

Feedback control of unstable flows: a direct modelling approach using the Eigensystem Realisation Algorithm

Thibault L. B. Flinois^{1,†} and Aimee S. Morgans¹

¹Department of Aeronautics, Imperial College London, South Kensington, London SW7 2AZ, UK

(Received 5 August 2015; revised 21 December 2015; accepted 7 February 2016;
first published online 14 March 2016)

Obtaining low-order models for unstable flows in a systematic and computationally tractable manner has been a long-standing challenge. In this study, we show that the Eigensystem Realisation Algorithm (ERA) can be applied directly to unstable flows, and that the resulting models can be used to design robust stabilising feedback controllers. We consider the unstable flow around a D-shaped body, equipped with body-mounted actuators, and sensors located either in the wake or on the base of the body. A linear model is first obtained using approximate balanced truncation. It is then shown that it is straightforward and justified to obtain models for unstable flows by directly applying the ERA to the open-loop impulse response. We show that such models can also be obtained from the response of the nonlinear flow to a small impulse. Using robust control tools, the models are used to design and implement both proportional and \mathcal{H}_∞ loop-shaping controllers. The designed controllers were found to be robust enough to stabilise the wake, even from the nonlinear vortex shedding state and in some cases at off-design Reynolds numbers.

Key words: instability control, low-dimensional models, wakes

1. Introduction

1.1. *The need for low-order models*

Over the last few decades, the aerospace and automotive industries among others have developed a keen interest in flow control. Indeed, great promise lies in the modification of the dynamics of fluid flows for drag reduction, stabilisation of fluctuations, lift enhancement, mixing optimisation, etc. A plethora of both passive strategies (with no energy input) and active strategies (with an external source of energy) have been applied with great success in a large spectrum of applications. The reader is referred to reviews such as Gad-el Hak (2000), Kim & Bewley (2007) and Choi, Jeon & Kim (2008) for a more detailed account of some of these attempts.

Closed-loop control is a type of active control whereby information is measured in real time in the flow field, in order to automate the response of the actuation. In

† Email address for correspondence: t.flinois11@imperial.ac.uk

some cases, feedback control can be applied without a model for the input–output dynamics. For instance, it is possible to optimise parameters of an open-loop control strategy, by using methods like extremum or slope seeking control (e.g. Beaudoin *et al.* 2006*a,b*; Becker *et al.* 2007; Henning *et al.* 2007, 2008; Pastoor *et al.* 2008). This allows the controller to track reference output values across a range of operating conditions. However, these quasi-steady controllers respond slowly by design and are therefore intrinsically unable to directly interact with the flow dynamics. Alternatively, a deep understanding of the flow structure and control objective sometimes makes it possible to design controllers based on ‘intuition’. These usually target specific flow features and can thus be very powerful (e.g. Pastoor *et al.* 2008; Joe, Colonius & MacMynowski 2010). Unfortunately, for most flows of interest, even a good understanding of the flow physics is not sufficient to guide the design of effective controllers. Another approach is to take advantage of machine learning techniques. These methods have also been used to control flow fields and have provided promising results in several studies (e.g. Milano & Koumoutsakos 2002; Sengupta, Deb & Talla 2007; Gautier *et al.* 2015). However, a large number of simulations and/or experiments is typically required for the algorithms to converge, making these methods prohibitive for many flows.

An issue with model-free control methods is that they are often either too specific or too general to provide an efficient means of designing reliable controllers across a broad range of fluid flows. The availability of a model can make a large set of powerful and mature control design tools available. These tools are not in general directly applicable to the Navier–Stokes equations, which are infinite-dimensional, nonlinear, partial differential equations. If a low-order linear approximation of the input–output dynamics of the flow is available, however, linear feedback control methods can directly be applied.

Optimal control is one of the most popular techniques in flow control. For instance, the linear–quadratic–Gaussian (LQG) framework has been successfully used in numerous studies (e.g. Åkervik *et al.* 2007; Huang & Kim 2008; Bagheri *et al.* 2009*a,c*; Bagheri, Brandt & Henningson 2009*b*; Barbagallo, Sipp & Schmid 2009, 2011; Semeraro *et al.* 2011; Illingworth, Morgans & Rowley 2011, 2012; Barbagallo *et al.* 2012; Dadfar *et al.* 2013; Juillet, Schmid & Huerre 2013; Fabbiane *et al.* 2014, 2015). It is appealing for its theoretical optimality and intuitive structure, based on the design of a dynamic observer (Kalman filter), which optimally estimates the state of the system, and a linear–quadratic regulator (LQR), which minimises a chosen cost function. However, LQG control does not have any stability margin guarantees (Doyle 1978) and requires detailed information about the noise and uncertainty in the system, which may not be available.

Another common optimal control approach is model predictive control (MPC), whereby the control waveform that minimises a cost function over a chosen horizon is evaluated at every time step (e.g. Goldin *et al.* 2013; Fabbiane *et al.* 2014; Ghiglieri & Ulbrich 2014). This approach can also be used as an off-line adjoint-based optimisation technique if applied to the full (potentially nonlinear) Navier–Stokes equations (e.g. Bewley, Moin & Temam 2001; Protas & Styczek 2002; Joe *et al.* 2010; Flinois & Colonius 2015). In this case, the results are useful for revealing effective control mechanisms, which can for instance be exploited using intuition-based controllers (e.g. Joe *et al.* 2010).

Robust or \mathcal{H}_∞ control tools have also been successfully implemented in several studies. Having the ability to optimise the robustness of the controller to model uncertainties and disturbances is particularly appealing when its design is based on a

linear low-order model that neglects a large part of the flow dynamics and does not account for nonlinearities. As with optimal control, it is possible to identify either the waveform (Bewley, Temam & Ziane 2000) or the observer–controller pair (Bewley & Liu 1998; Lauga & Bewley 2004) that minimises a chosen cost function. Unlike with optimal control, however, robustness is improved by assuming that the worst-case disturbances are present in the flow.

The model-based methods introduced above mainly rely on the minimisation of a cost function so the design procedure does not allow much user intervention. Instead, some authors have turned to classical loop-shaping to control fluid flows. With this approach, lead–lag compensators and the like can be used to modify the Nyquist/Bode diagram and manually shape the dynamics of the controller (e.g. Kwong & Dowling 1994; Dowling & Morgans 2005; Morgans & Dowling 2007; Morgans & Stow 2007; Dahan, Morgans & Lardeau 2012).

Approaches that combine the flexibility of classical control with the optimality and robustness of modern control have also been applied in fluid mechanics. For example, in the mixed-sensitivity approach, the designer is able to specify closed-loop performance specifications by constructing weights that influence the robust controller being synthesised (e.g. Becker *et al.* 2005; Henning & King 2007; Henning *et al.* 2007; Williams *et al.* 2010). The \mathcal{H}_∞ loop-shaping framework introduced by McFarlane & Glover (1992) provides even more control over the closed-loop system's dynamics. In this case, the gain of the transfer function (for single-input–single-output (SISO) systems) is first shaped as in classical loop-shaping before the robustness of the controllers is optimised. In this article, we use this procedure and show that it is particularly well suited to the robust stabilisation of fluid flows using projection-free models based on the Eigensystem Realisation Algorithm (ERA).

1.2. Model reduction

In the previous section, we have given a brief overview of the large spectrum of tools that can be used once a linear low-order model of the relationship between the inputs and the outputs of the system has been obtained. In order to construct such a model, two dominant approaches can be identified from previous studies. The first approach is based on the Galerkin projection of known (linearised) flow equations onto a set of modes, in order to reduce the dimension of the system. These modes can hold valuable information, which can help both to understand the flow structure and to design the controller, for instance by guiding actuator–sensor placement. Additionally, they provide a way to reconstruct the full flow field from the reduced-order model (ROM) state. This in turn facilitates the interpretation of the effect of the controller on the flow and enables the use of full-state feedback algorithms (e.g. LQR).

Choosing the nature of the modes onto which the dynamics are projected is the first crucial step of this approach. ‘Global modes’ (e.g. Åkervik *et al.* 2007; Henningson & Åkervik 2008; Barbagallo *et al.* 2011; Ehrenstein, Passaglia & Gallaire 2011) and ‘proper orthogonal modes’ (e.g. Aubry *et al.* 1988; Podvin & Lumley 1998; Graham, Peraire & Tang 1999; Ravindran 2000*a,b*, 2002, 2006; Prabhu, Collis & Chang 2001; Ma & Karniadakis 2002; Noack *et al.* 2003; Gloerfelt 2008; Siegel *et al.* 2008; Barbagallo *et al.* 2009, 2012; Tadmor *et al.* 2010) can yield successful ROMs, but ‘balanced modes’ have been shown to lead to superior performance in terms of robustness and required model order in a number of studies (e.g. Willcox & Peraire 2002; Ilak & Rowley 2008; Bagheri *et al.* 2009*c*; Barbagallo *et al.* 2009; Dergham *et al.* 2011). If the system is stable, such modes can be approximately identified

at a low computational cost, by using a method (Willcox & Peraire 2002; Rowley 2005) referred to as approximate balanced truncation or balanced proper orthogonal decomposition (BPOD), which relies on snapshots from a set of impulse responses of the forward and associated adjoint system (one of each for SISO systems). This technique assumes the flow is evolving around a stable equilibrium point (base flow) and that the impulse response of the (linearised) forward and adjoint systems can be computed. It has been successful in modelling the dynamics of a large range of flows (e.g. Willcox & Peraire 2002; Rowley 2005; Bagheri *et al.* 2009a,b; Ahuja & Rowley 2010; Dergham *et al.* 2011; Semeraro *et al.* 2011; Tu & Rowley 2012).

One drastic limitation of BPOD is the fact that it is restricted to stable flows. Extensions have therefore been designed to circumvent this issue, but these require either the evaluation of all the unstable global modes of the flow field (Barbagallo *et al.* 2009; Ahuja & Rowley 2010) or the inversion of large matrices (Dergham *et al.* 2011), making them often computationally intractable. These restrictions have recently been lifted, as it was shown by Flinois, Morgans & Schmid (2015) that the snapshot-based approach can in fact be used directly to obtain ROMs of unstable flows, as long as the snapshots are chosen appropriately. This is therefore one of the approaches that is considered in this study. An outline of the BPOD method is provided in § A.1.

1.3. System identification

The second main approach for obtaining low-order linear models is system identification (for an overview, see Ljung 1987). In this case, only information collected by sensors and chosen actuator inputs is used to obtain the model. These methods are therefore more readily applicable in experiments than the projection approach introduced in § 1.2. Indeed, they do not rely on full-state information or knowledge of the governing equations, and hence do not require a linearised or an adjoint solver. Although no modal or sensitivity information is readily available, input–output models are sufficient for the purposes of controller design.

In a limited number of cases, it is possible to first model each physical phenomenon affecting the system independently. The resulting submodels can then be combined in order to create the full input–output model. For instance, Illingworth *et al.* (2012) considered the behaviour of the shear layer, the acoustics, the scattering and the receptivity of a cavity flow separately in order to construct an overall model. For most flow set-ups, however, this is not practical, and more general approaches are required. To this end, the relationship between the actuators and the sensors is often assumed to have a general mathematical description with unknown parameters. These parameters are then calibrated by minimising the error between a representative set of input–output data and the model predictions. This can be done either in the frequency domain or in the time domain. In the former case, the data is fitted directly to a chosen transfer function structure (e.g. Kwong & Dowling 1994; Zhu, Dowling & Bray 2005; Dahan *et al.* 2012). In the latter case, ‘prediction–error methods’ are often used, whereby the output is considered to be a linear combination of earlier inputs, outputs and a moving average of an error term. In flow control, only some of these components are typically necessary to obtain an accurate model, as for instance in Morgans & Dowling (2004), Zhu *et al.* (2005), Huang & Kim (2008), Hervé *et al.* (2012), Juillet, McKeon & Schmid (2014), Roca *et al.* (2014) and Gautier *et al.* (2015).

System identification can also be performed using ‘subspace identification’ methods, which result in an approximation for the system in state-space form directly. Subspace

identification has been successfully implemented in some flow control studies (for instance Huang & Kim 2008; Juillet *et al.* 2013; Guzman Inigo, Sipp & Schmid 2014) and a review of these methods can be found in Qin (2006). Despite the more advanced mathematical framework, one of the advantages of this approach is that it yields noise covariance matrices that can be used in an LQG framework.

Finally, a method that has attracted an increasing amount of attention in recent years is the ‘Eigensystem Realisation Algorithm’ (ERA) introduced by Juang & Pappa (1985) and subsequently used in numerous flow control studies (e.g. Ma, Ahuja & Rowley 2011; Illingworth *et al.* 2011, 2012; Dahan *et al.* 2012; Belson *et al.* 2013; Brunton, Rowley & Williams 2013; Dadfar *et al.* 2013; Brunton, Dawson & Rowley 2014; Illingworth, Naito & Fukagata 2014). It is straightforward to implement and results in state-space models that have recently been shown to be theoretically equivalent to those obtained with BPOD for stable systems (Ma *et al.* 2011). The focus of this article is the ERA, and an outline of the method is thus included in § A.2.

As with most system identification techniques (and BPOD), the ERA was originally only designed for stable systems. One solution is to first stabilise the system with an *ad hoc* controller, then identify the (stable) closed-loop dynamics, from which the plant’s open-loop (unstable) dynamics can finally be extracted. This approach has been successful in some studies (e.g. Morgans & Dowling 2007; Illingworth *et al.* 2011, 2012), but is not always practical, as it can be challenging to find a stabilising controller without a model. Here, we instead use the theoretical equivalence of the ERA and BPOD and the fact that BPOD can be applied to unstable systems to build ERA models directly from the impulse response of unstable systems.

1.4. Article structure

In this paper, we build on the findings of Ma *et al.* (2011) and Flinois *et al.* (2015) to show that the ERA is a practical and computationally cheap approach to obtain low-order models of unstable flows that are useful for feedback control. We confirm that such models are in practice very similar to the ones obtained with BPOD and also show that they can be obtained directly from the nonlinear system dynamics, without the need for a linearised or an adjoint solver. As a result, the method may even be applicable in some experimental set-ups.

After introducing the numerical approach and set-up in § 2, we consider the unstable flow over a D-shaped body and compare three different ROMs in § 3: the first is obtained with BPOD (§ 3.2), the second with the ERA (§ 3.3), and the third also with the ERA but based on the flow’s nonlinear dynamics (§ 3.4). We show that the three models are very similar, as expected, and then design controllers based on the last set of models using proportional control and \mathcal{H}_∞ loop-shaping in § 4. We show that the controllers are able to stabilise the nonlinear flow, even from the vortex shedding limit cycle and in some cases at off-design Reynolds numbers. We compare two sensor configurations, where in one case the sensor is located in the wake and in the other the sensor is body-mounted. Concluding remarks are included in § 5.

