Proceedings of the Institute of Acoustics

MECHANICAL HEALTH MONITORING OF A SATELLITE IN ORBIT

P P Collins (1) and S J C Dyne (2)

1 European Space Research and Technology Centre, Nordwijk, The Netherlands
2 Institute of Sound and Vibration Research, University of Southampton

Abstract

When the Olympus satellite was launched on 12 July 1989, it carried a vibration monitoring experiment called PAX. The aim of PAX was to quantify the vibration environment and to build up a database of the vibration characteristics of the on-board equipment. The signals were to be transmitted to Earth as continuous time-histories for ground processing, with the result that large volumes of data had to be stored and analysed. The task reported in this paper is the approach adopted in reducing those data to the minimum required to describe adequately the vibration environment and how the use of PAX broadened into a mechanical health monitor as a result of the development of an automatic data analysis system. Following a preliminary assessment of the data, the first task was to identify and isolate a particular event from the background noise. A monitor was built which could detect events using the autocorrelation of the signal, showing how the necessary large-scale reduction of the data could be achieved. The second step was to find a simple way of describing an event using a small number of parameters. These had to be presented in a way which allowed trends to emerge, so that potential problems could be anticipated, unexpected changes could be diagnosed and designers could make best use of the new information about the vibration environment. The methods developed yielded a practical system for isolating specific events in real-time with a low computing overhead, paving the way for conventional condition monitoring and statistical analysis of the vibration environment in space.

Introduction

Putting Olympus into orbit cost around $30,000/kg, excluding development costs. This posed several problems for the designers of the PAX vibration monitor. Severe constraints on mass and volume not only put restrictions on the device itself, but also limited the power available to drive it. There was no immediate commercial justification for PAX, as its output would be of use only to the designers of future missions. Finally, PAX required the most expensive commodity of all; a full audio-range transmission bandwidth to downlink the data. It is against this background that one can see why a long-term vibration monitor had not been flown on a satellite before and one can begin to appreciate that the design requirements were far from simple [1].

On board a spacecraft, there are many vibration sources, including thrusters, stabilising reaction wheels, solar array drives and antenna pointing mechanisms. Little was known about the expected level of vibration or whether there would be significant differences between one measurement direction and another at a given location.

On the Olympus satellite, it was decided to use a tri-axial accelerometer with two sensitivity ranges monitoring a single measurement point (Table 1). The data were carried on the side-band of a TV channel, a much wider bandwidth than normally available, making PAX the first system capable of sending live, continuous audio-band data about the vibration environment on an orbiting satellite. The intention had been to monitor briefly at intervals the vibration signature under various running conditions, near the beginning of life, middle and end of life of the satellite. Initial attempts were made to analyse the entire time-history captured using Fourier analysis, but it was soon found that this was too indiscriminate, producing large volumes of output which were difficult to interpret [2].

Table 1: PAX Equipment Specifications

mass	2.37 kg
power consumption	6.3 W
dimensions	200 x 134 x 110 mm
sample rate	3125 samples/s
bandwidth	0.5 - 1000 Hz
ADC resolution	12 bits
range	100 mg coarse mode
	10 mg fine mode

When the signal was coupled to a D/A converter and an audio system, however, it was found that the ear was able very easily to separate the sounds of the various devices and identify periods of interest. The PAX accelerometers were acting like a stethoscope attached to the satellite.

However, it was clear that it would be impossible to monitor aurally and analyse even a single channel of data, 24 hours per day. As a result, the task was broken into two parts: firstly, the data had to be reduced down from 2.3 GBytes/day to a manageable level and secondly, very specific parameters had to be identified so that analysis could be limited to carefully chosen areas of interest.

Data Reduction

The accelerometer measurements on three axes and the temperature at the PAX location were transmitted to ground at a sampling rate of 3125 samples/second, providing a usable bandwidth of around 1 kHz. Observations of the three axes revealed that while there were differences between them, the basic spectral shape was very similar, with all the major components visible in all three directions. The out-of-plane Z-channel was chosen for analysis, as this gave a response level approximately 10 dB higher than the two in-plane channels. In periods of high activity on Olympus, the transient events of interest were seconds or more apart and in normal operation, the period between transients was nearly 15 seconds. The continuous vibration from devices such as the reaction wheels was also of interest, but this only needed to be recorded at intervals. Clearly, major reduction of the remaining data-stream could be achieved by extraction of event signatures from the background and this would enable individual mechanisms to be monitored. The key to continuous 24-hour monitoring was the automation of this process.

