Accurate Low-Dimensional Approximation of the Linear Dynamics of Fluid Flow

BRIAN F. FARRELL

Department for Earth and Planetary Sciences, Harvard University, Cambridge, Massachusetts

PETROS J. IOANNOU

Section of Astrophysics, Astronomy and Mechanics, Department of Physics, University of Athens, Panepistimiopolis, Zografos, Athens, Greece

(Manuscript received 4 October 1999, in final form 30 January 2001)

ABSTRACT

Methods for approximating a stable linear autonomous dynamical system by a system of lower order are examined. Reducing the order of a dynamical system is useful theoretically in identifying the irreducible dimension of the dynamics and in isolating the dominant spatial structures supporting the dynamics, and practically in providing tractable lower-dimension statistical models for climate studies and error covariance models for forecast analysis and initialization. Optimal solution of the model order reduction problem requires simultaneous representation of both the growing structures in the system and the structures into which these evolve. For autonomous operators associated with fluid flows a nearly optimal solution of the model order reduction problem with prescribed error bounds is obtained by truncating the dynamics in its Hankel operator representation. Simple model examples including a reduced-order model of Couette flow are used to illustrate the theory. Practical methods for obtaining approximations to the optimal order reduction problem based on finite-time singular vector analysis of the propagator are discussed and the accuracy of the resulting reduced models evaluated.

1. Introduction

The dynamical system of fluid flow can be expressed as a set of partial differential equations and while the formal dimension of the state space of these equations is infinite, projection of the dynamics onto a sufficiently large, but finite, basis of functions suffices for an adequate representation of the dynamics. Nevertheless, the dimension, N, of this finite basis may be very large, for example, $N = O(10^7)$ for present forecast models, and this large dimension is an obstacle to implementing statistically optimal state estimation methods for forecast initialization, including the Kalman filter which requires the perturbation error covariance of dimension N^2 $[O(10^{14})]$ be obtained (Farrell and Ioannou 2001). Moreover, it is of fundamental theoretical interest to determine the minimal dynamical dimension of a fluid system such as a storm track model and also to determine the minimal structures required to accurately represent the dynamics and the dynamical relation among these structures.

The system governing the dynamics of perturbations to the fluid equations is linear, and if the nonlinear state to which the perturbation is introduced is stationary, then a finite-dimensional linear autonomous system governs the perturbation dynamics. We study the problem of finding a dynamical system of lower dimension that best approximates the full perturbation dynamics in such a finite-dimensional linear autonomous system. This model order reduction requires for an accurate representation of the dynamics that both the set of optimally growing perturbations at an initial time, and the set of perturbations into which these evolve at the future (optimizing) time (the optimals and the evolved optimals or equivalently the left and right singular vectors of the linear system propagator), be faithfully represented in the dynamics. For normal systems this problem has a simple solution: truncation retaining a set of least damped eigenmodes of the system. But the perturbation dynamics of fluid flow is governed in general by nonnormal operators which, unlike normal operators, have optimals and evolved optimals that change with the optimizing time and are distinct both from each other and from the eigenmodes of the dynamical operator. Finding a way to represent the dynamics of both the optimals and the evolved optimals in a balanced manner is required for solution of the problem of model order reduction in nonnormal systems.

Traditionally in fluid mechanics model order reduction has been performed by an ad hoc truncation accomplished by reducing the number of collocation

Corresponding author address: Dr. Brian F. Farrell, Department for Earth and Planetary Sciences, Pierce Hall, 29 Oxford St., Harvard University, Cambridge, MA 02138.
E-mail: farrell@deas.harvard.edu

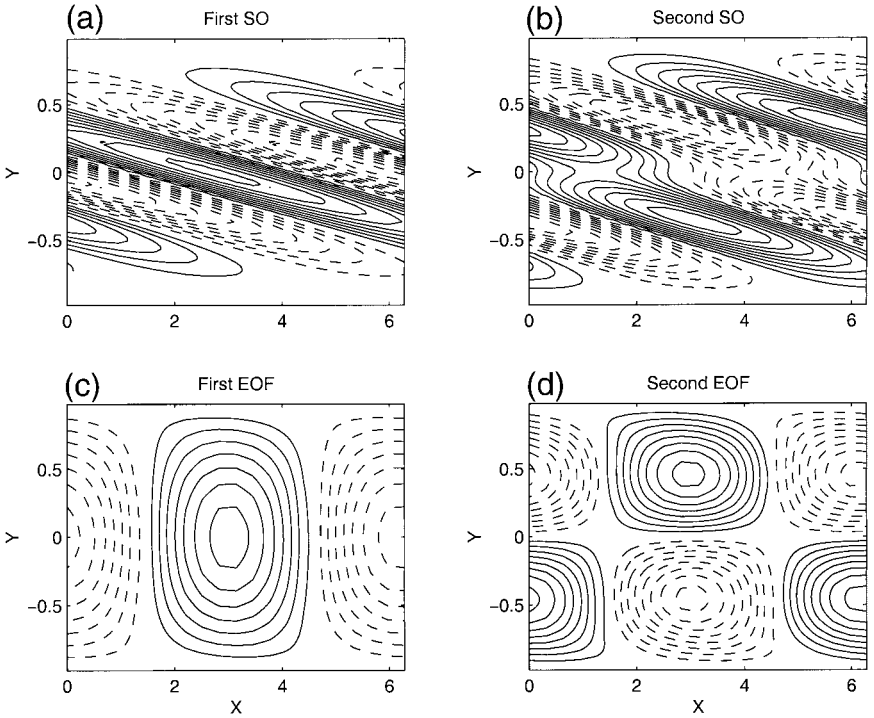

FIG. 1. The structure of the streamfunction of the first two stochastic optimals which are the eigenfunctions of **Q** with the largest eigenvalues. (a) First SO, which is responsible for producing 41% of the maintained variance. (b) Second SO, which is responsible for producing 25% of the variance. (bottom) The structure of streamfunction of the first two EOFs, which are the eigenfunctions of **P** with the largest eigenvalues. (c) First EOF, which is responsible for producing 57% of the maintained variance. (d) Second EOF, which is responsible for producing 17% of the variance. For the Couette flow at Re = 800, $k = 1$.

points in a gridpoint representation or modes in a spectral representation of the dynamical equations. This method is straightforwardly implemented but reduction to a low-order system by simple truncation of a grid point or spectral representation lacks error bounds and may even introduce instabilities. A more dynamically motivated truncation is truncation to an ordered set of least damped eigenmodes of the system. This method of reducing the model order at least has the advantage of retaining in the reduced-order system the stability properties of the original system but except for normal systems truncation to a set of least damped modes is suboptimal and in practice produces disappointing results as the examples to be shown demonstrate.

An alternative basis for the dynamics is obtained by exciting the system with unbiased forcing and projecting the dynamics onto the empirical orthogonal functions (EOFs) of its response; a closely related basis is obtained by projecting onto the observed EOFs of a nonlinear system (Hasselmann 1998; Holmes et al. 1997). Each of these methods is an attempt to introduce information about the dynamics into the truncation. However, while projection onto EOFs is easily implemented and represents the evolved optimals well, it is not suited to representing the optimals themselves (Farrell and Ioannou 1993a,b).

We wish to find a method for reducing the order of the dynamics that retains both the growing structures and the structures into which these evolve. At the least we would like to have assurance that the stability properties of the original system are preserved, but desirable would be explicit bounds on the approximation error. In the case of discrete multiple input multiple output (MIMO) discrete parameter systems optimal truncation of the dynamics has been obtained through balanced truncation of the Hankel operator representation of the dynamics (Moore 1981; Glover 1984; Zhou and Doyle 1998). We seek to extend this balanced truncation method of model order reduction to provide an optimal truncation of the continuous dynamics of fluid flow.

The method of model order reduction presented in this paper not only demonstrates constructively that the linear dynamics are not optimally approximated by reducing the dimension using either EOFs or singular vectors or the least damped modes of the system but also provides the spatial and temporal structure of the errors incurred in the approximation and a constructive a method of determining the effective dimension of the dynamical system.

In the following, methods of model order reduction are discussed, the balanced truncation is introduced, a

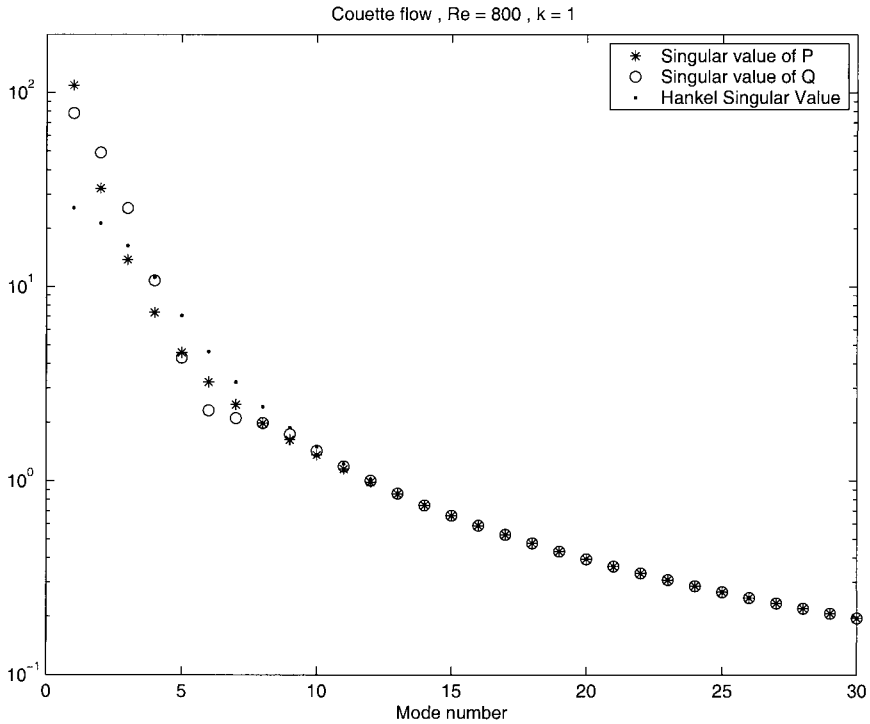

FIG. 2. The Hankel singular values (dots), the singular values (or eigenvalues) of **P** (stars) and the singular values (or eigenvalues) of **Q** in descending order of their magnitude. The corresponding structures of the first SO and EOF are shown in Fig. 1. The Hankel singular values are the square roots of the eigenvalues of **PQ**.

comparison is made among these methods of model reduction, and an example problem is examined.

2. Methods for reducing the order of a linear dynamical systems

a. Optimal reduction of the order of a matrix and reduction of order for normal dynamical systems

Consider the linear dynamical system

$$\frac{d\psi}{dt} = \mathbf{A}\psi, \tag{1}$$

in which, if necessary, the continuous operators have been discretized so that the dynamical operator is a finite-dimensional matrix. Our goal is to reduce further the order of this finite-dimensional matrix, retaining desired accuracy of the solution.

