Engineering Data Summaries for Space Missions

Michael A. Swartwout
Space Systems Development Laboratory
Durand 250
496 Lomita Mall
Stanford, CA 94305-4035
650/725-6794 (rcktboy@leland.stanford.edu)

Abstract— New paradigms in spacecraft design are leading to radical changes in spacecraft operations. Increased constraints on resource usage and greater focus on operations costs require new approaches. One such method, beacon-based health monitoring, automates the task of routine health monitoring and migrates the process from the ground to the spacecraft.

The performance of this automated method is further improved by a supplemental approach that monitors the long-term health of the spacecraft. Called "engineering data summarization", this process has the responsibility of creating an on-board summary of the spacecraft state of health, tracking notable sensor values and trends as appropriate. Every few weeks, the summary is transmitted to ground operators. The purpose of the summary is to provide operators with context about the spacecraft's state.

Building on the systems for spacecraft operations being developed at Stanford's Space Systems Development Laboratory, this paper is the first step in developing a methodical study of engineering data summarization. Five simple solutions are proposed, and each is examined using newly-established, competitive metrics. Analysis of the solutions' characteristics leads to a definition of the problem's "solution space." These results point to the next steps needed for a thorough characterization of the engineering data summarization problem.

TABLE OF CONTENTS

1. INTRODUCTION

Two pressures – performance and cost – are pushing a revolution in methods of spacecraft mission operations. The advent of large, interdependent constellations like global satellite communication networks increases the scope of operations by an order of magnitude, creating system-wide effects not present in single-spacecraft missions. The scale and nature of these effects make traditional operator-intensive solutions intractable. New constraints on the Deep Space Network greatly increase the cost of communicating with deep-space vehicles, especially during the "idle" years of cruise phase. For these and other situations, the need is not merely to lower costs to make the mission more competitive; methods must be developed to enable the mission to be accomplished at all.

One proposed approach to improve spacecraft operations is beacon-based health monitoring, also called "beacon monitoring." Since more than 90% of all health assessment contacts do not require any response by the operators [1], automating the process of anomaly detection can significantly reduce cost. In this method, the spacecraft analyzes its own sensor data to assess its state of health; this abstracted state is broadcast in the form of a low-power, low-bandwidth beacon. Human operators are involved only to respond to abnormal conditions; the state-of-health beacon assists by indicating the neeeded changes in operational procedures.

There are, however, some long-term drawbacks to this approach. One of the primary goals of beacon monitoring is to reduce the amount of data sent to the ground, which is achieved by eliminating the download of telemetry data. But that telemetry set is used in other tasks as well. Operators gain intuition about the performance and characteristics of each spacecraft and component by examining the real-time telemetry and through simulations with the exhaustive data archive. They develop informal heuristics for troubleshooting a spacecraft and have learned to distinguish "quirks" from malfunctions. If the operators do not have some means of developing detailed understanding about the spacecraft, they cannot adequately fulfill their duties.

Therefore, in order to fully obtain the benefits of the beacon monitoring method, the "fast loop" of real-time health assessment must be supplemented by a "slow loop" to study the long-term behavior of the spacecraft. Moreover, there are operator tasks other than health management which would benefit from added information about the vehicle. Examples of these other responsibilities are command verification and understanding third-party payload usage.

The supplement to beacon-based health monitoring is generically called "engineering data summarization", or "Summary." This term is intended to encompass the family of implementations whereby the spacecraft creates a second set of abstractions about the sensor telemetry; this information is sent back to the ground to provide context for operators. The form of the data, the amount, and the frequency of summary downloads are all variables to be optimized, resulting in the least cost – in terms of communication bandwidth and operator effort – while maintaining the performance margins of operator-intensive missions.

This paper will outline the concept of engineering data summarization. It is intended to be the first step in the development of a methodology for long-term automation of spacecraft mission operations. Section 2 will define the problem, emphasizing the initial assumptions and constraints used in this scope. Particular attention will be paid to clearly defining the roles and relationships between beacon-based health monitoring and engineering data summarization. In Section 3, five strawman Summary solutions are proposed; the purpose of these solutions is to identify important issues in the Summary problem. Analysis of these solutions helps develop the performance metrics as described in Section 4. These assessments of effective solutions are based on competitive measures of quality, cost, and timeliness. In Section 5, study of the strawman approaches leads to the definition of the solution space. Section 6 provides conclusions of the work thus far and describes the next steps for developing the methodology.

2. PROBLEM STATEMENT

Because engineering data summarization is a fairly recent concept [2], the scope and nature of the problem is still not fully defined. Therefore, it is important to describe both what Summary is and what it is not. Specifically, engineering data summarization must be distinguished from beacon-based health monitoring.

Beacon-Based Health Monitoring

Beacon monitoring is, essentially, two migrations of the spacecraft anomaly detection task: the responsibility for routine health monitoring shifts from an operator to an automated process, and this process is moved from the ground to the vehicle. The mission saves operator man-hours and communications bandwidth because of each respective migration. Where there was once a continual data stream requiring teams of operators to analyze, beacon monitoring creates a small (on the order of a few bits) signal that requires operator attention only in the event of an anomaly.

