that have traditionally been manually overseen. Thus, spacecraft autonomy will become an even greater part of mission operations. While recent missions have made great strides in the ability to autonomously monitor and react to changing health and physical status of spacecraft, little progress has been made in responding quickly to science driven events. The new generation of space-based telescopes/observatories will see deeper, with greater clarity, and they will generate data at an unprecedented rate. Yet, while onboard data processing and storage capability will increase rapidly, bandwidth for downloading data will not increase as fast and can become a significant bottleneck and cost of a science program.

For observations of inherently variable targets and targets of opportunity, the ability to recognize early if an observation will not meet the science goals of variability or minimum brightness, and react accordingly, can have a major positive impact on the overall scientific returns of an observatory and on its operational costs. If the observatory can reprioritize the schedule to focus on alternate targets, discard uninteresting observations prior to downloading, or download them at a reduced resolution its overall efficiency will be dramatically increased.

We are investigating and developing tools for a science goal monitoring (SGM) system. The SGM will have an interface to help capture higher-level science goals from scientists and translate them into a flexible observing strategy that SGM can execute and monitor. SGM will then monitor the incoming data stream and interface with data processing systems to recognize significant events. When an event occurs, the system will use the science goals given it to reprioritize observations, and react appropriately and/or communicate with ground systems – both human and machine – for confirmation and/or further high priority analyses.

Keywords: science optimization, data communications, onboard analysis, spacecraft autonomy, science goals

1. INTRODUCTION

Fifteen years ago, the standard paradigm for obtaining observations was for astronomers to travel to the telescope. Now, with queue scheduling, digital imaging, and the Internet, instead of astronomers traveling to obtain their data, data is electronically delivered to their desktop. In this new paradigm, scientists involved with the project no longer have “their eye at the telescope.” Some astronomers argue^{1 2 3} that the astronomer at the telescope can make timely decisions and therefore achieve the best science. This implies that in a queue-scheduled/remote observing environment, we will be able to optimize the quality of our scientific returns if we can find ways to enable the science-driven dynamic response that is achieved by the “eye to the telescope” approach of traditional ground-based observing.

1.1 Current ground-based observatory operations

At ground-based observatories, the health and safety of the telescope is monitored by staff who are often geographically co-located with the observatory, so that they can respond to problems as fast as possible. The astronomer/observer is either at the telescope (the observation is not queue scheduled), or by using fast Internet links and other suitable technology, they may be geographically separated, but virtually at the telescope. Further, at ground-based telescopes there is a “reasonable amount” of scheduling flexibility for the observer to reprioritize the observing program to achieve his/her science goals. Thus, in the ground-based observatory scenario, even with queue scheduling, we can develop a virtual observatory where the astronomer is driving the observing strategy and schedule depending on science driven

events. With the number of space-based observations playing an important and growing role in astronomy, we need to find better ways to remotely achieve this same science-driven dynamic response as the traditional ground-based observatories.

1.2 Current space-based observatory operations

Recent missions have made great strides in the ability to autonomously monitor and react to changing health and physical status of spacecraft⁴, yet little progress has been made in responding quickly to science driven events. At present, after the observer submits his/her proposal, he/she has no involvement in the observation execution, until data are delivered often months later. For simpler observations, this may be desirable. However, because of the time lag between proposal submission and receipt of observations there could be many occasions when an observer wants to change parts of the observing program to respond to changing science events. Observers are often restricted in optimizing their observations close to the time of execution because of the rigidity in the observing schedule. If the observer makes changes to the observations after a program has been submitted, the observing program often loses its place in the long-range schedule, leading to an even longer delay in obtaining observations. Hence, changes to an observing program are done only if critically necessary. We note that the flexibility in making changes to an observing program and the time lag between proposal submission and data receipt is very much dependent on the spacecraft.

For observations of inherently variable targets, monitoring projects and targets of opportunity, the ability to recognize early if an observation will not meet the science goals, and react accordingly, can have a major positive impact on the overall scientific returns of an observatory and on its operational costs. Thus, for at least a class of targets there is a need for flexible scheduling to capture quality data to achieve science goals.

1.3 Factors influencing science driven operations

There are several factors that are essential to address if we are to even consider the possibility of reacting to science driven events in space-based observing. These include: (1) the ability to schedule observations flexibly, (2) the ability to capture science goals in a machine interpretable format to make event driven decisions, and (3) the ability of the observatory to adapt dynamically and autonomously to a changing schedule or set of observing priorities.

Flexible scheduling implies different things to different people. In this paper, we use the term to mean the ability to change the observing strategy within a given scheduling parameter space. For example, one can change the target, exposure length, number of exposures, etc. For a long-term variability study, it is desirable to be able to change the frequency with which a target is observed depending on the target's outburst phase. When an observer is at the telescope, they can evaluate the existing conditions and/or initial observations and adjust their subsequent observing parameters to adapt to those conditions. In order for an automated flexible scheduling system to be successful, it must accommodate these scientific parameters: it must capture the science goals of the project and autonomously set the optimal parameters to achieve those goals.

