Hydro-turbine monitoring: from self-learned equipment behavior to a single global deviation indicator

François Léonard, James Merleau, Dominique Tapsoba, Martin Gagnon

Hydro-turbine monitoring: from self-learned equipment behavior to a single global deviation indicator

François Léonard, James Merleau, Dominique Tapsoba, Martin Gagnon

Abstract:

In large machines, the possible failures are numerous and complex events. Furthermore, monitoring a multitude of measurement points leads to an accumulation of a large amount of data, making the task of tracking the equipment even more difficult. Often, there might even be an emerging problem that can influence several indicators simultaneously without exceeding the alert level of any one of them in particular. For end-users, the proposed monitoring system acts like a black box with an alert that comes on when the equipment's behaviour changes, much like the "check engine" of a car. The behavior of the equipment is tracked by a continuously updated statistical model and it is thus possible to compare future measurements with model predictions. Both the model and measurement uncertainties are taken into account by the procedure, and the confidence level can be determined by the user and is therefore a fixed quantity. In addition, the approach combines all the individual exceedances of each indicator in a single global deviation. When an alert is raised, the individual exceedances can be used in order to determine the source of the alert. Two different examples are presented where failures occurred on hydraulic turbine-generator units: the loss of a hydraulic turbine runner cone and a breakage of a thrust bearing.

Introduction

Large machines are generally constructed in small numbers and are custom made for each project. Therefore, they usually differ from one project to the next. Each time an equipment is refurbished, units which were initially similar in their construction become less and less

so. An observed defect on a unit will not show exactly the same symptoms as a similar defect observed on another unit, and furthermore, such events are fairly rare. As a consequence, we usually do not have a sufficiently large historical dataset of observed failures to feed machine learning methodologies such as neural networks [1]. On the other hand, to model a large machine in order to simulate failure modes is laborious and there is no guarantee that the full set of possible events will be explored. We are thus left with monitoring large machines to detect behavioural changes [2] which above a certain threshold establishes a compromise between the probability of missing a defect, a missed alert, and that of a false alert. In this paper, we present an alert system that uses available information in order to detect statistically significant behavioural changes. The main advantage of the approach is to increase the detection sensitivity to behavioural change while decreasing the false alert rate.

The proposed system consists of four modules: a data validating module, a modeling module, an interpolation module and a comparison module. We will call a snapshot the set **S** of indicators (see Fig. 1) obtained from data measured to characterise the response of an equipment as well as the prevailing operating conditions during the measurements. Let the i^{th} snapshot be given by $\mathbf{S}_i \equiv \mathbf{R}_i \cup \mathbf{O}_i$, where $\mathbf{R}_i \equiv \{r_{i,1}, r_{i,2}, \cdots, r_{i,n}, \cdots, r_{i,N}\}$ is the set of response characteristics, and $\mathbf{O}_i \equiv \{o_{i,1}, o_{i,2}, \cdots, o_{i,m}, \cdots, o_{i,M}\}$, the set of operating conditions. An operating condition can for example correspond to a given power set (MW and MVAR) at a given cooling fluid temperature, say $\mathbf{O}_i = \{280\text{ MW}, 40\text{ MVAR}, 10\text{ °C}\}$ for snapshot «*i*».

The data validation module retains snapshots for which operating conditions have been stable since a fixed amount of time and for measurements constrained to a range of acceptable values. For example, snapshots taken during transitory regimes will not be conserved and snapshots for which the measured cooling water temperature is below zero will also be eliminated. If there is no Boolean variable indicating that the equipment is operating, the module will detect that the machine is stopped to remove the corresponding snapshots.

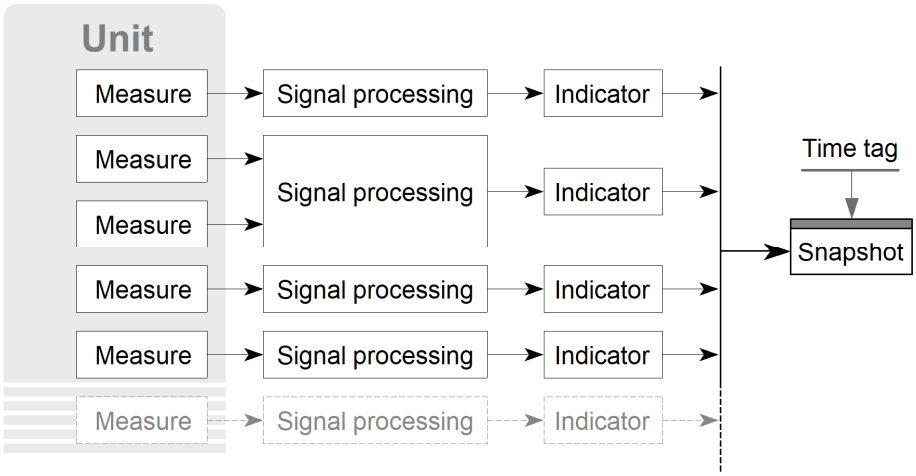

Figure 1: Snapshot definition.