Motivated by the opportunity to cut the cost of operations, beacon-based health monitoring is being studied for several projects. It is an essential element of JPL's proposed new Pluto/Europa/Sun missions [2], a technology demonstration for the New Millenium Program's Deep Space 1 mission, and has been proposed for the Air Force Satellite Control Network [3]. In 1998, Stanford University's SAPPHIRE microsatellite [4] will conduct validation experiments of this concept for JPL [5, 6].

These projects have also determined that operators need more information about the vehicle than what the small, simple health assessment flag can provide. It is not a matter of recording the full telemetry sets once an anomaly is detected, though that may be an important part of fault isolation and recovery. Instead, operators need additional information about the state of the spacecraft *before* the anomaly, so that they can perform all of their tasks. This additional information is the Summary.

Table 1 describes the differences between beacon-based health monitoring and the Summary. While both are essentially operational tools that convert on-board data into information useable by operators – in fact, one could argue that beacon monitoring is nothing more than a specific implementation of the Summary – other factors associated with implementation make them worth distinguishing. For example, the primary task of beacon monitoring is to fulfill the responsibility of health monitoring by providing an instantaneous indication of a need for a change in operations. By contrast, the Summary is intended to enhance the performance of the operators (and therefore the vehicle) within any given mode. A useful but limited analogy is to liken the Beacon to the function of health monitoring and the Summary to the function of examining the telemetry archive. Each utilizes the vehicle data for related, yet distinct, purposes.

Table 1 Comparison of Beacon and Summary

Element	Beacon Monitoring	Engineering Data Summarization
Time	Fast	Slow
Output Size	Very Small	(Undetermined)
Input	Telemetry	Telemetry
Assessment	Abstracted State	Abstracted States Processed Data Raw Data
Role in Operations	Indicates a Change in Operational Modes	Assists Effective Performance of Operators and Vehicle within a Mode
Analogous Function	Health Monitoring	Consulting Telemetry Archive

The purpose of this study is to develop a methodology for engineering data summarization in space missions. Research in beacon-based health monitoring is further discussed in the aforementioned references. While beacon monitoring is an important precursor to an effective Summary, the subjects are distinct enough to allow for separate investigation.

Example Scenarios

For a definition of the engineering data summarization problem, it is helpful to create two examples, drawn from two general classes of space missions. The first is an Earth-orbiting constellation of spacecraft, which involves issues of large-scale, complex systems, rapid response times, and high-performance payloads. The second is a single deep-space vehicle, which involves issues of long communications delays and robust performance. Between these two examples, most of the crucial needs of data summarization can be identified.

Scenario 1: Global constellation – Assume that a space constellation of several dozen communications satellites has implemented beacon monitoring. Given the usual variations in manufacturing quality and the complexity of these spacecraft, it is not surprising that the heat pipes for the batteries of Vehicle 28 perform slightly worse than expected, and this variation was not detected before launch. Thus, during the season when the orbital plane brings more sunlight onto that region of the spacecraft, the batteries are slightly hotter than average. The difference is small – certainly below the limits defined for an abnormal condition – and thus is not detected on-orbit, either.

When the nearby power regulator starts registering dangerously high temperatures, the beacon system detects the anomaly. The operator called to investigate has never directly operated this spacecraft before, and because the quirky heat pipes have not been identified, she may believe that the regulator is affecting battery temperature. This misunderstanding could lead to unneeded changes in operations to protect the battery, and at best confuses the task of isolating the heat effects. The operator would be greatly aided by the ability to know that the battery quirk is a long-standing phenomena and is probably not related to the overheated regulator. In other words, the chances for mission success would be greatly enhanced by providing some means of reproducing the long-term trend analysis (as used in "typical" operations of today) on the ground.

Scenario 2: Deep-Space Mission –Pluto Express is in its seventh year of a ten-year cruise to the outermost planet. Everyone originally involved in the development and check-out of the spacecraft has retired or been reassigned. The beacon-based health monitoring registers an alarm, and the operator on call discovers that one of the propulsion tanks has unusually low pressure. Since a leaking tank would require profound changes to future operations, it is imperative to determine if this is a true leak or a problem with the sensor. The operator would be greatly assisted in his tasks by the ability to look at the past history of the sensor, especially during maneuvers where general sensor performance could be predicted, and over the previous few days to identify performances characteristic of a leak.

Granted, these two examples are simplistic and the problems addressed could possibly be solved by better operator training, redundant sensors and spacecraft check-out procedures. But that is precisely the point: long-term functionality of a mission using beacon-based health monitoring requires additional operational procedures and possibly changes to the spacecraft architecture. Automated health monitoring alone cannot account for long-term vehicle health.

Still, it is unlikely that simple adjustments to operational procedures will ensure adequate performance. For example, putting tighter bounds on limit-checking, to catch the small discrepancies, leads to a higher false alarm rate with commensurately higher operating costs. Accurate, detailed models require additional efforts to maintain and update and are susceptible to unmodeled or unobservable inputs. Additional sensors cost mass, power, computational ability and operator effort. Regular telemetry downloads increases the cost of communication bandwidth.

While changes such as those listed above may indeed be elements of the best low-cost mission operations scenario, it is imperative to establish a careful methodology. Current trends in spacecraft operations laud the benefits of automation, but automation requires clear methods for proper implementation. The spacecraft operations business has, over the last forty years, invented many "rules of thumb" and informal heuristics. These heuristics often do not translate well into automated methods. In order to create effective automated solutions, these informal ideas must be formalized.