At present, science programs are specified by the mechanics of the observation. The underlying goals of the observation are not captured in machine interpretable format. For example, in an observing proposal, only the target, instrument/detector/filter combination and the exposure time are captured. *The underlying science goals of the observation are not captured even though they are often clearly stated in the scientific justification of the program.* For example, one such scientific goal would be the desire to observe a variable brightness target long enough to achieve a particular signal-to-noise ratio (SNR) instead of a fixed exposure time. Exposing for less time than necessary to reach the SNR will not achieve the science objective, and exposing longer than necessary is inefficient and could risk data overflows that damage the observation. Nor do the science programs capture contingency plans - if an event occurs that will disrupt the science plan, there is no mechanism to alter the plan to adjust for the new event. The results are periods of lost observatory time, lost/useless data, and lost opportunities for interesting science. This is especially the case for intrinsically variable targets. About 30% of the variability campaigns do not capture the phenomenon effectively. This is either because the observations could not be scheduled properly or the intrinsic behavior of the target did not become interesting until near the end of the campaign when no more observational time was available to the project. In addition to being able to capture science goals, we will need onboard capability to recognize when the goals are achieved (or determined to be un-achievable), and the ability to dynamically adjust the observing plan accordingly. At present spacecraft do not have the capability to adapt to changing observing conditions. There are, however, prototype missions

such as Techsat-21 and Three Corner Sat (3CS) that will fly in the next few years and will test the necessary hardware and software technologies.

The ability to capture science goals is not only important to obtaining observations, but can also be used at the time of data download. Many of the upcoming missions will not only have better “eyes,” they will generate vast volumes of data making ongoing communication and data downloading more difficult/expensive. Significant research and development efforts, such as the Ultra High Data Rate Communications⁵ project, are underway to increase the download capacities to meet upcoming demand, but the bandwidth may not be available fast enough and it will come at a cost. The cost of downloading every pixel of these immense datasets may exceed the feasibility for cost-capped mid-sized missions that must share existing communications resources thereby restricting their available bandwidth and contact time. The download challenges will be further compounded by upcoming constellations/fleets of high data rate spacecraft, each needing to download data over an increasingly crowded communications facility. Consequently, it will become increasingly important to be able to autonomously evaluate the data prior to downloading, and download only scientifically valuable data (i.e. we may have to make prioritization decisions because we might not be able to download all data.)

If we have the ability to determine the status/quality of an observation before transmission to earth, we can use our scant data communications resources efficiently. The quality of an observation can be determined using the science goals for the proposal. This strategy can be further refined if the observer/scientist has the ability to define simple analysis algorithms for onboard execution.

Download bandwidth and storage capacities are also often parameters that become fixed and immutable very early in the design process of missions, therefore giving little encouragement to exploit the capabilities from a science perspective. For example, the recent New Millennium mission, EO-1, can hold only five scenes at once, since the size of the average scene is 1.2 gigabytes, and it was designed with this storage limitation as a “given”. Consequently, they have a fairly fixed downlink schedule, which does not easily accommodate a dynamic volume in downlinks. This complicates the ability to test onboard dynamic image processing or scheduling that might increase science value of the data downloaded but may require that the download volume vary. For example, we may be able to test onboard algorithms to recognize clouds obscuring the current “scene” and dynamically adjust the schedule to cancel the current scene and move to other non-obscured scenes. This means that the downlink scheduling needs to be flexible as well (less data and possibly more data later).

With the existing operations paradigm, observations do not always proceed as required by the scientist close to the time of the observation, and the observer has little voice in how the observatory responds to these anomalies. The resulting observational data may be less than optimal, or even useless. If the observatory can reprioritize the schedule to focus on alternate targets, discard uninteresting observations prior to downloading, or download them at a reduced resolution, its overall efficiency will be dramatically increased. By giving astronomers tools to evaluate their programs in entirely new ways, we envision not simply productivity gains, but also the ability to optimize a program both to increase the level of scientific return, and to avoid observations where the quality of an observation does not live up to expectations.

2. ENABLING TECHNOLOGIES

While increasing data gathering capabilities are driving the need for goal driven observing, there are also other emerging technologies that will enable its development and implementation.

2.1 Onboard capabilities

The onboard processing and storage capabilities of missions are increasing at a significant rate. The next generation of missions will have the necessary onboard processing and storage for performing sophisticated calculations in near real-time. The Next Generation Space Telescope (NGST), for example, is currently base-lining a flight CPU capable of sustained operation of at least 80 MIPS, with an option to increase to 750 MIPS before launch in 2010⁶. NGST data volume studies recommend approximately 50 GB of onboard storage for science data⁷. With these capabilities, future missions will be able to conduct real-time science data analysis onboard.

