Feedback Flow Control in Experiment and Simulation Using Global Neural Network Based Models

Stefan Siegel

Department for Aeronautics, 2354 Fairchild Dr. Ste 6H27, United States Air Force Academy, CO 80840, USA

Abstract For feedback control of complex spatio-temporally evolving flow fields, it imperative to use a global flow model for both flow state estimation, as well as controller development. It is important that this model correctly presents not just the natural, unforced flow state, but also the interaction of actuators with the flow for both open and closed loop situations. In order to achieve this, a novel extension of POD is introduced in this chapter, which we refer to as Double POD (DPOD). This decomposition allows the construction of a POD basis that is valid for a variety of flow conditions, which may be distinguished by changes in actuation, Reynolds number or other parameters. While traditionally the velocity field has been used as input for POD, other variables, for example the pressure or density field, may be used as well. The mode amplitudes of the DPOD spatial modes are then used as input for a system identification process, the nonlinear ANN-ARX method is employed here. The result is a dynamic model that represents both the unforced, open loop forced and closed loop flow fields with good accuracy.

1 Introduction

One of the main purposes of flow control is the improvement of aerodynamic characteristics of air vehicles and munitions enabling augmented mission performance. An important area of flow control research involves the phenomenon of vortex shedding in the wake behind bluff bodies where the flow separates from the bluff body's surface. Shedding of counter-rotating vortices is observed in the wake of a two-dimensional cylinder above a critical Reynolds number of $Re_{crit} \approx 47$, non-dimensionalized with respect to freestream velocity and cylinder diameter. This phenomenon is often referred to as the von Kármán vortex street, shown schematically in Figure 1. The vortex shedding leads to a sharp rise in drag, noise and fluid-induced

vibration. The ability to control the wake of a bluff body could be used to reduce drag, increase mixing and heat transfer, and enhance combustion. The benchmark problem of a circular cylinder wake is studied at a Reynolds number of $Re_D \approx 100$, which corresponds to the range in which the wake is laminar and two-dimensional. When active open-loop forcing of the wake is employed, the vortices in the wake can be "locked" to the forcing signal. This also strengthens the vortices and consequently increases the drag. As opposed to the open-loop approach, the goal is to control the unsteady wake using a feedback controller in order to reduce or suppress the vortex shedding. The feedback control law is designed using a reduced order model of the unsteady flow. A common method used to substantially reduce the order of the model is Proper Orthogonal Decomposition (POD). This method is an optimal approach in that it will capture the largest amount of the flow energy for fixed amount of modes of any decomposition of the flow. However, a number of extensions to the POD decomposition are necessary in order to derive a numerical model suitable for feedback controller development.

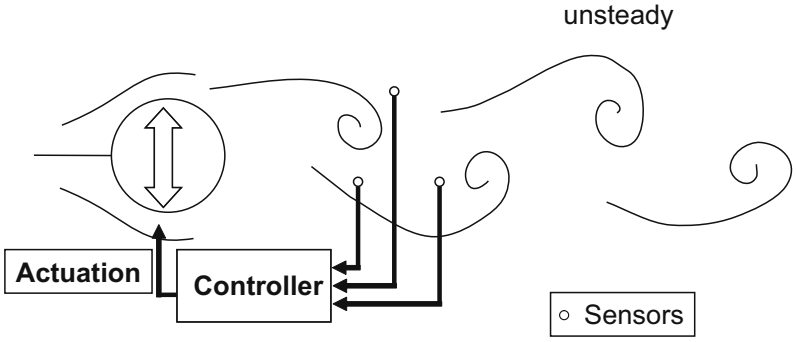

Figure 1. Feedback Control Setup

Low dimensional model development based on POD decomposition is a three step process, as shown in Figure 2. In the first step, data on the flow field to be modelled is gathered using either numerical or experimental methods. Selection of suitable data sets is a crucial step in model building and will be discussed in detail below. In the second step, spatial modes and their mode amplitudes are derived from the flow field data. The most often used approach is the method of snapshots developed by Sirovich (1987). Various methods of deciding how to precondition or cluster the data have been suggested in literature and are discussed below.

In conjunction with the spatial POD modes, the associated mode amplitudes may be calculated using an inner product or least square fit approach.

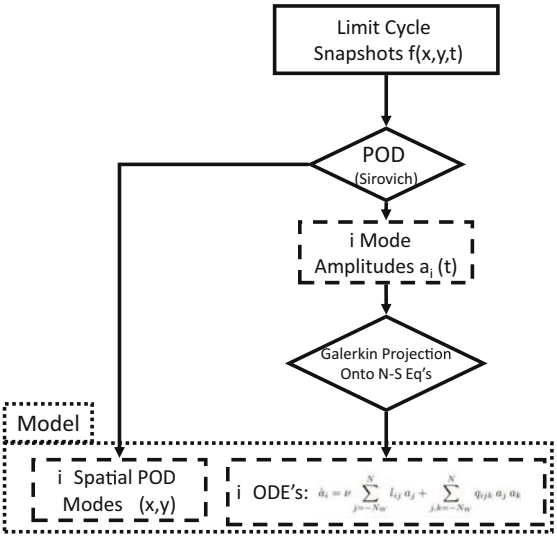

Figure 2. Flow Chart of traditional POD modeling process (see Chapter by Noack et al.)

To arrive at a low dimensional model, the mode set is truncated, usually based on an energy criterion (the eigenvalues in the Karhunen-Loève system represent twice the modal kinetic energy if POD is applied to the velocity field). The third and final step is the development of a model for the remaining mode amplitudes. This is commonly achieved by a Galerkin projection on the Navier-Stokes equations, which yields a system of equations that describes the evolution of the mode amplitudes over time. This set of equations can then be used to develop feedback control algorithms in a systematic fashion, or to test the performance of control algorithms against this model. All three steps of model development described above involve assumptions and potential problems, with many different solution proposed in literature, which will be discussed in the following.

Figure 3 gives an overview over the newly developed DPOD-ANN-ARX modeling approach discussed in this Chapter. The overview flow diagram illustrates the modifications made to the traditional approach shown in

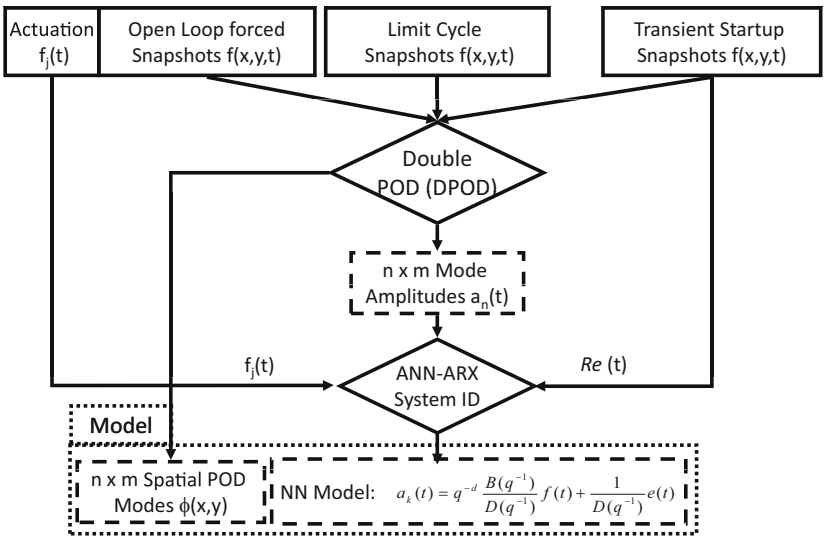

Figure 3. Flow Chart of DPOD-ANN-ARX modeling process

Figure 2. At the input data level, transient data from both a change in Reynolds number as well as open loop forced flow states is included. Then, a modified POD procedure which we refer to as Double Proper Orthogonal Decomposition (DPOD) is employed. The main benefit of this method lies in the ability to derive spatial POD modes that cover a large range of flow conditions with small estimation errors, for details refer to the DPOD section in the following. Instead of the traditional Galerkin projection, a nonlinear ARX (Auto Regressive eXternal input) system identification method based on Artificial Neural Networks (ANN) is used in order to develop a dynamic model of the flow. For the benchmark problem of the circular cylinder wake, it is shown that this approach yields a numerically stable model that is not just valid for the data used in its derivation, but also for a range of different Reynolds numbers, different open loop forcing conditions and, most importantly, feedback controlled flow states (Siegel et al., 2008). Further details on the DPOD-ANN-ARX modeling approach are included in the following subsections of this chapter, while the modeling results as well as feedback controlled flow simulations are discussed in the

later sections of this chaper.

