

Cross-spectral matrix denoising for beamforming in wind tunnels

Hald, Jørgen¹

Ginn, Kevin Bernard²

Brüel & Kjær Sound & Vibration Measurement A/S
Skodsborgvej 307, 2850 Nærum, Denmark

ABSTRACT

Microphone-array based beamforming measurements in wind tunnels exhibit reduced dynamic range due to flow-induced noise in the individual microphones. Even in open wind tunnels with the arrays outside the flow region, air turbulence will induce such flow noise. Assuming stationary signals and performing long-time averaging of a cross-spectral matrix (CSM), the noise contamination will be concentrated on the CSM diagonal. When the CSM is used for traditional frequency-domain beamforming, Diagonal Removal (DR) will avoid use of the contaminated diagonal. DR is effective at suppressing the noise, but it also often underestimates source strengths and removes weak sources. Other array processing methods must use the diagonal. Several algorithms that attempt to reconstruct the CSM diagonal have been published. It has been shown that remaining off-diagonal noise contributions will limit the performance of methods that operate only on the diagonal. A few algorithms exist that can overcome this limitation by performing modifications also on the off-diagonal elements of the matrix. The paper describes a few of these methods and describes their respective limitations and advantages. Results from simulated and real measurements are presented.

Keywords: Noise source identification, Wind tunnel measurements, Beamforming, Flow noise

I-INCE Classification of Subject Number: 72

1. INTRODUCTION

Flow-induced noise in microphones is a well-known problem when measuring outdoors or in a wind tunnel. The problem can be reduced using windscreens, but not avoided. Typically, noise source identification on a vehicle in a wind tunnel is performed with beamforming, using a microphone array. In an open tunnel, the array can be placed outside of the core flow, but still in a region with significant air turbulence. Based on recorded time signals from the microphone array, a Cross-Spectral Matrix (CSM) is first averaged, and in a second step the CSM is used typically for frequency-domain beamforming. If the measured flow-noise signals are independent stationary stochastic processes, then with increasing averaging time, flow-noise contributions will eventually

¹ Jorgen.Hald@bksv.com

² KevinBernard.Ginn@bksv.com

be present only 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 created around one windscreen may be picked up by nearby microphones, resulting in contributions outside the CSM diagonal.

Assuming the flow-noise contributions to be concentrated on the CSM diagonal, their effect on frequency-domain beamforming results can be removed by use of so-called Diagonal Removal (DR), where use of the diagonal elements is avoided [1]. Effectively, the diagonal elements are set to zero, which will often lead to the matrix having negative eigenvalues, which will in turn produce negative contributions to the focused source power map. The result will be underestimated source levels, cancellation of weak sources and perhaps areas with negative power in the contour maps produced.

Beamforming deconvolution techniques such as Non-Negative Least Squares (NNLS) [2], which are based on use of the so-called Point Spread Function (PSF), can be implemented with DR. This is because the PSF can be calculated using DR. CLEAN-SC [3] also includes a DR procedure, which however is much more complicated, requiring for each new identified point source an iterative solution of a system of non-linear equations. In Ref. [4] the iterative solution procedure was found to work nicely for the strong sources but failed to converge for the weaker sources. To overcome that limitation, the Diagonal Denoising (DD) algorithm [5] was adopted. With other beamforming algorithms, such as Functional Beamforming [6], the diagonal of the CSM is needed, so DR cannot be applied.

The DD method and related methods [7-8] subtract a maximum of signal power from the CSM diagonal while maintaining all off-diagonal elements unchanged and while retaining the matrix positive semidefinite (i.e. with no negative eigenvalues). It is shown in Ref. [5] that residual off-diagonal flow-noise contributions will limit the effectiveness of the method. If a very long averaging time has been used, ensuring that the off-diagonal flow-noise contributions have been reduced to be small, then the DD method will be effective. A long averaging time, however, requires the acquisition and storage of long time-data records for typically 100 to 200 array channels. This is expensive in terms of wind-tunnel time and data storage. Methods have therefore been developed, which can overcome the performance limitation set by the off-diagonal flow-noise contributions. These methods still retain the CSM as positive semidefinite, but they support small changes being applied to the off-diagonal elements. Examples of such methods include: Robust Principal Component Analysis (RPCA) [9-10], Probabilistic Factorial Analysis (PFA) [11] and Canonical Coherence Denoising (CCD) [12].

