


TECHNICAL REVIEW

NO. 1 – 2017



Removal of Incoherent Noise from an Averaged Cross-spectral Matrix

Active Vibration-based SHM System for Wind Turbine Blades:
Demonstration on an Operating Vestas V27 Wind Turbine

Brüel & Kjær 

BEYOND MEASURE

Previously issued numbers of Brüel & Kjær Technical Review

- 1 – 2016 Flight Test Validation of Noise Models for High Performance Military Aircraft Using Beamforming
Use of a Portable Flanged Impedance Tube for Absorber Design and Measurement
- 1 – 2015 Fast Wideband Acoustical Holography
Evaluation of SPC and BSS for Indoor Pass-by Noise Contribution Analysis
- 1 – 2014 Experimental Characterization of Operating Bladed Rotor Using HPS and SSI Techniques
Microphone Acoustic Impedance in Reciprocity Calibration
- 1 – 2013 Noise Test of Revised Notched Nozzle Using a Jet Engine
Heat Conduction Correction in Reciprocity Calibration of Laboratory Standard Microphones
- 1 – 2012 High-resolution Fly-over Beamforming
Clustering Approaches to Automatic Modal Parameter Estimation
- 1 – 2011 Performance Investigation of the Dual-Layer Array (DLA) at Low Frequencies
Calculating the Sound Field in an Acoustic Intensity Probe Calibrator – A Practical Utilisation of Boundary Element Modelling
Multi-field Microphone – When the Sound Field is Unknown
- 1 – 2010 Time Selective Response Measurements – Good Practices and Uncertainty
Measurement of Absorption Coefficient, Radiated and Absorbed Intensity on the Panels of a Vehicle Cabin using a Dual Layer Array with Integrated Position Measurement
ISO 16063 – 11: Uncertainties in Primary Vibration Calibration by Laser Interferometry – Reference Planes and Transverse Motion
- 1 – 2009 Use of Volume Velocity Sound Sources in the Measurement of Acoustic Frequency Response Functions
Turnkey Free-field Reciprocity System for Primary Microphone Calibration
- 1 – 2008 ISO 16063–11: Primary Vibration Calibration by Laser Interferometry: Evaluation of Sine Approximation Realised by FFT
Infrasound Calibration of Measurement Microphones
Improved Temperature Specifications for Transducers with Built-in Electronics
- 1 – 2007 Measurement of Normal Incidence Transmission Loss and Other Acoustical Properties of Materials Placed in a Standing Wave Tube
- 1 – 2006 Dyn-X Technology: 160 dB in One Input Range
Order Tracking in Vibro-acoustic Measurements: A Novel Approach
Eliminating the Tacho Probe
Comparison of Acoustic Holography Methods for Surface Velocity Determination on a Vibrating Panel

(Continued on cover page 3)

Technical Review

No. 1 – 2017

Contents

Removal of Incoherent Noise from an Averaged Cross-spectral Matrix 1

Dr Jørgen Hald

Active Vibration-based SHM System for Wind Turbine Blades: Demonstration on an Operating Vestas V27 Wind Turbine..... 25

Dmitri Tcherniak; Lasse L. Mølgaard

TRADEMARKS

PULSE is a trademark of Brüel & Kjær Sound & Vibration Measurement A/S

Copyright © 2017, Brüel & Kjær Sound & Vibration Measurement A/S

All rights reserved. No part of this publication may be reproduced or distributed in any form, or by any means, without prior written permission of the publishers. For details, contact:

Brüel & Kjær Sound & Vibration Measurement A/S, DK-2850 Nærum, Denmark.

Editor: Harry K. Zaveri

Removal of Incoherent Noise from an Averaged Cross-spectral Matrix

*Dr Jørgen Hald**

Abstract

Measured cross-spectral matrices (CSMs) from a microphone array will in some cases be contaminated by severe incoherent noise signals in the individual channels. A typical example is flow-noise generated in the individual microphones when measuring in a wind tunnel. Assuming stationary signals and performing long-time averaging, the contamination will be concentrated on the CSM diagonal. When the CSM is used for traditional frequency-domain beamforming, Diagonal Removal (DR) will avoid the use of the diagonal. DR is effective at suppressing the contamination effects, but it also has some side effects. With other beamforming algorithms and in connection with acoustic holography, however, the diagonal of the CSM is needed. This paper describes a method for removal of incoherent noise contamination from the CSM diagonal. The method formulates the problem as a Semidefinite Program, which is a convex optimization problem that can be solved very efficiently and with guaranteed convergence properties. A first numerical study investigates the question, whether the Semidefinite Program formulation will provide the desired output in all cases. A second numerical study investigates the limitations introduced by off-diagonal noise contributions due to finite averaging time. The findings of that study are backed up by results from a practical measurement.

Résumé

Les matrices interspectrales (cross-spectral matrices = CSM) calculées au moyen d'une antenne microphonique peuvent parfois être contaminées par des signaux de bruit incohérents apparaissant dans telle ou telle voie de mesurage. Ce peut par exemple être les bruits de flux générés dans les microphones lors de mesurages en

* Brüel & Kjær Sound & Vibration Measurement A/S, Skodsborgvej 307, DK-2850 Nærum, Denmark, Jorgen.Hald@bksv.com

soufflerie. Quand des signaux stationnaires sont moyennés sur une longue durée, cette contamination se concentre sur la matrice diagonale. Lorsque les CSM sont utilisées dans le cadre d'un mesurage beamforming traditionnel dans le domaine fréquentiel, l'outil Diagonal Removal (DR) permet d'éviter l'utilisation de la diagonale. Cet outil, efficace pour supprimer les effets de la contamination, s'accompagne toutefois d'effets indésirables. Et si ce sont d'autres algorithmes beamforming qui sont utilisés, et ce, conjointement avec une holographie acoustique, la matrice diagonale devient nécessaire. La présente communication décrit une méthode de suppression de la contamination de la matrice diagonale due aux signaux de bruit incohérents. Cette méthode fait intervenir, sous forme d'un Programme Semi-défini, une formulation telle du problème que celui-ci devient un problème d'optimisation convexe se résolvant très efficacement et avec des propriétés de convergence garanties. Une première approche chiffrée répond à la question de savoir si la formulation en Programme Semi-défini peut fournir le résultat souhaité dans tous les cas. Une seconde approche chiffrée étudie les limites qu'engendrent les contributions de bruit hors-diagonale, et qui sont liées au fait que le moyennage est de durée finie. Les conclusions de cette étude sont illustrées et complétées par les résultats d'un mesurage réel.

Zusammenfassung

Gemessene Kreuzspektralmatrizen (Cross-Spectral Matrices, CSM) von einem Mikrofonarray sind manchmal mit stark inkohärenten Rauschsignalen in den einzelnen Kanälen kontaminiert. Ein typisches Beispiel ist Strömungsrauschen, das bei Windkanalmessungen in den Mikrofonen erzeugt wird. Unter der Voraussetzung, dass stationäre Signale vorliegen und eine Langzeitmittelung erfolgt, konzentriert sich die Kontamination auf die Diagonale der CSM. Bei herkömmlichem Beamforming im Frequenzbereich kann durch Diagonal Removal (DR) die Verwendung der Diagonalen vermieden werden. DR unterdrückt effektiv die Kontaminationseinflüsse, hat jedoch Nebenwirkungen. Mit anderen Beamforming-Algorithmien und in Verbindung mit akustischer Holographie wird die Diagonale der CSM jedoch benötigt. Dieser Artikel beschreibt ein Verfahren, mit dem sich inkohärente Rauschkontamination aus der CSM-Diagonale entfernen lässt. Das Verfahren formuliert das Problem als semidefinites Programm, das heißt, ein konvexes Optimierungsproblem, das sehr effizient und mit garantierten Konvergenzeigenschaften gelöst werden kann. Eine erste numerische Studie untersucht die Frage, ob die Formel des semidefiniten Programms in allen Fällen den gewünschten Output ergibt. Eine zweite numerische Studie untersucht die

Einschränkungen, die sich aufgrund der endlichen Mittelungszeit durch Rauschbeiträge außerhalb der Diagonale ergeben. Die Ergebnisse dieser Studie werden durch Ergebnisse einer praktischen Messung unterstützt.

1. Introduction

Microphone array measurements performed outdoors, or in wind tunnels, will often suffer from self-induced flow-noise in the individual microphones [1]. Use of windscreens can reduce the effect, but not remove it. For frequency-domain beamforming applications, averaging of a Cross-Spectral Matrix (CSM) is typically performed first. If the measured flow-noise signals are (to a good approximation) independent stationary stochastic processes over the measured time interval, and if the averaging time is sufficiently long, then the flow-noise will contribute almost exclusively on the diagonal of the averaged CSM. Use of windscreens will reduce the noise, but if the screens are not much smaller than the microphone spacing, then the noise from one windscreen may be picked up by nearby microphones, resulting in contributions outside the CSM diagonal.

Assuming the flow-noise contributions are concentrated on the CSM diagonal, their effect on beamforming results can be almost completely removed by use of so-called Diagonal Removal (DR), where use of the diagonal elements is avoided [1]. This technique has a simple and robust implementation in connection with traditional frequency-domain beamforming, typically Delay and Sum. Unfortunately, it also has some side-effects, such as underestimation of source levels and areas with negative power in produced contour maps. With other beamforming algorithms, such as Functional Beamforming [2], and in connection with acoustic holography, the diagonal of the CSM is needed.

Beamforming deconvolution techniques based on the use of the so-called Point Spread Function (PSF), such as Non-Negative Least Squares [3], can be implemented with DR, because the PSF can be calculated using DR. Clean-SC [4] also includes a DR procedure, which is much more complicated, requiring an iterative solution of a system of non-linear equations for each new identified point source. In Reference [5] the iterative solution procedure was found to work well for the strong sources but failed to converge for the weaker sources. To overcome that limitation, the Diagonal Denoising (DD) algorithm of the present paper was adopted. The method was first introduced in Reference [6]. Compared with Reference [6], this paper includes simulated measurements of incoherent random flow-noise time signals, which allows the effects of residual off-diagonal elements

of CSMs, after a finite number of averages, to be investigated. These simulated measurements are used in support of the experimental results.

Dougherty [7] recently published a paper describing a method entitled Cross Spectral Matrix Diagonal Optimization, which uses the same basic idea as the method presented in this paper. The idea is to reduce as much as possible the auto-power elements on the CSM diagonal, while maintaining the matrix positive semidefinite (that is, with non-negative eigenvalues). All off-diagonal (cross-power) elements remain unchanged. Dougherty did not, however, identify the problem as being a Semidefinite Program problem that can be solved efficiently and with guaranteed convergence properties using Convex Optimization, see for example References [8] and [9]. The DD method in this paper employs Convex Optimization. Compared to Reference [7], this paper also contains a numerical investigation of the conditions under which the DD method removes all the incoherent noise added on the CSM diagonal and nothing else. There seems to be no mathematical theory to support this at the moment.

The DD method is outlined in Section 2. Section 3 presents a numerical investigation of the mentioned conditions for precise removal of noise contributions added only on the CSM diagonal. Based on simulated measurements, Section 4 investigates the effects of residual off-diagonal flow-noise contributions, due to the use of a finite averaging time. Results from comparable actual measurements with flow-noise are presented in Section 5, and Section 6 provides a summary.

2. Theory

2.1. *Basic Properties of Cross-spectral Matrices*

This section gives a short introduction on some important properties of CSM's, which will be useful in this paper. Cross-spectra are typically calculated from recorded time series using the Welch method, where weighting (for example Hanning) and Fast Fourier Transform (FFT) is applied to a sequence of data 'records' from overlapping time intervals, and an averaging is performed over the intervals. Denoted by \mathbf{p} , a vector of complex pressure values across all M array microphones from FFT's applied to a single time interval and with a single frequency line picked out, the CSM \mathbf{G} for that frequency line is calculated as:

$$\mathbf{G} = \overline{\mathbf{p}\mathbf{p}^H} \quad (1)$$

Here, H represents Hermitian (conjugate) transpose, and the bar over the right side represents averaging over time intervals. Clearly, the matrix is Hermitian, that is, conjugate symmetric.

Considering the contributions in \mathbf{G} from stationary stochastic signals that are incoherent between the microphones, all off-diagonal contributions will approach zero during the averaging process. After long-time averaging, the contributions can be considered to be only on the diagonal, but as we shall see in Section 4, this requires a lot of averages.

With \mathbf{G} calculated according to Eq. 1, \mathbf{G} will always have non-negative eigenvalues, that is, be positive semidefinite. This can, for example, be seen by arranging all the vectors \mathbf{p} as columns in a matrix \mathbf{P} , which allows \mathbf{G} to be written as $\mathbf{G} = \alpha \mathbf{P} \mathbf{P}^H$, where α equals one divided by the number of averages. Introduction of the Singular Value Decomposition of \mathbf{P} , $\mathbf{P} = \mathbf{U} \mathbf{S} \mathbf{V}^H$, where \mathbf{U} and \mathbf{V} are unitary matrices and \mathbf{S} a diagonal matrix, leads to:

$$\mathbf{G} = \alpha \mathbf{P} \mathbf{P}^H = \alpha \mathbf{U} \mathbf{S} \mathbf{V}^H \mathbf{V} \mathbf{S}^H \mathbf{U}^H.$$

Since \mathbf{V} is unitary, $\mathbf{V}^H \mathbf{V}$ equals the unit diagonal matrix, and therefore the expression can be written as:

$$\mathbf{G} = \mathbf{U} \mathbf{\Sigma} \mathbf{U}^H \quad (2)$$

with $\mathbf{\Sigma} \equiv \alpha \mathbf{S} \mathbf{S}^H$. Eq. 2 is an eigenvalue/-vector decomposition of \mathbf{G} . The eigenvalues σ_n , $n = 1, 2, \dots, M$, on the diagonal of $\mathbf{\Sigma}$ are equal to α times the squared amplitudes of the singular values from the diagonal of \mathbf{S} , showing that they are always non-negative.

