

Enhanced Telemetry Monitoring with Novelty Detection

José Martínez Heras, Alessandro Donati

■ Typically, automatic telemetry monitoring in space operations is performed by out-of-limits (OOL) alarms. This approach consists of defining an upper and lower threshold so that when a measurement goes above the upper limit or below the lower one, an alarm is triggered. We discuss the limitations of the out-of-limits approach and propose a new monitoring paradigm based on novelty detection. The proposed monitoring approach can detect novel behaviors, which are often signatures of anomalies, very early — allowing engineers in some cases to react before the anomaly develops. A prototype implementing this monitoring approach has been implemented and applied to several ESA missions. The operational assessment from the XMM-Newton operations team is presented.

The most widely extended approach for automatically detecting anomalous behavior in space operations is the use of out-of-limits (OOL) alarms. The OOL approach consists of defining an upper and lower threshold so that when a measurement goes above the upper limit or below the lower one, an alarm is triggered. Then engineers will inspect the parameter that is out of limits and determine whether it is an anomaly or not and decide which action to take (for example, run a procedure). This is the original out-of-limits concept.

The current OOL concept has evolved to cope with more situations such as distinguishing between soft and hard limits; for example, a soft OOL triggers a warning to pay attention, a hard OOL triggers an error that demands attention. Soft limits are contained within hard limits. In addition OOL thresholds (soft and hard) can be configured so that different thresholds are applicable in different situations (for example, depending on the working mode of a given instrument).

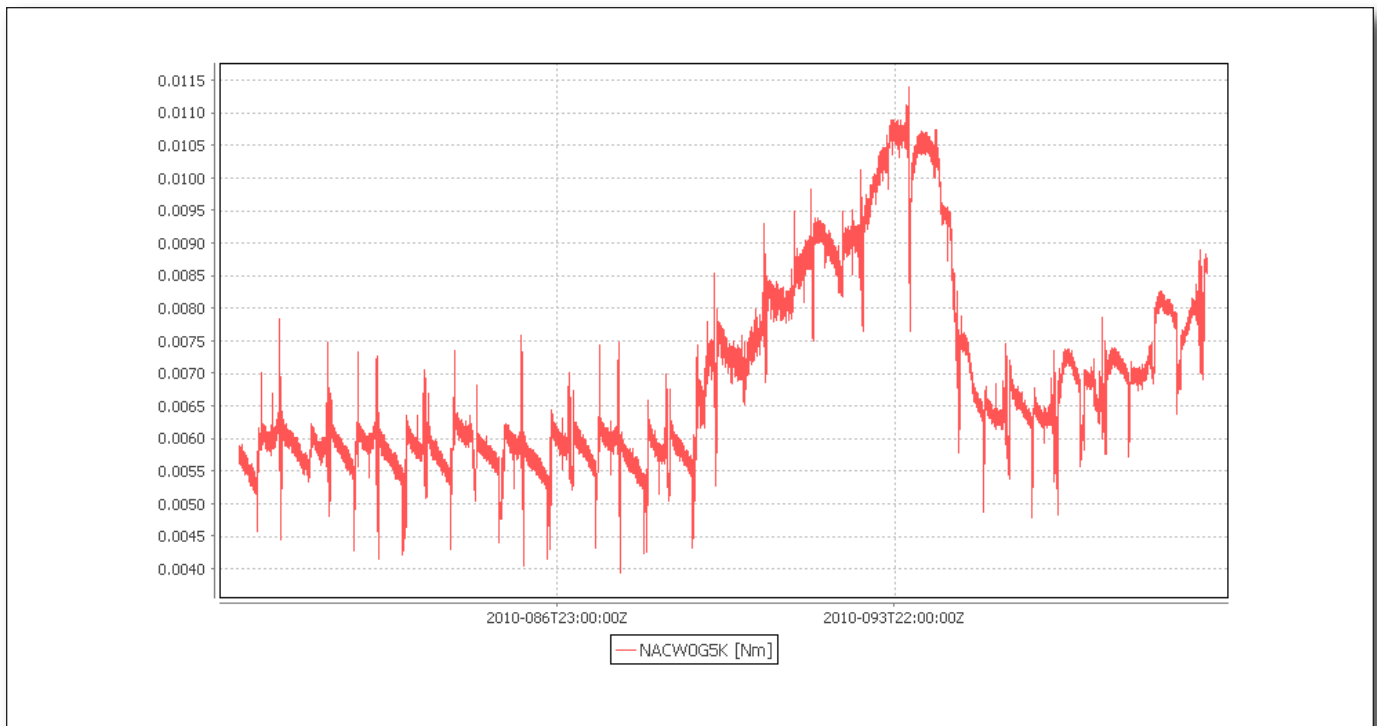


Figure 1. A Typical Example of Anomalous Behaviors.

Venus Express reaction wheel 4 friction increases starting in day of year 2010.89. This behavior was not detected by out-of-limits alerts as the upper limit was set to 0.02 at that time. In this case, this behavior was recognized by Venus Express engineers as they closely monitored the reaction wheels even if they did not trigger any out-of-limits alerts.