Assumptions

Because this is only a preliminary investigation into the nature of the engineering data summarization problem, or Summary Problem, it is helpful to reduce the study's scope. Several key assumptions will assist in creating a topic that, while not complete, exhibits the most important characteristics of the real-world scenarios.

It is assumed that beacon-based health monitoring has been included in the spacecraft specifications and design, fulfilling the functions defined earlier in this section. And while it is expected that the exact parameters of this beacon system would be optimized to fit the specific mission, some basic characteristics can be assumed.

(1) The beacon monitoring system performs automated health monitoring of the spacecraft, informing operators in a health state change that requires an operational procedures change.
(2) The beacon monitoring system is not required to perform more advanced autonomous functions, such as fault isolation and correction. Such functions may be incorporated into more advanced Summary techniques.
(3) The engineering data summarization element of the spacecraft and the beacon-based health monitoring system have access to the same reasoning and analysis capabilities.

The point of assumption (3) is to emphasize the existence of strong coupling in functions – and therefore designs – between the beacon monitoring and Summary systems. As will become evident, a vehicle using enhanced methods in health monitoring should also apply those methods to the Summary; similarly, a highly-capable Summary solution will allow for a more capable Beacon system.

(4) Human operators will be involved in the more complicated and/or unexpected functions of spacecraft operations.

Whether or not operators are involved – and to what degree they are involved – greatly affects the type of Summary solution created. The presence of human operators is one of the driving factors behind the Summary. Of course, it may eventually be feasible to completely eliminate humans from health monitoring and other operations functions. But given the reliability and capabilities of present-day space systems, it is more reasonable to assume that automated systems will face problems they are unable to handle, requiring the intervention of operators.

(5) The role of operators is limited to that of health management.

This assumption helps to clarify (4), and limits the scope of the study to the following operator tasks: identifying anomalies, isolating their sources, assessing the impact of the anomaly, taking action to recover from faults, and altering the understanding of the system based on the new information. The Summary does not perform these tasks, but provides the operator with information that was once directly available from vehicle telemetry.

(6) The communications equipment used to transmit the Summary will be no more capable than that used by "normal" spacecraft for telemetry downlinks.

This assumption emphasizes that the purpose of this study is not to rely on new technologies to send more information in less time, but to discover systems that enable operators to do more with less information. Mission designers will always push the communication capabilities to their limits; rather than seeking ways to create and manipulate more data, it is important to seek ways to take advantage of existing but underutilized knowledge of the system.

The Summary Problem

The general statement of the Summary Problem is this: **Engineering data summarization is to provide the necessary information for operators to carry out their tasks, while minimizing both the efforts of these operators and communications resources**. Granted, this statement is quite vague on what constitutes "necessary information," and definition of "minimal" is equally nebulous. The study of metrics in Section 4 will explore the latter subject; the rest of this section explains the former.

As highlighted in the examples, the operators need to be provided with context and history about the vehicle. However, they must glean this context from a reduced – or "summarized" – set of information. From the standpoint of contingency operations, therefore, the fundamental goal of engineering summarization is to be able to reproduce information about the vehicle for use by operators while improving performance metrics such as communications cost and operator efforts.

There are many methods to accomplish this goal; determining which method is "best" is done using the metrics described in a later section. Note that a Summary solution need not recreate the *full* telemetry set from a reduced supply of information; it must provide that information which operators need to fulfill their responsibilities. While effective solutions may indeed involve the ability to wholly or partially rebuild sensor data history, this is not an expected.

3. INITIAL SOLUTIONS

As an aid in understanding the nature of the solution space, and in order to provide starting points for developing solutions to the summary problem, five strawman solutions have been proposed. Each places exaggerated emphasis on one or more characteristics; the intent is to identify how each of the parameters affects the performance of the overall solution. Again, the emphasis of this study is not to propose viable Summary solutions, but to better understand the nature of the problem itself. The characteristics of each candidate are presented in Table 2, along with how they perform for the space mission examples.

Reduced Sample Rate

One of the simplest methods of summarization is to downlink a partial telemetry set. If the ground-based communications system has a maximum data downlink capability, then the solution is to prioritize the sensor information and transmit the maximum amount. The duration between transmits is altered to fall just within the specifications. The system adds no additional modeling or sensors beyond that which would normally be on board. The information saved for download would be what is considered most important: short bursts around interesting events and then snapshots spaced out to fill in the rest of the available time.

This approach is very simple to implement. However, one shortcoming is that it requires advance understanding of the spacecraft and an expectation of what are likely to be the most important events. Moreover, it does not reduce the cost or training of operators; this approach assumes they will continue to familiarize themselves with the reduced data set as was done with the complete set. And, this method does not consider that important information may be lost during the no-sample periods.

For example, in the Earth-orbiting constellation, operators could designate the basic kinds of information about each vehicle that helps them to perform their tasks. All vehicle Summaries could be configured to emphasize the relevant parameters in their download. It is not at all clear, however, that the warm battery problem of the example would be spotted in this selective Summary process – unless battery temperatures were already known to be unusually important.