The modelling module (see gray shading of Fig. 2) uses validated historical snapshots when the equipment is considered to be in good condition. The history of snapshots needs to be representative of most operating conditions usually imposed on the equipment since for monitoring purposes, it is difficult to quantify a behavioural discrepancy for an equipment that is operated under a new regime. The model captures the average behaviour of the machine responses and the statistical variations of the responses for given historical operating conditions. Numerical modelling being necessarily discrete, the number of operating conditions are limited and they thus need to be fairly uniformly distributed through the space of operating conditions historically explored.

The interpolation module (see Fig. 2) calculates the response and its statistical dispersion in regions unexplored by the available operating conditions of the model. Furthermore, this module evaluates the error magnitude of the interpolation computations. For example, an interpolation near one or many available operating conditions will be more precise than an interpolation, or extrapolation, far from available operating conditions.

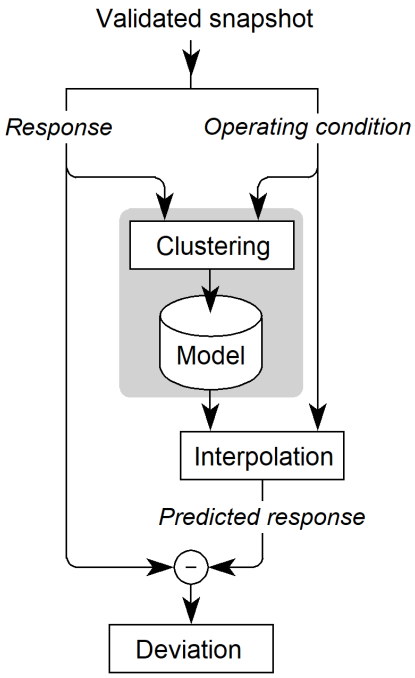

Figure 2: Monitoring process flowchart. The clustering module processes each new snapshot to construct the model. When the model has treated enough information on the machine's response, the interpolation module gives a response prediction for the current operating conditions and the observed response is compared to the predicted response which quantifies the behavioural deviation (comparison module).

The comparison module (see circled minus sign of Fig. 2) computes the deviation from normal behaviour by comparing the measured value with the one predicted by the model. The size of the deviation is put in relation with the quadratic sum of the model standard deviation which takes into account the interpolation error. The sensitivity of the procedure thus increases for measurements close to the available model historical operating conditions for which the machine has a reproducible response. The module has a deviation concentrator that calculates the quadratic sum of observed deviations on the set of monitored indicators. An indicator can be an instantaneous measuring channel value, a spectral line value, a camera pixel value, etc. Therefore, a scaling method is needed for each indicator before computing the quadratic sum of the deviations. For a given operating condition $\mathbf{O}_i$, the sum of the snapshot deviations has a constant, non-zero, magnitude when computed on several snapshots: the average deviation for this operating condition is written $\bar{d}(\mathbf{O}_i)$. When carrying out the comparison, the quadratic distance of the snapshot with the model prediction is compared with $\bar{d}(\mathbf{O}_i)$, in order for the average behaviour of distances to be centered at zero. Moreover, for a large number of indicators (> 30), the deviations from $\bar{d}(\mathbf{O}_i)$ will tend towards a normal (Laplace-Gauss) distribution.

Therefore, this leads to a single global indicator that is distributed according to a centered normal distribution, which facilitates establishing threshold values relative to sensitivity and a desired proportion of false alerts.

Clustering model

The model must capture the machine's response in normal operating conditions for a set of such conditions. This comes down to a mapping problem which should give a representation of an indicator as a function of the coordinates defined by the aforementioned conditions. Linear binning would be the simplest way to achieve this. For instance, in the above example where $\mathbf{O}_i = \{280 \text{ MW}, 40 \text{ MVAR}, 10 \text{ °C}\}$, in order to bin an operating condition, the individual domains could be partitioned in the following manner from 100 to 300 MW by increments of 10 MW, from -20 to 100 MVAR by increments of 10 MVAR, and from 0 to 35 °C by increments of 5 °C. Considering all the possibilities in these three dimensions given this binning scheme, the modelling support will then contain close to 2000 bins which will for most of them be empty because the accumulated historical data of running conditions will not have covered all the defined operating bins. Linear binning might be useful in one or two dimensions, but in multiple dimensions, as required for a hydro-electrical production unit, it is necessary to have narrower bins where the machine is most often operating, thus requiring nonlinear binning. Operating conditions defined by several indicators and a filling procedure based on historical data suggests using clustering to replace binning: an adequate clustering method automatically adapts the sampling grid in the domain of the operating conditions as a function of available information.

The proposed dynamic clustering approach [3] is similar to the *k*-means [4] methodology. The procedure provides a way to partition the space of the equipment's response but also to determine its reproducibility as a function of different operating conditions. In the multidimensional representation, a cluster can be understood as a group of points around a center of gravity, where each point corresponds to a snapshot. For a given cluster, the center of gravity is given by the means of the different indicators, taken over the snapshots