Representing by \mathbf{u}_n , $n = 1, 2, \dots, M$, the columns of the matrix \mathbf{U} , the vectors

$$\mathbf{p}_n \equiv \sqrt{\sigma_n} \mathbf{u}_n \quad (3)$$

are the so-called Principal Components related to the individual eigenvalues σ_n . These vectors are orthogonal, and together they represent the total measured sound pressure across the array:

$$\mathbf{G} = \sum_{n=1}^M \mathbf{p}_n \mathbf{p}_n^H \quad (4)$$

which can be verified through application of Eqs. 2 and 3. Clearly, the principal components contribute independently (additively) to all elements of \mathbf{G} , showing that they are mutually incoherent. A single principal component is coherent between all microphones, which can easily be verified from its partial CSM: $\mathbf{p}_n \mathbf{p}_n^H$.

Considering a special case with only coherent sources, the snapshot vectors \mathbf{p} in Eq. 1 will all be parallel (they differ only through a complex scaling factor), leading to a single non-zero eigenvalue σ . Considering another special case of a single loudspeaker excited simultaneously by several independent signal components, these signal components will contribute with parallel components to the snapshot vectors \mathbf{p} because the related steering vectors (transfer function vectors) are parallel. Therefore, in this case, all the vectors \mathbf{p} will also be parallel, producing a single non-zero eigenvalue σ . In general, the number K of non-zero eigenvalues is equal to the number of incoherent sources that can be distinguished by the array, based on the averaged CSM.

2.2. Description of the Denoising Algorithm

Let \mathbf{G} be an $M \times M$ element CSM contaminated by microphone self-noise contributions only on its diagonal. We would like to add an $M \times M$ element diagonal matrix \mathbf{D} with a priori unknown non-positive diagonal elements d_m , $m = 1, 2, \dots, M$, to cancel the self-noise. Arranging the unknown diagonal elements in an M -element vector \mathbf{d} , we write $\mathbf{D} = \mathbf{diag}(\mathbf{d})$. After the determination of \mathbf{D} , we will replace \mathbf{G} by $\mathbf{G} + \mathbf{D}$ in subsequent array processing. Cancelling the self-noise on the diagonal implies a reduction of the sum of the elements d_m . A lower limit is set by the fact that the remaining denoised CSM, $\mathbf{G} + \mathbf{D}$, must still be positive semidefinite. Therefore, in mathematical terms, the idea is to determine \mathbf{d} as the solution to:

$$\min_{\mathbf{d}} \sum_m d_m \quad \text{subject to} \quad \mathbf{G} + \mathbf{diag}(\mathbf{d}) \geq 0 \quad (5)$$

where ‘ ≥ 0 ’ for a matrix means that it has non-negative eigenvalues (is positive semidefinite). Eq. 5 has the form of a so-called Semidefinite Program, which can be solved efficiently and with guaranteed convergence properties using Convex Optimization methods, see for example, Reference [8].

Shown below are the five lines of Matlab code needed to solve Eq. 5 using the publicly available Matlab library CVX:

```
cvx_begin;
variable d(M);
G + diag(d,0) == hermitian_semidefinite(M);
minimize( sum(d) );
cvx_end;
```

The library can be downloaded from the website given in Reference [9].

An important question is whether the formulation above will always provide an output with exactly the self-noise removed from the diagonal. The minimization in Eq. 5 will proceed until one of the eigenvalues of $\mathbf{G}+\mathbf{D}$ reaches zero, showing that too much will be removed, if the desired matrix has only positive eigenvalues. To see this, assume the smallest eigenvalue of the desired noise-free target matrix $\mathbf{G}_0 = \mathbf{U}\Sigma_0\mathbf{U}^H$ is a positive number, $\epsilon > 0$. Then we can subtract ϵ from all eigenvalues and still be left with a positive semidefinite matrix. This subtraction corresponds to a subtraction of $\epsilon\mathbf{I}$ from the diagonal eigenvalue matrix Σ_0 , \mathbf{I} being the unit diagonal matrix. For \mathbf{G}_0 the result of the subtraction is:

$$\mathbf{U}[\Sigma_0 - \epsilon\mathbf{I}]\mathbf{U}^H = \mathbf{U}\Sigma_0\mathbf{U}^H - \epsilon\mathbf{U}\mathbf{U}^H = \mathbf{G}_0 - \epsilon\mathbf{I},$$

showing that ϵ has just been subtracted from all diagonal elements. The DD algorithm can therefore remove at least that amount from the diagonal of \mathbf{G}_0 , while still keeping it positive semidefinite. Repeated numerical checks have shown that at least one eigenvalue equals zero after use of the DD algorithm. The following section presents some numerical simulations to further investigate this and related issues.

3. Investigation Based on Synthesized Cross-spectral Matrices

The algorithm was tested by synthesis of $M \times M$ element CSM's from K random rank-1 matrices, $K \leq M$, representing K incoherent contributions:

$$\tilde{\mathbf{G}} = \sum_{k=1}^K \mathbf{q}_k \mathbf{q}_k^H \quad (6)$$

Here, the real and imaginary parts of each element in the $M \times 1$ vectors \mathbf{q}_k were generated as random numbers with zero mean and standard deviation equal to 1. Unlike a set of principal components, the vectors \mathbf{q}_k are not orthogonal. The resulting matrix $\tilde{\mathbf{G}}$ will in general have a set of K non-zero eigenvalues covering a wide range. To shape the eigenvalue spectrum, an eigenvalue/vector decomposition was made first:

$$\tilde{\mathbf{G}} = \mathbf{U}\tilde{\Sigma}\mathbf{U}^H \quad (7)$$

followed by a scaling of the K non-zero eigenvalues in $\tilde{\Sigma}$ to have a desired distribution, resulting in the diagonal matrix Σ_0 . The target ‘noise-free’ CSM \mathbf{G}_0 is then calculated as:

$$\mathbf{G}_0 \equiv \mathbf{U}\Sigma_0\mathbf{U}^H \quad (8)$$

Ideally, the DD algorithm should be able to recover the matrix \mathbf{G}_0 from a matrix \mathbf{G} with noise added on its diagonal:

$$\mathbf{G} = \mathbf{G}_0 + \mathbf{diag}(\mathbf{n}) \quad (9)$$

\mathbf{n} being a vector of squared, suitably scaled, random numbers. The convex optimization (Eq. 5) should have the solution $\mathbf{d} = -\mathbf{n}$.

Subsection 3.1 that follows investigates the error in the case where the matrix \mathbf{G}_0 has a variable number of equal, non-zero eigenvalues (the rest being equal to zero), while Subsection 3.2 investigates the probably more typical case of a decaying eigenvalue spectrum with none of the eigenvalues equal to zero.

3.1. Varying Number of Equal, Non-zero Eigenvalues

In this case, the largest eigenvalue $\tilde{\sigma}_{\max}$ in $\tilde{\Sigma}$ was first identified, and the eigenvalues $\sigma_{0,m}$ in Σ_0 were then set as:

$$\begin{aligned} \sigma_{0,m} &= \tilde{\sigma}_{\max} & m &= 1,2,\dots,K \\ \sigma_{0,m} &= 0 & m &= K+1,K+2,\dots,M \end{aligned} \quad (10)$$

Based on Σ_0 , the matrix \mathbf{G}_0 could be computed using Eq. 8. This synthesis of \mathbf{G}_0 simulates a measurement with K incoherent sources contributing equally across the array microphones.

The elements of the noise vector \mathbf{n} were initially generated as squared random variables with mean zero and standard deviation one. Subsequently, the vector \mathbf{n} was scaled to have an average element size a chosen number of decibels, X , higher than the average auto-power on the diagonal of \mathbf{G}_0 . Finally, \mathbf{G} could be generated through application of Eq. 9.

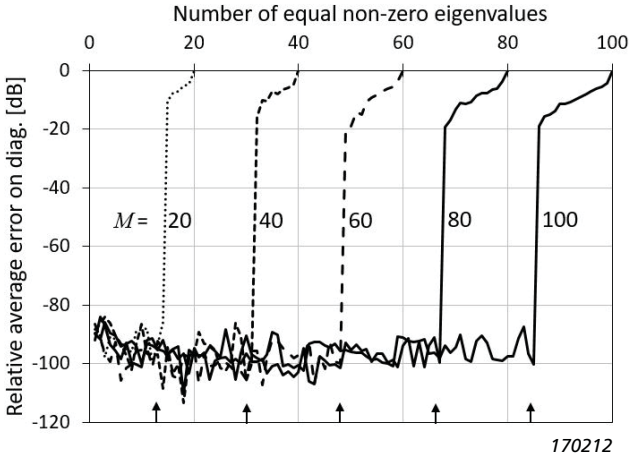
After the solution of the DD (Eq. 5), the deviation from the ideal solution $\mathbf{d} = -\mathbf{n}$ was quantified by the following relative average error function:

$$Err \equiv \frac{\sum_{m=1}^M |d_m + n_m|}{\sum_{m=1}^M G_{0,m,m}} \quad (11)$$

where d_m are the elements of \mathbf{d} , n_m are the elements of \mathbf{n} , and $G_{0,m,m}$ are the diagonal elements of \mathbf{G}_0 . Clearly, if Err equals zero, then the ideal solution has been reached with all auto-power values precisely reconstructed. Fig. 1 shows this relative average error in decibels as a function of the number K of equal, non-zero eigenvalues for five different values of the matrix dimension M used as annotation at the curves: $M = 20, 40, 60, 80$ and 100 .

In this case, the relative level X of the added noise was set to be 40 dB, but several levels were tested with almost identical result: $X = 0, 20, 40$ dB. Therefore, apparently, it is not the level of the added noise that determines the accuracy, but rather the eigenvalue spectrum of the desired matrix \mathbf{G}_0 .

Fig. 1. For $M = 20, 40, 60, 80$ and 100 microphones, the relative average error in decibel $10 \cdot \log_{10}(Err)$, of the reconstructed diagonal is shown as a function of the number of equal, non-zero CSM eigenvalues. The arrows at the bottom indicate (for each value of M) the x-axis value K_{max} given by Eq. 12



Based on a series of simulations and results like those in Fig. 1, it was found that an exact noise subtraction can be performed effectively when the number K of equal non-zero eigenvalues does not exceed a value K_{\max} given approximately as:

$$K_{\max}(M) \cong M - \sqrt{2.5 \cdot M} \quad (12)$$

Following the reasoning in the last paragraph of Section 2.1, this condition will be met, if the number of independent sources does not exceed K_{\max} . Close to the point where K exceeds K_{\max} , the relative error increases steeply to be in the range between -20 and -10 dB, and after that point, it increases smoothly to 0 dB at $K = M$. As mentioned above, the method cannot accurately handle the situation where all eigenvalues are positive. Intuitively, one might think that the method should work perfectly as long as \mathbf{G}_0 had just one eigenvalue equal to zero. The reason why a number of eigenvalues must equal zero is not known. A guess could be finite precision in the number handling, but it remains as an interesting open question.

3.2. Decaying Eigenvalue Distribution

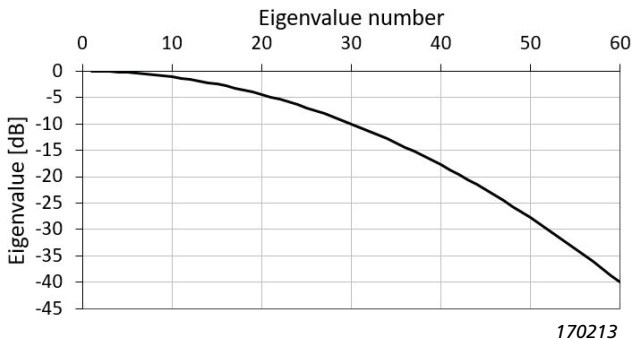
In practical applications, the desired CSM \mathbf{G}_0 will often have a more or less smoothly decaying spectrum of eigenvalues, generated by a set of independent sources with amplitudes covering a wide range. Here, we will consider eigenvalue spectra of the form shown in Fig. 2, which are synthesized from the formula:

$$\sigma_{0,m} = e^{-km^2/M}, \quad m = 1, 2, \dots, M \quad (13)$$

with k set to produce a specific attenuation, A , in decibels at $m = M$. In the case of Fig. 2, the attenuation is $A = 40$ dB. Since the largest (first) eigenvalue is very close to 1.0, A is also the range covered by the eigenvalue spectrum.

After calculation of \mathbf{G}_0 by Eq. 8, the noise vector \mathbf{n} was calculated using the same procedure as described in Section 3.1. For this, the level X of the noise relative to the average auto-power on the diagonal of \mathbf{G}_0 had to be chosen. Again, it turned out that, to a good approximation, the errors were independent of the noise level, depending almost entirely on the eigenvalue spectrum of the underlying noise-free CSM. Fig. 3 shows the relative noise removal error, calculated using Eq. 11, as a function of the eigenvalue range A for an array with $M = 60$ microphones.

Fig. 2. Simulated CSM eigenvalue distribution with 40 dB range for a 60-element array




The relative noise level was again $X = 40$ dB, that is, the noise was 40 dB higher than the average auto-power on the diagonal of \mathbf{G}_0 . Very similar error curves were found for other values of M .

The results show that the reconstruction error vanishes if a sufficient number of eigenvalues are sufficiently small, or in this case, if the eigenvalue range is sufficiently large.

An interpretation of the results in this section and in Section 3.1, is that the solution of Eq. 5 will very effectively remove noise added on the diagonal of the CSM, but depending on the eigenvalue spectrum of the noise-free matrix it will also remove/modify some of the diagonal of that matrix.

Fig. 3. Relative average reconstruction error, Err , for a 60-element array with CSM eigenvalue distribution of the type shown in Fig. 2. The eigenvalue range, A , is on the x-axis



If a sufficient number of eigenvalues of the noise-free matrix are zero, or very small, then the introduced error is negligible. This will be fulfilled, if the number of significant incoherent sources is somewhat smaller than the number of array microphones.

4. Investigation Based on Synthesized Flow-noise Time Data

The simulated measurements in Section 3 all dealt with the case where noise is added only on the diagonal of the CSM. In connection with real measurements there will be contributions also outside the diagonal. The level of these off-diagonal contributions can be minimized through long-time averaging, provided the noise is incoherent between the microphones. In this section we shall consider the case where only incoherent noise time signals are measured by the individual microphones, that is, no source signal components with coherence between microphones. With an increasing number of averages applied with the Welch method, the averaged CSM should eventually become diagonal, so all energy on the diagonal could be removed by the DD method.