The present paper gives in Sec. 2 a brief description of the different methods with a focus on DD and CCD. Results from simulated measurements are given in Sec. 3, while Sec. 4 presents results from a real wind-tunnel measurement. Sec. 5 contains a Discussion, providing a graphical overview of the principles and applications of the methods. Finally, Sec. 6 contains a summary.

2. THEORY

2.1 Problem statement

We model the sound pressure at M array microphones as coming from K mutually incoherent target sources and with independent noise added at each individual microphone. For snapshot number $l = 1, 2, \dots, L$, we consider the complex signals at a selected FFT line. Representing by \mathbf{s}_l the K -element complex source-signal vector, by \mathbf{n}_l the M -element complex noise-signal vector and by \mathbf{p}_l the microphone pressures, we can then write:

$$\mathbf{p}_l = \mathbf{H}\mathbf{s}_l + \mathbf{n}_l, \quad l = 1, 2, \dots, L, \quad (1)$$

where $\mathbf{H} \in \mathbb{C}^{M \times K}$ contains the transfer functions from sources to microphones. The averaging of the cross-spectral matrix can now be expressed as follows:

$$\mathbf{G} \equiv \frac{1}{L} \sum_{l=1}^L \mathbf{p}_l \mathbf{p}_l^H \equiv \overline{\mathbf{p}_l \mathbf{p}_l^H}, \quad (2)$$

where H represents Hermitian (conjugate) transpose, and where a bar on top represents snapshot averaging.

With a denoising algorithm we would ideally like to extract from \mathbf{G} the noise free matrix \mathbf{G}_0 defined as:

$$\mathbf{G}_0 = \overline{\mathbf{p}_{0,l} \mathbf{p}_{0,l}^H}, \quad (3)$$

where $\mathbf{p}_{0,l}$ is a vector containing the noise-free microphone signals:

$$\mathbf{p}_{0,l} = \mathbf{H}\mathbf{s}_l. \quad (4)$$

Using the above definitions, we can write:

$$\mathbf{G} \equiv \mathbf{G}_0 + \overline{\mathbf{n}_l \mathbf{n}_l^H} + \overline{\mathbf{p}_{0,l} \mathbf{n}_l^H} + \overline{\mathbf{n}_l \mathbf{p}_{0,l}^H}. \quad (5)$$

When the number of averages tends to infinity, the noise contribution $\overline{\mathbf{n}_l \mathbf{n}_l^H}$ will converge to a diagonal matrix $\mathbf{N}_\infty = \text{diag}(\sigma_n^2)$, while the cross terms $\overline{\mathbf{p}_{0,l} \mathbf{n}_l^H}$ and $\overline{\mathbf{n}_l \mathbf{p}_{0,l}^H}$ will vanish, because the source signals and the noise signals are all mutually incoherent. Here, σ_n is the noise variance and $\text{diag}(\sigma_n^2)$ is a diagonal matrix with diagonal elements σ_n^2 . However, as described for example in reference [5], this convergence is quite slow, so with a realistic number of averages, the off-diagonal elements of \mathbf{G} will typically still contain significant noise contributions. The amplitudes of the off-diagonal noise-noise cross spectra $N_{m,\tilde{m}}$ in the matrix \mathbf{N} will converge approximately as:

$$|N_{m,\tilde{m}}| \approx \sqrt{\frac{N_{m,m} N_{\tilde{m},\tilde{m}}}{L}} \approx \frac{\sigma_n^2}{\sqrt{L}} \quad \text{for } m \neq \tilde{m}, \quad (6)$$

and similar expressions will hold for the signal-noise cross terms in $\overline{\mathbf{p}_{0,l} \mathbf{n}_l^H}$ and $\overline{\mathbf{n}_l \mathbf{p}_{0,l}^H}$. When the source signals in $\mathbf{p}_{0,l}$ are stronger than the noise signals in \mathbf{n}_l , then the off-diagonal cross terms will be larger than the off-diagonal noise terms $N_{m,\tilde{m}}$. Use of for example 900 averages will only reduce the off-diagonal noise contributions by a factor 1/30, corresponding to 14.8 dB of attenuation. In addition, with many microphones there will be a much larger number of off-diagonal elements than the number of diagonal elements. This seems to increase the impact of these remaining off-diagonal noise contributions on the DD denoising algorithm to be described in the following section.

