

Spectral proper orthogonal decomposition

Moritz Sieber^{1,†}, C. Oliver Paschereit¹ and Kilian Oberleithner¹

¹Institut für Strömungsmechanik und Technische Akustik, Hermann-Föttinger-Institut, Technische Universität Berlin, Müller-Breslau Str. 8, D-10623 Berlin, Germany

(Received 13 August 2015; revised 10 November 2015; accepted 3 February 2016;
first published online 4 March 2016)

The identification of coherent structures from experimental or numerical data is an essential task when conducting research in fluid dynamics. This typically involves the construction of an empirical mode base that appropriately captures the dominant flow structures. The most prominent candidates are the energy-ranked proper orthogonal decomposition (POD) and the frequency-ranked Fourier decomposition and dynamic mode decomposition (DMD). However, these methods are not suitable when the relevant coherent structures occur at low energies or at multiple frequencies, which is often the case. To overcome the deficit of these ‘rigid’ approaches, we propose a new method termed spectral proper orthogonal decomposition (SPOD). It is based on classical POD and it can be applied to spatially and temporally resolved data. The new method involves an additional temporal constraint that enables a clear separation of phenomena that occur at multiple frequencies and energies. SPOD allows for a continuous shifting from the energetically optimal POD to the spectrally pure Fourier decomposition by changing a single parameter. In this article, SPOD is motivated from phenomenological considerations of the POD autocorrelation matrix and justified from dynamical systems theory. The new method is further applied to three sets of PIV measurements of flows from very different engineering problems. We consider the flow of a swirl-stabilized combustor, the wake of an airfoil with a Gurney flap and the flow field of the sweeping jet behind a fluidic oscillator. For these examples, the commonly used methods fail to assign the relevant coherent structures to single modes. The SPOD, however, achieves a proper separation of spatially and temporally coherent structures, which are either hidden in stochastic turbulent fluctuations or spread over a wide frequency range. The SPOD requires only one additional parameter, which can be estimated from the basic time scales of the flow. In spite of all these benefits, the algorithmic complexity and computational cost of the SPOD are only marginally greater than those of the snapshot POD.

Key words: computational methods, low-dimensional models, turbulent flows

1. Introduction and motivation

1.1. Contemporary methods for data reduction

Today’s high-fidelity computational fluid dynamics (CFD) and high-end experimental data acquisition systems tend to produce vast amounts of data that are getting harder

[†] Email address for correspondence: moritz.sieber@fd.tu-berlin.de

to interpret and overview. Methods to analyse such data are numerous and are always developing to stay in line with acquisition and computation systems. The most challenging data stem from turbulent flows that feature a huge range of temporal and spatial scales. A key challenge in turbulent flow data mining is the distinction of deterministic coherent motion from purely stochastic motion. Numerous methods exist that exploit the periodicity or energetic dominance of these coherent structures. These methods range from classic Fourier decomposition to dynamic mode decomposition (DMD) and proper orthogonal decomposition (POD). The most prominent among them are briefly introduced in the following.

POD has been widely used since its introduction by Lumley (1970) and Sirovich (1987). It was applied in nearly every fluid dynamic field. Beyond fluid dynamics, this method is also known as singular value decomposition, principal component analysis or Karhunen–Loève expansion (Berkooz, Holmes & Lumley 1993). The basic idea behind this method is to construct an optimal basis that represents most of the data variance with as few basis functions as possible. In the context of POD the variance is turbulent kinetic energy. Therefore, the POD searches for the most energetic modes whereby coherent structures with high energy content are likely to be represented by POD basis functions (Holmes *et al.* 2012).

Another classical approach is the linear stochastic estimation introduced by Adrian & Moin (1988), where the readings of different sensors are related via a linear mapping. This is closely related to the extended POD (Boree 2003), also described in a unified framework (observable inferred decomposition) by Schlegel *et al.* (2012). In recent extensions of linear stochastic estimation, the use of time delays between the different sensors and also the use of one sensor at multiple time instances are pursued to separate periodic coherent structures from turbulent fluctuations (Durgesh & Naughton 2010; Lasagna, Orazi & Iuso 2013). This approach was also used to improve the determination of harmonic POD modes from few pressure sensors (Hosseini, Martinuzzi & Noack 2015). These utilizations of data from various time instances are also related to the temporal constraint used for the POD extension proposed in this article.

Targeting the temporal periodicity of the coherent structures, spectral methods such as discrete Fourier transform (DFT) and the recently introduced DMD (Rowley *et al.* 2009; Schmid 2010) come into play. These methods commonly span the mode space according to fixed frequencies, which enables the identification of coherent structures within small spectral bandwidths. In contrast to DFT, DMD also distinguishes modes with respect to their linear amplification. The recently introduced extended DMD (Williams, Kevrekidis & Rowley 2015) tries to overcome the limitations encountered by the (linear) DMD approach when trying to decompose data from nonlinear systems. The idea is to use nonlinear functions that create observables of the data, which are exactly described by a linear system. This approach translocates the problem towards the identification of these nonlinear functions, which can be solved using the ‘kernel trick’ that is common in machine learning. This paper presents an alternative approach, which extends POD to account for temporal dynamics in addition to energetic optimality.

1.2. *Why yet another method?*

After this short and incomplete review of data processing methods, one may ask if there is need for another method. The answer is probably no, so we take the most used method (POD) and bring it up to date for present research issues. The

approach pursued here includes a simple yet effective extension to the classical POD, which leads to a more general method comprising POD and also DFT. This approach unifies existing methods, but also offers possibilities beyond these. From the authors' experience, the currently available methods often fail when applied to challenging flow data. These stem from flows with weak coherent structures where the recorded data have low signal-to-noise ratios, from flows with intermittent dynamics, or from flows featuring multi-modal interactions leading to frequency modulations, to name a few. In such cases, much effort is required to optimize the data processing until satisfactory results are obtained. The usual escape route is to focus on a certain spatial region or to apply suitable filters to pick out a certain wavelength or frequency range. This involves trial and error or requires prior knowledge of the investigated flow. There is also the danger of cutting off a substantial portion of the data, leading to false interpretations. These procedures can be collected under the heading 'identifying symmetries' as done by Holmes *et al.* (2012). The drawback of this approach is that the investigated flow must feature some symmetries and they must be known *a priori*. A recent application shows the huge variety of spatial and temporal filtering together with POD to separate different phenomena into different modes (Bourgeois, Noack & Martinuzzi 2013), exemplifying the complexity of this approach.

The usage of spectral methods for highly turbulent flows is even more challenging than POD. The variable frequency of single coherent structures and intermittent occurrence of different structures with the same frequency hinders a proper decomposition. In terms of the DFT, averaging of spectra from multiple measurements or sensors is essential to obtain reliable results. Analogously, for DMD, averaging over several events is an option to reject noise (Tu *et al.* 2014). Nonetheless, the results obtained with DFT and DMD suffer from limiting the temporal dynamics to single frequencies. Turbulent flows hardly ever feature discrete frequencies and it is not always valuable to restrict a single mode (flow phenomenon) to a single frequency. Coherent structures that feature significant phase jitter or frequency modulation are represented by many modes at similar frequencies. In contrast, the POD puts no temporal constraint on the modes. This can result in modes that represent flow phenomena occurring at largely different temporal scales. Thus, it is often hard to interpret these modes and draw meaningful conclusions from the temporal dynamics.

From our point of view, there is a big gap between the energetically optimal decomposition of POD and the spectrally clean decomposition of DFT or DMD. This gap will be bridged with the spectral proper orthogonal decomposition (SPOD) introduced in this article. This new method not only places itself somewhere in between these two extrema, but it allows for a continuous shifting from one to the other. The main idea is to apply a filter operation to the POD correlation matrix, which will force the POD towards clear temporal dynamics. Depending on the filter strength we continuously sweep from classic POD to DFT.

The remainder of this article is organized as follows: the proposed method is described in detail in § 2. The reader is guided from snapshot POD via an in-depth interpretation of the correlation matrix towards the general description of the SPOD. In addition, a method is explained to identify coupled mode pairs describing a single coherent structure, which becomes handy when working with SPOD. In § 3, the new method is demonstrated on three different experimental data sets. The results are compared against POD and DFT to point out the benefits of SPOD. In § 4 the capabilities of SPOD are summarized, based on the findings from the application to experimental data.

2. Description and interpretation of the proposed method

2.1. Classical snapshot POD

To introduce the method and the nomenclature, the snapshot POD approach is described first. We start off with a decomposition of a data set into spatial modes and temporal coefficients:

$$u(\mathbf{x}, t) = \bar{u}(\mathbf{x}) + u'(\mathbf{x}, t) = \bar{u}(\mathbf{x}) + \sum_{i=1}^N a_i(t) \Phi_i(\mathbf{x}). \quad (2.1)$$

Note that only the fluctuating part $u'(\mathbf{x}, t)$ is decomposed. It is split into a sum of spatial modes Φ_i and mode coefficients a_i . A set of M spatial points recorded simultaneously over N time steps is considered. To calculate the POD, the correlation matrix of this data set is needed. For data obtained from particle image velocimetry (PIV) or CFD, the number of spatial points is usually larger than the number of snapshots. The correlation matrix is then calculated between individual snapshots (temporal correlation). The alternative approach (spatial correlation) that applies to $M \ll N$ is detailed in appendix A. The correlation between two snapshots is calculated from an appropriate inner product $\langle \cdot, \cdot \rangle$, usually defined as the L^2 inner product

$$\langle u(\mathbf{x}), v(\mathbf{x}) \rangle = \int_V u(\mathbf{x})v(\mathbf{x}) \, dV, \quad (2.2)$$

where V specifies the spatial region or volume over which the correlation is integrated. The elements of the correlation matrix \mathbf{R} are given by

$$R_{i,j} = \frac{1}{N} \langle u'(\mathbf{x}, t_i), u'(\mathbf{x}, t_j) \rangle. \quad (2.3)$$

Matrix \mathbf{R} is of size $N \times N$.

The temporal coefficients $\mathbf{a}_i = [a_i(t_1), \dots, a_i(t_N)]^T$ and mode energies λ_i are obtained from the eigenvectors and eigenvalues of the correlation matrix:

$$\mathbf{R}\mathbf{a}_i = \lambda_i \mathbf{a}_i; \quad \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N \geq 0. \quad (2.4)$$

The subscript i refers to single eigenvalues, which are sorted in descending order. Since the \mathbf{a}_i are the eigenvectors of the real symmetric positive-definite matrix \mathbf{R} , they are orthogonal. Moreover, they are scaled with the energy of the single modes such that

$$\frac{1}{N} (\mathbf{a}_i, \mathbf{a}_j) = \lambda_i \delta_{ij}, \quad (2.5)$$

where (\cdot, \cdot) denotes the scalar product. The spatial modes are obtained from the projection of the snapshots onto the temporal coefficients:

$$\Phi_i(\mathbf{x}) = \frac{1}{N\lambda_i} \sum_{j=1}^N a_i(t_j) u'(\mathbf{x}, t_j). \quad (2.6)$$

These modes are orthonormal by construction, i.e.

$$\langle \Phi_i, \Phi_j \rangle = \delta_{ij}. \quad (2.7)$$

The formulation so far is perfectly in line with classical snapshot POD, which can also be computed by a singular value decomposition. However, since the SPOD requires a manipulation of the correlation matrix we retain the classical form.




FIGURE 1. (Colour online) (a) Pseudo-colour plot of the correlation matrix elements $R_{i,j}$ and (b) the corresponding correlation coefficient \hat{R} . The displayed data are selected from PIV measurements of a forced turbulent jet.

2.2. Properties of the correlation matrix

The SPOD described in this article is essentially a filter applied to the correlation matrix \mathbf{R} . To offer a better understanding of this approach, the structure of the correlation matrix \mathbf{R} is inspected first.