The coherence function $\gamma_{m,n}^2$ between two microphones m and n is defined as (see for example Eq. (3.43) in Reference [10]):

$$\gamma_{m,n}^2 \equiv \frac{|G_{m,n}|^2}{G_{m,m} \cdot G_{n,n}} \quad (14)$$

Here, $G_{m,n}$ are the elements of the CSM. Combining that definition with Eq. (11.21) in Reference [10], we get the following expression for the variance of the cross-power ($m \neq n$) amplitudes $|G_{m,n}|$ after J averages:

$$\text{Var} \left[|G_{m,n}| \right] = \frac{G_{m,m} \cdot G_{n,n}}{J} \quad (15)$$

Since the cross power will approach zero with an increasing number of averages, the variance in Eq. 15 is equal to an expected value of the squared cross-spectrum amplitude after J averages, $\text{Var} \left[|G_{m,n}| \right] \approx |G_{m,n}|^2$, see for example Reference [10].

Eq. 15 therefore leads to the following expected value of the residual coherence $\gamma_{m,n}^2$ between two microphones:

$$\gamma_{m,n}^2 \equiv \frac{|G_{m,n}|^2}{G_{m,m} \cdot G_{n,n}} \approx \frac{\text{Var} \left[|G_{m,n}| \right]}{G_{m,m} \cdot G_{n,n}} = \frac{1}{J} \quad (16)$$

Since an averaged CSM involves many coherences (between all microphone pairs), we will be using an average coherence defined as:

$$\bar{\gamma}^2 \equiv \frac{\sum_{m \neq n} |G_{m,n}|^2}{\sum_{m \neq n} G_{m,m} G_{n,n}} \quad (17)$$

Use of Eq. 15 and the relation $\text{Var} \left[|G_{m,n}| \right] \approx |G_{m,n}|^2$ in Eq. 17 leads to the result that the expected average coherence is also approximately $1/J$ when all microphone signals are incoherent:


$$\bar{\gamma}^2 \approx \frac{1}{J} \quad (18)$$

A set of simulated measurements were performed, in which incoherent noise signals were generated by a random-number generator, with a rectangular distribution over the interval from -1 to 1 . The random-number sequence for each one of the microphones was then used as a time signal with sampling frequency equal to 32.768 kHz. Thus, all microphones were exposed to the same noise level. Cross-spectral averaging was performed with a 200-line FFT and with 50% record overlap, providing approximately 128 averages per second of time signal.

For the case of 30 microphones, Fig. 4 shows average coherence spectra $10 \cdot \log_{10}(\bar{\gamma})$ obtained from Eq. 17 after 10, 20 and 30 s of averaging.

Also shown are the expected values from Eq. 18, with which the agreement is seen to be very good. The agreement has been checked with different microphone counts and found to be very good, the main difference being that the variance on the average coherence spectrum increases with decreasing number of microphones. The explanation is almost certainly the lower amount of random time data involved in the estimation.

Fig. 4. Average coherence spectra $10 \cdot \log_{10}(\bar{\gamma})$ from the simulated measurements and from Eq. 18 after 10, 20 and 30 second averaging for a 30-element array



When applying the DD algorithm to the averaged CSM, power will be subtracted from the diagonal elements until (at least) one eigenvalue is equal to zero, see Section 2.2. Considering the case of only two microphones, this implies only one non-zero eigenvalue, meaning that the remaining matrix represents two perfectly coherent signals.

The reduced diagonal elements $\tilde{G}_{1,1}$ and $\tilde{G}_{2,2}$ of that matrix will therefore fulfil the relation $\tilde{G}_{1,1} \cdot \tilde{G}_{2,2} = |G_{1,2}|^2$ (coherence equal to one), which leads to:

$$\frac{\tilde{G}_{1,1} \tilde{G}_{2,2}}{G_{1,1} G_{2,2}} = \frac{|G_{1,2}|^2}{G_{1,1} G_{2,2}} = \gamma_{1,2}^2 = \bar{\gamma}^2 \approx \frac{1}{J} \quad (19)$$

With signals of equal level applied to the two microphones, the diagonal elements of the CSM should be approximately equal, and for reasons of symmetry, the reduction factors achieved by DD for these elements should also be approximately equal. The reduction factor α achieved by DD for the sum of the diagonal elements can therefore be estimated by application of Eq. 19 with the following result:

$$\alpha \equiv \frac{\sum_m \tilde{G}_{m,m}}{\sum_m G_{m,m}} = \frac{\tilde{G}_{1,1} + \tilde{G}_{2,2}}{G_{1,1} + G_{2,2}} \approx \sqrt{\frac{\tilde{G}_{1,1} \tilde{G}_{2,2}}{G_{1,1} G_{2,2}}} = \tilde{\gamma} \approx \frac{1}{\sqrt{J}} \quad (20)$$

Results from simulated measurements have confirmed that the relative reduction α of the diagonal sum is almost identically equal to the average coherence $\tilde{\gamma}$ in the case of two microphones.


With a larger number of microphones, however, the relative reduction α was found to be smaller than $\tilde{\gamma}$. The difference turned out to increase with increasing microphone count, but to be almost independent of the number of averages. Fig. 5 shows the relative reduction α for the conditions represented in Fig. 4: 30 microphones with 10, 20 and 30 s of averaging.

Clearly, the reduction is significantly smaller than predicted by $\tilde{\gamma}$. The loss in reduction $10 \log_{10}(\alpha/\tilde{\gamma})$ is presented in Fig. 6 as a function of microphone count together with the following simple empirical approximation:

$$\alpha/\tilde{\gamma} \approx (M - 1)^{0.625} \quad (21)$$

The form of this approximation can be derived by assuming identical diagonal elements G , by assuming equal reduction by DD of these to become αG , and by

Fig. 5. Relative reduction $10 \cdot \log_{10}(\alpha)$ of the diagonal sum in the CSM obtained with DD for the case of 30 microphones and with 10, 20 and 30 s of averaging time



assuming that some generalized power must be retained on the diagonal for all the corresponding generalized power that exists off the diagonal. According to Eq. 16, all off-diagonal elements will, in that case, have amplitudes approximately equal to G/\sqrt{J} , so the generalized power assumption is formulated as:

$$M \cdot (\alpha G)^\mu \approx M(M-1) \cdot (G/\sqrt{J})^\mu \tag{22}$$

since there are M diagonal elements and $M(M-1)$ off-diagonal elements.

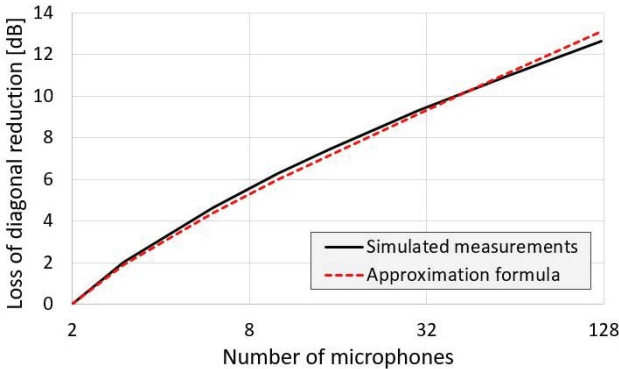
With $\mu=1.0$, both sides of the equation represent a simple power sum, while with $\mu=2.0$, they represent a sum of squared power. With $\mu=1.6$, Eq. 22 leads directly to Eq. 21. This is in no way a derivation. It just explains how the idea for the empirical expression (21) came up.

From Eqs. 18 and 21 we get the following estimate for the relative reduction α of noise on the CSM diagonal that can be achieved by DD:

$$\alpha \approx \frac{(M-1)^{0.625}}{\sqrt{J}} \tag{23}$$

In case only a subset of the array microphones is exposed to incoherent noise signals and these noise signals are of equal level, only the number of exposed microphones should be used in the above formulae. This can be verified by going through the derivations. The practical measurement to be presented in the following section has a rather small subset of microphones exposed to strong flow

Fig. 6. Diagonal reduction relative to the average coherence (entitled ‘loss of diagonal reduction’) as a function of the number of microphones



170217

and therefore to high-level flow-noise.

Considering a wind tunnel measurement, with all microphones subject to equal levels of self-induced flow-noise, we can use Eq. 23 to estimate the number of averages J required to achieve a reduction by a given factor α of the flow-noise contribution on the CSM diagonal:

$$J \approx \alpha^{-2}(M - 1)^{1.25} \quad (24)$$

Requiring for example a reduction by a factor $\alpha=0.1$ (10 dB) for an array with 100 microphones, the required number of averages will be approximately 31228, corresponding to an averaging time around 240 s, when using the Welsh method, with the parameters described earlier in this section.

In Section 3, the case of noise added only on the diagonal was treated. The presented data indicated that all such added noise will be removed by DD. However, unless the number of significant independent target sources is smaller than $M - \sqrt{2.5 \cdot M}$, some of the auto-power related to the target source signals will also be removed. This section 4 has considered the case of no target source, that is, only incoherent noise, but the effect of related residual off-diagonal contributions in the CSM due to finite averaging has been investigated. In practical applications both target source signals and flow-noise signals with residual off-diagonal contributions will be present. The residual off-diagonal flow-noise contributions will introduce signal components in the CSM, which will be ‘seen’ as target signals by DD and may bring the number of significant ‘target signals’ above the $M - \sqrt{2.5 \cdot M}$ limit, causing DD to remove parts of these signals. Therefore, part of the desired target-signal auto power may also be removed. This has not been investigated further.

5. Measurements

A series of measurements were taken with a 30-element flush-mounted pseudo-random array in front of a Brüel & Kjær Mouth Simulator Type 4227 (speaker), excited by white random noise. Fig. 7 shows the setup with the 30 cm diameter array at 40 cm distance from the sound source, and with the flow-outlet nozzle visible at the bottom right corner of the picture.

The nozzle emitted a flow-beam towards a bottom section of 5 to 7 microphones of the array. That section was exposed to strong close-to-laminar flow, while the flow speed was much lower at the remaining microphones. Recordings of 30 s length, with 32.768 kHz sampling frequency, were taken with and without

Fig. 7. Picture of the setup with the array (right), sound source (left) and flow-beam hose (bottom right)



the flow. Using a 200-line FFT for cross-spectral averaging, 3839 averages could be achieved with 50% overlap between records.

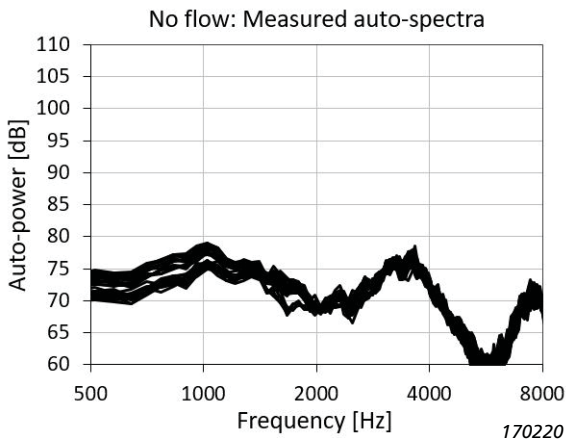
Measurements were performed under two sets of conditions: ‘speaker plus flow’; and ‘speaker only’. A third measurement should have been performed using ‘speaker off’ conditions (that is, where only the flow noise was present) as will be apparent in the following. The third measurement, however, was not taken, and subsequently the flow source has not been available.

Since the speaker signal and the flow-induced noise in the microphones are mutually incoherent, the measured CSM, \mathbf{G} , will ideally be the sum of their contributions: $\mathbf{G} = \mathbf{G}_0 + \mathbf{G}_{\text{flow}}$, \mathbf{G}_0 being the contribution from the speaker. However, due to finite averaging time, the equation will only hold approximately. With the first two of the three matrices measured, the flow noise contribution \mathbf{G}_{flow} could be calculated. However, you have to be aware of the previously mentioned approximate nature of the above equation, plus the fact that \mathbf{G}_{flow} is obtained as a difference between two non-simultaneously measured matrices.

Fig. 8 shows the 30 auto-power spectra (black curves) from the measurement with no flow, while the corresponding spectra with flow are shown in Fig. 9.

Clearly, the flow has generated some very strong broadband noise in the five ‘core’ microphones (indicated in red in Fig. 9), which are in the core beam and close to the nozzle.

Fig. 8. Auto-power spectra from all 30 microphones measured without flow



Six microphones have medium-level flow contributions (indicated by green dashed curves). These ‘peripheral’ microphones were in the outskirts of the flowbeam. The remaining microphones (indicated in black) exhibit insignificant flow-noise contributions in the displayed frequency range. Figs. 10 and 11 display auto-power spectra after application of the DD algorithm to the measurement with flow.

Fig. 9. Auto-power spectra measured with flow.

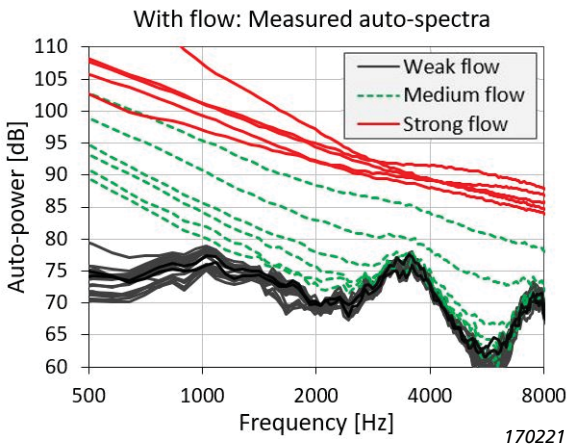


Fig. 10. Auto-power spectra measured with flow and subsequently denoised by applying DD to the entire averaged CSM. **Red:** Microphones in core flow with high flow-noise levels. **Dashed green:** Microphones in outskirts of core flow. **Black:** Microphones with insignificant flow noise


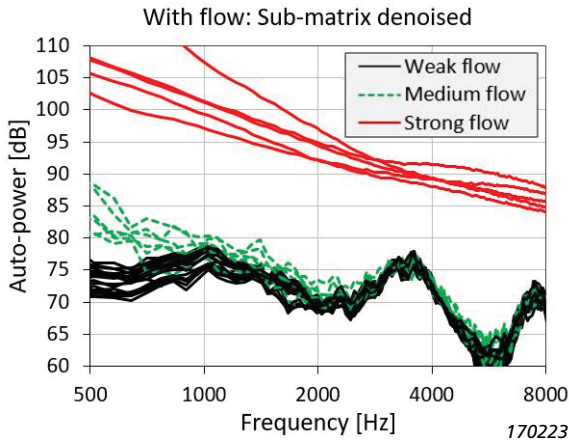


Fig. 11. Auto-power spectra measured with flow, and subsequently a sub-matrix excluding the five microphones in core flow (red color) has been denoised using DD. **Red:** Microphones in core flow with high flow noise levels. **Dashed green:** Microphones in outskirts of core flow. **Black:** Microphones with insignificant flow noise



In the case of Fig. 10, the algorithm has been applied to the full measured CSM. The flow-induced auto-power has been significantly reduced for both the core and the peripheral microphones, but not removed, and below 1000 Hz some auto-power has even been added to the remaining microphones, which originally had virtually no flow contributions.