contained in the cluster. The decision to include a snapshot in a given cluster or not is based on Euclidian distances calculated in the operating condition space **O** and the decisions are used to compute the mean and dispersion of each indicator. Given that the measurement scales are different for different signals (or indicators), a scaling operation is performed in the domain of the response indicators **R** as well as the domain of the operating conditions **O**. The set of multiplicative scaling factors is written as **M**. For example, one could have $\mathbf{M} = \{\mathbf{R}; \mathbf{O}\} = \{(500\mu\text{m})^{-1}, (1\text{g})^{-1}, (1\text{g})^{-1}, (100°\text{C})^{-1}; (300\ \text{MW})^{-1}, (100\ \text{MVAR})^{-1}, (30\ °\text{C})^{-1}\}$, corresponding to {displacement, acceleration, acceleration, temperature, active power, reactive power, temperature}, where the symbol « ; » indicates the separation between the scaling coefficients of the response and those of the operating conditions. The scaling problem is common to all monitoring systems that combine heterogeneous signals: a judicious scaling procedure is fundamental for the monitoring methodology to work well [5][6]. In practice, the maximum and minimum values of the indicators can be used for scaling since some of them will be known and others can be adjusted sequentially from the accumulated snapshots. The multiplicative factors can thus be adjusted from the difference of the maximum and minimum of each indicator, or determined by a weighing strategy for the different indicators. For example, decreasing a multiplicative factor of an indicator has the effect of decreasing the importance of that indicator.

The number of clusters *M* determines the sampling resolution fixed at the initialising stage. The population cluster is given by the number of snapshots included in the cluster. To each cluster corresponds an operating condition for which the equipment response and its dispersion are estimated. In order to adequately estimate the dispersion of a cluster, many snapshots are necessary, therefore the number of clusters is determined to reach a compromise between the cluster populations and the sampling resolution. This compromise also needs to take into account computation time and the required memory which increases with the number of clusters via the modelling and interpolation modules.
In the clustering procedure, the first *M* snapshots initially form the *M* clusters. Afterwards, the module groups each new snapshot with the closest cluster according to a Euclidian distance in the operating condition space **O**. If a new snapshot is too distant from the current clusters, the two closest clusters are merged into one and a new cluster is born, containing

the new snapshot. The result of the dynamic clustering procedure can be slightly different according to the arrival order of the snapshots: the method is suboptimal since it considers the snapshots sequentially and not the whole history of snapshots simultaneously. For a history with a large number of snapshots, it is not realistic to consider simultaneously all the snapshots. The clustering procedure stops when enough historical data has been processed to determine the equipment's initial response. Monitoring the equipment's behaviour can nonetheless start before the characterization of the equipment's response is completed as long as a population of some clusters gives a measure of their dispersion.

Interpolation by kriging

In the case of a machine that is always running close to the same operating condition, clustering can be summarised by a single cluster which provides a mean machine response and the dispersion about this response. In this case, an interpolation step is not required before the comparison stage. Otherwise, when the model captures different machine responses associated with different operating conditions, interpolation can efficiently use the discretized information support. For equipment monitoring, the snapshots have intrinsic properties which restrain the type of interpolation that can be used. An important of these properties is reproducibility: for two neighbouring operating conditions, even possibly the same operating conditions, the response is reproducible but not identical. Interpolating splines could therefore not be used in this given context since it would be impossible for the method to approach two responses with identical operating conditions; alternative methodologies would be smoothing splines or least-squares regression, which can treat the intrinsic variability of the responses. It is also necessary to consider that in the model, there are clusters which have large populations and others, small ones. The uncertainty of a response associated with a cluster is roughly given by the dispersion of the snapshots divided by the square root of the cluster population. The interpolation procedure therefore needs to take into account these sources of uncertainty. Kriging, as an unbiased linear interpolation method, seems to be a good solution. This methodology, used in a wide variety of disciplines, was first proposed by Danie Gerhardus Krige, a South African statistician and mining engineer in 1951. The method was formalised by the French

mathematician Mathéron [7] who showed that its application gives best linear unbiased estimates.

The important point is that kriging takes into account the distances between an interpolation point and the different clusters, but also the uncertainty associated with the clusters. Since an interpolated value does not have any meaning without an associated uncertainty, the interpolation module evaluates a standard deviation for each interpolation point. The module produces two vectors: $\mathbf{E}(\mathbf{O}_i)$, which gives the estimated response for each indicator, and $\mathbf{C}(\mathbf{O}_i)$, which gives the standard deviation of the interpolation procedure for the different indicators, all quantities expressed in the original units.

Multivariate comparison

The comparison module (circled minus sign in Fig.2) computes the distance from the normal machine behaviour by comparing the measured value with the predicted value from the model. A multidimensional distance is calculated in the space with dimensions defined by the indicators which characterise the machine's response. The total instantaneous response distance is given by

$$d_i = \left\| \left(\mathbf{R}_i - \mathbf{E}\left(\mathbf{O}_i\right) \right) \cdot \mathbf{M} \right\| = \sqrt{\sum_{n=1}^{N} \left(\left(r_{i,n} - e\left(\mathbf{O}_i\right)_n \right) \cdot m_n \right)^2} \quad (1)$$

where *N* is the number of indicators. It measures the discrepancy between the response $\mathbf{R}_i$ of snapshot «*i*» and the estimate **E** in the operating condition neighbourhood $\mathbf{O}_i$, with **M** representing the scaling vector. This total distance is the quadratic combination of the observed deviations of the different indicators and is compared to the average total response distance

$$\bar{d}_k = \frac{1}{\sum_{j \in \mathbf{A}} 1} \sum_{j \in \mathbf{A}} d_j \quad \text{with} \quad j \in \mathbf{A} \text{ if } \mathbf{O}_j \approx \mathbf{O}_k, \quad (2)$$