Figure 1(a) shows the structure of the correlation matrix for the data set of a forced turbulent jet. The data were acquired with PIV inside a 2-D plane aligned with the jet axis. The considered flow shows strong vortex shedding at the forcing frequency (the acquisition frequency is 25 times the forcing frequency). The presence of these periodic patterns in the flow, and their convection within the observed flow field, leads to a diagonal, wave-like structure of the matrix. This is closely related to the periodicity of the autocorrelation coefficient. In fact, if the individual elements of the correlation matrix \mathbf{R} are summed up along the diagonals, we obtain the spatially averaged autocorrelation coefficient

$$\hat{R}(\tau) = \frac{\int_{\tau}^T \langle u'(\mathbf{x}, t), u'(\mathbf{x}, t - \tau) \rangle dt}{\int_0^T \langle u'(\mathbf{x}, t), u'(\mathbf{x}, t) \rangle dt}, \quad (2.8)$$

where the upper bound T is the length of the measured sequence. This is depicted in figure 1(b), showing the same periodicity as the correlation matrix. The autocorrelation coefficient itself represents the spectral content of different time scales and wavelengths and it is directly related to the power spectral density of the underlying data. However, it contains no information on the phase of individual frequencies, due to the reference of the signal to itself. This is why the elements along the diagonals of \mathbf{R} look so similar, as they represent only relative changes with respect to the time step on the main diagonal. Thus, increased similarity along the diagonals of \mathbf{R} is equivalent to an increased similarity of the dynamics of the underlying signal. This property will be discussed in more depth in § 2.4. The obvious consequence from these findings is: if we want to obtain smooth dynamics from the POD, we have to enforce the diagonal similarity of the correlation matrix. This is where we step into spectral POD.

2.3. General description of the SPOD

The yet so simple but radical approach is a filter operating on the correlation matrix \mathbf{R} . To augment the diagonal similarity of \mathbf{R} a simple low-pass filter is applied along the diagonals. This results in a filtered correlation matrix \mathbf{S} with the elements given as

$$S_{i,j} = \sum_{k=-N_f}^{N_f} g_k R_{i+k,j+k}. \quad (2.9)$$

The filter above is just a symmetric finite impulse response filter with a filter coefficients vector \mathbf{g} of length $2N_f + 1$. The most simple approach would be a box filter, where all coefficients have the same value $g_k = 1/(2N_f + 1)$. In the examples discussed later, we use a Gaussian filter, which features a smooth response in time and frequency domain. Moreover, we choose a standard deviation such that the filter gives the same cutoff frequency as a box filter with half the length. In fact, any kind of digital finite impulse response filter can be used here.

The further procedure of the SPOD is the same as for the classical POD. From the filtered correlation matrix \mathbf{S} the temporal coefficients \mathbf{b}_i and mode energies μ_i are obtained from the eigendecomposition

$$\mathbf{S}\mathbf{b}_i = \mu_i \mathbf{b}_i; \quad \mu_1 \geq \mu_2 \geq \dots \geq \mu_N \geq 0. \quad (2.10)$$

The temporal coefficients are also scaled with the mode energy and they are still orthogonal:

$$\frac{1}{N} (\mathbf{b}_i, \mathbf{b}_j) = \mu_i \delta_{ij}. \quad (2.11)$$

The spatial modes are finally obtained from the projection of the snapshots onto the temporal coefficients,

$$\Psi_i(\mathbf{x}) = \frac{1}{N\mu_i} \sum_{j=1}^N b_i(t_j) u'(\mathbf{x}, t_j), \quad (2.12)$$

where these modes are no longer orthogonal. This property of the spatial modes is detailed in appendix B. The total energy of the data set is still represented by the decomposition ($\sum \lambda_i = \sum \mu_i$), but the energy per mode is less for the first modes. Hence, increasingly plain temporal dynamics is obtained at the expense of spatial orthogonality and a dispersed SPOD spectrum. Nevertheless, the decomposition (as in (2.1))

$$u(\mathbf{x}, t) = \bar{u}(\mathbf{x}) + \sum_{i=1}^N b_i(t) \Psi_i(\mathbf{x}) \quad (2.13)$$

is still exact if all N SPOD modes are used for the recomposition.

If the filter size is extended over the entire time series, the filtered correlation matrix converges to a symmetric Toeplitz matrix. This matrix has the form

$$S_{i,j} = \widehat{R}(\Delta t |i - j|), \quad (2.14)$$

with the diagonals given by the average correlation coefficient (2.8). This special matrix is also known as the covariance matrix and its eigenvalues trace out the power spectral density of the underlying time series (Wise 1955). This equality is a part of Szegő's theorem and it is valid for the limiting case where the number of samples

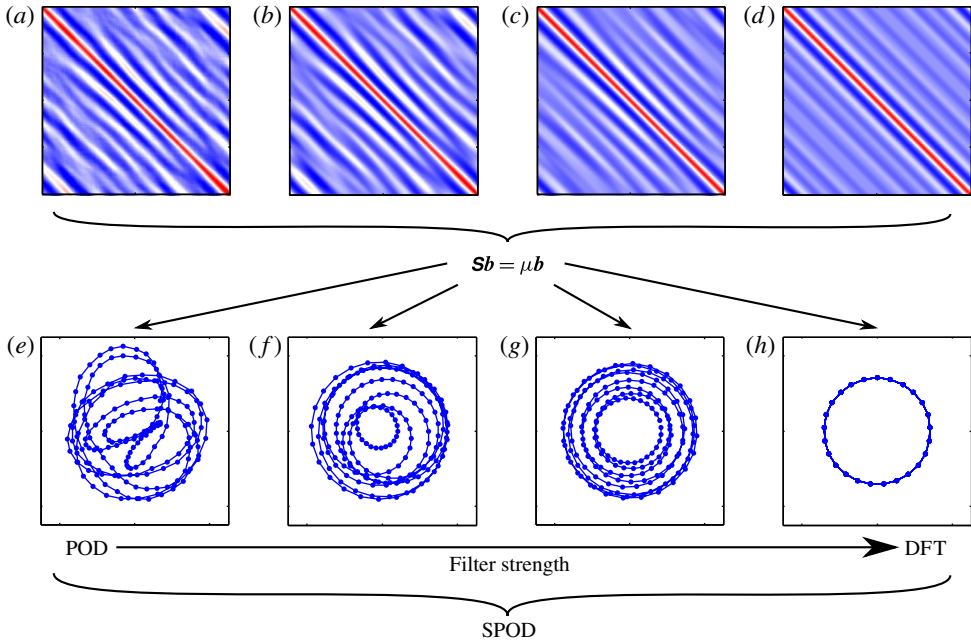


FIGURE 2. (Colour online) Schematic describing the main properties of the SPOD for increasing filter strength (from left to right): (a,e) $N_f = 0$; (b,f) $N_f = 25$; (c,g) $N_f = 100$; (d,h) $N_f = 200$. (a–d) Pseudo-colour plots of the filtered correlation matrix (\mathbf{S}). (e–h) The phase portraits of the corresponding first two modes (\mathbf{b}_1 and \mathbf{b}_2) that describe the dominant oscillations. The axes of the plots shown here are the same as for the plots in figures 1 and 16, respectively. The graphs are based on the data already presented in § 2.2 and the SPOD is calculated from 200 snapshots.

approaches infinity. To discuss this feature for finite series, the treatment of the start and end of the time series must be clarified. At the boundaries of \mathbf{R} , the filter operation is not properly defined, since the symmetric filter lacks elements before and after the finite series. Either these elements can be replaced by zeros or the time series is assumed to be periodic. For the zero padded boundary, $R_{i,j} = 0$ in (2.9) for any i or j that is outside of the domain $[1, N]$. In the case of periodic boundary conditions, indexing outside of the valid domain is mapped back into the domain by addition or subtraction of N (e.g. $R_{i,j} = R_{i-N,j}$ if $i > N$). For periodic conditions and a box filter of the same size as the number of snapshots, a symmetric circulant matrix is obtained, where the eigenvalues and eigenvectors are given by the Fourier transform of the first row (Gray 2005). Hence, DFT and SPOD produce the same decomposition for this limiting case. For the Gaussian filter this limit is only exactly reached with an infinitely large filter size, but practically, the decomposition remains nearly constant for values $N_f > N$. In the following example of application, the DFT is calculated with $N_f = N$ and a box filter.

The general SPOD approach and its placement with respect to the existing methods are graphically summarized in figure 2. The images in the first row show the filtered correlation matrix at different filter widths N_f . The images below depict the phase portraits of the leading two modes (compare figure 16). It is apparent that the increased diagonal similarity of the correlation matrix, which goes in hand with




FIGURE 3. (Colour online) Magnitude of the frequency spectrum for the first SPOD mode for different filter widths N_f . The analysed flow is the forced turbulent jet presented in figures 1 and 2.

the increased filter width, successively limits the temporal variations of the mode amplitude and frequency until a stable limit cycle is reached. In summary, the SPOD allows for continuous fading from the energetically optimal POD to a purely spectral DFT. What happens in between these two limits is very promising, as will be shown later in this article.

2.4. Spectral properties of the SPOD filter operation

Now provided with the continuous transition between POD and DFT, we may determine how the SPOD filter affects the resulting mode dynamics. In other words, we would like to know how the filter frequency response transfers into mode space. This enables the choice of an adequate filter size N_f , which is the crucial property of the introduced filter operation.

We consider the forced turbulent jet data already consulted in the previous section. In this flow the periodic shedding is strong and regular and it is represented by the same SPOD mode regardless of the filter setting. Furthermore, the flow is turbulent, showing sufficient jitter and broadband noise in the mode dynamics to demonstrate the effect of the SPOD filter. Figure 3 shows the spectrum of the most dominant SPOD mode coefficient b_1 computed for different filter sizes N_f . According to these graphs, the SPOD operation applies a band-pass filter to the modal dynamics with the band centred around the dominant frequency. The width of the band and the spectral attenuation are directly related to the filter width.

The spectral response of the filter coefficients \mathbf{g} that define the spectral constraint (2.9) of the SPOD are drawn in figure 4. As described in the previous section, these coefficients are chosen to represent a Gaussian filter with a typical low-pass characteristic, where the cutoff frequency is related to the filter width by $f_c \propto 1/N_f$.

To evaluate the frequency response of the entire SPOD method, we compute the spectral attenuation of the SPOD with respect to POD and compare this to the frequency response of the filter coefficients. The spectral attenuation is calculated from the ratio of spectral magnitudes of the first SPOD ($N_f = 25$) and POD ($N_f = 0$) coefficients b_1 . The results are shown in figure 5, where the graph of the filter frequency response must be shifted to compensate for the difference between low- and band-pass filtering. Accordingly, both curves show nearly the same attenuation as well




FIGURE 4. (Colour online) Magnitude of the SPOD filter frequency response for different filter widths N_f .




FIGURE 5. (Colour online) The spectral attenuation of the SPOD obtained from the ratio of the spectral magnitudes $|\mathcal{F}(b_{N_f=0})|/|\mathcal{F}(b_{N_f=25})|$ and the response of the filter coefficients $|\mathcal{F}(g)|$ from (2.9). For the filter response the frequency axis is offset to match the modes' centre frequency.

as the same characteristic stop-band ripples. In this adapted representation the spectral attenuation of the SPOD can be directly estimated from the filter coefficients. However, the SPOD shows a band-pass characteristic while the filter coefficients describe a low-pass filter. The frequency around which the band is centred is selected intrinsically by the SPOD.

The band-pass behaviour of the SPOD is a direct consequence of the smoothing of the diagonal elements of the correlation matrix. As elaborated in appendix C for linear system dynamics, variations of the diagonal elements relate directly to variations of the frequency and amplification of the modes. Hence, the SPOD filter operation constrains the temporal variation of the mode frequency and amplification rate, but sets no limitations on the temporal mean frequency and amplitude. Thus, the SPOD filter averages out frequency variations, which reduces the spectral bandwidth of the modes. With the selected low-pass filter coefficients, the smoothing intensity can be adjusted, which controls the bandwidth of the filter. However, the selection of a centre frequency for this band-pass filter is data driven, which means the filter 'snaps' to a dominant (coherent) fluctuation contained in the data.