Still, this approach highlights the fact that sensors have varied importance for different phases of operations. Some sensors convey more critical information than others, and some events where many items are changing at once, such

Table 2 Use of Initial Summary Solutions for Example Problems

Solution	Deep Space	Earth constellation
Reduced Sample Rate	Helps prioritize data Requires significant bandwidth May "lose" pressure sensor trend	Helps identify system-wide issues Requires significant bandwidth May "lose" battery history
On-Board Database	Hampered by long delays Helped by "Anticipating Summary" Requires very robust memory "Finds" pressure sensor trend	Helped by short transmit times Requires robust memory "Finds" battery history
Statistical Summary	Hampered by small "sample size" Relies on expectations May "lose" pressure sensor trend	Helped by large "sample size" "Finds" battery history
Alarm Threshold Summary	Very low bandwidth Requires accurate thresholds "Loses" pressure sensor trend	Very low bandwidth Requires accurate thresholds "Loses" battery history
Model-Following Summary	Helped by "stable" models Helped by very low bandwidth "Finds" pressure sensor trend	Hampered by many varying models "Finds" battery history

as a thruster firing, may require more study than a week of dormant cruise. Effective summary solutions will assess the relative worth of spacecraft sensors in order to select the most essential information.

On-Board Database

Another straightforward approach is to store all telemetry data on board the spacecraft. The telemetry archive is not lost; it is maintained in spacecraft memory instead of on the ground. The operators can request any information they want, but otherwise no data is ever sent to Earth.

While at first this approach seems ludicrous because it requires immeasurable amounts of on-board memory and assumes that operators will have the time to retrieve desired data, it does have some helpful ideas. It emphasizes the usefulness of contingency-generated Summaries that pay close attention to component(s) with anomalous conditions. But instead of waiting for a user request, the vehicle should anticipate the demand for additional information about these suspect parts, search its memory for related information and take additional samples to clarify the anomaly. It would thus prepare an initial report, with background information, for use by the operators in troubleshooting.

This candidate seems to be more appropriate for an Earth-orbiting mission where a spacecraft database can be rapidly and repeatedly queried. When the transmit time can be on the order of hours, though, it makes little sense to send requests for information. In the "anticipating Summary" case, it would be very useful for the deep-space vehicle to search for relevant information to download while the ground team is being assembled. If the pressure sensor was giving an anomalous reading, the Summary could call up its time-history and prepare it for download.

Statistical Summary

A slightly more involved approach is to perform statistical analyses on the sensor outputs. The time-history of sensor data can be examined for such items as trends, frequency spectra, and averages. In order to make use of this information, it is necessary to have expected performances of these sensors. However, the generation of expectations and comparisons can all be performed on the ground, where the baselines can be easily updated as new information becomes available. This method emphasizes the usefulness of expectations; understanding about the vehicle leads to predictions about its behavior, which can be used to identify both quirks and faults.

While promising to be a capable solution, using techniques that are well-understood in the scientific and engineering communities, this approach does not take advantage of the full information available. The relationships between components are completely ignored. For both the Earth-orbiting and deep-space missions, the use of statistical summaries might catch the battery and pressure sensor problems – but, especially in the case of the pressure sensor, actually making use of its statistics requires additional understanding about the vehicle in order to create reasonable expectations. On the other hand, in a constellation there are many spacecraft, which automatically increases the "sample size" for the statistics.

While such information may not be necessary for every effective Summary, it is worth investigating more model-based approaches to see if additional performance benefits or cost savings are possible. (Alternately, a statistical-based method that uses causal models, such as the one developed by Doyle [7], could be used. Basic information about how one component affects another is assembled into a relationship tree; statistics are kept on the responses of each sensor and this information is used to predict which components are behaving properly.)

Alarm Threshold Summary

Since most modern spacecraft already employ alarm thresholds for health monitoring, it makes sense to explore this technique for the Summary. A basic model of the spacecraft is created, with state- and mode-based abnormal and emergency alarm thresholds defined. Those sensor outputs which fall outside the normal limits would be stored for Summary download. Great savings in bandwidth would result, since only the "abnormal" information is kept. Like the statistics approach, the primary shortcoming in this method is the loss of valuable information about the relationship between components. One component within normal limits may point to a problem with a component that is out of limits, but the former data is ignored. In fact, for both examples, the vital information about the components would be "lost" because it was not in the abnormal range.

The lessons learned from this method are the importance of saving information that may seem to be "normal", and also the emphasis on using the already-assembled body of knowledge to simplify the approach.

Model-Following Summary

A more complex approach that promises great reduction in downloaded data is a model-based summary. Detailed input-output transfer functions are kept for every component of interest; identical models exist in the vehicle computer and on the ground. Given the inputs to the system, very precise outputs for each component are generated, and this information is compared to the sensor data. Only those points that reflect significant deviation from the model are stored, along with an the system configuration and inputs at the time of the discrepancy. Thus, the deviations from expected behavior and the context of the situation are explicitly identified and stored. Using this information, the ground-based model can faithfully simulate the performance of any component on board the spacecraft.

For a deep-space mission, especially during cruise when the external environment and vehicle modes are changing quite slowly, this approach has special appeal. The models of the spacecraft and its environment are stable over the long term, allowing for detailed simulations to be developed. The pressure sensor problem would be readily diagnosed, since its behavior would have been very closely followed. And the Summary need only pay attention to the discrepancies in the model – a relatively small amount of data to store for download.