As the early FFT analysis had shown, comparison of acceleration spectral density (ASD) measurements of nominally the same device at different times was fraught with difficulties. Not only did it require a skilled vibration engineer to assess the output, but it was often uncertain whether the data were related only to the event of interest or to a combination of factors. Furthermore, the output had to be available to specialists from other disciplines for trend-spotting or fault diagnosis and they may not be skilled at interpreting ASD measurements or vibration time-histories. Instead, the specialist may simply want to know if the vibration level from a particular thruster was about to cause a de-pointing of a TV antenna. Ideally, he would like the binary answer of *yes* or *no*.

The early analysis had shown that the emphasis had to be switched from the application of routine analysis techniques to the identification of specific features of interest; having defined clearly the parameters to be monitored, it is a relatively straightforward task to extract them [3].

The present case study highlights the scale of this change of emphasis. The early analysis had taken many man-months to perform and yet the extraction of useful data was completed in a matter of days; the basic algorithm for automatically logging a specific event, however, took only one and a half man-weeks to implement, while the refinement of the analysis and the presentation went on for over a year, enabling a clear picture to be built up of the event's behaviour.

Automatic Event Detection

The spur to the development of automatic event detection came when a screech-like noise was heard. Sometimes, this would start without warning and disappear just as suddenly. On other occasions, it would start and fade gradually and on others it would be present continuously, but at a much reduced level. The spectral features, however, were very similar on each occasion. Some simple testing confirmed that the noise came from the reaction wheels, devices used to maintain the orientation of the spacecraft. While it had no immediate effect on the control of the satellite, the phenomenon was unexpected and it was decided to monitor its occurrence and overall level. Simple rms level monitoring did not work, as a low threshold was triggered by several mechanisms, while a higher level missed a significant number of occurrences of the screech. Band-pass filtering improved matters slightly, but thruster firings contained enough energy in the pass-band to cause false alarms.

The key features of the screech centred on a unique cluster of high-frequency peaks between 600 and 750 Hz (Figure 1), but these could only be extracted from the ASDs by developing complex software. It was decided that these peaks were the simplest reliable indicator of the occurrence of the screech and that a detector should be developed which extracted them from the time-history. The required output was a *Present/Not Present* warning from what could be called a selective condition monitor.

Condition Monitoring

Health monitoring has become an established branch of vibration analysis [4]. It is widely used for predicting machine service requirements and even the likely time-to-failure of a component. Often, acceleration spectral density measurements are compared at regular intervals. In other fields, Cepstral analysis is used for identifying sources and changes in characteristics. One technique for predicting failure involves the use of time-series analysis and auto-regressive methods for highlighting periodic components or regularly occurring events in a signal. An example of this is the detection of the passing of a damaged tooth in a gearbox, where the autocorrelation function can be used to give the time delay between a reference marker and the cyclical passing of a damaged tooth, allowing its location to be determined.

The autocorrelation function, however, gives more than just time-domain information; it can show the relative importance of harmonic and broad-band components in a signal. As many machines produce low-frequency, harmonic components in normal operation and suffer from anomalies which produce high-frequency, broad-band noise, the autocorrelation function can be used as a means of recognising key features of a machine's signature and in particular, changes in that signature. For a typical AC coupled vibration signal from a mechanical device, the autocorrelation will have a fairly sharp peak at the origin, whose height is the rms level of the signal, with any high amplitude, low-frequency components appearing as periodic components along the time axis. The appearance of a few broad-band, high-frequency components will tend to increase the signal's resemblance to white noise, whose autocorrelation function is a delta function. This effect is revealed by the position of the first zero-crossing, which moves nearer to the origin as the broad-band content at high frequencies increases. Note that these properties are dependant only on the relative strength of the frequency components and not on absolute level and are therefore insensitive to rms level or minor changes in the frequency content of the broad-band terms. For event recognition, this makes the autocorrelogram a useful alternative to the technique of band-pass filtering and threshold detection, as a particular signature can be recognised at different signal levels.