A method for optimally reducing the order of a finite-dimensional matrix is immediately available from matrix theory; provided that the error is measured in the euclidean or L_2 norm.[1] This solution can be exploited to approximate the system propagator that advances the initial state of the system at $t = 0$ to its state at a fixed later time t because this propagator is a finite dimensional matrix: for **A** autonomous the explicit expression for the propagator matrix $\mathbf{\Phi}(t)$ at time t is $\mathbf{\Phi}(t) = e^{\mathbf{A}t}$. Optimal truncation of this matrix propagator can be constructed using singular value decomposition. The singular value dyadic expansion of the propagator is $\mathbf{\Phi}(t) = \sum_{i=1}^{n} \sigma_i u_i v_i^\dagger$, where n is the dimension of the state space, and the n vectors u_i and v_i define two orthonormal bases. The v_i are the optimals, the initial perturbations ordered in decreasing growth σ_i over time t; and the u_i are the evolved optimals, the structures to which the v's evolve in time t (sometimes referred to respectively as the left and right singular vectors of the matrix propagator). If **A** is normal (i.e., $\mathbf{A}\mathbf{A}^\dagger = \mathbf{A}^\dagger\mathbf{A}$) both the optimals and the evolved optimals coincide with the eigenvectors of the operator **A**. It can then be shown that in the L_2 norm the error in approximating the propagator at time t by any matrix **X** of rank less than k, with $k < n$, is necessarily at least equal to the $k + 1$st singular value of the propagator; that is, the error in the L_2 norm satisfies $\|\mathbf{\Phi}(t) - \mathbf{X}\| \geq \sigma_{k+1}$. The minimum error σ_{k+1} is attained (nonuniquely) by the matrix approximation

$$\mathbf{X}_k = \sum_{i=1}^{k} \sigma_i u_i v_i^\dagger, \tag{2}$$

[1] The optimal solution is known only for the class of unitarily invariant norms, i.e., norms that are invariant to unitary transformation of their vectors; the L_2 or the Frobenius norm are such norms (cf. Stewart and Sun 1990).

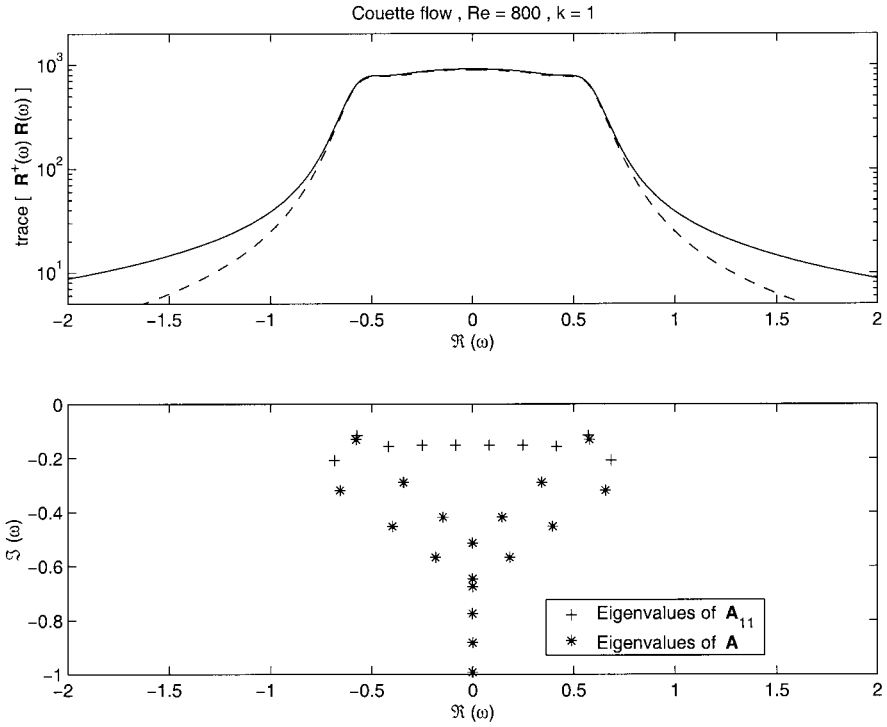

FIG. 3. (top) Frequency response trace $[\mathbf{R}^{\dagger}(\omega)\mathbf{R}(\omega)]$ as a function of real frequency ω. The continuous curve is the original system, the dash curve is the order 10 balanced truncation. (bottom) Frequency $[\mathcal{I}(\omega)]$and decay rate $[\mathcal{R}(\omega)]$ of the eigenvalues of the original system (stars), and of the order 10 balanced operator $\mathbf{A}_{11}$ (crosses). For Couette flow at Re = 800 and $k = 1$.

which is obtained by truncating the singular value decomposition of $\mathbf{\Phi}(t)$ to its first k largest singular values. This result is the Eckart–Schmidt–Mirsky (ESM) theorem (Schmidt 1907; Eckart and Young 1936; Mirsky 1960; Stewart and Sun 1990) and it provides an optimal truncation of the propagator in the euclidean norm at any fixed time t.

For a normal system the propagator of the full system is

$$\mathbf{X}_n = \sum_{i=1}^{n} e^{\lambda_i t} e_i e_i^{\dagger}, \tag{3}$$

where λ_i is the eigenvalue associated with eigenvector e_i. If the λ_i are ordered decreasing in their real part, $\mathcal{R}(\lambda_i)$, then truncation to the first k modes results in the order k propagator:

$$\mathbf{X}_k = \sum_{i=1}^{k} e^{\lambda_i t} e_i e_i^{\dagger}. \tag{4}$$

This propagator limits the dynamics to the subspace spanned by the first k least damped modes of the system. It is clear by appeal to the ESM theorem that reduction of the dimension of the system by Galerkin projection on the first k least damped modes produces the optimal order k truncation of the propagator at all times and that error in the propagator at time t is bounded by $e^{\lambda_{k+1} t}$. Also note that the modally reduced system inherits the stability properties of the original system.

However, modal reduction is not suitable for most fluid problems because the linearized equations associated with fluid flows are generally nonnormal and modal reduction is suboptimal for nonnormal systems. In order to obtain an approximation of the dynamics for nonnormal systems we seek to extend application of the ESM theorem to an arbitrary nonnormal system, obtaining a truncation that is otherwise similar to that resulting from modal truncation of a normal system. The necessary result was developed by Adamjan et al. (1971) and methods of model reduction based on it were developed in the context of discrete MIMO systems by Moore (1981) and Glover (1984) (see also Zhou and Doyle 1998).

The method of optimal model reduction utilizes the stochastic optimals and the EOFs (Farrell and Ioannou 1993a,b, 1996) and we start by discussing how these structures are obtained.

b. Reducing the order of a nonnormal system by projection on the space spanned by the EOFs and the stochastic optimals

We wish to educe the intrinsic dynamics of the linear system without influencing the system response by the

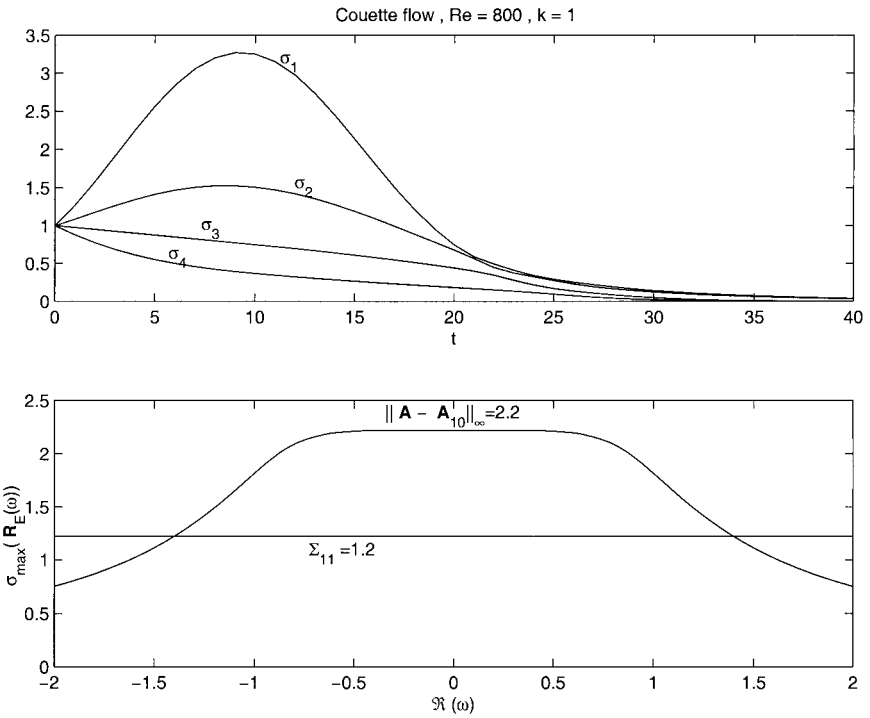

FIG. 4. (top) The top 4 singular values of the propagator, $e^{\mathbf{A}t}$, as a function of time for the Couette flow at Re = 800 and $k = 1$. The curves correspond to the full system and also (indistinguishably) to the order 10 approximate system obtained by balanced truncation. (bottom) A plot of the maximum singular value of the error system $\mathbf{A} - \mathbf{A}_{10}$ as a function of frequency. The system $\mathbf{A}_{10}$ is an order 10 approximation obtained from $\mathbf{A}$ by balanced truncation. The maximum of this curves is the H_∞ error of the order 10 balanced truncation which is found here to be 2.2. Also indicated with a straight line is the theoretical minimum error of an order 10 truncation, which equals the first neglected Hankel singular value $\Sigma_{11} = 1.2$. The balanced truncation is seen to be nearly optimal.

choice of forcing, therefore we excite the system in an unbiased manner:

$$\frac{d\psi}{dt} = \mathbf{A}\psi + \mathbf{F}w(t). \tag{5}$$

In (5) the operator $\mathbf{A}$ is assumed stable and is stochastically forced with white noise $w(t)$, and the structure of the forcing, $\mathbf{F}$, is unitary.[2] The statistically steady response (cf. Farrell and Ioannou 1996, hereafter F196) can be characterized by the covariance matrix $\mathbf{P} = \langle \psi\psi^\dagger \rangle$, where the brackets denote the ensemble average, which for unitary white noise forcing is also given by the expression:

$$\mathbf{P} = \int_0^\infty e^{\mathbf{A}t} e^{\mathbf{A}^\dagger t}\, dt. \tag{6}$$

As the covariance matrix, $\mathbf{P}$, is hermitian, its eigenvectors form an orthonormal basis. These are the EOFs that span most concisely the structures accounting for the variance; in the sense that the variance is distributed over the EOFs according to their associated eigenvalues, and the trace of the covariance matrix therefore is the total variance maintained by the stochastic forcing.

Because the EOFs are the most concise representation of the response of the system in the sense of accounting for the maximum variance with the fewest structures when the system is subjected to unbiased forcing, it is perhaps reasonable to expect that the EOFs also provide the preferred basis of functions for truncation of the system dynamics. The EOFs, after all, include in their structures information about the preferred responses of the system and should therefore be superior to an arbitrarily chosen orthogonal basis of functions in representing economically the state of the system. Indeed, a method of truncation based on observed covariance is used to produce lower-order systems for turbulent modeling (Holmes et al. 1997).