Based on these conclusions, an adequate filter size N_f may be selected prior to the SPOD. Ideally, it should be a time span for which the coherent structures of interest exhibit constant dynamics. In practice, we found that suitable selections for this time span are the characteristic time scales of the flow. These are either one period of a dominant oscillation or the time that a structure needs to be convected a typical length scale of the flow.

Based on the coefficient spectra shown in figure 3, it may seem that SPOD is just a method that is comparable to a subsequent band-pass filtering of POD modes, but the situation is substantially different. A posterior filtering of POD modes cuts out parts of the temporal dynamics that are ignored in further investigations. In the case of the SPOD, the filter is applied prior to the decomposition. Hence, the filter shifts the dynamic content removed from one mode to the other modes. Therefore, the entire dynamic content is still represented in the decomposition, but it is redistributed among the modes (2.13). Moreover, the SPOD may ‘find’ modes that were hidden in noise or mixed with other modes, which is not possible with a subsequent filter.

2.5. Identification of coupled modes

One crucial point in POD and SPOD is the identification of linked modes. Assuming the presence of periodic coherent structures, their dynamics is described by a pair of modes, analogues to the sine and cosine in the Fourier space or the real and imaginary parts of DMD modes. They constitute the real and imaginary parts of an analytical signal,

$$\tilde{b}(t) = b_i(t) + ib_j(t) = A(t)e^{i\phi(t)}, \quad (2.15)$$

where A is the amplitude and ϕ the phase of the signal ($i = \sqrt{-1}$). The coupling of such a mode pair is not given by the SPOD and it has to be identified *a posteriori*. Coupled modes typically show a similar amount of energy and pair in the POD spectrum (Oberleithner *et al.* 2011; Oberleithner, Rukes & Soria 2014). For more complex dynamics with multiple energetic modes, the pairs are not easily identified and visual inspection of Lissajous figures and spatial modes is required. This manual procedure is cumbersome and by no means objective.

To provide an alternative, we propose an unbiased approach that gives a quantitative measure of the dynamic coupling of individual modes. The idea is to evaluate the mode coefficients cross-spectra at a $\pi/2$ phase lag, in order to measure the spectral proximity. Mode pairs that describe a single structure must have the same spectral content, but are shifted by a quarter period. In the following, this spectral measure is also called harmonic correlation. Fourier decomposition and DMD were both considered to evaluate this harmonic correlation, but the DMD turned out to be more reliable for this task. The general procedure is schematically outlined in figure 6. From the depicted operations, the DMD and the mode coupling are discussed in this section.

For the DMD of the SPOD coefficients, their temporal evolution is assumed to be governed by a linear operator \mathbf{T} ,

$$\mathbf{b}(t + \Delta t) = \mathbf{T}\mathbf{b}(t). \quad (2.16)$$

To approximate this operator, the SPOD coefficients are arranged in two matrices $\mathbf{X} = [\mathbf{b}(0) \ \mathbf{b}(\Delta t) \ \dots \ \mathbf{b}((N-2)\Delta t)]$ and $\mathbf{Y} = [\mathbf{b}(\Delta t) \ \mathbf{b}(2\Delta t) \ \dots \ \mathbf{b}((N-1)\Delta t)]$ (following the notation of Tu *et al.* (2014)). Hereafter, the operator is given by

$$\mathbf{T} = \mathbf{Y}\mathbf{X}^{-1}, \quad (2.17)$$

where \mathbf{X}^{-1} is the Moore–Penrose pseudo-inverse of \mathbf{X} . Alternatively, this can be solved as a least squares problem, minimizing $\|\mathbf{T}\mathbf{X} - \mathbf{Y}\|$. To reject measurement noise in the

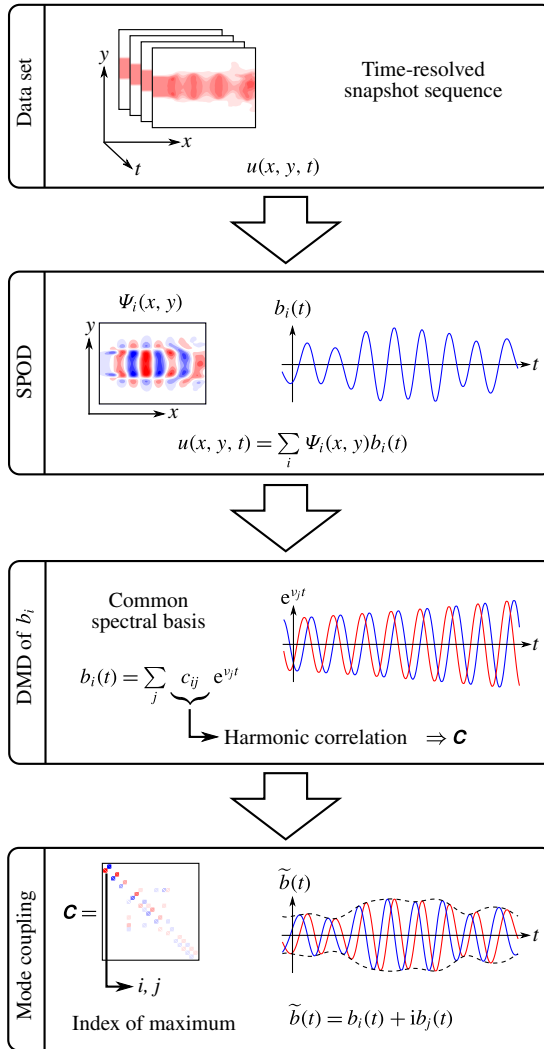


FIGURE 6. (Colour online) Schematic illustrating the main steps towards the identification of coupled modes (red and blue lines indicate real and imaginary parts of an analytic signal). The data displayed here were derived from measurements of a forced turbulent jet (see also figure 1).

identification procedure, only the ‘physical part’ of the SPOD modes is considered for the calculation of the operator \mathbf{T} . That means only the modes with acceptable signal-to-noise ratio should be considered. Therefore, the number of retained modes is calculated from the energy resolved by the SPOD, truncated after N_c modes, with

$$\mathcal{E}(N_c) = \frac{\sum_{k=1}^{N_c} \mu_k}{\sum_{k=1}^N \mu_k}. \tag{2.18}$$

In the examples shown later the modes are truncated around $\mathcal{E}(N_c) = 0.95$. This value depends on the signal-to-noise ratio of the considered measurement, which can be estimated from the POD spectrum (Raiola, Discetti & Ianiro 2015). Note that the number of retained modes increases for wider SPOD filters and correspondingly flatter SPOD spectra (μ_j).

The DMD modes are obtained by the eigendecomposition of matrix \mathbf{T} as

$$\mathbf{T}\mathbf{c}_i = \nu_i\mathbf{c}_i. \quad (2.19)$$

The eigenvalues ν_i comprise the frequencies ω_i and amplification rates σ_i of the operator \mathbf{T} , which are related to the logarithm of the eigenvalue via $\ln(\nu_i)/\Delta t = \sigma_i + i\omega_i$. The eigenvectors \mathbf{c}_i hold the relative spectral content of the single SPOD coefficients with respect to these frequencies. More precisely, the element $c_{i,j}$ of vector \mathbf{c}_i is the spectral content of the single mode coefficient b_j with respect to ν_i . The actual modal representation is given by

$$b_j(t) = \sum_{i=1}^{N_c} c_{i,j} e^{(\sigma_i + i\omega_i)t}. \quad (2.20)$$

It must be noted that this decomposition is only exact if $N_c = N$, whereas in the current approach $N_c < N$. Nevertheless, the decomposition gives a common spectral basis, which allows the ranking of spectral similarity of the temporal coefficients $\mathbf{b}(t)$. The developed proximity measure is given by

$$C_{i,j} = \text{Im} \left\{ \sum_{k=1}^{N_c} c_{k,i} c_{k,j}^* \text{sgn}(\text{Im}(\nu_k)) \right\}, \quad (2.21)$$

where $*$ indicates the complex conjugate and the coefficients are normalized so that $(\mathbf{c}_i, \mathbf{c}_i) = 1$. The sign function (sgn) in this expression accounts for the conjugate pairs that appear in the DMD spectrum (mirrored at the real axis).

For two modes to be coupled, they must have a similar spectral content, which is shifted a quarter period forward or backward. This implies a positive or negative imaginary part of the harmonic correlation (2.21), respectively, and coupled modes appear as peaks in the matrix \mathbf{C} . Hence, the row and column indices of the maximum of \mathbf{C} identify the first coupled SPOD modes. The corresponding row and column in \mathbf{C} are then set to zero and the next lower maximum is identified. This procedure is repeated until all modes are paired. It has to be noted that this approach also creates weakly correlated and possibly unphysical mode pairs.

Together with the identification of coupled modes, the procedure gives an average frequency of the coherent structure represented by the mode pair. Therefore, the eigenvalues ν_k of the matrix \mathbf{T} are sorted in descending order with respect to their content for the identified mode pair $\tilde{c}_k = c_{k,i}^2 + c_{k,j}^2$. The frequency is given as the weighted sum of the eigenvalues,

$$f = \frac{\sum_{i=1}^n \text{Im}\{\ln(\nu_i)\} \tilde{c}_i}{2\pi\Delta t \sum_{i=1}^n \tilde{c}_i}. \quad (2.22)$$

The weighting accounts for the relative energy content of a mode pair with respect to the single frequencies. In fact, only the most relevant eigenvalue ($n = 1$) can be chosen

to determine the frequency, but for practical application it is recommended to use more than one eigenvalue as noise may corrupt them. For the examples discussed in the next chapter, we used an average over three eigenvalues ($n = 3$) to obtain accurate results.

The coupled SPOD modes are considered as one complex mode (see (2.15) and figure 6) similar to the Fourier mode. The relative energy content of the identified modes is computed as

$$K = \frac{\mu_i + \mu_j}{N_c \sum_{k=1} \mu_k}, \quad (2.23)$$

where i and j again refer to the indices of the coupled SPOD modes.

3. Applications to experimental data

In this chapter the SPOD is applied to three different data sets. The three examples originate from very different engineering problems, demonstrating the capability and broad applicability of SPOD. We consider the flow of a swirl-stabilized combustor, the wake of an airfoil with Gurney flap and the flow field of a sweeping jet generated with a fluidic oscillator. All three flows were recorded with the same PIV measurement system. It consists of a Photron Fastcam SA 1.1 high-speed camera (1 Mpixel at 2.7 kHz double frame) and a Quantronix Darwin Duo laser (30 mJ at 1 kHz). The PIV data were processed with PIVview (PIVTEC GmbH) using standard digital PIV processing (Willert & Gharib 1991) enhanced by iterative multi-grid interrogation (Soria 1996) with image deformation (Huang, Fiedler & Wang 1993; Raffel *et al.* 2007, pp. 146–158). Analysing the present data sets with existing POD, DFT or DMD approaches caused some difficulties. As will be demonstrated, the SPOD is able to handle these shortcomings. The DMD and the DFT equally suffer from the restriction of the modes to narrow frequency bands, therefore we limit the following presentation to DFT. This choice is particularly useful as the DFT is a limiting case of the SPOD. We also performed DMD of the data presented here, where we found all DMD eigenvalues to lie in close proximity of the unit circle (distance $< 10^{-6}$). This indicates that the DMD is nearly equivalent to the DFT. Rowley *et al.* (2009) stated that a DMD is identical to a DFT for periodic time series, but from our observations it appears as if they are also almost identical for statistically stationary time series. This supposition is based on the fact that none of the analysed series is entirely periodic, but all are statistically stationary.