While out-of-limits approaches are useful and they successfully trigger alarms when parameter readings go out the defined thresholds, they suffer from some limitations. First, some behaviors are anomalous even if they are within the defined limits. A typical example is shown in figure 1. The parameter values reach the upper limit but do not hit it. Since engineers don't know when this will happen they have to monitor key telemetry parameters closely even if in most cases everything would be nominal. Paradoxically, sometimes the anomalous behavior is more in limits than the nominal one. Figure 5, depicted later on in this article, shows an example of this situation. More information about this anomaly will be discussed later. Second, OOL bounds are not defined for every parameter. Engineers only define OOL for a subset of parameters for which they want to receive alarms if they exceed the limits. Therefore, OOL is not systematic in the sense that it does not cover every parameter. Third, and quite often, engineers receive OOL alarms that are completely expected. A typical example is the OOL bounds defined for the automatic control gain (AGC) during a pass. At acquisition of signal (AOS) and loss of signal (LOS) the AGC goes outside limits. However, it is expected to happen and in every pass these two OOL alarms will

be raised. Ideally, engineers should only have to investigate real potential anomalies. Finally, it requires effort to adapt OOL alerts to useful values as the mission goes through different phases or simply degrades with time.

The novelty detector project has been developed to cope with the current OOL limitations. The novelty detector main goal is to automatically detect anomalies and report them to engineers for further investigation. The ideal novelty detector should take into account that parameters can behave nominally in several different ways. In addition, it should not make any assumption on what kind of behavior or how many different behaviors a parameter will have. This will allow it to work with any parameter without the need of prior knowledge.

Approaches to Automatic Anomaly Identification

There are mainly three approaches to identify novel behavior (possibly anomalies): supervised, unsupervised, and semisupervised. The supervised approach consists of having labeled examples of both nominal and anomalous behavior. This approach works quite well if we need to recognize previously labeled nom-

inal and anomalous behavior. Its major drawback is that it can only identify anomaly occurrences for the anomaly types that it knows already. This is a major limitation since the anomalies that generally will affect operations the most are the ones that are happening for the first time. Being able to recognize first-time-ever anomalies quickly allows flight control teams to take action immediately.

The unsupervised approach consists of having unlabeled examples of data — no prior knowledge is provided. The implicit assumption made by the systems using a nonsupervised approach is that anomalies happen far less often than nominal behaviors. So they attempt to automatically distinguish what is nominal and what is anomalous. The major drawback is the risk of missing anomalies: if an anomaly happened several times in the past, a nonsupervised system may consider it a normal behavior and not report it in the future.

The semisupervised approach is a combination of the supervised and unsupervised approaches. It consists of providing only nominal behaviors examples. The advantages of this approach are that engineers are in full control to specify what should be considered as nominal, repeated anomalies can be detected since they are not in the nominal set, and since no assumptions are made about the possible behavior of the anomalies, any anomalous behavior can be potentially detected.

The proposed monitoring paradigm follows a semisupervised approach to perform anomaly detection. We will use the term *novelty detection* instead of *anomaly detection* since the only thing that can be said is that a behavior is novel when compared to a set of behaviors known to be nominal. The new behavior might well be also nominal but so far not present in the nominal set. The decision of classifying a new behavior as nominal or anomalous is left to the flight control engineers.

Novel Behavior Detection

To characterize behaviors we compute four statistical features (average, standard deviation, maximum, and minimum) of fixed-length periods. The duration of the time period is chosen so that it represents a natural time span (for example, orbit period or time covered by the short-term planning). The exact duration is not critical; however, it should be long enough to allow behaviors to develop and not so long that many different behaviors happen in it.

While there are other ways to characterize behavior in a given time period (for example, Fourier transformations, wavelets, and so on) we used statistical features because they are robust to sampling rate changes and behavior order, and work even if very few samples are available. In addition, they are compatible with the future European Space Operations Centre (ESOC) infrastructure data archive (DARC).

DARC precomputes and make available these statistical features.

Figure 2 shows representation of how these fixed time periods look in a side-by-side two dimensional comparison. We are showing this representation in this document only as an example. In reality, four dimensions (average, standard deviation, maximum, and minimum) are used simultaneously.

Periods Distance

Once we have defined the representation of a time period for a given parameter we need to be able to compare time periods. We need a distance measurement so that we can say that for a given parameter A , the period X is closer to the period Y than to the period Z . Mathematically: $d(X, Y) < d(X, Z)$. We use the Euclidean distance as distance measurement (equation 1):

$$d(X, Y) = \sqrt{(avg_x - avg_y)^2 + (std_x - std_y)^2 + (max_x - max_y)^2 + (min_x - min_y)^2}$$

Outlier Detection

We make use of outlier detection techniques to find which periods have anomalous behavior. The general assumption is that anomalous behaviors will have greater distances to known nominal behaviors than known nominal behaviors among them. The question is how big the distance should be so that it can be considered a novel behavior. If the distance is too small many false anomalies will be reported. If the distance is too big then some anomalies will be missed.