However, such a truncation is not optimal in general even for linear operators because, as was argued in Farrell and Ioannou (1993b), the linear operator $\mathbf{A}$ arising in fluid flow perturbation problems is generally nonnormal so that the optimal responses are structurally distinct from the corresponding optimal excitations, and consequently a lower-order system approximation of the

[2] That is, $\mathbf{F}\mathbf{F}^\dagger = \mathbf{I}$, where $\mathbf{I}$ the identity and $\mathbf{F}^\dagger$ the Hermitian transpose of $\mathbf{F}$.

dynamics needs to include in a balanced way both the optimal excitation structures as well as their distinct optimal responses. The time-integrated optimal excitations of the system are identified as the stochastic optimals that are the eigenfunctions of the Hermitian matrix

$$\mathbf{Q} = \int_0^\infty e^{\mathbf{A}^\dagger t} e^{\mathbf{A}t}\, dt. \tag{7}$$

This identification stems from writing the maintained variance of the stochastically forced system (5) as

$$\langle \psi^\dagger \psi \rangle = \text{trace}(\mathbf{F}^\dagger \mathbf{Q} \mathbf{F}), \tag{8}$$

where brackets denote the ensemble average. The eigenvectors of the positive definite hermitian matrix $\mathbf{Q}$ are the stochastic optimals (SOs) that order the forcing structures (the columns of $\mathbf{F}$) according to their contribution to exciting the variance under the assumption of unbiased forcing. The variance excited in statistical equilibrium by each eigenvector is given by the corresponding eigenvalue of $\mathbf{Q}$ and the total variance is the trace of $\mathbf{Q}$.

The importance of retaining a balance between the leading EOFs and the leading SOs in a truncation of the dynamical operator $\mathbf{A}$ is demonstrated by considering the following nonnormal system:

$$\mathbf{A} = \begin{pmatrix} -1 & -1/\epsilon \\ \epsilon & -2 \end{pmatrix}. \tag{9}$$

Because $\mathbf{A}$ is asymptotically stable, the covariance matrix $\mathbf{P}$, given by Eq. (6), is readily found by solving the Lyapunov equation (FI96):

$$\mathbf{A}\mathbf{P} + \mathbf{P}\mathbf{A}^\dagger = -\mathbf{I}, \tag{10}$$

where $\mathbf{I}$ is the identity. The solution is

$$\mathbf{P} = \frac{1}{18}\begin{pmatrix} 7 + 1/\epsilon^2 & 2\epsilon - 1/\epsilon \\ 2\epsilon - 1/\epsilon & 4 + \epsilon^2 \end{pmatrix}. \tag{11}$$

For small ϵ the eigenvalues of $\mathbf{P}$ approach the values $1/(18\epsilon^2)$ and $1/6$, and the leading EOF associated with the first eigenvalue approaches $[1, 0]^T$. As $\epsilon \to 0$ this eigenvector accounts for nearly all the variance. One may expect then that truncation to retain this first EOF would provide a good description of the dynamics. Indeed, this is the case for a normal dynamical operator $\mathbf{A}$. However, in the case of perturbation dynamics in fluid flows, for which $\mathbf{A}$ is generally nonnormal, this truncation of the system is demonstrably wrong in the sense that if the system is truncated to a single equation by projection on the dominant EOF the dynamics are poorly represented. The neglected EOF, which accounts for an asymptotically vanishing fraction of the total variance, is in fact the first SO, that is, the structure that is responsible asymptotically for producing all the variance. The SO is obtained by eigenanalysis of the matrix $\mathbf{Q}$, which can be found by solving the Lyapunov equation:

$$\mathbf{A}^\dagger \mathbf{Q} + \mathbf{Q}\mathbf{A} = -\mathbf{I}. \tag{12}$$

Because of the structure of $\mathbf{A}$ in this example, we can obtain $\mathbf{Q}$ by replacing ϵ by $-1/\epsilon$ in the expression for $\mathbf{P}$ in (11) to obtain

$$\mathbf{Q} = \frac{1}{18}\begin{pmatrix} 7 + \epsilon^2 & \epsilon - 2/\epsilon \\ \epsilon - 2/\epsilon & 4 + 1/\epsilon^2 \end{pmatrix}. \tag{13}$$

The eigenvalues of $\mathbf{Q}$ for $\epsilon \ll 1$ approach the eigenvalues of $\mathbf{P}$, but the eigenvector corresponding to the largest eigenvalue, the SO, approaches $[0, 1]^T$, a structure orthogonal to the first EOF as $\epsilon \to 0$. This structure, almost exclusively responsible for producing the variance, remarkably accounts only for a vanishing fraction of the variance. Truncation to retain only this leading SO would, however, also lead to an inaccurate approximation of the dynamical system as the dynamics would then not include the dominant response of the system, the leading EOF. This example demonstrates that for a good approximation of the dynamics both the leading SOs and the leading EOFs must be retained in the truncated dynamics.

If the system is normal the SOs and the EOFs coincide (they are identical to the eigenmodes of the system), and in that case an optimal k order truncation corresponds to retaining the least k damped modes of the system as we have already seen. The truncated operator is then $\mathbf{A}_k = \mathbf{E}_k \mathbf{D}_k \mathbf{E}_k^\dagger$ where $\mathbf{E}^k$ is the column matrix consisting of the first k eigenmodes of $\mathbf{A}$ and $\mathbf{D}_k$ is a $k \times k$ diagonal matrix with its diagonal elements equal to the eigenvalues corresponding to the eigenvectors in $\mathbf{E}_k$.

This procedure is not applicable to nonnormal systems in which the SOs, EOFs and eigenvectors are not identical. But if there were a coordinate system in which the EOFs and the SOs become identical, then in that coordinate system we could proceed with the truncation as in normal systems. Such a transformation exists and this procedure is called balancing and the coordinates in which both the $\mathbf{P}$ and $\mathbf{Q}$ are transformed to a diagonal matrix $\boldsymbol{\Sigma}$ is called the balanced realization (Moore 1981; Zhou and Doyle 1998).

Because the original dynamics are represented in coordinates chosen for physical relevance, we demand also that the transformed coordinates preserve the inner product of the original coordinates; for instance, the state ψ may be such that the euclidean inner product corresponds to energy. Consider now the transformation

$$x = \mathbf{T}\psi. \tag{14}$$

In order that the inner product in the new coordinate system, x, satisfies $(x, x) = \psi^\dagger \psi$ the inner product in the x space must be defined as

$$(x, x) = x^\dagger \mathbf{T}^{-1\dagger}\mathbf{T}^{-1}x. \tag{15}$$

The transformed dynamical equations (5) are then

$$\frac{dx}{dt} = \tilde{\mathbf{A}}x + \mathbf{T}\mathbf{F}w(t), \tag{16}$$

where

$$\tilde{\mathbf{A}} = \mathbf{T}\mathbf{A}\mathbf{T}^{-1}. \tag{17}$$

We obtain the transformed $\tilde{\mathbf{Q}}$ by making use of the invariance of the inner product as follows:

$$\begin{aligned}\langle \psi^\dagger \psi \rangle &= \text{trace}(\mathbf{F}^\dagger \mathbf{Q}\mathbf{F}) \\ &= \langle x^\dagger \mathbf{T}^{-1\dagger}\mathbf{T}^{-1}x \rangle \\ &= \text{trace}(\mathbf{F}^\dagger \tilde{\mathbf{Q}}\mathbf{F}),\end{aligned} \tag{18}$$

where

$$\tilde{\mathbf{Q}} = \int_0^\infty e^{\tilde{\mathbf{A}}^\dagger t}\mathbf{T}^{-1\dagger}\mathbf{T}^{-1}e^{\tilde{\mathbf{A}}t}\,dt. \tag{19}$$

Using (17) and the definition of $\mathbf{Q}$ given in (7) we have

$$\tilde{\mathbf{Q}} = \mathbf{T}^{-1\dagger}\mathbf{Q}\mathbf{T}^{-1}. \tag{20}$$

We obtain now the transformation law for $\mathbf{P}$. From the forced solution of (16) we have

$$\tilde{\mathbf{P}} \equiv \langle xx^\dagger \rangle = \int_0^\infty e^{\tilde{\mathbf{A}}t}\mathbf{T}\mathbf{T}^\dagger e^{\tilde{\mathbf{A}}^\dagger t}\,dt, \tag{21}$$

and again using (17) and the definition of $\mathbf{P}$ given in (6) we have

$$\tilde{\mathbf{P}} = \mathbf{T}\mathbf{P}\mathbf{T}^\dagger. \tag{22}$$

The coordinate transformation $\mathbf{T}$, which simultaneously diagonalizes the $\mathbf{P}$ and $\mathbf{Q}$ matrices, is obtained as follows (Moore 1981; Glover 1984). Define the Hermitian matrix

$$\mathbf{R} = \mathbf{P}^{1/2}\mathbf{Q}\mathbf{P}^{1/2}, \tag{23}$$

which requires for its calculation the determination of the square root[3] of the Hermitian and positive definite covariance matrix $\mathbf{P}$. Determine the unitary matrix $\mathbf{U}$ that diagonalizes the Hermitian matrix $\mathbf{R}$:

$$\mathbf{U}^\dagger \mathbf{R}\mathbf{U} = \mathbf{\Sigma}^2, \tag{24}$$

with $\mathbf{U}^\dagger\mathbf{U} = \mathbf{I}$. It can then be verified that the coordinate transformation

$$\mathbf{T} = \mathbf{\Sigma}^{1/2}\mathbf{U}^\dagger \mathbf{P}^{-1/2}, \tag{25}$$

simultaneously diagonalizes the $\mathbf{P}$ and $\mathbf{Q}$ matrices, transforming $\mathbf{P}$ to $\tilde{\mathbf{P}} = \mathbf{\Sigma}$ and $\mathbf{Q}$ to $\tilde{\mathbf{Q}} = \mathbf{\Sigma}$. The structure of the forcing $\mathbf{TF}$ (which is no longer unitary) and the nonnormality of the transformed operator $\tilde{\mathbf{A}}$ are such as to render the resulting $\tilde{\mathbf{P}}$ and $\tilde{\mathbf{Q}}$ matrices in these coordinates equal. The diagonal elements of $\mathbf{\Sigma}$ in descending order are called the Hankel singular values. In these coordinates the SOs and the EOFs are the same and the system is truncated at order k by retaining the first k Hankel singular values, as if the ESM theorem were applicable. In the H_∞ matrix norm (which will be introduced subsequently) the error incurred by the reduction of the system will be shown to be bounded below by the first neglected Hankel singular value and above by twice the sum of the all the neglected Hankel singular values (Glover 1984). In summary, the method is to reduce the system to a kth-order system by retaining the first k SO and EOFs, which are both equal to the first k Cartesian basis vectors in these new coordinates.