2.2 Diagonal Denoising (DD)

The Diagonal Denoising algorithm assumes that enough averaging has been performed that all off-diagonal noise contributions have been effectively removed:

$$\mathbf{G} \cong \mathbf{G}_0 + \text{diag}(\sigma_n^2). \quad (7)$$

To remove the noise contributions on the diagonal, the algorithm subtracts a maximum of power from the diagonal elements of \mathbf{G} , while retaining the resulting matrix positive semidefinite:

$$\max_{\mathbf{d}} \sum_m d_m \quad \text{subject to} \quad \mathbf{G} - \text{diag}(\mathbf{d}) \geq 0. \quad (8)$$

Here, \mathbf{d} is a vector with elements d_m , $\text{diag}(\mathbf{d})$ is a diagonal matrix with elements d_m , and “ ≥ 0 ” for a matrix means that it has non-negative eigenvalues (is positive semidefinite). As pointed out in reference [5], the optimization problem in Equation (8) has the form of a so-called Semidefinite Program, which can be solved efficiently and with guaranteed convergence properties using Convex Optimization methods. Once the problem has been solved, the denoised matrix \mathbf{G}_{DD} is obtained as:

$$\mathbf{G}_{\text{DD}} = \mathbf{G} - \text{diag}(\mathbf{d}). \quad (9)$$

It follows from Equation (8) that the algorithm will continue subtracting power from the diagonal until at least one eigenvalue of \mathbf{G}_{DD} equals zero. Thus, if the noise-free matrix \mathbf{G}_0 has only positive eigenvalues, then the DD method will subtract too much. This happens when the number of incoherent sources is equal to or larger than the number of microphones. In reference [5] it was shown by simulated measurements that accurate subtraction of the noise contribution on the diagonal is achieved, if the number of incoherent target sources does not exceed a limit around $M - \sqrt{2.5M}$. Above that limit, the deviation between \mathbf{G}_0 and \mathbf{G}_{DD} increases quickly, probably because too much is subtracted.

So far, we have assumed that all noise contributions were removed from the off-diagonal elements by averaging. In reference [5] it is investigated based on simulated measurements, how remaining off-diagonal noise contributions will limit the amount of noise power that the DD algorithm can remove from the diagonal. The investigation assumed identical noise level in all microphones, and the noise signals were assumed to be much stronger than the source signals, implying that off-diagonal cross terms $\overline{\mathbf{p}_{0,l}\mathbf{n}_l^H}$ and $\overline{\mathbf{n}_l\mathbf{p}_{0,l}^H}$ could be neglected. The off-diagonal noise-related contributions will cause the smallest eigenvalue of the matrix $\mathbf{G} - \text{diag}(\mathbf{d})$ to reach zero at an earlier stage, when maximizing the noise subtraction from the diagonal in Equation (8), providing therefore less noise reduction. The simulations in Ref. [5] showed that the achievable reduction factor α for the noise contributions on the diagonal could be approximated as:

$$\alpha \approx \frac{(M-1)^{0.625}}{\sqrt{L}}. \quad (10)$$

Equation (10) shows, as mentioned in the previous section, that large arrays will have poorer noise reduction on the diagonal than small arrays, probably because of the relatively larger number of off-diagonal matrix elements with disturbing noise contributions.

Figure 1 shows the achievable reduction in decibels as a function of the number L of averages for three different microphone counts M . The graphs show that for large arrays, a huge number of averages is needed to provide a significant noise reduction on the CSM diagonal.