The solution to overcome the problem of having to define an outlier distance is to use local density outlier detection techniques. The most widely used is called local outlier factor (LOF) (Breunig et al. 2000). LOF computes a factor that gives an indication of the degree of outlierness (novel behavior). It takes into account the density of the k closest points. If they are very dense, little distance is required to consider a new behavior a novelty. If the neighbors are sparse a bigger distance is required to consider a new behavior a novelty.

The major disadvantage of LOF is that the resulting factor values are quotient values and hard to interpret. A value of 1 or less indicates a clear inlier, but there is no clear rule for when a point is an outlier. In one data set, a value of 1.1 may already be an outlier; in another data set and parameterization (with strong local fluctuations) a value of 2 could still be an inlier. These differences can also occur within a data set due to the locality of the method.

To overcome the LOF limitations we will use local outlier probabilities (LoOP) (Kriegel et al. 2009). It is a relatively new method derived from LOF. It uses inexpensive local statistics to become less sensitive to the choice of the parameter k . In addition, the resulting values are scaled to a value range of [0:1] (Kriegel et al. 2009), that can be directly interpreted as the

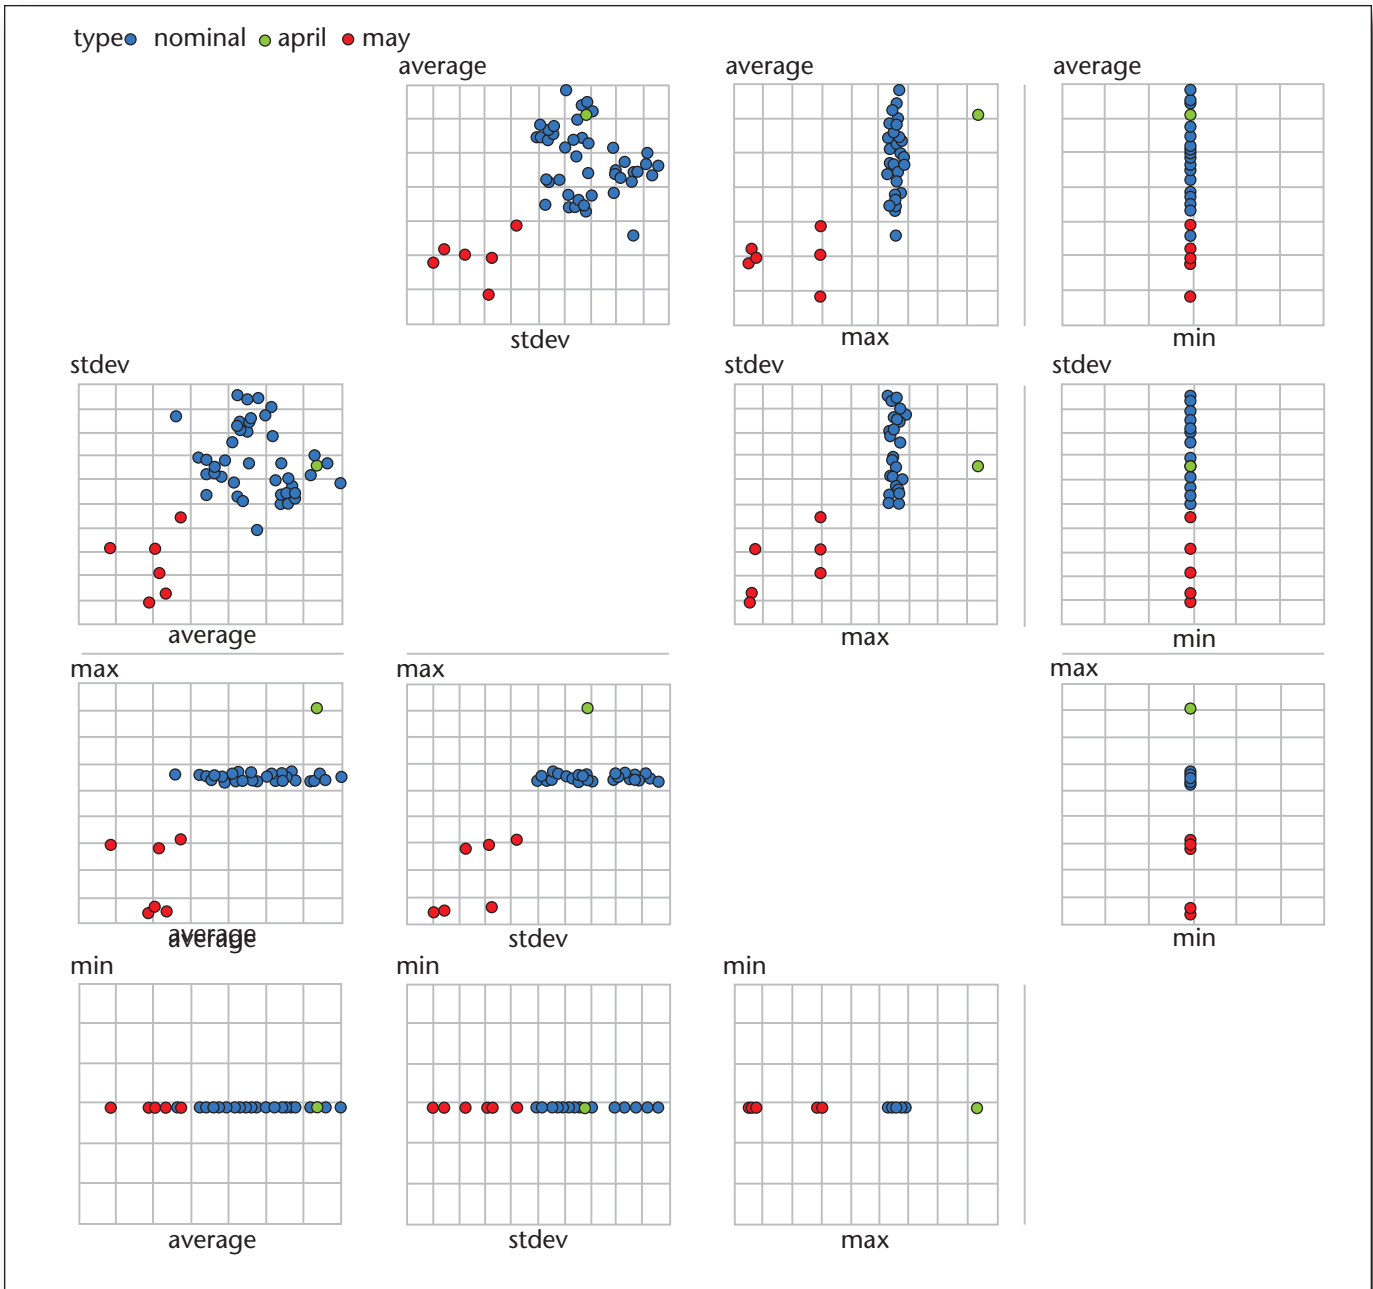


Figure 2. Statistical Features.

Feature extraction and side-by-side comparison of the XMM anomaly shown in figure 5. Every point represents a time period; its position in the chart is given by its statistical features. Legend: blue consists of nominal periods (1 January 2009 – 31 March 2009), green represents a spike in April, and red represents the thermostat dithering anomaly³. This two-dimensional example is only an easy way to visualize this representation. In reality, four dimensions (average, standard deviation, maximum, and minimum) are used simultaneously.

probability of being a new behavior. Figure 3 shows an example using two dimensions (for example, average and standard deviation). LoOP has the advantage of being more robust and providing a much more intuitive output. By using LoOP we can rank novel behaviors by novelty probability showing first the behaviors with higher chances to be truly novel.

Novel Behavior Detection — Summary

We use LoOP (Kriegel et al. 2009) to find whether a new behavior is novel with respect to a set of given behaviors. If the set of given behaviors consists of nominal behaviors only, and LoOP finds that the