2.2 Enabling software technologies

As onboard processing, memory, and storage capacities increase, so do the power of the onboard software tools. There have been significant advancements such as the emergence of such tools as Embedded Linux, and the Java 2 Micro-Edition (J2ME), both of which are designed to run on small, power and memory constrained devices with little or no direct user-interfaces. This will allow a jump in the flexibility of onboard software. Both Linux and Java have well-developed, multi-threaded capabilities that contain well respected security precautions built into their architectures. Such capabilities are readily available now, and should be aboard spacecraft within the next few years. These technologies will substantially increase the onboard software capabilities, and ease the software development costs for developing applications because they are well-known development environments.

2.3 Flexible onboard scheduling

One desired outcome of detecting science events is the modification of the observing plan/schedule. The onboard planner and scheduler must be flexible enough to accommodate unexpected changes while maintaining a consistent plan and schedule. NASA's Deep Space One spacecraft pioneered the use of onboard scheduling with the Remote Agent (RA) software. RA includes an onboard planner and scheduler that is able to re-plan based on changing spacecraft conditions⁸. Another JPL product, CASPER (Continuous Activity Scheduling Planning Execution and Replanning) is emerging as a popular choice for future spacecraft including Three Corner Sat (3CS) and Techsat 21. CASPER is an onboard planner and scheduler that aims to achieve very high responsiveness to plan changes⁹. It uses iterative repair to support continuous modification and updating of the working plan as conditions change. For example, the spacecraft may image a target volcano with its radar and generate a reflectivity image onboard. The onboard science software compares this new image with the previous image to detect changes in the lava field. Based on the changes, the software can interact with the CASPER scheduler to arrange for a new high-resolution image centered on the area of interest.

2.4 Event/Feature recognition algorithms

The advent of increased processing power, onboard storage, and main memory in conjunction with sophisticated software environments have enabled the capability to detect observation events in a timely manner on a spacecraft. Previously, the amount of processing power to analyze the science data looking for significant features would have been prohibitive or impossible to attain. The Techsat-21 mission provides an example of an onboard event detection algorithm¹⁰ where a new image is compared with a previous image onboard, to detect changes/differences. Another example is the cloud cover detection algorithm. This algorithm determines the image quality to determine how much of the imaged earth surface is covered by clouds. Once the cloud coverage has been determined, other satellites and instrumentation can be redirected appropriately. The above techniques of acquiring a new image or the creation of a new science product and comparing it to previous acquisitions are expected to be common and effective onboard algorithms. Processing of this type has necessarily occurred on the ground in the past, but could now be migrated onboard to provide shorter response times for both the detection of changes and for the responses to those changes.

The upcoming SWIFT mission also reflects a step towards scientific autonomy. It will recognize gamma-ray bursts (GRB), and evaluate their intensity using a fixed onboard analysis to determine a figure-of-merit (FOM). If merited, the telescope will automatically slew to make additional observations on the GRB without waiting for human or ground-based guidance. This is a careful and important first step towards goal-oriented observing. We believe much greater returns are possible if, at the time scientists define their observing program, they have the ability to specify science goals and the observatory's reaction if those goals are met.

3. FUTURE TECHNOLOGIES

3.1 Dynamic algorithm definition

Emerging standards such as MathML and the runtime protections of Java may allow a huge leap in the flexibility of online processing by providing the ability for scientists to specify their own analytical algorithms to be run onboard as part of the goal analysis. The scientist could then perform this analysis in near real-time as the observation is being made, instead of waiting until all the observations have been made, the data downloaded, cataloged, archived, and finally delivered. Much work needs to be done to develop this concept: management and overseeing the amount of computation consumed, ensuring the protection of the spacecraft's software systems, pre-campaign validation of the algorithms, development of standard tools and data access interfaces within the scientist's dynamic "runtime"

environment, etc. However this could yield a huge increase in the flexibility a scientist has and could enable a corresponding leap in the quality and efficiency of observing time.

3.2 Natural Language

Ideally, science goals should be captured at the highest level, i.e. in natural language. Efforts towards making machines understand human languages and/or act as them resulted in different generations of formal languages, with each generation trying to get closer to the level of natural languages. Making computers understand natural languages is ongoing research in the field of artificial intelligence.

Domain knowledge is implicit in the way scientists describe their science goals in the scientific justification of their program. For example, a specific terminology or word can have different meanings in different science domains. This is driving research efforts to make the World Wide Web a "Semantic Web". One approach is through the use of 'Ontologies'. "An ontology defines the terms used to describe and represent an area of knowledge"¹¹. Having a standard and common understanding between the user and the way a machine interprets user's input is the key to our goal capturing and definition efforts. Further, we need software agents that can "understand" and recognize these relationships and then react to them. The ontology of basic concepts in the various space science domains and the relationships between these concepts needs to be developed.