2. Numerical set-up

The work presented in this article is based on the fast multi-grid immersed-boundary fractional step (IBFS) algorithm introduced in Taira & Colonius (2007) and Colonius & Taira (2008) for two-dimensional direct numerical simulations of incompressible flows. Various versions of the IBFS finite-volume code have been rigorously tested




FIGURE 1. Input/output and flow set-up for the simulations. The actuators (dark grey) are disk-shaped body forces with a diameter of 10% of the body height, acting in opposite horizontal directions as shown by the arrows. The first sensor measures the vertical velocity of the flow along the symmetry plane, two body heights downstream of the trailing edge (black triangle). The second sensor (light grey) is a distributed sensor, which measures the antisymmetric component of the force acting on the base of the body.

and used in many studies (e.g. Taira & Colonius 2007; Colonius & Taira 2008; Ahuja & Rowley 2010; Joe *et al.* 2010; Joe, Colonius & MacMynowski 2011; Ma *et al.* 2011; Tu & Rowley 2012; Brunton *et al.* 2013, 2014; Choi, Colonius & Williams 2015; Flinois & Colonius 2015) so only an overview of the formulation of this code is given here. The reader is referred to the cited studies and the references therein for further details.

The incompressible Navier–Stokes equations in vorticity form are considered and discretised on a uniform Cartesian grid. A set of discrete immersed-boundary forces is defined along the surface of the body and regularised onto the Cartesian grid in order to enforce the no-slip boundary condition on the surface. The flow equations are solved sequentially on a number of nested grids at each time step. Each grid is identical to the one it is nested into, but has half the physical extent and hence is twice as fine. The flow state on the next grid level is used to obtain boundary conditions at the edges of each grid. For the largest grid, a far-field boundary condition is used (zero vorticity). In order to solve for the linearised and adjoint dynamics of the flow, the unstable steady state is identified using selective frequency damping (Åkervik *et al.* 2006; Jordi, Cotter & Sherwin 2014), and the equations are linearised about this state. We emphasise that no time-averaged flow is used in this study: all linearised and adjoint simulations evolve about an unstable steady base flow, which is a solution of the Navier–Stokes equations. The time-continuous adjoint equations are then derived from the spatially discretised vorticity equations. The linearised and associated adjoint equation solvers are based on the same unstable base flow and both use the same discretisation and time-marching schemes as the nonlinear equations.

The flow around a D-shaped body is used as a test problem in this article. The body is in the shape of a half-ellipse with a blunt vertical base. The length of the body is twice as long as its height (see figure 1). The Reynolds number $Re = U_\infty H / \nu$ (where U_∞ is the incoming flow velocity, H is the body height and ν is the kinematic viscosity) is 80, corresponding to an unstable flow. A time step of 0.005 convective time units was used; the finest grid is of dimensions $30H \times 10H$ and contains 1500×500 grid points. Three nested grids were used in total. The trailing edge of the body is located at the centre of the finest grid since the ‘flow’ is advected upstream in the adjoint simulations. Note that, although the body shape is different, the grid used is the same as in Tu & Rowley (2012) and Flinois *et al.* (2015), where the flow over a cylinder was considered in both cases.

In order to measure the fluctuations in the wake, two sensor configurations are considered. The first measures the vertical velocity at a point located along the

symmetry plane, two body heights downstream of the trailing edge. The second is a body-mounted, distributed sensor that measures the antisymmetric component of the force on the base, as illustrated in figure 1. The first sensor is used in §3 to demonstrate the modelling procedure, and both sensors are used in §4, in order to compare the performance of the resulting closed-loop systems. The flow is actuated using two disk-shaped horizontal body forces of diameter $0.1H$ and centred $0.15H$ above and below the trailing edge corners, directly in line with the base. The two forces are of equal magnitude but opposite sign – thus this is a SISO system. The input forces are acting purely in the x direction and the total force on the flow is always zero but with a non-zero antisymmetric component.

3. Comparison between modelling approaches

In this section, three different approaches are used to obtain ROMs for the input–output dynamics between the actuator and the wake velocity sensor described above. First, the linearised forward and adjoint impulse responses are computed and the projection-free snapshot-based balanced truncation (BPOD) approach is used to obtain a ROM. We use this as our starting point, as this method has been shown (Flinois *et al.* 2015) to yield balanced ROMs for unstable systems. Second, the ERA is applied to the linearised impulse response and we show that we obtain a very similar model to the BPOD-based ROM. Finally, we show that another effectively identical model can be obtained for the linearised dynamics about this steady state, simply by considering the early response of the nonlinear system to a small impulsive input. This is a key point, as it shows that it is possible to obtain ERA models for unstable systems with standard (nonlinear) computational fluid dynamics codes and even potentially with experiments.

3.1. Unstable steady state and unforced flow

The flow over the D-shaped body considered in this article was also studied computationally by Palei & Seifert (2007), who identified the critical Reynolds number of the unforced flow to be $Re_c \approx 63$. As the simulations in the present work focus on $Re = 80$, only one complex conjugate pair of eigenvalues is expected to be unstable, since this is only slightly higher than the critical value (it will become apparent that this is indeed the case in later sections of this article). In order to validate the mesh used for our simulations, the vortex shedding Strouhal number ($St = fH/U_\infty$, where f is the dimensional frequency) and the root mean square (r.m.s.) value of the lift coefficient fluctuations C'_L were compared to those found in Palei & Seifert (2007) for several Reynolds numbers. The results are summarised in table 1. The agreement is very good, and this provides a good indication that our grid is sufficiently well resolved at all the considered Reynolds numbers.

The unstable base flow at $Re = 80$, obtained using selective frequency damping (Åkervik *et al.* 2006; Jordi *et al.* 2014), is compared to the unforced fully developed vortex shedding flow field (also at $Re = 80$) in figure 2. The length of the recirculation bubble in the base flow was found to be $x_{rec} = 3.58H$ and the drag coefficient was found to be $C_D = 1.138$. The mean drag coefficient of the unforced vortex shedding flow was found to be $C_D = 1.234$.

3.2. The BPOD model

In order to compute the balanced truncation ROM, the impulse response of the linearised forward and adjoint systems were computed (with the wake velocity




FIGURE 2. (Colour online) (a) Unstable base flow, obtained using selective frequency damping. (b) Snapshot of the unforced, fully developed limit-cycling flow. Vorticity contours are shown; negative contours are bounded by a line. Contour levels are from -1.8 to 1.8 in increments of 0.4 .




FIGURE 3. (Colour online) (a) First stored snapshot of the forward impulse response: horizontal velocity contours are shown as well as actuator locations as circles. (b) First stored snapshot of the adjoint impulse response: adjoint 'vertical velocity' contours are shown as well as the sensor location as a triangle in the wake. Negative contours are bounded by a line. Contour levels are from -0.09 to 0.09 in increments of 0.02 for panel (a) and from -0.9 to 0.9 in increments of 0.2 for panel (b).

Reynolds number (Re)	80	90	100	150	
Lift coefficient r.m.s. (C_L)	0.084	0.109	0.131	0.240	Present study
	0.082	0.107	0.130	0.236	Palei & Seifert (2007)
Strouhal number (St)	0.142	0.150	0.156	0.180	Present study
	0.144	0.151	0.158	0.182	Palei & Seifert (2007)

TABLE 1. R.m.s. of lift fluctuations and vortex shedding Strouhal number obtained in the present study and by Palei & Seifert (2007) for the considered D-body geometry at a range of Reynolds numbers.

sensor). State snapshots were stored every 0.2 convective time units (40 time steps) and the corresponding ROM was computed (see § A.1 for more details). The first stored snapshot of each simulation is shown in figure 3.

3.2.1. Actuator and sensor placement

As discussed for instance in Bagheri *et al.* (2009c) and Chen & Rowley (2011), a feedback controller can only stabilise an unstable flow if all unstable modes are

controllable and observable using the chosen actuators and sensors, respectively. In other words, the spatial location of the actuators and sensors must overlap with the spatial distribution of the adjoint and forward global modes, respectively. Furthermore, the flow structures that are generated with the least amount of input energy by the actuators are given by the leading eigenmodes of the controllability Gramian. It is therefore important to ensure that sensors can measure the state of the system where the magnitude of these modes is large. Conversely, the flow structures that lead to the most energetic output signal correspond to the leading eigenmodes of the observability Gramian. It is therefore important to ensure that actuators can influence the state of the system where the magnitude of these modes is large. Finally, the leading balanced modes inform us about the states that contribute the most to the input–output dynamics of the flow.

In the case of unstable systems, it was shown in Flinois *et al.* (2015) that, using projection-free BPOD (and assuming the unstable modes are observable and controllable), the direction of the unstable balanced modes tends to the direction of the unstable eigenmodes of the Gramians as $t_\infty \rightarrow \infty$. Additionally, the subspace spanned by these modes converges to the subspace spanned by the unstable global modes (note that the modes themselves are not identical in general in this case). Physically, this corresponds to the fact that the behaviour of unstable modes is eventually dominant for an uncontrolled linear system, as the influence of any initial transient growth becomes negligible as $t_\infty \rightarrow \infty$. In figure 4, we therefore show the first two primal and adjoint balanced modes, identified with a large final simulation time of $t_\infty = 100$ convective time units. Note that the ERA and most other system identification methods do not usually provide this type of full-state or adjoint information.

Figure 4(c,d) confirms that there is an overlap between the chosen actuators and the unstable modes and figure 4(a,b) confirms that the location of the vertical velocity sensor in the near wake will allow it to measure the growth of the unstable modes. Figure 4(a,b) also shows that, qualitatively, the unstable modes are dominant in the wake, downstream of the body. This suggests that the base (rear face) of the body is a justified location for a body-mounted sensor.

However, the leading global modes and Gramian eigenmodes do not provide sufficient information to select actuators and sensors that guarantee good robustness and performance in a closed-loop setting, especially in the presence of disturbances (Chen & Rowley 2011). On the one hand, considering global modes individually (even unstable ones) is not adequate for highly non-normal systems since the input–output dynamics may be strongly dependent on the interaction between several modes (Sipp *et al.* 2010). On the other hand, Gramian eigenmodes decouple the effect of actuators and sensors: the controllability Gramian is computed without any knowledge of the sensor matrix \mathbf{C} , while the observability Gramian is computed without any knowledge of the actuator matrix \mathbf{B} . Gramian eigenmodes thus cannot take into account the effect of time delays between actuation and sensing, which can have a large impact on the performance of feedback control (Chen & Rowley 2011). In the case we are considering here, placing the velocity sensor further downstream of the body than $x = 2$ would provide a better observability of the unstable modes, as indicated by figure 4(a,b). However, in practice, it was found that the design of robust stabilising controllers became increasingly difficult as the sensor was shifted further downstream.

In order to find a compromise between these conflicting requirements, Chen & Rowley (2011) suggested that a ‘structural sensitivity’ analysis, resulting in the




FIGURE 4. (Colour online) Balanced modes associated with unstable dynamics. (a,b) Primal modes, shown as vertical velocity contours. The vertical velocity sensor in the wake is shown as a triangle and the body-mounted force sensor is shown as a line along the base. (c,d) Adjoint modes, shown as adjoint ‘vorticity’ contours. The actuators are shown as circles. Negative contours are bounded by a line. Contour levels are from -0.09 to 0.09 in increments of 0.02 for panels (a,b) and from -9 to 9 in increments of 2 for panels (c,d).

‘wavemaker’ region introduced by Giannetti & Luchini (2007), might provide a promising starting point for actuator–sensor placement in a closed-loop setting. This region is identified by taking the pointwise product of the forward and adjoint global modes. Using projection-free BPOD, however, the unstable balanced modes only converge to a subspace spanned by the global modes. We therefore show the average of the two wavemaker regions predicted by the unstable modes in order to estimate the actual wavemaker of the system in figure 5. It is encouraging that the identified region is in good qualitative agreement with the one resulting from the structural sensitivity analysis performed by Giannetti & Luchini (2007) for the circular cylinder at similar Reynolds numbers. Figure 5 shows that the wake sensor is at the centre of the wavemaker region and also predicts that sensors located further downstream than $x \approx 4$ – 5 may not perform well. Figure 5 also shows that the actuators and the force sensor are located as close as possible to the region where the structural sensitivity is high, while remaining ‘body-mounted’.

3.2.2. Reduced-order model

As explained in Flinois *et al.* (2015), it is necessary to select an adequate final simulation time t_∞ in order to strike a balance between the convergence of the model and the exponential growth of the unstable modes. Four final simulation times were tested of 12.4, 25, 50 and 100 convective time units. Hankel singular values (HSVs) are a measure of how significant different states are in the BPOD ROM (see § A.1 for further details) and can be used to choose the required order for the model. HSV distributions for the first 20 singular values are shown in figure 6. As expected, the singular values of the two unstable modes grow as the final simulation




FIGURE 5. (Colour online) Approximation of the wavemaker region in the D-shaped body wake, constructed by averaging the product of the velocity magnitudes of the forward and adjoint unstable balanced modes for the two unstable modes. Contour levels are from 0.01 to 0.06 in increments of 0.01.

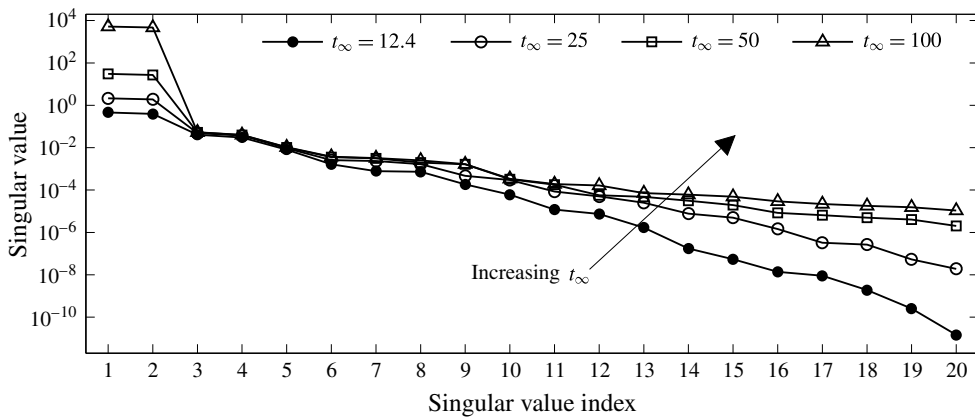


FIGURE 6. Hankel singular values corresponding to ROMs created using BPOD with different final simulation times t_∞ (given in convective time units).

time t_∞ increases, while the other singular values converge. HSVs of low dynamical importance are also affected by t_∞ , as described in Flinois *et al.* (2015): as the unstable modes start to dominate the response, information related to these less significant modes is lost. The corresponding singular values then increase as a whole with t_∞ as opposed to converging.