Figure 3. LoOP.

Local outlier probabilities outlier detection technique graphical example with two dimensions. Points with high outlier probability indicate a potential novel behavior.

new behavior is an outlier with high probability, it can mean only two things: the new behavior is either an anomaly or it is nominal in a new way (for example, it behaves differently than other nominal behaviors in the nominal set). The more nominal time periods the nominal set has, the more the chances that the reported novelties are really anomalies.

Behaviors are characterized by simple statistical features (average, standard deviation, maximum, and minimum) in a fixed-size time period (for example,

duration of an orbit). This means performing LoOP in four dimensions. Euclidean distance is used to measure how different two time periods are. The same procedure is applied to all parameters, and those with high probability of being novelties are notified to engineers.

To be mode independent of the choice of the parameter k (number of closest points to compute density) we try several $k = \{10, 20, 30\}$ and use as outlier probability the minimum value of these differ-

ent runs. The assumption is that if a new behavior is really an outlier it should be an outlier independently of the number of closest points (k) used to compute the local density. This procedure minimizes the chances of getting false alarms. It has limitations, however, because this approach does not consider novelties in the combination of two or more parameters; it works on a parameter-by-parameter basis only. It is, however, systematic in the sense that it can be applied to every parameter.

Prototype

If flight control engineers would be able to look every day at every parameter they would be able to identify all novelties. Unfortunately, they cannot. There are way too many parameters (in the order of several thousands) and the trend is that this number will increase in future missions. The objective of this prototype is to automate the process of noticing novel behavior at the parameter behavior level.

New behaviors are often signatures of anomalies either happening now or in the way to develop. Noticing them early is of utmost importance for planning corrective measurements and keeping the spacecraft healthy. We should take into account that not every new behavior corresponds to an anomaly: it could be related to a new environmental condition (for example, extremely high radiation) or be totally expected as the result of planned new operations (for example, Venus orbit insertion).