for a similar operating condition. In practice, if the group response has dissimilar dispersion for different operating conditions, it is recommended to consider a single average deviation for the whole set of operating conditions. This leads to a faster monitoring initialisation

and also a faster algorithm. The distance of interest thus becomes $\bar{d}$, the observed average total response distance for the whole set of operating conditions, i.e. the average over *k* in equation (2). The quantity to be monitored is the relative deviation

$$w_i = d_i - \bar{d} \tag{3}$$

which represents the total instantaneous response distance of the last snapshot *i* relative to the observed average total response distance. For a response made up of several indicators and when the machine behaviour is reproducible, the distribution of the quantity w_i tends to a centered Laplace-Gauss variate with its standard deviation given by

$$\sigma = \sqrt{\frac{1}{\sum_{j\in\mathbf{B}} 1} \sum_{j\in\mathbf{B}} w_j^2} \quad . \tag{4}$$

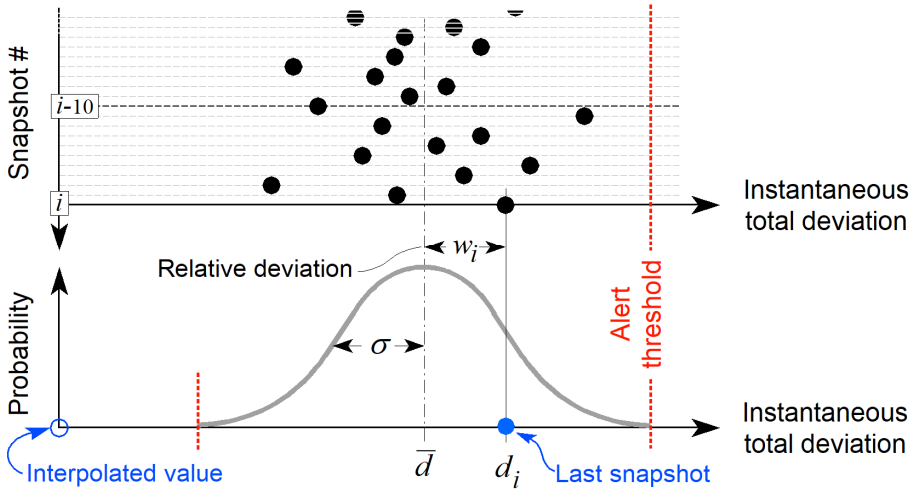

Figure 3: Sequence of total instantaneous response distance d_i (top) and the corresponding probability density function (bottom) with the most recent relative deviation for snapshot *i*. Here, the alert threshold is set at about $2.5\sigma'$ (see equation (5)).

The top part of Figure 3 illustrates the last computed total instantaneous response distances (equation (1)) while the bottom part shows the last relative deviation with respect to the probability density function of the observed total deviations. To take into account the

interpolation standard deviations $\mathbf{C}(\mathbf{O}_i)$, it is added quadratically to the standard deviation in (4) such that

$$\sigma' = \sqrt{\sigma^2 + \left|\mathbf{C}(\mathbf{O}_i)\cdot\mathbf{M}\right|^2} \tag{5}$$

becomes the standard deviation used in the comparison module to detect anomalies. A multiple of the global standard deviation σ' will thus determine the alert threshold (see Figure 3). The introduction of the interpolation standard deviation in the expression of the global standard deviation takes into account the precision of the interpolation, which depends on the dispersion of the historical data summarised by clusters and also the interpolation distance relative to the clusters. The set **B** in equation (4) can leave out data for which the relative deviation from the average total response distance is excessively large; it is thus possible to avoid introducing a bias that could lead to an increase of the standard deviation σ with the appearance of a defect. This can be done by using for example the criterion

$$\left|w_j\right| > 3\cdot\sigma' \quad . \tag{6}$$

Indeed, in the presence of a real defect, of a measuring problem or a transcription error in the database, the relative deviation increases abruptly and should be reported to the operators and be discarded in the computations. Also, the presence of a progressive defect can lead to a drift in the computations of the model and the calculated distances. Therefore, it is highly recommended to stop the clustering procedure when the machine response is reasonably modeled, or in other words, has captured most of the operating conditions.

From a multivariate perspective, for an operating condition $\mathbf{O}_i$, the machine response without noise will be close to the center of the scatter plot corresponding to the snapshots obtained during the given operating conditions. With the increase of the number of snapshots, the center of the scatter plot will tend towards a response without the random contribution, which follows from the law of large numbers. Contrary to what is usually accepted, for example with principal component analysis [8][9], the multivariate scatter plot is not full but hollow. As illustrated in the bottom part of figure 3, the probability density function of having a snapshot at a certain distance from the model prediction is

given by a “hyperspherical” shell for a fixed operating condition. In the case of changing operating conditions, in a multivariate context, the domain of the machine’s response is a manifold and the scatter plot resulting from the snapshots will form a shell around the manifold [10]. The corresponding probability density function, perpendicular to the manifold, is Gaussian at a constant distance from the manifold (see bottom part of figure 3).

Case studies

Loss of a hydraulic turbine runner cone

During an inspection on April 24th 2015, the loss of a hydraulic turbine runner cone was recorded (see Fig. 4).

Figure 4: Picture of the cone of the propeller seen from the draft tube, before (left) and after (right) loss.

