Focused Real-time Systems Monitoring Based on Multiple Anomaly Models

Richard J. Doyle, Steve A. Chien, Usama M. Fayyad and **E. Jay Wyatt**
AI Group, M/S 525-3660
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109-8099, U.S.A.
{rdoyle,chien,fayyad,wyatt}@aig.jpl.nasa.gov

Abstract

In real-time monitoring situations, more information is not necessarily better. When faced with complex emergency situations, operators can experience information overload and a compromising of their ability to react quickly and correctly. We describe an approach to focusing operator attention in real-time systems monitoring based on a set of empirical and model-based measures for detecting different kinds of anomalies and for determining the relative importance of sensor data. This approach has been evaluated on data from the life support system testbed of Space Station Freedom.

1 Introduction: Sensor Selection

Mission Operations personnel within NASA have the task of determining, from moment to moment, whether a space platform is exhibiting behavior which is any way anomalous, which could disrupt the operation of the platform, and in the worst case, represent a loss of ability to achieve mission goals. A traditional technique for assisting mission operators in space platform health analysis is the establishment of alarm thresholds for sensors, typically indexed by operating mode, which summarize which ranges of sensor values imply the existence of anomalies. However, experienced mission operators reason about more than alarm threshold crossings to detect anomalies: they may ask whether a sensor is behaving differently than in the past, whether a current behavior may lead to a global perturbation or whether a current behavior may lead to – the particular bane of operators – a rapidly developing alarm sequence.

A fault which propagates through a system faster than the sensor polling rate can create a situation where, between one sampling and the next, the number of sensors in alarm goes from zero to tens or more. Information about the ordering of events is lost. In this kind of emergency situation, operators can experience information overload and a compromising of their ability to interpret sensor data.

Our approach to introducing automation into real-time systems monitoring is based on two observations: 1) mission operators employ multiple methods for recognizing anomalies, and 2) mission operators do not and should not interpret all sensor data all of the time. The subject of this paper is an approach to determining from moment to moment which subset of the available sensor data for a system is most informative about the presence of, or potential for, anomalies occurring within the system. We term this process *sensor selection* and we have implemented a prototype selective monitoring system called SELMON [Doyle et al 89, Doyle et al 92a, Doyle et al 92b].

The SELMON system has its origins in a sensor planning system called GRIPE [Doyle et al 86] which plans information gathering activities to verify the execution or robot task plans. Other model-based monitoring systems include Dvorak's MIMIC, which performs robust discrepancy detection for continuous dynamic systems [Dvorak & Kuipers 89, Dvorak 92], and DeCoste's DATMI, which infers system! states from incomplete sensor data [DeCoste 90]. The SELMON work complements other work within NASA which has focused on empirical and model-based methods for fault diagnosis of aerospace platforms [Abbott 88, Muratore et al 89, Scarl et al 88].

2 Approach: Sensor Ordering

How does an intelligent agent, observing a complex system, decide when something went wrong? To quantify this notion, we have developed an approach to focusing operator attention in real-time monitoring. This approach involves defining a set of sensor importance measures. Each of these measures embodies a different viewpoint on why, at a particular moment, one sensor may be more worthy of operator attention than others. The measures are based on concepts from model-based reasoning, statistics, and information theory. Some of the measures utilize sensor value predictions generated by simulating a causal model of the system being monitored.

During each timestep all sensors are scored according to these measures. The scores are used as a basis for an ordering on the sensors. See Figure 1. These scoring measures are divided into two categories. The first set – empirical methods – rely upon current and historical data to determine importance. These measures include *surprise*, *alarm*, *anticipate alarm*, and *value change*. The second set uses a causal model of the system and a simulation capability to reason about expected current and future system performance to determine sensor importance. These methods include *deviation*, *sensitivity*, and *cascading alarms*.

After describing each of these measures, we describe how these measures are combined into an overall importance score for each sensor.