Such an increase of noise auto-power in some microphones could be prevented by adding constraints in the optimization process of Eq. 5. However, that would not overcome the underlying problem of the remaining off-diagonal flow-noise contributions, and it would imply a somewhat reduced effect of DD on the microphone signals with the highest flow-noise levels. To be able to compare the achieved reduction of the noise auto-power with a prediction from the empirical model of Eq. 23, we need to specify a number of microphones with equal noise exposure, the remaining having no noise. With five core microphones and six peripheral ones, a reasonable choice could be $5+6/2=8$. Use of $M=8$ in Eq. 23 leads to a predicted reduction equal to $10 \cdot \log_{10}(\alpha) = -12.6$ dB, which agrees quite well for example at 3 kHz.

The auto-power spectra in Fig. 11 are the result of applying the DD algorithm to a sub-matrix of the measured CSM: The rows and columns related to the microphones in core flow were removed before application of DD and added again subsequently. Clearly, the DD algorithm has been more successful in reducing the flow noise from all the non-core microphones because the large off-diagonal flow-noise contributions from the core microphones no longer disturb.


Fig. 12 shows the average coherence $\bar{\gamma}$ obtained by application of Eq. 17 to the estimated flow-noise CSM, G_{flow} .

The figure can be compared to the corresponding Fig. 4 for the simulated measurement. Above 1.2 kHz the agreement is very good, but below that frequency, a small increase is observed in the average coherence. The reason has not been analyzed. Possible explanations are:

- Background noise with high coherence across the array, for example, from the electric motor and the fan generating the flow
- The previously mentioned errors in the calculation of the flow-noise CSM


Fig. 13 shows the achieved reduction in decibels of the flow-noise-related auto-power sum (the sum of the diagonal elements in $G_{\text{flow}} = \mathbf{G} - \mathbf{G}_0$), when applying the DD algorithm to the full matrix \mathbf{G} . The reduction is shown in decibels for the cases, where 10, 20 and 30 s averaging has been applied, and the results are compared against the values obtained from Eq. 23, with the assumption that all flow-noise is distributed equally across eight microphones, as explained above.

Fig. 12. Average coherence $10 \cdot \log_{10}(\bar{\gamma})$ for the measured flow-noise CSM calculated using Eq. 17 after 10, 20 and 30 s of averaging. The dashed lines represent theoretical values obtained from Eq. 18



Except at frequencies below 1.2 kHz, the agreement is very good. The fact that DD was applied to \mathbf{G} instead of \mathbf{G}_{flow} apparently did not significantly degrade the flow-noise auto-power removal. A main reason could be the fact that the flow-noise in the microphones in core flow was much higher than the speaker signal.

Fig. 13. Reduction of flow-noise auto-power on the CSM diagonal, obtained using the DD method of Section 2, to the full matrix, ref. Fig. 10. The theoretical values were obtained from Eq. 23, assuming all flow-noise to be equally distributed across eight microphones



Also, the total number of significant incoherent signals (speaker, flow-noise in core microphones and flow noise in peripheral microphones) is in this case much smaller than the number of microphones. Thus, the DD method will not remove speaker-generated auto-power, despite the strong residual off-diagonal contributions from the flow-noise.

Clearly, remaining flow-noise related off-diagonal contributions in \mathbf{G} will limit the reduction of the flow-noise auto-power that can be achieved by use of the DD algorithm. The degree to which this will be visible in beam formed maps depends on the level of the flow-noise relative to the level of the source signal of interest. This was investigated in Reference [6] and will not be further explored in this paper.

6. Summary

The subject of this paper has been methods to reduce the impact on array-processing results of incoherent noise in individual measurement channels, typically flow-noise in connection with measurements in wind tunnels or outdoors. Diagonal Removal (DR) is an established technique for noise reduction with conventional frequency-domain beamforming, avoiding completely the use of the CSM diagonal. However, DR also has some limitations and with some array processing methods, it does not apply.

This paper has described a method called Diagonal Denoising (DD) to remove the incoherent noise contributions from the diagonal of the CSM. The idea is to minimize the sum of the CSM diagonal elements, while leaving all off-diagonal elements unchanged, and while keeping the CSM positive semidefinite. The problem has the form of a so-called Semidefinite Program, which can be solved very efficiently and with guaranteed convergence properties using Convex Optimization methods.

The properties and limitations of the method were first investigated by applying it to computer synthesized CSM's with noise added only on the diagonal. It turned out that with up to approximately $M - \sqrt{2.5 \cdot M}$ incoherent target sources, M being the number of microphones, DD is accurate within the numerical accuracy almost independent of the level of the noise added on the diagonal. In connection with real measurements, however, there will be off-diagonal residual contributions from the incoherent noise signals because of finite averaging time. These off-diagonal contributions will limit the auto-power noise reduction achievable by DD. Based on simulated measurements with equal noise levels added in all channels, an approximate empirical model of the impact was developed.

According to the model, the number of averages required to reduce the noise auto-power by a factor α is approximately $\alpha^{-2}(M-1)^{1.25}$.

Measurements have been performed on a small loudspeaker with a section of the array exposed to airflow. The measurement results agree well with the predictions from the empirical model.

7. References

- [1] R.P. Dougherty. *Aeroacoustic Measurements: Beamforming in Acoustic Testing*, edited by T.J. Mueller (Springer-Verlag Berlin Heidelberg 2002, pp. 83-86).
- [2] R.P. Dougherty. *Functional Beamforming*, Proceedings of Berlin Beamforming Conference 2014, paper BeBeC-2014-01.
- [3] K. Ehrenfried and L. Koop. *A Comparison of Iterative Deconvolution Algorithms for the Mapping of Acoustic Sources*, Proceedings of AIAA Aeroacoustics Conference 2006, Paper 2006-2711.
- [4] P. Sijtsma. *CLEAN Based on Spatial Source Coherence*, Proceedings of AIAA Aeroacoustics Conference 2007, Paper 2007-3436.
- [5] J. Hald. *Mapping of Contributions from Car-exterior Aerodynamic Sources to an In-cabin Reference Signal Using Clean-SC*, Proceedings of Inter-Noise 2016, Paper 512.
- [6] J. Hald. *Cross-spectral Matrix Diagonal Reconstruction*, Proceedings of Inter-Noise 2016, Paper 518.
- [7] R.P. Dougherty. *Cross Spectral Matrix Diagonal Optimization*, Proceedings of Berlin Beamforming Conference 2016, Paper BeBeC-2016-S2.
- [8] L. Vandenberghe and S. Boyd. *Semidefinite Programming*, SIAM Review, Vol. 38, No. 1, pp. 49-95, March 1996.
- [9] M. Grant and S. Boyd. *CVX: Matlab Software for Disciplined Convex Programming, Version 2.1*, <http://cvxr.com/cvx>, 2014.
- [10] S.B. Bendat and A.G. Piersol. *Engineering Applications of Correlation and Spectral Analysis* John Wiley & Sons, Inc., New York, 1980, pp. 264-291.

Active Vibration-based SHM System for Wind Turbine Blades: Demonstration on an Operating Vestas V27 Wind Turbine

Dmitri Tcherniak^{*}; *Lasse L. Mølgaard*[†]

Abstract

This study presents a structural health monitoring (SHM) system that is able to detect structural defects on wind turbine blades, such as cracks, leading/trailing edge openings, or delaminations. It is shown that defects as small as 15 cm in size can be detected remotely, without stopping the wind turbine. The SHM system presented is vibration-based: mechanical energy is artificially introduced by means of an electromechanical actuator, whose plunger periodically hits the blade. The induced vibrations propagate along the blade and are picked up by accelerometers mounted along the blade. The vibrations in mid-range frequencies are utilized: this range is above the frequencies excited by blade-wind interaction, ensuring a good signal-to-noise ratio. At the same time, the corresponding wavelength is short enough to deliver required damage detection resolution and long enough to be able to propagate the entire blade length.

This paper demonstrates the system on a Vestas V27 wind turbine. One blade of the wind turbine was equipped with the system, and a 3.5 month monitoring campaign was conducted while the turbine was operating normally. During the campaign, a defect – a trailing edge opening – was artificially introduced into the blade and its size was gradually increased from the original 15 cm to 45 cm.

Using a semi-supervised learning algorithm, the system was able to detect even the smallest amount of damage, while the wind turbine was operating under different weather conditions. This paper provides detailed information about the instrumentation and the measurement campaign, and explains the damage detection algorithm.

^{*} Brüel & Kjær Sound & Vibration Measurement A/S, Skodsborgvej 307, Nærum 2850, Denmark, Email: Dmitri.Tcherniak@bksv.com

[†] Department of Applied Mathematics and Computer Science, Technical University of Denmark, Denmark

Résumé

Cette étude concerne un système SHM (système de suivi de la santé des structures) capable de repérer les altérations sur la structure des pales d'éoliennes, telles que fissures, ouvertures sur les bords d'attaque et de fuite, ou délaminages. Il y est montré que ces altérations de la structure, même de taille très faible (de l'ordre de 15 cm) peuvent être détectées à distance sans avoir à stopper l'éolienne. Ce système SHM est un système vibratoire : la vibration est transmise au moyen d'un actionneur électromécanique dont le poussoir frappe la pale de manière périodique. Les vibrations induites, qui se propagent le long de la pale, sont captées par des accéléromètres montés sur celle-ci. Ce sont les vibrations dans la gamme des fréquences moyennes qui sont ici utiles: cette gamme, supérieure à celle des fréquences générées par l'interaction pale-vent, garantit un bon rapport signal/bruit. La longueur d'onde correspondante est à la fois suffisamment courte pour que la détection bénéficie d'une bonne résolution, et suffisamment longue pour assurer la propagation sur toute la longueur de la structure. Le système a été installé sur une des pales d'une turbine éolienne Vestas V27 à l'occasion d'une campagne de suivi de trois mois et demi, la turbine étant alors en cycle de fonctionnement normal. Une dégradation - un interstice sur le bord de fuite - a artificiellement été apportée à la structure et progressivement agrandie, de sa taille initiale de 15 cm jusqu'à 45 cm. Grâce à un algorithme d'apprentissage semi-supervisé, le système a pu détecter les altérations même les plus faibles subies par la structure pendant que l'éolienne fonctionnait sous différentes conditions météorologiques. La présente communication fournit des informations détaillées sur l'instrumentation utilisée et sur le programme de mesures, ainsi qu'une description de l'algorithme de détection des défauts.

Zusammenfassung

Diese Studie präsentiert ein Structural Health Monitoring System (SHM), das Strukturdefekte an Windturbinenrotorblättern erkennen kann, wie Risse, Öffnungen an der Blattvorderkante/-hinterkante oder Delaminationen. Es wird gezeigt, dass selbst 15 cm große Defekte aus der Ferne erkennbar sind, ohne die Windturbine zu stoppen. Das hier vorgestellte SHM-System ist schwingungsbasiert: Mechanische Energie wird durch einen elektromechanischen Aktuator künstlich eingeleitet, dessen Stößel in regelmäßigen Abständen auf das Rotorblatt schlägt. Die induzierten Schwingungen breiten sich entlang des Blattes aus und werden von Beschleunigungsaufnehmern erfasst, die am Blatt entlang montiert sind. Genutzt werden die Schwingungen im mittleren Frequenzbereich:

Dieser Bereich liegt über den Frequenzen, die beim Zusammenwirken von Rotorblatt und Wind angeregt werden, was ein gutes Signal-Rausch-Verhältnis ergibt. Gleichzeitig sind die entsprechenden Wellenlängen ausreichend kurz, um die erforderliche Auflösung zur Erkennung des Defekts zu liefern, und ausreichend lang, um sich über die gesamte Blattlänge auszubreiten. Dieser Artikel demonstriert das System an einer Windturbine vom Typ Vestas V27. Ein Rotorblatt der Windturbine wurde mit dem System ausgestattet und 3,5 Monate lang überwacht, während die Turbine normal in Betrieb war. Im Rahmen der Kampagne wurde das Rotorblatt künstlich mit einem Defekt - einer Öffnung an der Blatthinterkante – versehen und dessen Größe schrittweise von ursprünglichen 15 cm bis auf 45 cm erweitert. Mithilfe eines semi-überwachten Lernalgorithmus konnte das System selbst den kleinsten Defekt erkennen, während die Windturbine unter verschiedenen Wetterbedingungen betrieben wurde. Dieser Artikel bietet detaillierte Informationen zur Geräteausstattung und Messkampagne und erläutert den Algorithmus für die Defekterkennung.

1. Introduction

Blades of modern wind turbines are designed for 20 to 25 years of service under severe weather conditions, and during this period damage is unavoidable. With a high probability, a small blade defect may develop into a bigger failure, and if no countermeasures are taken, may become critical, causing catastrophic consequences. Repair of a small defect is significantly cheaper than repair of a bigger one, or replacement of an entire blade. Therefore, wind turbine operator companies pay close attention to structural health monitoring of the blades. Today this is done by periodical visual inspections conducted every one-to-two years, but many in the industry realize that a better approach is needed. Many approaches have been suggested, attacking the problem from different angles (see Reference [1]). Besides using more robust blade design and special surface treatments to protect the blades, the new approaches include, for example, facilitating visual monitoring by means of transportable ground-based optical systems, by drones equipped with high-resolution video cameras, using thermography and many others techniques.

One of the most promising ways is fitting wind turbines with vibration sensors and monitoring the blades' integrity via permanent monitoring of their vibration, [7], [11]. This approach is already adopted for monitoring the mechanical components of wind turbines, such as gearboxes and bearings. The main advantage

of such a system is that the operator/owner is notified about the occurrence of damage almost immediately after it has happened and not after one or two years later, when it is detected by visual inspection.