The cone is located at the end of the hydraulic profile of the propeller to minimize losses. The cone also reduces the amplitude of the vortex rope, a physical phenomenon that decreases the performance of turbines and causes fluctuations in pressure and electrical power. The loss of a cone results in a decrease of the efficiency of the order of 0.6% and a slight increase of the vibration.

Figure 5 shows the results obtained with the proposed monitoring system. In the main screen of Figure 5, the relative deviation (equation (3)) is plotted in white for roughly 80000 snapshots, the red and green traces represent the upper and lower confidence bounds (or alert thresholds) fixed here at $\pm 4\sigma'$ (see equation (5)). The path of the relative deviations indicates an emerging problem with the cone from roughly the 57000th snapshot onwards, where the 30 bolts are most probably giving away one after the other, with an acceleration of the process when fewer bolts are left (see snapshot 68000) with the loss of the cone around snapshot 71000. The variations of the $\pm 4\sigma'$ confidence intervals (red and green traces) for the first 10 000 snapshots can be explained by an important contribution from the interpolation standard deviations, $\mathbf{C}(\mathbf{O}_i)$, which stems from a fairly short history of snapshots treated by the clustering module.

Take the example of a motor vehicle wheel having an imbalance. If we add on the rim a very small weight, depending on the position of this weight, it can increase as well as decrease the imbalance. However, when one of the first bolts holding the propeller cone was lost, it was a reduction in the imbalance that occurred: a reduction in vibration was observed for 22 days (Figure 6, roughly snapshots 66500-68000). At the next bolt, the vibration recovered close to the usual level (Figure 6, roughly snapshots 68000-70000) and two bolts lost later, the vibration increased (Figure 6, after snapshot 70000).

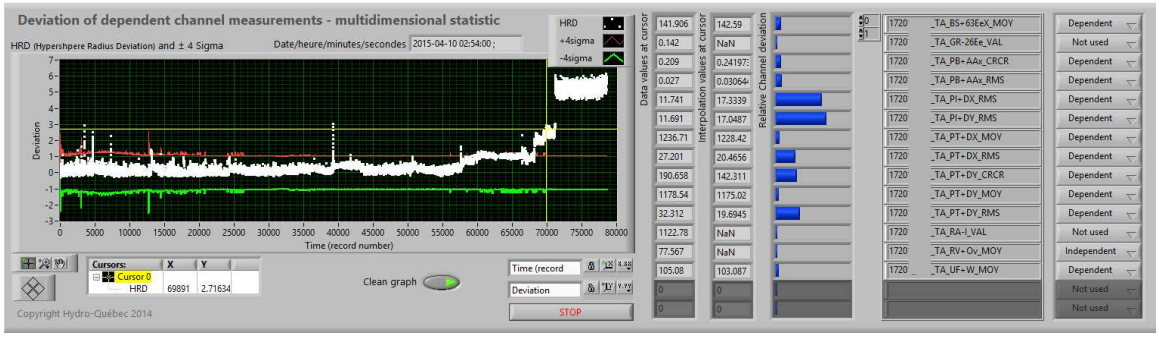

Figure 5: Relative multidimensional deviation observed over two years, from August 2013 to June 2015, where cone loss occurs on April 23, 2015.

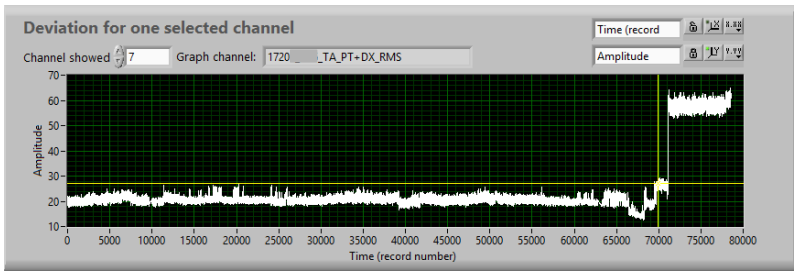

Figure 6: RMS vibration observed on the turbine bearing channel in the "x" direction (power house axis seen from the downstream side) for the same recordings as those shown in Figure 5 with the same cursor position.

Figure 7, which zooms in on the last sequence of snapshots of Figure 5 (67000 to 73152), shows exceedances larger than 8σ'(a deviation of about 2 on the y scale) which start around snapshot 68200 and last for more than 26 days in a row before the cone is lost (a deviation of 5 on the y scale). The exceedances of the global deviation are a result of several indicators and this can be appreciated by the contributions of the different indicators given by the blue rectangles on the right side of Figures 5 and 7. Indeed, Figure 7 shows that several vibration indicators (RMS, CR-CR) and the mean shaft position (MOY in Figure 7).