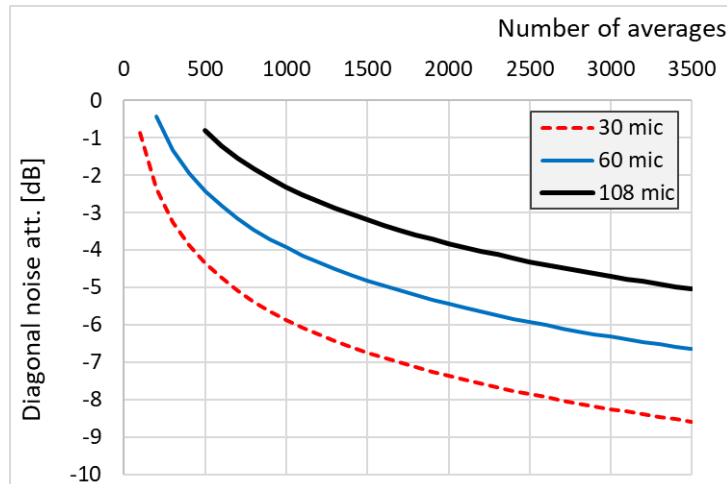


Figure 1. Achievable reduction in decibels of the noise on the diagonal versus the applied number of averages for arrays with 30, 60 and 108 microphones.

When the source signals have an amplitude higher than the noise signals, then the cross terms $\overline{\mathbf{p}_{0,l}\mathbf{n}_l^H}$ and $\overline{\mathbf{n}_l\mathbf{p}_{0,l}^H}$ will add an even larger off-diagonal noise contribution, which

will tend to further reduce the noise reduction that can be achieved on the diagonal. However, then the source signals will add positive contributions to the CSM eigenvalues, which will to some degree support a larger diagonal noise subtraction in Equation (8) before one of the eigenvalues reaches zero. The achievable reduction will depend on the actual source configuration.

2.3 Canonical Coherence Denoising (CCD)

The previous section showed that DD and related methods that subtract noise power from the CSM diagonal, while maintaining the matrix positive semidefinite, and while retaining all off-diagonal elements unchanged, will have very limited effect with realistic amounts of averaging. To overcome that limitation without generating denoised CSM's with negative eigenvalues, methods that can modify the off-diagonal elements based on some reasonable model/principle will be needed. This section and the following two sections describe such methods.

The present section describes the Canonical Coherence Denoising method of reference [12], which is based on dividing the array elements in two equally large groups X and Y, followed by extraction of the signal sub-space that is coherent between the two groups. The output from that basic method is the part $\tilde{\mathbf{G}}$ of the matrix \mathbf{G} , which represents the coherent sub-space. If enough averaging had been performed to concentrate the noise on the diagonal of \mathbf{G} , this operation would remove all the noise from the diagonal. However, the matrix $\tilde{\mathbf{G}}$ cannot represent more than $M/2$ incoherent sources. Thus, if the number of target sources is larger than $M/2$, then only a partial representation will be achieved. The solution described in reference [12] to overcome that limitation is to iteratively apply the coherent sub-space extraction to the residual matrix $\mathbf{G} - \tilde{\mathbf{G}}$ using each time a new X/Y grouping and add the result to $\tilde{\mathbf{G}}$. Three iterations were found to provide good results.

It is demonstrated in reference [12] that after three iterations, this basic iterative approach will generate an output $\tilde{\mathbf{G}}$ close to the output \mathbf{G}_{DD} from Diagonal Denoising. This is because a full modelling of all off-diagonal elements, including the residual noise contributions, is gradually enforced during the iteration. The following adaptations were introduced to optimize the performance and remove noise-related off-diagonal contributions:

- 1) Signal sub-spaces with a low coherence between the X and Y groups were removed in each iteration.
- 2) The number of iterations was reduced, if the noise level was estimated to be high. The noise level was estimated based on the relative average reduction of the diagonal elements in the first iteration (from \mathbf{G} to $\tilde{\mathbf{G}}$).
- 3) After each iteration, the number of significant sources was estimated from the eigenvalue-spectrum of $\tilde{\mathbf{G}}$. If the number of sources was found to be relatively small, then the number of iterations was reduced.

The final output will be written as \mathbf{G}_{CCD} . Details of the algorithm can be found in reference [12].

2.4 Robust Principal Component Analysis (RPCA)

The RPCA method models the measured CSM as the sum of a sparse noise-CSM and a low-rank signal-CSM. The model adaptation is performed by minimizing a weighted sum of the nuclear norm (sum of the eigenvalues) of the signal-CSM and the 1-norm of the noise-CSM, while keeping the sum of the two partial CSM's equal to the measured CSM. It is shown in reference [11] that the relative weighting of the two terms is very sensitive and difficult to choose, in particular when the number of incoherent sources is

not small. As in the CCD method, parts of the off-diagonal elements can be identified as noise and removed.