Selective Condition Monitoring

Using an FFT based analyser to calculate the autocorrelation function of the PAX data, it was observed that a significant movement of the first zero-crossing towards the origin occurred when the screech was present and that no other device or combination of devices caused the same shift. Figures 2 and 3 show the PAX signal autocorrelation function under normal conditions and with screech respectively. The inset figures show an expansion of the first few points of the autocorrelation function. The autocorrelation data were downloaded into an HP 9000 computer and the screech detection algorithm looked for a first zero-crossing in the first three data points. If found, a warning was displayed. Loss of PAX data was an occasional problem and this could be detected by making use of a small DC offset in the data: if the PAX signal were lost, the autocorrelogram would always be positive. In practice, if the first zero-crossing was beyond the seventh point, the PAX signal was no longer present.

The autocorrelogram could be used to detect other mechanisms, such as the very strong 150 Hz tone of the antenna pointing mechanisms, but these were not included in the first system.

The monitor worked well, but the rapid fluctuations in the screech meant that the log recorded many hundreds of short bursts, making interpretation difficult. After a few days of collecting data, it was found that only the long bursts were of interest. The routine was modified to group short bursts together and to ignore brief losses of screech in an otherwise long recording. The log was greatly simplified by this approach and over a period of weeks the output listings were further refined as trends began to emerge in the output. Many options were explored for plotting the data and once a format had been established, a dedicated monitor was built based on a 386 computer and a digital signal processing card. The result was a compact monitor system, which produced a graphic display of both occurrence and rms amplitude of the screech [5,6,7]. The system automatically produced backups of these data to tape, so that analysis could be made entirely off-line, in order to minimise the chances of human error interrupting the monitor.

Discrimination versus Robustness

The problem faced in writing event recognition algorithms is the trade-off between discrimination and robustness. While discrimination provides great precision in identifying the source, it will only allow a positive identification if all the specified conditions are met. Making the algorithm more robust provides higher tolerance to pollution from other sources, but usually results in a less discriminating algorithm. In a remote monitoring system, robustness is important, as failure to detect an event could be critical. It may also prove difficult to spot the failure of the algorithm, allowing recurrence of the event to go unrecorded for a significant time.

The screech algorithm was designed therefore to be robust and it worked perfectly for 17 months. Its lack of discrimination (it could not distinguish between the screeches from the three reaction wheels, for example) was less important. However, it eventually did fail, when new, high-amplitude, low-frequency harmonics appeared in the wheel signatures. These altered the relative strengths of the broad and narrow-band components in the signal, moving the zero-crossing away from the origin. This failure was important for two reasons. Firstly, it was only detectable using the audio monitor. This re-emphasised the important role of aural monitoring of the signal and indicated the value of being able to transmit the original time-history. Even if the transmission bandwidth is only a few bytes per second, such as on a scientific mission, an original time-history may be the only sure way to detect such an algorithm failure. Secondly, the failure of the algorithm, as much as its success, has shown the way forward in algorithm development. In order to achieve a given reliability of detection, it would be better to link a larger set of simple algorithms than to use a small series of general ones linked by a complex set of inter-relations. By starting with simple algorithms looking for specific parameters, it had been possible to see quickly the reason for the failure. If more general feature-extracting algorithms had been linked by some form of neural net, this simple but vital conclusion, based on a single failure in seventeen months, would have been lost in the overall success rate of the network. Solving the problem with a more complex network would only have resulted in an unwieldy system in which it became progressively more difficult to identify problems with the event detector algorithms.

It may prove necessary to develop two or more independent networks for detecting the same event, so that the failure of one segment can be identified by another before the signature changes too drastically. This cross-checking would be particularly important on scientific missions where continuous audio monitoring at the ground station would not be possible.