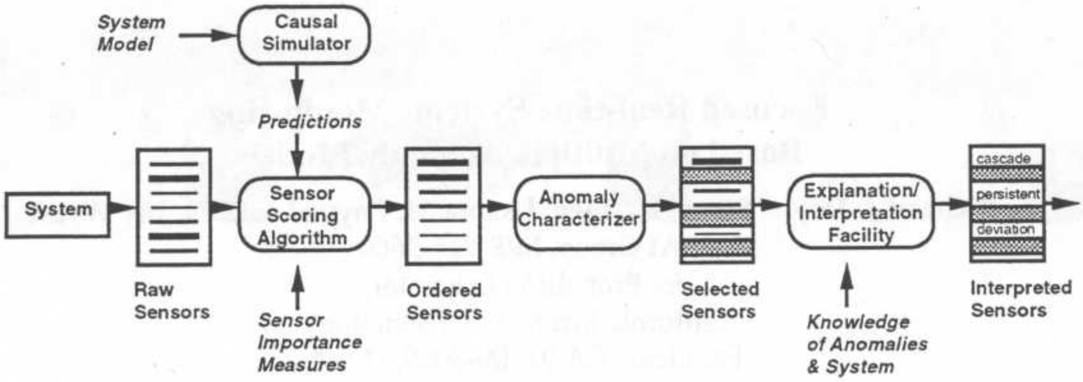

Figure 1: SELMON Architecture.

2.1 Empirical Sensor Scoring

In this section, we describe the empirical measures that are used in determining the overall importance score assigned to each sensor. This part of the score is based on four measures: *surprise*, *alarm*, *anticipate alarm*, and *value change*. These measures use knowledge about each individual sensor, independently of any knowledge about the interconnectedness of the sensors.

Surprise

In order to obtain an ordering on the set of sensors, we need to quantify the following notions: How reliable is a sensor? How stable is it? How often does it go into an alarm state?

From an information theoretic point of view, a change in the value of a sensor gives us a certain amount of information (usually measured in bits) about the system state. Assume we have two sensors, S_A and S_B. Further assume that sensor S_A's value has been wildly changing over the last 100 readings, while sensor S_B's value has been constant. If we are told that according to the latest update, the values of both sensors have increased by 25%, which do we consider a more informative event? Clearly the fact that S_B's value changed is more informative since it is more unusual. Prior to the latest reading, if we were asked to predict the values of S_A and S_B, then based on previous data, we would naturally guess that S_A's value is likely to have changed while S_B's value is likely to have remained constant. Then the fact that S_B changed value tells us something that we did not know or expect.

For each sensor, a cumulative histogram of its values is maintained for each system operating mode. This is done by dividing its range into a fixed number of bins. The boundaries between bins are determined through specific knowledge of the sensor and of the "interesting" subranges in its range. This histogram is then used to determine two measures of the interestingness of the most recent value returned by a sensor.

Denote the range of sensor S by $Range(S)$. If S is a continuous-valued sensor, we can discretize its range into a set of mutually exclusive ranges $\{R_1(S), R_2(S), \ldots, R_K(S)\}$, where

$$Range(S) = \bigcup_{i=1}^{K} R_i(S)$$

With each range $R_i(S)$ we associate a frequency measure $f_i(S)$ that gives the proportion of time that S's value has been in this range. Thus $f_i(S)$ is an estimate of the probability of the value of S falling in range $R_i(S)$ and

$$\sum_{i=1}^{K} f_i(S) = 1$$

To quantify the degree to which sensor S is stable in its reading, we apply the notion of information entropy. The entropy of the values of a sensor S, denoted by $VEntropy(S)$, is defined by

$$VEntropy(S) = -\sum_{i=1}^{K} f_i(S) \cdot \log f_i(S)$$

where $VEntropy(S)$ is maximum when all ranges of values of S are equally likely. It is minimum when the values of S have all been in one range $R_i(S)$, thus $f_i(S) = 1$ (for some i, $1 \leq i \leq K$). It can easily be shown that $0 \leq VEntropy(S) \leq \log K$. We are now ready to define the average value informativeness of sensor S, denoted by $VInform(S)$, to be

$$VInform(S) = 1 - \frac{VEntropy(S)}{\log K}$$

where $VInform(S)$ takes on values between 0 and 1. A value of 1 indicates that S normally rarely changes its value, while a value of 0 indicates that S's value is equally likely to be in any of its ranges.