For stable systems, there exists an upper bound on the \mathcal{H}_∞ norm of the difference between the full-order model G_n (where n is the full system order) and the reduced-order model G_r (of order r) obtained with balanced truncation (e.g. Rowley 2005): $\|G_n - G_r\|_\infty < 2 \sum_{j=r+1}^n \sigma_j$, where σ_j is the j th HSV of the system. This bound holds for exact balancing procedures but can in practice be observed by using a sufficient number of snapshots in the BPOD approach. In the case of unstable systems, the \mathcal{H}_∞ error norm is infinite by definition. An upper bound on the \mathcal{H}_∞ error of the system's stable subspace can nevertheless be recovered if a projection approach is used to first identify and project out the unstable subspace. This can be done either by first identifying the unstable modes as in Barbagallo *et al.* (2009) and Ahuja & Rowley (2010), or by using the extension to the projection-free BPOD approach suggested in Flinois *et al.* (2015), whereby a first set of impulse responses is computed to identify the unstable balanced modes only. This is not performed in the present work.

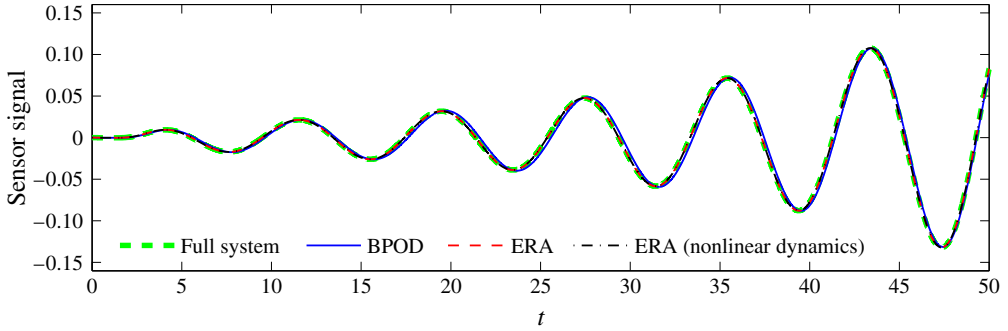


FIGURE 7. (Colour online) Impulse response of the full system and the ROMs obtained with BPOD, with the ERA based on the linearised dynamics, and with the ERA based on the nonlinear dynamics. Note that all four lines overlap almost perfectly.

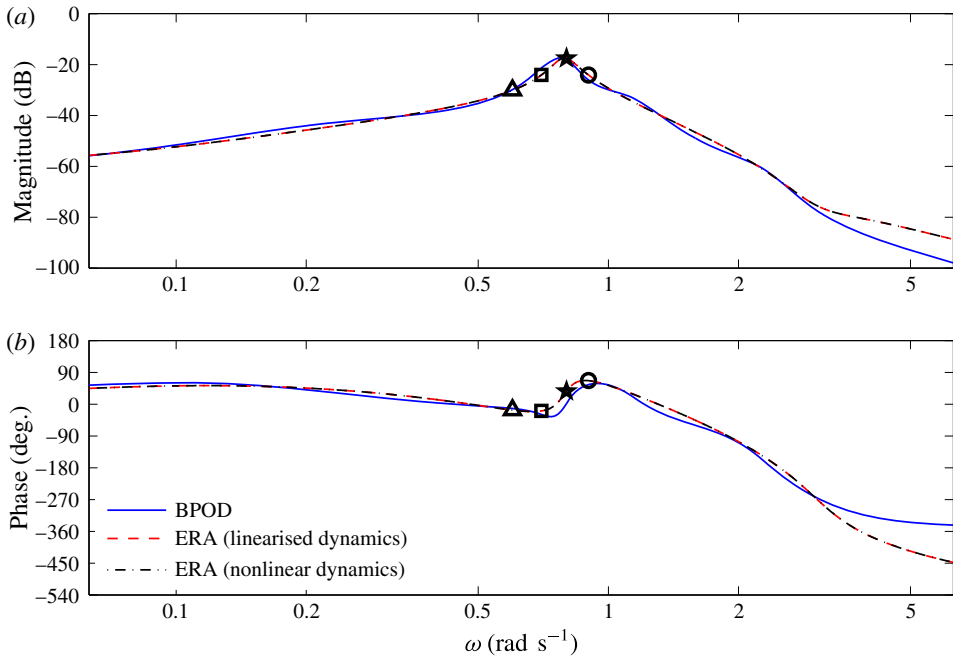


FIGURE 8. (Colour online) Bode diagram showing the transfer functions of the ROMs obtained with BPOD, with the ERA based on the linearised dynamics, and with the ERA based on the nonlinear dynamics. Note that the curves corresponding to the two ERA models overlap almost perfectly. The locations corresponding to frequencies of $\omega = 0.6 \text{ rad s}^{-1}$ (Δ), 0.7 rad s^{-1} (\square), 0.8 rad s^{-1} (\star) and 0.9 rad s^{-1} (\circ) are also shown to enable comparisons with the Nyquist diagrams in figure 12.

In the present case, it was found by trial and error that an accurate 10th-order ROM could be obtained with only $t_\infty = 25$ convective time units. In figure 7, we show the impulse response of the model as a thin solid line, superimposed on the full impulse response (thick dashed line). Clearly, the long-term response is well very predicted, even for much longer times than the final simulation time used to identify the model. The transients are also well approximated by the model. Figure 8 is a

Bode plot showing the transfer function of the ROM as a thin solid line. It can be used to design controllers using classical or \mathcal{H}_∞ loop-shaping methods. This plot is also useful to compare this model with the ROMs obtained in §§ 3.3 and 3.4 in the frequency domain.

3.3. The ERA model

The construction of ERA models is computationally very efficient as it relies only on the sensor measurement of the impulse response of the forward linearised equations. The procedure for obtaining such models is described in § A.2. As with BPOD, this method was originally designed exclusively for stable systems since the output becomes unbounded as $t_\infty \rightarrow \infty$ if the system is unstable. However, as shown by Ma *et al.* (2011) and summarised in § A.2, the ERA and BPOD are essentially the same algorithm and can thus theoretically yield exactly identical models. Furthermore, as shown in Flinois *et al.* (2015), the snapshot-based BPOD method also leads to balanced ROMs for unstable systems. As a result, the ERA must also lead to balanced ROMs, even for unstable systems.

In order to construct a ROM that is comparable to the BPOD ROM of § 3.2, every 40th sensor measurement and a final simulation time of $t_\infty = 50$ convective time units were used (this matches the size of the Hankel matrices on which the two models are based, as shown in §§ A.1 and A.2). The computational cost for obtaining the ROM is negligible, since we are simply stacking the output signal from the sensor into two matrices and performing a singular value decomposition of a (125×125) matrix. Much larger Hankel matrices could thus be used at virtually no additional expense if required.

Using again a 10th-order model, the ERA ROM impulse response and transfer function are compared with the BPOD ones in figures 7 and 8, respectively (dashed lines). Clearly there is a good match between the two models, in both the time and frequency domains. We provide a quantitative measure of how similar the two models are in § 4.

The models are not exactly identical, mainly due to small numerical and algorithmic differences in their implementation. The adjoint simulations are not the exact discrete adjoint of the forward simulations due to the multi-grid solver, which is not self-adjoint despite being used in both solvers, as noted by Ahuja & Rowley (2010). Additionally, the continuous version of BPOD was used in the present work and the forward and adjoint responses were not computed exactly in the same manner: in the adjoint simulation, the initial state was set to $\mathbf{z}(1) = \mathbf{C}^\dagger$ and this state is shown in figure 3(b). On the other hand, in the forward simulation, the \mathbf{B} matrix was approximated by a one-step pulse at the first time step of the forward simulation. This was done partly for convenience and partly because the same simulation was used for the BPOD and ERA models, where we attempt to keep the procedure pseudo-experimental. The flow state, just as the pulse is applied, is shown in figure 3(a) in this case.

3.4. ERA model based on nonlinear dynamics

3.4.1. Obtaining the reference model

Using the ERA with unstable systems as above allows balanced ROMs to be obtained for unstable systems, without the need for an adjoint solver. However, in many relevant flow control scenarios, the linearised impulse response cannot be computed directly either. This is the case if a linear solver is not available or in

experimental set-ups. As a third approach, therefore, we show that ERA models can be readily obtained directly from the nonlinear code: if the impulse is sufficiently small, the early response of the system is approximately linear until disturbances grow enough for nonlinear effects to become significant. It is therefore theoretically possible to obtain an arbitrarily long (approximately) linear response by imposing a correspondingly small impulse. Of course, in any realistic situation, having an excessively small impulse will result in significant issues due to the associated low signal-to-noise ratio of the output signal. In practice, an excessively long impulse response is not desirable anyway for unstable flows, as t_∞ is limited by the exponential growth of the unstable modes.

Nevertheless, if a sensor measurement can be recorded with a large enough portion in the linear regime and an acceptable signal-to-noise ratio, it is straightforward to scale the response back to recover the output that would result from a full impulse applied to the linearised system. The results are shown again in figures 7 and 8 as dot-dashed lines for a ROM obtained with exactly the same procedure as in §3.3, but scaled by a factor of 400 to retrieve the full-sized signal from the small impulse response of the nonlinear flow. The two ERA models are indistinguishable in both the time domain and the frequency domain. The model obtained in this section was thus used as the nominal or reference model for controller design in §4.

3.4.2. Effect of nonlinearity

In a linear setting, the effect of letting $t_\infty \rightarrow \infty$ on the quality of balanced models was investigated in Flinois *et al.* (2015). In a nonlinear setting, the appearance of a limit cycle (or other nonlinear effects) may impose a more stringent upper bound on t_∞ than the exponential growth of unstable global modes if the applied impulse is large. It is therefore relevant to examine how nonlinearity affects the ERA models. To this end, the system is subjected to a large impulse (200 times greater than the one used to obtain the model above) in order to promote the saturation of the unstable modes. Figure 9(a) shows the evolution of the sensor signal in both the linear and nonlinear simulations.

The main purpose of obtaining reduced-order models in this study is to use them to design stabilising feedback controllers for the flow. The models we obtain are therefore only useful if they are able to accurately represent the behaviour of the unstable subspace of the actual system (this is a necessary but not always a sufficient condition). The influence of the final simulation time on the extent to which this requirement is satisfied is shown in figure 9(b,c): recalling that the considered linearised system has two unstable modes, we analyse how the real and imaginary parts of the poles of the second-order ERA model obtained from the nonlinear system's response (to the large impulse) are affected by t_∞ . For approximately the first 50 convective time units, both the real and imaginary parts of the poles converge to the values obtained with the linearised impulse response. For larger values of t_∞ , the growth rate of the poles tends to zero and the frequencies tend to the limit-cycle frequency. Note that the maximum acceptable t_∞ value depends on the size of the impulse applied to the system.

In figure 10, we compare the transfer function gain of the ERA models from figure 9(b,c) for several t_∞ values to the 'reference' transfer function gain, corresponding to the 10th-order model described above (obtained with a small impulse) and used in §4 for controller design. At $t_\infty = 20$ convective time units the model is not yet fully converged. The approximately converged second-order model at $t_\infty = 45$ convective time units and the reference model are in reasonable agreement in the




FIGURE 9. (Colour online) Effect of nonlinearity on the quality of ERA models. (a) The impulse response of the linearised system (red dashed line) and response of the nonlinear system, subject to a large impulse (black solid line). (b) Real part/growth rate and (c) imaginary part/frequency of the poles of the second-order ERA model identified from the nonlinear impulse response with a range of final simulation time t_∞ . In panels (b) and (c), the growth rate and frequency of the converged linearised system poles (dashed line) and of the saturated limit cycle (dot-dashed line) are also shown. Note that the two poles are a complex conjugate pair, so they have opposite frequencies and identical growth rates.

frequency range of the unstable modes. As t_∞ is further increased, the frequency of the peak tends to the limit-cycle frequency and it becomes increasingly sharp as the growth rate tends to zero.

As mentioned above, if the system's unstable poles are not accurately identified by the model, it cannot be expected to be of satisfactory quality for controller design. Figure 9 shows that the poles move away from their 'converged' values shortly after the linear and nonlinear responses begin to differ. Unlike in the linear setting, the saturation of unstable modes affects the (dominant) unstable modes first as opposed to the least significant modes, so the loss of accuracy of the models can be expected to be more sudden in this case, even if a higher-order model is used. This therefore suggests that the ERA should only be applied in the approximately linear portion of the impulse response in order to obtain a useful model for controller design purposes. As the maximum acceptable t_∞ value for a given impulse response can be identified directly from figure 9, it is relatively straightforward to enforce this condition.




FIGURE 10. (Colour online) Gain of the transfer function of the second-order ERA models identified from the response of the nonlinear system, subject to a large impulse, with $t_\infty = 20$ (red thin dashed line), $t_\infty = 45$ (red thin solid line), $t_\infty = 100$ (red thick dashed line) and $t_\infty = 200$ (red thick solid line) convective time units. The gain of the ‘reference’ 10th-order model (black thin solid line), the frequency of the converged linearised system poles (black vertical dashed line) and the frequency of the saturated limit cycle (black vertical solid line) are also shown.

4. Feedback control

In this section, we use the (linear) model obtained from the nonlinear impulse response in § 3.4 (converted to continuous time) to design linear controllers using \mathcal{H}_∞ loop-shaping (§ 4.1) and proportional control (§ 4.2). We also show how robust control tools can be used to analyse model uncertainty and controller robustness in § 4.3. Although convenient numerically, velocity sensors in the middle of the wake are not always realistic. In § 4.4, we therefore also design controllers using the body-mounted force sensor on the base of the body (as described in § 2) and compare the results obtained with the two sensor configurations in § 4.5.

The block diagram for the control set-up is shown in figure 11. It consists of a linear system, whose input signals $u(s)$ (s is the Laplace variable) and output signals $y(s)$ are related by the model $G(s)$, the system’s open-loop transfer function: $y(s) = G(s)u(s)$. In closed loop, the measurements $y(s)$ are used to automate the actuation $u(s)$ using a mathematical law $K(s)$. This controller can be designed for instance to improve the system’s tracking performance or disturbance rejection capabilities. In the case of an unstable $G(s)$, the controller can be used to robustly stabilise the system. The full closed-loop transfer function from $r(s)$ to $y(s)$ is

$$G_{CL}(s) = \frac{y(s)}{r(s)} = \frac{G(s)}{1 - G(s)K(s)}. \quad (4.1)$$

The modelling procedure from the previous sections makes it possible to design controllers for unstable systems in a systematic way, without prior flow stabilisation. Note also that the entire procedure only requires a standard nonlinear solver: the linear model is obtained from an approximately linear impulse response, the controller is designed using this model, and it is applied directly to the nonlinear flow. Although the controllers are also applied to the linearised Navier–Stokes equations in § 4.5.1, this is only done to compare their response with theoretical predictions obtained with the ROM, and hence this step is not strictly necessary.




FIGURE 11. Block diagram of the input–output arrangement. Note that a positive feedback sign convention is used here.