The associated truncation of $\tilde{\mathbf{A}}$ is truncation to the submatrix $\tilde{\mathbf{A}}_{11}$ of the first k rows and columns of $\tilde{\mathbf{A}}$. The reason is that in these coordinates the Lyapunov equation (10) satisfied by the covariance matrix $\tilde{\mathbf{P}} = \mathbf{\Sigma}$ is

$$\tilde{\mathbf{A}}\mathbf{\Sigma} + \mathbf{\Sigma}\tilde{\mathbf{A}}^\dagger + \mathbf{T}\mathbf{T}^\dagger = 0. \tag{26}$$

Because $\mathbf{\Sigma}$ is diagonal, $\tilde{\mathbf{A}}_{11}$ solves the Lyapunov equation:

$$\tilde{\mathbf{A}}_{11}\mathbf{\Sigma}_1 + \mathbf{\Sigma}_1\tilde{\mathbf{A}}_{11}^\dagger + \mathbf{T}_1\mathbf{T}_1^\dagger = 0, \tag{27}$$

where $\mathbf{\Sigma}_1$ is the diagonal matrix with the first k Hankel singular values and $\mathbf{T}_1$ the transformation matrix in which only the first k rows have been retained. Consequently, the reduced $k \times k$ operator $\tilde{\mathbf{A}}_{11}$ accounts for the variance, which is equal to the sum of the retained first k Hankel singular values. The reduced low-order dynamical system is

$$\frac{dz}{dt} = \tilde{\mathbf{A}}_{11}z, \tag{28}$$

where z is a k dimensional vector. The physical state is recovered by $\psi = \mathbf{T}_1^{-1}z$, where $\mathbf{T}_1^{-1}$ is the $n \times k$ submatrix of $\mathbf{T}^{-1}$ obtained by retaining the first k columns.

An important general property of the truncated system (28) is that $\tilde{\mathbf{A}}_{11}$ inherits the stability properties of $\mathbf{A}$, and is therefore stable.[4] This follows from a general property of the Lyapunov equation and the fact that in Eq. (27) the matrix $\mathbf{T}_1\mathbf{T}_1^\dagger$ is positive definite (Zhou and Doyle 1998).

[3] The square root of a matrix $\mathbf{P}$ is the matrix $\mathbf{P}^{1/2}$ with the property $\mathbf{P} = \mathbf{P}^{1/2}\mathbf{P}^{1/2}$. Clearly the square root of a diagonal matrix is the diagonal matrix consisting of the square root of the diagonal elements. If $\mathbf{P}$ is diagonalizable with eigenvalues forming the diagonal matrix $\mathbf{\Lambda}$ the square root of $\mathbf{P}$ is defined as follows: diagonalize $\mathbf{P} = \mathbf{S}\mathbf{\Lambda}\mathbf{S}^{-1}$, then the generally multivalued $\mathbf{P}^{1/2}$ is given by $\mathbf{P}^{1/2} = \mathbf{S}\mathbf{\Lambda}^{1/2}\mathbf{S}^{-1}$. If $\mathbf{P}$ is Hermitian and positive definite then the square root is unique. In this case it is numerically advantageous to calculate the square root using its singular value decomposition: $\mathbf{P} = \mathbf{S}\mathbf{\Sigma}\mathbf{S}^\dagger$. Then $\mathbf{P}^{1/2} = \mathbf{S}\mathbf{\Sigma}^{1/2}\mathbf{S}^\dagger$. For nondiagonalizable matrices a generalized square root can be defined based on the Jordan block decompositions (cf. Horn and Johnson 1991, p. 459).

[4] The balanced transformation can be also applied to unstable systems (Glover 1984).

c. Deterministic interpretation of **P** and **Q**

Because of their importance in the theory of model order reduction and their application in predictability theory we give alternative deterministic interpretations of the matrices **P** and **Q**.

The time integrated square magnitude of the response of the system to an initial perturbation is obtained by analysis of the matrix $\mathbf{Q} = \int_0^\infty e^{\mathbf{A}^\dagger t} e^{\mathbf{A}t}\, dt$. This is because the initial condition ψ_0 will evolve to $\psi(t) = e^{\mathbf{A}t}\psi_0$ with time-integrated square magnitude:

$$\int_0^\infty \psi(t)^\dagger \psi(t)\, dt = \psi_0^\dagger \left(\int_0^\infty e^{\mathbf{A}^\dagger t} e^{\mathbf{A}t}\, dt \right) \psi_0 \equiv \psi_0^\dagger \mathbf{Q} \psi_0. \quad (29)$$

It follows that eigenanalysis of the positive definite Hermitian matrix **Q** orders the initial perturbations according to the resulting time integrated response.

The covariance matrix $\mathbf{P} = \int_0^\infty e^{\mathbf{A}t} e^{\mathbf{A}^\dagger t}\, dt$ can also be given a deterministic interpretation. Consider square integrable deterministic forcings $u(t)$ that force the system for past times, that is, for $t < 0$:

$$\frac{d\psi}{dt} = \mathbf{A}\psi + u(t). \quad (30)$$

Define the functional, $J(u) = \int_{-\infty}^0 u(t)^\dagger u(t)\, dt$, to be the integrated square magnitude of the imposed forcing, $u(t)$. We wish to determine the forcing $\tilde{u}(t)$ of minimal $J(u)$ that produces a given state ψ_0. Luenberger (1969) showed that this minimal $J(\tilde{u})$ is determined from **P** through the relation: $J(\tilde{u}) = \psi_0^\dagger \mathbf{P}^{-1} \psi_0$. To show this we form

$$\frac{d\psi(t)^\dagger \mathbf{P}^{-1} \psi(t)}{dt} = \psi(t)^\dagger (\mathbf{A}^\dagger \mathbf{P}^{-1} + \mathbf{P}^{-1}\mathbf{A}) \psi(t) + u(t)^\dagger \mathbf{P}^{-1} \psi(t) + [\mathbf{P}^{-1}\psi(t)]^\dagger u(t), \quad (31)$$

then integrate using (10) and complete the square to obtain

$$\begin{aligned} \psi_0^\dagger \mathbf{P}^{-1} \psi_0 &= \int_{-\infty}^0 u(t)^\dagger u(t)\, dt \\ &\quad - \int_{-\infty}^0 [u(t) - \mathbf{P}^{-1}\psi(t)]^\dagger [u(t) - \mathbf{P}^{-1}\psi(t)]\, dt \\ &= J(u) - \int_{-\infty}^0 [u(t) - \mathbf{P}^{-1}\psi(t)]^\dagger \\ &\quad \times [u(t) - \mathbf{P}^{-1}\psi(t)]\, dt. \end{aligned} \quad (32)$$

This shows that the integrated square magnitude $J(u)$ for all past forcings $u(t)$ that produce state ψ_0 satisfies the inequality

$$J(u) \geq \psi_0^\dagger \mathbf{P}^{-1} \psi_0. \quad (33)$$

Inspection of (32) also reveals that the minimum integrated forcing is attained for

$$\tilde{u}(t) = \mathbf{P}^{-1}\psi(t). \quad (34)$$

Substituting this minimum forcing in (30) we obtain that $\psi(t)$ obeys for $t < 0$ the equation

$$\frac{d\psi}{dt} = (\mathbf{A} + \mathbf{P}^{-1})\psi, \quad (35)$$

which for $t < 0$ and $\psi(0) = \psi_0$ has the solution

$$\psi(t) = e^{(\mathbf{A}+\mathbf{P}^{-1})t}\psi_0 = e^{-\mathbf{P}\mathbf{A}^\dagger \mathbf{P}^{-1} t}\psi_0 = \mathbf{P} e^{-\mathbf{A}^\dagger \mathbf{P}^{-1} t}\psi_0, \quad (36)$$

where (10) was used to write $\mathbf{A} + \mathbf{P}^{-1} = -\mathbf{P}\mathbf{A}^\dagger\mathbf{P}^{-1}$. This proves that the minimum forcing that produces ψ_0 is

$$\tilde{u}(t) = e^{-\mathbf{A}^\dagger \mathbf{P}^{-1} t}\psi_0, \quad (37)$$

and the time-integrated square magnitude of this forcing is $\psi_0^\dagger \mathbf{P}^{-1} \psi_0$.

We pause now to consider an implication of this result for forecast. Consider perturbations to the initial state, which represent uncertainty in knowledge of the initial state. Singular value decomposition of the forward propagator of the perturbation forecast equations provides the structures contributing to forecast uncertainty after a chosen interval of time and orders these structures in their contribution to producing this uncertainty. As we have seen, truncating the singular value decomposition of the propagator provides an explicit optimal order reduction of the propagator of forecast error over the chosen interval of time. In contrast (37) provides the minimum past forcing, $\tilde{u}(t)$, distributed over space and time that produces a given perturbation at $t = 0$. From the forecast perspective examining $\tilde{u}(t)$ provides information on sources of uncertainty in the forecast distributed over space and time. One implication of this result is that the forecast at $t = 0$ would be most affected by model error distributed in space and time as $\tilde{u}(t)$ in (37) with ψ_0 chosen to be the eigenfunction of $\mathbf{P}^{-1}$ with smallest eigenvalue, that is, the first EOF. Another implication is that the first EOF is the most likely error to be produced by the forecast model. A third implication is observing resources may be best deployed at locations and times indicated by this $\tilde{u}(t)$.

We have argued that accurate representation of the dynamics requires that both the amplifying structures and the resulting responses be resolved. Based on the interpretations of matrices **P** and **Q** we can quantify this requirement by demanding that the structures ψ_0 contributing maximally to the functional,

$$H[\psi_0] = \frac{\psi_0^\dagger \mathbf{Q} \psi_0}{\psi_0^\dagger \mathbf{P}^{-1} \psi_0}, \quad (38)$$

be well approximated. The $H[\psi_0]$ is referred to as the Hankel quotient hereafter. For initial condition ψ_0, $H[\psi_0]$ is the ratio of the time-integrated square magnitude of the response of the system to ψ_0 (for $t > 0$) to the time-integrated square magnitude of the minimum forcing (for $t < 0$) required to produce ψ_0. This quotient measures in a joint manner both sensitivity to initial conditions (determined by the **Q** matrix) and the inte-

grated forcing required to produce initial conditions (determined by the **P** matrix). The stationary values of this quotient are the eigenvalues of **PQ** which are also the eigenvalues of **R** defined in (23) and used to cast the system in balanced coordinates by simultaneously diagonalizing **P** and **Q**. The stationary values of the Hankel quotient are thus the squares of the Hankel singular values, which indicate the accuracy of the model order reduction.

d. The Hankel operator representation of the dynamics

The ESM theorem solves the problem of optimally truncating the propagator in the Euclidean norm at any fixed time, but it does not directly provide a solution to the problem of truncating the underlying dynamical system. In order to address this question we must apply the ESM theorem to the linear map that describes the dynamics, that is, the map that connects all forcings to all responses. The reduced model provides another linear map connecting all forcings to all responses. The closeness of these two maps determines the accuracy of the approximation, which from the ESM theorem can be as small as the largest singular value of the linear map which is removed in the truncation. Therefore, in order to determine the optimal truncation and assess its performance we must first find the singular values of the linear map describing the dynamics.