The functionality of being able to automatically detect anomalies has been the driver for this project. However, we understood that we could not build such system. The closest we can get is identifying a new behavior as novel when compared to a set of known behaviors. Hence the name novelty detection.

In order to get closer to our goal of being able to automatically detect anomalies, we choose the known behavior set so that it only contains nominal behaviors. With this configuration, when a new behavior is classified as novel it can only mean two things: it is either an anomaly or a new nominal behavior. As time passes, the set of known nominal behaviors will grow. This has a positive impact in reducing the number of novelty alerts, as many behaviors will be classified as nominal.

The novelty detection prototype makes use of mission utilities and support tools (MUST) (Martínez-Heras, Baumgartner, and Donati 2005; Baumgartner et al. 2005) as housekeeping telemetry and ancillary data provider. The MUST's performance allows the performance of advanced monitoring with novelty detection efficiently.

Functionalities

We will now discuss the two major functionalities of the novelty detection prototype. The underlying

principle is the same, but one can achieve one functionality or the other depending on which configuration is used.

Identification of Potential Anomalies

The main purpose of the novelty detection prototype is to detect potential anomalies. For fulfilling this objective we will use as a known periods set the collection of all known nominal behaviors. This way, when a new behavior is classified as novel with a high probability, it is very likely that it would be an anomaly. It could be still a new kind of nominal behavior but, as time passes, this should happen less and less frequently.

Verification of Expected New Behavior

In addition to identification of potential anomalies, the same novelty detection technique can be used to verify the occurrence of an expected new behavior. For instance, let's say that certain behavior is expected as a consequence of a maneuver and we would like to verify it. A way of doing it with the novelty detection prototype is by using the recent past as the known behavior set. The output of the novelty detection will be the novel behaviors compared with the recent past. These novelties should contain the expected new behaviors and possibly other parameters. These parameters that were not initially foreseen can be considered side effects of the expected behaviors.

Input

Two inputs are required to run the novelty detection prototype: periodicity and the set of known behaviors.

Periodicity is the statistical features needed to characterize a fixed-length time period and has to be computed over a large enough time period. A typical example is to use the periodicity of the orbit or the amount of time that the short-term planning covers.

Set of known behaviors: the novelty detection will detect whether a new behavior is novel as compared with the set on known behaviors specified as input. Two options are recommended: (1) use all nominal behaviors: this is ideal to perform anomaly detection; and (2) use the recent past (this should be used to verify expected new behavior).

Output

The output consists of a text file that contains the list of parameters that are believed to be novel. They are grouped by period and by novelty probability within periods. Figure 4 shows an example of such an output file. The same information is stored in a database for further analysis by client applications.

Operational Validation and Current Usage

To prove the feasibility of the monitoring paradigm with novelty detection we applied it to an already documented anomaly that the ESA satellite XMM-

```

0.935057, no-Gaps, 2010-10-14T00:00:00Z [2010.287.00.00.00], NTTX1000, Loop Error Term. 1
0.817176, no-Gaps, 2010-10-14T00:00:00Z [2010.287.00.00.00], NACW0G15, RW4_SWR estimat friction
0.935141, no-Gaps, 2010-10-15T00:00:00Z [2010.288.00.00.00], NTTX1000, Loop Error Term. 1
0.804666, no-Gaps, 2010-10-15T00:00:00Z [2010.288.00.00.00], NACW0R08, AUT_GUID Cmd Quater Vy

```

Figure 4. Example Output File.

Novelties found for Venus Express using the novelty detection prototype for verifying expected new behavior. The format of the file output is the probability of being an outlier, whether this parameter had data gaps during this period, the start time of the period (in two time formats), parameter mnemonic, and parameter description.

Newton experienced in 2009 and checked when the novelty detection prototype would have been able to detect the anomaly. Here is an excerpt of the paper that describes the anomaly and measures taken by the FCT (flight control team) to cope with it (Pantaleoni et al. 2010):

We noticed that the thermostat T6073 started to have a strange behavior since mid-May 2009, already 2 months before the failure was spotted. The thermostat range reduced the temperature where it opened and started to decrease, a sign of a deterioration of the high threshold, even if the bottom limit was respected quite well, until mid-July, when the upper limit and the lower limit went very close to each other. The thermostat started to oscillate, in a narrow temperature range, until it did not close anymore at the correct temperature, and it let the temperature go down to almost 22 deg. This first temperature drop was not spotted because it did not generate any OOL. After that the thermostat had some cycles with a nominal behavior, but on 13 July 2009 the temperature went down deeper, to 21.25, triggering an OOL and allowing the FCT to spot the problem.

We configured the novelty detection prototype to consider data in the range (January 2009, March 2009) as nominal. We used as time period 48 hours since it is the duration of an XMM-Newton orbit. Then we run the novelty detection prototype for the period (April 2009, July 2009). The results is that the novelty detection prototype managed to find unusual behavior 2 months before the out-of-limit triggered. This is remarkable not only because it allows to react to anomalies early, but also because it matches flight control engineers diagnosis results and mimics the effect of having somebody looking every day at every parameter and noticing if something new is happening. Figure 5 shows where the OOL triggered and where the novel behavior was found.

