Attention Focusing and Anomaly Detection in Systems Monitoring

Richard J. Doyle
Artificial Intelligence Group
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109-8099
rdoyle@aig.jpl.nasa.gov

Abstract

Any attempt to introduce automation into the monitoring of complex physical systems must start from a robust anomaly detection capability. This task is far from straightforward, for a single definition of what constitutes an anomaly is difficult to come by. In addition, to make the monitoring process efficient, and to avoid the potential for information overload on human operators, attention focusing must also be addressed. When an anomaly occurs, more often than not several sensors are affected, and the partially redundant information they provide can be confusing, particularly in a crisis situation where a response is needed quickly.

The focus of this paper is a new technique for attention focusing. The technique involves reasoning about the distance between two frequency distributions, and is used to detect both anomalous system parameters and "broken" causal dependencies. These two forms of information together isolate the 1ocus of anomalous behavior in the system being monitored.

1 Introduction

Mission Operations personnel at NASA have the task of determining, from moment to moment, whether a space platform is exhibiting behavior which is in any way anomalous, which could disrupt the operation of the platform, and in the worst case, could 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. Another established technique for anomaly detection is the comparison of predicted values from a simulation to actual values received in telemetry. However, experienced mission operators reason about more than just alarm threshold crossings and discrepancies between predicted and actual sensor values: they may ask whether a sensor is behaving differently than it has in the past, or whether a single behavior is resulting in—the particular bane of operators—a rapidly developing alarm sequence.

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. We seek an approach for determining from moment to moment which of the available sensor data is most informative about the presence of anomalies occurring within a system. The work reported here extends the anomaly detection capability in the SELMON monitoring system [2, 3] by adding an attention focusing capability. This work complements other work within NASA on empirical and model-based methods for fault diagnosis of aerospace platforms [4, 5].

2 Background: The SELMON Approach

Abnormal behavior is always defined as some kind of depart ure from normal behavior. Unfortunately, there appears to be no single, crisp definition of "normal" behavior. In the traditional monitoring technique of limit sensing, normal behavior is predefined by nominal value ranges for sensors, A fundamental limitation of this approach is the lack of sensitivity to context. In the other traditional monitoring technique of discrepancy detection, normal behavior is obtained by simulating a model of the system being monitored. This approach, while avoiding the insensitivity to context of the limit sensing approach, has its own limitations, The approach is only as good as the system model. It can be difficult to distinguish genuine anomalies from errors in the mode].

Noting the limitations of the existing monitoring techniques, we have developed an approach to monitoring which is designed to make the anomaly detection process more robust, i.e., to reduce the number of undetected anomalies. Towards this end, we introduce *multiple* anomaly models, each employing a different notion of "normal" behavior.

2.1 Anomaly Detection Methods

In this section, we briefly describe some of the methods that we use to determine when a sensor is reporting anomalous behavior. These measures use knowledge about each individual sensor, without knowledge of any relations among sensors.

Surprise

An appealing way to assess whether current behavior is anomalous or not is via comparison to past behavior, This is the essence of the *surprise* measure. It is designed to highlight a sensor which behaves other than it has historically. Specifically, *surprise* uses the historical frequency distribution for the sensor in two ways: To determine the likelihood of the given current value of the sensor *(unusualness),* and to examine the relative likelihoods of different values of the sensor *(informativeness).* It *is* those sensors which display unlikely values when other values of the sensor arc more likely which get a high *surprise* score. *Surprise is* not high if the only reason a sensor's value is unlikely is that there are many possible values for the sensor, all equally unlikely.

Alarm Anticipation

The *alarm anticipation* measure in SELMON performs a simple form of trend analysis to decide whether or not a sensor is expected to be in alarm in the future. A straightforward curve tit is used to project when the sensor will next cross an alarm threshold, in either direction. A high score means the sensor will soon enter alarm or will remain there. A low score means the sensor will remain in the nominal range or emerge from alarm soon.

Value Change

A change in the value of a sensor may be indicative of an anomaly. In order to better assess such an event, the *value change* measure in SELMON compares a given value change to historical value changes seen on that sensor. The score reported is based on tbe proportion of previous value changes which were less than the given value change. It is maximum when the given value change is the greatest value change seen to date on that sensor. It is minimum when no value change has occurred in that sensor.