3.1. Swirling jet undergoing vortex breakdown

First, we consider the flow field of a swirl-stabilized combustor. Swirling jets are widely used in the gas turbine industry due to their capability of obstacle-free flame anchoring and enhanced mixing. The experimental set-up to study these flows is sketched in figure 7. Swirl is generated by injecting fluid tangentially into a mixing tube that terminates in the combustion chamber. The cylindrically shaped chamber is made of quartz glass to allow optical access for PIV. Flow measurements are conducted in the meridional section as indicated in the schematic. The case investigated here is non-reacting at a Reynolds number of 58 000 based on the nozzle diameter D and the bulk velocity at the nozzle exit. Additional details about the experimental set-up can be found in Reichel, Terhaar & Paschereit (2015).




FIGURE 7. (Colour online) Experimental set-up of the swirl-stabilized combustor. Air enters from the left, passes a swirl generator and exits into the combustion chamber. Flow field measurements with PIV are conducted in the meridional plane indicated by the dashed square (ROI).




FIGURE 8. (Colour online) Swirl-stabilized combustor flow: time-averaged flow field depicted by (a) contours of velocity magnitude and streamlines, and (b) spatially averaged power spectral density.

The mean flow field and the spatially averaged power spectral density (PSD) are depicted in figure 8, with the Strouhal number based on the same length and velocity scale as the Reynolds number ($St = fD/u_{bulk}$). The flow exhibits a strong recirculation zone in the centre, surrounded by an annular, strongly diverging jet. The cross-sectional jump at the combustion entrance creates an additional external recirculation zone between the jet and the confining walls. The spectral content of the flow is spread over all time scales and it decreases with increasing Strouhal number. The spectrum gives no indication of any dominant coherent structure even though these flows typically feature helical global modes (Oberleithner *et al.* 2011).

Figure 9(a–c) illustrate the results from the SPOD for the filter lengths $N_f = 0$, 10 and 2000, respectively. Note that the limiting cases $N_f = 0$ and $N_f = 2000$ produce results equivalent to those obtained with classical POD and DFT respectively, while the case in between represents the SPOD. Hence, this particular presentation concisely demonstrates the difference between POD, SPOD and DFT.

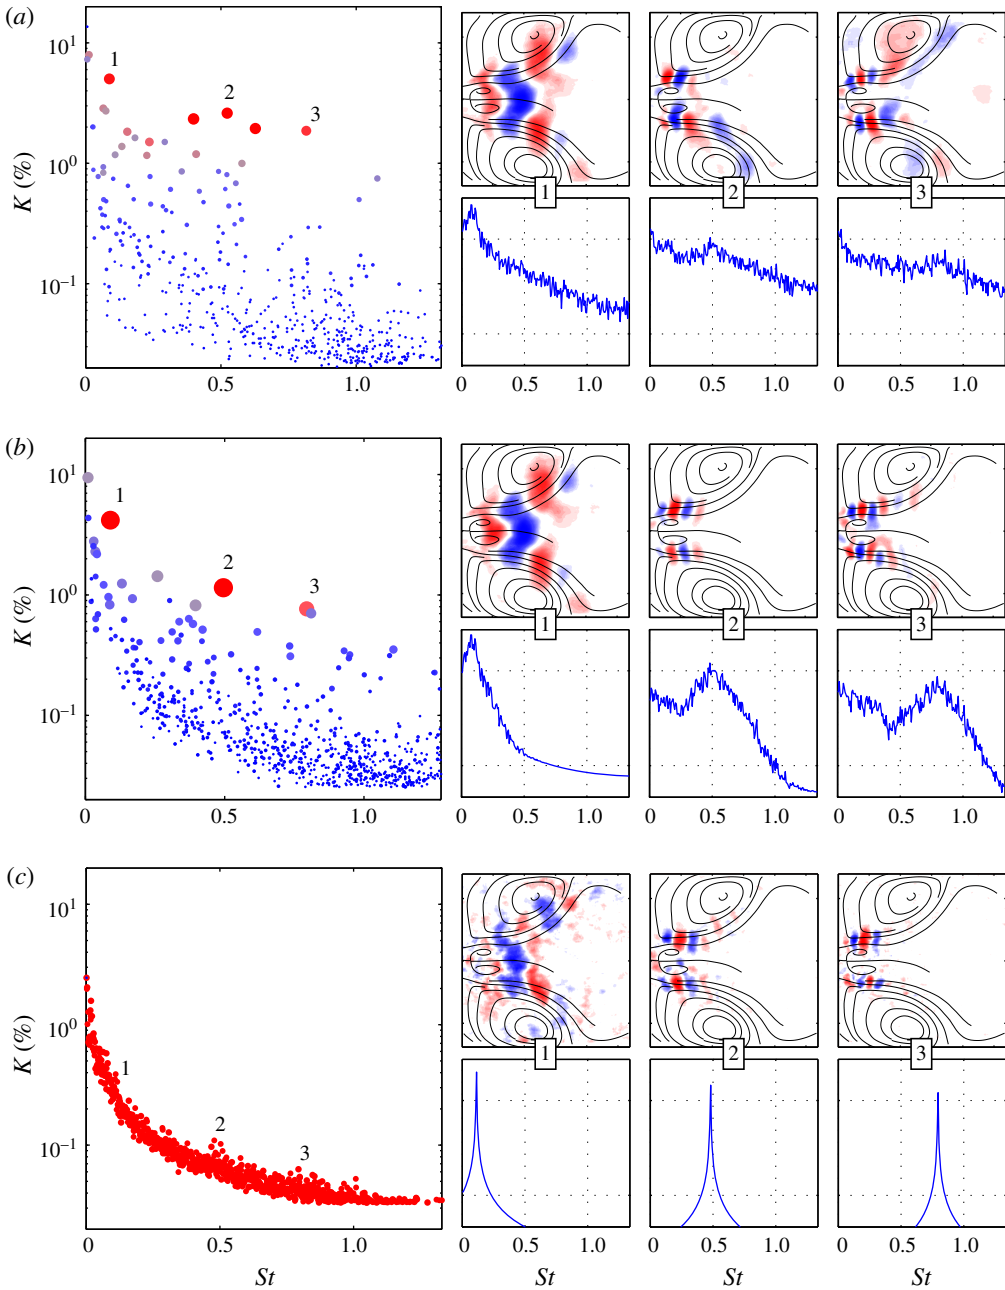


FIGURE 9. (Colour online) Swirling jet: results from SPOD for different filter lengths (a) $N_f = 0$ (POD), (b) $N_f = 10$ (SPOD) and (c) $N_f = 2000$ (DFT). For each filter length the SPOD spectrum is displayed as a scatter plot (left), where a single dot indicates one mode pair (size and colour $C_{i,j}$ in (2.21)). For three selected pairs the spatial modes (upper row) and PSD of the temporal coefficient (lower row) are depicted. They are indicated by numbers in the SPOD spectrum, as well as between the small mode plots.

Each of the three cases in figure 9 shows the so-called SPOD spectrum, where every mode pair is represented by a single dot. The size and colour of the dots indicate the harmonic correlation of a mode pair $C_{i,j}$, according to (2.21). The frequency of a mode pair is determined according to (2.22) and the energy from the two eigenvalues relative to the total energy from (2.23). On the right side of every case in figure 9, three spatial modes $\Psi_i(\mathbf{x})$ and the corresponding temporal coefficients $b_i(t)$ are plotted above each other. The spatial modes are visualized by the cross-wise velocity component (in the y -direction) together with streamlines of the time-averaged flow. They are numbered likewise in the SPOD spectrum and in the small mode plots. The plots are given without axis labels to allow a compact representation of the data; the section is the same as for the mean flow shown in figure 8(b). The time coefficients are represented by their power spectral density, where the time series is split into five (50% overlapping) sections, which are Fourier transformed and averaged. The horizontal dotted lines in the PSD plots indicate a spacing of three orders of magnitude (10^3) and the y -axis is scaled logarithmically. The spectral averaging was also applied to the power spectra shown in figure 8(b) and in the subsequent PSD plots.

The POD (figure 9a) yields a broad spectrum of modes, where modes at lower Strouhal numbers have more energy. There are several modes with high harmonic correlation (diameter and colour of the points) and high energy content K . The spatial shape of the low-frequency mode (label 1) shows clear spatial symmetry and a limited spectral bandwidth ($St \approx 0.1$). This mode is frequently observed in swirl-stabilized combustors and it is associated with a global hydrodynamic instability (Terhaar, Oberleithner & Paschereit 2015). From the four additional outstanding modes between $St = 0.3$ and $St = 0.8$, we choose two for further investigation. Their spatial structures cannot be precisely assigned to a symmetric or anti-symmetric mode and they indicate mixtures of several spatial wavelengths. Accordingly, the mode spectra are broad and show only a slight hump at the frequencies indicated by the SPOD spectrum. The other modes around $St = 0.5$, which are not shown here, show similar spatial and spectral content.

Overall, the POD indicates the presence of a single mode at low frequency, together with other coherent structures that are not properly resolved. The most energetic mode (not inspected here) corresponds to a low-frequency, large-wavelength fluctuation, as indicated by the SPOD. Such slow changes of the mean flow are usually called shift modes (Luchtenburg *et al.* 2009; Hosseini *et al.* 2015). In this particular case the shift mode stems from weak movements of the vortex breakdown bubble.

Consider now the SPOD in figure 9(b), with a filter length $N_f = 10$; from the SPOD spectra we identify three peaks at $St = 0.09$, 0.5 and 0.8. The first mode is the same as the one already identified by the POD, but its spectral content at higher frequencies is reduced. It describes a single-helical structure in the wake of the recirculation zone. The second identified mode exhibits a broad spectral peak at $St = 0.5$. Its spatial structure and Strouhal number match the global mode identified by Oberleithner *et al.* (2011). It is a single-helical mode linked to the precessing motion of the recirculation zone. The spatial structure of mode three suggests a double-helical shape. It is not a harmonic of the second mode, as their frequencies are not related.

When the filter size is extended to its maximum, the DFT (figure 9c) and SPOD spectrums converge to the spatially averaged PSD (figure 8b). The temporal coefficients converge to sines and cosines and all mode pairs show full harmonic correlation (uniform dot size in the SPOD spectrum). Since selection based on harmonic correlation is impossible, we resort to the frequencies already known from

the SPOD at $N_f = 10$. The spatial structures resemble the ones from figure 9(b), but they are corrupted by noise. Moreover, the spatial symmetries are no longer as clear as for the $N_f = 10$ case. Note that the corresponding spectral peaks are broadened due to the averaging procedure, which is applied here only for consistency.

From this first example, we can point out some striking features of the SPOD. The SPOD is able to separate coherent fluctuation from stochastic turbulent fluctuations even though they both have the same energy content (see the SPOD spectrum in figure 9c). The classical POD yields partially mixed structures that cannot be assigned to distinct flow phenomena, whereas the SPOD properly separates these structures and identifies them from harmonic correlations. The DFT instead shows the same structures at the identified frequencies, but they are corrupted with noise and the method itself would give no clue about the frequencies of interest.

The structures identified with the SPOD may also be found with the POD if the decomposition is applied to a subsection of the measured domain. Moreover, the exploitation of spatial symmetries prior to the POD decomposition usually provides good results for this type of flows (Terhaar 2015). Nevertheless, these alternative approaches would require prior knowledge of the shape or spatial extent of the structures, whereas the SPOD requires none of these.

All modes identified by the SPOD show clear spatial symmetries and distinct spectral peaks. The frequency and shape of the first mode coincide well with previous experimental observations in swirl-stabilized combustors (Terhaar *et al.* 2015). The second mode is very similar to the one observed in unconfined swirling jets (Oberleithner *et al.* 2011). However, the presence of these different modes in a single flow configuration raises the question regarding their physical relevance. To deal with this issue, we conducted a linear stability analysis of the underlying mean flow, following the procedure outlined by Oberleithner *et al.* (2015). This analysis similarly delivered three unstable modes whose frequencies and azimuthal and axial wavenumbers match the SPOD modes surprisingly well. To limit the scope of this paper the analysis is not further detailed here. One important parameter of the SPOD is the filter size N_f , which is twice the period of the second mode. The experiences gained throughout the first application show that a filter size of one to two periods of the mode of interest gives the best results.