FIGURE 12. (Colour online) Controller design for the wake velocity sensor, showing the Nyquist diagram of the (linear) ERA ROM, based on the nonlinear dynamics G (scaled by a factor of -21), the manually shaped model $(-WG)$, the robust model $(-K_0G)$ and the reduced robust model $(-KG)$. The point with coordinates $(-1, 0)$ is shown by a cross. The locations corresponding to frequencies of $\omega = 0.6 \text{ rad s}^{-1}$ (Δ), 0.7 rad s^{-1} (\square), 0.8 rad s^{-1} (\star) and 0.9 rad s^{-1} (\circ) are also shown. Note that the minus signs are due to the positive feedback sign convention.

4.1. \mathcal{H}_∞ loop-shaping

As mentioned in § 1.1, the method chosen here is a two-step procedure. The end goal is to obtain a robust controller that stabilises the system despite noise and/or disturbances, nonlinearities and model inaccuracies. Ideally, we would like to be able to stabilise the system from an off-design state like the unforced vortex shedding wake or at increased Reynolds numbers. \mathcal{H}_∞ loop-shaping provides a framework for obtaining this type of robust controller (McFarlane & Glover 1992).

The \mathcal{H}_∞ loop-shaping methodology requires the user to first manually ‘shape’ the gain of the open-loop frequency response or transfer function using dynamic weights

$W(s)$ that ensure the resulting response gain has the desired properties. In the second step, an optimisation problem is solved (using MATLAB), which generates a controller $K_\infty(s)$ that modifies mainly the phase of the shaped plant in order to make the closed-loop system robust to a large class of disturbances. Finally, the two components are joined together to form the final controller: $K(s) = K_\infty(s)W(s)$. Further details are included in appendix B.

4.1.1. Manual loop-shaping step

For this system, the first iterative step was used to improve the robustness of the final controller as much as possible and to ensure the shaped model was stabilised in a desired manner. Without this step, it was found that \mathcal{H}_∞ optimisation often returned stabilising controllers that were themselves unstable. Although this is not necessarily an issue, it should be avoided if possible for practical reasons.

The result of this procedure is shown in figure 12 (where the minus signs are due to the positive feedback sign convention used in figure 11). The thick solid line shows the Nyquist diagram of the open-loop model $G(s)$, scaled by a factor of -21 (for reasons explained below and in §4.2). The Nyquist diagram is the locus of a transfer function in the complex plane for values of $s = i\omega$ and for all real frequencies ω . The thin dashed line is the Nyquist diagram of the shaped plant $-WG$. In this case, $W(s)$ is a first-order lag compensator.

The standard stability criterion for a system with two unstable poles is that, following the Nyquist curve from $\omega = -\infty$ to $\omega = +\infty$ rad s⁻¹, the point with coordinates $(-1, 0)$ should be ‘encircled’ twice in the anticlockwise direction for the controller to stabilise the plant. Figure 12 shows that this is the case for the scaled open-loop system, which shows that there exists a so-called ‘gain window’, where proportional control stabilises the ROM. We therefore also implement a proportional controller and compare its performance with the robust controller. The design of the proportional controller is discussed in §4.2. The shaped system also stabilises the ROM in this case, although we emphasise that this is only an intermediate step and that the aim of this initial procedure is to improve the robustness of the final controller iteratively and not to obtain the best possible controller using only classical loop-shaping.

4.1.2. Robustness optimisation step

The aim of the second step is to optimise the shaped system’s robustness, as quantified by the generalised stability margin b , defined formally in appendix B. If the controller does not stabilise the plant, then $b = 0$, and as the robustness increases, $b \rightarrow 1$. A minimum value of $0.2 \leq b \leq 0.3$ is usually considered to be acceptable. As implied by its name, b is a generalisation of the standard gain and phase margins used in classical control. It gives a measure of both the robust performance and robust stability characteristics of a given (shaped) model and is also applicable to multiple-input–multiple-output systems.

The shaped plant WG was therefore used as an initial condition for the \mathcal{H}_∞ loop-shaping algorithm. This resulted in a robust 11th-order controller $K_0(s) = K_\infty(s)W(s)$. High-order controllers are less easily implemented in practice and also less reliable, so it is desirable to keep the order of the controller as low as possible without excessively compromising the closed-loop performance. The K_∞ part of the controller was thus subsequently reduced using MATLAB’s `balreal` and `modred` commands, resulting in a final eighth-order controller $K(s) = K_{\infty,r}(s)W(s)$. The generalised stability margins for the optimised and reduced controllers are $b(WG, K_\infty) = 0.427$

and $b(WG, K_{\infty,r}) = 0.422$, respectively, which shows first that the robustness of the \mathcal{H}_{∞} controller is satisfactory, and second that the order reduction process only deteriorated this robustness very marginally.

4.2. Proportional control

In order to compare the \mathcal{H}_{∞} loop-shaping approach to the proportional control approach used in several classical studies such as Roussopoulos (1993) and Park, Ladd & Hendricks (1994), we also design a proportional controller: $u(s) = K_P y(s)$ if $r(s) = 0$, where K_P is simply a constant. As mentioned above, it was found that there exists a ‘gain window’, where proportional control can be expected to stabilise the ROM. In a similar manner to Illingworth *et al.* (2014), we use our ERA model to select an adequate K_P value. Here, we choose to maximise robustness, as quantified by $b(K_P G, 1)$, in order to show that this stability and performance parameter is useful even for very simple controller design strategies. The optimal value was found to correspond to $K_P \approx 21$, for which $b(K_P G, 1) = 0.180$. This value is smaller than $b(WG, K_{\infty,r}) = 0.422$, indicating that the \mathcal{H}_{∞} loop-shaping controller can be expected to be more robust than the proportional controller. This is further discussed in the next section.

4.3. Controller robustness and model uncertainty

In order to obtain the reference model that was used for controller design in §§4.1 and 4.2, we selected one of three modelling algorithms, a final simulation time of $t_{\infty} = 50$ convective time units and a model order of $r = 10$. In this section, we show that robust control tools can quantify the model uncertainty associated with these choices. A related approach was used by Jones *et al.* (2015) in a grid refinement study.

A useful tool that naturally fits in the \mathcal{H}_{∞} loop-shaping framework is the ν -gap metric introduced by Vinnicombe (1993) and denoted by δ_{ν} . It provides a measure of how differently two systems shaped with the same weights $W(s)$ can be expected to behave in closed loop. The ν -gap takes a value between 0 and 1 ($\delta_{\nu} = 0$ between identical models) and it is formally defined in appendix B. This metric can be used to provide stability guarantees for uncertain systems. Specifically, if $b(WG, K_{\infty}) > \delta_{\nu}(WG, WG_i)$ for given W and K_{∞} , and for a ‘perturbed’ model G_i , then the controller $K_0 = K_{\infty}W$ also stabilises G_i . For any perturbed model where $b(WG, K_{\infty}) < \delta_{\nu}(WG, WG_i)$, the stability of the perturbed closed-loop model can be checked by calculating $b(WG_i, K_{\infty})$.

As summarised in table 2, several ‘perturbed’ open-loop models G_i were constructed by varying the model order, final simulation time and modelling method used. For all models, the ν -gap metric corresponding both to the \mathcal{H}_{∞} loop-shaping weights $\delta_{\nu}(WG, WG_i)$ and to the proportional controller $\delta_{\nu}(K_P G, K_P G_i)$ were calculated and are shown in table 2.

The ν -gap values in table 2 demonstrate the advantage of using \mathcal{H}_{∞} loop-shaping: although the ν -gap associated with all models is similar for both controllers, $b(WG, K_{\infty}) > \delta_{\nu}(WG, WG_i)$ for all i , whereas this is not the case with the proportional controller. The robust controller is thus guaranteed to stabilise all the models considered here. On the other hand, it was necessary to verify that the proportional controller in fact does stabilise G_{BPOD} by checking that $b(K_P G_{BPOD}, 1) > 0$ since $b(K_P G, 1) < \delta_{\nu}(K_P G, K_P G_{BPOD})$. It was found that $b(K_P G_{BPOD}, 1) = 0.144$.

Table 2 also shows that models obtained with a larger final simulation time or order have a very small ν -gap so they can be expected to behave in a similar manner to the

Name	Modelling method	t_∞	r	$\delta_v(WG, WG_i)$	$\delta_v(K_p G, K_p G_i)$
G (ref.)	ERA (nonlinear)	50	10	0	0
G_{BPOD}	BPOD	50 (equiv.)	10	0.179	0.195
$G_{ERA,lin}$	ERA (linear)	50	10	6.54×10^{-6}	7.18×10^{-6}
$G_{t_\infty=25}$	ERA (nonlinear)	25	10	0.0280	0.0291
$G_{t_\infty=75}$	ERA (nonlinear)	75	10	2.11×10^{-4}	2.20×10^{-4}
$G_{r=6}$	ERA (nonlinear)	50	6	0.343	0.138
$G_{r=14}$	ERA (nonlinear)	50	14	1.82×10^{-3}	1.30×10^{-3}

TABLE 2. Parameters used to generate the considered family of ‘perturbed’ models G_i for the wake velocity sensor, and values of the corresponding ν -gap metric, calculated for both the robust controller weights and the proportional controller.

reference model, whereas larger differences appear with lower t_∞ and r values. This suggests that our chosen t_∞ and r values are allowing us to obtain a converged model and resolve most of the important dynamics of the input–output system. Similarly, the distance between the two ERA models is negligible, which confirms that the early impulse response of the nonlinear flow provides adequate information to obtain a useful ERA model. Finally, a non-negligible value of δ_v is obtained with G_{BPOD} . This indicates that the ROMs may be sensitive to the modelling approach in this case and that, with the chosen weights W and K_p , some differences can be expected between the responses of the two models in closed loop. Nevertheless, the model G_{BPOD} is still stabilised by both controllers.

We emphasise that these tests are all performed *a priori*, in the sense that no closed-loop simulation is required. They suggest that the full-order flow and the ROMs are likely to behave in a similar manner since the stability and performance of the closed loop are suitably robust to the modelling technique and model parameters. This analysis is equally instructive for controllers designed explicitly using robust control methods such as $K(s)$ and for much simpler controllers such as K_p .

4.4. Body-mounted force sensor: modelling and controller design

The same procedure was then followed using the body-mounted force sensor introduced in § 2: first a 10th-order ERA model was obtained by scaling the response of the full nonlinear flow field to a small impulse. Note that the two sensor signals were recorded at the same time so it was not necessary to run a new simulation to obtain this model. In fact, the measurements from an arbitrary number of candidate sensors can be stored simultaneously for a given actuator configuration. As a result, there is effectively no added cost associated with the construction and comparison of models obtained with a large number of sensors. An ERA model was therefore also obtained from the fully linear response using this sensor.

A robust stabilising controller was then designed by shaping the Nyquist diagram with a first-order lag filter and a first-order low-pass filter. In this case, we note that, at this intermediate stage in the robust controller design process, the chosen weights $W(s)$ in fact do not stabilise the open-loop model. The robustness of the shaped system WG was then improved using \mathcal{H}_∞ loop-shaping. The stability margins corresponding to the resulting 14th-order controller and the eighth-order reduced controller are $b(WG, K_\infty) = 0.535$ and $b(WG, K_{\infty,r}) = 0.526$, respectively.




FIGURE 13. (Colour online) Controller design for the body-mounted force sensor, showing the Nyquist diagram of the (linear) ERA ROM in series with the proportional controller ($-K_p G$), the manually shaped model ($-W G$), the robust model ($-K_0 G$) and the reduced robust model ($-K G$). The point with coordinates $(-1, 0)$ is shown by a cross. The locations corresponding to frequencies of $\omega = 0.6 \text{ rad s}^{-1}$ (Δ), 0.7 rad s^{-1} (\square), 0.8 rad s^{-1} (\star), and 0.9 rad s^{-1} (\circ) are also shown. Note that the minus signs are due to the positive feedback sign convention.

The corresponding Nyquist diagrams are shown in figure 13. Note that the robust shaped systems $K_0 G$ and $K G$ have a peak at $\omega \approx 0.78 \text{ rad s}^{-1}$ so their Nyquist diagrams have a large magnitude close to this frequency.

Using the force sensor configuration, it was noticed that the range of stabilising proportional gains was greatly reduced due to the fact that proportional control triggered instabilities at frequencies close to the time-stepping frequency. In order to circumvent this numerical issue, the proportional controller was first coupled with a simple second-order low-pass filter. The filter has a cutoff frequency of $\omega = 200 \text{ rad s}^{-1}$, and hence has no effect on the frequencies where physical phenomena take place in the flow. The final controller can therefore still be considered to be effectively a proportional controller. With this added filter, it was possible to identify a ‘gain window’ where a stability margin of $b(K_p G, 1) = 0.182$ was obtained by using a scaling factor of 555. The model shaped with this controller K_p is also shown in figure 13.

Table 3 shows the ν -gap corresponding to the robust controller weights W and the proportional controller K_p for all considered models for the force sensor. The results are similar to those in table 2, although the ν -gap values are overall slightly higher. As with the velocity sensor, $b(W G, K_{\infty, r}) > \delta_\nu(W G, W G_i)$ for all G_i with the robust controller. With the proportional controller, on the other hand, $b(K_p G, 1) < \delta_\nu(K_p G, K_p G_{t_\infty=25})$. The stability of the perturbed closed-loop model was thus checked and it was found that $b(K_p G_{t_\infty=25}, 1) = 0.202 > 0$.

This analysis again confirms that the closed loop is not excessively sensitive to the chosen modelling technique and model parameters. We can therefore confidently proceed to implement the controllers in the full-order flow.




FIGURE 14. (Colour online) Control using the wake velocity sensor connected in closed loop with the two considered controllers: (a) output and (b) input corresponding to the closed-loop impulse response predicted by the ROM from § 3.4 and compared with the full linearised impulse response at $Re = 80$. In all cases the system is subject to an impulse at $t = 0$ and the controller is switched on after $t = 50$ convective time units.

Name	Modelling method	t_∞	r	$\delta_v(WG, WG_i)$	$\delta_v(K_p G, K_p G_i)$
G (ref.)	ERA (nonlinear)	50	10	0	0
$G_{ERA,lin}$	ERA (linear)	50	10	3.89×10^{-3}	4.23×10^{-3}
$G_{t_\infty=25}$	ERA (nonlinear)	25	10	0.233	0.238
$G_{t_\infty=75}$	ERA (nonlinear)	75	10	2.27×10^{-3}	2.45×10^{-3}
$G_{r=6}$	ERA (nonlinear)	50	6	0.158	0.165
$G_{r=14}$	ERA (nonlinear)	50	14	0.0832	0.0755

TABLE 3. Parameters used to generate the considered family of ‘perturbed’ models G_i for the body-mounted force sensor, and values of the corresponding ν -gap metric, calculated for both the robust controller weights and the proportional controller.

4.5. Controller performance

Equipped with robust stabilising controllers, three increasingly demanding tests were performed for the four considered control configurations. The results are presented in this section.