The limitations of this system are related to the complexity of the modeling approach. Inaccurate models or unobservable conditions lead to holes in the summary, and effort is required to maintain the models as they change over time. Also, the ability to recreate the entire telemetry set probably indicates that excessive, unusable data will be downloaded.

4. PERFORMANCE METRICS

As demonstrated by the examples, most of the strawman methods are viable Summary techniques. With a little work, some sort of reasonable solution could be achieved. However, some options seem to be more appropriate, more effective, than others. A methodical approach to choosing the most effective solution requires the introduction of performance metrics.

Performance metrics should reflect the overall goals of the Summary Problem. Additionally, they should favor the solutions that are "competitive" in engineering practice [8], that is, those that increase the quality of the product, lower the overall cost, and reduce the time needed to produce the product. Summary metrics are presented in Table 3. Not surprisingly, these sets of goals are often in conflict: increase in quality comes at an added cost. How to choose between metrics is discussed following the metric descriptions.

Quality Measures

The "quality measure" is a reflection of the usefulness of the Summary. The fundamental issue of quality is whether or not an operator is able to acquire the information he needs. Determining a quantitative measure for total summary quality has proved elusive; at the least, it can be measured by comparing the performance of anomaly management tasks using the complete data, and then using the Summary. High-quality Summaries will duplicate the results of procedures performed with the original telemetry.

A related, though indirect, measure of quality is the effort required to use the summary. Since the man-hours spent on operations is more appropriate as a measure of *cost*, an applicable *quality* measure is the level of training. High-quality Summary solutions would not require extensive operator training. An admittedly inadequate means to quantify the level of training is the operator's salary, since that indirectly reflects his or her training and background.

A third quality measure is how effectively the system responds to anomalies. Solutions that perform well will reduce or eliminate the loss of payload operations due to uncorrected faults, because the summary helps the operator to efficiently recover the vehicle. One way of measuring this response is to keep track of the lost opportunities to perform payload operations because the spacecraft is not available. The Summary approach that results in the least vehicle "down time" is of highest quality.

As a whole, the quality measures apply similarly to both classes of Summary problems. Constellations want to reduce operator training because of the number of operators and spacecraft involved; deep space missions cannot afford to retain skilled operators during the years in cruise phase. For deep-space programs, down time relates completely to the phase of the mission, because obviously the loss of performance during a flyby is much greater than during cruise. A commercial Earth-orbiter, by contrast, relies on high-performance payloads to turn a profit; down time is extremely costly.

Time Measures

In a product development context, the "time measure" is an indication of how quickly the product can be delivered to the market; in the "faster, cheaper, better" spacecraft

Table 3 Engineering Data Summarization Performance Metrics

Type	Description	Metric	Measurements	Deep Space	Earth-Orbiting
Quality	Usefulness	*Operator Confidence*	*???*	Vital	Vital
	Operator Effort	Operator Training	*Salary*	Important	Important
	Down Time	Spacecraft Down Time	Lost bits/sec	Conditionally Important	Vital
Time	Time to Implement	Time to Check-Out	Man-hours	Desirable	Important
Cost	Operator Resources	Operator Man-Hours	Average hours/week	Vital	Important
	Communication Resources	Length of Contact Data Downloaded Time Between Contacts	Carraway and Squibb's Metric	Vital	Important
	Spacecraft Resources	On-Board Storage On-Board Processing Sensor Requirements	Bytes Bits/sec Price, Mass, Power, Size	Desirable	Desirable
	Ground Resources	Cost of Equipment	Dollars	Unimportant	Important

context, the time measure is an indication of how quickly the Summary service is available, that is, how quickly the Summary system could be implemented and functional in a mission operations concept. This is measured in the time taken to develop models, perform simulations, and carry out initial orbit checkouts to confirm that the system is working properly. Obviously, those systems which can be rapidly implemented and validated have better time measures than those that cannot. Missions with long checkout and/or cruise phases are less concerned with time measures, compared to a commercial payload.

Cost Measures

The two major "cost measures" involved in the Summary Problem are the time it takes for operators to perform their tasks, and the communications resources that are utilized. The former is rather straightforward to measure; the average total number of man-hours per week the operators devote to health management tasks indicates the effort required. This metric is particularly useful since managing the cost of operations was a fundamental motivation for establishing beacon monitoring, and thus Summaries. It is important for constellations to control the number and cost of operators. It is extremely important for a deep-space mission to slash the cost of operating the spacecraft during the long idle phases.

The communications resources cost measure is more difficult to quantify, since it includes several factors. The true cost of using a communications network is hard to determine since the systems are often shared between many programs and involve both fixed and recurring costs. However, the total cost of communications is a reflection of the amount of data being received, the duration of the contact, and how often the contacts occur. Since all three of these elements are easy to measure, a metric composed of them would be a good indication of communications cost. It is assumed that other factors – such as payload data – will be the drivers in defining communications hardware; the Summary will use whatever is already available to the vehicle. This assumption is a driving factor in developing deep-space Summaries, since communication involves long delay times and expensive equipment. It is less important, though still important, to an Earth-orbiting constellation that wishes to conserve its resources.