Space limitations preclude describing additional SELMON anomaly measures which reason about individual sensors and about system interactions through the use of a causal model.

2.2 Previous Results

In order to assess whether SELMON increased the robustness of the anomaly detection process, we performed the following experiment: We compared SELMON performance to the performanceof the traditional limit sensing technique in selecting critical sensor subsets specified by a Space Station Environmental Control and Life Support System (ECLSS) domain expert, sensors seen by that expert as useful in understanding episodes of anomalous behavior in actual historical data from ECLSS testbed operations.

The experiment asked the following specific question: How often did SELMON place a "critical" sensor in the top half of its sensor ordering based on the anomaly detection measures?

The performance of a random sensor selection algorithm would be expcctcd to be about 50%; any particular sensor would appear in the top half of the sensor ordering about half the time. Limit sensing detected the anomalies 76.3% of the time. SELMON detected the anomalies 95.190 of the time.

These results show SELMON performing considerably better than the traditional practice of limit sensing. They lend credibility to our premise that the most effective monitoring system is one which incorporates several models of anomalous behavior. Our aim is to offer a more complete, robust set of techniques for anomaly detection to make human operators more effective, or to provide the basis for an automated monitoring capability.

The following is a specific example of the value added of SELMON. During an episode in which the ECLSS pre-heater failed, system pressure (which normally oscillates within a known range) became stable. This "abnormally normal" behavior is not detected by traditional monitoring methods because the system pressure remains firmly in the nominal range where limit sensing fails to trigger. I 'urthermore, the fluctuating behavior of the sensor is not modeled; the predicted value is an averaged stable value which fails to trigger discrepancy detection,

3 Attention Focusing

A robust anomaly detection capability provides the core for monitoring, but only when this capability is combined with attention focusing dots monitoring become both robust *and* efficient. Otherwise, the potential problems of information overload and too many false alarms may defeat the utility of the monitoring system.

Although many anomalies can be detected by applying anomaly models to the behavior reported at individual sensors, monitoring also requires reasoning about interactions occurring in a system and detecting anomalies in behavior reported by several sensors.

The attention focusing technique developed here uses two sources of information: historical data describing nominal system behavior, and causal information describing which pairs of sensors are constrained to be correlated, due to the presence of a dependency. The intuition is that the origin and extent of an anomaly can be determined if the misbehaving system parameters *and* the misbehaving causal dependencies can be identified.

3.1 Two Additional Measures

While SELMON runs, it computes incremental frequency distributions for all sensors being monitored. These frequency distributions can be saved as a method for capturing behavior from any episode of interest. Of particular interest are historical distributions which correspond to nominal system behavior.

To identify an anomalous sensor, we apply a distance measure, defined below, to the frequency distribution which represents recent behavior to the historical frequency distribution representing nominal behavior. We call the measure simply *distance.* To identify a "broken" causal dependency, wc first apply the same distance measure to the historical frequency distributions for the cause sensor and the effect sensor. This reference distance is a weak representation of the correlation that exists between the values of the two sensors due to the causal dependency. This reference distance is then compared to the distance between the frequency distributions based on recent data of the same cause sensor and effect sensor. The difference between the reference distance and the recent distance is the measure of the "brokenness" of the causal dependency. We call this measure *causal distance.*

3.2 Sonic Definitions

Define a distribution D as the vector d_i such that

$$\forall i, 0 \leq d_i \leq 1$$

and

$$\sum_{i=0}^{n-1} d_i = 1$$

For a sensor S, we assume that the range of values for the sensor has been partitioned into n contiguous subranges which exhaust this range. Wc construct a frequency distribution as a vector D_S of length n, where the value of d_i is the frequency with which S has displayed a value in the ith subrange.

We define two special types of frequency distribution. I-et F be the random, or flat distribution where Vi, $d_i = \frac{1}{n}$. I-d S_i be the set of "spike" distributions where d_i = 1 and $\forall j \neq i. d_j$ = O.

' If our aim was only to compare different frequency distributions of the same sensor, wc could use a distance measure which required the number of partitions, or bins in the two distributions to be equal, and the range of values covered by the distributions to be the same. However, since our aim is to be able to compare the frequency distributions of different sensors, these conditions must be relaxed.