On the other hand, the quantity

$$VUnusual(S) = 1 - f_i(S)$$

gives the unusualness of sensor S's value being in the i-th bin. $VUnusual(S)$ is computed each time S reports a value, and the i used is the index of the bin containing the reported value. This measure can assign the same degree of unusualness in fundamentally different situations. For instance, it does not distinguish between a value having a probability of $\frac{1}{K}$ occurring when all other values have an equal probability of $\frac{1}{K}$ each, and a value with probability $\frac{1}{K}$ when only one other value has probability $(1 - \frac{1}{K})$ with the remaining values having probability 0. In the first case, the value is just as likely as any other. In the second case, the interesting event is that the most likely value did not occur. To make this distinction we combine the unusualness and value entropy measures to obtain the *surprise* score:

$$Surprise(S) = VInform(S) \cdot VUnusual(S).$$

This measure takes on the maximum value of 1 when one bin in the histogram has probability one and the sensor registers a value in another bin. It has a minimum value of zero when all bins in the histogram are equally likely.

Accounting for Alarm Thresholds

Alarm thresholds for sensors, indexed by operating mode, typically are established through an off line analysis of system design. SELMON makes use of alarm threshold information in the following way: A sensor whose value traverses the safety threshold is said to go into a state of alarm. The predicate $In\_Alarm(S)$ captures this notion:

$$In\_Alarm(S) = \begin{cases} 1 & \text{if } S \text{ is outside its safety range} \\ 0 & \text{if } S \text{ is within its safety range} \end{cases}$$

We compute the value of an alarm score for S as follows:

$$Al\_Score(S) = In\_Alarm(S) \cdot \left[1 + Trav(S)\right].$$

where $Trav(S)$ is the proportion of the alarm range traversed.

We consider alarms as interesting events whose importance decreases with time. Thus a sensor that persists in alarm state for prolonged periods of time should gradually fade from our attention. To achieve this we add an exponential decay factor. Let $t_A(S)$ be the time at which sensor S last entered into alarm. At any time t, the alarm score is computed as follows:

$$Alarm\_Score(S) = \frac{1}{2} Al\_Score(S) e^{-\beta(t - t_A(S))}$$

where $\beta > 0$ is the time decay constant. β is chosen small so the decay will not be too fast; typically $\beta \leq 0.1/\text{second}$.

Given the recent values of S, one may conduct a simple form of trend analysis to decide whether or not sensor S is anticipated to be in alarm soon. The measure $Predict\_Alarm(S)$ is a curve-fitting prediction of when the sensor will enter alarm. This measure has a minimum of 1 and a maximum of infinity if the curve fit indicates that the sensor will never enter alarm. If the sensor is currently in alarm, $Predict\_Alarm(S)$ measures when the sensor is predicted to leave alarm. This measure is used to compute a score *Anticipate Alarm* as follows:

$$Anticipate\_Alarm(S) = \begin{cases} 1/Predict\_Alarm \\ 1 - 1/Predict\_Alarm \end{cases}$$

The first case applies when S is within its safety range. The second case applies when S is outside its safety range.

By examining this definition, the reader may determine how the boundary cases of immediate alarm threshold crossings or indefinite persistence in the nominal or an alarm range are handled.

Quantifying Value Change

A change in the value of a sensor is considered to be an event of interest. The surprise measure described above measures the degree of interestingness of a sensor taking on a certain value. Another aspect of sensor behavior to measure is the most recent change in value of the sensor that brought it to its current reading. However, absolute change magnitude is not interesting in and of itself. What is interesting is the probability of the most recent change taking place. Hence we need a scheme for normalizing the absolute change in value of a sensor.

The scheme we use assigns a score to each change in the value of a sensor that is an estimate of the proportion of all previous value changes for that sensor that had value changes strictly less than the change under consideration. Suppose we get a change in value of the sensor equal to Δ. Furthermore, suppose that 60% of the previous value changes for this sensor in the current operating mode have been less than Δ. In this case, we assign a score of 0.6 to the change Δ. Changes with magnitude greater than Δ will get higher scores.

This scheme seems to require that we keep track of a sorted sequence of all value changes of each sensor. This is neither feasible nor necessary. An approximation of this value can be obtained by keeping a constant number of values, say W, in a sorted sequence. Let the total number of changes in the values of a sensor so far be $C(S)$. Rather than storing all $C(S)$ values, we store only $W < C(S)$ values. With the arrival of a new change in value for sensor S, we increment the count of changes $C(S)$ and then we decide whether to replace one of the W values we are storing or simply ignore the current value change. The decision criterion is to generate a random number in [0, 1] according to a uniform distribution, and replace one of the W values if and only if that random number is less than $\frac{W}{C(S)}$. It can be proven that this algorithm is equivalent to one that stores all $C(S)$ values, randomly samples W of them, and returns as score the proportion of the W elements that have value less than the change under consideration.