4.5.1. Closed-loop response predictions

First, the closed-loop response predicted by each ROM was compared to that of the full linearised system: the flow was forced with an impulse and allowed to evolve for 50 convective time units, thus allowing the unstable modes to develop before the controllers are switched on. In figure 14 the response predicted by the ROM (velocity sensor) is compared to the full linearised system’s response in closed loop. With both controllers, the two responses match very closely, even after the controller is switched on. Note also that the proportional controller leads to a slower stabilisation of the flow in this case.

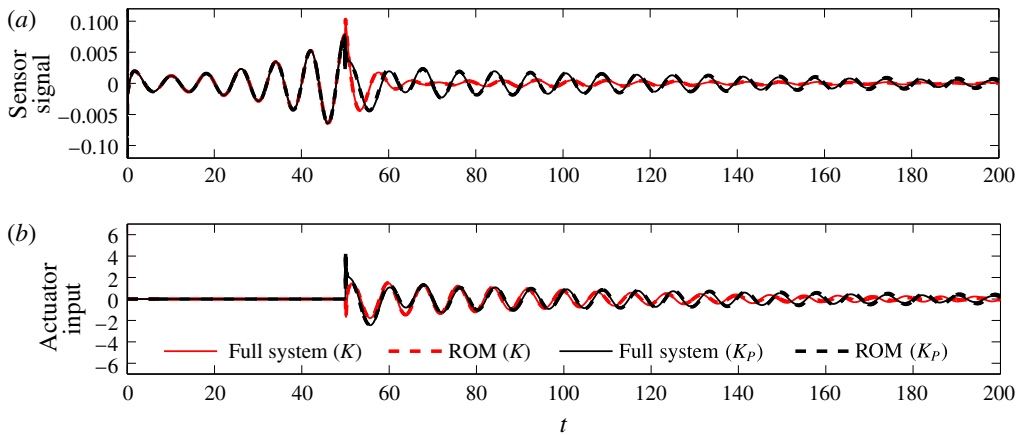


FIGURE 15. (Colour online) Control using the body-mounted force sensor connected in closed loop with the two considered controllers: (a) output and (b) input corresponding to the closed-loop impulse response predicted by the ROM (identified with the ERA based on nonlinear dynamics and with the body-mounted force sensor) and compared with the full linearised impulse response at $Re=80$. In all cases the system is subject to an impulse at $t=0$ and the controller is switched on after $t=50$ convective time units.

Similar conclusions can be drawn from the results obtained with the body-mounted force sensor, shown in figure 15. In this case, the convergence rate of the output signal is lower and it therefore takes longer for the flow to return to its unstable equilibrium state. On the other hand, the initial response of the controller is less aggressive than with the wake velocity sensor. Again, the stabilisation of the flow is faster with K than with K_p .

4.5.2. Robustness to nonlinear initial conditions

The robustness of the controllers to nonlinear effects was then tested. After triggering the instability with an impulse, the nonlinear flow was left to evolve to its limit-cycling (vortex shedding) state. The controllers were only switched on when vortex shedding was fully established (after 200 convective time units). We emphasise that neither the models nor the controllers incorporate any information about the limit-cycle dynamics, which do not even necessarily evolve about the same state as the base flow.

However, as shown in figures 16 and 17, the two proportional controllers and the two \mathcal{H}_∞ controllers are robust enough to stabilise the flow. Figure 18 shows the evolution of the forces on the body in all four configurations: starting from their base flow values, the instability is triggered and the forces evolve as the flow reaches the limit cycle. Once the controllers have been switched on, the forces quickly return to the base flow values.

In fact, the entire closed-loop flow returns to a state that is close to the low-drag unstable steady state in all cases. Figure 19 shows snapshots of the flow field (controlled using the robust controller associated with the body-mounted force sensor) at the moment when the control is switched on (after $t=200$ convective time units), and then as the flow is gradually stabilised. After 500 convective time units (figure 19d), the vortex shedding is effectively fully suppressed and the flow is similar to the base state in figure 2.




FIGURE 16. (Colour online) (a) Output and (b) input corresponding to the closed-loop response of the full nonlinear system at $Re = 80$, using the wake velocity sensor and the two corresponding controllers. The system is subject to an impulse at $t = 0$ and the controllers are switched on after $t = 200$ convective time units.




FIGURE 17. (Colour online) (a) Output and (b) input corresponding to the closed-loop response of the full nonlinear system at $Re = 80$, using the body-mounted force sensor and the two corresponding controllers. The system is subject to an impulse at $t = 0$ and the controllers are switched on after $t = 200$ convective time units.

For the wake velocity sensor, the convergence of the forces takes place at a similar rate with both controllers. As in the linear framework in § 4.5.1, the flow takes longer to converge to its steady state with the body-mounted force sensor configuration. However, in this case, the proportional controller leads to a faster flow stabilisation than the robust controller.

Overall, these results thus show that the designed controllers are robust to significant nonlinear effects for all four configurations.

4.5.3. Robustness to changes in the Reynolds number

The controllers were then challenged further by repeating the test from § 4.5.2 but with larger (and hence off-design) Reynolds number values of $Re = 90$ and 100. This introduces additional modifications in the dynamics of the flows that the controllers

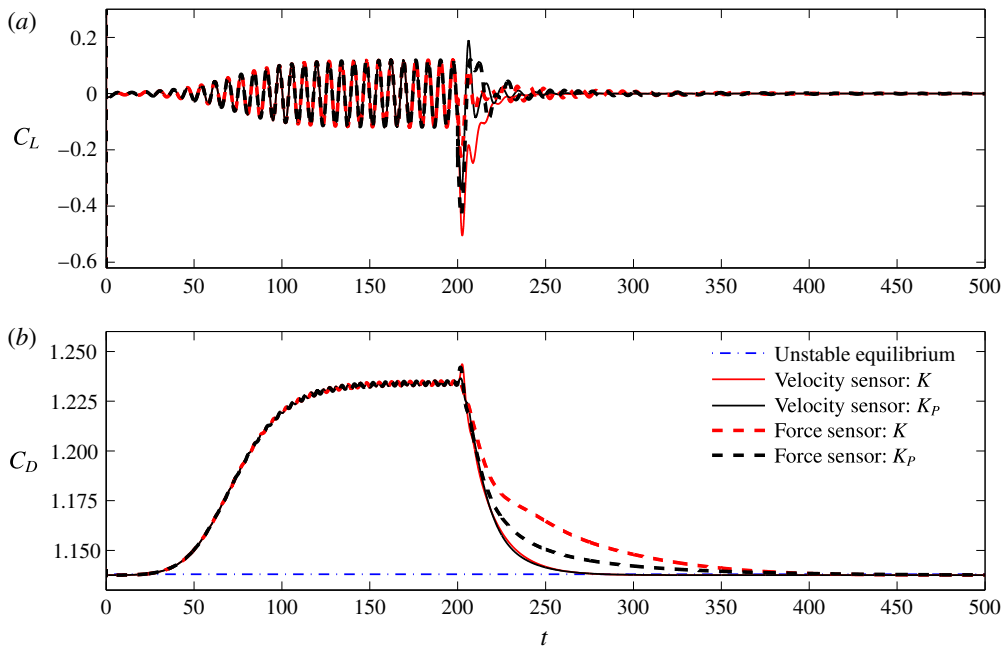


FIGURE 18. (Colour online) (a) Lift and (b) drag force coefficients corresponding to the closed-loop response of the full nonlinear system at $Re = 80$. The drag coefficient of the base flow is also shown in panel (b). The system is subject to an impulse at $t = 0$ and the controllers are switched on at $t = 200$ convective time units. The results obtained with all four sensor–controller configurations are shown.

are attempting to stabilise. In particular, the flow becomes more unstable at higher Re values and is hence harder to control.

It was found that, using the body-mounted force sensor, neither controller was able to fully stabilise the flow at $Re = 90$ or 100 . With the velocity sensor configuration (and with both of the associated controllers), the flow was fully stabilised from the limit cycle at $Re = 90$ and 100 . Further simulations were run at $Re = 110$, 130 and 150 : the flow was found to be stabilised from the limit cycle at $Re = 110$ with the robust controller and almost fully with the proportionally controller. At $Re = 130$, neither controller fully stabilised the flow.

Even in cases where the flow was not fully stabilised, the control was found to result in a positive net energy gain, as measured by the total average power required to drive the body through the fluid (after transients). The instantaneous total power can be expressed as the sum of the power required to overcome the drag $D(t)$ and the power supplied to the fluid by the actuators (see e.g. Frohnapfel, Hasegawa & Quadrio 2012). In non-dimensional form, this can be written as

$$\begin{aligned}
 \frac{P_{total}(t)}{\frac{1}{2}\rho U_\infty^3 H} &= \frac{D(t)U_\infty}{\frac{1}{2}\rho U_\infty^3 H} + \frac{P_{control}(t)}{\frac{1}{2}\rho U_\infty^3 H} \\
 &= C_D + \int_{A_F} \left[\frac{\mathbf{f}(t)}{\frac{1}{2}\rho U_\infty^2 H} \cdot \frac{\mathbf{u}(t)}{U_\infty} \right] dA_F,
 \end{aligned} \tag{4.2}$$




FIGURE 19. (Colour online) Snapshots of the controlled flow field (using the body-mounted force sensor and the robust controller) after (a) $t = 200$ convective time units when the control is turned on, and then at (b) $t = 300$, (c) $t = 400$, and (d) $t = 500$ convective time units. Vorticity contours are shown; negative contours are bounded by a line. Contour levels are from -1.8 to 1.8 in increments of 0.4 .

where $P_{total}(t)$ and $P_{control}(t)$ are the dimensional total and control power per unit width, respectively, ρ is the fluid density, A_F is the area where the forcing is applied, \mathbf{f} is the local dimensional force per unit area and \mathbf{u} is the local dimensional velocity. Note that, in practice, an actuator model would be desirable in order to obtain a more realistic estimate of the required input power.

In our case, the control power (second term) tends to zero when the flow is fully stabilised. Moreover, in all cases where the control input converges to a finite value, it was found that the time-averaged value of $P_{control}$ after transients was at most of the order of a hundredth of a per cent of P_{total} and usually much less. This negligible value is due to the fact that the flow is almost symmetric at the actuation location while the forcing is purely antisymmetric. As a result, changes in the converged drag coefficient can be directly linked to changes in the total required power. The results are shown in figure 20: clearly in all cases, the required time-averaged converged power is reduced significantly. For the wake velocity sensor, the robust controller leads to a great power reduction than the proportional controller. This general trend seems to be inverted for the body-mounted force sensor.

The results from this test suggest that the velocity sensor configuration is in general more robust to an increase in the Reynolds number than the force sensor configuration, despite the fact that the stability margin is larger with the body-mounted sensor. This implies that the dynamics of the shaped system experience larger changes with the force sensor than with the velocity sensor. As the Reynolds number is increased, Giannetti & Luchini (2007) have shown that the wavemaker region in a cylinder wake at low Reynolds numbers becomes elongated and shifts further downstream. Since the body-mounted sensor is located in a less observable region of the flow (see figure 4) and is already on the upstream edge of the wavemaker region at the design

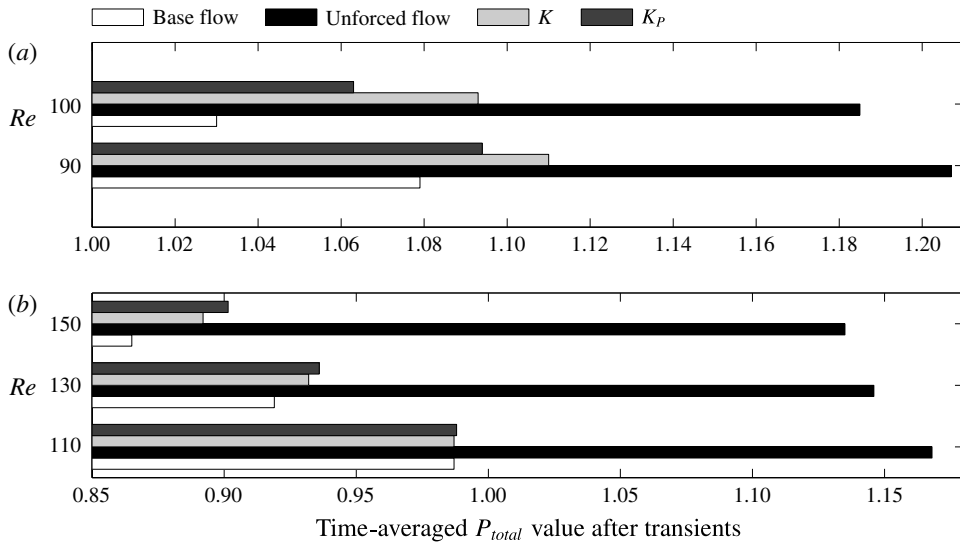


FIGURE 20. Comparison of the total power required to drive the body through the fluid in the unstable equilibrium flow state, the unforced vortex shedding state and with the two controllers, for cases where the flow is not fully stabilised with (a) the body-mounted force sensor and (b) the wake velocity sensor.

Reynolds number (see figure 5), this could explain why the associated input–output dynamics are more strongly affected by an increase in the Reynolds number.

4.5.4. Further comments

In §4.3, it was found that \mathcal{H}_∞ loop-shaping yields controllers that are expected to be considerably more robust than proportional controllers (as quantified by *b*). Additionally, the window of stabilising proportional gains eventually disappears for sufficiently large Re and in this case a dynamic controller is necessary to stabilise the flow (e.g. Illingworth *et al.* 2014). Nevertheless, the results from this section indicate that, in cases where a gain window does exist, choosing the proportional control gain that maximises the generalised stability margin can result in controllers whose robustness and performance are in practice similar to that of \mathcal{H}_∞ loop-shaping controllers. The generalised stability margin b therefore appears to be a useful parameter to guide the design of stabilising controllers for fluid flows in general. We note however that b should not be used blindly: although the force sensor robust controller provided in the largest stability margin, this system was not the most robust to an increase in the Reynolds number.

Compared to other studies where proportional and/or classical control was used (e.g. Roussopoulos 1993; Park *et al.* 1994), the availability of an accurate input–output model enables the use of b to select a proportional controller gain value that provides good robustness. This information is not available when tuning a controller manually.

It is also important to note that the controllers used in this study were designed in part manually. This iterative procedure is different for each actuator–sensor configuration and may be better optimised in this case with one of the sensor configurations than the other. As a result, the comparison between the two sensors cannot be made quantitatively.

Additionally, recall that the purpose of the simulations performed here was to test the robustness of fixed controllers as the Reynolds number takes on larger off-design values. The aim was therefore not to identify the highest value of the Reynolds number for which flow stabilisation is possible, as optimal control techniques would be better suited for such a study (see for instance Lauga & Bewley 2003).

5. Conclusions

This paper has shown that it is straightforward and justified to obtain linear low-order models for the dynamics of fluid flows about an unstable equilibrium state, by directly applying the ERA to the system's impulse response. The resulting models were shown to be capable of predicting the transient and long-term flow dynamics in open loop and in closed loop. They were also found to be sufficiently accurate to design robust stabilising controllers.