2.5 Probabilistic Factorial Analysis (PFA)

This method assumes the same number of unknown sources as the number of microphones. Then, the complex transfer functions from the sources to the microphones and the complex amplitudes of the source and noise signals across snapshots are considered as random variables with assigned probability density functions. The parameters of that statistical model are inferred by use of the measured data. Again, the noise model will cover parts of the CSM outside the diagonal, and these parts will be removed. It is shown in reference [11] that the method performs very well, providing in most cases at least 5 dB better noise reduction than DD. The main drawback of the method is a very long computational time.

3. SIMULATED MEASUREMENTS




Figure 2. Illustration of the set-up used for the simulated measurements. Blue circles: microphones. Red diamonds: sources (18 in this case).

A series of simulated measurements have been performed in Matlab using the setup illustrated in Fig. 2. The 108-element array has an optimized half-wheel geometry with a diameter of 4.1 metre. In combination with its image in a rigid ground plane at $Y=0$, it constitutes a full wheel. Because of the reflective ground, the transfer matrices \mathbf{H} from the simulated sources to the 108 microphones must include the ground reflection. The sound source consists of a number of equidistant, incoherent, monopole point sources along a line at 4 metre distance, covering a 4-metre span and with equal amplitudes of all sources. The real and imaginary parts of the complex source signals in the snapshot vectors \mathbf{s}_l , $l = 1, 2, \dots, L$, were generated with standard normal distributions, implying that the signals were circularly-symmetric complex Gaussian. The noise signals in the vector \mathbf{n}_l were generated in the same way. All noise signals were of equal level, SNR (Signal to Noise Ratio) decibels below the average level of the signals in $\mathbf{p}_{0,l}$. The simulated measurements to be presented used SNR values between -10 and 10 dB, with $\text{SNR} = 10$ dB representing the average sound pressure level in $\mathbf{p}_{0,l}$ being 10 dB stronger than the common noise level.

Only the error on the reconstructed matrix diagonal will be investigated. The relative average error level E for the “Raw” averaged matrix \mathbf{G} was calculated as:

$$E \equiv 10 \cdot \log_{10} \left(\frac{\|\text{diag}(\mathbf{G}-\mathbf{G}_0)\|_2}{\|\text{diag}(\mathbf{G}_0)\|_2} \right), \quad (11)$$

where the operator $\text{diag}(\cdot)$ extracts the diagonal elements of a matrix into a vector. The corresponding errors E_{DD} and E_{CCD} for the reconstructed matrices \mathbf{G}_{DD} and \mathbf{G}_{CCD} , respectively, were calculated in the same way. Only results from these two algorithms will be presented, and only a relatively high frequency of 10 kHz will be considered.

Figure 3a shows the relative average errors as functions of the SNR for the case of 18 incoherent sources and 1000 averages. On average, DD provides a noise reduction around 3.1 dB, while CCD provides 8 to 11 dB of reduction. The DD noise reduction predicted by Equation (10) is 2.3 dB.




Figure 3. Relative average diagonal error for the case of 1000 averages. (a) 18 sources, variable SNR. (b) Variable number of sources, 0 dB SNR.

Figure 3b displays the diagonal error level as a function of the number of sources for the case of 1000 averages and 0 dB SNR. CCD provides between 7 and 12 dB of noise reduction, while for DD it is between 3 and 4 dB. Notice the weak step in the CCD curve that occurs, where the number of sources reaches 30. With few sources, only 1 iteration is taken. At approximately 30 sources, the algorithm switches to 2 iterations, and beyond 67 sources, 2 or 3 iterations are used. As described above, the switching is controlled by a detected relative noise level and a detected number of significant sources. When the number of physical sources approaches the number of microphones, then it becomes difficult to both represent all the target sources accurately and to get good noise suppression.

4. WIND-TUNNEL MEASUREMENT

The simulated measurements of the previous section considered only the accuracy in the reconstruction of the noise-free cross-spectral matrix elements. To investigate, how CCD works in practice with beamforming, the method was applied to data acquired with a 108-element half-wheel array in FCA’s open semi-anechoic wind tunnel in Auburn Hills, Michigan. The array (with geometry as given in Sec. 3) was placed around 1 metre outside the core flow region, where there is still a significant amount of air turbulence. With 15 seconds of time data, 959 averages could be made. Only data from a single measurement with 140 km/h wind speed will be presented. Figure 4 shows aeroacoustic noise-source maps (sound intensity) for the frequency range 5-6 kHz obtained from three different beamforming calculations applied to the same measurement:

- 1) NNLS [2] with DR
- 2) CLEAN-SC [3] with DD
- 3) CLEAN-SC with CCD.