3.2. Airfoil with Gurney flap

The second flow configuration considered here is the flow behind a pitched airfoil equipped with a Gurney flap. The experimental set-up is shown in figure 10. It illustrates the working principle of the Gurney flap deployed at the pressure side of the airfoil. The flap creates additional lift (and drag), which can be used to locally control varying loads on large wind turbine airfoils (Bach *et al.* 2014, 2015a). The flow features around the Gurney flap are characterized by a single vortex that develops upstream of the flap and periodic shedding in its wake. The vortex upstream of the flap continuously grows up to a critical size, then it is shed into the wake, and it starts growing again. Here, a FX 63-137 airfoil at 5° angle of attack is investigated at a Reynolds number of 180 000 based on chord length. The reference length for the Strouhal number is the flap height, which is 3.6% of the chord. The measured region comprises only the wake of the airfoil capturing the dominant vortex shedding. More details about the experimental set-up can be found in a previous publication on these data (Bach *et al.* 2015b). The Strouhal number in the following results is calculated with the flap height h and the free stream velocity.




FIGURE 10. (Colour online) Schematic of the airfoil equipped with a Gurney flap at the trailing edge. Streamlines indicate the surrounding flow and the vortex upstream of the flap. The measured section (ROI) is a streamwise cut in the wake of the airfoil, capturing the periodic shedding behind the flap.




FIGURE 11. (Colour online) Wake of the airfoil with Gurney flap: time-averaged flow field depicted by (a) contours of velocity magnitude and streamlines, and (b) spatially averaged power spectral density. The origin of the coordinate system is located at the trailing edge.

The mean flow shown in figure 11(a) reveals a velocity deficit in the wake of the Gurney flap, which generates the vortex shedding. The PSD (figure 11b) indicates strong oscillations at $St = 0.105$ with a weak higher harmonic. In a previous investigation it was found that the vortex, which is shed from upstream of the flap, causes an alteration of the periodic vortex shedding behind the flap (Troolin, Longmire & Lai 2006). Hot-wire measurements in the wake of the Gurney flap supported this assumption. The combination of a strong periodic flow pattern and the intermittent short-time events provides a formidable benchmark for the SPOD.

The presentation of the decomposition with the different methods is organized in the same way as for the previous example. The classic POD decomposition is shown in figure 12(a). The vortex shedding is represented by the most energetic POD mode with the highest harmonic correlation. The remaining modes show weak harmonic correlations and no distinct peak in the SPOD spectrum. The modes labelled 2 and 3 exhibit a broad spectral content with a spatial extent limited to the vicinity of the flap. There are additional modes with similarly compact spatial extent located further downstream. These compact modes describe the intermittent alteration of the vortex

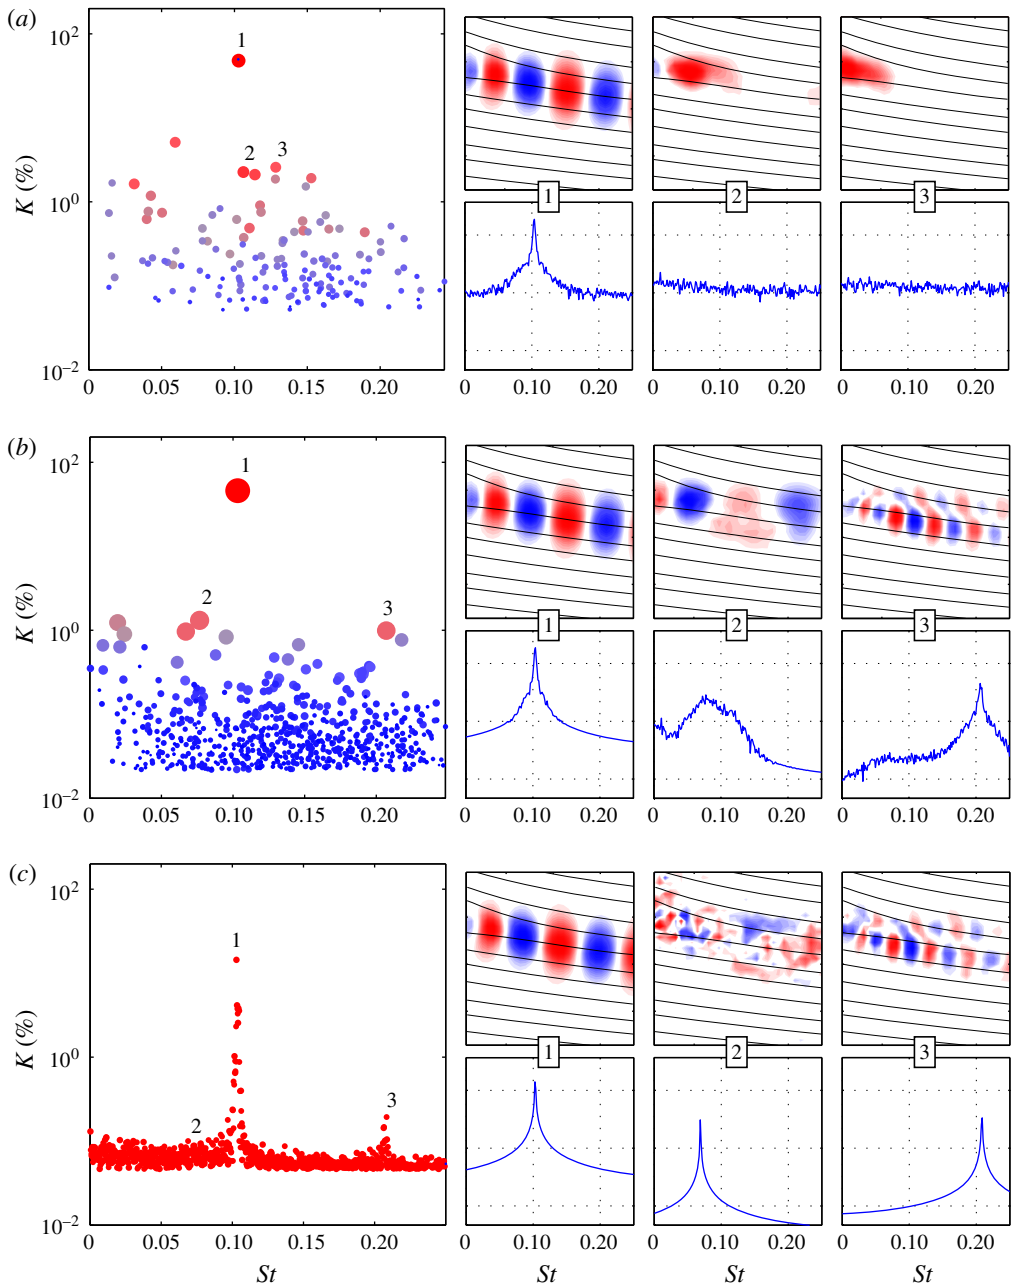


FIGURE 12. (Colour online) Airfoil with Gurney flap: results from SPOD for different filter lengths (a) $N_f = 0$ (POD), (b) $N_f = 15$ (SPOD) and (c) $N_f = 2000$ (DFT). For each filter length the SPOD spectrum is displayed as a scatter plot (left), where a single dot indicates one mode pair (size and colour $C_{i,j}$ in (2.21)). For three selected pairs the spatial modes (upper row) and PSD of the temporal coefficient (lower row) are depicted. They are indicated by numbers in the SPOD spectrum, as well as between the small mode plots.

shedding during the passage of the single vortex that was generated upstream of the Gurney flap. An inspection of their time coefficients (not shown) reveals that these modes are only active one after another during one shedding period. Depending on the phase lag between the natural oscillation and the shedding of the upstream vortex, the developing wake vortex is either strengthened or weakened. The convection of this altered vortex is described by a spatial series of modes. The modes 2 and 3 represent only the upstream part of this convection, while the additional modes located further downstream (not shown) describe the subsequent motion. This behaviour is indicated by a complex interaction of these modes and the periodic shedding modes, which is hard to identify in the POD expansion.

The SPOD yields a much clearer set of modes (figure 12*b*). In addition to the shedding mode, the SPOD also uncovers three other modes, which are offset from the rest. The two modes that appear at similar frequencies capture the alteration of the vortex shedding during the passage of the single vortex. Their mode shape is similar to the shedding mode, but with larger spatial wavelengths and lower frequencies (see mode 2 in figure 12*b*). The interaction of the upstream vortex with the vortex shedding increases the vortex size and thus the wavelength in the wake. Assuming a constant convection speed in the flow, this mode consequently has a lower frequency. In the case of the SPOD the alteration of the vortex shedding is captured by a single mode (pair) and is thus much easier to interpret. The third mode represents the second harmonic of the vortex shedding, with a clear spectral peak and clean spatial mode with twice the wavelength of the shedding mode. This higher harmonic is completely missed in the POD. The SPOD filter size N_f is equivalent to three shedding periods, which is approximately equal to the travelling time through the measurement domain.

The DFT shown in figure 12*c*) reproduces the spectrum shown in figure 11*b*). The natural mode and its higher harmonic can be identified from the spectral peaks. The corresponding mode shapes are similar to the SPOD, although the higher harmonic is corrupted with noise, resulting in a fragmented spatial mode. The DFT at the frequency of the second SPOD mode gives no indication of the structure identified before and the vortex–vortex interaction is completely missed. This is attributed to the fact that this phenomenon is highly intermittent with varying frequencies and amplitudes, which cannot be represented by a single-frequency mode. The same dilemma applies for the DMD.

For this example, the SPOD has shown its ability to separate dynamics with similar spatial structures and frequencies but very different energies. The spectral proximity and spatial similarity of the involved dynamics impede the application of POD. The modulation of the natural vortex shedding is represented by a natural mode with several intermittent modes. The DFT, however, with its single-frequency modes does not capture the modulation of the shedding at all. The frequency constraint imposed by the SPOD is sharp enough to split the natural shedding from the modulation and soft enough to allow for frequency and amplitude variations. Hence the SPOD again gives easy access to dynamic features of the flow, which cannot be found with other common methods of similar algorithmic complexity. There may be feature tracking approaches capable of identifying the dynamics in this case, but they usually require more computational effort and might not be as versatile as the SPOD.

3.3. Fluidic oscillator

In this example, SPOD is applied to the flow field of the sweeping jet generated from a fluidic oscillator. This device is essentially a nozzle with feedback channels,




FIGURE 13. (Colour online) Schematic of the experimental set-up with the fluidic oscillator. Air enters from the left, passes the oscillator and exits into the unconfined ambient air. The angle of the jet leaving the oscillator sweeps periodically up and down. The measured region (ROI) captures the meridional plane of the jet's near field. The oscillator has a square nozzle, hence the thickness of the jet normal to the plane is also D .




FIGURE 14. (Colour online) Fluidic oscillator: time-averaged flow field depicted by (a) contours of velocity magnitude and streamlines, and (b) spatially averaged power spectral density.

which cause a self-sustained oscillation of the jet. Figure 13 shows the approximate geometry of this device and indicates the meandering shape of the sweeping jet. The shape and motion of the jet resemble a travelling wave. These devices are used for active flow control applications, where the sweeping motion of the jet allows a much wider actuator spacing, resulting in less energy consumption (Woszidlo *et al.* 2014). The data presented here stem from an experimental set-up investigating the spreading and entrainment of sweeping jets (Ostermann *et al.* 2015b; Woszidlo *et al.* 2015). The data are recorded at a Reynolds number of 37 000 based on the nozzle diameter D and the mean velocity in the nozzle. These scales are also used for later calculation of the Strouhal number. The mean velocity in figure 14(a) shows that the PIV domain is moved off the jet centre towards the lower half of the jet. Data points closer than $x/D = 2$ were distorted due to laser light reflections. The spectral content averaged over the PIV domain (figure 14b) shows a narrow dominant peak and at least five higher harmonics. The narrow peaks indicate a stable operation at the fundamental frequency, while the additional peaks suggest more complex dynamics. The key

challenge of this data set is to accurately reconstruct the sweeping jet dynamics from a truncated measurement domain.