Carraway and Squibb [9] have proposed a set of metrics to indicate the level of autonomy for spacecraft. One metric in particular, the Spacecraft Engineering Analysis, provides an effective means of combining the above cost measures:

$$\begin{array}{c} Spacecraft \\ Engineering \quad Analysis \end{array} = \left(\frac{NT}{DWN \times WORK}\right) \times \left(\frac{NT}{T}\right) \tag{1}$$

Where T = duration of track (hours)
NT = time between tracks (hours)
DWN = data downloaded (bits)
WORK = operator effort (man-hours)

This metric promises to be quite useful, since it penalizes large downloads, operator effort, and the duration of contacts. It also doubly rewards long no-contact periods, which reflects the overhead involved with scheduling and pre-calibration for a contact. In other words a few, long-duration communication passes are better than several shorter ones.

A third cost measure reflects the requirements a Summary technique levies on other vehicle subsystems. This question is somewhat harder to quantify, since Summary will often be an enabling technology – meaning that the mission *must* perform a Summary or the mission cannot be accomplished. If a particular solution requires extensive computing power, then that becomes a vehicle requirement. When comparing between viable methods, however, those solutions which cause the least impact to the other subsystems while providing similar communication and operator performance are clearly superior. Standard measurements include data throughput, data storage, and the price, mass, power, and sampling requirements of the sensors.

Similarly hidden in the background of the Summary Problem is the use of the ground equipment such as computers to store and further analyze the Summary data.

Like vehicle impact, this issue is somewhat clouded by the fact that Summary is required, and thus the equipment is required. Often, the ground equipment is fixed, pre-existing infrastructure and so the issue is not what equipment to choose, but what can be done with what is available. The main measurement for this parameter will be the cost involved with creating and maintaining the equipment.

Using the Metrics

Given the specific goals for each mission, these individual performance metrics are weighted to reflect the most important issues. For example, a specific deep-space missions may emphasize the cost of the communication link because of the effort involved to receive weak signals. But a global communications network will be very concerned about the response time since the loss of payload functionality means a loss of business. These metrics can be ranked in priority, or combined into a weighted sum, to provide one general measurement of system performance.

Of course, the driving factor in developing the beacon-based health monitoring problem – which in turn drives the Summary Problem – is cost. Summary solutions which maintain the level of cost savings promised by beacon monitoring will be given overwhelmingly priority. Therefore, in general, the two most important metrics are operator cost and communications cost.

Since this paper is intended to investigate the major issues of the Summary Problem, it is impossible to rank the metrics with any greater precision. Instead, the solutions that are proposed will be measured up against each individual metric. Furthermore, true quantitative measurements have not yet been established, and thus the solutions in this paper must be qualitatively examined.

5. SOLUTION ELEMENTS

Having made some initial attempts to create candidate solutions, and having established the major metrics to evaluate performance, it is now possible to develop the "solution space" of the problem. As defined, there are two main components of a Summary solution: the data that is generated, and the manner in which it is sent to the ground. These categories are further subdivided into the main characteristics that can be adjusted to solve the problem. The major, trade-level elements of the Summary Problem are listed in Table 4.

Transmission Characteristics

The main questions about the data transmission are "How much data?" and "How often is it sent?" Each of these questions are directly reflected in the metrics. As explained below, the "How much?" question is a function of the Summary techniques used by the spacecraft. It can be effectively ignored for the purposes of developing transmission schemes, since the impact to the communications subsystem will be small compared to that required for "normal" telemetry downloads.

The "How often?" question, then, is the fundamental issue in the transmission aspect. In development of the Summary Problem, it was assumed that very infrequent – on the order of weeks – summaries was the best plan. This may not be true; regular, very short, daily updates may in fact provide better performance. Whatever the result, the

Table 4 Engineering Data Summary Solution Elements

Component	Elements and Sub-Elements		Description	Examples	Costs
Data Transmission	**No-Contact Duration**		How long to wait between downloads	Never Weekly Monthly On-Demand	On-Board Storage
Data Generation	**On-Board Model**		How the spacecraft estimates its behavior	None Yellow/Red Limits State-Based Limits Transfer Functions	Throughput Reasoning Capability
	Sensor Selection		What information the system provides about itself	None Standard Sensors Commands	Sensor Requirements Data Sampling
	Summary Technique	**Sample Rate**	How often to poll each sensor	Every Minute Event-Based	On-Board Storage Data Transmit Rate Throughput Reasoning Capability
		Comparison	How to relate model and sensor outputs	Alarm Thresholds Error Bounds	
		Manipulation	Number-crunching the data	Statistics	
		Abstraction	Reasoning about the data	Abstracted States	
		Data Selection	What data gets sent to the ground	Nothing Raw Telemetry Abstracted States	

parameter to be traded is the amount of time between data downloads. This parameter is especially important because it directly contributes to the communication cost metric.

As explained in assumption (6), other aspects of the data transmission are not expected to significantly contribute to the performance of a Summary solution and are ignored for the purposes of this study. Since the metrics favor vast reductions in data size, not just transmission speed, such factors such as transmitter frequency, data rates, and similar characteristics are not considered.