The system dynamics compose a transformation of past forcings to future responses. The time translational symmetry of the linear autonomous system can be exploited to fix the time origin at $t = 0$ so that the dynamics can be viewed as a transformation of past forcings over t on $(-\infty, 0]$ to future responses over t on $(0, \infty)$. This linear transformation is called the Hankel operator and it is defined as the linear map $\mathbf{H}:\mathcal{L}^2(-\infty, 0] \rightarrow \mathcal{L}^2(0, \infty)$, from square integrable forcings, $u(t)$, in the interval $(-\infty, 0]$ to square integrable responses of the system, $\psi(t)$, in the interval $[0, \infty)$:

$$\mathbf{H}u(T) = \psi(t) = e^{\mathbf{A}t} \int_{-\infty}^{0} e^{-\mathbf{A}s} u(s)\, ds. \tag{39}$$

Note that the Hankel operator can be found by the action of two linear maps: **N**, which maps past forcings to the state of the system at $t = 0$, called ψ_0; and **M**, which maps ψ_0 to its state at later time, $\psi(t)$; that is, the Hankel operator is $\mathbf{H} = \mathbf{MN}$, where

$$\psi(t) = \mathbf{M}\psi_0 = e^{\mathbf{A}t}\psi_0, \quad \text{and} \tag{40}$$

$$\psi_0 = \mathbf{N}u = \int_{-\infty}^{0} e^{-\mathbf{A}s} u(s)\, ds. \tag{41}$$

Both **N** and **M** are operators of rank n where n is the dimension of the system and consequently the Hankel operator is also of rank n.

In order to obtain the singular values of the linear operator **H** the inner product of the function spaces it operates on must be specified. The Euclidean inner product will be used: for the responses, which are functions limited to $t > 0$, it is $(f, g) = \int_0^\infty f(t)^\dagger g(t)\, dt$, while for the forcings, which are functions limited to $t < 0$, it is $(f, g) = \int_{-\infty}^0 f(t)^\dagger g(t)\, dt$. The squares of the singular values are obtained as the stationary values of the ratio of the square of the norm of the responses to the square of the norm of the forcings:

$$H[\psi_0] = \frac{\displaystyle\int_0^\infty \psi(t)^\dagger \psi(t)\, dt}{\displaystyle\int_{-\infty}^0 u(t)^\dagger u(t)\, dt}, \tag{42}$$

which can be found by sequential maximization over $u(t)$ in the subspace that is orthogonal to that spanned by the previously obtained stationary forcings. Instead of maximizing this quotient over $u(t)$ we can equivalently maximize it over the initial states ψ_0 and specify $\tilde{u}(t)$ to be the forcing of minimum square amplitude, which by (37) is $\psi_0 \mathbf{P}^{-1} \psi_0$. It is now clear that the stationary values of the Hankel quotient $H[\psi_0]$, introduced in (38), are also the stationary values of quotient (42), and that the singular values of the Hankel operator are the Hankel singular values Σ, which were defined as

$$\Sigma = \sqrt{\Lambda(\mathbf{PQ})} = \sqrt{\lambda(\mathbf{P}^{1/2}\mathbf{Q}\mathbf{P}^{1/2})}, \tag{43}$$

where $\Lambda(\mathbf{B})$ denotes the eigenvalues of a matrix **B**. We can write the singular value decomposition of the Hankel operator as

$$\mathbf{H} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\dagger, \tag{44}$$

where $\boldsymbol{\Sigma}$ is the diagonal matrix of the Hankel singular values Σ; and **U**, **V** are unitary operators.

Just as for linear transformations in which the L_2 norm of vectors induces a matrix norm equal to the largest singular value of the matrix representation of the transformation, so also here we can define the L_2 induced norm of the linear operator **H** as the largest Hankel singular value. We can then define the Hankel norm of the dynamical operator matrix **A** as $\|\mathbf{A}\|_H = \|\mathbf{H}\|_2$; that is, the Hankel norm of a dynamical operator matrix is the L_2 induced norm of its associated Hankel operator. From the ESM theorem the optimal truncation of **H** in the L_2 induced norm of its Hankel operator is the $\mathbf{H}_k$ of rank k with Hankel singular values equal to the first k Hankel singular values.

We have truncated the singular value decomposition of the Hankel operator but if we instead truncated the singular value decomposition of **M** in (40), which maps initial conditions to responses, instead of the Hankel operator, then we would retain the top **M**-singular values of the system and obtain again by appealing to the ESM theorem an optimal truncation of the system in the associated **M** norm of **A**. Alternatively, optimality in the ESM sense could be obtained if the dynamics were trun-

cated in the **N** map [cf. (41)]. So there is a need to measure and compare the accuracy of truncations in a norm that does not depend on the representation of the dynamics. One such norm is the time integral of the L_2 norm of the propagator, another would be the H_∞ norm, which is defined for a stable system as the maximum singular value of the resolvent over all real frequencies ω, that is,

$$\|\mathbf{A}\|_\infty \equiv \max_{\omega \in \Re} \sigma_{\max}(\mathbf{R}(\omega)), \tag{45}$$

where $\sigma_{\max}$ is the largest singular value of the resolvent (cf. FI96):

$$\mathbf{R}(\omega) = (i\omega\mathbf{I} - \mathbf{A})^{-1}. \tag{46}$$

The H_∞ norm is the optimal response of the system to sinusoidal forcing. This generally accepted notation for the H_∞ norm of a dynamical system governed by dynamical operator **A** should not be confused with the notation of the L_∞ matrix norm of **A**, which is the largest row sum of the entries of the matrix.

The relation among the Hankel norm $\|\mathbf{A}\|_H$ of an operator **A**, which is equal to its largest Hankel singular value Σ_1, and other norms of the operator (Zhou and Doyle 1998) is

$$\|\mathbf{A}\|_H = \Sigma_1 \le \|\mathbf{A}\|_\infty \le \int_0^\infty \|e^{\mathbf{A}t}\|_2 \, dt \le 2 \sum_{i=1}^{n} \Sigma_i. \tag{47}$$

From (47) and the ESM theorem it follows that the approximation error for reduction to rank k as evaluated in the H_∞ norm must be at least

$$\Sigma_{k+1} \le \|\mathbf{A} - \mathbf{A}_k\|_\infty, \tag{48}$$

where Σ_{k+1} is the first neglected Hankel singular value of **A**. The method for calculating the H_∞ error is described in appendix A. As measured in the H_∞ norm no reduction technique can give an approximation better than this bound. Some reduction techniques also come with an upper bound on the approximation error but in practice in choosing the approximation technique we must also consider its computational complexity. The balanced truncation introduced in the previous section, although it is not Hankel and is suboptimal in the Hankel norm, is a good compromise. It gives an error bound that can be seen from (47) to be twice the sum of the neglected Hankel singular values (Zhou and Doyle 1998); that is, for balanced truncations the error satisfies the inequality

$$\Sigma_{k+1} \le \|\mathbf{A} - \mathbf{A}_k\|_\infty \le 2 \sum_{i=k+1}^{n} \Sigma_i, \tag{49}$$

which is easily computed. In practice the error is found to be close to the lower bound.[5]

[5] Optimal Hankel norm approximations with tighter infinity norm error bounds can be constructed iteratively (Glover 1984), but they are cumbersome and probably unsuitable for implementation in the large systems that arise in flow problems. They involve N^4 operations compared to the N^3 operations required for balanced truncation and reduce the upper bound of the error in (49) by a factor of 2.

3. Reduced-order approximation of the linear operator associated with shear flow

Consider perturbations to a time-independent zonally homogeneous barotropic Couette flow, $U(y) = y$, confined to a channel in which y is the northward direction and x is the zonal direction. Choosing the inverse of the mean shear and the half-channel width as characteristic time- and space scales, the nondimensional barotropic vorticity equation for the meridionally and temporally varying component of the streamfunction $\psi(x, y, t) = \psi(y, t)e^{ikx}$ is given by

$$\frac{\partial \mathbf{D}^2\psi}{\partial t} = -iky\mathbf{D}^2\psi + \frac{1}{\mathrm{Re}}\mathbf{D}^4\psi, \tag{50}$$

where k is the zonal wavenumber. The operator $\mathbf{D}^{2n}$ is defined as

$$\mathbf{D}^{2n} \equiv \left(\frac{d^2}{dy^2} - k^2\right)^n, \tag{51}$$

and these continuous operators have been cast in finite-dimensional matrix form by discretizing and approximating the derivatives with central differences. If a grid-point discretization is sufficiently fine, we are assured that the solution of the matrix approximation approaches asymptotically the solution of the continuous system (50) (Ince 1956). The Reynolds number is $\mathrm{Re} = L^2/(\alpha\nu)$ where L is the channel half-width, α the shear, and ν the viscosity. The boundary conditions at the channel walls are $\psi(\pm 1, t) = 0$ and the nonslip condition $\psi_y(\pm 1, t) = 0$. The perturbation dynamics are governed by the classical Orr–Sommerfeld equation (Drazin and Reid 1981).

We transform to generalized velocity coordinates, ϕ, defined as $\phi = (-\mathbf{D}^2)^{1/2}\psi$. It can be shown that in the Euclidean inner product and the chosen boundary conditions that $-\mathbf{D}^2$ is Hermitian and positive definite and its square root, $(-\mathbf{D}^2)^{1/2}$, is uniquely defined. The square root of the finite-dimensional matrix approximation to $(-\mathbf{D}^2)^{1/2}$ is easily obtained as explained in footnote (3). In generalized velocity coordinates the Euclidean inner product $\phi^\dagger\phi$ is proportional to perturbation energy density. The dynamical system in generalized velocity is

$$\frac{d\phi}{dt} = \mathbf{A}\phi, \tag{52}$$

with the dynamical operator[6]

$$\mathbf{A} = (-\mathbf{D}^2)^{1/2}\mathbf{D}^{-2}\left(-iky\mathbf{D}^2 + \frac{1}{\mathrm{Re}}\mathbf{D}^4\right)(-\mathbf{D}^2)^{-1/2}. \tag{53}$$

Discretization on 100 grid points approximates well the spectrum and the evolution dynamics of the full continuous operator at $\mathrm{Re} = 800$ (convergence to the con-

[6] The inverse of the matrix representation of $\mathbf{D}^2$ is defined to be the matrix approximation of the operator $\mathbf{D}^{-2}$.