Currently the problem of capturing science goals in natural language is vast, and with the current state of information technology, it is not possible to capture all conceivable science goals in machine interpretable format. But we believe there is a subset of science problems that we can and should be able to capture and reproduce. One example is to capture goals to represent an effective strategy for observing time-variable phenomena. Therefore, in the Science Goal Monitor project (SGM; described below), we are confining our objectives to more feasible goals: developing a visual interface that can capture a subset of the "science goals", i.e. goals that can be defined as measurable objectives; allowing a user to define contingency plans when appropriate; interfacing with the science data stream to determine if a goal is met; and then notifying the relevant person/scheduler regarding changes in priorities.

4. RISK MITIGATION

Introduction of flexible scheduling and autonomously reacting to science driven events inherently infuses automation technologies into mission operations. Another reason for introducing automation technologies into mission operations centers is to reduce operations cost. Fewer personnel are needed for mission support. Recent GSFC's Advance Architectures and Automation Branch studies suggest that automation has yielded success in operations staff reductions, particularly for small, simple missions, but there remain barriers that impede further progress for complex, high profile missions such as Terra and HST^{12 13}. Missions, especially complex high-profile missions, are more culturally and politically averse to risk when it comes to automation. The trade-off between scientific gains for successful risks versus the damaging publicity for failed risks is difficult to ignore.

In addition, onboard flight software has historically followed very conservative development and implementation methodologies due to the critical nature of the role it plays in the overall success of a mission. The skepticism is not limited to the flight software teams; scientists themselves remain leery of expert/automation systems and are not yet convinced that their unique goals can be effectively and accurately captured and executed. Yet, reducing the perceived mission risk associated with increased use of automation is of great importance if we are to realize efficiency gains and science optimization.

Clearly, the capture of science goals rather than the mechanics of an observation while developing observing programs, and the subsequent automatic analysis of the data stream to determine if goals are met, represents not just a leap forward in automation, but a wholesale change in the operations paradigm. We are in the early stages of the Science Goal Monitor project to develop prototypes, evaluate their effectiveness, and understand the risks.

5. THE SCIENCE GOAL MONITOR PROJECT

The Science Goal Monitor (SGM) project is a research effort funded by NASA Code R to determine if it is feasible to “get the eye back to the telescope” and to develop prototype software that will enable this. In the first phase of our study, we have determined that there is a subset of space science problems that are conducive to the philosophy of science driven observing. These problems are related to objects that are time variable and are often monitored for long periods. Using these problems as test cases for our prototype, we are in the process of designing and developing the SGM system which is a set of tools *that have the ability to capture the underlying science goals of an observation, translate them into a machine interpretable format, and then autonomously recognize and react in a timely fashion when goals are met.* SGM will provide astronomers with visual tools to capture their scientific goals in terms of measurable objectives and be able to autonomously monitor the data stream in near-real time to see if these goals are being met. Our prototype is designed for use in a distributed environment where some analysis can be performed onboard a spacecraft, while other analyses can be performed on the ground.

From our experience with the Scientist’s Expert Assistant (SEA) project, we know that infusion of software is easier after the test bed is built, and there is something tangible to demonstrate rather than “vaporware.”^{14 15} To test bed our proof-of-concept prototype, we have determined that an observatory that has the following conditions would get the most benefit from an SGM like system:

- *A typical target is intrinsically variable.* Science observations of inherently variable targets and targets of opportunity by their very nature need contingency plans to maximize the quality of science data acquired. There may be occasions when the target is too faint and proposed observations will generate useless data or that the source behavior is interesting and more observations are needed, or sufficient data have already been acquired. An effective strategy to dynamically reprioritize the observing will be essential. An important step in that process is to acquire, analyze and monitor the science goals. Because of the expected data volumes, this will have to be an automated process so that the telescope can react in “real” time.
- *Targets are monitored for long periods of time.* In long visits, to make sure that the best quality observations are obtained, it will be essential to monitor the target behavior, and dynamically schedule the observatory. To achieve this it will be essential to know and monitor the science goals for each target. This can be done manually for small data volumes but not with large data volumes and small operations staff. A SGM on the ground would be critical.
- *Spacecraft designed for scheduling flexibility.* Changes in observing schedule can easily be implemented.

Our initial partner mission to provide a focus for our prototype was *Kronos*. *Kronos* was proposed to the NASA Office of Space Science (OSS) as a Mid-sized Explorer (MIDEX) mission that would exploit the intrinsic variability of accreting sources, to map the extreme environments of black holes and other compact objects. Since *Kronos* observatory would be designed specifically for long-term, high time-resolution observations of accretion-driven sources, it was an ideal test-bed for our proof-of-concept. Unfortunately, *Kronos* was not selected for 2002 MIDEX selections, but we are actively evaluating other opportunities to provide a test-bed for our vision. Many Earth Science missions acquire data and images and then invest significant effort understanding changes in them. Such missions may provide ideal collaborative environments for testing and evaluating SGM’s concepts. *By researching strategies and developing tools to improve the ability to monitor and react to the changing scientific status of observations, we will be enabling a step towards spacecraft autonomy.*