2 Wake stabilization benchmark

The approach to model development discussed in this chapter can be used with either experimental or numerical data, since it is not dependent on e.g. solving the Navier Stokes equations inverse in time or similar processes that cannot be performed in an experiment. However, due to the abundance of data available from simulations, all results presented here have been obtained from DNS simulations. It is worth while noting that most CFD solvers, in particular commercially available codes, are not well suited to time resolved simulations. This is manifest either in excessively slow calculation times, or the inability to properly resolve the dynamic flow behavior due to artificial viscosity inherent in code, grid or both leading to nonphysical simulation results. Thus, it is paramount to check the performance of the candidate CFD solver not just for time averaged quantities like lift or drag, but more importantly for dynamic quantities like shedding frequencies (as done below) or even better, transient flow time scales. All numerical simulations reported here were conducted with Cobalt Solutions COBALT solver for direct numerical solution of the Navier-Stokes equations with second order accuracy in time and space. A structured two-dimensional grid with 63,700 nodes and 31,752 elements was used Figure 4. The grid extended from $\frac{x}{D} = -16.9$ to $\frac{x}{D} = 21.1$ in the x (streamwise) direction, and $\frac{|y|}{D} \leq 19.4$ cylinder diameters in the y (flow normal) direction. While this size of a domain is necessary to obtain accurate CFD results, the region of interest for feedback flow control is much smaller as indicated in Figure 4. While truncation in the spatial domain would pose major problems for many other model development approaches due to boundary condition issues, our approch does not involve boundary conditions as will be described later, and thus the truncation to a much smaller spatial domain is possible.

For validation of the computations of the unforced cylinder wake at $Re = 100$, the resulting value of the mean drag coefficient, c_d, was compared to experimental and computational investigations reported in the literature. Experimental data reported by Oertel (1990) point to c_d values between 1.26 and 1.4. Furthermore, Min and Choi (1999) report on several numerical studies that obtained drag coefficients between 1.34 and 1.35. The current simulations yield $c_d = 1.35$, which compares well with the reported literature. Another important benchmark parameter is the non-dimensional shedding frequency (Strouhal number, $St = \frac{fD}{U}$ for the unforced cylinder wake. Experimental results presented by Williamson (1996) point to values of $St \in [0.167, 0.168]$. The computations used in this effort result in

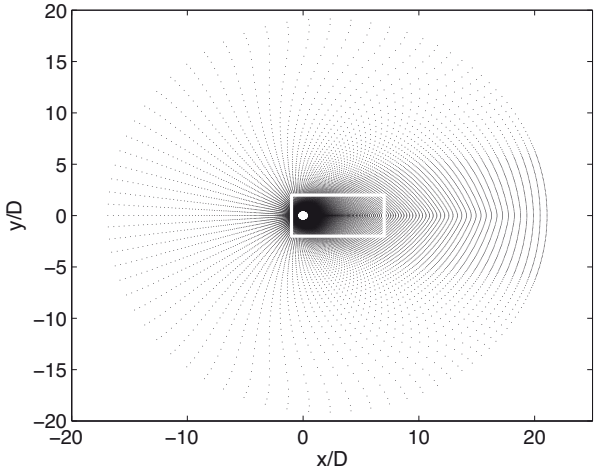

Figure 4. Circular Cylinder Grid and POD domain

$St = 0.163$, which also compares well with the reported literature. Beyond the unforced flow condition, cylinder displacement normal to the flow direction was used to investigate the interaction between this type of actuation and the flow field. A typical transient forcing signal is shown in Figure 5, where the cylinder displacement is sinusoidal in time with a frequency f and peak amplitude A. The flow response in Figure 5 shows a lock-in behavior which leads to a fixed phase relationship between forcing and the oscillating lift force after a few shedding cycles. At the beginning and end of the forcing, the flow exhibits transient behaviour. These dynamic adjustment periods are a manifestation of the underlying dynamics of the flow field, and as such contain important information for the development of a dynamic model of the flow behavior.

As the aim is to develop an effective numerical model of the fluid dynamic behavior of the flow when subject to various forcing inputs where lock-in can be achieved, a number of simulations to cover different frequency and forcing amplitudes were conducted Figure 6 shows the selected parameter combinations, along with the boundaries beyond which lock in cannot be achieved any more. Nine different data sets for the open loop forced cases, were obtained using forcing amplitudes of 10, 15, 20, 25 and 30% cylinder displacement. Some of the cases use a 5-10 % lower or higher frequency at 30% displacement, which is still within the lock-in region. Lock-in is defined

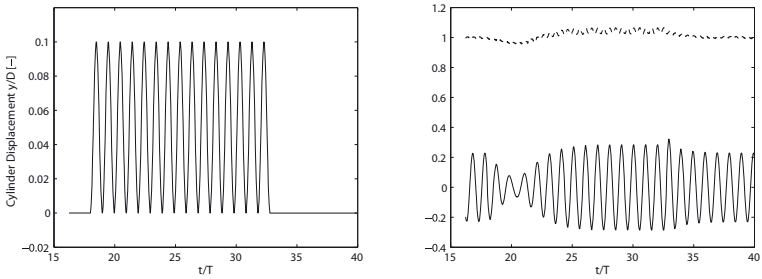

Figure 5. Left, Cylinder displacement. $f/f_0 = 1$; $A/D = 0.05$. Forcing activated at $t/T = 18$ and stopped at $t/T = 33$, after 15 full forcing cycles. Right, transient forced normalized lift coefficient c_l/c_{d0} (solid line) and drag coefficients c_d/c_{d0} (dashed line).

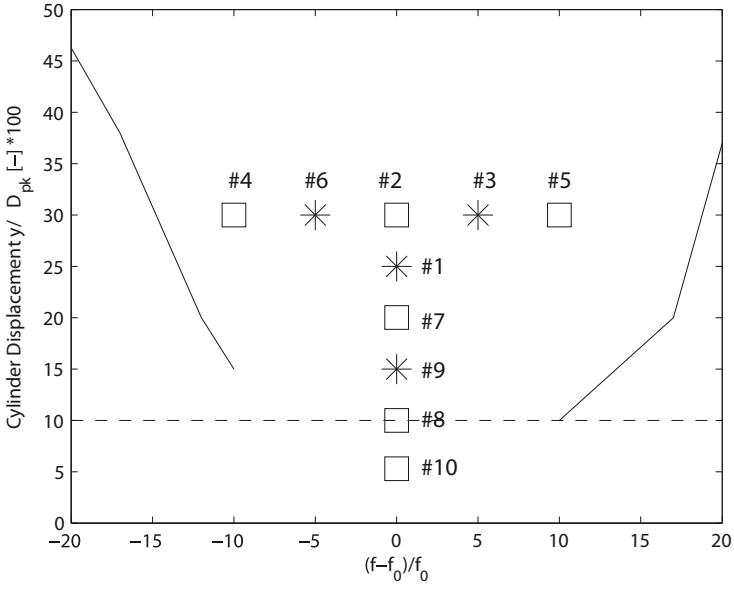

Figure 6. "Lock-In" Envelope with Design and Off-Design Cases

as the flow develping a fixed phase releationship between the actuation and the vortex shedding. The simulations indicated as squares were used as design cases for model development, while the cases indicated as stars, are utilized for model validation. The goal is to develop a model based on the design cases, that still yields accurate results for the validation cases.

3 Description of transient coherent structure via DPOD

The POD decomposition has been introduced in earlier chapters in this book, and is used it here employing the method of snapshots introduced by Sirovich (1987).

$$u(x, y, t) = \sum_{i=1}^{I} a_i(t)\phi_i(x, y). \tag{1}$$

However, in all data presented in this chapter, the mean flow is never subtracted from the snapshot set beforehand. Thus, Equation 1 will yield the mean flow as the first mode of index i=1. This is an important requirement for the development of a mode bases that tracks changes in the mean flow, as will be shown in the following sections.

3.1 Short Time POD - SPOD

In order to gather some intuitive understanding of the effect of the number of snapshots used in order to derive spatial POD modes, data sets of a 2D CFD simulation of the circular cylinder wake at a Reynolds number of 100 were used. The sampling interval was 1/25th or 4% of a shedding cycle, and a maximum of four shedding cycles was available to calculate the spatial POD modes. These four shedding cycles were obtained from the steady state vortex shedding and identical to each other. Proper Orthogonal Decomposition was performed on snapshot ensembles with varying length consisting of 2 to 100 snapshots of the flow field. This corresponds to 0.05 to 4 shedding cycles. Figure 7 shows the mean difference between the POD modes obtained from these snapshot ensembles compared to the POD modes from all 100 snapshots or four shedding cycles. The first mode, representing the mean flow, is reaching close to zero error for every integer number of cycles. It can also be seen that the mean flow error reaches a maximum for snapshot ensembles containing 1.5, 2.5 and further multiples plus half shedding cycles. These maxima decrease in amplitude with increasing number of shedding cycles. While the Kármán shedding modes show the same trend of decreasing error maxima with increasing number of

shedding cycles, they feature a periodicity of twice that of the mean flow mode. For all modes, using a snapshot ensemble of less than one shedding cycle leads to a large increase in mode error. From the error plots in Figure 7, we conclude that exactly one shedding cycle is the smallest possible snapshot ensemble that will yield spatial POD modes comparable to those obtained from a large number of identical shedding cycles. This can be compared to Nyqyuists theorem that states that at least two samples per shedding cycle are needed to avoid aliasing problems. It can also be observed that the error minima at integer shedding cycles are very small, demonstrating the importance of selecting the snapshot ensemble to be exactly one shedding cycle in length. Inspecting the spatial distribution of the mean flow mode, shown in Figure 8, for a very short snapshot ensemble compared to a snapshot ensemble comprising exactly one shedding cycle, it can be seen that the short snapshot ensemble results in a mean flow mode that is not symmetric around the x axis, which is what would be expected from a large snapshot ensemble (which looks like the one shown in Figure 8 top right). The snapshot ensemble consisting of exactly one shedding cycle, however, results in a perfectly symmetric mean flow. Inspecting the 1st von Kármán shedding mode shown in Figure 8, one can observe the same effect of a small ensemble of a non-integer number of shedding cycles leading to an asymmetric mode distribution. This investigation into the effect of snapshot ensemble length demonstrates the possibility of selecting snapshot ensembles of integer number of shedding cycles in length in order to obtain the same spatial POD modes that would be obtained from snapshot ensembles that are many shedding cycles in length. For time periodic data sets this finding can be used to minimize the computational effort by minimizing the snapshot ensemble size. More importantly, however, for transient data sets which slowly change on a time scale much larger than a shedding cycle, one can obtain spatial POD modes for each individual shedding cycle and thus track the change in spatial modes over time. However, the steep increase in POD modal error for non-integer numbers of shedding cycle snapshot ensembles poses a problem for real life application of this short term POD decomposition. In order to use a snapshot ensemble of exactly one shedding cycle, a data set relatively finely resolved in time is needed. Additionally, an effective method of isolating exactly one shedding cycle is required. If either of these requirements is not met, a steep increase in error of the POD mode is encountered.