Another tests related with caging was performed with the novelty detection prototype. In the XMM FCT's words:

XMM reaction wheel 1 faces unexpected increases in torque and current consumption during stable point-

ing. The reaction wheel 1 of XMM suffers from an increment of friction. The possible cause is cage instability due to under-lubrication, *Rosetta*, another spacecraft, also suffered from the same problem. The Rosetta manufacturer pointed to under-lubrication as a possible cause. Since this anomalous phenomenon has been spotted, the damaged reaction wheel has been closely monitored. The XMM FCT wanted to know if the novelty detection could have been able to detect the caging earlier than the flight control team did?

For this test, nominal data was defined as the month of May 2005 when no caging was present. May 2010 was investigated for caging on parameter A5249. The parameter A5249 refers to the torque of the reaction wheel 1. Caging has effectively been detected and the probabilities computed are high enough to be seriously considered. In the end, if the XMM flight control team could have used novelty detection beforehand, the caging phenomenon would have been detected and closely monitored earlier. The lifetime of the damaged wheel could have been saved thanks to an earlier relubrication.

Currently, the novelty detection prototype checks every day around 2000 XMM-Newton housekeeping telemetry parameters and reports which of them, if any, has a new behavior. The results are sorted by probability of certainty of being a new behavior. The novelty detection for XMM is integrated in a wider scope project, XMM early warning system (XNEWS) (Kirsch et al. 2012). XNEWS is developed to perform near real-time trend analysis of spacecraft parameters in order to detect early degradation of components. XNEWS will enable the mission to perform early countermeasures in case degradation is detected.

In addition, the novelty detection prototype has been integrated as a plug-in of WebMUST (Oliveira et al. 2012). WebMUST is a web-based interface to access telemetry and ancillary data in the MUST (Martínez-Heras et al. 2005, Baumgartner et al. 2005) Repository. WebMUST allows users to very efficiently create plots and reports. WebMUST can also sup-




Figure 5. Nominal Behavior.

February – April, OOL on 13 July. This thermostat has been properly working showing the same behavior for 10 years. However, it started to have a strange behavior since mid-May 2009 and it was only noticed two months after (July 2009) when it crossed the lower limit. For this type of anomaly, the out-of-limits checks are not effective because, paradoxically, the behavior of the anomaly was “more in limits” than before. The proposed novelty detection monitoring technique could find this anomaly two months before the out-of-limit alarm triggered.

port characterizations and anomaly investigation using the DrMUST (Martínez-Heras et al. 2012, Martínez-Heras et al. 2009) plug-in. A screen capture of the novelty detection display for an expected new behavior is shown in figure 6.

Related Work

The problem of automatically finding unusual behavior has been addressed by other researchers in a number of fields, both space and nonspace.

For instance, A. Patterson-Hine and colleagues (Patterson-Hine et al. 2001) uses a model-based approach to detect anomalies in helicopters; however, model-based approaches require a big upfront engineering effort. In this project we have focused on approaches that require as little engineering effort as possible.

E. Keogh and colleagues (Keogh, Lin, and Fu 2005)

describe the algorithm HOT SAX in order to automatically and efficiently find time series discords. Time series discords are defined as subsequences of longer time series that are maximally different to all the rest of the time series subsequences. Its major advantage is its simplicity as it just requires a single input: the length of the subsequence (in our case it would be the period length). While HOT SAX is successful at finding the ranking of discords for time series, it is of little use to spacecraft engineers that need to understand if these discords are relevant or not, especially when they are monitoring thousands of parameters of different nature. In addition, it is difficult to know how important a top-ranked discord of a certain parameter is in relation to the other parameters' top-ranked discords.

D. L. Iverson and colleagues (Iverson et al. 2012) present the inductive monitoring system (IMS), which uses clustering to characterize normal itera-