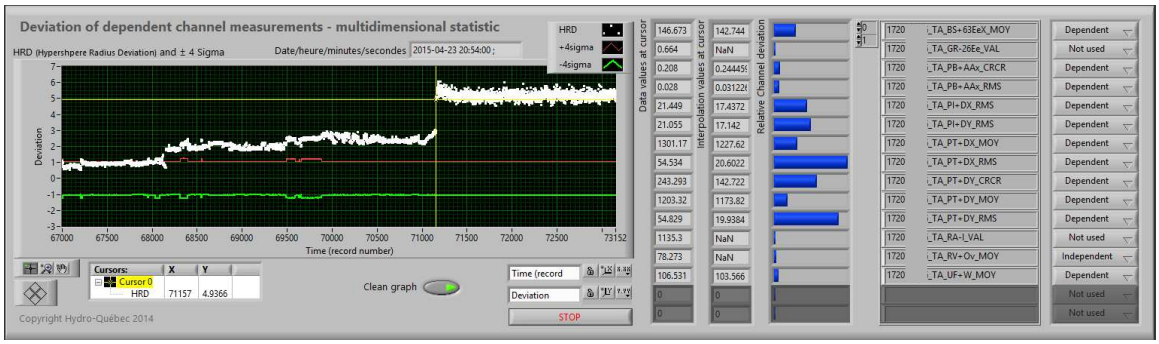

Figure 7: Zoom of Fig. 5 with the cursor placed on records located shortly after the loss of the cone.

Coming back to Figure 5, it should be noted that a few large “punctual” exceedances (in the relative deviation) are present between 0 and 5000, 5000 and 10000, and between 35000 and 40000. These exceedances can be explained by a cold-start after the machine was stopped for several hours: these could be discarded with the validation module by

extending the idle time delay when the machine is restarted. Figure 8 shows a close-up of the large exceedance between snapshots 35000 and 40000, which features two slopes of stabilization. The first one is due to the equilibrium temperature of the alternator ($\tau \approx$2h) and the second one to the equilibrium temperature of the structure ($\tau >$12h). For these transitory states when the machine restarts, it is suggested to increase the alert threshold according to a relevant profile, for example the sum of exponentials with decay constants of τ=2h and τ=12h with a delay of an hour. These "punctual" exceedances can necessitate adjusting the alert thresholds at high values or use a profile as described previously during transitory operations. The proposed global monitoring system could thus have adjusted alert thresholds in order to better keep track of the machine's behaviour even in transitory regimes.

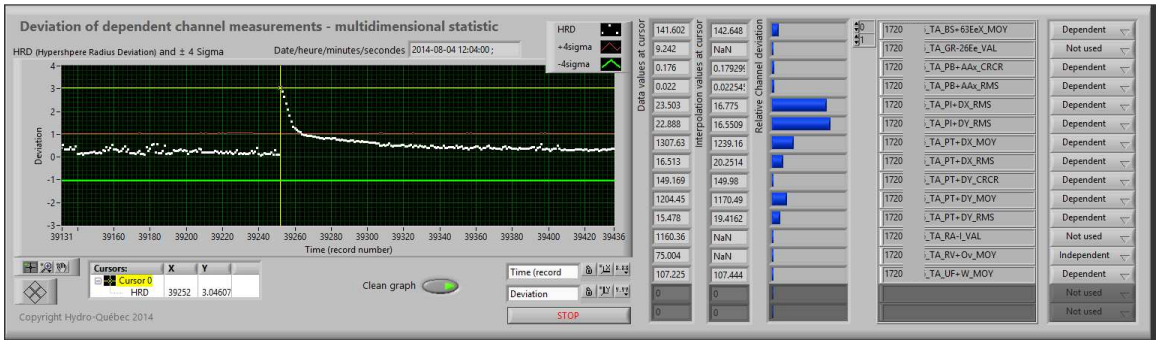

Figure 8: Momentary overrun of the 4σ threshold explained by the group's warming response following a start occurring after being stopped for several hours.

In another case study, dating back nearly 30 years on a hydraulic machine, a mechanical slack appeared on the turbine guide bearing with the result that the relative displacement measurement between the bearing and the shaft was reduced considerably since the guide bearing moved with the shaft. The increase of the vibration on the two other guide bearings was not enough to exceed the threshold of alarm: it is the manager of the power station which, when walking by the machine, noted a worrisome knocking noise. It seems relevant to consider not only increases in vibration but also abnormal decreases. With the proposed system, the increase as well as the decrease of an indicator has a similar statistical weight in quantifying a change in behavior.

Breakage of a thrust bearing

The damaged generating unit is a 390 MVA vertical shaft Francis runner type. The most likely interpretation of the results is that the breakage took place in two steps: a degradation of the thrust bearing followed several months later by the total destruction of the bearing Babbitt. Figure 9 shows that the relative deviation exceeds $10\sigma'$ (a deviation of roughly 2 on the y scale) after the first event and that the final breakage of the bearing leads to a relative deviation which exceeds $50\sigma'$ (a deviation of roughly 17 on the y scale), close to 15 times the alert threshold of $4\sigma'$. It should be noted that the machine underwent an emergency shutdown at the end of the snapshots because of an extremely high bearing temperature.