Both “Non-negative Least Squares” (NNLS) and “CLEAN based on Source Coherence” (CLEAN-SC) apply conventional frequency domain beamforming, but in a second “deconvolution” step they derive a point source model. The sound intensity maps are obtained from that point source model by assuming omni-directional radiation. Unfortunately, the actual aerodynamic source distribution cannot be accessed, so we can only compare the output from different methods.

The Diagonal Removal (DR) applied with NNLS is very effective at suppressing noisy spikes in the map due to flow noise, but as described in Sec. 1 it is known to produce underestimation and to remove weak sources, in particular at high frequencies. The NNLS map in Fig. 4 is nice and smooth with no noise spikes inside the 20 dB display range, but weak sources - such as the side mirror and the rear edge of the front light - almost disappeared. Although the source at the side mirror is relatively weak, it has a very important contribution to the in-cabin noise, and thus its level is important. NNLS and related methods are nevertheless quite widely used, making the investigation shown here even more important. The CLEAN-SC calculation with DD was performed using only 12 dB of dynamic range to limit the flow-noise spikes in high-frequency maps. The CLEAN-SC calculation with CCD was performed with an increased 15 dB of dynamic range, and even then, it produces significantly lower noise spikes than DD in the 5-6 kHz map. The results in Fig. 4 show that CCD removes a very large part of the noise spikes, while at the same time retaining all the important target sources.




Figure 4. Three calculations applied to the same wind-tunnel measurement with a wind speed of 140 km/h. The same colour scale and threshold “T” is used in all plots. (With kind permission from FCA).

Figure 5 shows the sound power spectra obtained by area-integration of sound intensity plots as those in Fig. 4 over the full map area. As expected, the NNLS calculation provides somewhat lower power values at high frequencies than the other two methods because DR was applied. CLEAN-SC with DD produces the highest high-frequency levels, probably mainly because of the many high-level noise spikes in the maps. On that background, and because of the good results in simulated measurements, CLEAN-SC with CCD is believed to provide the most accurate results. Some small fluctuations are

observed on the CCD based spectrum. Some of these are related to changes in the number of applied CCD iterations.

The CCD denoising performed during the calculation of the 163 frequency lines represented in Fig. 5 took 1.15 s of CPU time on a Dell Latitude E6540 laptop when executing a serial (not multi-threaded) Fortran implementation from a Visual Studio debugger. Matrix computations were performed using the Intel Math Kernel Library (MKL). In the commercial implementation used to produce the results in Figs. 4 and 5, multi-threading is applied, and the software is not executed from a debugger. As a result, the CPU time spent on the CCD denoising is insignificant, being around 0.1 s.

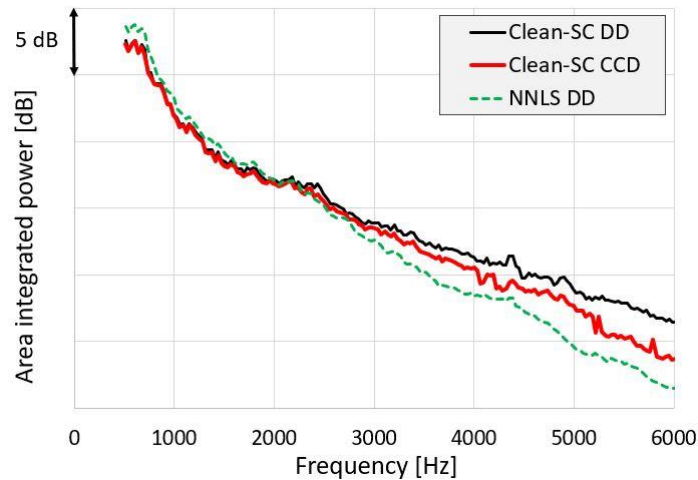


Figure 5. Area-integrated sound power spectra from the three calculations of Fig. 4.