The method was applied to the unstable flow over a D-shaped body, with body-mounted actuators, and two sensor configurations. We first showed that, if projection-free BPOD is used to obtain the model as in Flinois *et al.* (2015), then the analysis of actuator–sensor placement can be facilitated by extracting relevant information from the balanced modes. We then confirmed that the ERA can also be used at a much reduced cost and without the need for an adjoint solver, in order to obtain theoretically identical models, directly from the unstable system's impulse response. In experiments or if a linearised solver is not available, it was shown that accurate models can also be obtained using the ERA, from the early, approximately linear, part of the impulse response of the full nonlinear flow. Such ERA ROMs, based on the nonlinear flow dynamics, were then used to design stabilising feedback controllers for two sensor configurations: one located in the wake and one mounted on the body surface.

Both proportional controllers and robust controllers based on the \mathcal{H}_∞ loop-shaping procedure of McFarlane & Glover (1992) were designed. In both cases, it was found that robust control tools such as the generalised stability margin and the ν -gap metric are particularly well suited to the design of robust stabilising controllers for unstable flows. These enable the analysis of the sensitivity of the closed-loop behaviour to the chosen modelling method and to model parameters such as the model order. They are central to the \mathcal{H}_∞ loop-shaping framework and were also found to provide an effective means to optimise the robustness of proportional controllers using the ERA models.

Comparing the closed-loop response of the full-order linearised flow to the predictions made by the models showed that the ERA ROMs were sufficiently accurate to design stabilising controllers for the actual flow and also capable of predicting its closed-loop behaviour accurately. Furthermore, all controllers led to the full stabilisation of the nonlinear flow from the (unmodelled) vortex shedding limit cycle. Finally, it was found that the wake sensor was more robust to an increase in the Reynolds number than the body-mounted sensor.

All the key steps followed throughout this study can be performed with a standard nonlinear flow solver. This is a crucial point, as linearised and adjoint solvers are not available in a large number of relevant flow control scenarios. It also suggests that some experimental set-ups may benefit from this approach. In order to obtain models for unstable systems, previous studies have relied either on expensive computational techniques or on the difficult task of finding a stabilising controller for the flow, without a model for the dynamics. In contrast, the ERA is straightforward to implement and has a low computational cost. Here, we have shown that it can be used directly to construct models of high enough quality to design and implement robust stabilising controllers for unstable flows.

Acknowledgements

This research is funded by the UK Engineering and Physical Sciences Research Council (EPSRC). The authors also thank Professor T. Colonius from the California Institute of Technology for the immersed-boundary fractional step code on which the simulations presented here are based. All readers who are interested to obtain the data used in this paper are invited to email t.finois11@imperial.ac.uk.

Appendix A. Eigensystem Realisation Algorithm for unstable flows

A.1. Projection-free approximate balanced truncation for unstable flows

This method assumes that the evolution of the system is linear-time-invariant (LTI) and can hence be described in state-space form. We will assume the system is discrete to make the comparison between this method and the ERA more straightforward in § A.2, but the problem can equivalently be formulated in continuous form. In fact, the continuous approach is the one implemented in the present work.

We thus assume the system is governed by:

$$\left. \begin{aligned} \mathbf{x}(k+1) &= \mathbf{A}\mathbf{x}(k) + \mathbf{B}\mathbf{u}(k), \\ \mathbf{y}(k) &= \mathbf{C}\mathbf{x}(k). \end{aligned} \right\} \tag{A 1}$$

Here $\mathbf{x} \in \mathbb{C}^{n_x}$ is the state vector, $\mathbf{u} \in \mathbb{C}^{n_u}$ is the input vector, $\mathbf{y} \in \mathbb{C}^{n_y}$ is the output vector and the realisation of the system is defined by the matrices $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ of appropriate dimensions. The dynamics of the associated adjoint system can also be expressed in state-space form:

$$\left. \begin{aligned} \mathbf{z}(k+1) &= \mathbf{A}^\dagger \mathbf{z}(k) + \mathbf{C}^\dagger \mathbf{v}(k), \\ \mathbf{w}(k) &= \mathbf{B}^\dagger \mathbf{z}(k). \end{aligned} \right\} \tag{A 2}$$

The \dagger superscript refers to the complex conjugate transpose and in this case $\mathbf{z} \in \mathbb{C}^{n_x}$ is the adjoint state vector, $\mathbf{v} \in \mathbb{C}^{n_y}$ is the adjoint input vector and $\mathbf{w} \in \mathbb{C}^{n_u}$ is the adjoint output vector. Note that, in general, the adjoint formulation is obtained based on a non-trivial inner-product matrix, but for clarity we assume the identity matrix is used here.

The realisation $(\mathbf{A}, \mathbf{B}, \mathbf{C})$ of the primal system in (A 1) is not unique. Any appropriate coordinate transformation defined by the matrices \mathbf{T} and \mathbf{S} such that $\mathbf{T} = \mathbf{S}^{-1}$ can be applied to it, resulting in the transformed realisation $(\mathbf{SAT}, \mathbf{SB}, \mathbf{CT})$, with unmodified input–output dynamics. The purpose of balanced truncation is to find the transformation that results in a transformed state vector $\hat{\mathbf{x}} = \mathbf{T}\mathbf{x}$, where the states are sorted in decreasing order of dynamical significance. The significance of a state in this context is affected both by the amount of input energy required to reach it from $\hat{\mathbf{x}} = \mathbf{0}$ (its controllability) and the energy of the output signal generated by letting the system evolve with no input from that state (its observability). If such a transformation is found, the least significant states can easily be discarded by simply truncating the state vector and system matrices.

Willcox & Peraire (2002) and Rowley (2005) introduced a snapshot-based procedure that is based on the work of Moore (1981) and results in an approximation of the required transformation. It provides a set of so-called Hankel singular values (HSVs) σ_i , sorted so that $\sigma_1 > \dots > \sigma_{n_x}$. They indicate how significant the corresponding state is. The algorithm is also referred to as balanced proper orthogonal decomposition (BPOD). The first step consists in storing state snapshots from the impulse response of

the forward system into a matrix \mathbf{X} and state snapshots from the impulse response of the adjoint system into another matrix \mathbf{Z} . Assuming that a snapshot is saved every p time steps and that $(N + 1)$ snapshots are saved in each case, the matrices are structured as follows:

$$\mathbf{X} = [\mathbf{x}(1) \quad \mathbf{x}(1+p) \quad \dots \quad \mathbf{x}(1+pN)] = [\mathbf{B} \quad \mathbf{A}^p \mathbf{B} \quad \dots \quad \mathbf{A}^{pN} \mathbf{B}], \quad (\text{A } 3)$$

$$\mathbf{Z} = [\mathbf{z}(1) \quad \mathbf{z}(1+p) \quad \dots \quad \mathbf{z}(1+pN)] = [\mathbf{C}^\dagger \quad \mathbf{A}^{\dagger p} \mathbf{C}^\dagger \quad \dots \quad \mathbf{A}^{\dagger pN} \mathbf{C}^\dagger]. \quad (\text{A } 4)$$

Here, $\mathbf{x}(k) = \mathbf{A}^{(k-1)} \mathbf{B} = [\mathbf{A}^{(k-1)} \mathbf{b}_1 \quad \dots \quad \mathbf{A}^{(k-1)} \mathbf{b}_{n_u}]$ for $k \geq 1$ is an $(n_x \times n_u)$ matrix and \mathbf{b}_i is the i th column of \mathbf{B} , so that $\mathbf{A}^{(k-1)} \mathbf{b}_i$ is the forward impulse response from the i th input. Similarly, $\mathbf{z}(k) = \mathbf{A}^{\dagger(k-1)} \mathbf{C}^\dagger$ is an $(n_x \times n_y)$ matrix and $\mathbf{A}^{\dagger(k-1)} \mathbf{c}_i^\dagger$ is the adjoint impulse response from the i th input. The forward (respectively adjoint) impulse response snapshots can be obtained either by letting the state evolve from the initial condition $\mathbf{x}(1) = \mathbf{b}_i$ ($\mathbf{z}(1) = \mathbf{c}_i^\dagger$) for each column of \mathbf{B} (\mathbf{C}^\dagger) or by starting the simulation at $\mathbf{x}(0) = \mathbf{0}$ ($\mathbf{z}(0) = \mathbf{0}$) and forcing the system with an impulsive input at $k=0$.

Once the snapshots have been collected, the following singular value decomposition (SVD) is performed to obtain the transformation matrices:

$$\mathbf{Z}^\dagger \mathbf{X} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\dagger = [\mathbf{U}_1 \quad \mathbf{U}_2] \begin{bmatrix} \mathbf{\Sigma}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{\Sigma}_2 \end{bmatrix} \begin{bmatrix} \mathbf{V}_1^\dagger \\ \mathbf{V}_2^\dagger \end{bmatrix}, \quad (\text{A } 5a)$$

$$\hat{\mathbf{T}} = \mathbf{X} \mathbf{V}_1 \mathbf{\Sigma}_1^{-1/2}, \quad \hat{\mathbf{S}} = \mathbf{\Sigma}_1^{-1/2} \mathbf{U}_1^\dagger \mathbf{Z}^\dagger. \quad (\text{A } 5b,c)$$

Here $\hat{\mathbf{T}}$ is an $(n_x \times r)$ matrix and $\hat{\mathbf{S}}$ is an $(r \times n_x)$ matrix. These matrices are used to simultaneously balance and truncate the system and thus obtain an r th-order ROM. The elements of the diagonal matrix $\mathbf{\Sigma}$ are the HSVs σ_i mentioned above. The singular values in $\mathbf{\Sigma}_2$ identified in (A 5a) are zero or negligible since more snapshots than there are dynamically significant states are usually stored. The number of truncated states can be chosen based on the HSV distribution. The transformed and truncated system is then given by

$$\left. \begin{aligned} \hat{\mathbf{x}}_1(k+1) &= \hat{\mathbf{A}} \hat{\mathbf{x}}_1(k) + \hat{\mathbf{B}} \mathbf{u}(k), \\ \hat{\mathbf{y}}(k) &= \hat{\mathbf{C}} \hat{\mathbf{x}}_1(k) \approx \mathbf{y}(k), \end{aligned} \right\} \quad (\text{A } 6)$$

where $(\hat{\mathbf{A}}, \hat{\mathbf{B}}, \hat{\mathbf{C}}) = (\hat{\mathbf{S}} \hat{\mathbf{A}} \hat{\mathbf{T}}, \hat{\mathbf{S}} \hat{\mathbf{B}}, \mathbf{C} \hat{\mathbf{T}})$, $\hat{\mathbf{x}}_1 \in \mathbb{C}^r$ is the ROM state and $\hat{\mathbf{y}} \in \mathbb{C}^{n_y}$ is the ROM output, which approximates the full system output \mathbf{y} .

This method is particularly appealing because it is applicable to even very large systems since only matrix multiplications and an SVD of an $(N + 1) \times (N + 1)$ matrix are required regardless of the actual dimension of the full system. However, it was originally designed for stable systems only: clearly, as $t \rightarrow \infty$ the state vector $\mathbf{x} \rightarrow 0$ if the system is stable. As a result, if snapshots are stored over a long enough period of time, the matrices $\hat{\mathbf{T}}$ and $\hat{\mathbf{S}}$ converge and so does the identified balanced ROM.

If the system is unstable, on the other hand, the state vector becomes unbounded for large t , so it is not immediately obvious that the transformation matrices converge, that they lead to a converged ROM or that this ROM is balanced. For this reason, a number of methods have been developed to circumvent this issue (e.g. Barbagallo *et al.* 2009; Ahuja & Rowley 2010; Dergham *et al.* 2011). They require additional computations that render the algorithm significantly more costly and often completely

intractable if the system dimension is very large (as for instance in three-dimensional flows).

In a recent study, Flinois *et al.* (2015) showed that the unmodified algorithm does in fact theoretically lead to a fully converged balanced ROM. In practice, finite-precision arithmetic cannot handle both the small but important transient information and the unbounded long-term growth of the response if snapshots are stored until a large final simulation time $t_\infty \rightarrow \infty$. As a result, if the final simulation time is too large, the information related to the stable modes is lost. It was therefore shown that a compromise must be found: on the one hand, t_∞ must be large enough to obtain a converged model; and on the other, it must be small enough to retain the information about stable modes. In Flinois *et al.* (2015), accurate models were obtained both for the unstable Ginzburg–Landau equation and for the unstable flow around a cylinder. This approach therefore provides a way of obtaining balanced ROMs even for potentially very large unstable systems, and it is used in § 3.2 of this article.

A.2. The Eigensystem Realisation Algorithm

The ERA was developed by Juang & Pappa (1985). In this method, the impulse response output of the system given in (A 1) is recorded and stored in the following two Hankel matrices:


$$\begin{aligned}
 \mathbf{H}_1 &= \begin{bmatrix} \mathbf{y}_1 & \mathbf{y}_{1+p} & \cdots & \mathbf{y}_{1+Np} \\ \mathbf{y}_{1+p} & \mathbf{y}_{1+2p} & \cdots & \mathbf{y}_{1+(N+1)p} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{y}_{1+Np} & \mathbf{y}_{1+(N+1)p} & \cdots & \mathbf{y}_{1+2Np} \end{bmatrix}, \\
 &= \begin{bmatrix} \mathbf{C}\mathbf{B} & \mathbf{C}\mathbf{A}^p\mathbf{B} & \cdots & \mathbf{C}\mathbf{A}^{Np}\mathbf{B} \\ \mathbf{C}\mathbf{A}^p\mathbf{B} & \mathbf{C}\mathbf{A}^{2p}\mathbf{B} & \cdots & \mathbf{C}\mathbf{A}^{(N+1)p}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{C}\mathbf{A}^{Np}\mathbf{B} & \mathbf{C}\mathbf{A}^{(N+1)p}\mathbf{B} & \cdots & \mathbf{C}\mathbf{A}^{2Np}\mathbf{B} \end{bmatrix}, \tag{A 7}
 \end{aligned}$$

$$\begin{aligned}
 \mathbf{H}_2 &= \begin{bmatrix} \mathbf{y}_2 & \mathbf{y}_{2+p} & \cdots & \mathbf{y}_{2+Np} \\ \mathbf{y}_{2+p} & \mathbf{y}_{2+2p} & \cdots & \mathbf{y}_{2+(N+1)p} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{y}_{2+Np} & \mathbf{y}_{2+(N+1)p} & \cdots & \mathbf{y}_{2+2Np} \end{bmatrix}, \\
 &= \begin{bmatrix} \mathbf{C}\mathbf{A}\mathbf{B} & \mathbf{C}\mathbf{A}^{(p+1)}\mathbf{B} & \cdots & \mathbf{C}\mathbf{A}^{Np+1}\mathbf{B} \\ \mathbf{C}\mathbf{A}^{(p+1)}\mathbf{B} & \mathbf{C}\mathbf{A}^{(2p+1)}\mathbf{B} & \cdots & \mathbf{C}\mathbf{A}^{((N+1)p+1)}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{C}\mathbf{A}^{(Np+1)}\mathbf{B} & \mathbf{C}\mathbf{A}^{((N+1)p+1)}\mathbf{B} & \cdots & \mathbf{C}\mathbf{A}^{(2Np+1)}\mathbf{B} \end{bmatrix}, \tag{A 8}
 \end{aligned}$$

where $\mathbf{y}_k = \mathbf{C}\mathbf{A}^{(k-1)}\mathbf{B}$ is an $(n_y \times n_u)$ matrix.