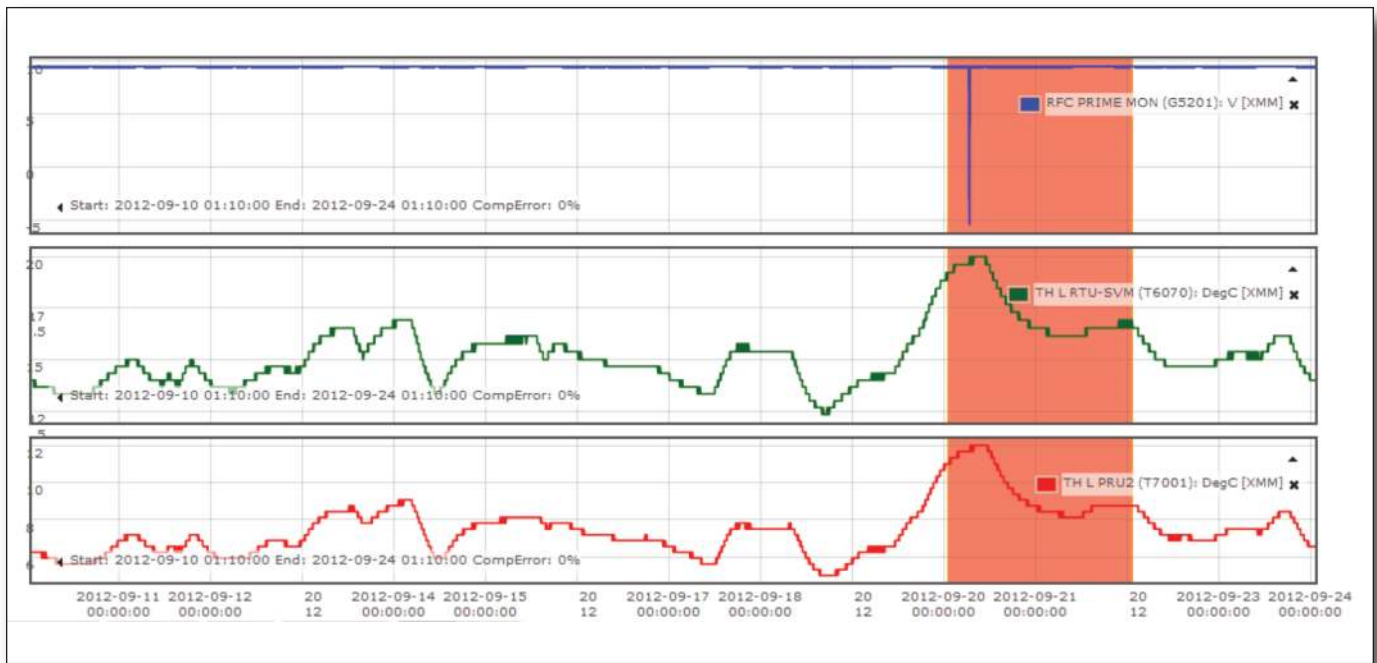


Figure 6. Novelty Detection Display for an Expected New Behavior.

Screen capture of the novelty detection plug-in integrated into WebMUST. It shows an expected new behavior for XMM. The highlighted area corresponds to the period when the novel behavior was detected.

tions between parameters. Engineers provide a vector describing which measurements make sense to monitor together (for example, pressures, valve positions, temperatures) and examples of nominal behavior. Then IMS builds clusters based in the nominal data. The learning process is influenced by three learning parameters: the maximum cluster radius, the initial cluster size, and the cluster growth percent. After the clusters are created, the IMS performs monitoring by comparing the current vector data with the clusters database. If the distance is 0 or close to 0, its behavior is considered nominal. As the distance increases, the current behavior can be considered more anomalous. The clear advantage of IMS over the proposed novelty detection monitoring technique is that it can detect unusual behaviors in a combination of parameters. However, it has some disadvantages with respect to the proposed novelty detection technique: the grouping of which parameters need to be monitored together, apart from requiring engineering effort, determines the kind of anomalies the system will be able to detect. In this work we are concerned with the detection of novel behaviors as soon as possible even if engineers do not think this will be possible. Another disadvantage is that the anomaly score is not intuitive; if it is 0 or close to 0 it is nominal, but it is not clear when engineers should start paying attention and performing investigations. A further disadvantage is the amount of tuning that is required

to have the IMS work properly: the weight in the vector components and the values of the three learning parameters have a strong impact in the creation of the cluster database and, therefore, in the results from the monitoring phase.

Conclusions

We have introduced a new monitoring paradigm based on novelty detection. In this approach, every day every telemetry parameter is automatically scanned and a list of the parameters that exhibit a novel behavior is reported to flight control engineers. New behaviors are often signatures of anomalies either happening now or in the way to develop. Noticing them early is of utmost importance for planning corrective measurements and keeping the spacecraft healthy.

A clear advantage of the proposed monitoring paradigm is the little amount of engineering effort required: the only inputs required consist of the period duration (for example, 1 day) and ranges of times to be used as examples of nominal behavior. During the validation phase, users really appreciated that it generates very few false alarms: the fact that it uses a local density outlier detection technique avoids the need of using a distance threshold to detect new behaviors. Therefore, this approach does not suffer from the problem of having to define a threshold

where having it too small would lead to many false alarms and having the threshold too big will lead to missing new behaviors.

The proposed novelty detection monitoring approach has been successfully validated with XMM anomalies, finding them before they were triggered by the out-of-limits alarms, sometimes as early as two months in advance. Currently, the XMM mission uses the novelty detection prototype to detect new behaviors on about 2000 parameters on a daily basis. We would like to highlight that monitoring with novelty detection is not mission specific but generic. We can easily adapt it to any ESOC controlled mission since we use MUST (Martínez-Heras et al. 2005, Baumgartner et al. 2005) as the data provider.

We believe that every mission will benefit from the adoption of the novelty detection monitoring paradigm as complement to the classic out-of-limits mechanism. Being able to know which few parameters (out of several thousands) exhibit a new behavior helps flight control engineers to efficiently direct their monitoring efforts. The ESA's patents group has decided to protect the proposed monitoring paradigm by filing an international patent (WO2013 010569).