In summary, by applying POD decomposition to snapshot ensembles of varying length, an integer numbers of shedding cycles is needed in order to obtain spatial POD modes that are identical to those obtained from snapshot ensembles that are many cycles in length. Snapshot ensembles

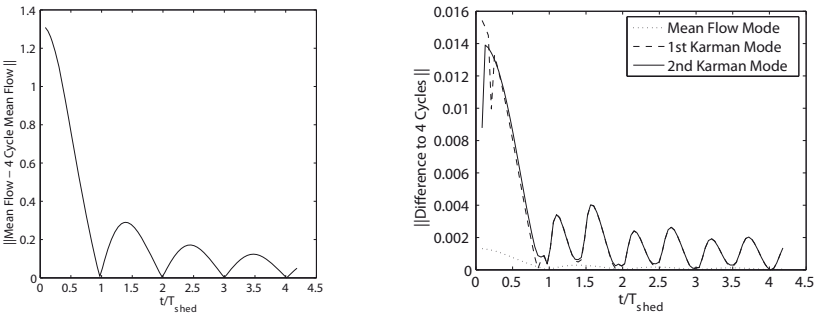

Figure 7. Average difference between Mean Flow Mode and the first two von Kármán Modes calculated from different fractions of a shedding cycle to mean flow calculated using all four shedding cycles

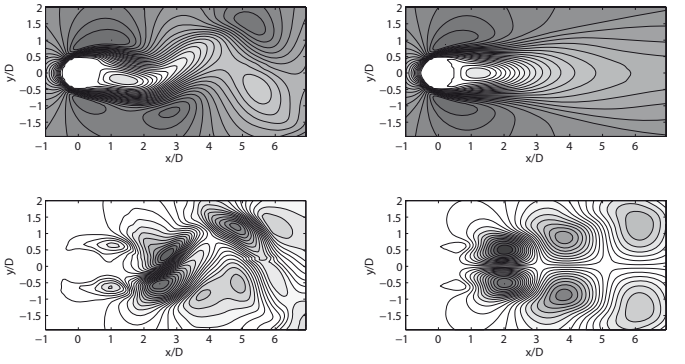

Figure 8. Mean Flow Modes (Top) and 1st Kármán Modes (Bottom) calculated from 4% of a shedding cycle (left) and 100% of a shedding cycle (right).Flow from left to right

that are non-integer multiples of shedding cycles lead to distorted POD modes. No useful POD modes could be obtained for snapshot ensembles that are less than one shedding cycle in length.

Gillies (1995) and Siegel et al. (2005) have shown that for time periodic flows, modes identical to those obtained from snapshot ensembles containing a large number of shedding cycles can be obtained using snapshot ensembles of small integer number of cycles, down to a minimum of one shedding cycle. Siegel et al. (2005) demonstrate that the difference between a spatial mode obtained from integer numbers of shedding cycles is minimal compared to a POD decomposition obtained from a large number (in the limit infinite) number of shedding cycles. Similar behaviour is observed in a fast Fourier transformation (FFT). hile in an FFT the error due to non-integer numbers of cycles can be alleviated using windowing functions, this approach does not appear to work for POD decompositions (Siegel et al., 2005). Siegel et al. (2005) refer to POD of only a single oscillation cycle as short time POD or SPOD, due to its similarity to procedures like Short Time Fourier decomposition. SPOD allows for a decomposition of time evolving flow fields with some approximate periodicity into (k) individual events of exactly one cycle of the dominant frequency,

$$u^{(k)}(x,y,t) = \sum_{i=1}^{I} a_i^{(k)} \phi_i^{(k)}(x,y). \tag{2}$$

The result is a collection of K cycles in individual bins. Note that these bins may contain a different number of samples in time, or span slightly varying time intervals as the period of a cycle changes. However, since SPOD yields K bins of spatial POD modes that are valid for one individual cycle of a transient flow change, it is not as low dimensional as one would wish: The result of SPOD is one entire mode set for each period of the flow. It should also be noted that modes obtained from an individual cycle are a priori not orthogonal to modes from other cycles. In fact, if the data is completely periodic, the modes obtained from different bins are identical if the number of snapshots is constant per cycle.

3.2 Double POD - DPOD

Building on the resulting spatial modes of a SPOD decomposition, one could conceive the following mode construction procedure: if the modes in two consecutive cycles vary only slightly, it should be possible to obtain a representation of the modes of the second cycle as the corresponding mode of the first cycle plus a small shift. This procedure, borne from the aforementioned mean flow mode or shift mode idea, can be formalized by

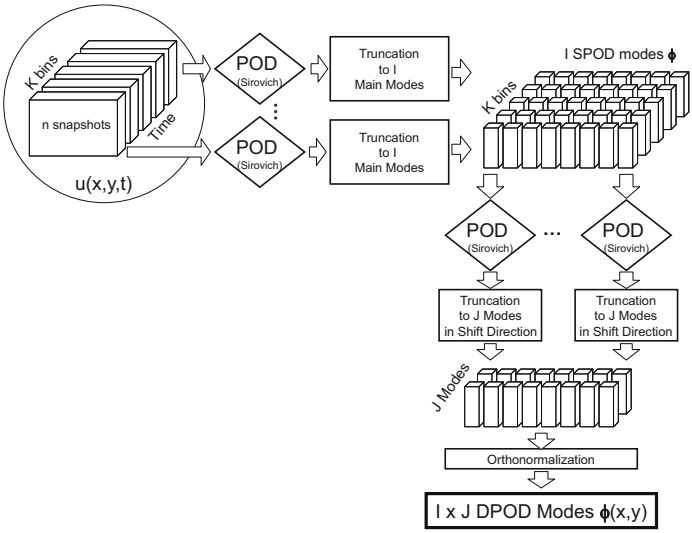

Figure 9. Flowchart of DPOD decomposition process.

realizing that mode i of all bins (k) from the SPOD procedure can be viewed as the input to a second POD (the bins now act similar to time above) as shown in

$$u^{(k)}(x,y) = \sum_{j=1}^{J} w_{ij}^{(k)} \Phi_{ij}(x,y) \tag{3}$$

This leads again to an optimal representation of all SPOD main modes i. Equation 4 summarizes the Double POD (DPOD) decomposition of the velocity vector field u:

$$u(x,y,t) = \sum_{i=1}^{I} \sum_{j=1}^{J} a_{ij}(t) \Phi_{ij}(x,y) \tag{4}$$

This DPOD formulation takes the concept of the shift mode one step further: we can now develop a shift mode, even a series of higher order shift modes, for all main modes i by applying the POD procedure to the

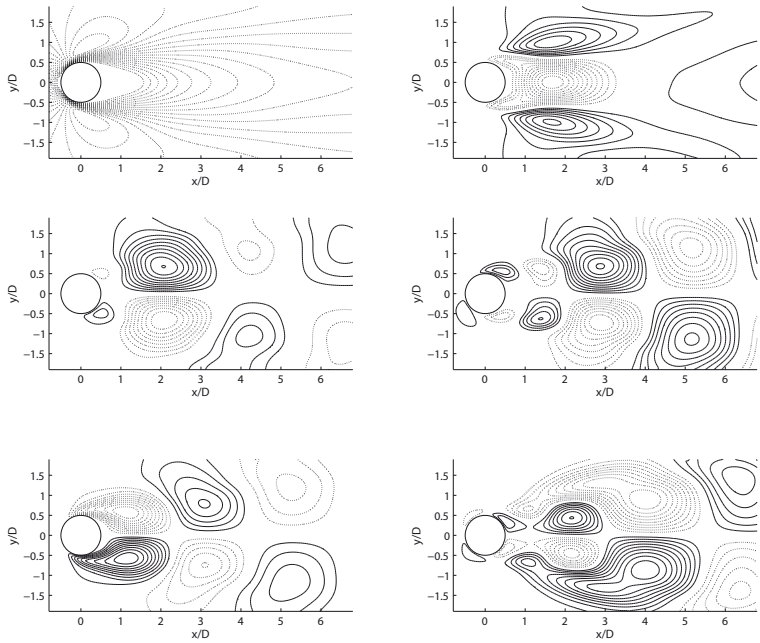

Figure 10. Transient forced DPOD spatial mode set using one shift mode for each main mode, the first 3×2 DPOD modes are shown. Iso-contours of streamwise velocity are shown, solid lines are positive, dashed lines negative.