5.1 The Science Goal Monitor (SGM) project overview

The Science Goal Monitor (SGM) system will interact with not only the data processing pipeline for a mission, but when used on board the spacecraft, it will also interact with the raw data from the detector. Figure 1 shows the high-level concept of the SGM. At present, we are concentrating on designing the Science Goal Capture Tool (SGCT). We hope to design and develop an underlying architecture and framework for the SGM system (see Figure 1). The three key modules include:

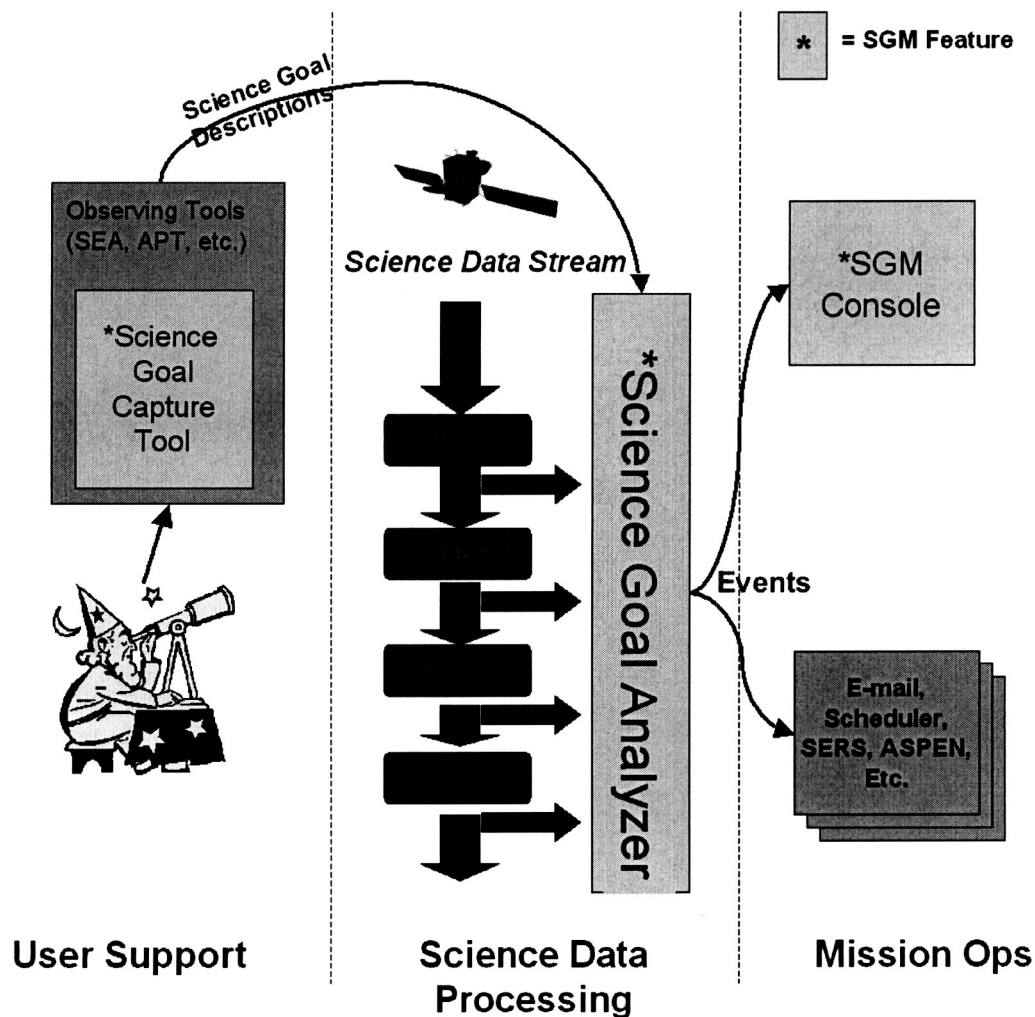


Figure 1: SGM System Overview (ground-based)

- The Science Goal Capture Tool (SGCT) – an interface to be used by the astronomer to help identify and specify science goals in a manner that can be programmatically evaluated.
- The Science Goal Analyzer (SGA) – will interpret the science data stream from the detector/observatory and interact with its processing pipeline to analyze whether or not a set of science goals are being satisfied. When goals are met, the SGA will fire events notifying any registered observer of the event and provide access to the details.
- The Science Goal Monitor Console (SGMC) – this will provide a visual interface and console for mission operations to query, monitor, and interact with the SGA.

There are already several event notification systems in development or operation. We propose to provide an interface to one or more of these systems, but do not plan to develop one specifically for the SGM.