3.3 The Distance Measure

The distance measure is computed by projecting the two distributions into the two-dimensional space $[f, s]$ in polar coordinates and taking the euclidian distance between the projections.

Define the "flatness" component $f(D)$ of a distribution as follows:

$$\sum_{i=0}^{n-1} \frac{1}{2} \left| \frac{1}{n} - d_i \right|$$

This is simply the sum of the bin-by-bin differences between the given distribution and F. Note that O $\leq f(D) \leq 1$. Also, $f(S_i) \rightarrow 1$ as n -i ∞.

Define the "spikeness" component s(D) of a distribution as:

$$\sum_{i=0}^{n-1} \phi \frac{i}{n-1} d_i$$

This is simply the centroid value calculation for the distribution. The weighting factor ϕ will be explained in a moment. Once again, 0< s(D) ≤ 1.

Now take $[f, s]$ to be polar coordinates [r, θ]. This maps F to the origin and the S_i to points along an arc on the unit circle. Sec Figure 1.

Note that we take $\phi = \frac{\pi}{3}$. This choice of ϕ guarantees that $\Delta(S_0, S_{n-1})$ = A(F, SO) = A(F, S_{n-1}) = 1 and all other distances in the region which is the range of A are by inspection $\leq$ 1.

Insensitivity to the number of bins in the two distributions and the range of values encoded in the distributions is provided by the [f,s] projection function, which abstracts away from these properties of the distribut ions.

Additional details on desired properties of the distance measure and how they are satisfied by the function A may be found in [1].

3.4 Results

In this section, we report on the results of applying the distribution distance measure to the task of focusing attention in monitoring. The distribution distance measure is used to identify misbehaving nodes (distance) and arcs (*causal distance*) in the causal graph of the system being monitored, or equivalently, detect and isolate the extent of anomalies in the system being monitored.

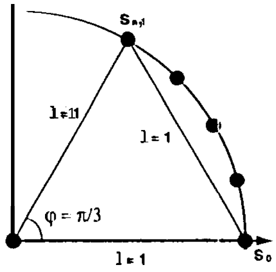

Figure 1: The function $\Delta(D_1, D_2)$.

Figure 2 shows a causal graph for a portion of the Forward Reactive Control System (FRCS) of the Space Shuttle. SELMON was run on seven episodes describing nominal behavior of the FRCS. The frequency distributions collected during these runs were merged. Reference distances were computed for sensors participating in causal dependencies.

SELMON was then run on 13 different fault episodes, representing faults such as leaks, sensor failures and regulator failures. Due to space limitations, only one of these episodes will be examined here; results were similar for all episodes. In each fault episode, and for each sensor, the distribution distance measure was applied to the incremental frequency distribution collected during the episode and the historical frequency distribution from the merged nominal episodes. These distances were a measure of the "brokenness" of nodes in the causal graph; i.e., instantiation of the *distance* measure.

New distances were computed between the distributions corresponding to sensors participating in causal dependencies. The differences between the new distances and the reference distances for the dependencies were a measure of the "brokenness" of arcs in the causal graph; i.e., instantiation of the *causal distance* measure.

The episode of interest involves a leak affecting the first and second manifolds (jets) on the oxidizer side of the FRCS. The pressures at these two manifolds drop to vapor pressure. The dependency between these pressures and the pressure in the propellant tank is severed because the valve between the propellant tank and the manifolds is closed, Thus there are two anomalous system parameters (the manifold pressures) and two anomalous mechanisms (the agreement between the propellant and manifold pressures when the valve is open).