POD mode sets. The resulting mode ensemble in its untruncated form has as many main modes I as there were snapshots in the smallest SPOD bin, and as many shift modes J as there were bins. It can then be truncated in both i and j, leading to a mode ensemble that is $I_M \times J_M$ in size. We will thus refer to the size of the truncated DPOD mode sets by indicating the truncation indices $I_M \times J_M$ in the following. A pictorial representation of the DPOD procedure is given in Figure 9.

Starting in the top left corner, the data is split into K bins and each bin is used as an input data set for its individual POD procedure. The resulting SPOD modes are then collected across the bins and POD is applied again to obtain the shift modes. The resulting Eigenfunctions can be truncated in both I and J in the same way as a regular POD decomposition. After orthonormalization, the decomposition is again optimal in the sense of POD. In the limit of $J = 1$, the original POD decomposition is recovered. While

the different modes distinguished by the index i remain the main modes described above, the index j identifies the transient changes of these main modes: For $J > 1$, the energy optimality of the POD decomposition in that direction leads to modes that are the optimum decomposition of a given main mode as it evolves throughout a transient data set. If $J = 2$, then modes $\phi_{1,1}$ and $\phi_{1,2}$ are similar to the mean flow and its shift mode or mean flow mode as described by Noack et al. (2003) and Siegel et al. (2003), respectively. Thus the modes with indices $j > 1$ can be referred to as first, second and higher order shift modes that allow the POD mode ensemble to adjust for changes in the spatial modes. We will refer to all of these additional modes obtained by the DPOD decomposition as shift modes, since they modify a given main mode to match a new flow state due to either a recirculation zone length or formation length change. This may be due to effects of forcing, a different Reynolds number, feedback or open loop control or similar events. Thus, in the truncated DPOD mode ensemble for each main mode, one or more shift modes may be retained based on inspection of energy content or spatial structure of the mode.

In Figure 10, a set of DPOD modes is shown that covers both the unforced time periodic vortex shedding state of the circular cylinder wake, as well as the low amplitude forced flow within the lock-in region. This mode ensemble will thus cover not just the limit cycle, but also the influence of forcing onto the vortex shedding process. The corresponding mode amplitudes contain the entire dynamics of the flow, and can thus be used to derive a set of equations describing the global flow dynamics for these flow states. In the following chapter, this is achieved using a system identification technique.

4 Artificial neural network system identification to develop a numerical plant model

With the DPOD spatial mode basis, developed in the previous section, covering a range of both Reynolds numbers and forcing conditions, the entire time dependent global dynamic behaviour of the flow is captured in the corresponding mode amplitudes. Thus, the next goal is to develop a set of equations describing the dynamic behaviour of these mode amplitudes. These equations are needed both for development of control algorithms, as well as for testing of these controllers. Traditionally, Galerkin projections of various types have been used to project the mode amplitudes onto the Navier Stokes equations. However, this approach has led to a variety of problems which are discussed both in the introduction and the following section, leading to the use of a different modelling approach as described in

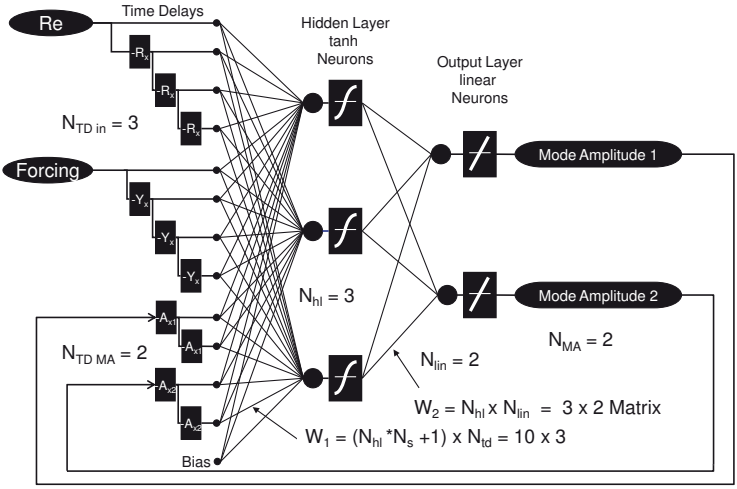

Figure 11. Example ANN-ARX system identification network topology. The model displays a two-layer (tanh and linear), 3 neuron, 2 mode output, 12 input neural network. The periodic sampling rates Rx , Yx , Axn , account for the number of time delays between sampled data points. Refer to table 1 for actual time delays and number of inputs for the 2×3 DPOD-ANN-ARX model.

the following sections.

In an alternative approach to assuring model stability, the ARX (Auto Regressive eXternal input) dynamic model structure, which is very widely used in the system identification community, is incorporated. A salient feature of the ARX predictor is that it is inherently stable even if the dynamic system to be modelled is unstable. This characteristic of ARX models often lends itself to successful modelling of unstable processes (Nelles, 2001). While system identification for a circular cylinder wake is demonstrate here, the nonlinearity of this flow is typical for a wide range of flow fields of technical interest, and thus this approach is applicable to other fluid dynamic modeling problems as well. The mathematical decscription of the linear ARX system indentification is as follows:

$$a(t) = q^{-d}\frac{B(q^{-1})}{D(q^{-1})}f(t) + \frac{1}{D(q^{-1})}e(t) \tag{5}$$

where a(t) is the state vector representing the POD mode amplitudes $a_i(t)$ shown Equation 4. f(t) describes the external input, which in the current effort is the vertical displacement of the cylinder and $e(t)$ is the white noise vector. For the above case, B and D are matrix polynomials in q^{-1}. The time delay operator is defined as

$$q^{-d}a(t) = a(t-d) \tag{6}$$

where d is a multiple of the sampling period. The parameter matrix, θ, and the regression vector, $\phi(t)$, are respectively defined as

$$\theta = [d_{ij}b_{ij}]^T \tag{7}$$

$$\phi = [a(t-1), ...a(t-n), f(t-d), ...f(t-d-m)]^T \tag{8}$$

As can be seen in Equation 8, the vector $\phi(t)$ is comprised of past states and past inputs. The ARX predictor (Ljung, 1999) may then be written as

$$\hat{a}(t|\theta) = q^{-d}B(q^{-1})f(t) + [1 - D(q^{-1})]a(t) = \phi^T(t)\theta \tag{9}$$

Equation 9 represents an algebraic relationship between the prediction, given on the left hand side, and past inputs and states, summarized by $\phi(t)$. The parameter matrix, θ, is determined during the estimation process. The main advantage of the ARX predictor is that it is always stable, even when the dynamic plant (the flow field in this case) being estimated is unstable. This feature is of utmost importance when modelling an unstable system such as the absolutely unstable cylinder wake flow.

The main drawback of this approach is that it is limited to modelling of linear systems which, as described above, is insufficient for modelling of unstable limit cycles. A general representation of nonlinear system identification, based on a hybrid ANN-ARX approach (Norgaard et al., 2003), may be written as:

$$a(t|\theta) = g[\phi(t), \theta] + e(t), \tag{10}$$

where θ is the matrix containing the weights of the ANN that are estimated by a back propagation algorithm using a training data set (Norgaard et al., 2003), and g is the nonlinear mapping realized by the feed-forward structure of the ANN.

The ANN-ARX predictor can then be expressed as

$$\hat{a}(t|\theta) = g[\phi(t), \theta], \tag{11}$$

The ANN-ARX algorithms used in this effort are a modification of the toolbox developed by Norgaard et al. (2003). We performed three main modifications to the toolbox. The first extends the toolbox for use in simulations, as opposed to one step prediction, of MIMO (multi-input, multi-output) systems. Secondly, the implementation of the time tapped delay system allows for the periodic sampling rate of inputs to the network. This helps to decrease network training times and also extends the time history of the inputs while keeping the number of inputs low. Thirdly, the network was decoupled, meaning the modes could be trained separately for individual mode amplitudes, and compiled to one large network at the end. This allowed for greater flexibility in training each individual mode. Although this resulted in much larger, more complex final network, simulation times are negligible when compared to training times. A schematic representation of the feed-forward ANN-ARX network topology is presented in Figure 11. After the DPOD mode amplitudes were obtained from the CFD data as described in the previous section, a single hidden layer ANN-ARX architecture is selected. The training set comprised input-output data obtained from CFD simulations. The model is validated for off-design cases and if the estimation error is unacceptable, then the ANN architecture is modified. This cycle was repeated until estimation errors were acceptable for all off-design cases.