We call this score the *percentile value change* score. It is used to assign a normalized score in the range [0, 1] for each change of value that occurs in each sensor. By definition, this score is maximum when the change is the maximum change of value seen so far for a particular sensor. It is minimum when no change occurs in the value of a sensor.

2.2 Model-Based Measures

SELMON also uses a model of the monitored system to determine sensor importance. This model is used to compute three scores: deviation, sensitivity, and cascading alarms. This section describes how each of these scores is computed. Simulation is performed by a discrete-event based simulator which operates on causal models with quantitative behavior models [Rouquette et al 92].

Deviation

The deviation measure uses a model of the monitored system to make predictions of expected current sensor readings. The concept of the deviation score is that sensor readings deviating significantly from the predicted values are anomalous and should be reported to the operator.

The deviation score is computed in the following manner. First, the raw deviation is computed as the difference between the predicted and observed sensor scores. This raw deviation is entered into a normalization process identical to that used for the value change score, and the resultant score in the range [0,1] is the overall deviation score.

2.2.1 Model-Based Analysis

The SELMON system also uses the causal model of the monitored system to reason about future effects of current quantity changes. These future effects are considered in two causal-based measures. First, *sensitivity* measures the effect of predicted changes in quantities on the overall state of the system. This is done by projecting each predicted change in a quantity individually forward as a perturbation of the system, and measuring the overall change in the system. Those currently occurring changes which have a greater effect on the future state of the system are likely to be more important and thus receive high scores to be displayed to the operators. The second causal reasoning measure is *cascading alarms*, which measures the potential for observed changes to result in rapidly developing alarm sequences. The cascading alarms measure uses the same perturbation analysis used in the sensitivity analysis and measures the number of alarms triggered and how quickly alarms occur. Those predicted changes which are expected to trigger large numbers of alarms are scored highly and thus will be selected to be displayed to operators.

Sensitivity Analysis

Sensitivity analysis measures the sensitivity of other quantities in the monitored system to changes in each quantity in the model. This is performed as follows. Beginning with a simulation of the system in its current state and time $T_{current}$, simulate forward one timestep (i.e. until the next time sensors are expected to be polled). For each quantity Q, choose ΔQ_{pred} as the current 50th percentile value change recorded for the given sensor.

Then, for each quantity Q, run a simulation beginning again with the current system state, perturbing Q by ΔQ_{pred}, propagating this change to other quantities in $All\_Quantities$ (the set of all quantities in the model) as dictated by the model. For each such changed quantity Q' in $All\_Quantities$, for each time $time$ that the quantity changes during the simulation, collect a sensitivity score proportional to the amount of change in Q' normalized to the size of the nominal range of the sensor but also modified by a decreasing function of $time$. This calculation captures the characteristic that delayed and less direct effects are more likely to be controllable and less likely to occur. Thus, a change which affected a quantity Q' but occurred slowly is considered less important. This simulation proceeds for a predetermined amount of simulated time. Then, for each changed quantity Q', take the maximum of the collected $change\_scores$ for that quantity. The sensitivity score for Q is the sum of these maximums for all the Q's. Thus, for each quantity Q, a simulated change produces a set of $change\_scores$ for each other quantity in the model. The sensitivity score for Q is the sum of the respective maximums of each of these sets. If there are no changes to a quantity, this set is empty and the quantity receives a zero score.

Cascading Alarms Analysis

Cascading alarms analysis measures the potential for change in a single quantity to cause a large number of alarm states to occur, thus causing information overload and confusion for operators. In the cascading alarms score, the same simulation used in the sensitivity score computation is used to also determine the number of alarms triggered by the observed change. In the cascading alarms score, for each quantity Q, the number of alarms triggered by a perturbation of Q by ΔQ_{pred} is computed.

The alarm count is then normalized for the total number of possible alarms and the weight of each alarm state triggered is also decreased as a function of the time delay from the initial change event to the alarm. This has the effect of focussing this measure on quickly developing cascading alarm sequences which are the most difficult to interpret and diagnose.

2.3 Computing a Total Sensor Score

We use the *surprise* score to modulate the percentile *value change* associated with a sensor. This accounts for the unusualness of a sensor value as well as the change in the sensor value that brought it to its current reading. The percentile value change score is also used to modulate the scores obtained by the causal analysis of the system: the *sensitivity* score and the *cascading alarms* score. These are modulated by the percentile *value change* because they are computed based on an analysis of the effect of a perturbation in the value of the sensor on the overall system. The remainder of the score combinations are simple sums. See Figure 2.