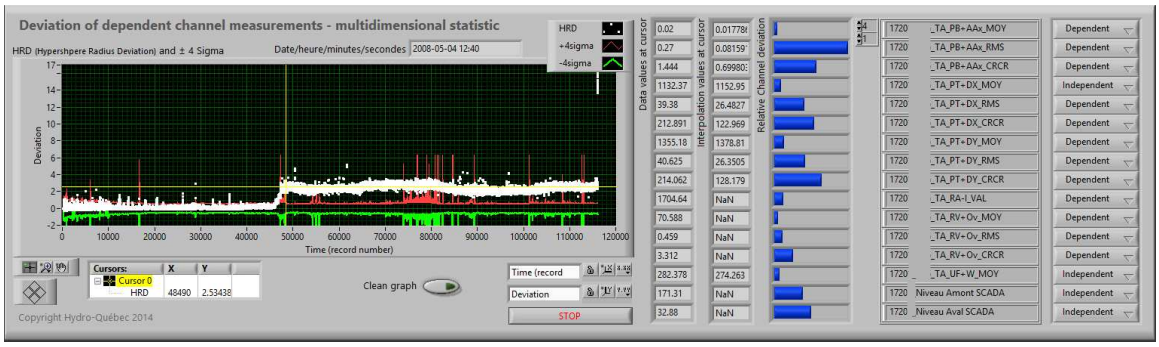

Figure 9: Relative multidimensional deviation observed over two years.

Figure 10 gives a sequential account of the machine's evolution as measured by the relative deviation and the contributions of the different indicators to this quantity. Panel (a) shows the contributions of the indicators for snapshots before the first breakage (snapshot around 45000), panel (b) gives the contributions after the first breakage (snapshot around 50000), and panel (c) shows the contributions after the total destruction of the bearing Babbitt (snapshot around 110000). The deviation from normal behaviour (see Figures 10 b and 10 c) is clearly manifest for all vibration measurements and some pressure measurements, as can be seen by the contributions of the different indicators given by the blue rectangles on the right side of these Figures. The vibration patterns appear to be similar for several of the vibration indicators, as illustrated in Figure 11. However, the most obvious signal to show signs of deterioration is associated with the vertical machine displacement, *i.e.* the vertical acceleration and dynamic water pressure in the hydraulic flow channel. For the sequence

of later measurements, the vibration level increases 5 to 6-fold from the earlier levels while an increase of roughly 50% is observed for the lateral vibration indicators at guide bearings. Since only the thrust bearing was damaged and a defect in this support system causes mainly vertical vibration, we conclude that the source of the vibration at the first breakage also came from this bearing.

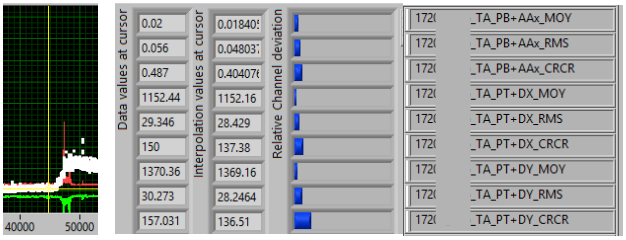

Figure 10 a: Vertical acceleration, shaft positions and vibration levels before first breakage.

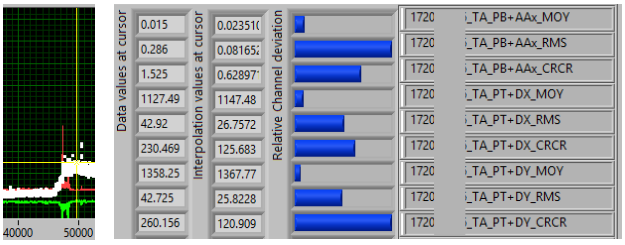

Figure 10 b: Vertical acceleration, shaft positions and vibration levels after first breakage.

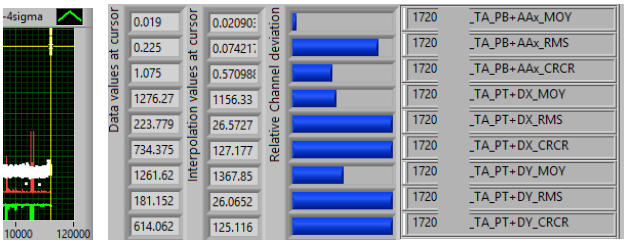

Figure 10 c: Vertical acceleration, shaft positions and vibration levels after the final breakage.

It needs to be pointed out that the alarm thresholds of the machine's protection system are adjusted for starting conditions and for vortex rope operating condition (small wicket gate opening): the vibration levels and observed temperatures between the two breakages did not exceed the aforementioned alarm thresholds. This indicates clearly the advantage of treating many indicators simultaneously in a global deviation measure. Taken individually, each indicator approaches an alert threshold without exceeding it, while when they are grouped together, the indicators lead to a global deviation 10 standard deviations from the historical behaviour.