The ANN-ARX predictor is inherently stable because, although the modelling approach is nonlinear, the algebraic relationship between the prediction and past states and inputs is preserved. This is extremely important when dealing with nonlinear systems represented by PDEs like the Navier Stokes equations, since the stability problems are more severe than in linear systems. The ANN-ARX approach is an ideal choice when the system to be modelled is deterministic and the signal to noise ratio (SNR) of the data is good (Norgaard et al., 2003).

The choice of the specific artificial neural network (ANN) architecture was based on three main design criteria. The first concerns the number of hidden layers. This was selected as one, i.e. a single hidden layer, since it is the simplest form that allows for a universal approximator (Cybenko, 1989) and its effectiveness for system identification problems has been shown by Norgaard et al. (2003). The second decision concerns the number of nodes. If the number of nodes in the hidden layer is small, the resulting error is unacceptable. As the number of nodes is increased, the error is reduced at the cost of computational complexity until a number of nodes is reached beyond which no further improvement in error is observed. The third design criterion is the choice of the network inputs time histories and delays. Larger sampling periods (more data inputs) allow for a more dynamical based

network. Networks with larger time histories are more accurate because of their ability to predict future outputs based on a better knowledge of the past. The trade off is that training time is greatly increased with larger input vectors. The time tapped delay technique allows for a much broader sampling period while keeping the number of inputs low, thus improving the training speed. For example, instead of using every single past delay input for a certain period of time, the time tapped delay can be set to 2 and the network will skip every other data point, thus reducing the number of inputs by 50% while maintaining the same total length of the sampling history. The time tapped delays for the Reynolds number input (Re), actuator input (Y) and previous mode amplitudes ($A \times n$) are presented in Table 1. Figure 11 represents a simple example of an ANN-ARX model.

Input	Number of past inputs or outputs to neural net	Delay between inputs		Total time history
Reynolds number	1	R_x	10	10
Actuator position	4	Y_x	2	8
Mode 1,1	1	$A_{x1,1}$	1	1
Mode 2,1	3	$A_{x2,1}$	8	24
Mode 3,1	3	$A_{x3,1}$	8	24
Mode 1,2	1	$A_{x1,2}$	1	1
Mode 2,2	1	$A_{x2,2}$	12	12
Mode 3,2	1	$A_{x3,2}$	12	12

Table 1. Network topology representing inputs and their time delays of the 3×2 DPOD-ANN-ARX model. The number of past inputs per signal and the delay which is the number of time steps in between sampling periods (R_x, Y_x, A_{xn}) are shown. The product of these two is the total time history.

Initially, neural networks were designed and trained for the 6 mode (3×2 DPOD) model. However, training times were very excessive, leading to the realization that new techniques needed to be established. The network was split into smaller sub-networks which could be trained much more efficiently. These smaller networks can be superimposed and compiled into a larger, overall network. Thus, many parameters exist in the design of ANN-ARX models. The resulting ANN has the following features:

4.1 Input Layer

Two network input parameters, the normalized cylinder displacement and the Reynolds number. In addition to these readings, in order to obtain a

strong representation of the dynamics of the system, the input layer includes past outputs of the 6 modes and past inputs for each of the two inputs (Reynolds number and cylinder displacement) as described in the toolbox developed by Norgaard et al. (2003), Table 1 presents the actual input / output time delays. The number of time delays for the past outputs was about one shedding cycle. On the other hand, the number of time delays for the past inputs was about half of a shedding cycle. The selection of time delays for past inputs/outputs was based on a sensitivity study which investigated the trade-off between estimation accuracy and network complexity. Therefore, the final configuration of the input layer chosen includes six mode outputs, namely, the first three main DPOD modes $i = 1, j = 1$; $i = 2, j = 1$; $i = 3, j = 1$ and their shift modes with a maximum of twenty-four time delays; two inputs, the Reynolds number and the actuator position for a maximum of 10 time delays; and one bias input. The total number of inputs to the net is therefore 15 (see Table 1).

4.2 Hidden Layer

One hidden layer consisting of 130 neurons. The activation function in the hidden layer neurons is the tanh function. A single bias input has been added to the output from the hidden layer.

4.3 Output Layer

6 outputs, namely, the 6 states representing the DPOD mode amplitudes of the 3×2 DPOD spatial mode basis. The output neurons have linear activation functions.

4.4 Weighing Matrices

The weighing matrix between the input layer and the hidden layer is of size (130×16), whereas the weighing matrix between the hidden layer and the output layer is of size (6×131). These weighing matrices are initialized randomly before training.

4.5 Training the ANN

Back propagation, based on the Levenberg-Marquardt algorithm, was used to train the ANN using the toolbox of Norgaard et al. (2003). The training data (21,790 time steps) is comprised of output from multiple CFD simulations. The first portion of the training consists of open loop forced transient simulations comprising design cases 1, 4, 5, 8, 10. The forcing of case 10 is repeated four times, with different starting phase angles of

0, 90, 180 and 270 degrees between the shedding of the vortices in the wake and the actuation signal at the start of the actuation. This leads to very different transient adjustment in the wake, which was found to be an essential feature for network training. During these open loop actuation simulations, the Reynolds number is fixed at $Re = 100$. The second portion of the training data set is comprised of transients in Reynolds number, while the actuation is kept at zero level. The Reynolds number transient training data include ramping from $Re = 100$ to $Re = 40$, from $Re = 100$ to $Re = 60$ and from $Re = 100$ to $Re = 160$.

Along with the 3×2 DPOD mode amplitudes, the Reynolds number as well as the cylinder displacement were provided as inputs to the network. The training procedure converged after 50 to 150 iterations depending on which mode was being trained.

As a measure of the model quality, Figure 12 shows both the mode amplitudes derived from the CFD simulation data, as well as the output of the ANN-ARX model. This is one of the open loop forced cases that were not used for model development, and the fact that there is good agreement between model and simulation data is an indication that this model is usable for controller development. In Siegel et al. (2008) more detailed comparison and error estimates for this model can be found.

5 Feedback control wake stabilization results

The amplitudes of the DPOD modes were used both for feedback, as described in the following section, as well as for evaluation of the effectiveness of the controller. In order to be used for feedback control, the amplitudes of the modes that were used for feedback were estimated using a least square fit of the sensor data onto the spatial POD modes. The sensor locations, developed by Siegel et al. (2005), were chosen to evenly cover the entire portion of the flow of interest for control, using 35 flow sensors.

The controller development is based on a Proportional and Differential (PD) controller. As mentioned earlier, active forcing is introduced into the wake by displacement of the cylinder in the flow normal direction as shown schematically in Figure 1. The full state estimator provides estimates for all 15 DPOD amplitudes which may then be used as input to a full state controller. Since a neural network model of the flow has been developed, several indirect control designs can be employed. The indirect design is very flexible and applicable in real-time for the problem at hand. Relevant concepts include approximate pole placement, minimum variance and predictive control. The approximation is based on instantaneous linearization, which is a popular method for control of ANN models. Following the effort

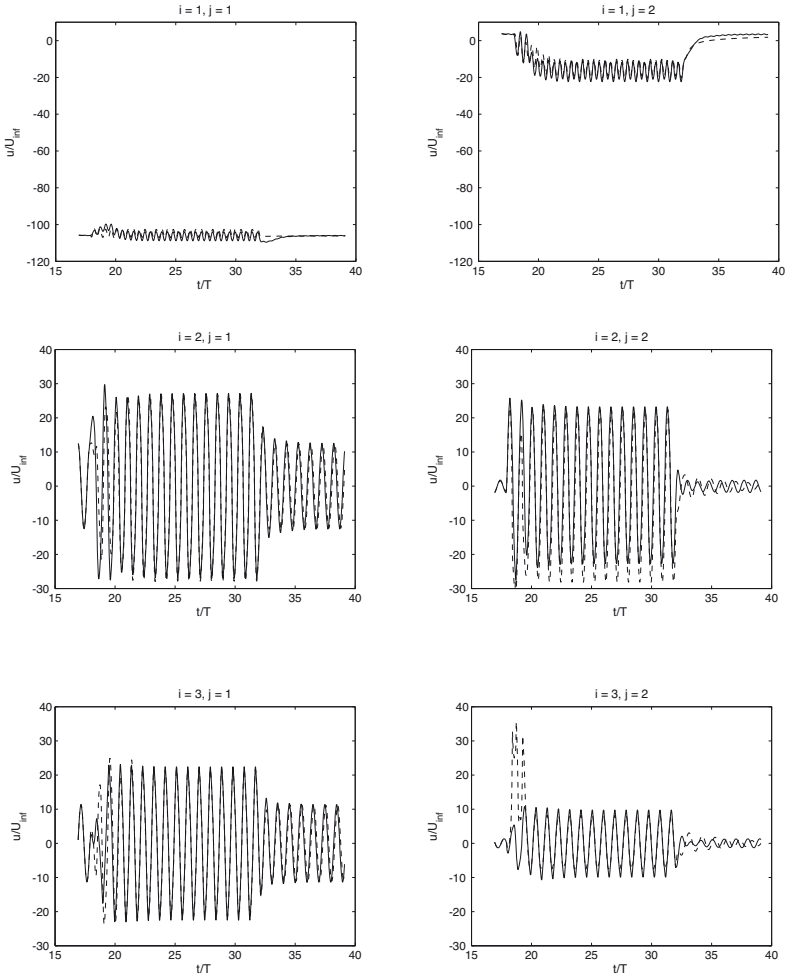

Figure 12. Mode amplitudes a_{ij} of the first 6 DPOD modes for the off-design case 3, forcing with $f/f_0 = 1.05$ and A/D =0.30. This data was not used for training of the ANN-ARX network. Lines, mode amplitudes from the CFD simulation; dot-dash lines, mode amplitude estimation from the ANN-ARX model.