3 Application Domain

Our application domain is the hardware testbed of the water side of the Environmental Control and Life Support System (ECLSS) for Space Station Freedom. The water side of ECLSS consists of three principal systems: Multifiltration (MF), Vapor Compression and Distillation (VCD), and the Volatile Removal Assembly (VRA). Using a combination of analysis of system description documents, consultation with testbed engineers, and actual hardware testbed data, we have constructed models of these subsystems. Each subsystem model contains 30-50 quantities and 15-30 mechanisms. Work in elaborating fault models is ongoing. This model has been validated by comparison against actual data from the subsystem testbed at the Marshall Space Flight Center (MSFC) in Huntsville, Alabama. We are also in the process of extending our model to cover the ECLSS air side subsystems.

4 Performance Evaluation

The output of the SELMON algorithms is dynamically computed each time the sensors are polled. SELMON produces a total ordering on the set of sensors according to the sensor importance measures outlined above. In order to assess whether SELMON usefully focuses operator attention, we assessed SELMON sensor ordering in the light of critical sensor subsets specified by an ECLSS domain expert as useful in understanding episodes of anomalous behavior in actual historical data from ECLSS testbed operations.

In one experiment, we asked the specific question: How often did SELMON place a "critical" sensor in the top half of the sensor ordering? The performance of a random sensor selection algorithm would be expected to be 50.0%. Table I shows the results of our experiment. The first column identifies one of the episodes specified by the domain expert. The second column identifies the number of timesteps in the episode. The third column shows the overall "hit" rate for that episode: the percentage of time SELMON placed the given sensor in the top half of the sensor ordering based on the total

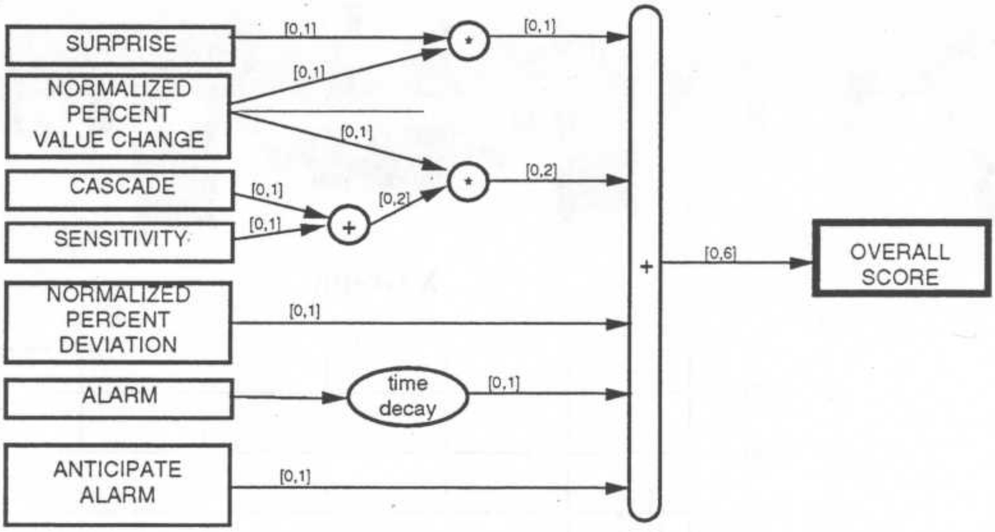

Figure 2: SELMON Sensor Scoring Algorithm.

sensor score. The fourth column reports performance using the alarm score component only. Finally, the fifth column reports performance based on using the maximum SELMON sensor score component only.

EPISODE	Timesteps	Total	Alarm	Max
High Flow Rate	91	94.0	98.2	83.8
Sensor Malfunction	80	100.0	100.0	100.0
Unibed Loading	530	56.6	48.9	90.5
Pre-Heater Off	328	100.0	100.0	81.1
Emergency Shutdown	9	100.0	100.0	100.0
Pressure Fluctuations	71	100.0	99.2	89.8
High Pressure	67	95.6	92.0	90.7
All	1176	79.7	76.3	87.9

Table I: Performance at selecting critical sensor data.