Figure 15 shows the results from the SPOD for filter lengths $N_f = 0, 30$ and 2000. As in the foregoing examples, these filter settings span the range between both limiting cases (from the POD to the DFT). The spectrum attained with the POD (figure 15a) reveals distinct modes at the fundamental frequency of the oscillator (labelled 1) and at higher harmonics. The mode at the third harmonic frequency (labelled 2) shows a surprisingly high harmonic correlation. The PSD of the mode coefficients reveals that each mode is not limited to a single frequency. The additional peaks in the PSD are partly attributed to the fact that only part of the jet is measured. During one oscillation period, the jet leaves and enters the measurement domain, creating sharp transitions in the time domain and thus a series of higher harmonics in the frequency domain. Due to the purely statistical POD approach, these higher harmonics appear in every mode coefficient, which contradicts the idea of a proper modal decomposition. The mode that seems to represent the fifth harmonic in the SPOD spectrum (labelled 3) shows no distinct peak at all in the PSD of the coefficient. Thus, the POD of this data set does not provide a proper separation of the fundamental and higher harmonic contributions.

If the SPOD is applied instead, the fundamental and harmonic modes line up perfectly (figure 15b). Now, the harmonics are separated clearly up to seventh harmonics. The spectral content and spatial shape are further examined for the fundamental, third and seventh harmonics. The PSDs of the mode coefficients reveal narrow spectral bands. The corresponding mode shapes show an appropriate spatial symmetry, although the PIV domain is cropped shortly above the symmetry line. It is worth mentioning that the broad peak in the PSD of the seventh harmonic indicates considerable frequency jitter, while the mode shape remains remarkably smooth and symmetric.

The results obtained with the DFT are presented in figure 15(c). The peaks in the SPOD spectrum clearly indicate the fundamental and the first five higher harmonics. Their spatial shapes agree well with the SPOD modes, which is not surprising as these modes have narrow spectral bands. Note, however, that each peak is split into several DFT modes, which indicates slight frequency variation. This becomes crucial for the higher harmonics, where the frequency jitter is significant and the mode energy is low. For the seventh harmonic, the DFT fails to reproduce the structure seen in the SPOD (figure 15b). The frequency variations detected by the SPOD are simply too high and the mono-frequent energy content too low. This emphasizes the superior noise rejection of the SPOD.

In this example, SPOD is superior to POD and DFT. The energy-ranked POD modes primarily suffer from the incompleteness of the data set. This is of immense importance as the relevant domain size is rarely known prior to a set of experiments or POD analysis. The frequency-sharp DFT is insensitive to the domain size, but it fails to reconstruct weak modes with substantial noise and frequency jitter. The soft frequency constraint of the SPOD filter operation combines the advantages of both methods and generates a clear mode space. The SPOD generates modes with distinct frequencies and mode shapes for modes even weaker than the overall noise level.

For the fluidic oscillator, the DFT modes are nearly as accurate as those derived with the SPOD. The advantages of the SPOD are more obvious for less mono-frequent flow dynamics (see the previous examples). However, an additional advantage of the SPOD over the DFT is that it provides a reliable estimate of the oscillatory phase by accounting for the frequency jitter. Similar approaches, which also produce satisfactory

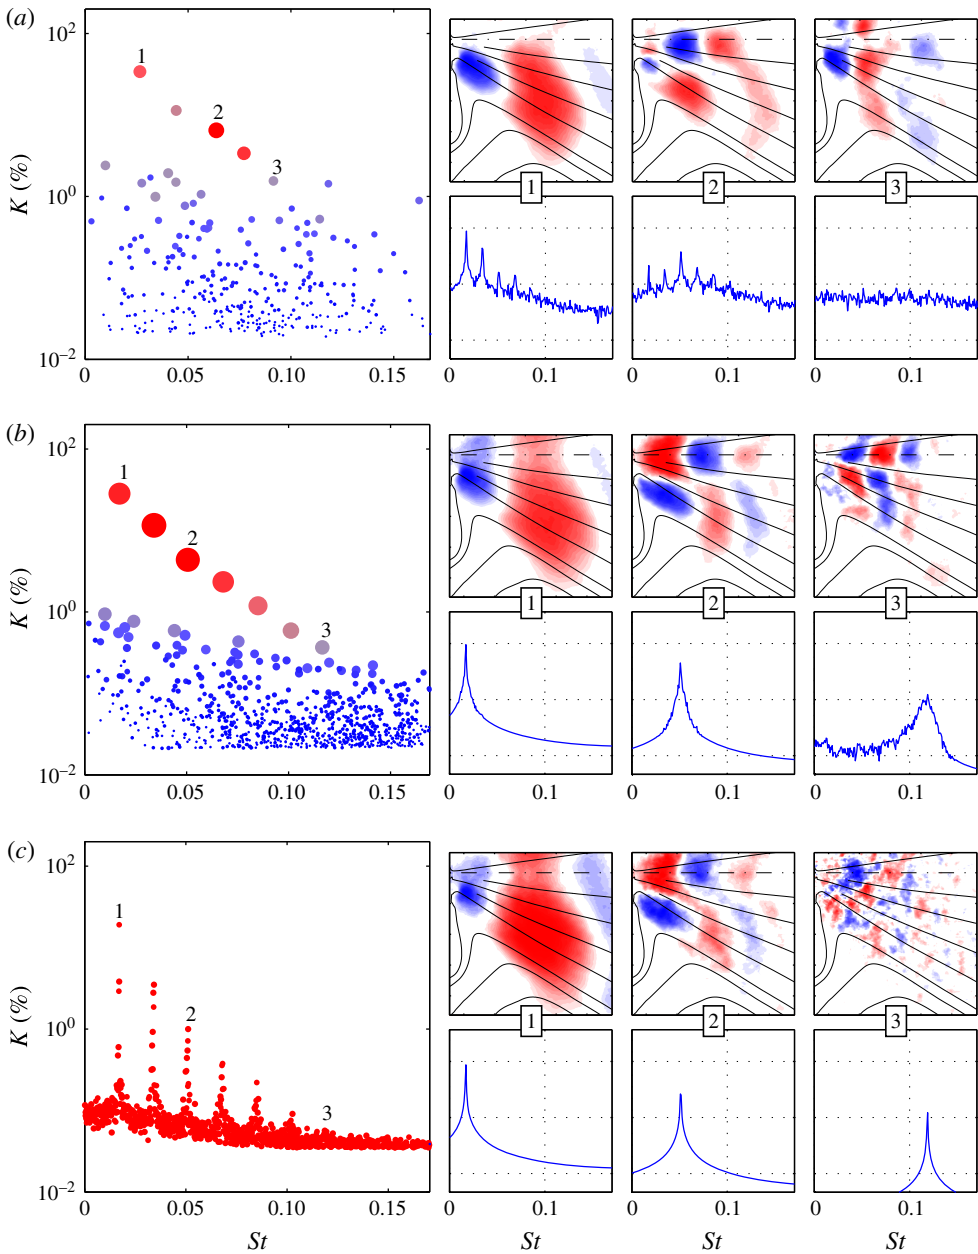


FIGURE 15. (Colour online) Fluidic oscillator: results from SPOD for different filter lengths (a) $N_f = 0$ (POD), (b) $N_f = 30$ (SPOD) and (c) $N_f = 2000$ (DFT). For each filter length the SPOD spectrum is displayed as a scatter plot (left), where a single dot indicates one mode pair (size and colour $C_{i,j}$ in (2.21)). For three selected pairs the spatial modes (upper row) and PSD of the temporal coefficient (lower row) are depicted. They are indicated by numbers in the SPOD spectrum, as well as between the small mode plots. The centreline is indicated by a dash-dotted line.




FIGURE 16. (Colour online) Phase portraits of first temporal coefficients from (a) POD, (b) SPOD and (c) both methods against each other.

results for the current case, are described by Ostermann *et al.* (2015a), but again, the scope of SPOD is beyond this particular application. Figure 16(a,b) shows the phase portraits (Lissajous figures) of the temporal coefficients of the two most energetic POD (a_i) and SPOD (b_i) modes, respectively. The trajectory of the POD coefficients does not follow a clear circle that would indicate the limit cycle. It rather follows a third of a circle and then collapses at one point. The coefficient of the SPOD modes follows a clear circle and the instantaneous phase and instantaneous frequency can consistently be deduced. A comparison of the first mode coefficient from both methods is shown in figure 16(c). It reveals that half of the period is cut out for the POD (a_1). This corresponds to the sweeping jet leaving the measurement domain, where the energy-based POD ‘sees’ no jet. The SPOD properly recovers the temporal dynamics over the entire oscillation period. Furthermore, note that the SPOD produces coefficients with smooth temporal dynamics, while the POD coefficients show rather erratic movements. This is particularly important for reduced-order modelling (Luchtenburg *et al.* 2009), phase averaging and extended POD (Boree 2003). Practically, most of the further processing is eased if there is less noise.

4. Conclusions

4.1. Properties and capabilities of the proposed method

SPOD is introduced as an extension of the POD for time-resolved data. This novel method involves a filter operation on the snapshot correlation matrix. The procedure is closely related to the classic snapshot POD with a negligible increase of algorithmic complexity and numerical costs. The SPOD filter allows for a continuous fading from the energetic optimality of POD to the spectral purity of DFT. It is conceptualized in a general form, with POD and DFT as the limiting cases. The concept of SPOD was developed through our experience with experimental data processing, and not from a constraint optimization problem. It arose from the great need for a method that applies to a wide range of turbulent flows with minimum user input. The SPOD is motivated by theoretical considerations, where it is interpreted as a short-time linearization of the flow dynamics.

The key feature of SPOD is the smoothing along all diagonals of the correlation matrix. This filter operation is shown to constrain the amplitudes and frequencies of the SPOD modes. By setting the filter width, one gains control over the spectral bandwidth of the single modes. When the filter is set to the maximum length, the modes are assumed to be strictly periodic and the SPOD converges to the DFT.

The application of SPOD to the flow field of a swirl-stabilized combustor, an airfoil with Gurney flap and a fluidic oscillator revealed different advantageous features of the SPOD in comparison to POD or DFT. For every single case there exist other suitable methods which may perform equally well as the SPOD, but none of them is as versatile as the proposed method.

The main advantages can be summarized as follows:

Separation of structures. The soft spectral constraint of the SPOD allows for a much better separation of individual fluid dynamic phenomena into single modes, whereas POD or DFT mix or spread them among several modes.

Noise rejection. SPOD is insensitive to noise and even recovers dynamics that is weaker than the overall noise level.

Missing data compensation. SPOD can eliminate the degradation of temporal dynamics of partially recorded phenomena, resulting from an improper choice of the measured region.

Plain dynamics. The mode coefficients are smooth in time and they feature adjustable variations of frequency and amplitude (set by the filter size).

The characteristics of the SPOD modes ease further processing such as the identification of linked modes, comparisons against other simultaneously acquired measurements (phase averaging or extended POD) and the identification of reduced-order models. The SPOD may also provide a better basis for modal representation of snapshots as input for a DMD, as pursued in § 2.5.

4.2. Concluding remarks

The SPOD has proven to be a reliable method for extracting distinct phenomena from turbulent flows. It was not derived from purely mathematical considerations, but rather evolved from practical data processing. Nonetheless, the method has well-defined upper and lower bounds and generates modes that can be easily interpreted. As shown in the considered examples, SPOD is a simple way to extract coherent structures from turbulent flows, where the POD or the DFT failed to provide valuable results. The SPOD constrains the spectral content, but retains the modal sparsity of the POD.

There are certainly plenty of other cases where this new method will ease the identification of hidden coherent structures. Its true benefit lies in the fact that only one assumption is made about the investigated flow dynamics, which is the filter time scale. This can also be understood as an inertia imposed on the mode dynamics, limiting the rate of change of the frequency and amplitude. The choice of this time scale can be assessed from the flow's dominant frequency or convective time scale, as shown in this article. The authors hope that the SPOD will give access to new fluid dynamical phenomena and enrich currently available methods.