by Siegel et al. (2005), a similar PD feedback control strategy is employed for the single mode feedback control law:

$$y_{cyl} = K_{p21} a_{21} + K_{d21} \frac{da_{21}}{dt}, \tag{12}$$

Instead of directly specifying the K_p and K_d gains, these can be expressed in terms of an overall gain K and a phase advance ϕ for mode i:

$$K_{pi} = K_i \cos(\phi_i), \tag{13}$$

$$K_{di} = \frac{K_i \sin(\phi_i)}{2\pi f}. \tag{14}$$

with f being the natural vortex shedding frequency. Equation 12 refers to the single-input closed-loop control based on feedback using an estimate of a_{21}. We will refer to this control law as a Single Input Single Output (SISO) controller in the following, since only one mode is used for determining the feedback response. The control law may be modified to enable dual-input from two modes, and a simple example is adding proportional control of the shift mode of the von Kármán shedding mode, a_{22}. This control law which essentially constitutes a multi input single output (MISO) PD controller is:

$$y_{cyl} = K_{p21} a_{21} + K_{d21} \frac{da_{21}}{dt} + K_{p22} a_{22} + K_{d22} \frac{da_{22}}{dt} \tag{15}$$

where it is possible to substitute the proportional and differential gains in the same fashion as for the SISO controller, with an overall amplitude gain K and a phase shift ϕ for each of the two modes independently.

5.1 Single Mode Feedback - SISO

This section presents the results both of single mode feedback, as well as multi mode feedback based on the DPOD model. The effectiveness of the single mode feedback was evaluated by keeping the overall gain K constant at $K = 5 \cdot 10^{-4}$, while varying the phase ϕ of the feedback from 0 to 360 degrees. The controller amplitude was chosen such that the overall cylinder displacement would remain within the range of validity of model, i.e. limiting the maximum displacement $y/D < 0.3$. The resulting closed loop lift and drag force are shown in Figure 13. It can be seen that for the range of feedback phases between about 30 and 250 degrees the drag is increased compared to the unforced flow field, while between 250 and 30 degrees a decrease in drag is observed. The normalized lift and drag forces for both $\phi = 90°$ (drag increase) and $\phi = 330°$ (drag decrease) is shown in Figure 14. In both cases the wake is stabilized at its new state, with about

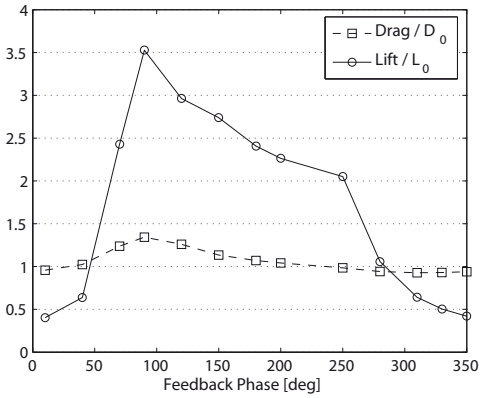

Figure 13. Non-dimensional RMS amplitudes of the stabilized lift and drag forces using single mode feedback (L_0 and D_0 are the RMS amplitudes of the unforced flow's lift and drag force, respectively).

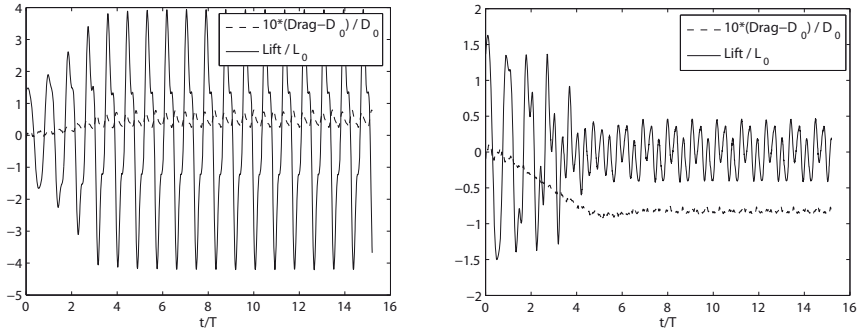

Figure 14. Lift and drag for feedback using $\phi_{21} = 90°$ (left) and $\phi_{21} = 330°$(right), SISO feedback

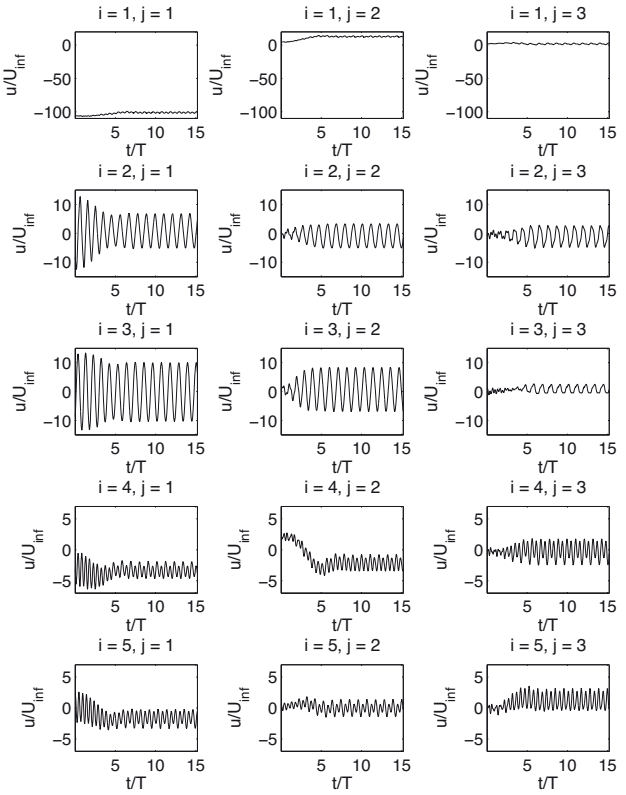

Figure 15. DPOD Mode Amplitudes for SISO feedback of a_{21}. $K_{21} = 5 \cdot 10^{-4}$, $\phi_{21} = 330°$. The controller is activated at $t/T = 0$.

10% drag decrease for the feedback phase of 330 degrees. The DPOD mode amplitudes for this case are shown in Figure 15. The amplitude of the mode used for feedback, a_{21}, is greatly reduced by the effect of feedback control. However, the amplitude of its first shift mode, a_{22}, is actually increased as a result of feedback. This observation led to the implementation of multi mode feedback, where both a_{21} and a_{22} are fed back with individual gains K and phases ϕ.

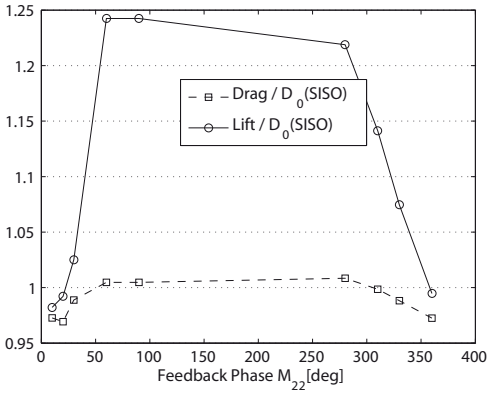

Figure 16. MISO feedback of a_{21} and a_{22}, with $K_{21} = K_{22} = 5 \cdot 10^{-4}$ and a fixed phase $\phi_{21} = 330°$. The feedback phase of a_{22} is varied, and the resulting lift and drag forces are normalized by those of the SISO feedback of a_{21} with $K_{21} = K_{22} = 5 \cdot 10^{-4}$ and $\phi_{21} = 330°$

.

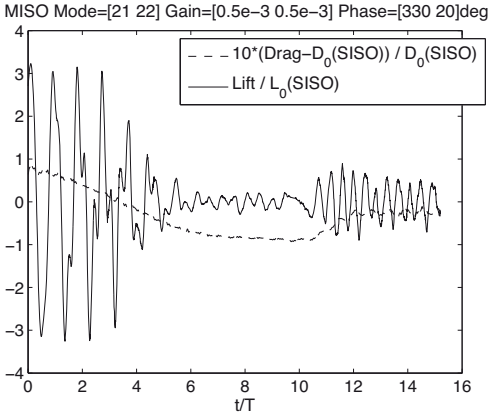

Figure 17. MISO feedback of a_{21} and a_{22}, with K_{21}=$K_{22} = 5 \cdot 10^{-4}$ and phases $\phi_{21} = 330°$ and $\phi_{22} = 20°$.