Structural health monitoring via vibration monitoring may be based on different physical phenomena. One of the popular vibration-based approaches is detecting changes in modal parameters: loss of structural integrity leads to reduction of stiffness, which can be detected by monitoring modal parameters. However, this approach cannot achieve the required damage resolution since the modal parameters are not very sensitive to damage [10]. Another modal-based technique is FE model updating [19]. It can provide more detailed information about detected damage, with respect to its location and identification [12], [8], [9], but the method is prone to numerical instability due to the ill-conditioned system of equations required to be solved when updating the parameters [5].

Another well-known vibration approach is based on guided-waves [13]: a piezoelectric exciter generates stress waves, which propagate through the structure and get picked by another piezoelectric sensor. Typically, a network of active sensors (which can measure and generate vibrations) is used. Blade damage can be detected and localized by monitoring how the vibration propagates from the actuators to the sensors. The guided-waves approach has much better damage resolution, but requires high sensor density, since the high-frequency oscillations quickly decay with the propagation distance. Using a large number of sensors adds complexity to the SHM system and negatively influences its cost, making it less attractive for the end users.

In a previous study [18] the authors introduced another technique (patent pending), which is similar to the guided-waves technique, but has inherent differences: the excitation is introduced by an electromechanical actuator, and the utilized frequency range is much lower compared to the guided-waves approach. The introduced vibrations are picked up by an array of accelerometers. The waves at the lower frequency (around 1 kHz) can propagate longer distances, thus the technique requires far fewer sensors. At the same time, the frequency is high enough to ensure sufficient damage detection resolution (at least 15 cm size). Structural damage changes the properties of the energy propagation between the actuator and the accelerometers; this can be detected by comparing the vibration pattern in a reference (healthy) state with the damaged state.

The important feature of the suggested approach is that it is possible not only to detect damage, but also to follow its development, [18]. Additionally, in studies

[6], [20] and [21] the possibility to use the technique for damage localization was demonstrated.

In [18] the method was applied to an SSP34m blade (34 m long), mounted on a test rig. This study reports the results when the same technique was used on an operating wind turbine.

2. System Implementation on Vestas V27 Wind Turbine


In [18], the authors described the experiment conducted on an SSP34m blade mounted in a test rig. Test rig facilities greatly simplified the experiment: since the blade did not move and was located indoors, it required much less effort to mount accelerometers, actuators and cabling. The experiment proved that the proposed approach performs well on a modern blade, with feasible actuator location and using a reasonable number of sensors. The system managed to detect a realistic blade fault (trailing edge opening) and follow up on its progression. However, using the test rig, we could not evaluate the robustness of the method against noise. Indeed, when operating, the wind turbine blade is subjected to wind excitation and excitation from the hub and nacelle mechanisms, which mask the signal from the actuator. In [18] some artificial noise (recorded on the blades of another wind turbine in operation) was mixed with the measured signals from the actuator: we had to admit that the selected signal-to-noise ratio was very much a guess. In addition, to be able to demonstrate technical feasibility of the proposed system, it was important to do it on a real operating wind turbine.

A Vestas V27 wind turbine was selected for the experiment due to its availability. The wind turbine stands in the grounds of Technical University of Denmark (DTU), Department of Wind Energy (formerly known as Risø), in Denmark, near the town of Roskilde. The Vestas V27 is a relatively old wind turbine, with 27 m rotor diameter and 225 kW rated power. However, this wind turbine can be considered representative of many modern wind turbines, it is an upwind, pitch regulated, horizontal axis wind turbine. In contrast to modern wind turbines, its blades are relatively stiff, and it has only two speed regimes: 32 and 43 rpm.

For blade excitation, the same actuator was used as for the SSP34m blade experiment (see Fig. 1).

The actuator is a simple electromechanical device: a coil is mounted on a steel base; driven by an electrical pulse, the coil ‘shoots’ the plunger towards the

Fig. 1. Actuator design



structure; after the hit, the plunger retracts to the initial position by means of a spring.

Due to the size of the blade, it was not possible to install the actuator inside the blade (as was done on the SSP34m blade). Instead, the actuator was installed outside the blade, on its upwind side about one metre from the root (see Fig. 2), covered by a waterproof lid and secured with a strap (see Fig. 3). The vibrations were measured by accelerometers. The blade was fitted with twelve monoaxial

Fig. 2. Actuator location, circled on figure



Fig. 3. Installation on the blade



piezoelectric accelerometers (Brüel & Kjør Type 4507-B); their location on the blade is shown in Fig. 4.

The nominal sensitivity of accelerometers #5 to #15 was 10 mV/ms^{-2} (Type 4507-B-004) and accelerometer #16, located near the actuator, had nominal sensitivity 1 mV/ms^{-2} (Type 4507-B-001).

For mounting the accelerometers, we used plastic mounting clips, which were glued directly to the blade (no special alignment was performed, the accelerometer's measurement direction was perpendicular to the blade surface). To protect the accelerometers, they were covered by silicon, then 'helicopter tape' (polyurethane tape) was applied on top to give the silicon a smooth shape (see Fig. 5). The accelerometers and cables were placed on the downwind side of the

Fig. 4. Installation of accelerometers/actuator on the blade. Red circles on the blade contour indicate the location of the accelerometers, the green circle indicates the actuator position




Fig. 5. Accelerometer and cable fitted to the blade (helicopter tape visible)



blade. The accelerometer cables ran from the accelerometers towards the trailing edge and then along the trailing edge towards the blade root (Fig. 6).

The cables were glued to the blade with silicon and covered with helicopter tape.

From experience we knew that long-term measurement campaigns on the same wind turbine could last several months [17], and that this set-up was sufficient for the planned campaign. (While this set-up is sufficient for the planned campaign, it

Fig. 6. Accelerometer cables running along the trailing edge of the blade



is obviously not good enough to survive on the blade for several years, so in that case, other arrangements must be developed.)

The accelerometers were connected to a data acquisition system (Brüel & Kjær Type 3660-C, with two LAN-XI modules, a 12-channel input module Type 3053-B-120 and 4-channel input/output module Type 3160-A-042), see Fig. 7.

Two piezoresistive DC accelerometers Type 4574-D mounted in the spinner were used to estimate the rotor azimuth, with a possibility to derive the rotational speed of the rotor. In addition, the pitch angle was also measured.

The actuator was controlled by the signal from the signal generator built into one of the data acquisition modules. The generated rectangular pulse triggered the actuator's electronics, making a 1000 μF capacitor discharge through the coil. Then the capacitor was charged again to 48 V using a DC/DC converter to be ready for the next shot.

The data acquisition system and the electronics were placed in a waterproof box (dimensions $60 \times 45 \times 20 \text{ cm}^3$ and weight 25 kg), which was mounted to the inner surface of the spinner (Fig. 8). The equipment was powered by 24 V from the nacelle via a slip ring.

The measured data (in total 16 signals sampled with 16.384 kHz frequency) was wirelessly transmitted from the rotating part to the nacelle via two Cisco wireless access points, one located inside the waterproof box and another installed in the nacelle.

Fig. 7. LAN-XI modules installed in waterproof box



Fig. 8. Waterproof box mounted inside the spinner – cables not connected



When the turbine is operating, the line of sight between the hub and nacelle might be blocked by the steel parts of the hub. To keep an uninterrupted wireless connection, two pairs of antennas were employed: two omnidirectional antennas attached to the hub and two directional antennas mounted inside the nacelle.

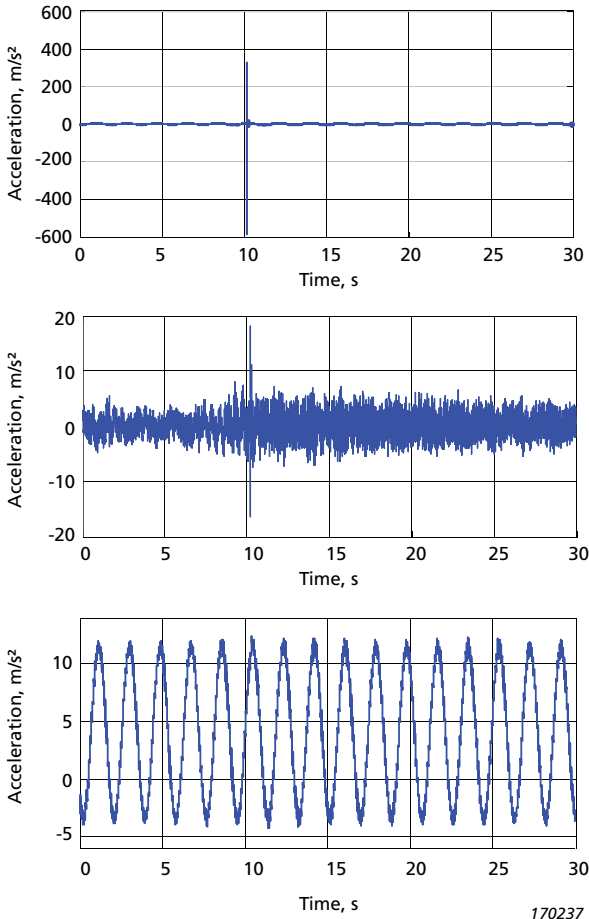
The data acquisition system was controlled by Brüel & Kjær PULSE™ LabShop software. The software was programmed to start data acquisition, record 10 s, initiate an actuator hit and record for another 20 s. Then acquisition was stopped and the system waited for four and a half minutes and initiated again. Thus, 12 actuator hits and corresponding datasets were produced every hour. Typical signals are shown in Fig. 9.

Simultaneously with the vibration data, meteorological data was collected from a weather mast located a few hundred metres away. The weather data included temperature, wind speed and direction, wind turbulence at different altitudes, atmospheric pressure, precipitation, etc. The data from the mast was delivered averaged within one-minute intervals and the power production data and yaw angle (the wind direction seen from the nacelle) was also available from the wind turbine system.

3. Experiment

The measurement campaign lasted 104 days. With 12 actuator hits per hour, data from 24,693 actuator hits were collected. During this time, the wind turbine was

Fig. 9. Typical signals – **Top:** Accelerometer #16 (20 cm from the actuator);
Centre: Accelerometer #15 (8.5 m from the actuator); **Bottom:** DC accelerometer



170237

subjected to different weather conditions. The monitoring period covered about a third of the year, thus no season-related events were observed. During the campaign, the turbine was in its normal power production regime, governed by its controller. However, following the agreement with the wind turbine owner, in the damaged state, we could only operate the turbine under visual surveillance, that is,

during working hours. For nights, weekends and holidays, the wind turbine was set to idling (no power production), though the SHM system was kept running.

4. Damage Implementation

For validating the capabilities of the proposed SHM system, an artificial defect was introduced on the instrumented blade. The following considerations were taken into account:

- Input from wind turbine manufacturers and service companies regarding blades' typical defects and their location
- Reparability of the defect: it should be possible to repair the blade inexpensively after the end of the experiment
- Risk of the artificial damage developing to critical should be minimal.

In addition to this, we planned to test another property of the proposed SHM system, the indication of damage progression. For this reason, we planned to gradually increase the size of the defect.

Taking the above mentioned into account, the 'trailing edge opening' type of damage was selected. This is a typical defect for blades manufactured using this technology.

Commercial sources, for example [22], inspection reports and technical papers such as [1] and [4], indicate that trailing edge failures are frequently observed in blades. Besides this, such a failure is easy to introduce, extend and repair, and according to experience, the probability that it can progress uncontrollably is very low.

The initial artificial damage was introduced at the start of the measurement campaign by technicians from service company Total Wind Group, (see Fig. 10).

The trailing edge was opened and extended to simulate a crack. The length of the opening was 15 cm. The opening was covered by helicopter tape to prevent atmospheric water from coming into contact with unprotected inner blade material.

After a period, the length of the opening was extended to 30 cm, and then extended even further to 45 cm (Fig. 11). After the measurement campaign, the defect was repaired.

The exact position of the defect is shown on the overall view of the blade (Fig. 4), and a close-up is shown in Fig. 12.

Fig. 10. Implementation of artificial blade damage - initial 15 cm trailing edge opening





Fig. 11. Implementation of artificial blade damage - extended to 30 cm (left); extended to 45 cm (right)



Fig. 12. Damage location on the blade relative to the accelerometers and its development: 15 cm > 30 cm > 45 cm correspond to black > dark grey > grey



5. Data Processing and Damage Detection

5.1. Classification

The damage detection approach discussed in this paper is a case of anomaly or novelty detection, and the vast knowledge accumulated using this paradigm can be utilized for solving SHM problems, see, for example, [16] and [3]. Using the terminology in the anomaly detection field, semi-supervised anomaly detection describes our approach. A supervised approach to damage detection would require recordings of a normal (healthy) operating state, as well as recordings of the blades in operation with the damage that should be detected by the system. A supervised classifier could then be trained to distinguish between the different observed operating states.

In practice it would only be feasible to obtain operational data for a limited number of damage types, constraining the usefulness and versatility of the SHM system. Employing a semi-supervised view of the problem we assume that only the normal (healthy) state is known, and significant deviations from this state are associated with damage.

The damage detection procedure therefore includes two phases: the training phase and the detection phase. During the training phase, we assume that the structure is undamaged, here, we collect a number of samples, characterizing the normal state under different operating regimes and establish a statistical model (or models, one for every wind turbine regime) of the normal state. In the detection phase, every newly acquired sample is compared to the model of the normal state. If a significant deviation is detected, we declare that the blade is damaged.

To prepare the recordings from the accelerometers for statistical modelling, two steps must be performed: pre-processing and calculation of a feature vector. The steps are considered in the following sections.

5.2. *Data Processing*

In the presented SHM system prototype, the data was not processed in real time (as one would expect from a commercial SHM system), but was post-processed later, when the data from several hundred actuator hits became available.

The data analysis started with processing the two DC accelerometers' signals (example shown in Fig. 9, bottom graph). Note that the mean of the signal is not at zero due to the centrifugal acceleration, which is due to the rotation of the rotor. By detecting and counting the peaks, it is possible to obtain the rpm and azimuth profiles, and derive the rpm and azimuth values at the moment of the actuator hit.

For each actuator hit, the time history from the accelerometers, pitch angle, derived rotor rpm and rotor azimuth information were combined with the weather data and saved into a database to facilitate data access.