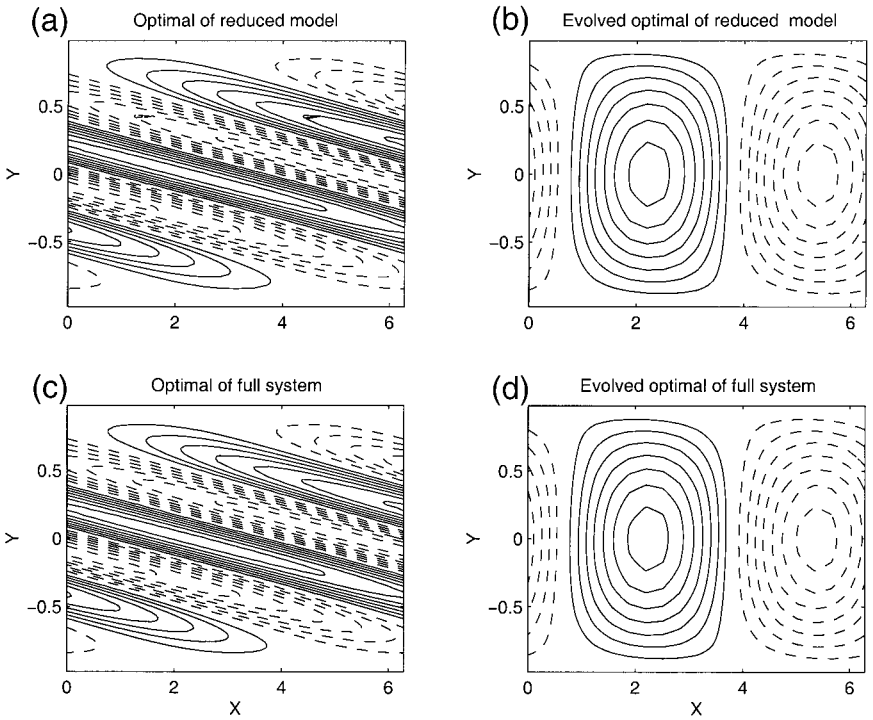

FIG. 5. (left) The structure of the streamfunction of the optimal perturbation that leads to the greatest energy growth at $t = 10$, and the evolved optimal streamfunction, (right) the structure that these optimals evolve into at the optimizing time $t = 10$. The bottom panels are for the full system while the top panels are for the order 10 balanced truncation. The difference in the structure of the streamfunction is less than 1%.

tinuous operator was verified by doubling the number of discretization points), so at the start we consider (52) to be a 100 dimensional dynamical system. The performance of various truncations will be assessed for Re $= 800$ and the single wavenumber $k = 1$.

a. Balanced truncations

We start with the balanced truncation of the full system. The operator $\mathbf{Q}$ is calculated from (12) and operator $\mathbf{P}$ from (10). The eigenfunctions of $\mathbf{Q}$ are the stochastic optimals. The first two SOs are shown in the two top panels of Fig. 1. The eigenfunctions of $\mathbf{P}$ are the EOFs of the dynamical system forced white in space and time. The first two EOFs are shown in the bottom panels of Fig. 1. The first 8 SOs are responsible for exciting 90% of the maintained variance. In this example the first 8 EOFs also account for 90% of the maintained variance. The eigenvalues (or singular values) of $\mathbf{P}$ and $\mathbf{Q}$ are shown in Fig. 2 along with the Hankel singular values obtained by applying coordinate transformation $\mathbf{T}$ given by (25). This transformation renders $\mathbf{P}$ and $\mathbf{Q}$ diagonal. The diagonal elements of these matrices are the Hankel singular values, and in these coordinates the singular vectors are the canonical basis. A k rank balanced truncation is obtained by retaining the first k singular vectors of the diagonalized $\mathbf{P}$ and $\mathbf{Q}$ while setting the remaining singular vectors to zero. The resulting dynamical operator $\tilde{\mathbf{A}}_{11}$ consists of the first k columns and rows of $\tilde{\mathbf{A}}$.

The low-order dynamical system is constructed so as to reproduce as closely as possible the time development of the original system. This implies that it will also closely approximate the frequency response of the operator. The frequency response can be quantified by the magnitude of the resolvent (46). The integral of trace $[\mathbf{R}^{\dagger}(\omega)\mathbf{R}(\omega)]$ over all frequencies is the variance maintained under spatially and temporally white forcing (cf. FI96). The magnitude of the frequency response depends in part on the damping of the modes of the operator, but for nonnormal systems it also depends on the nonnormality of the operator, which can lead to variance far in excess of that anticipated from the rate of damping alone (cf, Farrell and Ioannou 1994; Ioannou 1995). The spectrum of the 10th-order operator obtained from balanced truncation and that of the original system are shown in the bottom panels of Fig. 3. We see that these operator spectra are very different, they do not even have the same least damped mode, and yet they produce nearly identical frequency responses, especially for the range of frequencies for which the response is large. The truncated system also reproduces very accurately the structure and growth of individual perturbations. In the top panel of Fig. 4 the four largest singular values of the propagator of both the original system and its

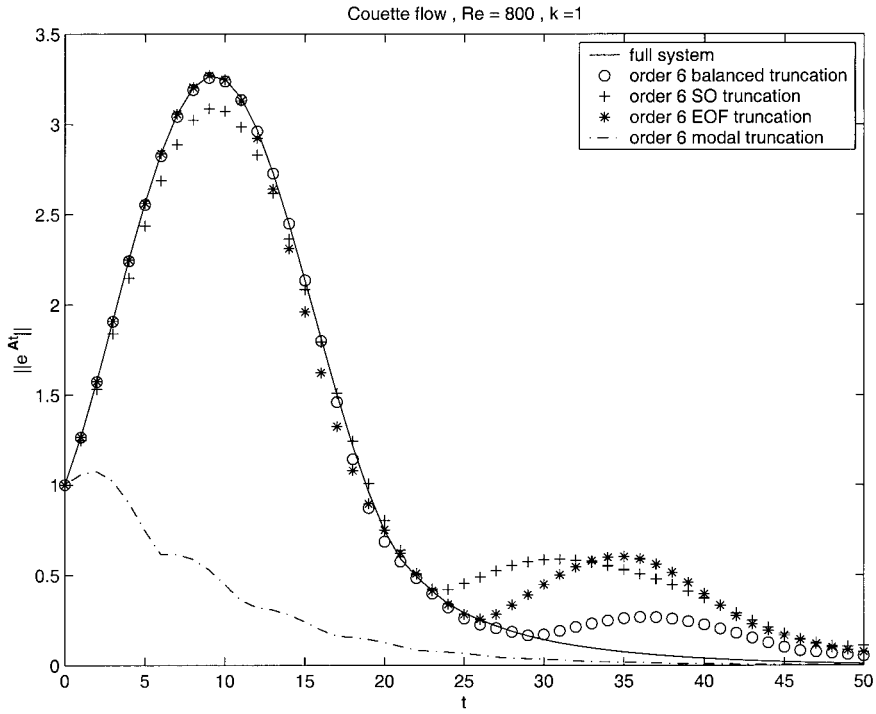

FIG. 6. Optimal growth, $\|e^{\mathbf{A}t}\|$, as a function of time for the Couette flow at Re = 800 and k = 1. Shown is the optimal growth for the full system (solid); optimal growth produced by an order 6 approximate system obtained by balanced truncation of the full system (circles); the optimal growth for an order 6 approximate system in which the top 6 stochastic optimals have been retained (crosses); the optimal growth by an order 6 approximate system in which the top 6 EOFs have been retained (stars); and the optimal growth by an order 6 approximate system in which the top 6 least damped modes have been retained (dash–dot).

order 10 truncation are shown as a function of time. The time development of these growing disturbances in the truncated and the untruncated system is indistinguishable. The structures of the optimals in the two systems are also indistinguishable. The structure of the optimal perturbation that leads to the largest energy growth in the 10th-order reduced system and in the full system are shown in Fig. 5. The evolved optimals are also shown in Fig. 5. The structures are indistinguishable. The maximum possible error in the truncation can be found by calculating the H_∞ response of the error system (see appendix A), which identifies the excitation sinusoidal in time that leads to the greatest discrepancy between the response of the original system and its approximant. We have seen in the previous section that the H_∞ error is bounded below by the first neglected Hankel singular value and above by twice the sum of all the neglected Hankel singular values, as in (49). For the order 10 balanced truncation the bounds are not tight: the lower bound is $\Sigma_{11} = 1.2$ while the upper bound is 30.5. In the bottom panel of Fig. 5 is shown the largest singular value of the resolvent of the error system as a function of frequency (see appendix A). The maximum of this curve gives the H_∞ error of the approximation which is calculated to be $\|\mathbf{A} - \mathbf{A}_{10}\|_\infty = 2.2$. This shows that the balanced truncation is almost optimal in the sense of having error in the H_∞ norm close to the lower bound, which in this case is 1.2. We have found this to be true in all cases we have investigated.

b. Comparison with other truncations

We compare the balanced truncation with truncations in which only the top EOFs are retained, and to truncations in which only the top stochastic optimals are retained. We also compare to modal reductions in which the first least damped eigenmodes of the system are retained. These truncations correspond to Galerkin projection of the dynamics onto EOFs, SOs, and eigenvectors, respectively.

Truncations based on the top EOFs are equivalent to approximating the $\mathbf{N}$ map in (41) because the singular values of $\mathbf{N}$ coincide with those of $\mathbf{P}$. The reduction follows the same steps as for the balanced truncations: introduce a change of coordinates $\mathbf{T}$ but with $\mathbf{T} = \mathbf{U}^\dagger$ where $\mathbf{U}$ is the unitary matrix of the eigenvectors of $\mathbf{P}$ arranged in columns; that is, the columns of $\mathbf{U}$ are the EOFs of the system when forced with spatially and temporally white noise. This transformation renders $\mathbf{P}$ diagonal and transforms the dynamical operator to $\tilde{\mathbf{A}}$. A k order system is obtained by limiting the dynamics to

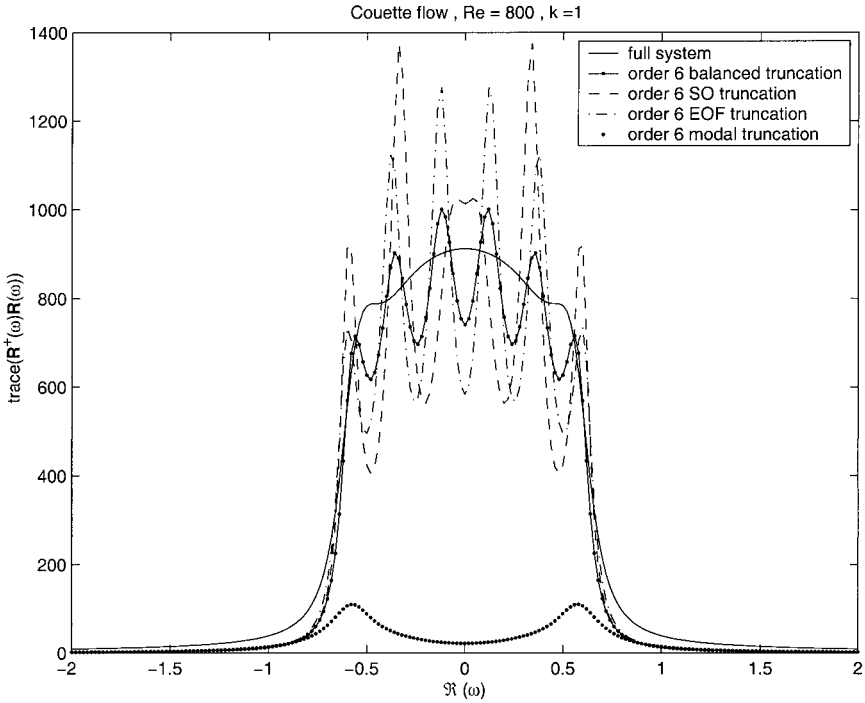

FIG. 7. Power spectrum, trace $[\mathbf{R}^\dagger(\omega)\mathbf{R}(\omega)]$, where $\mathbf{R}(\omega)$ is the resolvent (46) of the operator. Shown is the power spectrum of the full system (solid); power spectrum of an order 6 approximate system obtained from the full system by balanced truncation (line–dot); the power spectrum of an order 6 approximant obtained by retaining the top 6 SO's (dash); the power spectrum of an order 6 approximant obtained by retaining the top 6 EOF's (dash–dot); and the power spectrum of an order 6 approximant obtained by retaining the 6 least damped modes (dot). The case is Couette flow at Re = 800 and $k = 1$.