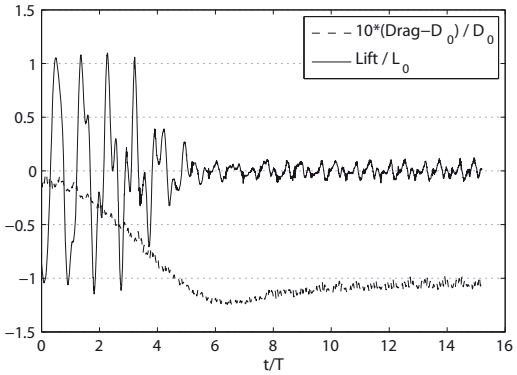

Figure 18. MISO feedback of a_{21} and a_{22}, with $K_{21} = 5 \cdot 10^{-4}$, $K_{22} = 3.75 \cdot 10^{-4}$and phases $\phi_{21} = 330°$ and $\phi_{22} = 20°$.

5.2 Two Mode Feedback - MISO

Figure 16 shows a parameter scan varying the feedback phase of a_{22}, while keeping all other gains as well as the phase of a_{21} constant. Again a range of detrimental phases exists, where the drag is compared to the new baseline of SISO feedback with 330 degrees phase. However, for a small range of feedback phases between $\phi = 310°$ and $\phi = 50°$, a further reduction in drag beyond SISO control can be observed. This demonstrates the benefit of using MISO control over SISO control. However, the detailed analysis of the MISO control run that led to the largest reduction in drag. Figure 17 shows that there is a lack of stabilization in this type of feedback control. While the drag initially is decreased by more than 6% compared to the SISO level, it is followed by an increase in drag to a stable level about 3% below the baseline. Analyzing the DPOD mode amplitudes for this case (not shown), it turns out that feedback of a_{22} actually destabilizes a_{21} later in the simulation.

With some fine tuning of the gains applied to a_{21} and a_{22}, the flow field can be stabilized using multi mode feedback. Figure 18 and Figures 19 demonstrate this, where the flow is stabilized at a lift fluctuation level more than one order of magnitude smaller than the unforced flow, and at a reduction of drag of more than 10%. However, in this situation the fluctuating amplitude of the shift modes is increased as well not just for the von Kármán modes, but also for the higher harmonic modes. From a flow physics perspective, this indicates a shift of the vortex formation further downstream as the controller becomes effective in stabilizing the near wake.

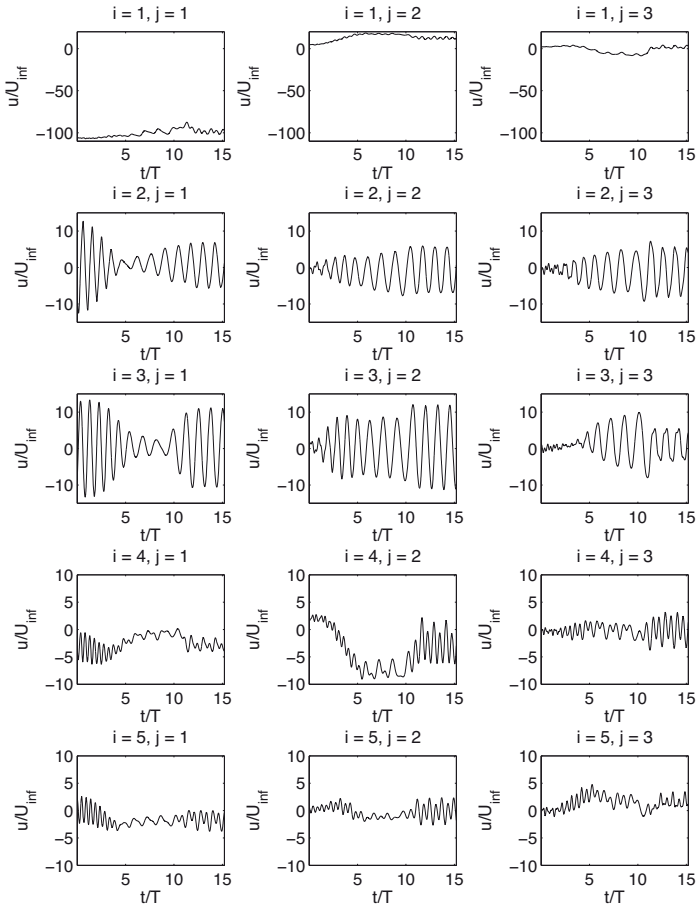

Figure 19. DPOD Mode Amplitudes for MISO feedback. $K_{21} = K_{22} = 5 \cdot 10^{-4}, \phi_{21} = 330°, \phi_{22} = 20°$. The controller is activated at $t/T = 0$.

At the same time, the actuation which remains at the cylinder location becomes less and less efficient in controlling this remaining vortex shedding far away from the cylinder. We demonstrate CFD results of feedback control of the two-dimensional circular cylinder at $Re = 100$ based on one or two modes of a DPOD model developed using transient data truncated to 15 Modes. The objective of the control is stabilization of the von Kármán vortex street. Two different controllers are investigated, a Single Input Single Output (SISO) feeding back the von Kármán mode amplitude, and a Multi Input Single Output controller (MISO), feeding back the von Kármán mode and its shift mode. The SISO control leads to a reduction of the overall drag force of about 10% for the best combination of proportional and differential gains. It is worth mentioning that all investigated SISO feedback gains lead to a stabilized wake, which is not the case for simpler POD models that do not include transient data. However, while the amplitude of the von Kármán mode used for feedback is reduced, as a result of feedback, the amplitude of its shift mode is increased. This observation leads to the use of the MISO controller, where in addition to PD feedback with the best combination of gains found in SISO control the shift mode, a_{21}, is used for feedback as well. The MIMO control is able to further reduce the drag compared to the best SISO case temporarily by more than 6 percent, in its stabilized state by more than 3 percent. The investigated MISO gains so far did lead to some partial destabilization of the flow. The overall conclusion of this investigation when compared to the results achieved by Siegel et al. (2005) is that a DPOD model is able to greatly increase the stability and performance of the feedback controller. This demonstrates the importance of obtaining a model that covers the entire parameter range that the feedback controller covers, rather than a point design type of a model that is only valid for the time periodic uncontrolled flow. While demonstrated for the benchmark case of a laminar circular cylinder wake, we expect similar trend for more complex flows which we will investigate in the future.

6 Key enablers / strategy

The focus of this chapter - as well as of this book in general - is on low dimensional modeling, as opposed to control theory. This already demonstrates that the main enabling strategy lies in in understanding and properly representing the flow field to be controlled, rather than conceiving elaborate control strategies. In the previous chapter, simple linear controllers were sufficient to stabilize the flow, while at the same time a nonlinear model was necessary to properly and accurately represent the flow physics. Thus,

the first and foremost enabling strategy lies in accurate, global and nonlinear modeling of the flow physics, without which the results shown here cannot be achieved.

The main idea behind the modeling approach introduced in this chapter is a shift away from the Galerkin projection of data onto the Navier Stokes equations. While this might seem counter intuitive, it resolves many physical and mathmatical problems. As any data obtained from either numerical simulations or experiments does not constitute an exact solution of these equations, the projection process will always be fraught with peril. For this reason, a more generic approach will have less numerical problems in that aspect. However, as has been shown above, one can still obtain a global model of complex flow fields, and in addition, establish fairly accurate input-output type relationships between actuation, Reynolds number or other parameters governing the flow. These input-output relationships are very difficult to include in a Galerkin projection due to the limits on boundary conditions that can be implemented. The quality of these relationships is what ultimately determines the useability of a model for feedback flow control controller development and testing.

6.1 Integration of Experiments, Modelling and Simulation

Beyond the use of simulation data for the purpose of deriving a numerical model, simulations offer many advantages at the start of a feedback flow control investigation. Typically, it is faster to develop a numerical grid for a given flow geometry than for example, to design and build a wind or water tunnel model with actuation capabilites. This makes simulations a great tool for development of the actual experiment, for example, to investigate the performance of an actuation system without having to actually build it (Siegel et al., 2009). On the other hand, once a physical model for the flow field of interest is available, performing parameter studies covering different actuation conditions and Reynolds numbers is done much faster in an experiment than in a simulation. Especially for high Reynolds number turbulent flow fields, the computation cost is often in the order of thousands of CPU hours for investigating a single forcing condition, whereas a parameter study of actuation parameters can be achieved in a wind tunnel in a matter of seconds to minutes. This makes experiments a preferred option to determine for example the lock-in range for a flow under investigation. However, experiments will always only yield very limited amounts of flow data, as even with state of the art PIV measurement techniques at most a two dimensional slice of the flow can be measured simultaneously. This circumstance makes experimental data of limited use when a global model of

the flow field is to be developed. Therefore, a typical work flow would use the experiment to determine particular flow conditions of interest, which then can be investigated in greater detail using CFD simulations.

6.2 Applicability to Other Flow Fields

In this chapter, the applicability of the DPOD-ANN-ARX modeling approach to the benchmark problem of a two dimensional circular cylinder wake at a Reynolds number of 100 has been shown. In this section, the applicability of the method to more complex flows is discussed. While the aim is to avoid unsupported speculation as much as possible, it is beneficial to share some insight that has been gained into this issue from preliminary investigations.