5. DISCUSSION

Table 1 illustrates in the left column, how the three fundamentally different denoising methods (DR, DD and CCD) operate on an averaged CSM. The operations performed on and off the diagonal are shown.

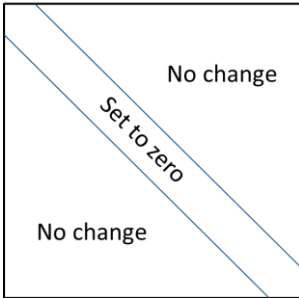
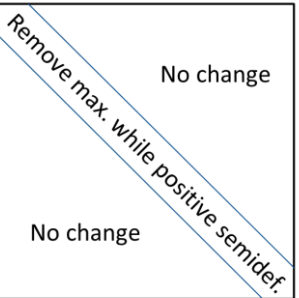
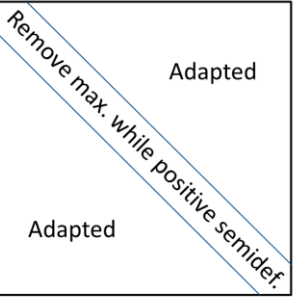
Diagonal Removal sets all diagonal elements to zero and re-scales the output in such a way that the correct level is obtained, when focusing on a single monopole point source. However, when focusing in other directions, negative power will often occur due to negative eigenvalues of the matrix. When there are several sources, such negative power contributions will partially cancel other sources, leading to underestimation. If some of these sources are weak, they may be completely cancelled.

Diagonal Denoising subtracts a maximum of signal power from the diagonal of the matrix, while keeping off-diagonal elements unchanged and retaining the matrix positive semidefinite. Thus, the problem with negative eigenvalues is avoided. The main limitation of the method is that remaining off-diagonal noise contributions will typically limit the noise power that can be removed from the diagonal before the smallest eigenvalue reaches zero. If a very long averaging time is applied, then the method will be effective, but with realistic averaging times, the benefit is very limited.

Canonical Coherence Denoising (and RPCA and PFA) overcomes the limitation in DD by supporting modifications also outside the matrix diagonal. This requires a model that can identify noise-related off-diagonal component. CCD identifies noise as having very low coherence between two equally large groups of the array microphones. Remaining components (with medium-to-high coherence) are retained. If the number of target sources is larger than half of the microphone count, then some of the target signal components will be discarded. To overcome that limitation, several iterations must be performed with different groupings of the microphones.

The right column in Table 1 gives examples of the array processing applications, with which the three denoising methods can be applied. As already mentioned, DR is associated with frequency domain beamforming (FDB), providing correct level when focusing on a single monopole point source. FDB is known to provide poor resolution at low frequencies and to produce high sidelobes (ghost sources). Deconvolution methods such as NNLS attempt to identify a distribution of incoherent monopole point sources that would produce the same FDB map as the actual measurement. This is done using the known response of the FDB beamformer to each one of the point sources – the so-called Point Spread Function (PSF). The PSF can be calculated using DR on a simulated measurement, and thus DR can be applied to the PSF-based deconvolution methods.

Table 1. Left column: Illustrations showing, how the denoising methods operate on a CSM. Right column: Typical array processing methods, where the three denoising methods can be applied. The methods used in this paper are shown with bold font.

Denoising method	Typical array processing applications
<p data-bbox="403 801 619 828">Diagonal Removal</p> 	<ul data-bbox="715 891 1246 1070" style="list-style-type: none"> • Frequency Domain Beamforming (Delay and sum) • Related deconvolution methods based on use of a Point Spread Function, for example NNLS.
<p data-bbox="403 1171 619 1198">Diagonal Denoising</p> 	<ul data-bbox="756 1238 1241 1462" style="list-style-type: none"> • Frequency Domain Beamforming • NNLS • CLEAN-SC • Functional Beamforming • MUSIC • Acoustic Holography
<p data-bbox="403 1541 619 1599">Canonical Coherence Denoising</p> 	<ul data-bbox="756 1612 1241 1836" style="list-style-type: none"> • Frequency Domain Beamforming • NNLS • CLEAN-SC • Functional Beamforming • MUSIC • Acoustic Holography

DD and CCD can basically be used as a pre-processing step before any array-based noise source identification. This includes FDB and NNLS (PSF-based deconvolution),

but also other beamforming methods such as CLEAN-SC, Functional Beamforming and Multiple Signal Classification (MUSIC) [13-14]. The MUSIC algorithm associates the largest eigenvalues of a CSM with the target signal and the remaining with noise. If the noise is too high, the method fails. CCD would be able to remove some noise before application of MUSIC. CCD is of course not restricted to beamforming applications. For example, the denoised CSM could be used for any type of Near-field Acoustic Holography, see for example reference [15].