Once these matrices have been formed, we take the SVD of \mathbf{H}_1 . Noting that $\mathbf{Z}^\dagger\mathbf{X} = \mathbf{H}_1$, this results in $\mathbf{H}_1 = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\dagger$ as in (A 5a). The reduced system matrices are then given by

$$\left. \begin{aligned}
 \mathbf{A}_r &= \mathbf{\Sigma}_1^{-1/2}\mathbf{U}_1\mathbf{H}_2\mathbf{V}_1\mathbf{\Sigma}_1^{-1/2}, \\
 \mathbf{B}_r &= \text{first } n_u \text{ columns of } \mathbf{\Sigma}_1^{1/2}\mathbf{V}_1^\dagger, \\
 \mathbf{C}_r &= \text{first } n_y \text{ rows of } \mathbf{U}_1\mathbf{\Sigma}_1^{1/2}.
 \end{aligned} \right\} \tag{A 9}$$

FIGURE 21. Block diagram considered for \mathcal{H}_∞ loop-shaping.

In Ma *et al.* (2011), it was shown that the ROM obtained with both BPOD and the ERA are in fact identical. Indeed $\mathbf{Z}^\dagger \mathbf{A} \mathbf{X} = \mathbf{H}_2$, and hence

$$\begin{aligned} \mathbf{A}_r &= \boldsymbol{\Sigma}_1^{-1/2} \mathbf{U}_1 (\mathbf{Z}^\dagger \mathbf{A} \mathbf{X}) \mathbf{V}_1 \boldsymbol{\Sigma}_1^{-1/2} \\ &= \hat{\mathbf{S}} \hat{\mathbf{A}} \hat{\mathbf{T}} \\ &= \hat{\mathbf{A}}. \end{aligned} \quad (\text{A } 10)$$

Similarly, using (A 5a) we have $\boldsymbol{\Sigma}_1^{1/2} \mathbf{V}_1^\dagger = \boldsymbol{\Sigma}_1^{-1/2} \mathbf{U}_1^\dagger \mathbf{H}_1$ and since the first n_u columns of \mathbf{H}_1 are equal to $\mathbf{Z}^\dagger \mathbf{B}$, we have $\mathbf{B}_r = \hat{\mathbf{B}}$. Finally $\mathbf{U}_1 \boldsymbol{\Sigma}_1^{1/2} = \mathbf{H}_1 \mathbf{V}_1 \boldsymbol{\Sigma}_1^{-1/2}$ and again since the first n_y rows of \mathbf{H}_1 are equal to $\mathbf{C} \mathbf{X}$, we have $\mathbf{C}_r = \hat{\mathbf{C}}$.

This approach therefore also provides a way of obtaining balanced ROMs, but at an even lower computational cost, and without the need for an adjoint solver as it is only based on sensor measurements from as little as one impulse response simulation. This article focuses mainly on the ERA and models are obtained using this method in §§ 3.3 and 3.4.

Appendix B. \mathcal{H}_∞ loop-shaping

B.1. The generalised stability margin

In \mathcal{H}_∞ loop-shaping, the control system shown in figure 21 is considered. Here, G_w is a shaped linear model, with input and output disturbances d_1 and d_2 , respectively; K_∞ is a controller, which is subject to input and output noise sources n_1 and n_2 , respectively. The nominal inputs and outputs to the system are u and y , respectively. Note that a positive feedback sign convention is used here.

The transfer functions from the disturbance and noise sources to the input and output signals are given by

$$\begin{pmatrix} u \\ y \end{pmatrix} = \begin{bmatrix} K_\infty \\ \mathbf{I} \end{bmatrix} (\mathbf{I} - G_w K_\infty)^{-1} [G_w \quad \mathbf{I}] \begin{pmatrix} d_1 \\ d_2 \end{pmatrix} + \begin{bmatrix} \mathbf{I} \\ G_w \end{bmatrix} (\mathbf{I} - K_\infty G_w)^{-1} [K_\infty \quad \mathbf{I}] \begin{pmatrix} n_1 \\ n_2 \end{pmatrix}, \quad (\text{B } 1)$$

where \mathbf{I} is the identity matrix. The generalised stability margin of the feedback configuration $[G_w, K_\infty]$ is defined as the inverse of the \mathcal{H}_∞ norm of the second set of transfer functions if $[G_w, K_\infty]$ is internally stable (i.e. if all transfer functions from n_1 and n_2 to u and y are stable) and to 0 if not:

$$b(G_w, K_\infty) = \begin{cases} \left\| \begin{bmatrix} \mathbf{I} \\ G_w \end{bmatrix} (\mathbf{I} - K_\infty G_w)^{-1} [K_\infty \quad \mathbf{I}] \right\|_\infty^{-1} & \text{if } [G_w, K_\infty] \text{ is stable,} \\ 0 & \text{otherwise.} \end{cases} \quad (\text{B } 2)$$

The generalised stability margin can only take values between 0 and 1. Maximising $b(G_w, K_\infty)$ can be shown (Vinnicombe 2001; Jones *et al.* 2015) to maximise the robust stability of the system with respect to perturbations to the normalised coprime factors of the model, which can be used to represent a large set of realistic model uncertainties. In fact, it can be shown (Vinnicombe 2001) that the norms of both sets of transfer functions are equal:

$$\left\| \begin{bmatrix} K_\infty \\ I \end{bmatrix} (I - G_w K_\infty)^{-1} \begin{bmatrix} G_w & I \end{bmatrix} \right\|_\infty = \left\| \begin{bmatrix} I \\ G_w \end{bmatrix} (I - K_\infty G_w)^{-1} \begin{bmatrix} K_\infty & I \end{bmatrix} \right\|_\infty. \quad (B 3)$$

Hence, maximising $b(G_w, K_\infty)$ minimises the influence of the worst-case disturbances entering the model from all positions in the feedback loop and can be seen as giving a measure of robust performance as well as robust stability (Jones *et al.* 2015).

B.2. Controller synthesis

In the \mathcal{H}_∞ loop-shaping procedure proposed by McFarlane & Glover (1992), the designer first chooses frequency-dependent weights W to shape a nominal model G_0 . For SISO systems like the ones considered here, the shaped model is simply $G_w = WG_0$. The weights are chosen to enforce the performance criteria for the given application.

In the second step of the design procedure, the controller K_∞ that maximises $b(G_w, K_\infty)$ is identified, typically using the `ncfsyn` command in MATLAB. If the resulting value of $b(G_w, K_\infty)$ is acceptably large (usually 0.2–0.3 is sufficient (Vinnicombe 2001)), then the controller can be expected to provide reasonable robustness to the closed-loop system. If not, one must redesign the weights W and/or modify the actual system and/or the nominal model G_0 . In the final step, once an acceptable K_∞ has been designed, the final controller is generated by absorbing the weights within the controller so that $K = K_\infty W$.

B.3. The ν -gap metric

The ν -gap metric, introduced by Vinnicombe (1993), is a measure of uncertainty that fits naturally into the \mathcal{H}_∞ loop-shaping framework. The ν -gap between two SISO linear rational models G_1 and G_2 can be evaluated directly from their frequency responses,

$$\delta_\nu(G_1, G_2) = \sup_\omega \frac{|G_1(i\omega) - G_2(i\omega)|}{(1 + |G_1(i\omega)|^2)^{1/2} (1 + |G_2(i\omega)|^2)^{1/2}}, \quad (B 4)$$

where ω is the angular frequency, if the following conditions are satisfied:

$$\left. \begin{aligned} &1 + G_2(-i\omega)G_1(i\omega) \neq 0, \quad \forall \omega, \\ &\text{wno}[1 + G_2(-i\omega)G_1(i\omega)] - \eta[G_1] - \eta[G_2] = 0. \end{aligned} \right\} \quad (B 5)$$

Here, given a linear rational SISO model whose frequency response is $G(i\omega)$, $\eta[G]$ is its number of open right half-plane (unstable) poles and $\text{wno}[G]$ is its winding number, referring to the number of anticlockwise encirclements of the origin of $G(i\omega)$ in the complex plane. If the conditions above are not satisfied, then $\delta_\nu(G_1, G_2) = 1$. Note that $\delta_\nu(G_1, G_2) = \delta_\nu(G_2, G_1)$ and that, for identical systems, $\delta_\nu(G_1, G_1) = 0$.

In order to use the ν -gap within the \mathcal{H}_∞ loop-shaping procedure outlined above, one typically first obtains a nominal model G_0 and a family of ‘perturbed’ models that collectively represent an uncertain system, where the i th model is G_i . Then, after designing the weights W as discussed above, the ν -gap is calculated for each shaped perturbed model: $\delta_\nu(WG_0, WG_i)$.

The shaped nominal model is then used to obtain $K = K_\infty W$. At this stage, the ν -gap can be used to provide guarantees regarding the stability of the model family. As described by Vinnicombe (1993), given weights W and a controller K_∞ that result in a stability margin of $b(WG_0, K_\infty) = \beta$, two scenarios can arise for each G_i :

- (i) $\delta_\nu(WG_0, WG_i) < \beta$, in which case WG_i is guaranteed to be stabilised by K_∞ ;
- (ii) $\delta_\nu(WG_0, WG_i) > \beta$, in which case there exists a controller that stabilises WG_0 with a stability margin of at least β but destabilises WG_i .

Note that the shaped models are again used in all cases. The second scenario is problematic since it is possible that our particular choice of K_∞ does not stabilise WG_i . In this case, one option is to alter the system/models in order to reduce the associated uncertainty. Alternatively, the weights W can be changed in order to ensure $b(WG_0, K_\infty) > \delta_\nu(WG_0, WG_i)$ for all G_i . An important note therefore is that the ν -gap only provides information regarding the similarity of two models in feedback given the weights W and is not a measure of the similarity between the two open-loop models in general. In other words, modifying W is likely to have a significant impact on the calculated ν -gap.

REFERENCES

- AHUJA, S. & ROWLEY, C. W. 2010 Feedback control of unstable steady states of flow past a flat plate using reduced-order estimators. *J. Fluid Mech.* **645**, 447–478.
- ÅKERVIK, E., BRANDT, L., HENNINGSON, D. S., HØPEFFNER, J., MARXEN, O. & SCHLATTER, P. 2006 Steady solutions of the Navier–Stokes equations by selective frequency damping. *Phys. Fluids* **18**, 68102.
- ÅKERVIK, E., HØPEFFNER, J., EHRENSTEIN, U. & HENNINGSON, D. S. 2007 Optimal growth, model reduction and control in a separated boundary-layer flow using global eigenmodes. *J. Fluid Mech.* **579**, 305–314.
- AUBRY, N., HOLMES, P., LUMLEY, J. L. & STONE, E. 1988 The dynamics of coherent structures in the wall region of a turbulent boundary layer. *J. Fluid Mech.* **192**, 115–173.
- BAGHERI, S., ÅKERVIK, E., BRANDT, L. & HENNINGSON, D. S. 2009a Matrix-free methods for the stability and control of boundary layers. *AIAA J.* **47**, 1057–1068.
- BAGHERI, S., BRANDT, L. & HENNINGSON, D. S. 2009b Input–output analysis, model reduction and control of the flat-plate boundary layer. *J. Fluid Mech.* **620**, 263–298.
- BAGHERI, S., HENNINGSON, D. S., HØPEFFNER, J. & SCHMID, P. J. 2009c Input–output analysis and control design applied to a linear model of spatially developing flows. *Appl. Mech. Rev.* **62**, 020803.
- BARBAGALLO, A., DERGHAM, G., SIPP, D., SCHMID, P. J. & ROBINET, J.-C. 2012 Closed-loop control of unsteadiness over a rounded backward-facing step. *J. Fluid Mech.* **703**, 326–362.
- BARBAGALLO, A., SIPP, D. & SCHMID, P. J. 2009 Closed-loop control of an open cavity flow using reduced-order models. *J. Fluid Mech.* **641**, 1–50.
- BARBAGALLO, A., SIPP, D. & SCHMID, P. J. 2011 Input–output measures for model reduction and closed-loop control: application to global modes. *J. Fluid Mech.* **685**, 23–53.
- BEAUDOIN, J.-F., CADOT, O., AIDER, J.-L. & WESFREID, J.-E. 2006a Bluff-body drag reduction by extremum-seeking control. *J. Fluid Struct.* **22**, 973–978.
- BEAUDOIN, J.-F., CADOT, O., AIDER, J.-L. & WESFREID, J.-E. 2006b Drag reduction of a bluff body using adaptive control methods. *Phys. Fluids* **18**, 085107.