Acknowledgements

The authors kindly acknowledge the stimulating discussions with L. Rukes and the generous provision of experimental data from A. Bach and F. Ostermann. The funding from the German Research Foundation (DFG Project PA 920/30-1) is also gratefully acknowledged.

Appendix A. The spatial correlation version of SPOD

The original POD can be calculated from either a spatial or a temporal correlation, which allows a computationally efficient calculation by restricting the size of the

problem to the number of snapshots or the number of grid points. Similarly, the SPOD has a spatial correlation counterpart, which is computationally more efficient if the number of snapshots is much larger than the number of grid points. This approach is slightly more complex and less intuitive than the snapshot version. Nevertheless, it is very valuable if long time series of few sensors are supposed to be decomposed into proper modes to perform an extended POD or to derive the phase of an oscillatory mode from the measurements. Assume a simultaneous multi-point pressure measurement that shall be decomposed with SPOD. This series is decomposed as

$$P(x_i, t) = \bar{p}(x_i) + \sum_{s=1}^N b_s(t) \Psi_s(x_i), \quad (\text{A } 1)$$

where the number of measured positions M is much smaller than the number of samples N . The number of samples may easily reach a million or more, which complicates the solution in terms of memory requirements for the composition of the temporal correlation matrix and in terms of computational time for the solution of the eigenvalue problem. Therefore the temporal correlation described in § 2.3 is not feasible in this case. Instead, the spatial correlation should be employed, as outlined in this section. The multi-time-shift correlation tensor for the spatial SPOD reads

$$\mathcal{S}_{i,j,k,l} = \frac{\sqrt{g_k g_l}}{MN \Delta t} \int p(x_i, t - k \Delta t) p(x_j, t - l \Delta t) dt \quad (i, j = 1 \dots M; k, l = -N_f \dots N_f), \quad (\text{A } 2)$$

where $p = P - \bar{p}$ is the fluctuating part of the pressure and g_k are the filter coefficients. For numerical implementation this is reshaped to a matrix such that $\mathcal{S}_{i,j,k,l} = \tilde{\mathcal{S}}_{(i+kM), (j+lM)}$, but for the theoretical description the tensor notation is retained. The correlation tensor is decomposed into eigenvalues and eigenvectors, such that

$$\sum_{l=-N_f}^{N_f} \sum_{j=1}^M \mathcal{S}_{i,j,k,l} \tilde{\Psi}_s(x_j, \tau_l) = \mu_s \tilde{\Psi}_s(x_i, \tau_k); \quad \mu_1 \geq \mu_2 \geq \dots \geq \mu_N \geq 0, \quad (\text{A } 3)$$

where $\tau_k = k \Delta t$. The eigenvector $\tilde{\Psi}_s$ constitutes a discrete convolution filter, which is applied to the time series to obtain the mode coefficients

$$b_s(t) = \sum_{k=-N_f}^{N_f} \sum_{i=1}^M \sqrt{\frac{g_k}{M}} \tilde{\Psi}_s(x_i, \tau_k) p(x_i, t - \tau_k). \quad (\text{A } 4)$$

The spatial mode is the central part ($\tau_k = 0$) of the convolution filter $\Psi_s(x_i) = \tilde{\Psi}_s(x_i, 0)$. The entire eigenvectors $\tilde{\Psi}_s$ can be understood as a data-driven filter bank, which allows for decomposition of time series into modal contributions. It might be applied to a single sensor, where each mode represents a certain spectral band of the signal. The principal approach is comparable to the empirical mode decomposition (Huang *et al.* 1998), but the SPOD can also handle multiple sensors. In the case of multiple sensors, it gives excellent results when the phase of a dominant oscillation has to be reconstructed from pressure measurements. The approach outlined in this section is similar to the multi-time-delay POD phase estimation pursued by Hosseini *et al.* (2015).

In contrast to the snapshot version, the computational cost of the spatial version of SPOD scales with the filter size. It is only more efficient than the snapshot approach if $M(2N_f + 1) < N$.

Appendix B. Properties of the SPOD modes

In §2.3 it was briefly mentioned that the spatial SPOD modes are no longer orthogonal, which is only part of the true picture. If the spatial mode $\tilde{\Psi}$ in (A3) together with all of the temporally shifted instances are considered, they are orthonormal:

$$\frac{1}{M} \sum_{l=-N_f}^{N_f} \sum_{k=1}^M \tilde{\Psi}_i(x_k, \tau_l) \tilde{\Psi}_j(x_k, \tau_l) = \delta_{i,j}. \quad (\text{B } 1)$$

The snapshot-based calculation introduced in §2.3, however, is restricted to the zero-time-delay ($\tau = 0$) part of the spatial mode. Furthermore, the decomposition of the data into modal contributions is only feasible for time-independent spatial modes. This limitation to one part of the spatial mode $\Psi_s(x_i) = \tilde{\Psi}_s(x_i, \tau = 0)$ introduces some imperfections. The selected modes for the decomposition are neither normal nor orthogonal,

$$\frac{1}{M} \sum_{k=1}^M \Psi_i(x_k) \Psi_j(x_k) \neq \delta_{i,j}. \quad (\text{B } 2)$$

The loss of normality is a fact, whereas the norm of the spatial modes gives further insights into the data set. The norm

$$\zeta_i = \sqrt{\frac{1}{M} \sum_{k=1}^M \Psi_i^2(x_k)} \quad (\text{B } 3)$$

indicates how well a single mode is represented by the investigated data set.

With the application of the filter (2.9), an idealized correlation matrix is constructed that delivers modes which are more or less captured by the initial data set. This fact is reflected by the deviation of the mode norm ζ_i from one. Consider for example the measurement of the sweeping jet. There, the fundamental mode is only partially captured, as shown in figure 16. With SPOD, the missing data are completed and a SPOD mode pair with equal energy levels μ_i is obtained. However, for the construction of the spatial modes the coefficients are projected onto the original data (2.12). There, the imperfect representation of one of the two modes re-enters the processing. For the sweeping jet's leading mode pair the norm ζ_i of one mode is clearly below the other, but they approximately add up to one. Therefore, the eigenvalues μ_i describe the idealized energy content of the single modes, and the norm of the spatial mode ζ_i corrects the deficits in comparison to the actual data set. The limiting SPOD cases (POD and DFT) do not show such deficits. The POD modes are already normalized to $\zeta_i = 1$, and for the DFT the modes pair perfectly, while the norm of these pairs (i, j) exactly add up to one ($\zeta_i + \zeta_j = 1$).

Appendix C. SPOD of linear dynamics

In this article a relation between the diagonal similarity of the correlation matrix and the dynamics of the underlying system is presumed but never discussed in detail. This section provides a better comprehension of the stated dynamics, which causes the diagonal similarity of the correlation matrix. Therefore, the term 'similar dynamics'

is replaced by ‘linear dynamics’, whereby the temporal evolution of the investigated flow is locally governed by a linear, time-invariant model. The term locally refers to a short and finite temporal extension, and it is closely related to the filter size N_f . The temporal evolution of this system is given by

$$\frac{\partial \mathbf{u}(t)}{\partial t} = \mathbf{L}\mathbf{u}(t), \quad (\text{C1})$$

where the matrix \mathbf{L} is the system matrix describing the entire dynamics and the spatial points of the velocity field are organized as rows in $\mathbf{u}(t) = [u(\mathbf{x}_1, t), \dots, u(\mathbf{x}_M, t)]^T$. This approach may also be understood as continuous-time DMD. The linearized dynamics allows for the calculation of the velocity field at any time step from a reference snapshot $\mathbf{u}_0 = \mathbf{u}(t=0)$ by the matrix exponential

$$\mathbf{u}(t) = e^{\mathbf{L}t} \mathbf{u}_0. \quad (\text{C2})$$

To allow further simplifications, we require the matrix \mathbf{L} to be normal. This allows the decomposition $\mathbf{L} = \mathbf{U}\mathbf{D}\mathbf{U}^*$, where \mathbf{U} is a unitary matrix (eigenvectors of \mathbf{L}), \mathbf{D} is a diagonal matrix (eigenvalues of \mathbf{L}) and $*$ means the conjugate transpose of a matrix (adjoint). Hereafter, (C2) becomes

$$\mathbf{u}(t) = e^{\mathbf{U}\mathbf{D}\mathbf{U}^*t} \mathbf{u}_0 = \mathbf{U}e^{\mathbf{D}t} \mathbf{U}^* \mathbf{u}_0. \quad (\text{C3})$$

The diagonal elements of the matrix \mathbf{D} contains the complex eigenvalues d_k of the system. Each of these eigenvalues contains the amplification rate σ and the frequency ω of the related mode, $d_k = \sigma_k + i\omega_k$. The diagonal matrix can be decomposed into the amplification rate $\mathbf{\Sigma}$ and frequency $\mathbf{\Omega}$, thus $\mathbf{D} = \mathbf{\Sigma} + i\mathbf{\Omega}$.

With the formulation in (C3), the inner product which forms the elements of the correlation matrix is simplified to

$$\begin{aligned} \langle u(\mathbf{x}, t_1), u(\mathbf{x}, t_2) \rangle &= \langle \mathbf{U}e^{\mathbf{D}t_1} \mathbf{U}^* \mathbf{u}_0, \mathbf{U}e^{\mathbf{D}t_2} \mathbf{U}^* \mathbf{u}_0 \rangle \\ &= \langle \mathbf{U}^* \mathbf{u}_0, e^{\mathbf{\Sigma}(t_1+t_2)+i\mathbf{\Omega}(t_2-t_1)} \mathbf{U}^* \mathbf{u}_0 \rangle. \end{aligned} \quad (\text{C4})$$

An inspection of the exponent in (C4) reveals that the inner product only depends on the sum and difference of time steps. According to this equation, the velocity fields are projected onto the subspace spanned by the linear operator $\mathbf{U}^* \mathbf{u}_0$ and changes of the inner product are only governed by the eigenvalues of the system.

Within the context of the correlation matrix, we use the abbreviated nomenclature for the snapshots $\mathbf{u}_i = \mathbf{u}(i\Delta t)$ and for the projected velocity $\tilde{\mathbf{u}} = \mathbf{U}^* \mathbf{u}_0$. The correlation matrix constructed around the neighbourhood of snapshots \mathbf{u}_0 yields

$$\begin{aligned} \mathbf{S}_{sub} &= \begin{pmatrix} \langle \mathbf{u}_{-1}, \mathbf{u}_{-1} \rangle & \langle \mathbf{u}_0, \mathbf{u}_{-1} \rangle & \langle \mathbf{u}_1, \mathbf{u}_{-1} \rangle \\ \langle \mathbf{u}_{-1}, \mathbf{u}_0 \rangle & \langle \mathbf{u}_0, \mathbf{u}_0 \rangle & \langle \mathbf{u}_1, \mathbf{u}_0 \rangle \\ \langle \mathbf{u}_{-1}, \mathbf{u}_1 \rangle & \langle \mathbf{u}_0, \mathbf{u}_1 \rangle & \langle \mathbf{u}_1, \mathbf{u}_1 \rangle \end{pmatrix} \\ &= \begin{pmatrix} \langle \tilde{\mathbf{u}}, e^{-2\mathbf{\Sigma}\Delta t} \tilde{\mathbf{u}} \rangle & \langle \tilde{\mathbf{u}}, e^{-\mathbf{\Sigma}\Delta t - i\mathbf{\Omega}\Delta t} \tilde{\mathbf{u}} \rangle & \langle \tilde{\mathbf{u}}, e^{-2i\mathbf{\Omega}\Delta t} \tilde{\mathbf{u}} \rangle \\ \langle \tilde{\mathbf{u}}, e^{-\mathbf{\Sigma}\Delta t + i\mathbf{\Omega}\Delta t} \tilde{\mathbf{u}} \rangle & \langle \tilde{\mathbf{u}}, \tilde{\mathbf{u}} \rangle & \langle \tilde{\mathbf{u}}, e^{\mathbf{\Sigma}\Delta t - i\mathbf{\Omega}\Delta t} \tilde{\mathbf{u}} \rangle \\ \langle \tilde{\mathbf{u}}, e^{2i\mathbf{\Omega}\Delta t} \tilde{\mathbf{u}} \rangle & \langle \tilde{\mathbf{u}}, e^{\mathbf{\Sigma}\Delta t + i\mathbf{\Omega}\Delta t} \tilde{\mathbf{u}} \rangle & \langle \tilde{\mathbf{u}}, e^{2\mathbf{\Sigma}\Delta t} \tilde{\mathbf{u}} \rangle \end{pmatrix}. \end{aligned} \quad (\text{C5})$$