6. SUMMARY

The paper has described a couple of methods for denoising of an averaged cross-spectral matrix (CSM) measured with an array of microphones. The target application of the denoised matrix is here noise-source identification on a vehicle in a wind tunnel using beamforming. Three fundamentally different approaches have been compared:

- 1) **Diagonal Removal (DR)** which will set all elements on the matrix diagonal to zero. The method can be used with standard frequency domain beamforming and with associated deconvolution methods based on use of a Point Spread Function. The method is effective at suppressing noise effects, but due to reduced or even negative eigenvalues in the modified CSM, severe underestimation and removal of weak sources will result.
- 2) **Diagonal Denoising (DD)** which subtracts a maximum of auto-power from the CSM diagonal, while keeping off-diagonal elements unchanged and while retaining the matrix positive semidefinite. The method is stable and computationally efficient but remaining off-diagonal noise contributions will seriously limit the amount of noise that can be removed, unless an extremely long averaging time has been used.
- 3) Methods that can identify noise contributions in the full matrix (on- and off-diagonal) and subtract these, while keeping the matrix positive semidefinite. Several such methods exist, but only one based on Canonical Coherence between groups of microphones has been reviewed and tested. That **Canonical Coherence Denoising (CCD)** method is fast, and it seems to work in general well with an automated selection of a coherence threshold and a number of iterations.

These methods have been compared based on simulated measurements on monopole point sources and based on an actual array measurement from a wind tunnel. The simulated measurements showed that CCD can provide much better noise suppression than DD with a typical number of averages equal to 1000. The results from the wind-tunnel measurement illustrated the underestimation by DR and the poor noise suppression provided by DD, pointing at CCD as the best option. CCD provided good noise suppression while avoiding the suppression of important secondary sources. CCD was also found to be computationally very efficient.

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. 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.

3. P. Sijtsma, "CLEAN based on spatial source coherence," *Proceedings of AIAA Aeroacoustics Conference 2007*, paper 2007-3436.
4. 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.
5. J. Hald, "Removal of incoherent noise from an averaged cross-spectral matrix," *J. Acoust. Soc. Am.* **142**(2), 846-854 (2017).
6. R.P. Dougherty, "Functional Beamforming," *Proceedings of Berlin Beamforming Conference 2014*, paper BeBeC-2014-01.
7. R.P. Dougherty, "Cross spectral matrix diagonal optimization," *Proceedings of Berlin Beamforming Conference 2016*, paper BeBeC-2016-S2.
8. Q. Leclère et.al., "Extraction of the acoustic part of a turbulent boundary layer from wall pressure and vibration measurements," *Proceedings of Novem 2015*.
9. J. Wright et.al., "Robust principal component analysis: Exact recovery of corrupted low-rank matrices via convex optimization," *Proceedings of Advances in Neural Information Processing Systems 22*, 2080-2088 (2009).
10. A. Finez, A.A. Pereira and Q. Leclère, "Broadband mode decomposition of ducted fan noise using cross-spectral matrix denoising," *Proceedings of Fan Noise 2015*.
11. A. Dinselmeyer, J. Antoni et. al., "On the denoising of cross-spectral matrices for (aero)acoustic applications," *Proceedings of Berlin Beamforming Conference 2018*, paper BeBeC-2018-S2.
12. J. Hald, "Denoising of cross-spectral matrices using canonical coherence," *J. Acoust. Soc. Am.*, to appear.
13. R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Transactions on Antennas and Propagation* **34**, 276-280 (1986).
14. D.H. Johnson and D.E. Dudgeon, *Array Signal Processing: Concepts and Techniques*, P T R Prentice Hall, Englewood Cliffs, New Jersey 07632, 1993.
15. J. Hald, "Fast wideband acoustical holography," *J. Acoust. Soc. Am.* **139**(4), 1508-1517 (2016).