These results suggest that SELMON performs much better than random at replicating the attention focusing of an ECLSS domain expert. In addition, when only the maximum score component is reported, SELMON performs significantly better than random and considerably better than alarm thresholds only. This result suggests that the most effective monitoring system is one which incorporates several models of anomalous behavior. This supports our approach of taking multiple views of sensor importance.

SELMON is intended to assist operators in efficient anomaly detection – the first step towards diagnosis. Another planned experiment will investigate how sensor selection supports diagnostic reasoning. In addition to the ECLSS subsystem models which describe nominal behavior, a number of ECLSS fault models have been developed. After implementing a diagnostic reasoning algorithm, we will determine how this algorithm performs at correctly diagnosing faults from behavior traces resulting from simulation of these fault models. We will then test the performance of the diagnostic reasoning algorithm when it is given SELMON-selected sensor data. Finally, we will test the performance of this algorithm when it is given the same number of sensor data randomly selected. Some degradation of performance may occur in the diagnostic reasoning algorithm using SELMON-selected data. A measure of success will be a significantly greater loss of performance with randomly selected data. Such a result will suggest that SELMON selects sensor data relevant to diagnostic reasoning.

5 Discussion

In this section, we elaborate on the experimental results reported in the previous section by making general observations about the utility of the SELMON approach while citing a specific example of how SELMON has reacted to anomalous behavior episodes. We also discuss areas for improvement in the current realization of the SELMON concept.

The best, most experienced mission operators are remarkably effective at knowing when something is going wrong on a space platform using only traditional anomaly detection techniques such as alarm thresholds. Our aim is to offer a more complete, more robust set of techniques for anomaly detection, based on multiple models of what constitutes an anomaly. These techniques will make mission operators even more effective, or can be the basis of an automated monitoring capability. The following example illustrates how SELMON highlights a subtle manifestation of an anomaly which the traditional alarm threshold approach fails to detect.

During an episode when the ECLSS pre-heater failed, system pressure, which normally oscillates within a known range, became more stable. This "abnormally normal" behavior is not detected by alarm thresholds because the system pressure remains firmly in the nominal range. However, the SELMON informativeness measure rises during such an episode. Informativeness rises when the frequency distribution across the range of values for a sensor departs from a flat distribution. A suddenly stable system pressure causes one of the value ranges for system pressure to begin to dominate the frequency distribution. See Figures 3 and 4. SELMON provides the means of detecting and reasoning about this kind of subtle anomaly.

SELMON will be brought closer to an operational monitoring system by adopting known techniques for sensor data preprocessing. Currently, we perform only the most rudimentary