The actual properties of this complex expression are highlighted by showing just the factors for $\mathbf{\Sigma}$ and $i\mathbf{\Omega}$ in the exponents, given as

$$\mathbf{\Sigma}: \begin{bmatrix} -2 & -1 & 0 \\ -1 & 0 & 1 \\ 0 & 1 & 2 \end{bmatrix} \Delta t; \quad i\mathbf{\Omega}: \begin{bmatrix} 0 & -1 & -2 \\ 1 & 0 & -1 \\ 2 & 1 & 0 \end{bmatrix} \Delta t. \quad (\text{C6a,b})$$

It is perfectly visible that the changes of the correlation matrix along the diagonals are caused by amplification of modes, while changes along the anti-diagonals are caused by the modes' frequency. Therefore, the changes of the correlation matrix along the diagonal are related to the modes' amplitude and the cross-wise changes are caused by the modes' phase. The amplification rate can be assumed to be much smaller than the frequency, which allows a change in phase to be distinguished from a change in amplitude. Therefore, the correlation matrix generated by a linear system exhibits similar diagonals. Hence, a flow governed by linear dynamics shows the observed diagonal similarity in the correlation matrix.

If the amplification rate $\sigma = 0$ and the frequency $\omega = \text{constant}$, a periodic oscillation with fixed amplitude is obtained. This constitutes the limiting DFT case, where the correlation matrix is perfectly constant along the diagonals. Thus, constant coefficients of the underlying system result in a perfectly diagonal matrix. In order to represent the variations in frequency and amplitude observed in the SPOD modes, the system parameters σ and ω may change in time. These changes, however, will break the diagonal similarity of the correlation matrix. Therefore, the nonlinearity and non-normality of the Navier–Stokes equations contribute to parameter variations of the linear system. The flow is assumed to behave like a linear normal system within the time scale of the filter and all nonlinear and non-normal dynamics are represented by variations in σ and ω . This kind of approach is also pursued in the generalized mean field model of Luchtenburg *et al.* (2009), where the mode interaction via the mean flow is represented by an interaction of linear oscillators through nonlinear coupling of model parameters. Hence, the SPOD filter might be understood as a smoothing (linearization) of the dynamics (σ ω) of the underlying system.

In this system dynamic framework, a relation between SPOD and DMD can be drawn. For the DMD the entire analysed sequence is considered from the perspective of a linearized system. In the sense that a low-pass filter may also be described as a moving average, the SPOD may be understood as a moving DMD. The filter with N_f is the corresponding linearization horizon. Therefore, the eigenvalues of a short-term DMD may relate to the instantaneous dynamics (σ ω) of the SPOD modes.

REFERENCES

- ADRIAN, R. J. & MOIN, P. 1988 Stochastic estimation of organized turbulent structure: homogeneous shear flow. *J. Fluid Mech.* **190**, 531–559.
- BACH, A. B., BERG, R., PECHLIVANOGLU, G., NAYERI, C. & PASCHEREIT, C. O. 2015a Experimental investigation of the aerodynamic lift response of an active finite gurney flap. In *53rd AIAA Aerospace Sciences Meeting, AIAA SciTech*. American Institute of Aeronautics and Astronautics.
- BACH, A. B., LENNIE, M., PECHLIVANOGLU, G., NAYERI, C. N. & PASCHEREIT, C. O. 2014 Finite micro-tab system for load control on a wind turbine. *J. Phys.: Conf. Ser.* **524**, 012082.
- BACH, A. B., PECHLIVANOGLU, G., NAYERI, C. & PASCHEREIT, C. O. 2015b Wake vortex field of an airfoil equipped with an active finite Gurney flap. In *53rd AIAA Aerospace Sciences Meeting, AIAA SciTech*. American Institute of Aeronautics and Astronautics.
- BERKOOZ, G., HOLMES, P. & LUMLEY, J. L. 1993 The proper orthogonal decomposition in the analysis of turbulent flows. *Annu. Rev. Fluid Mech.* **25**, 539–575.
- BOREE, J. 2003 Extended proper orthogonal decomposition: a tool to analyse correlated events in turbulent flows. *Exp. Fluids* **35** (2), 188–192.
- BOURGOIS, J. A., NOACK, B. R. & MARTINUZZI, R. J. 2013 Generalized phase average with applications to sensor-based flow estimation of the wall-mounted square cylinder wake. *J. Fluid Mech.* **736**, 316–350.

- DURGESH, V. & NAUGHTON, J. W. 2010 Multi-time-delay LSE-POD complementary approach applied to unsteady high-Reynolds-number near wake flow. *Exp. Fluids* **49** (3), 571–583.
- GRAY, R. M. 2005 Toeplitz and circulant matrices: a review. *Foundations Trends Commun. Inform. Theory* **2** (3), 155–239.
- HOLMES, P., LUMLEY, J. L., BERKOOZ, G. & ROWLEY, C. W. 2012 *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, 2nd edn. Cambridge University Press.
- HOSSEINI, Z., MARTINUZZI, R. J. & NOACK, B. R. 2015 Sensor-based estimation of the velocity in the wake of a low-aspect-ratio pyramid. *Exp. Fluids* **56** (1), 1–16.
- HUANG, H. T., FIEDLER, H. E. & WANG, J. J. 1993 Limitation and improvement of PIV: part II: particle image distortion, a novel technique. *Exp. Fluids* **15** (4–5), 263–273.
- HUANG, N. E., SHEN, Z., LONG, S. R., WU, M. C., SHIH, H. H., ZHENG, Q., YEN, N.-C., TUNG, C. C. & LIU, H. H. 1998 The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proc. R. Soc. Lond. A* **454** (1971), 903–995.
- LASAGNA, D., ORAZI, M. & IUSO, G. 2013 Multi-time delay, multi-point linear stochastic estimation of a cavity shear layer velocity from wall-pressure measurements. *Phys. Fluids* **25** (1), 017101.
- LUCHTENBURG, D. M., GÜNTHER, B., NOACK, B. R., KING, R. & TADMOR, G. 2009 A generalized mean-field model of the natural and high-frequency actuated flow around a high-lift configuration. *J. Fluid Mech.* **623**, 283–316.
- LUMLEY, J. L. 1970 *Stochastic Tools in Turbulence*, Applied Mathematics and Mechanics, vol. 12. Academic Press.
- OBERLEITHNER, K., RUKES, L. & SORIA, J. 2014 Mean flow stability analysis of oscillating jet experiments. *J. Fluid Mech.* **757**, 1–32.
- OBERLEITHNER, K., SIEBER, M., NAYERI, C. N., PASCHEREIT, C. O., PETZ, C., HEGE, H.-C., NOACK, B. R. & WYGNANSKI, I. 2011 Three-dimensional coherent structures in a swirling jet undergoing vortex breakdown: stability analysis and empirical mode construction. *J. Fluid Mech.* **679**, 383–414.
- OBERLEITHNER, K., STÖHR, M., SEONG, H. I., ARNDT, C. M. & STEINBERG, A. M. 2015 Formation and flame-induced suppression of the precessing vortex core in a swirl combustor: experiments and linear stability analysis. *Combust. Flame* **162** (8), 3100–3114.
- OSTERMANN, F., WOSZIDLO, R., GAERTLEIN, S., NAYERI, C. & PASCHEREIT, C. O. 2015a Phase-averaging methods for the natural flowfield of a fluidic oscillator. *AIAA J.* **53** (8), 2359–2368.
- OSTERMANN, F., WOSZIDLO, R., NAYERI, C. & PASCHEREIT, C. O. 2015b Experimental comparison between the flow field of two common fluidic oscillator designs. In *53rd AIAA Aerospace Sciences Meeting, AIAA SciTech*. American Institute of Aeronautics and Astronautics.
- RAFFEL, M., KOMPENHANS, J., WERELEY, S. T. & WILLERT, C. E. 2007 *Particle Image Velocimetry: A Practical Guide*, 2nd edn. Springer.
- RAIOLA, M., DISCETTI, S. & IANIRO, A. 2015 On PIV random error minimization with optimal POD-based low-order reconstruction. *Exp. Fluids* **56** (4), 1–15.
- REICHEL, T. G., TERHAAR, S. & PASCHEREIT, O. 2015 Increasing flashback resistance in lean premixed swirl-stabilized hydrogen combustion by axial air injection. *Trans. ASME J. Engng Gas Turbines Power* **137** (7), 071503.
- ROWLEY, C. W., MEZIC, I., BAGHERI, S., SCHLATTER, P. & HENNINGSON, D. S. 2009 Spectral analysis of nonlinear flows. *J. Fluid Mech.* **641**, 115–127.
- SCHLEGEL, M., NOACK, B. R., JORDAN, P., DILLMANN, A., GRÖSCHEL, E., SCHRÖDER, W., WEI, M., FREUND, J. B., LEHMANN, O. & TADMOR, G. 2012 On least-order flow representations for aerodynamics and aeroacoustics. *J. Fluid Mech.* **697**, 367–398.
- SCHMID, P. J. 2010 Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* **656**, 5–28.
- SIROVICH, L. 1987 Turbulence and the dynamics of coherent structures: part I: coherent structures. *Q. Appl. Maths* **45** (3), 561–571.
- SORIA, J. 1996 An investigation of the near wake of a circular cylinder using a video-based digital cross-correlation particle image velocimetry technique. *Exp. Therm. Fluid Sci.* **12** (2), 221–233.
- TERHAAR, S. 2015 Identification and modeling of coherent structures in swirl stabilized combustors at dry and steam diluted conditions. PhD thesis, Technische Universität Berlin.

- TERHAAR, S., OBERLEITHNER, K. & PASCHEREIT, C. O. 2015 Key parameters governing the precessing vortex core in reacting flows: an experimental and analytical study. *Proc. Combust. Inst.* **35** (3), 3347–3354.
- TROOLIN, D. R., LONGMIRE, E. K. & LAI, W. T. 2006 Time resolved PIV analysis of flow over a NACA 0015 airfoil with Gurney flap. *Exp. Fluids* **41** (2), 241–254.
- TU, J. H., ROWLEY, C. W., LUCHTENBURG, D. M., BRUNTON, S. L. & KUTZ, J. N. 2014 On dynamic mode decomposition: theory and applications. *J. Comput. Dyn.* **1** (2), 391–421.
- WILLERT, C. E. & GHARIB, M. 1991 Digital particle image velocimetry. *Exp. Fluids* **10** (4), 181–193.
- WILLIAMS, M. O., KEVREKIDIS, I. G. & ROWLEY, C. W. 2015 A data-driven approximation of the Koopman operator: extending dynamic mode decomposition. *J. Nonlinear Sci.* **25** (6), 1307–1346.
- WISE, J. 1955 The autocorrelation function and the spectral density function. *Biometrika* **42** (1–2), 151–159.
- WOSZIDLO, R., OSTERMANN, F., NAYERI, C. N. & PASCHEREIT, C. O. 2015 The time-resolved natural flow field of a fluidic oscillator. *Exp. Fluids* **56** (6), 1–12.
- WOSZIDLO, R., STUMPER, T., NAYERI, C. & PASCHEREIT, C. O. 2014 Experimental study on bluff body drag reduction with fluidic oscillators. In *52rd AIAA Aerospace Sciences Meeting, AIAA SciTech*. American Institute of Aeronautics and Astronautics.