Our initial proof-of-concept prototype will be a distributed ground-based solution, but our architecture will be designed to support a mixed space-based and ground-based system. An ideal operational scenario would include science goal

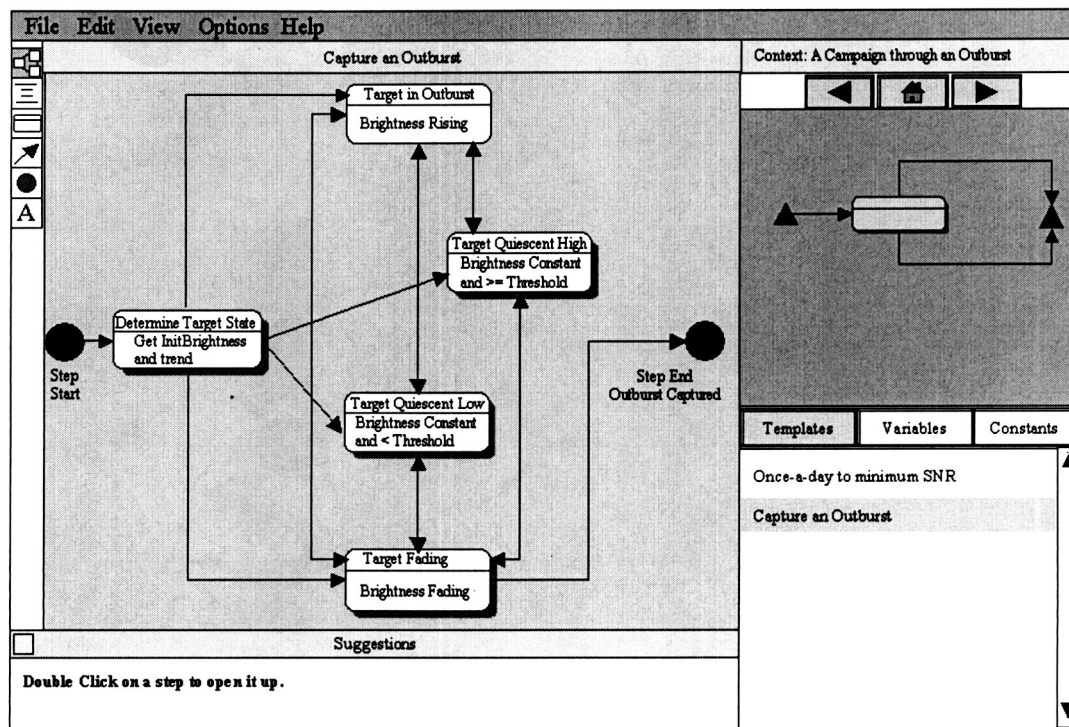


Figure 2 SGM Science Goal Capture Tool

monitoring systems both onboard and on the ground. Analysis could be moved between the two locations depending on the type of analysis and desired response time. For example, simple photon count analysis for a single image could be done onboard, while more complicated analysis that, perhaps, compares the data to existing archives would need to be done on the ground. Further, the system can provide support for the onboard monitor to communicate directly with the ground monitor, allowing the onboard and ground monitors to work in cooperation with each other. The onboard monitor may detect a possible event in a "quick" review of the data stream, and then communicate with the ground monitor where more extensive analyses can be initiated.

5.1.1 Science Goal Capture Tool (SGCT)

The primary purpose of the Science Goal Capture Tool (SGCT) is to capture the user's science goals in an intuitive and easily understandable manner while simultaneously storing them in a format for easy machine processing. To achieve this mix of capabilities we are modeling the SGCT after "state diagrams" commonly used in computer science. State diagrams provide a simple visual way to represent complex behavior focusing on process flow rather than distracting volumes of words.

Our users will express the goals of their observing campaign as a set of hierarchical diagrams which model the conditional flow of exposure sequencing within the user's overall observation strategy. Transitions from one state (or phase of the campaign) to the next are expressed as conditional statements or goals which, when met, indicate the need to change the observing strategy.

We are in the initial prototype stages of this interface and are working with scientists to assess and adjust the system to their intuitive reactions to the interface. An example of an early design overview of the SGCT interface is shown in Figure 2 SGM Science Goal Capture Tool.

Within the SGCT, there are three main areas:

- The **Step Diagram** on the left-hand side is the area of primary focus. The user develops the step diagram here. The Step Diagram area is also used to enter additional information such as the statements to process within the step, or expressions for the exit criteria from a step to some other step.
- The **Context Area** on the upper-right displays a simplified representation of the step diagram and it is used to indicate what portion of the overall diagram is being presented in the step diagram area. Individual steps may be comprised of sub-steps. These steps in turn may be subdivided into further sub-steps. The Context area not only shows the user where he is in the overall diagram, it also acts as a navigational aid. The user can double-click on a step and the step diagram will update to represent what is inside that step.
- The **Templates and Variables Area** is in the lower-right. Scrolling lists of predefined items are presented here. The user can choose from a list and drag the item to the step diagram. This saves the user from having to completely define every item. Complex, predefined template steps can easily be inserted into the user's diagram. Additionally, the area will also show detailed information about step variables. For example, the variable's name, current value, and range of values as appropriate.

Once the diagram is complete, it is saved in an internal format, which can be processed by the Science Goal Analyzer (SGA). Of course, the diagram can be loaded, edited, and resaved at any time.