The circular cylinder wake at $Re = 100$ is not an entirely two dimensional flow, as has been shown in literature for example by Williamson (1996) for the unforced wake, or by Seidel et al. (2006) for the feedback controlled wake. The latter investigation demonstrates how feedback control can cause a spanwise phase shift of the vortex shedding until the two dimensional control approach is only effective in a narrow vicinity of the two-dimensional sensing plane that was used in this investigation. This poses the question if the DPOD-ANN-ARX approach can be extended to three dimensional flows. The answer is an unequivocal yes, since there is no limitation of the DPOD decomposition to two dimensions as was pointed out in a previous section.

As the Reynolds number of the circular cylinder flow is increased beyond $Re \approx 180$, secondary instabilities lead to the formation of streamwise vortices (Williamson, 1996). While modeling these flow features will lead to additional DPOD modes, there is no apparent obstacle in applying POD based modeling to flow fields that contain both two and three-dimensional features, as has been show in the hybrid approach proposed by Ma and Karniadakis (2002). Extending their work to use the DPOD procedure introduced here, one could derive dynamic models capturing the effect of actuation and/or changes in Reynolds number for these types of flows. This could then be used to develop feedback controllers to suppress the von Kármán type vortex shedding. Based on results of Cohen et al. (2003), where feedback of the von Kármán mode only suppressed higher order harmonic modes, there is hope that suppressing the Kármán vortex street might eliminate the streamwise vortices as well, since they are the result of a secondary instability that only exists in the presence of the von Kármán vortex shedding. If this conjecture is in fact true remains to be shown.

At yet higher Reynolds numbers ($Re > 3900$), the von Kármán vortices break down into smaller and smaller turbulent structures that ultimately dissipate their energy into heat. These smaller structures can be quite energetic and thus more and more POD modes will need to be retained in order to model a given fraction of the overall flow energy content. This behaviour of POD is due to the energy optimality of the procedure, and DPOD inherits this property from POD. As a result, both POD and DPOD models will become inherently large for flows that break down into turbulence. If the purpose of model development is feedback control, however, there may not be a need to model the small turbulent eddies in order to capture the dynamic behaviour of the large vortical structures. In the case of the circular cylinder wake, one may only be interested in modeling the von Kármán type shedding for the reasons outlined in the previous paragraph. Thus, an approach where the flow data is subjected to either spatial or temporal filtering may be applied, as has been pursued by the authors with good preliminary results (Siegel et al., 2007). The approach proposed in this work removes small scale turbulent structures from the data used for POD mode derivation while retaining the large structures (i.e. von Kármán vortices) that are of interest for feedback controller development. This approach is much like the use of spatial filtering in large eddy simulations employed in state of the art CFD solvers. With this approach, one introduces a choice of how much or how little of the smaller structures are included in the model. Thus, one can derive models with relatively few modes that nonetheless capture the dynamics of the flow that is relevant for feedback control. The DPOD-ANN-ARX approach is particularly suited to this type of modeling, since no turbulence model is required: As the entire model development is data driven and does not include projection onto the Navier Stokes equations, no closure equations are needed. The approach can thus be used as introduced here, with the only added step being a filtering process before the derivation of the DPOD modes. However, more detailed investigations into filter kernel type, size and cutoff wave length are needed.

Another approach to modeling of turbulent flows, in this case a turbulent shear layer, has been demonstrated by Seidel et al. (2009). Instead of removing small vortex structures by spatial or temporal filtering, the enabling approach to modeling of the large vortex structures shown there is based on using a different flow quantity altogether. By deriving POD spatial modes from the density field, all small vortical structures are removed from the flow field while preserving the large shear layer vortices of interest. While the DPOD-ANN-ARX approach allows for the development of a global dynamic model based on density, a traditional Galerkin projection is not suitable for deriving such a model that does not even attempt to describe the velocity

field.

An important question pertaining to the application of DPOD to flow fields with multiple equally dominant frequencies is the selection of appropriate bin boundaries. In the present investigation, higher frequency content was small in amplitude compared to the fundamental frequency of the vortex shedding, and thus the lift force with a simple peak detection algorithm was suitable for bin segmentation. If several dominant frequencies coexist, there are different possible approaches to segmentation. Using a phase accurate temporal notch filter as a preprocessing step, one may recover the fundamental frequency and determine bin boundaries in the same fashion as introduced in this work. Alternatively, it is conceivable to use open loop forcing to elevate the amplitude of one of the dominant frequencies at a time, thus allowing for discovery of the spatial flow features related to each frequency using multiple SPOD procedures, one for each of the frequencies of interest. Having outlined possible pathways of how the DPOD-ANN-ARX approach may be applied and extended to flow fields at Reynolds number of technical interest, the question remains how applicable this approach may be to other flow geometries. We consider the circular cylinder wake a prototype flow featuring separated free shear layers that develop instabilities leading to vortex shedding. As such, there are similarities to many other flows of technical interest that contain free shear layers, featuring both simpler and more complex flow behavior. Examples that have been investigated by the authors are the separated flow over a stalled airfoil, free shear layers formed behind a D shaped cylinder, and the wake of an axisymmetric bluff body. While we yet have to apply the DPOD-ANN-ARX approach to these flows, we consider them promising candidates since they all feature large coherent structures resulting from instabilities. The interaction of these instabilities and their resulting structures with flow actuators of various kinds is of great technical interest, both for open and closed loop flow control. DPOD-ANN-ARX models may be used to investigate this interaction in a structured and quantitative fashion.

In summary, this approach was developed with the intent to use the resulting models for controller development in order to achieve control of the formation of large structures caused by flow instabilities. From a technical perspective, these types of flows are the most promising candidates for feedback flow control since instabilities can be influenced with relatively small amounts of actuation energy. This is important in the context of the power limitations inherent in state of the art of dynamic flow actuators. Our approach supports flow fields with many different modes present, and can also accommodate multiple actuator interaction allowing for MIMO control. We did not intend it to be used for random turbulent flows, but find that there

are many technical applications where this limitation is of no importance.

Bibliography

K. Cohen, S.G. Siegel, T. McLaughlin, and E. Gillies. Feedback control of a cylinder wake low-dimensional model. *AIAA Journal*, 41(7):1389–1391, 2003.

G. V. Cybenko. Approximation by superpositions of a sigmoidal function. *Maths Control Signals Syst*, Vol. 2:303–314, 1989.

E. A. Gillies. *Low-dimensional characterization and control of non-linear wake flows*. PhD thesis, Faculty of Engineering, University of Glasgow, UK, 1995.

H. Oertel. Wakes behind blunt bodies. *Ann. Rev. Fluid Mech.*, 22:539–564, 1990.

L. Ljung. *System Identification Theory for the User*. Prentice-Hall, 2nd edition, 1999.

X. Ma and G. Karniadakis. A low-dimensional model for simulating three-dimensional cylinder flow. *J. Fluid Mech.*, 458:181–190, 2002.

C. Min and H. Choi. Suboptimal feedback control of vortex shedding at low Reynolds numbers. *J. Fluid Mech.*, 401:123–156, 1999.

O. Nelles. *Nonlinear System Identification*. Springer, 2001.

B. R. Noack, K. Afanasiev, M. Morzynski, G. Tadmor, and F. Thiele. A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *J. Fluid Mech.*, 497:335–363, 2003.

M. Norgaard, O. Ravn, N. K Poulsen, and L. K Hansen. *Neural Networks for Modelling and Control of Dynamic Systems*. Springer Series Advanced Textbooks in Control and Signal Processing, Springer, London, 2003.

J. Seidel, S.G. Siegel, K. Cohen, V. Becker, and T. McLaughlin. Simulations of three dimensional feedback control of a circular cylinder wake. *AIAA Paper* 2006-1404, 2006.

J. Seidel, S.G. Siegel, and T. McLaughlin. Computational investigation of aero-optical distortions in a free shear layer. *AIAA Paper* 2009-0362, 2009.

S. Siegel, K. Cohen, and T. McLaughlin. Feedback control of a circular cylinder wake in experiment and simulation (invited). *AIAA Paper* 2003-3569, 2003.

S. G. Siegel, K. Cohen, J. Seidel, and T. McLaughlin. Short time proper orthogonal decomposition for state estimation of transient flow fields. *AIAA Paper* 2005-0296, 2005.

S.G. Siegel, S. Aradag, J. Seidel, K. Cohen, and T. McLaughlin. Low dimensional POD based estimation of a 3D turbulent separated flow. *AIAA Paper* 2007-0112, 2007.

S.G. Siegel, J. Seidel, C. Fagley, D.M. Luchtenburg, K. Cohen, and T. McLaughlin. Low-dimensional modelling of a transient cylinder wake using double proper orthogonal decomposition. *J. Fluid Mech.*, 610:1–42, 2008.

S.G. Siegel, J. Seidel, and T. McLaughlin. Experimental study of aero-optical distortions in a free shear layer. *AIAA Paper* 2009-0361, 2009.

L. Sirovich. Turbulence and the dynamics of coherent structures. Part I: Coherent structures. *Quarterly of Applied Mathematics*, 45(3):561 – 571, 1987.

C. H. K. Williamson. Vortex dynamics in the cylinder wake. *Ann. Rev. Fluid Mech.*, 28:477–539, 1996.