CONCLUSIONS

Many important roles for satellite vibration monitoring are beginning to emerge, of which fault diagnosis and health monitoring are two. The next phase should extend detection to other devices using simple algorithms, paving the way to a more all-embracing approach based on the general signal features which emerge. Future systems on-board satellites will provide more information if the available channels are fed by a dispersed array of single-axis accelerometers. Bandwidth could be saved by reducing the resolution of the data. Further work is needed to establish a reasonable upper frequency limit to enable typical mechanical sources to be characterised. It is also not clear how much raw data needs to be stored to permit effective historical analysis in the period leading up to a new anomaly, when new detection algorithms will need to be developed. The recording interval and the sample length needs to be determined. The final objective of this work is clear: The reduction to a minimum of the data required to describe an event. In order to fly a monitor on a scientific mission, the mass, power consumption, computing power and above all, transmission bandwidth will have to be reduced to a minimum.

REFERENCES

[1] D TUNBRIDGE, The PAX Experiment on Olympus, ESA Bulletin 64, 1990
[2] G NUGENT, PAX Ground Software System, CARA, Dublin, May 1990
[3] J N PINDER et al., ESA Contract 9360/91/NL/DG(MD) Final Report, Section 3, ISVR, University of Southampton, March 1992
[4] S BRAUN (ed.), Mechanical Systems Analysis: Theory and Applications, London Academic, 1986
[5] S J C DYNE and J M A HALL, ISVR Olympus PAX Monitor User Manual, University of Southampton, April 1992
[6] S J C DYNE and J M A HALL, ISVR Olympus PAX Monitor Technical Reference Manual, University of Southampton, April 1992
[7] S J C DYNE, Satellite Mechanical Health Monitoring, IEE Colloquium: Advanced Vibration Measurements, Techniques and Instrumentation for the Early Prediction of Failure, May 1992

Figure 1: Typical PAX acceleration spectra

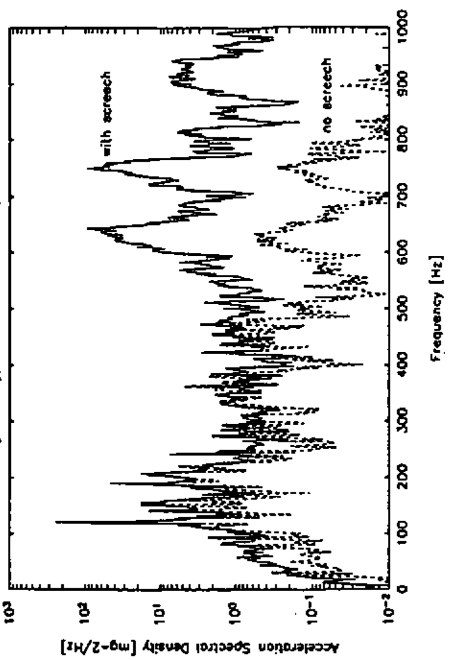

Figure 3: Autocorrelogram of signal with screech

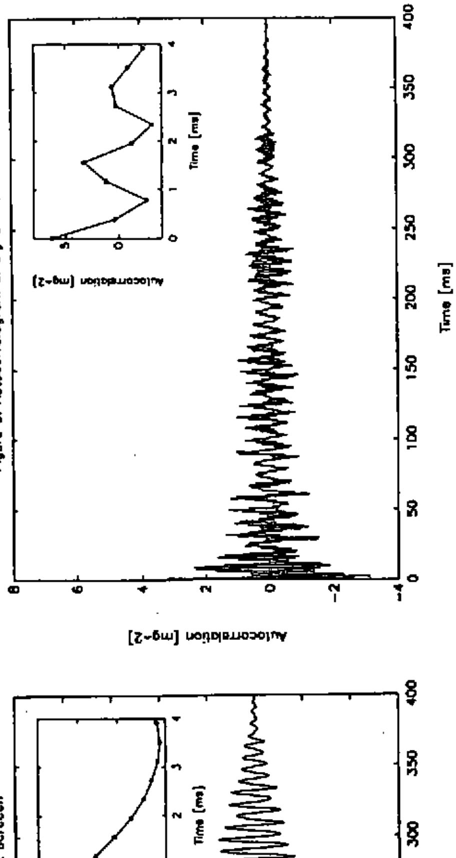

Figure 2: Autocorrelogram of signal without screech