- BECKER, R., GARWON, M., GUTKNECHT, C., BARWOLFF, G. & KING, R. 2005 Robust control of separated shear flows in simulation and experiment. *J. Process Control* **15**, 691–700.
- BECKER, R., KING, R., PETZ, R. & NITSCHKE, W. 2007 Adaptive closed-loop separation control on a high-lift configuration using extremum seeking. *AIAA J.* **45**, 1382–1392.
- BELSON, B. A., SEMERARO, O., ROWLEY, C. W. & HENNINGSON, D. S. 2013 Feedback control of instabilities in the two-dimensional Blasius boundary layer: the role of sensors and actuators. *Phys. Fluids* **25**, 054106.
- BEWLEY, T. R. & LIU, S. 1998 Optimal and robust control and estimation of linear paths to transition. *J. Fluid Mech.* **365**, 305–349.
- BEWLEY, T. R., MOIN, P. & TEMAM, R. 2001 DNS-based predictive control of turbulence: an optimal benchmark for feedback algorithms. *J. Fluid Mech.* **447**, 179–225.
- BEWLEY, T. R., TEMAM, R. & ZIANE, M. 2000 A general framework for robust control in fluid mechanics. *Physica D* **138**, 360–392.
- BRUNTON, S. L., DAWSON, S. T. M. & ROWLEY, C. W. 2014 State-space model identification and feedback control of unsteady aerodynamic forces. *J. Fluids Struct.* **50**, 253–270.
- BRUNTON, S. L., ROWLEY, C. W. & WILLIAMS, D. R. 2013 Reduced-order unsteady aerodynamic models at low Reynolds numbers. *J. Fluid Mech.* **724**, 203–233.
- CHEN, K. K. & ROWLEY, C. W. 2011 H₂ optimal actuator and sensor placement in the linearised complex Ginzburg–Landau system. *J. Fluid Mech.* **681**, 241–260.
- CHOI, H., JEON, W.-P. & KIM, J. 2008 Control of flow over a bluff body. *Annu. Rev. Fluid Mech.* **40**, 113–139.
- CHOI, J., COLONIUS, T. & WILLIAMS, D. R. 2015 Surging and plunging oscillations of an airfoil at low Reynolds number. *J. Fluid Mech.* **763**, 237–253.
- COLONIUS, T. & TAIRA, K. 2008 A fast immersed boundary method using a nullspace approach and multi-domain far-field boundary conditions. *Comput. Meth. Appl. Mech. Engng* **197**, 2131–2146.
- DADFAR, R., SEMERARO, O., HANIFI, A. & HENNINGSON, D. S. 2013 Output feedback control of Blasius flow with leading edge using plasma actuator. *AIAA J.* **51**, 2192–2207.
- DAHAN, J. A., MORGANS, A. S. & LARDEAU, S. 2012 Feedback control for form-drag reduction on a bluff body with a blunt trailing edge. *J. Fluid Mech.* **704**, 360–387.
- DERGHAM, G., SIPP, D., ROBINET, J.-C. & BARBAGALLO, A. 2011 Model reduction for fluids using frequential snapshots. *Phys. Fluids* **23**, 064101.
- DOWLING, A. P. & MORGANS, A. S. 2005 Feedback control of combustion oscillations. *Annu. Rev. Fluid Mech.* **37**, 151–182.
- DOYLE, J. C. 1978 Guaranteed margins for LQG regulators. *IEEE Trans. Autom. Control* **23**, 756–757.
- EHRENSTEIN, U., PASSAGGIA, P.-Y. & GALLAIRE, F. 2011 Control of a separated boundary layer: reduced-order modeling using global modes revisited. *Theor. Comput. Fluid Dyn.* **25**, 195–207.
- FABBIANE, N., SEMERARO, O., BAGHERI, S. & HENNINGSON, D. S. 2014 Adaptive and model-based control theory applied to convectively unstable flows. *Appl. Mech. Rev.* **66**, 060801.
- FABBIANE, N., SIMON, B., FISCHER, F., GRUNDMANN, S., BAGHERI, S. & HENNINGSON, D. S. 2015 On the role of adaptivity for robust laminar flow control. *J. Fluid Mech.* **767**, 1–12.
- FLINOIS, T. L. B. & COLONIUS, T. 2015 Optimal control of circular cylinder wakes using long control horizons. *Phys. Fluids* **27**, 087105.
- FLINOIS, T. L. B., MORGANS, A. S. & SCHMID, P. J. 2015 Projection-free approximate balanced truncation of large unstable systems. *Phys. Rev. E* **92**, 023012.
- FROHNAPFEL, B., HASEGAWA, Y. & QUADRIO, M. 2012 Money versus time: evaluation of flow control in terms of energy consumption and convenience. *J. Fluid Mech.* **700**, 406–418.
- GAUTIER, N., DURIEZ, T., AIDER, J.-L., NOACK, B. R., SEGOND, M. & ABEL, M. W. 2015 Closed-loop separation control using machine learning. *J. Fluid Mech.* **770**, 442–457.
- GHIGLIERI, J. & ULBRICH, S. 2014 Optimal flow control based on POD and MPC and an application to the cancellation of Tollmien–Schlichting waves. *Optim. Meth. Softw.* **29**, 1042–1074.
- GIANNETTI, F. & LUCHINI, P. 2007 Structural sensitivity of the first instability of the cylinder wake. *J. Fluid Mech.* **581**, 167–197.
- GLOERFELT, X. 2008 Compressible proper orthogonal decomposition/Galerkin reduced-order model of self-sustained oscillations in a cavity. *Phys. Fluids* **20**, 115105.

- GOLDIN, N., KING, R., PÄTZOLD, A., NITSCHKE, W., HALLER, D. & WOIAS, P. 2013 Laminar flow control with distributed surface actuation: damping Tollmien–Schlichting waves with active surface displacement. *Exp. Fluids* **54**, 1478.
- GRAHAM, W. R., PERAIRE, J. & TANG, K. Y. 1999 Optimal control of vortex shedding using low-order models. Part I – Open-loop model development. *Intl J. Numer. Meth. Engng* **44**, 945–972.
- GUZMAN INIGO, J., SIPP, D. & SCHMID, P. J. 2014 A dynamic observer to capture and control perturbation energy in noise amplifiers. *J. Fluid Mech.* **758**, 728–753.
- GAD-EL HAK, M. 2000 *Flow Control: Passive, Active and Reactive Flow Management*. Cambridge University Press.
- HENNING, L., BECKER, R., FEUERBACH, G., MUMINOVIC, R., KING, R., BRUNN, A. & NITSCHKE, W. 2008 Extensions of adaptive slope-seeking for active flow control. *Proc. Inst. Mech. Engrs* **222**, 309–322.
- HENNING, L. & KING, R. 2007 Robust multivariable closed-loop control of a turbulent backward-facing step flow. *J. Aircraft* **44**, 201–208.
- HENNING, L., PASTOOR, M., KING, R. & NOACK, B. R. 2007 Feedback control applied to the bluff body wake. In *Active Flow Control* (ed. R. King), pp. 369–390. Springer.
- HENNINGSON, D. S. & ÅKERVİK, E. 2008 The use of global modes to understand transition and perform flow control. *Phys. Fluids* **20**, 031302.
- HERVÉ, A., SIPP, D., SCHMID, P. J. & SAMUELIDES, M. 2012 A physics-based approach to flow control using system identification. *J. Fluid Mech.* **702**, 26–58.
- HUANG, S. C. & KIM, J. 2008 Control and system identification of a separated flow. *Phys. Fluids* **20**, 101509.
- ILAK, M. & ROWLEY, C. W. 2008 Modeling of transitional channel flow using balanced proper orthogonal decomposition. *Phys. Fluids* **20**, 034103.
- ILLINGWORTH, S. J., MORGANS, A. S. & ROWLEY, C. W. 2011 Feedback control of flow resonances using balanced reduced-order models. *J. Sound. Vib.* **330**, 1567–1581.
- ILLINGWORTH, S. J., MORGANS, A. S. & ROWLEY, C. W. 2012 Feedback control of cavity flow oscillations using simple linear models. *J. Fluid Mech.* **709**, 223–248.
- ILLINGWORTH, S. J., NAITO, H. & FUKAGATA, K. 2014 Active control of vortex shedding: an explanation of the gain window. *Phys. Rev. E* **90**, 043014.
- JOE, W. T., COLONIUS, T. & MACMYNOWSKI, D. G. 2010 Optimized waveforms for feedback control of vortex shedding. In *Active Flow Control II* (ed. R. King), pp. 391–404. Springer.
- JOE, W. T., COLONIUS, T. & MACMYNOWSKI, D. G. 2011 Feedback control of vortex shedding from an inclined flat plate. *Theor. Comput. Fluid Dyn.* **25**, 221–232.
- JONES, B. L., HEINS, P. H., KERRIGAN, E. C., MORRISON, J. F. & SHARMA, A. S. 2015 Modelling for robust feedback control of fluid flows. *J. Fluid Mech.* **769**, 687–722.
- JORDI, B. E., COTTER, C. J. & SHERWIN, S. J. 2014 Encapsulated formulation of the selective frequency damping method. *Phys. Fluids* **26**, 034101.
- JUANG, J.-N. & PAPPA, R. S. 1985 An eigensystem realization algorithm for modal parameter identification and model reduction. *J. Guid.* **8**, 620–627.
- JUILLET, F., MCKEON, B. J. & SCHMID, P. J. 2014 Experimental control of natural perturbations in channel flow. *J. Fluid Mech.* **752**, 296–309.
- JUILLET, F., SCHMID, P. J. & HUERRE, P. 2013 Control of amplifier flows using subspace identification techniques. *J. Fluid Mech.* **725**, 522–565.
- KIM, J. & BEWLEY, T. R. 2007 A linear systems approach to flow control. *Annu. Rev. Fluid Mech.* **39**, 383–417.
- KWONG, A. H. M. & DOWLING, A. P. 1994 Active boundary-layer control in diffusers. *AIAA J.* **32**, 2409–2414.
- LAUGA, E. & BEWLEY, T. R. 2003 The decay of stabilizability with Reynolds number in a linear model of spatially developing flows. *Proc. R. Soc. Lond. A* **459**, 2077–2095.
- LAUGA, E. & BEWLEY, T. R. 2004 Performance of a linear robust control strategy on a nonlinear model of spatially developing flows. *J. Fluid Mech.* **512**, 343–374.

- LJUNG, L. 1987 *System Identification – Theory for the User*. (PTR Prentice Hall Information and System Sciences Series), vol. 198. Prentice Hall.
- MA, X. & KARNIADAKIS, G. E. 2002 A low-dimensional model for simulating three-dimensional cylinder flow. *J. Fluid Mech.* **458**, 181–190.
- MA, Z., AHUJA, S. & ROWLEY, C. W. 2011 Reduced-order models for control of fluids using the Eigensystem Realization Algorithm. *Theor. Comput. Fluid Dyn.* **25**, 233–247.
- McFARLANE, D. & GLOVER, K. 1992 A loop-shaping design procedure using H-infinity synthesis. *IEEE Trans. Autom. Control* **37**, 759–769.
- MILANO, M. & KOUMOUTSAKOS, P. 2002 A clustering genetic algorithm for cylinder drag optimization. *J. Comput. Phys.* **175**, 79–107.
- MOORE, B. C. 1981 Principal component analysis in linear systems: controllability, observability, and model reduction. *IEEE Trans. Autom. Control* **26**, 17–32.
- MORGANS, A. S. & DOWLING, A. P. 2004 Low-order modeling for transonic helicopter noise. *AIAA J.* **42**, 2416–2428.
- MORGANS, A. S. & DOWLING, A. P. 2007 Model-based control of combustion instabilities. *J. Sound. Vib.* **299**, 261–282.
- MORGANS, A. S. & STOW, S. R. 2007 Model-based control of combustion instabilities in annular combustors. *Combust. Flame* **150**, 380–399.
- NOACK, B. R., AFANASIEV, K., MORZYSKI, M., TADMOR, G. & THIELE, F. 2003 A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *J. Fluid Mech.* **497**, 335–363.
- PALEI, V. & SEIFERT, A. 2007 Effects of periodic excitation on the flow around a D-shaped cylinder at low Reynolds numbers. *Flow Turbul. Combust.* **78**, 409–428.
- PARK, D. S., LADD, D. M. & HENDRICKS, E. W. 1994 Feedback control of von Karman vortex shedding behind a circular cylinder at low Reynolds numbers. *Phys. Fluids* **6**, 2390.
- PASTOOR, M., HENNING, L., NOACK, B. R., KING, R. & TADMOR, G. 2008 Feedback shear layer control for bluff body drag reduction. *J. Fluid Mech.* **608**, 161–196.
- PODVIN, B. & LUMLEY, J. L. 1998 A low-dimensional approach for the minimal flow unit. *J. Fluid Mech.* **362**, 121–155.
- PRA BHU, R. D., COLLIS, S. S. & CHANG, Y. 2001 The influence of control on proper orthogonal decomposition of wall-bounded turbulent flows. *Phys. Fluids* **13**, 520–537.
- PROTAS, B. & STYCZEK, A. 2002 Optimal rotary control of the cylinder wake in the laminar regime. *Phys. Fluids* **14**, 2073–2087.
- QIN, S. J. 2006 An overview of subspace identification. *Comput. Chem. Engng* **30**, 1502–1513.
- RAVINDRAN, S. S. 2000a A reduced-order approach for optimal control of fluids using proper orthogonal decomposition. *Intl J. Numer. Meth. Fluids* **34**, 425–448.
- RAVINDRAN, S. S. 2000b Reduced-order adaptive controllers for fluid flows using POD. *J. Sci. Comput.* **15**, 457–478.
- RAVINDRAN, S. S. 2002 Control of flow separation over a forward-facing step by model reduction. *Comput. Meth. Appl. Mech. Engng* **191**, 4599–4617.
- RAVINDRAN, S. S. 2006 Reduced-order controllers for control of flow past an airfoil. *Intl J. Numer. Meth. Fluids* **50**, 531–554.
- ROCA, P., CAMMILLERI, A., DURIEZ, T., MATHELIN, L. & ARTANA, G. 2014 Streakline-based closed-loop control of a bluff body flow. *Phys. Fluids* **26**, 047102.
- ROUSSOPOULOS, K. 1993 Feedback control of vortex shedding at low Reynolds numbers. *J. Fluid Mech.* **248**, 267–296.
- ROWLEY, C. W. 2005 Model reduction for fluids, using balanced proper orthogonal decomposition. *Intl J. Bifurcation Chaos* **15**, 997–1013.
- SEMERARO, O., BAGHERI, S., BRANDT, L. & HENNINGSON, D. S. 2011 Feedback control of three-dimensional optimal disturbances using reduced-order models. *J. Fluid Mech.* **677**, 63–102.
- SENGUPTA, T. K., DEB, K. & TALLA, S. B. 2007 Control of flow using genetic algorithm for a circular cylinder executing rotary oscillation. *Comput. Fluids* **36**, 578–600.

- SIEGEL, S. G., SEIDEL, J., FAGLEY, C., LUCHTENBURG, D. M., COHEN, K. & MCLAUGHLIN, T. 2008 Low-dimensional modelling of a transient cylinder wake using double proper orthogonal decomposition. *J. Fluid Mech.* **610**, 1–42.
- SIPP, D., MARQUET, O., MELIGA, P. & BARBAGALLO, A. 2010 Dynamics and control of global instabilities in open-flows: a linearized approach. *Appl. Mech. Rev.* **63**, 030801.
- TADMOR, G., LEHMANN, O., NOACK, B. R. & MORZYSKI, M. 2010 Mean field representation of the natural and actuated cylinder wake. *Phys. Fluids* **22**, 034102.
- TAIRA, K. & COLONIUS, T. 2007 The immersed boundary method: a projection approach. *J. Comput. Phys.* **225**, 2118–2137.
- TU, J. H. & ROWLEY, C. W. 2012 An improved algorithm for balanced POD through an analytic treatment of impulse response tails. *J. Comput. Phys.* **231**, 5317–5333.
- VINNICOMBE, G. 1993 Frequency domain uncertainty and the graph topology. *IEEE Trans. Autom. Control* **38**, 1371–1383.
- VINNICOMBE, G. 2001 *Uncertainty and Feedback*. Imperial College Press.
- WILLCOX, K. & PERAIRE, J. 2002 Balanced model reduction via the proper orthogonal decomposition. *AIAA J.* **40**, 2323–2330.
- WILLIAMS, D. R., KERSTENS, W., PFEIFFER, J., KING, R. & COLONIUS, T. 2010 Unsteady lift suppression with a robust closed-loop controller. In *Active Flow Control II*, pp. 19–30. Springer.
- ZHU, M., DOWLING, A. P. & BRAY, K. N. C. 2005 Transfer function calculations for aeroengine combustion oscillations. *Trans. ASME J. Engng Gas Turbines Power* **127**, 18–26.