$\tilde{\mathbf{A}}_k$, which is equal to the first k columns and rows of $\tilde{\mathbf{A}}$. Because in these coordinates the truncated system satisfies the following Lyapunov equation:

$$\tilde{\mathbf{A}}_{11}\mathbf{S}_k + \mathbf{S}_k\tilde{\mathbf{A}}_{11}^\dagger + \mathbf{T}_1\mathbf{T}_1^\dagger = 0, \tag{54}$$

where $\mathbf{S}_k$ is the diagonal matrix with the first k eigenvalues of $\mathbf{P}$, and $\mathbf{T}_1$ is the submatrix of $\mathbf{T}$ with the first k rows, $\tilde{\mathbf{A}}_{11}$ inherits the asymptotic stability of the full operator. Just as for balanced truncations, any reduction of the dynamics to the top EOFs leads to a dynamical system with guaranteed asymptotic stability.

We proceed similarly with reduction of order retaining the top stochastic optimals. This is equivalent to truncating the map $\mathbf{M}$ in (40) because the singular values of $\mathbf{M}$ are the same as those of $\mathbf{Q}$. For this reduction we transform with $\mathbf{T} = \mathbf{U}^\dagger$ where the columns of $\mathbf{U}$ are the stochastic optimals, which are the eigenvectors of $\mathbf{Q}$. We truncate the dynamical system by retaining the desired number of columns and rows of the transformed dynamical operator. Again, because the reduced dynamical operator satisfies a Lyapunov equation, the stability of the approximant is guaranteed.

We consider now reduction of the original system to an order 6 system. The optimal growth is shown as a function of time for a balanced truncation, an EOF based truncation, a SO based truncation, and a modal truncation in Fig. 6. All these low-order systems have dimension 6. We observe that, despite the small dimensionality of the low-order system, all the reduced systems perform rather well, except the one that was obtained from modal reduction. It is also clear that among the reductions the balanced is the best, followed by the EOF based reduction, which in turn is better than the SO based.[7] The power spectrum of the response of the reduced operators to white stochastic forcing, given by trace $[\mathbf{R}^\dagger(\omega)\mathbf{R}(\omega)]$, where $\mathbf{R}(\omega)$ is the resolvent of the corresponding reduced operator, are compared in Fig. 7. Again the balanced reduction is closest to the response of the full system. Also again the EOF reduction performs better than the SO based reduction. The eigenmode reduction at this order provides a poor approximation to the dynamics. The truncated eigenmode system inherits the first k elements of the spectrum and

[7] Note that all the truncations break down at approximately the same time. Because a small number of modes has been retained in the approximation the power spectrum of the reduced system shown in Fig. 7 reveals a natural beat frequency related to the difference frequency of the poles of the truncated dynamical operators. The unapproximated operator lacks these undulations. The breakdown of the approximation occurs at the beat period of the difference in the maxima in the frequency response which is at the difference frequency of the poles.

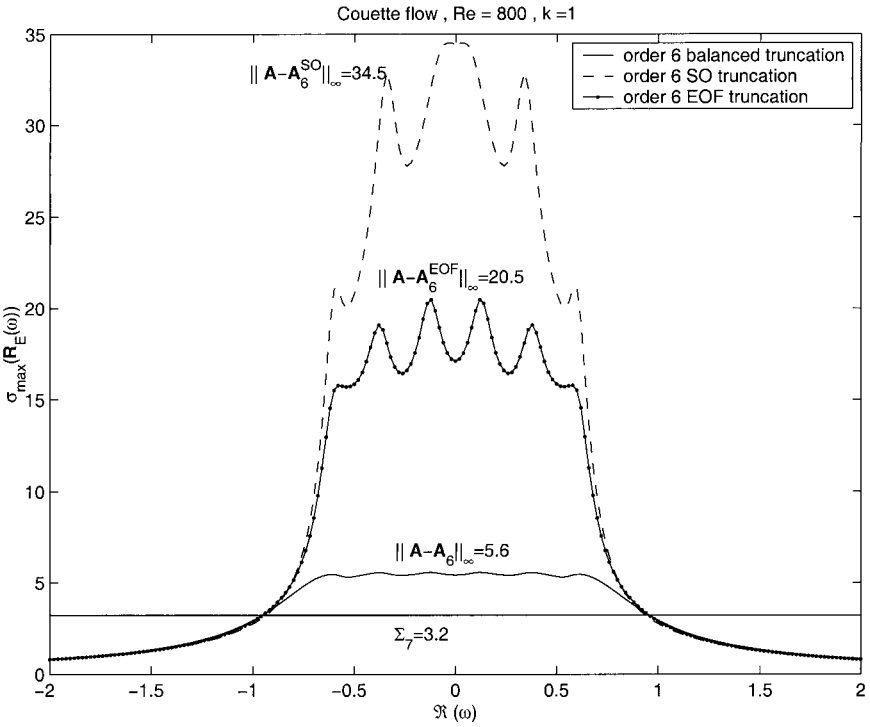

FIG. 8. Maximum singular value of the order 6 error systems $\mathbf{A} - \mathbf{A}_6$ as a function of frequency. Maxima of these curves gives the H_∞ error of the approximation. Shown is an order 6 balanced truncation of $\mathbf{A}$, which is seen to give $\|\mathbf{A} - \mathbf{A}_6\|_\infty = 5.6$ (solid). This is very close to the theoretical minimum error, which is equal to the first neglected Hankel singular value, $\Sigma_7 = 3.2$, indicated by the horizontal line. Also shown is an order 6 approximation in which the top 6 SOs have been retained, which gives an error $\|\mathbf{A} - \mathbf{A}_6^{\text{SO}}\|_\infty = 34.5$ (dash) and an order 6 approximation in which the top 6 EOFs have been retained, giving an error $\|\mathbf{A} - \mathbf{A}_6^{\text{EOF}}\|_\infty = 20.5$ (line–dot).

the nonnormality of the first k eigenvectors of the original system, so in order to approximate the system well a greater number of modes is needed; enough so that all the important nonorthogonal eigenvectors are included. The balanced truncations exploit the freedom to relocate the spectrum in order to better approximate the maintained variance, and this added freedom results in improved performance. Of course, as the dimension of the reduced system increases the differences among the various reduction methods decrease, but in all cases we find that the balanced truncation outperforms the other reduction methods. We can quantify the performance of the various reductions best by calculating the H_∞ norm of the truncation error. As discussed earlier this error must exceed the first neglected Hankel singular value, which for reduction to an order 6 system is $\Sigma_7 = 3.2$ while according to (49) the error cannot exceed 48.6. The maximum singular value of the resolvent of the error system for the three reductions, balanced, EOF based, and SO based, are compared in Fig. 8. The maximum of the corresponding curves over all real frequencies determines the H_∞ norm of the error of the various reductions. For the balanced reduction we find in this way that $\|\mathbf{A} - \mathbf{A}_6\|_\infty = 5.6$, which is nearly optimal. The error for the EOF based reduction is calculated to be 20.5 while for the SO optimal based truncation it is 34.5. The improved performance of the balanced reduction over the other reduction methods is of particular importance when we want to reduce the system as much as possible.

c. *Reductions based on balancing the optimal vectors*

Up to this point we have presented methods of model order reduction based on retaining a small number of EOFs and SOs. We saw that the best way to proceed was to transform into a coordinate system in which the EOFs and SOs are identical and diagonal, and then truncate the system in these balanced coordinates. This procedure guarantees stability and comes with explicit error bounds and in practice is found to approximate the dynamics well. However, because it is numerically costly to obtain the $\mathbf{P}$ and $\mathbf{Q}$ matrices for large systems, one may wonder if similar performance could be obtained with truncations based on retaining the first optimals and evolved optimals (right and left singular vectors) of the propagator at a specific time. From the earlier discussion it is clear that such a procedure does not formally approximate the dynamics, since by retaining a few of these optimal vectors we only approximate the

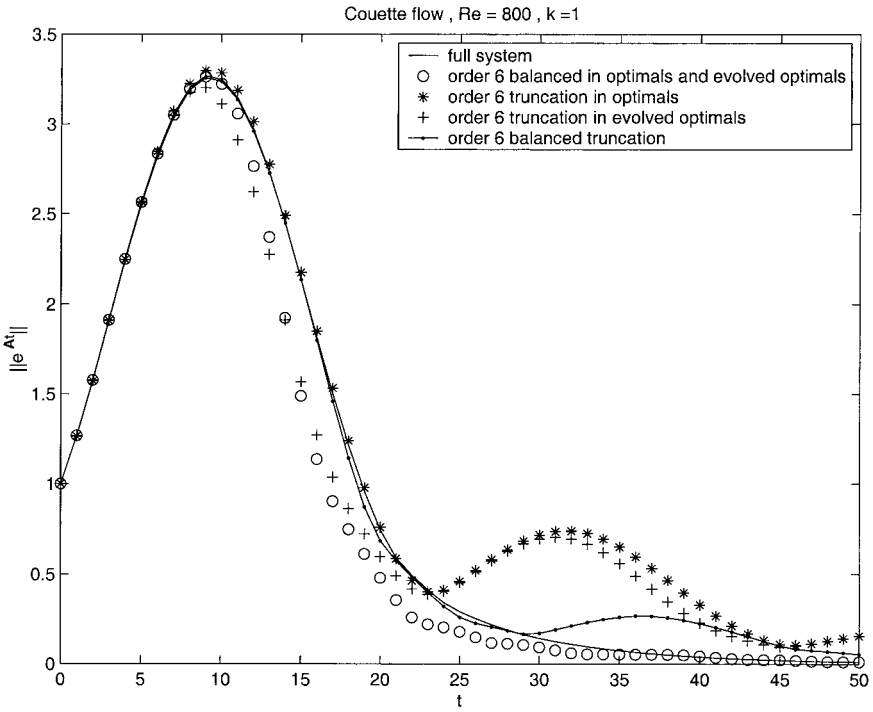

FIG. 9. Optimal growth, $\|e^{\mathbf{A}t}\|$, as a function of time for the Couette flow at Re = 800 and $k = 1$. Shown is the optimal growth for the full system (solid); the optimal growth by an order 6 approximate system obtained by balanced truncation of the full system (line dot); the optimal growth by an order 6 approximate system obtained by balancing the optimals and evolved optimals associated with the propagator of the system at $t = 10$ (circles); the optimal growth by an order 6 approximate system in which the top 6 evolved optimals of the propagator for $t = 10$ have been retained (stars); and a similar truncation in which the 6 top optimal perturbations of the propagator for $t = 10$ have been retained (crosses).

propagator at a specific time. On the other hand **P** and **Q** are the integral over time of $\mathcal{P}(t) = e^{\mathbf{A}t}e^{\mathbf{A}^\dagger t}$, and of $\mathcal{Q}(t) = e^{\mathbf{A}^\dagger t}e^{\mathbf{A}t}$, and one approximation of **P** and **Q** is simply to retain one element of the Riemann sum of this integration, which within a constant factor would be the value $\mathcal{P}(t)$ and $\mathcal{Q}(t)$ for some particular time *t*. If the reduction is done by retaining the leading eigenvectors of $\mathcal{P}(t)$ and $\mathcal{Q}(t)$ this is equivalent to retaining the dominant optimal and evolved optimal vectors, since the eigenvectors of $\mathcal{P}(t)$ are the evolved optimals of the propagator at time *t*, and those of $\mathcal{Q}(t)$ are the optimals of the propagator at this time.