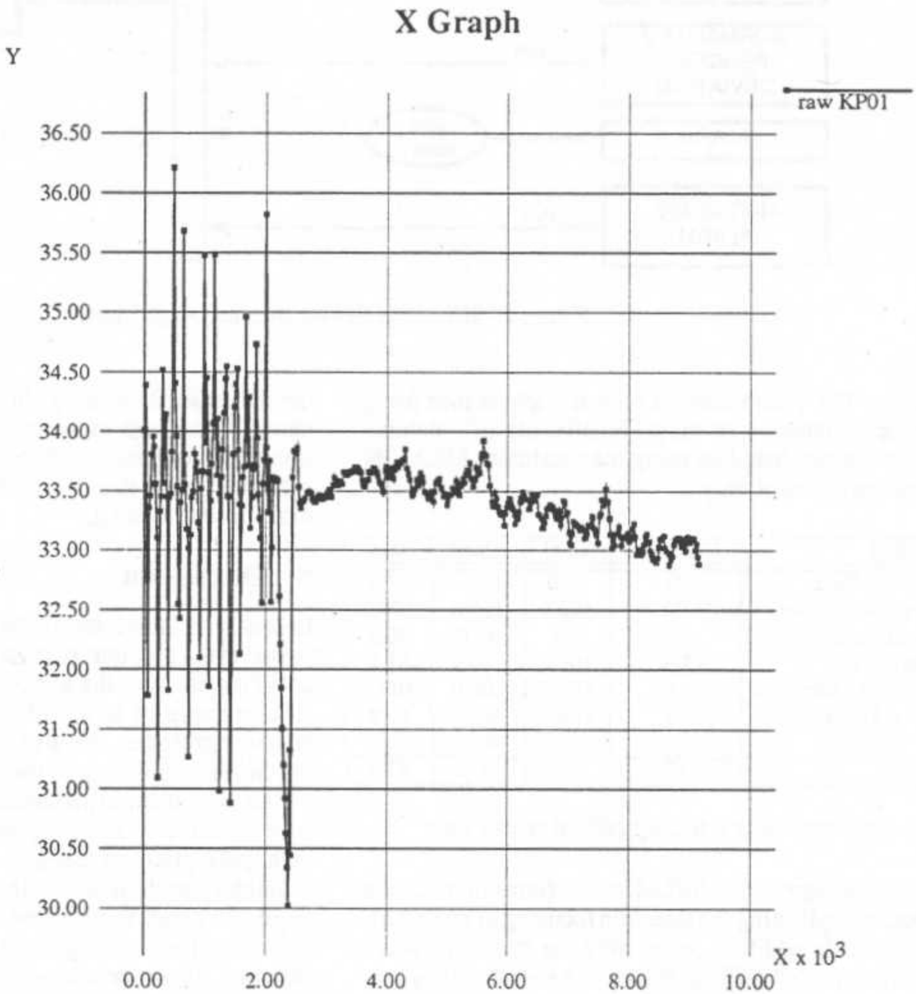

Figure 3: ECLSS System Pressure During Abnormally Normal Episode.

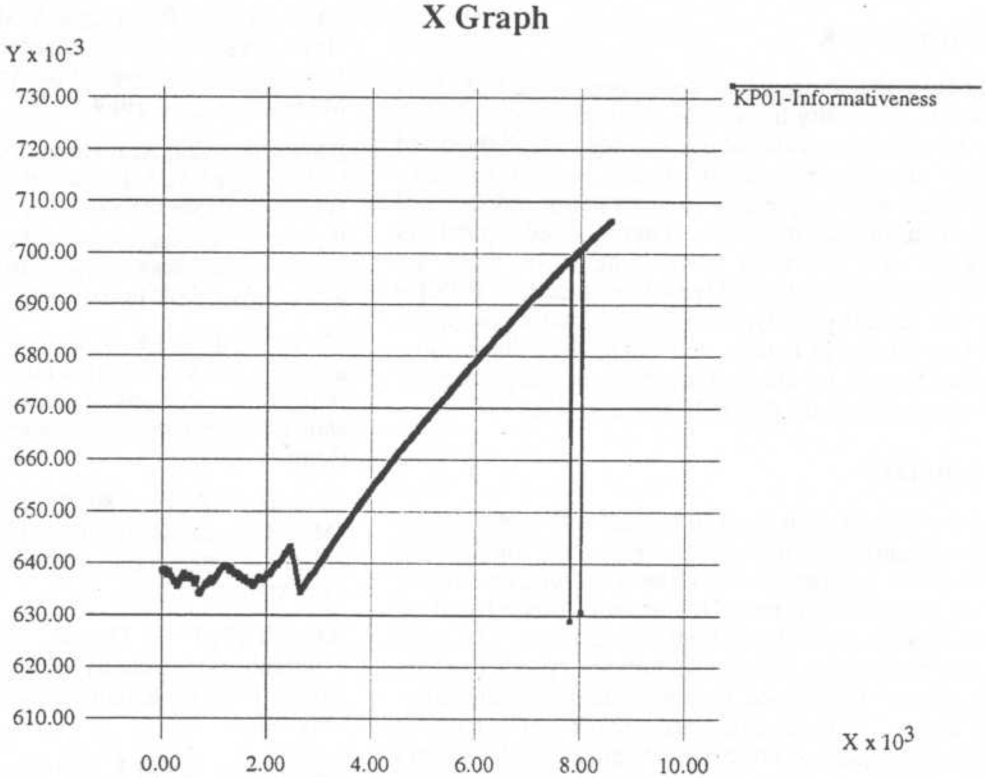

Figure 4: Informativeness of ECLSS System Pressure during Abnormally Normal Episode.

filtering of raw data. Since some of the ECLSS system parameters do display oscillatory behavior, proven techniques for suppressing fluctuations in the time domain which are manifestations of well-defined frequency components can be applied.

An unresolved area in SELMON is developing a well-founded method for utilizing all the individual sensor importance component scores. Our working concept is to compose the individual measures into a total sensor importance score. A theoretical or empirical analysis may provide insight as to the most appropriate composition function and weighting scheme.