Summary Generation Characteristics

The engineering data summary is the product of three elements. Only a few elements may be part of a particular solution, but every solution will typically have all to some degree. First, a set of sensors provides input to the system. These inputs are used in an on-board model to predict expected behavior. The sensor data and model predictions are then compared using a summary technique, which may perform additional processing. The output of the summary technique is the data to be downloaded.

Sensor Selection – This element describes the information available to the Summary system. In general, this refers to the numbers, types, and locations of the sensors used by the vehicle. Often, the spacecraft sensors are determined and positioned according to the requirements of other spacecraft subsystems, and thus the Summary method may not the ability to make trades. However, it is confidently expected that the ability to choose and place telemetry sensors will greatly enhance a Summary's performance, and thus sensor selection will become a necessary element of effective Summary solutions.

In addition to "normal" sensors, such as temperature, voltage, and current, it may prove helpful to include other information about the vehicle, such as the inputs from cameras, the available CPU memory, and state information such as what components are active at a given time.

The decisions made in sensor selection directly impact the spacecraft design in terms of adding hardware; each sensor has mass, volume, power, thermal and data handling requirements associated with it. Adding non-traditional sensors may also require added computational capabilities.

On-Board Model – This element describes the ability of the vehicle to determine the detailed state of each of its components. From a Summary standpoint, this parameter reflects the level of on-board reasoning and modeling. For example, the standard practice of defining "abnormal" and "emergency" alarm thresholds for sensor outputs is one form of modeling; the output is abstracted into five ranges (dangerously low, abnormally low, normal, abnormally high, and dangerously high), where the boundaries are based on an understanding of the expected output. A slightly more sophisticated on-board model would be a set of adjustable limits, altered for different modes – such as sunlight and eclipse – and changed over time to reflect natural environmental degradation. A much more sophisticated model would consist of a set of transfer functions for each component that predicted outputs, given various inputs and the current state of the vehicle.

It must be pointed out that "sophisticated" does not imply "better". Increasing the quality of one parameter does not necessarily indicate a more optimal Summary solution. Increasing the complexity of the on-board model impacts the rest of the spacecraft, which may affect the performance or feasibility of other designs. For example, detailed on-board models require highly capable computers to handle the throughput and reasoning requirements. On-board models also require varying degrees of data storage. The issue of "better" is resolved using the metrics.

Summary Technique – The actual data manipulation is called the Summary technique; the information provided by sensors and the model are compared and transformed in order to create the Summary. There are several key parameters to the Summary technique, all of which are highly coupled. Their functions are distinct enough to be described separately, even if altering one sub-element causes profound changes in the others.

One of the basic issues to be addressed is the *sample rate* of the sensed data. This decision, which may differ for each sensor, indicates the amount of information available to the Summary. It also affects the ways in which data is used to create the Summary.

Another question is what to do with the information generated by the on-board model. *Comparisons* should be made between the model and the sensor outputs, but what sorts? The answer relies greatly on the nature of the model; abnormal and emergency alarm thresholds are designed for specific comparisons, but for a transfer function model, it must be determined to what degree the model output and the sensor output should match.

Also, given the model outputs and the sensor outputs, there are many data *manipulations* that can be performed. Statistics such as extrema, averages, and frequency spectra can indicate whether a component is behaving as expected. A simple curve fit could be derived to compress a time-history into a few parameters. This area is one of the most-explored in engineering data summarization, as the tools are readily available and much of the work done for anomaly isolation applies.

In addition, there may be *abstractions* to be made about this data. Perhaps the performance of a component or a subsystem would be summarized in a single parameter.

Finally, once all this information has been generated, it is necessary for a *data selection* process to determine which of the information is saved. All or part of the data created by the other processes is put into long-term spacecraft memory for eventual download.

In all, Summary techniques contribute to the cost of operations in several ways, most notably the amount of data that must be generated. In addition, design of these elements levy requirements on the spacecraft processing capabilities and reasoning abilities.

Shortcomings to This Method

The primary shortcoming in this approach for defining the elements of the Summary problem is that the boundaries between the elements are not clearly distinguishable. For example, the existence of abnormal and emergency alarm thresholds is considered part of the on-board model, but use of the thresholds is part of the comparison sub-element of the Summary technique. This distinction is quite arbitrary. This problem is not simply a matter of choosing better categories; the complex nature of the Summary problem defies simple classification. However, these categories do provide an effective means of creating more workable sub-problems, an import first step in methodically addressing the Summary Problem.

6. CONCLUSIONS & FUTURE WORK

Beacon-based health monitoring extends the capabilities of spacecraft mission operations and promises to significantly reduce the cost of operations. However, in order to fully realize these costs, additional work is needed to ensure that in the long term, operators retain understanding about the vehicles. Engineering data summarization is a necessary part of beacon monitoring, giving operators the context they need about the specific components of each spacecraft.

Work has begun on a methodology for creating effective Summary approaches. Quality, timeliness, and cost measures have been established, with particular emphasis on the cost of using human operators and communications resources. The necessary elements of a Summary solution have been identified and further subdivided into their composite parts.

As shown in Table 5, five "strawman" solutions have been proposed, classified according to their solution space elements, and related according to their metrics. Actual evaluation of these solutions would depend on the specific qualities of the mission in question.