5.1.2 Science Goal Analyzer (SGA)

The Science Goal Analyzer module will be the primary component that monitors both the observing campaign underway and the state of the campaign goals defined by the user. While usability will be the key in the capture tool, the analyzer will place a premium on efficiency and speed. As exposures are obtained, the SGA will monitor their progress, see if the active goals of the campaign have been met, and make requests of the onboard scheduler for changes to the campaign as defined by the observer. The SGA will also react to messages and events coming from the onboard scheduler. We anticipate that SGA will work in cooperation with an onboard scheduler such as CASPER which is being developed at the Jet Propulsion Lab. SGA will be hierarchical, multi-threaded, and have the ability to adjust its processing priorities to further evaluate interesting finds. In addition, an onboard SGA will work as a side process with the instrument and spacecraft flight software. As such it will not actually control the craft or the instruments, but rather will be able to make requests of the flight software or onboard scheduler (such as end this exposure, or change the time between exposures). The SGA will also monitor messages and data from the instrument/spacecraft flight software and update the status of its monitored campaigns accordingly. This partnership will allow SGM to introduce progressive autonomy and dynamic behavior while the instrument and flight software continue to provide their traditional safety and control checks.

5.1.3 Science Goal Monitor Console (SGMC)

The operations console will provide an interface for an observatory's operations team to view the status of the SGA and oversee the event handler. We currently anticipate that this feature will be used primarily as a testing module to observe and validate the workings of the SGA and also as the "reference implementation" of an event monitor. A 'reference implementation' is an initial fully functional implementation of a standard or interface that provides a feasibility and compatibility baseline for ensuring that a standard or design is feasible.

5.1.4 Event Notification System

An important component of the SGM will be the ability to send alerts to interested scientists and/or operations managers, alerting them when SGM has recognized an interesting event. There are several systems already in existence that can relay events and notify operators. One such tool is the Spacecraft Emergency Response System (SERS) developed at NASA Goddard Space Flight Center. For our initial prototype, we plan to simply route messages to the SGMC. In later stages, we plan to link to SERS and provide a graduated scale of information for the end-user from simple text messages, to partial image downloads and possibly other user-definable information. Our goal is not to reinvent or replace such systems, but rather to interface with these existing systems and use them to notify operators or to send commands to other systems.

5.2 Determine applicability of SGM to other domains

While our initial focus has been astronomical observing campaigns of variable targets, SGM is applicable to a large pool of potential domains where there is a desire to dynamically change what to do next, based on some recognized feature of a current action. For earth science imaging satellites: the ability to recognize that the upcoming suites of targets are cloud-covered and therefore not worthwhile; for solar observatories: the ability to recognize the start of a flare and stay on that target instead of continuing to the next target; for constrained download environments: the ability to perform initial quality analysis on an image and if certain objectives are not met, then download only a highly compressed lower-resolution copy of the image and archive aboard the full image for retrieval only if specifically requested (or for discard after some period of time). Or, vice versa, if a satellite is not scheduled to perform a major data download for some period of time, then an onboard evaluation might trigger a high priority (and high cost) download and/or trigger priority messages to the operations center. Similarly, an initial onboard detection algorithm might trigger a priority request for search of archival data on the ground, helping scientists recognize high-value data quicker than their routine data analysis techniques might achieve.

As data archives become larger, there is a need to search the archives for observations that satisfy the researcher's science goals. In this context, it is essential to be able to capture the science goals and analyze data from the archives to determine which data sets satisfy the researcher's needs. Hence, although the proposed tools are initially focusing on science operations applications, they also have significant potential for use in archival analysis. The value of the tool will have to be evaluated in the context of data mining techniques.

6. IMPACT OF THE SCIENCE GOAL MONITOR

6.1 For Observers: the ability to influence the actual execution of their observations.

Currently, only the observation mechanics are carried into the mission execution. SGM will capture additional information in terms that the scientist will find useful ("Has the desired signal been achieved yet?", "Is the object bright enough to provide reasonable signal?", "Given an anomaly, can the science program be salvaged?"). SGM will also give the observer influence in how the observation is actually executed. It will capture the user's plan for the observation execution, including instructions for what to do in case certain conditions occur. These contingency plans could be triggered when the SGM analyzer detects science events that match the observer's criteria, or can be used in the scheduling of observations. The basic idea is to get the "eye back to the telescope".

6.2 For Observatories: reduced response times to science events, increased science return.

In a resource-constrained environment, where there is minimal observatory staff, SGM can also be used to automatically determine the status and quality of the data by comparing the observations with the desired scientific goals of the program. This is especially useful when data volumes are large and manual checks are not possible. By automating some of these manual tasks, SGM can help contain science operations costs. If we can successfully capture goals and implement onboard scheduling, we can design spacecraft for optimized resource utilization along with maximized scientific returns.