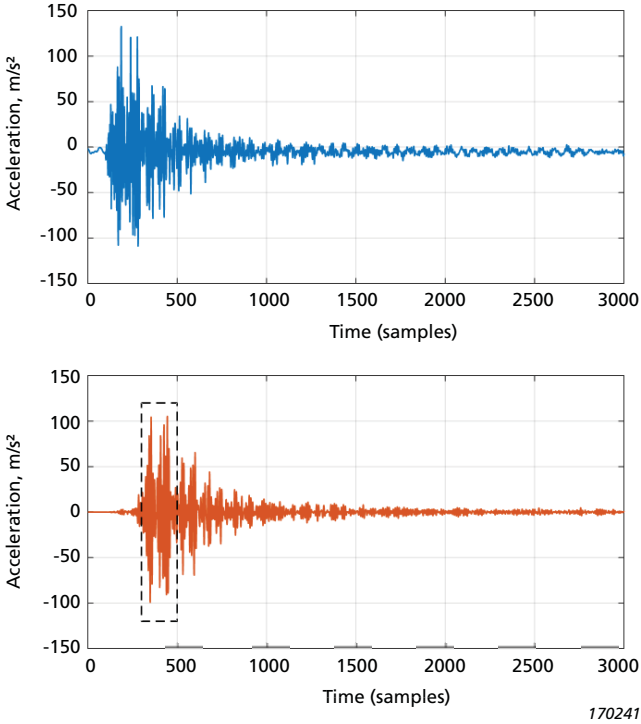
The three main operating regimes were identified: idling, operating at 32 rpm and operating at 43 rpm. It was recognized that the vibrational data from the regimes cannot be compared directly, and further analysis was conducted separately for each regime.

Further preprocessing steps included:

- 1) Selecting the part of the signal around the actuator hit. This was done by detecting the beginning of the peak in the actuator control signal. In this step, the originally recorded 30 s of vibration data (about 500,000 time samples) are reduced to 3,000 time samples per measured channel.
- 2) Fine alignment of the signals from different hits. This was done using the signal from accelerometer #16, closest to the actuator, since it is least affected by the noise.
- 3) Bandpass filtering. As mentioned in the introduction, a medium frequency range (around 1 kHz) is used. The bandpass filter (BPF) was designed around this frequency (700–1200 Hz).
- 4) Finally, the short part of the filtered signal was extracted for the following processing: only 201 time samples between time sample 300 and 500 (inside the dashed rectangle in Fig. 13, bottom graph) was retained and used for damage feature calculation.

Fig. 13 illustrates the steps in the process. It must be noted that the BPF filter parameters were found by trials, using data from the undamaged and damaged states and trying to maximize the performance of the damage detection.

Fig. 13. **Top:** Typical accelerometer signal (32 rpm, accelerometer #5); **Bottom:** same signal after bandpass filtering. The dashed rectangle indicates the final trim



170241

This approach is not feasible in the real life scenario, where the data from the damaged blade is not available.

However, based on the authors' experience gained from processing Vestas V27 and SSP34m blade data, the performance of the algorithm is only slightly affected by the fine tuning of the pre-processing parameters.

Using general recommendations, one can design a quite sensitive and robust SHM system. Such recommendations are: design a BPF with the centre frequency around 1 kHz and width 500–600 Hz and retain 200–400 time samples, where all the signals attain the highest magnitude. Apparently the original design can be further improved by either testing the system mounted on a blade on a test rig, or by creating 'virtual test environments' based on a detailed FE model of the blade, which can simulate the acceleration responses to the actuator hits. Such 'virtual

environments' can be a great tool for optimizing sensors and actuator location, to tune the performance of the system towards the most realistic failures.

5.3. Feature Vector

The collection of preprocessed signals for each actuator hit must be represented using a smaller number of quantities, termed features. This representation must retain sufficient information to separate 'normal' and 'anomalous' operation for the subsequent statistical modelling.

Following [14], the feature vector is based on the cross-covariance between all pairs of sensor time series. Such a covariance matrix characterizes the current state of the blade.

An acquired structural defect will change the energy propagation from the actuator to the sensors, which will affect the vibration pattern (relative magnitude and phase) of the measured acceleration signals. Since the cross-covariance function is a measure of similarity between two signals, the changes in the vibration pattern will be reflected as a change in the cross-covariance matrix.

For a fixed time-lag (in this study, we used zero time lag), the covariance matrix is an $N \times N$ symmetric matrix, where N is the number of sensors selected for the analysis (either the full set, or a subset of the sensors in Fig. 4). The number of distinct elements in the matrix is $N(N+1)/2$. If all 12 sensors are selected, the dimensionality of the feature vector is 78. However, the values in this feature vector will exhibit multicollinearity, which will impact the subsequent fitting of a statistical model. To counter the multicollinearity, dimensionality reduction using principal component analysis (PCA) is employed. The PCA technique is a statistical method that determines a lower-dimensional orthogonal representation from a set of vectors such that the extracted dimensions retain maximal variance of the data. The PCA representation of the data is determined by using the covariance vectors obtained during the training phase. In this work the number of vectors available for training is around 200 (see Section 6).

The feature vector used to represent the state of the blade for the i -th actuator hit is obtained by:

- Calculating the cross-covariance matrix and reshaping it into an $N(N+1)/2$ length feature vector, which consists of the distinct elements of the cross-covariance matrix
- Project the feature vector into the K -dimensional PCA-representation to produce the final compressed feature vector \mathbf{x}_i of length K , where K is smaller than $N(N+1)/2$

Choosing a small number of PCA dimensions to use in the subsequent modelling of the normal state can be crippling for the ability to detect damage. In this study K was chosen such that 99% of the variance in the training data is retained.

5.4. Normal State and Damage Index

From each of M actuator hits obtained in the healthy state, a feature vector is extracted. The collection of M feature vectors (which is often called training set) are gathered to a matrix $\mathbf{X}=[\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M]$, $\mathbf{X} \in R^{K, M}$.

This matrix forms the basis for the statistical model of the healthy blade, or in other words the normal, or reference state. The feature vectors \mathbf{x}_i are often called samples, as they reflect the state of the system at the time of the i -th actuator hit. When the actuator strikes the structure and a new sample \mathbf{y} arrives, the difference between the sample and the statistical model of the healthy state characterizes the health of the structure. The Mahalanobis distance is a convenient metric to quantify this difference. The Mahalanobis distance between a sample \mathbf{y} and the dataset \mathbf{X} is given by:

$$d(\mathbf{y}, \mathbf{X}) = \sqrt{(\mathbf{y} - \mu_{\mathbf{x}})^T (\Sigma_{\mathbf{x}})^{-1} (\mathbf{y} - \mu_{\mathbf{x}})} \quad (1)$$

where $\mu_{\mathbf{x}} \in R^K$ is the mean of the samples in \mathbf{X} , and $\Sigma_{\mathbf{x}} \in R^{K, K}$ is the covariance between the samples. This metric is selected as a *damage index* in this study. If the damage index is relatively small and does not exceed some threshold D , we declare that the system is in undamaged state. Conversely, if $d(\mathbf{y}, \mathbf{X}) > D$, we declare state \mathbf{y} as damaged.

Following the semi-supervised strategy, we base the choice of the threshold D exclusively on the observations from the training phase. A naïve choice of threshold would include all points in the training set:

$$D = \max_i(d(\mathbf{x}_i, \mathbf{X})) \quad (2)$$

but this choice is sensitive to outliers that could be present in the training set. Instead we shall estimate the probability of a sample being an outlier based on the training data. The value of the threshold D , can then be obtained as the $(100-R)^{\text{th}}$ percentile of the cumulative distribution function, which means that R percent of the samples from the healthy state may exceed this value. In other words, it means that we allow R percent of false alarms in the training set, which results in a more conservative choice of the threshold. Further, in the study we use this approach to control the threshold, and R is called *allowed false alarm rate*.

In practice, all values $d_i=d(\mathbf{x}_i, \mathbf{X})$ are calculated for the selected training set \mathbf{X} and sorted such that $d_j < d_{j+1}, j = 1..M - 1$. Then the threshold D_R is selected from the sorted sequence as:

$$D_R = d_k \quad (3)$$

where k is the nearest integer less than or equal to $M(100 - R)/100$. The important feature of the damage index (1) is that (in most of the cases) its value increases with the damage development. This allows one to identify if the damage appeared, but then stabilized, or if it keeps progressing.

6. Results

As mentioned before, Vestas V27 wind turbine has a controller which operates the wind turbine in one of the three regimes:

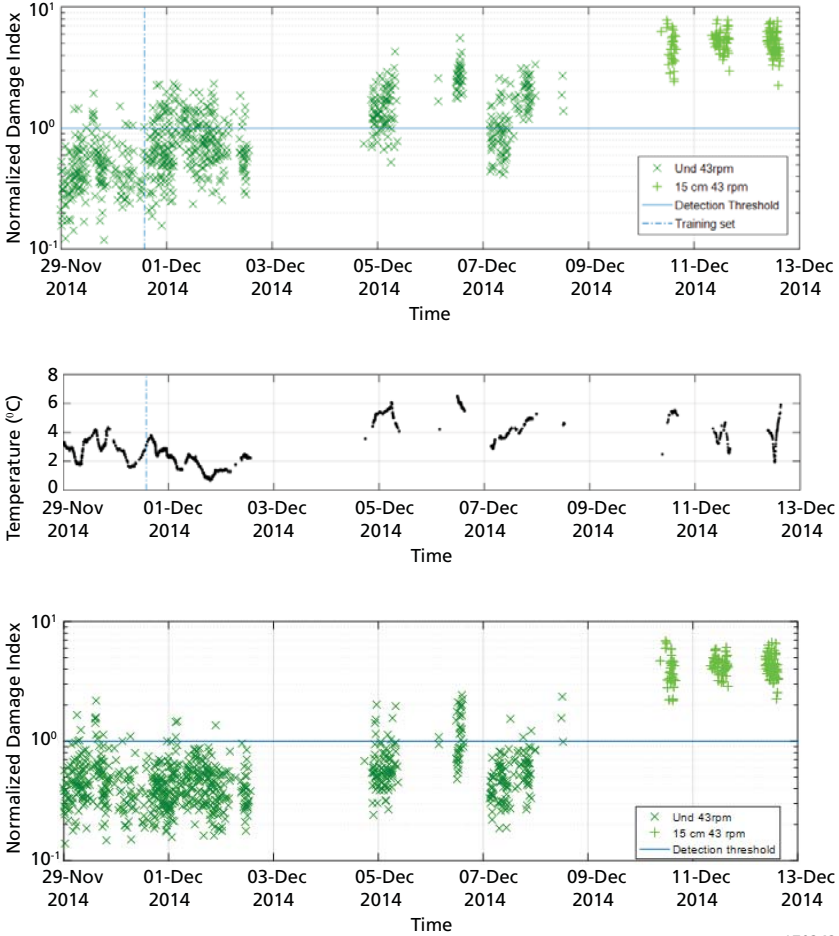
- 1) Idle: the blades' pitch is about 90° , so the wind turbine is not producing power. The rotor revolves due to wind shear. This regime is active if the wind is too weak (its speed is below the cut-on speed). This regime is activated if the wind turbine is broken, or 'switched off'.
- 2) Low speed production regime: the controller keeps the rotor speed at a constant 32 rpm. The pitch angle for all three blades is about 0° .
- 3) High speed production regime: the controller keeps the rotor speed at 43 rpm, the pitch angle is about 0° , though slight variations may be present.

Apparently, for SHM purposes, the latter two regimes are of the most interest. Most of the time wind turbines are operating, and it is unreasonable to stop them for SHM purposes. The present study addresses only the regimes when the wind turbine is operating.

6.1. Sensitivity to Weather Conditions

As it is well known from other literature, environmental conditions influence the dynamic response of the structure and this may seriously affect the performance of an SHM algorithm. Fig. 14 demonstrates this, using the data from leading- and trailing-edge accelerometers, from the 43 rpm regime. Following the scheme described in [23], the first third of the available 856 samples (for the healthy state at 43 rpm regime) was used to train the algorithm. These samples are located to the left of the blue vertical dotted line in Fig. 14 (top). Accepting 5% of false alarms in the training set ($R=5\%$), the model of the healthy state is generated, and the threshold is found (solid horizontal line, the points below the line indicate the

Fig. 14. **Top:** Damage index (43 rpm regime), training samples are picked from the beginning of the training set, left of the blue dotted line; accepted false alarm rate 5%; test false alarm rate 48.6%. **Centre:** Temperature when the samples were taken. **Bottom:** Damage index, training samples are randomly picked; accepted false alarm rate 5%; test false alarm rate 8.1%



170242

healthy state). The samples below the threshold line are classified as healthy state, while the samples above the threshold are declared as damaged state.


All states corresponding to the smallest amount of damage (the 15 cm crack) are identified correctly (denoted by '+'). However, validating the model against the healthy samples not included in the training set, one finds that the algorithm produces 48.6% of false alarms, which is much more than the allowed 5%.

One can readily correlate the false alarms with the peaks and troughs of the temperature, Fig. 14 (centre). Indeed, the model of the healthy state was generated when the temperature was between 2°C and 4°C, and this model fails when the temperature is outside this range.

Instead of taking the first samples, the model can be based on an equal number of samples randomly picked from the healthy state. This significantly improves the performance of the model. Fig. 14 (bottom) illustrates this for some realization of the random sequence. Now the test false alarm rate is 8.1%, which is much closer to the allowed 5% rate.

Fig. 15 demonstrates the application of this approach to all available data from the undamaged and damaged states, for both 32 and 43 rpm operating regimes collected during the entire monitoring period. The considered datasets contain 1684 samples for the healthy case, 224 samples for the 15 cm crack case, 310 samples for the 30 cm crack, and 237 samples for the 45 cm crack. The difference in the sample count for the different damaged states is due to the different duration of the periods, when the wind turbine was operating in these damaged states. This was due to the availability of the technicians who climbed

Fig. 15. Damage index for all undamaged and damage cases for 32 and 43 rpm cases. Allowed false alarm rate is 5%, test false alarm rate is 7.3%



170243

the blade to extend the blade opening. It was also dependent on the weather, when this operation was possible. The different number of samples in the damaged states does not affect the classification results, as every sample from the damaged states is considered independently. In the same way as described above, the model of the healthy state was generated based on a third of the available samples from the healthy state, which were randomly selected. The remaining two-thirds of the samples were used to validate the correctness of the healthy state classification and compute the false alarm rate.