As mentioned above, this paper is the start of an ongoing study of engineering data summarization. The next step is to better quantify the metrics involved, and to apply these metrics to the initial candidates. Once the relationships are established between solution parameters and the various metrics, true trades can be performed to identify which elements of the Summary solution space are most vital to effective solutions. It will be insightful to use information from existing and proposed programs to generate real metrics.

Another look is warranted into how the solution space has been partitioned, since the coupling and nebulous distinctions between different elements impedes the ability to perform trade studies. Perhaps a few more candidates could help identify the key, independent elements.

Given those developments, true trade studies could be performed on candidate solutions for engineering data summarization. Test cases must be created to provide realistic examination of various Summary approaches. Stanford's SAPPHIRE microsatellite and its "Al Wood" engineering prototype are candidates for true operational studies.

More work is also needed on the initial problem statement. Are these all the tasks operators perform? Is the ability to distinguish quirks from faults the only real requirement of a Summary? These answers will be found through interviews with spacecraft operators from a number of organizations.

Another assumption worth challenging is the performance of new communications technologies. Perhaps something like a laser communications link, used sparingly, could restore much of the telemetry archive, thereby alleviating much of the Summary Problem.

Table 5 Comparison of Initial Summary Solutions Using Qualitative Cost Metrics

Solution	Significant Elements	Favorable Metrics	Unfavorable Metrics
Reduced Sample Rate	No-Contact Duration Sensor Selection Sample Rate Data Selection	On-Board Storage On-Board Processing Sensor Requirements	Operator Man-Hours Length of Contact Data Downloaded
On-Board Database	No-Contact Duration	Time Between Contacts Data Downloaded	Operator Man-Hours Length of Contact On-Board Storage
Statistical Summary	Manipulation	Length of Contact Data Downloaded	Operator Man-Hours
Alarm Threshold Summary	On-Board Model Comparison	On-Board Storage Time Between Contacts Data Downloaded Length of Contact	Operator Man-Hours
Model-Following Summary	On-Board Model Comparison Selection	Operator Man-Hours Length of Contact Data Downloaded	On-Board Processing Cost of Equipment

Finally, this paper assumes that most of the Summary is performed on board the spacecraft. But is there any benefit to a second Summary performed on the ground? Perhaps some of the long-term trends and model corrections could be automatically performed by ground systems using the Summaries downlinked over several months.

ACKNOWLEDGEMENTS

The author would like to thank his friends and colleagues in the Space Systems Development Laboratory, especially the past and current members of the SAPPHIRE development team. Special thanks goes to Christopher Kitts for his tireless devotion to making clear sense out of these incoherent ideas, and to Professor Robert Twiggs for his continued efforts to keep this laboratory moving. This work was performed in partial satisfaction of graduate studies at Stanford University.

REFERENCES

[1] Mike Farmer and Randy Culver, "The Challenges of Low-Cost, Automated Satellite Operations," *Proceedings of the 31st International Telemetering Conference*, Las Vegas, NV, October 30-November 2, 1995

[2] E. Jay Wyatt and John B. Carraway, "Beacon Monitoring Approach to Operations," *Reducing the Cost of Spacecraft Ground Systems and Operations* , Oxfordshire, U.K., September 27-29, 1995.

[3] Schultz, Michael, "Lifeline: A Concept for Automated Satellite Supervision," *Proceedings of the Software Technology for Satellite Autonomy Workshop*, June 22-25, 1993.

[4] Richard A. Lu, Tanya A. Olsen, and Michael A. Swartwout, “Building ‘Smaller, Cheaper, Faster’ Satellites Within the Constraints of an Academic Environment,” *Proceedings of the 9th Annual AIAA/USU Conference on Small Satellites*, September 19-22, 1995.

[5] Michael A. Swartwout and Christopher A. Kitts "A Beacon Monitoring System for Automated Fault Management Operations," *Proceedings of the Tenth Annual AIAA/USU Small Satellite Conference*, September 16-19, 1996.

[6] Michael A. Swartwout and Christopher A. Kitts, "Automated Health Operations for the SAPPHIRE Spacecraft," *ITC/USA '97: Proceedings of 33rd Annual International Telemetering Conference*, October 30-November 1, 1997.

[7] Richard J. Doyle, "Determining the Loci of Anomalies Using Minimal Causal Models," *International Joint Conference on Artificial Intelligence*, August 1995.

[8] Christopher A. Kitts, "The ASSET Client Interface: Balancing High Level Specification with Low Level Control," *Proceedings of the 1998 IEEE Aerospace Conference*, March 21-28, 1998.

[9] John B. Carraway and Gael F. Squibb, "Autonomy Metrics," *Fourth International Symposium on Space Mission Operations and Ground Data Systems*, September 16-20, 1996.

***Michael Swartwout** is a Doctoral Candidate in Aeronautics & Astronautics at Stanford University. His research is in the area of engineering data summary. He is involved in SSDL's Automated Space Systems Experimental Testbed (ASSET) operations research system, developing the system architecture and managing the blackboard tool that forms the basis of ASSET. He is also the Project Leader for SAPPHIRE, SSDL's first student-built spacecraft. On SAPPHIRE, he is co-investigator of the Beacon-Based Health Monitoring Experiment. He earned a BS in Aeronautical & Astronautical Engineering from the University of Illinois at Urbana-Champaign in 1991, and a MS in the same, from the same, in 1992.*