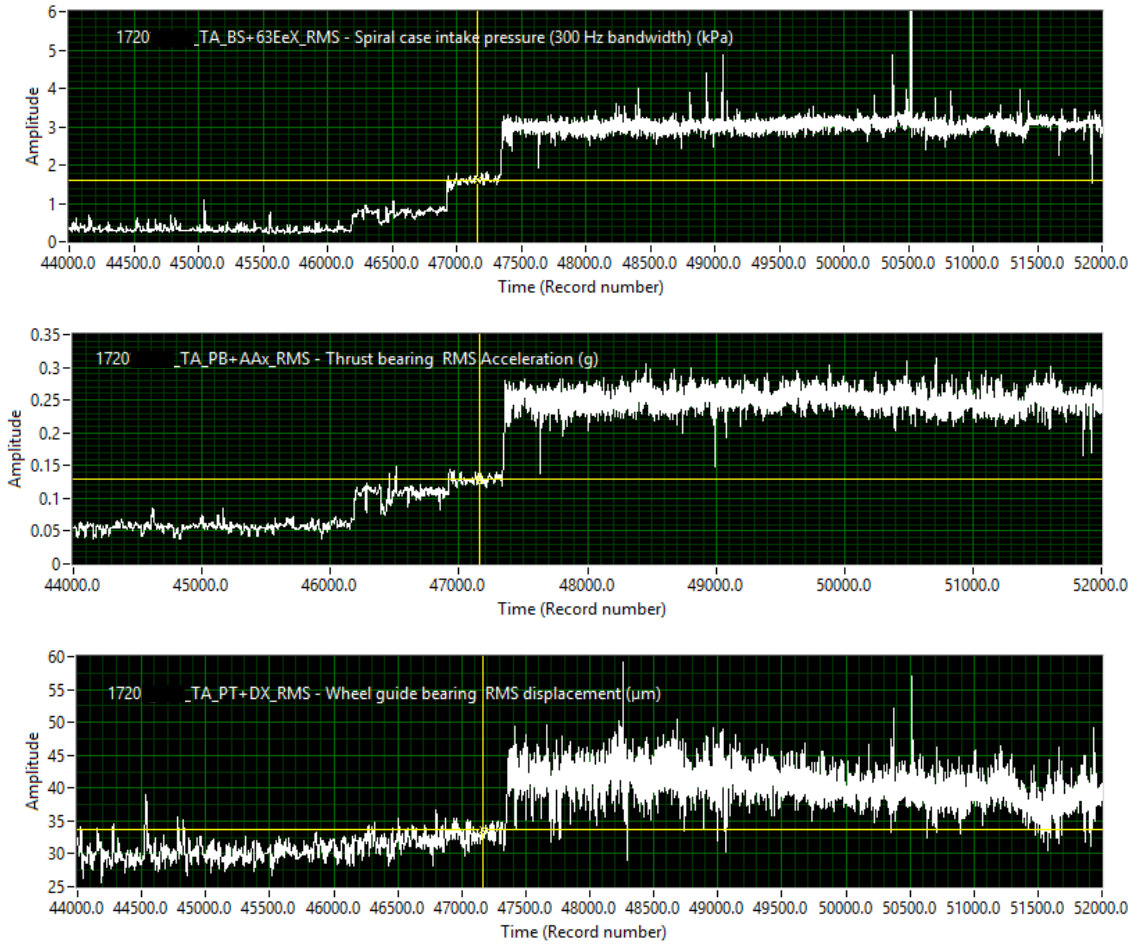

Figure 11: RMS amplitude of pressure at the inlet of the spiral case (top), acceleration (middle) and displacement at the turbine guide bearing (bottom).

Conclusion

We have presented a new global deviation indicator to monitor a machine's behaviour and detect abnormal functioning in order to alert an operator. This global deviation indicator, which combines information from several indicators, gives a result that is statistically sound and more sensitive in order to quantify deviations from normal behaviour. Furthermore, both an increase and a decrease of an indicator contribute a statistical weight similar with regards to the quantification of a change in behaviour: a failure mode which manifests itself by a decrease, or an increase, in one or several of the indicators will contribute to increase the global deviation indicator.

With the proposed global deviation indicator, exceedances due to machine starts and abnormal machine behaviour are better identified which leads to a better monitoring solution. At the present time, the methodology has been used to process collected historical data, but the next task is to apply it as a monitoring system in an operational setting with real-time observations. In future work, some aspects of the data validating module will be improved to avoid occasional false alerts due to outliers or erroneous data. Furthermore, with the accumulation of case studies, defects identified by the proposed monitoring methodology will be analysed in order to provide probable explanation for the failure mode.

References :

[1] Rafael E. Bourguet and Panos J. Antsaklis, “Artificial Neural Networks in Electric Power Industry”, Technical Report of the ISIS Group at the University of Notre Dame, ISIS-94-007, April, 1994.
URL : https://www3.nd.edu/~pantsakl/Publications/161-ISIS94.pdf

[2] Asgeir Øen Åsnes, “Condition Monitoring of Hydroelectric Power Plants”, Master of Science memories, Norwegian University of Science and Technology, May 2018.

[3] Francois. Léonard, “Dynamic clustering of transient signals”, patent US 2014/0100821, priority 2011.

[4] J. B. MacQueen, “Some Methods for classification and Analysis of Multivariate Observations” in *Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability* , no. 1, p. 281–297, 1967.

[5] Alexei V. Nikitin, Mark G. Frei, Naresh C. Bhavaraju, Ivan Osorio, Ruslan Davidchack, “Method, computer program, and system for automated real-time signal analysis for detection, quantification, and prediction of signal changes”, patent US 6,904,390 B2, priority 2000.

[6] Jason E. Lakomiak, Gilles Lanthier, “Condition monitoring parameter normalization system and method”, patent US 8,290,630 B2, priority 2008.

[7] G. Matheron, “Principles of geostatistics”, Economic Geology, 58, p. 1246–1266, 1963.

[8] I. Joliffe, “Principal Component Analysis”, Springer Verlag New York, Inc, 1986.

[9] Keith Landells; Zaid Rawi, “Abnormal event detection using principal component analysis”, US 8,121,817 B2, 2012.

[10] François Léonard. “Noisy data clusters are hollow”, Joint Statistical Meetings 2015, Seattle, JSM 2015 Proceedings, August 2015.