6 Future Work

We recognize that an important component of the SELMON approach is the ability to provide explanations or interpretations of why a particular sensor has been highlighted and is worthy of operator attention. Future work in the SELMON project will complement existing sensor ordering and anomaly detection capabilities with model-based capabilities for characterizing anomalies, partitioning sensors which are causally related, and selecting sensors according to multiple viewpoints (causal priority, proximity to control points, potential for irreversible damage, and magnitude of anomaly).

In related work, we are also investigating the problem of sensor placement during design [Chien et al 91].

7 Summary

We are developing techniques to support real-time monitoring through sensor selection, the moment to moment focusing of attention on a subset of the available sensor data. Sensor selection is based on a set of importance criteria based on different models of what constitutes an anomaly. The computational realizations of these importance criteria draw on concepts from model-based reasoning, statistics, and information theory. Experimental results show that our sensor selection techniques are effective at highlighting the sensors deemed critical by a domain expert for understanding actual anomalous episodes from the Space Station Freddom ECLSS life support system testbed. These results also suggest that a monitoring system which employs multiple models of anomalous behavior is more effective than one based on the concept of alarm thresholds only.

8 Acknowledgements

Others who have worked recently on the SELMON project include Prasanta Bose, Leonard Charest, Harry Porta and Nicolas Rouquette.

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[Abbott 88] K. H. Abbott, "Robust Operative Diagnosis as Problem Solving in a Hypothesis Space," *7th National Conference on Artificial Intelligence*, St. Paul, Minnesota, August 1988.

[Chien et al 91] S. A Chien, R. J. Doyle, and N. F. Rouquette, "A Model-based Reasoning Approach to Sensor Placement for Diagnosability," *2nd International Workshop on the Principles of Diagnosis*, Milan, Italy, October 1991.

[DeCoste 90] D. DeCoste, "Dynamic Across-Time Measurement Interpretation," *8th National Conference on Artificial Intelligence*, Boston, Massachusetts, August 1990.

[Doyle et al 86] R. J. Doyle, D. J. Atkinson, and R. S. Doshi, "Generating Perception Requests and Expectations to Verify the Execution of Plans," *5th National Conference on Artificial Intelligence*, Philadelphia, Pennsylvania, August 1986.

[Doyle et al 89] R. J. Doyle, S. M. Sellers, and D. J. Atkinson, "Predictive Monitoring Based on Causal Simulation," *11th International Conference on Artificial Intelligence*, Detroit, Michigan, August 1989.

[Doyle et al 92a] R. J. Doyle, D. Berleant, L. K. Charest, Jr., U. M. Fayyad, L. S. Homem de Mello, H. J. Porta, and M. D. Wiesmeyer, "Sensor Selection in Complex Systems Monitoring Using Information Quantification and Causal Reasoning," in *Recent Advances in Qualitative Physics*, B. Faltings and P. Struss (eds.), MIT Press, 1992.

[Doyle et al 92b] R. J. Doyle, S. A. Chien, U. M. Fayyad, and H. J. Porta, "Attention Focusing and Anomaly Detection in Real-time Systems Monitoring'," *3rd International Workshop on the Principles of Diagnosis*, Eastsound, Washington, October 1992.

[Dvorak & Kuipers 89] D. L. Dvorak and B. J. Kuipers, "Model-Based Monitoring of Dynamic Systems," *11th International Conference on Artificial Intelligence*, Detroit, Michigan, August 1989.

[Dvorak 92] D. L. Dvorak, "Monitoring and Diagnosis of Continuous Dynamic Systems using Semiquantitative Simulation," Ph.D. dissertation, The University of Texas at Austin, May 1992.

[Muratore et al 89] J. Muratore, T. Heindel, T. Murphy, A. Rasmussen, and R. McFarland, "Space Shuttle Telemetry Monitoring by Expert Systems in Mission Control," in *Innovative Applications of Artificial Intelligence*, H. Schorr and A. Rappaport (eds.), AAAI Press, 1989.

[Rouquette et al 92] N. F. Rouquette, S. A. Chien, and L. K. Charest, Jr., "Event-driven Simulation in SELMON: An Overview of EDSE," JPL Publication 92-23, August 1992.

[Scarl et al 88] E. A. Scarl, J. R. Jamieson, and E. New, "Deriving Fault Location and Control from a Functional Model," *3rd IEEE Symposium on Intelligent Control*, Arlington, Virginia, 1988.