6.3 A step towards spacecraft autonomy

The SGM project is a proof-of-concept to determine if we can effectively and efficiently obtain reliable and relevant data from scientists to make science driven scheduling changes. If successful, this will help maximize the science gained from the new generation of space-based telescopes/observatories and effectively manage scientific priorities. In SGM, there will be new, rapid, flexible, and autonomous approaches to analyzing the quickly growing stream of data. *The tools being developed in the SGM will help to improve the ability to monitor and react to the changing scientific status of observations. Such tools will be enablers for spacecraft autonomy.*

7. CONCLUSIONS: SMART SCIENCE AND SPACECRAFT AUTONOMY

NASA's future plans include constellations, formations, federations, sensor webs etc. All of these will be economically infeasible without a substantial reduction in mission operations staff and costs. One of the proposed strategies to achieve these cost reductions will be to infuse automation technologies into the missions. But missions and scientists are more

culturally and politically averse to risk when it comes to automation. Unless we develop strategies that will help reduce the perceived risk associated with increased use of automation, we will not be able to contain costs.

Furthermore, just automating the spacecraft's technical operations will not intelligently handle the increasing volume of scientific data. We must begin automating both scientific data analysis and reactions to that analysis in a timely and still scientifically valid manner. In other words, we must teach our platforms to dynamically understand, recognize, and react to the scientists' goals. SGM will help progress towards building intelligent spacecraft.

ACKNOWLEDGEMENTS

This work is funded by NASA Code R under the Computing, Information and Communication Technologies (CICT) program.

REFERENCES

- ¹ T. A. Boroson, D. L. Harmer, A. Saha, P. S. Smith, D. W. Willmarth, and D. R. Silva, "The WIYN Queue: Theory Meets Reality", 1998 SPIE 3349 41.
- ² A. Saha, D. Harmer, P. Smith, and D. Willmarth, "NOAO queue-observing experiment on the WIYN Telescope", 2000 SPIE 4010 25.
- ³ P. Massey, M. Guerrieri, and R. Joyce, "The number of publications used as a metric of the NOAO WIYN queue experiment", New Astronomy, Volume 5, Issue 1 (2000), pp. 25-33.
- ⁴ D. Bernard, G. Dorais, et al., "Spacecraft Autonomy Flight Experience: The DS1 Remote Agent Experiment", Proceedings of the AIAA 1999, Albuquerque NM.
- ⁵ G. Prescott, and K. Bhasin, "Ultra-High Data Rate Communications", 2001 NASA Earth Science Enterprise Technology Planning Workshop, available online at: http://nmp.jpl.nasa.gov/workshop-eo4/proceedings/ESE_Wkshp_Comm.pdf.
- ⁶ Next Generation Space Telescope (NGST) Space Segment Single Board Computer (SBC) Preliminary Requirements Document (PRD) Rev. 2a, November 2001.
- ⁷ Isaacs, J. C., NGST Data Volume and Communications Study, February 2001.
- ⁸ S. Chien, N. Muscettola, K. Rajan, B. Smith, and G. Rabideau, "Automated Planning and Scheduling for Goal-Based Autonomous Spacecraft", IEEE Intelligent Systems, pages 50-55, September/October 1998.
- ⁹ S. Chien, R. Knight, A. Stechert, R. Sherwood, and G. Rabideau, "Integrated Planning and Execution for Autonomous Spacecraft," Proceedings of the IEEE Aerospace Conference (IAC), Aspen, CO, March 1999.
- ¹⁰ S. Chien, R. Sherwood, M. Burl, R. Knight, G. Rabideau, B. Engelhardt, A. Davies, P. Zetocha, R. Wainright, P. Klupar, P. Cappelaere, D. Surka, B.C. Williams, R. Greeley, V. Baker and J. Doan, "The Techsat-21 Autonomous Spacecraft Constellation", International Symposium on Artificial Intelligence, Robotics and Automation in Space, St-Hubert, Canada, June 2001.
- ¹¹ E. Miller, "The Semantic Web", WWW2002, available online at: <http://www.w3.org/2002/Talks/www2002-w3ct-swintro-em>.
- ¹² M. Cooter, G. Davis, J. Mackey, and M. Rackley, "Alternative Approaches to Automation at NASA's Goddard Space Flight Center", AAAI Press, 2001.

-
- ¹³ L. Maks, J. Breed, and M. Rackley, "Current Level of Automation at NASA's Goddard Space Flight Center", AAAI Press, 2001.
- ¹⁴ J. Jones, C. Burkhardt, M. Fishman, S. Grosvenor, A. Koratkar, L. Ruley, K. Wolf, "Lessons Learned from the Scientist's Expert Assistant Project", Proceedings of SPIE Vol. 4010, pp 107-117.
- ¹⁵ A. Koratkar, C. Burkhardt, M. Fishman, S. Grosvenor, J. Jones, R. Lucas, L. Ruley, K. Wolf, "NGST's Scientist's Expert Assistant: Evaluation Results", Proceedings of SPIE Vol. 4010, pp 225-230.