References

- Baumgartner, A.; Martínez-Heras, J.; Donati, A.; Quintana, M. 2005. MUST — A Platform for Introducing Innovative Technologies in Operations. Paper presented at the 2005 International Symposium on Artificial Intelligence, Robotics and Automation for Space, Munich, Germany, 5–9 September.
- Breunig, M. M.; Kriegel, H.-P.; Ng, R. T.; Sander, J. 2000. LOF: Identifying Density-Based Local Outliers. In *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data* (ACM SIGMOD Record), 93–104. New York: Association for Computing Machinery. dx.doi.org/10.1145/342009.335388
- Iverson, D. L.; Martin, R.; Schwabacher, M.; Spirkovska, L.; Taylor, W.; Mackey, R.; Castle, J. P.; Baskaran, V. 2012. General Purpose Data-Driven Monitoring for Space Operations. *Journal of Aerospace Computing, Information, and Communication* 9(2): 26–44. In *Proceedings of the 5th IEEE International Conference on Data Mining (ICDM 2005)*, 226–233. Piscataway, NJ: Institute of Electrical and Electronics Engineers. dx.doi.org/10.1109/ICDM.2005.79
- Keogh, E.; Lin, J.; and Fu, A. 2005. HOT SAX: Efficiently Finding the Most Unusual Time Series Subsequence. In *Proceedings of the 5th IEEE International Conference on Data Mining (ICDM 2005)*, 226–233. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Kirsch, M.; Martin, J.; Pantaleoni, M.; Southworth, R.; Schmidt, F.; Webert, D.; Weissmann, U. 2012. Cage Instability of XMM-Newton's Reaction Wheels Discovered During the Development of an Early Degradation Warning System. Paper presented at the 12th International Conference on Space Operations (SpaceOps 2012), Stockholm, Sweden, 11–15 June.
- Kriegel, H. P.; Kröger, P.; Schubert, E.; and Zimek, A. 2009. LoOP: Local Outlier Probabilities. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, 1649–1652. New York: Association for Computing Machinery. dx.doi.org/10.1145/1645953.1646195
- Martínez-Heras, J.; Baumgartner, A.; Donati, A. 2005. MUST: Mission Utility and Support Tools. Paper presented at the Data Systems in Aerospace Conference (DASIA 2005), Edinburgh, Scotland, 30 May–2 June.
- Martínez-Heras, J. A.; Donati, A.; Sousa, B.; Fischer, J. 2012. DrMUST — A Data Mining Approach for Anomaly Investigation. Paper presented at the 12th International Conference on Space Operations (SpaceOps 2012), Stockholm, Sweden, 11–15 June.
- Martínez-Heras, J. A.; Yeung, K.; Donati, A.; Sousa, B.; Keil, N. 2009. DrMUST: Automating the Anomaly Investigation First-Cut. Paper presented at the IJCAI-09 Workshop on Artificial Intelligence in Space. Pasadena, California, USA, 17–18 July.
- Oliveira, H.; Lais, A.; Francisco, T.; Donati, A. 2012. Enabling Visualization of Large Telemetry Datasets. Paper presented at the 12th International Conference on Space Operations (SpaceOps 2012), Stockholm, Sweden, 11–15 June.
- Pantaleoni, M.; Kirsch, M.; Martin, J.; Weissmann, U.; Krusenstiern, N. 2010. The New Operational Strategy for XMM-Newton's Tank Thermal Control, After Low Temperature Protection Thermostat Failure. Paper presented at the 10th International Conference on Space Operations (SpaceOps 2010), Huntsville, AL, 25–30 Apr.
- Patterson-Hine, A.; Hindson, W.; Sanderfer, D.; Deb, S.; Domagala, C. 2001. A Model-Based Health Monitoring and Diagnostic System for the UH-60 Helicopter. In *Proceedings of the American Helicopter Forum*, Volume 57, No. 2, 1552–1564. Alexandria, VA: American Helicopter Society, Inc.

José Martínez-Heras is a software engineer/researcher at the Advanced Operations Concepts Office at the European Space Agency. His main interest consists of enhancing monitoring and diagnostics in space operations. His research includes the enhancement of remote observability, availability of space data, reduction of information overload, and data mining and machine learning applied to space operations. In recent years he was the main coauthor of four patents related to data compression, diagnostics, and monitoring. Martínez-Heras holds a master of computer science engineering degree from the University of Málaga, Spain.

Alessandro Donati has 25 years of experience in satellite operations at the European Space Operations Centre of ESA in Darmstadt, Germany. In 2001 he founded and has since led a skunk-work style managed team devoted to introducing and exploiting advanced enabling technology in support of innovative mission operations concepts. The operational tasks covered by the applied research include planning, scheduling, diagnostic, resources management, behavior modeling, and forecasting. The technologies of interest include artificial intelligence, data mining, and soft computing, among others. In recent years he was coauthor of three patents in the area of data compression and diagnosis. Donati holds a master of electronic engineering degree from the University of Rome La Sapienza.