As mentioned previously, two models were used, one for each regime, and to be comparable, the damage indices d are normalized by the corresponding threshold values $D_{5\%}$. For the given random sequence, the overall false alarm rate is 7.3%, while the correct detection rate for all three crack sizes is 100%. Fig. 15 also illustrates that the damage index value generally increases with the size of the crack: for the undamaged blade the value is below 1 (for the false alarm cases it reaches 2); for the 15 cm crack the value is between 1 and 8; for the 30 cm crack it is between 2 and 15; and for the 45 cm crack it is between 4 and 80. It is not possible to make a correlation between the damage index value and the crack size (as it naturally depends on the location of damage) but one can conclude whether the crack has stabilized, or continues to develop.

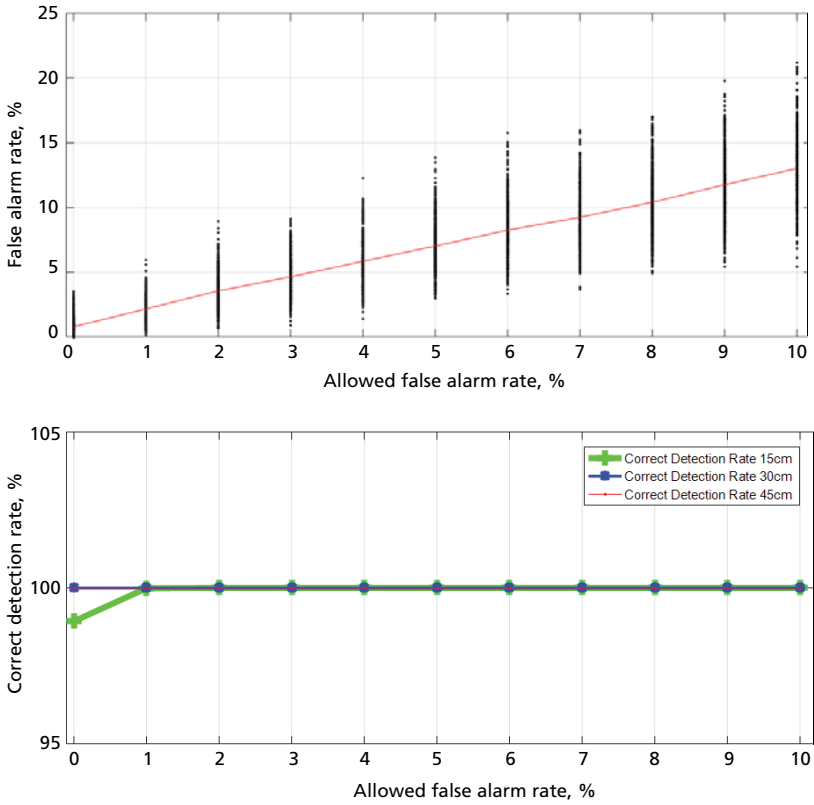
As mentioned before, the choice of the allowed false alarm rate R defines the detection threshold, thus affecting the resulting false alarm rate (False Positives) and the correct detection rate (True Positives), which in classification studies is typically illustrated by a Receiver Operating Characteristic (ROC) curve. Here we use a slightly different approach plotting the false alarm rate and the correct detection rate as functions of the allowed false alarm rate R (Fig. 16).

Since the samples of the training set are randomly selected from the healthy samples, the false alarm rate and correct detection rate depend on realization of the random sequence. Fig. 16 shows that the allowed false alarm rate R was iterated from 0 to 10%; for each case, 500 random realizations of the training set were generated and the corresponding models were applied to the data.

Fig. 16 (top) shows the development of the false alarm rate, where the dots represent the value for each random sequence realization and the line is the mean value. As expected, the mean false alarm rate is generally higher than the allowed false alarm rate, however there exist random sequence realizations where the false alarm rate is significantly higher or significantly lower than the mean value.

Fig. 16 (bottom) shows the averaged correct detection rate. The figure is based on the data from the 43 rpm case.

Fig. 16. **Top:** Test false alarm rate for different allowed false alarm values for 500 random realizations of the training set and the averaged value. **Bottom:** Averaged correct detection rate for 500 realizations of the training set



170244

6.2. *Influence of Different Sensor Configurations*

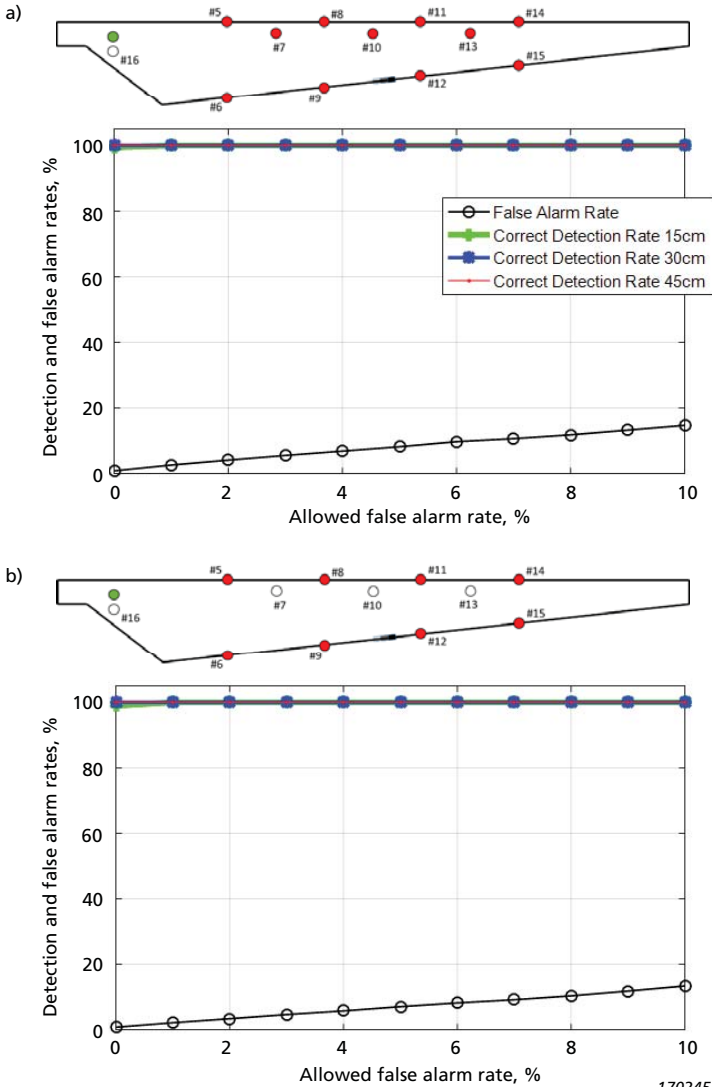
The measured datasets include the data from 12 accelerometers, as shown in Fig. 4. Accelerometer #16 is located close to the actuator, and its signal is only used to align time histories from the different actuator hits. Sensors #5 thru #15 are distributed along the blade. In the implementation of a real SHM system, using the least number of sensors is desirable from system cost considerations.

Optimization of the number of sensors and their location is out of the scope of this paper, however, in this section we provide some results of ‘what-if’ scenarios, selecting different sensor configurations and providing the corresponding detection results and the false alarm rate.

Fig. 17 thru Fig. 20 compare the results for different sensor configurations computed for the 32 rpm regime. In Fig. 17 (a) the results when using all 11 blade accelerometers are shown. Excluding the spar sensors (Fig. 17 (b)) improves the results, indicating that the spar sensors’ signals do not contain any information about the trailing edge opening. Indeed, using only spar sensor for detection (Fig. 17 (c)) shows extremely bad detection results. It is interesting to note that using the four leading edge sensors (Fig. 17 (d)), the bigger amount of damage (namely, 30 and 45 cm cracks) is detectable but the algorithm fails to detect the 15 cm crack. Using the four trailing edge sensors (Fig. 17 (e)), the 15 cm crack can be detected with a higher certainty; here we observe an interesting non-typical phenomenon: the larger area of damage (30 cm crack) produces a smaller damage index, thus the longer crack can be detected with lower certainty than the bigger one.

The energy of the mechanical impact provided by the actuator propagates along the blade towards its tip, and numerous reflections from the elements of the structure and blade’s tip define a complex unique vibration pattern sampled by the few accelerometers. Since most of the energy propagates from the actuator towards the tip, it could be expected that a structural fault will mainly affect the readings of the accelerometers located behind the damage, and in much lesser degree the readings of the accelerometers between the actuator and the fault. Fig. 17 (f) thru (h) supports this suggestion: using four sensors (two on the leading edge and two on the trailing edge) the detection rate is excellent when the sensors are behind the fault (Fig. 17 (h)), only the 30 and 45 cm cracks are detectable when the sensors surround the fault (Fig. 17 (g)), and the worst results are obtained when the sensors are located between the fault and the actuator. Using this observation, it is possible to roughly locate the fault by ‘scanning’ the blade by employing different sensor combinations.

Fig. 17. Test false alarm rate and correct detection rate as a function of allowed false alarm rate for different sensor configurations. Each graph has a corresponding sensor configuration shown above it with the number of accelerometers used shown as red dots



170245

Fig. 18. Test false alarm rate and correct detection rate as a function of allowed false alarm rate for different sensor configurations (cont.)

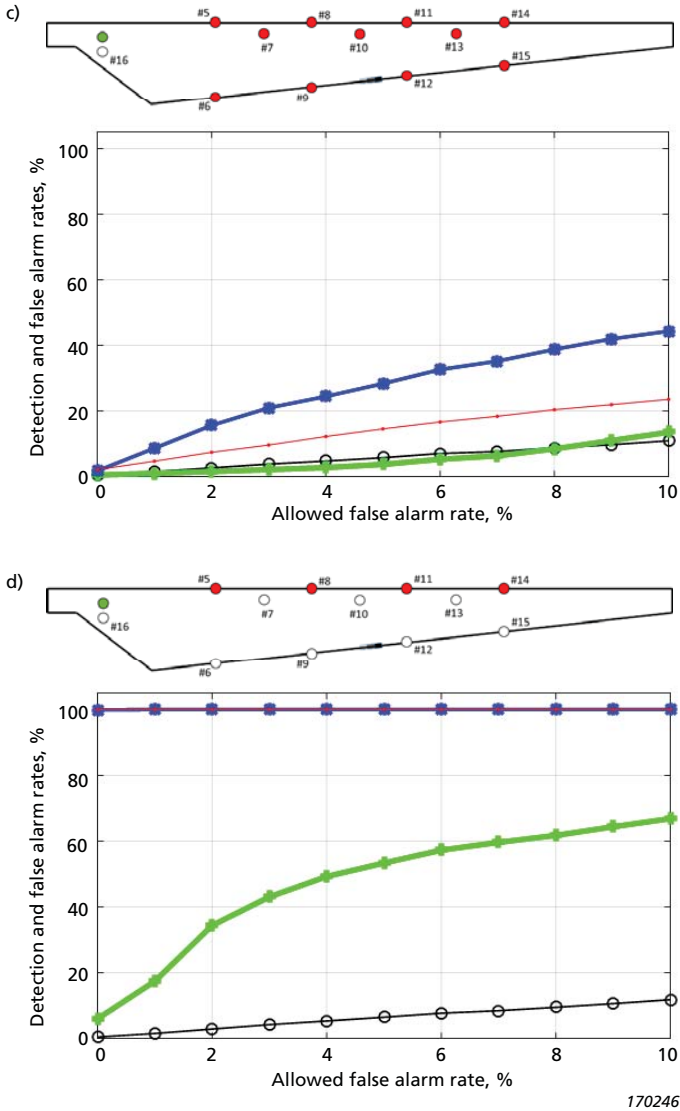


Fig. 19. Test false alarm rate and correct detection rate as a function of allowed false alarm rate for different sensor configurations (cont.)

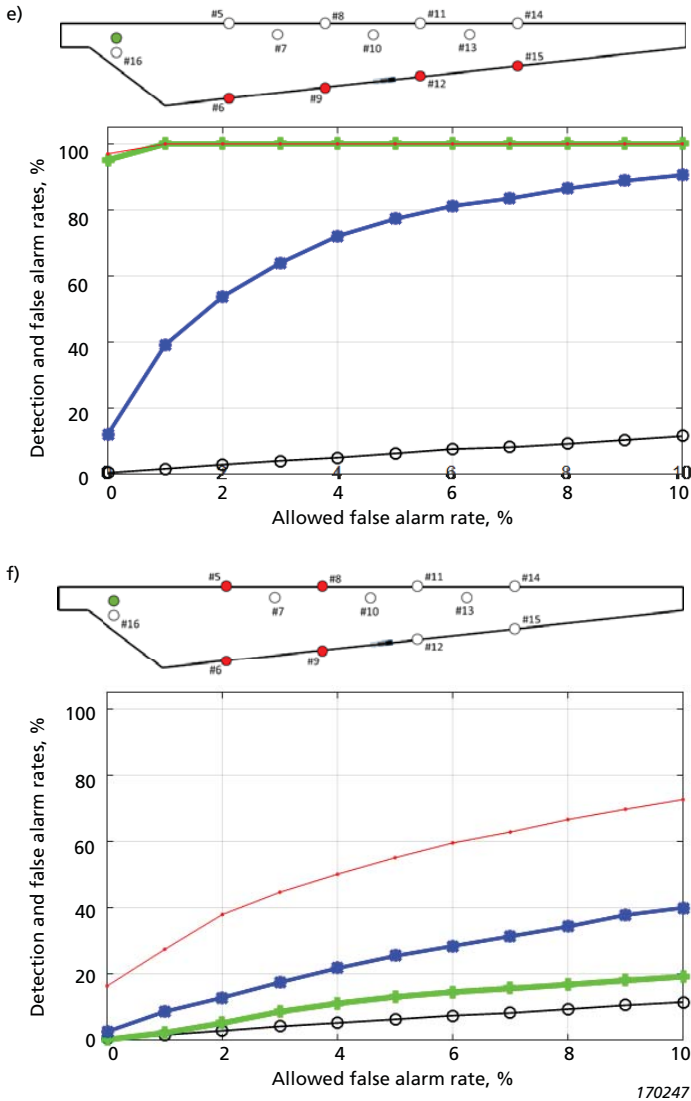

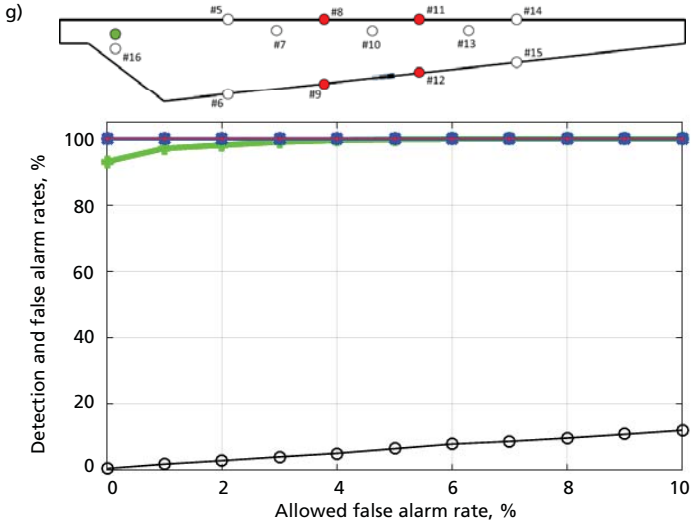


Fig. 20. Test false alarm rate and correct detection rate as a function of allowed false alarm rate for different sensor configurations (cont.)