The *distance* and *causal distance* measures computed for nodes and arcs in the FRCS causal graph reflect this faulty behavior, See Figure 3. (To visualize. how the distribution distance measure circumscribes the extent of anomalies, the coloring of nodes and the width of arcs in the figure are correlated with the magnitudes of the associated *distance* and *causal distance* scores, respectively). The apparent anomaly at the third manifold is due to a known flaw in the training simulator which generated the data. The explanation for the apparent helium tank temperature anomaly is more interesting: in response to the leak, the valve between the propellant tank and the manifolds closes. The closed system now has a smaller volume, and since the pressure remains the same, temperature must rise according to the ideal gas law. SELMON

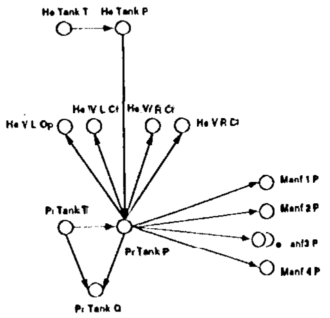

Figure 2: Causal Graph for the Forward Reactive Control System (FRCS) of the Space Shuttle.

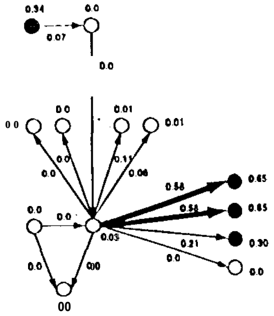

Figure 3: A leak fault.

flags this behavior as anomalous, even *though the relevant causal dependency" was not available in the model.* In this case, SELMON helped debug an incomplete model. This helium tank temperature behavior was present in the data for all six leak episodes.

4 Towards Applications

The approach described in this paper has usability advantages over other forms of model-based reasoning. The overhead involved in constructing the causal and behavioral model of the system is minimal. The behavioral model is derived directly from actual data; no offline modeling is required. The causal model is of the simplest form, describing only the existence of dependencies. For the Shuttle RCS, a 198-node causal graph was constructed in a single one and one half hour session between the author and the domain expert.

SELMON is being applied at the NASA Johnson Space Center as a monitoring tool for Space Shuttle Operations and Space Station Operations. Early applications include the one for the Propulsion (PROP) flight control discipline reported on here, and ones for the Thermal (EECOM) and Mechanical (MMACS) flight control disciplines. An operational SELMON prototype is available for evaluation by all flight control disciplines, only requiring that a list of sensors "owned" by that discipline be provided.

At the Jet Propulsion Laboratory, we are looking at the problem of onboard downlink determination for the Pluto Fast Flyby project, now in its early design phase. The spacecraft will have limited communications capacity and it will not be possible to transmit all onboard-collected sensor data. Only four hours of coverage from the Deep Space Network will be available per week. The challenge is to devise a method for constructing a suitable summary of a week's worth of sensor data guaranteed to report on any anomalies which occurred. The anomaly detection and attention focusing capabilities of SELMON may bc well-suited to this task.

5 Summary

We have described the properties and performance of a distance measure used to identify misbehavior at sensor locations and across mechanisms in a system being monitored. The technique enables the locus of an anomaly to be determined. This attention focusing capability is cornbincd with a previously reported anomaly detection capability in a robust, efficient and informative monitoring system, which is being applied in mission operations at NASA.

6 Acknowledgements

The members of the SELMON team are. Len Charest, Dennis DeCoste, Nicolas Rouquette and Jay Wyatt. Harry Porta provided valuable mathematical and counterexample insights during the development of the distance measure. Matt Barry and Charles Robertson provided valuable discussion. Matt Barry also served invaluably as domain expert for the Shuttle FRCS .

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

1] R. Doyle, "A Distance Measure for Anomaly Detection and Attention Focusing in Systems Monitoring," *8th International Workshop on Qualitative Reasoning about Physical Systems,* Nara, Japan, June 1994.

2] R. Doyle, L. Charest, Jr., N. Rouquette, J. Wyatt, and C. Robertson, "Causal Modeling and Event-driven Simulation for Monitoring of Continuous Systems," *Computers in Aerospace 9,* American Institute-of Aeronautics and Astronautics, San Diego, California, October 1993.

[3] R. Doyle, S. Chien, U. Fayyad, and J. Wyatt, "Focused Real-time Systems Monitoring Based on Multiple Anomaly Models," *7th International Workshop on Qualitative Reasoning about Physical Systems,* Eastsound, Washington, May 1993.

[4] T. Hill, W. Morris, and C. Robertson, "Implementing a Real-time Reasoning System for Robust Diagnosis," *6th Annual Workshop on Space Operations Applications and Research,* Houston, Texas, August 1992.

5] 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.