170248

7. Discussion

At 32 and 43 rpm regimes, the damage detection algorithm demonstrates comparable performance. However, in the idling case the performance drops. It was found that this case required taking the rotor azimuth angle into account; apparently the samples measured at different blade positions are not directly comparable. A possible reason for this is that while idling, the blade pitch is about 90° and the plunger hit direction lies in the rotor plane. Therefore, the strength of the hit is affected by gravity and depends on the azimuth angle (Fig. 3). When operating, the blade's pitch is around zero, the actuator direction is always perpendicular to the vector of gravity (Fig. 2), and the hit strength is not affected by the rotor position.

8. Technical Implementation

As mentioned earlier, the described system is a prototype, which was implemented in order to prove the concept of an actuator-based, vibration-based SHM system. This section discusses some technical aspects of a real-life implementation of such a system.

Perhaps the most challenging part of such a system, if implemented in reality, is the accelerometers. Due to possible lightning strikes, wind turbine manufacturers avoid placing any metal elements (except the components of the lightning protection system) into the blade further than one-third of its length. Apparently, the use of conventional piezoelectric accelerometers connected by copper wires is not an option for monitoring the entire blade, and they have to be replaced, for example, by metal-free optical accelerometers connected by optical fibre cables. This solution requires optical data acquisition systems. As the specifications of the currently available fibre optic accelerometers are considerably inferior to their piezoelectric counterparts, their performance can significantly affect the performance of the overall system, and this must be thoroughly investigated. Still, conventional piezoelectric accelerometers can be an option if only the root section of blades needs to be monitored.

Regarding the amount of data that the system has to handle, one can note it is quite moderate despite the high sampling frequency (in the described implementation, 16 kHz, but this can be reduced to 2 to 3 kHz without any effect on accuracy). Indeed, the system is not measuring constantly but activated periodically, for example once per hour (the five-minute interval used in the presented study is due to the need to collect a lot of data in a short period of time). The amount of data will depend on the practical implementation of the system,

namely the split of the responsibilities between the embedded system and the server. For example, if the bandpass filtering, signal alignment, calculation of the covariance matrix and the dimensionality reduction are performed by the embedded system, then each sample is expressed by a vector of 20 to 30 floating-point numbers, which is up to 120 bytes. This amount of data has to be transferred to the server for classification and decision making; if the embedded system is capable of producing these steps (which involves Mahalanobis distance calculation), then the embedded system needs to transfer only a few bytes to the server when damage is detected. During the training phase it is desirable to collect the raw data (time histories), which will provide more flexibility, for example, for simple damage localization, as described in previous sections, or if a few sensors are lost during system use.

Another crucial parameter characterizing the SHM system is its overall cost, which includes the cost of the hardware, cost of installation (either retrofitting or installation during blade manufacturing) and maintenance and operational costs. The apparently high cost of such a system makes it economically feasible only for multi-megawatt wind turbines, where the blade length reaches 80 to 100 metres. Monitoring of the full length of such a blade requires 20 to 30 accelerometers per blade, though the data acquisition channels can be shared between the blades using multiplexing. For long blades, a multi-actuator design (as described in [18]), appears promising, thus the blade can be monitored by sections. This allows a significant reduction of the data acquisition channels.

Concluding, we can note that the projected cost of such a system may be considerably high, however, with a proper design it can be notably reduced.

9. Conclusion and Future Research

This study presents an active vibration-based SHM system that utilizes an electromechanical actuator (automatic hammer) and an array of accelerometers. As a damage feature, a covariance matrix between the measured acceleration signals was used. The paper describes a three-and-a-half-month measurement campaign, when the system was installed on an operating Vestas V27 wind turbine. The ability of the system to detect an artificially introduced failure (blade's trailing edge opening) was investigated. It was demonstrated that a 15 cm long opening can be detected without stopping the wind turbine. It can be concluded that the actuator-based approach, in combination with covariance-based

damage feature can be used for successful detection of typical blade defects, while using a feasible hardware setup and semi-supervised learning algorithm.

The authors see the following directions of future development. The first direction concerns creating virtual test environments, which will allow simulation of the vibration response of the blade to an actuator hit. Such environments should be based on a 3D model of the blade detailed enough to realistically reproduce the time history of the acceleration responses with good resolution. The environments are a convenient tool for designing and optimizing the SHM system, as it facilitates selecting the number and location of the actuator and accelerometers, and testing the design against different types and locations of blade faults. Such a tool can also be useful to determine the minimum characteristics of the data acquisition chain. For example, using such a tool, one can evaluate the suitability of metal-free optical accelerometers in the blade SHM context.

The second direction is the improvement of the classification algorithm. First of all, the algorithm should be able to take into account the weather conditions. There are already a number of promising approaches minimizing the effect of weather conditions on the damage detection results. Another approach to improve the certainty of the classification is to utilize the sequence of observations. Indeed, the SHM system produces a binary result (the blade is either healthy or damaged) after every actuator hit. However, there is no need to make a conclusion solely based on one observation, while many sequential observations are available. Utilizing this sequence may considerably improve the certainty of the decision making. For example, the Sequential Probability Test Ratio (SPRT) approach looks like a promising tool to benefit from the observation sequence, see [15].

10. Acknowledgements

The work was partly supported by EUDP (Danish Energy Technology Development and Demonstration Programme), grant number 64011-0084 Predictive Structure Health monitoring of Wind Turbines. The authors would like to thank DTU Wind Energy for granting access to the test object.

11. References

- [1] Ataya, S., Ahmed, M.M.Z. *Damages of Wind Turbine Blade Trailing Edge: Forms, Location, and Root Causes*. Engineering Failure Analysis 35, 480-488, 2013.
- [2] Ciang, C. C., J. R. Lee, and H. J. Bang. *Structural Health Monitoring for a Wind Turbine System: a Review of Damage Detection Methods*. Measurement Science and Technology 19 (21), 2008.
- [3] Fassois, S.D., and J.S. Sakellariou. *Time-series Methods for Fault Detection and Identification in Vibrating Structures*. Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences, 365 411-448, 2006.
- [4] Haselbach, P.U., Branner, K. *Initiation of Trailing Edge Failure in Full-scale Wind Turbine Blade Test*. Engineering Fracture Mechanics 162 136-154, 2016.
- [5] Hua, X.G., Y.Q. Ni, and J.M. Ko. *Adaptive Regularization Parameter Optimization in Output-error-based Finite Element Model Updating* Mechanical Systems and Signal Processing 23(3) 563-579, 2009.
- [6] Garcia, D., D. Tcherniak, and I. Trendafilova. *Damage Assessment for Wind Turbine Blades Based on a Multivariate Statistical Approach*. International Conference on Damage Assessment of Structures. Ghent, Belgium, 2015.
- [7] Griffith, D. T., N. C. Yoder, B. Resor, J. White, and J. Paquette. *Structural Health and Prognostics Management for the Enhancement of Offshore Wind Turbine Operations and Maintenance Strategies*. Wind Energy 17 (11) 1737-1751, 2013.
- [8] Lam H.F. and Yang J.H. *Bayesian Structural Damage Detection of Steel Towers Using Measured Modal Parameters*. Earthquakes and Structures 8(4) 935-956, 2015.
- [9] Lam H.F., Yang J.H. and Au S.K. *Bayesian Model Updating of a Coupled-slab System Using Field Test Data Utilizing an Enhanced Markov Chain Monte Carlo simulation algorithm*. Engineering Structures 102 144-55, 2015.

- [10] Larsen, G. Chr., P. Berring, D. Tcherniak, P. H. Nielsen, and K. Branner. *Effect of Damage to Modal Parameters of a Wind Turbine Blade*. European Workshop on Structural Health Monitoring. Nantes, France, 2014.
- [11] McGugan, M., G. Pereira, B. F. Sørensen, H. Toftegaard, and K. Branner. *Damage Tolerance and Structural Monitoring for Wind Turbine Blades*. Royal Society of London. Philosophical Transactions A. Mathematical, Physical and Engineering Sciences 373, 2015.
- [12] Ni Y.Q., Zhou H.F., Chan K.C., and Ko J.M. *Modal Flexibility Analysis of Cable Stayed Ting Kau Bridge for Damage Identification*. Computer Aided Civil and Infrastructure Engineering 23(3) 223-236, 2008
- [13] Park, G., G. T. Taylor, K. M. Farinholt, and C. R. Farrar. *SHM of Wind Turbine Blades Using Piezoelectric Active Sensors*. European Workshop on Structural Health Monitoring. Sorrento, Italy, 2009.
- [14] Parker, D. L. *Multi-objective Design Optimization Framework for Structural Health Monitoring*. PhD Thesis, Mississippi State University, 2011.
- [15] Sohn, H., D.W. Allen, K. Worden, and C.R. Farrar. *Statistical Damage Classification Using Sequential Probability Ratio Tests*. Structural Health Monitoring Vol.2(1) 0057-74, 2003.
- [16] Tarassenko, L., D. Clifton, P. Bannister, S. King, and D. King. *Novelty Detection*. In *Encyclopaedia of Structural Health Monitoring*, 653-675. John Wiley and Sons, Inc., 2009.
- [17] Tcherniak, D., and G. C. Larsen. *Application of OMA to an Operating Wind Turbine: Now Including Vibration Data From the Blade*. International Operational Modal Analysis Conference (IOMAC). Guimarães, Portugal, 2013.
- [18] Tcherniak, D., and L. L. Mølgaard. *Vibration-based SHM System: Application to Wind Turbine Blades*. Damage Assessment of Structures. Ghent, Belgium, 2015.
- [19] Teughels, A., and G. De Roeck. *Structural Damage Identification of the Highway Bridge Z24 by FE Model Updating*. Journal of Sound and Vibration 278 589-610, 2004.

- [20] Ulriksen, M. D., D. Tcherniak, and L. Damkilde. *Damage Detection in an Operating Vestas V27 Wind Blade by use of Outlier Analysis*. IEEE Workshop on Environmental, Energy and Structural Monitoring Systems (EESMS). Trento, Italy, 2015.
- [21] Ulriksen, M. D., D. Tcherniak, L. M. Hansen, R. J. Johansen, L. Damkilde, and L. Frøyd. *In-situ Damage Localization for a Wind Turbine Blade Through Outlier Analysis of SDDLV-induced Stress Resultants*. Structural Health Monitoring 1-17, 2016.
- [22] Wind Energy Update. *The Wind Energy Operations & Maintenance Report 2016*. Industry Report, Wind Energy Update (<http://analysis.windenergyupdate.com/>), 2016.
- [23] Worden, K. & Manson, G. *The Application of Machine Learning to Structural Health Monitoring*. Phil. Trans. R. Soc. A 365, 515-537. (doi:10.1098/rsta.2006.1938), 2007.

Previously issued numbers of Brüel & Kjær Technical Review

(Continued from cover page 2)

- 1 – 2005 Acoustical Solutions in the Design of a Measurement Microphone for Surface Mounting
Combined NAH and Beamforming Using the Same Array
Patch Near-field Acoustical Holography Using a New Statistically Optimal Method
- 1 – 2004 Beamforming
- 1 – 2002 A New Design Principle for Triaxial Piezoelectric Accelerometers
Use of FE Models in the Optimisation of Accelerometer Designs
System for Measurement of Microphone Distortion and Linearity from Medium to Very High Levels
- 1 – 2001 The Influence of Environmental Conditions on the Pressure Sensitivity of Measurement Microphones
Reduction of Heat Conduction Error in Microphone Pressure Reciprocity Calibration
Frequency Response for Measurement Microphones – a Question of Confidence
Measurement of Microphone Random-incidence and Pressure-field Responses and Determination of their Uncertainties
- 1 – 2000 Non-stationary STSF
- 1 – 1999 Characteristics of the vold-Kalman Order Tracking Filter
- 1 – 1998 Danish Primary Laboratory of Acoustics (DPLA) as Part of the National Metrology Organisation
Pressure Reciprocity Calibration – Instrumentation, Results and Uncertainty
MP.EXE, a Calculation Program for Pressure Reciprocity Calibration of Microphones

Special technical literature

Brüel & Kjær publishes a variety of technical literature that can be obtained from your local Brüel & Kjær representative.

The following literature is presently available:

- Catalogues
- Product Data Sheets

Furthermore, back copies of the Technical Review can be supplied as listed above. Older issues may be obtained provided they are still in stock.

Front cover: A half-wheel beamforming array being used in the wind tunnel at FCA US Headquarters, Chrysler World Headquarters and Technology Center in Auburn Hills, Michigan

Each Brüel & Kjær Technical Review contains a collection of technical, scientific articles that describe theory, measurement techniques and instrumentation, which are specifically aimed at acousticians and vibration engineers.

See more on bksv.com/Knowledge-center, Technical Reviews

2017-12

ISSN 0007-2621

BV 0003-11

Brüel & Kjær Sound & Vibration Measurement A/S

DK-2850 Nærum · Denmark · Tel: +45 77 41 20 00 · Fax: +45 45 80 14 05

www.bksv.com · info@bksv.com

Local representatives and service organisations worldwide