We have defined a coordinate transformation **T** that simultaneously diagonalizes the **P** and **Q** matrices; if the same procedure is followed to bring $\mathcal{P}(t)$ and $\mathcal{Q}(t)$ into a balanced representation in which both are diagonal, then by the same steps we can reduce the system in this balanced representation. However, because now $\mathcal{P}(t)$ and $\mathcal{Q}(t)$ do not satisfy a Lyapunov equation the reduced system does not have guaranteed stability, and indeed for severe truncations and for small *t* the reduced system is found to be unstable. Still one may hope to find a choice of optimizing time for which the performance of a balanced truncation of the optimals and evolved optimals is good.

The performance of such a truncation for optimizing time $t = 10$, which corresponds to the time of the global optimum, is shown in Fig. 9. We plot in this figure the optimal growth as a function of time for an order 6 truncation of the dynamics in which the optimal and evolved optimals are truncated in their balanced representation, which the truncation proceeds by considering only the optimal vectors (as was done previously for the SOs) and also in which only the evolved optimals are considered in the truncation (as was done previously for the EOFs). For comparison we also plot the optimal growth of the full system and of an order 6 balanced truncation. We observe first that the unbalanced reductions based on the optimals or the evolved optimals produce approximations inferior to the reduction obtained by balancing the optimals and the evolved optimals. While the balanced truncation based on the optimal and evolved optimal vectors is inferior to that obtained by balanced truncation of **P** and **Q**, the difference is rather small; the power spectrum of the response of the system to white noise stochastic forcing is shown in Fig. 10.

It is of interest to estimate the H_∞ error of truncations obtained by using the optimal vectors at various optimizing times as a function of the optimizing time, and

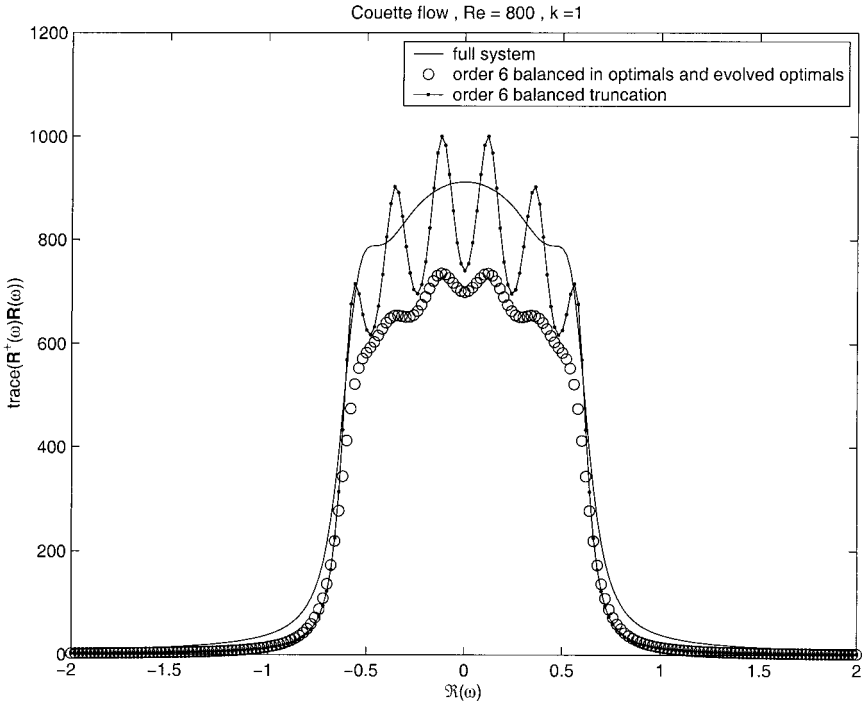

FIG. 10. Power spectrum given by trace $[\mathbf{R}^{\dagger}(\omega)\mathbf{R}(\omega)]$, where $\mathbf{R}(\omega)$ is the resolvent of the operator (46). Shown is the power spectrum of the full system (solid); power spectrum of an order 6 approximate system obtained from the full system by balanced truncation (line–dot); the power spectrum of an order 6 approximant obtained by balancing the optimals and evolved optimals associated with the propagator of the system at $t = 10$ (circles). The case is Couette flow at Re $= 800$ and $k = 1$.

compare it to the H_∞ error of the balanced truncation. For an order 6 balanced truncation the H_∞ error was found in the previous section to be 5.2, while the minimum H_∞ error possible is equal to the 7th Hankel singular value, which is $\Sigma_7 = 3.2$. The H_∞ error as a function of optimizing time for balanced truncation of the optimals and evolved optimals of the propagator at the optimizing time, and truncations using only the optimals or only the evolved optimals of the propagator at the optimizing time is shown in Fig. 11. The minimum error of 6.4 is attained by the balanced truncation of the optimals and evolved optimals of the propagator for $t = 10$. Truncations based only on the optimals or only on the evolved optimals do not perform particularly well. This result suggests that when balanced truncations are not feasible, truncations based on balancing the optimal and evolved optimal vectors can perform very well if the optimizing time is selected to be near the time of the global optimal at which time evolved optimals resemble the EOFs of the system.

4. Discussion and conclusions

We have examined methods for reducing the dimension of a linear autonomous dynamical system while retaining an accurate approximation of the dynamics. Reduction of the dimension of non-normal systems such as are associated with perturbation dynamics in fluid flow require, in order to accurately model the dynamics, that the disturbances spanning the perturbation variance and the distinct perturbations producing the variance be retained in a balanced manner when the system is truncated. This can be most efficiently accomplished by expressing the dynamics in the finite-dimensional Hankel operator form, balancing the operator between stochastic optimals and EOFs, and then truncating the Hankel operator in the balanced representation. This method of truncating assures stability of the truncated system if the original system is stable. It also has error in the H_∞ norm bounded between the first neglected Hankel singular value and twice the sum of the neglected Hankel singular values. It is possible to derive a still more accurate reduced-order model by using a more complex iteration algorithm with error bounded above by only the sum of the neglected singular values (Glover 1984) but this method appears to be of limited utility in fluid flow problems especially given that straightforward truncation of the balanced Hankel operator is found to be nearly optimal in practice.

Implementation of the method of Hankel norm balanced truncation requires only standard matrix algebra and is directly applicable for dimensions up to a few

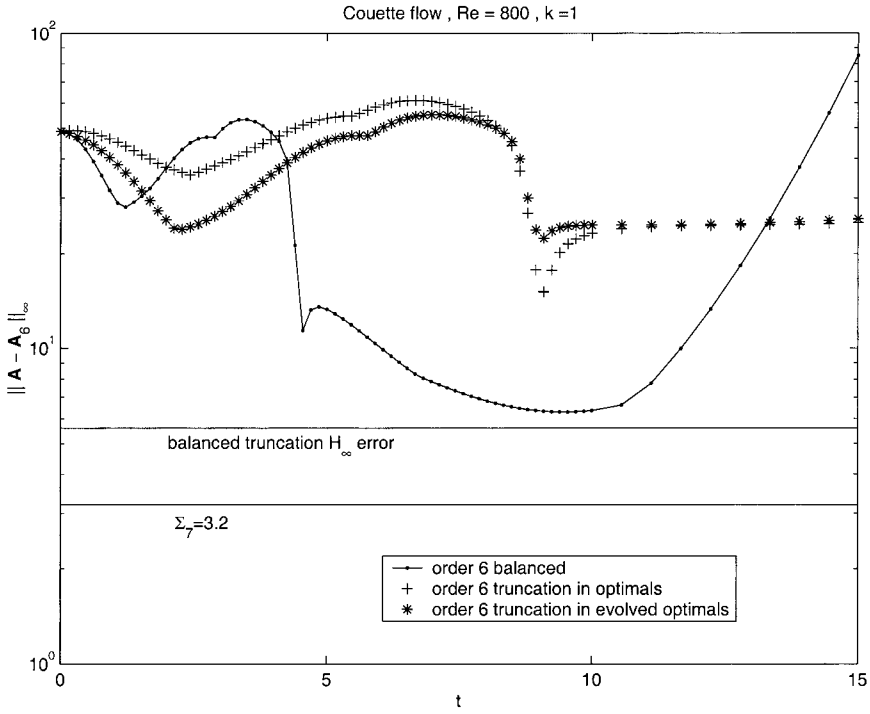

FIG. 11. The H_∞ error of an order 6 reduced system based on truncations to the top optimals and/or evolved optimals of the propagator for time t as a function of the optimizing time t. Shown is the error when the truncation has been done after balancing the optimals and the evolved optimals (line–dot); the error of a truncation to the top 6 evolved optimals of the propagator (stars); the error of a truncation to the top 6 optimals of the propagator (crosses). Notice that the truncation based on balancing the optimals and evolved optimals is a good approximation to the dynamics when the optimizing time is selected to be near the global optimal time: a minimum error of 6.4 is found for optimizing time 10, which is close to the error obtained from a balanced truncation of the system (this error is 5.2, and is indicated by a horizontal line). The minimum possible error is equal to the first neglected Hankel singular value, which is $\Sigma_7 = 3.2$, (also indicated by a horizontal line).

thousand but for systems of larger dimension these matrix algorithms must be replaced by approximations to the singular value decomposition of the propagator (Lanczos 1950) and Riemann sum approximation of the $\mathbf{P}$ and $\mathbf{Q}$ matrices. Examples of systems truncated using these methods suggest that they can be successful even for systems of very high dimension such as a forecast model so long as the Hankel singular values fall sufficiently rapidly.

The study of optimal order reduction has led to a deeper understanding of the role of nonnormality in linear dynamical systems. Application of these results to the problem of obtaining an accurate linear stochastic model of storm track statistics by reducing a climate model to the minimum dimension responsible for maintaining eddy variance will be reported elsewhere. Also to be reported is a second area of application in which model order reduction is being used to facilitate integration of the error covariance matrix required for implementing Kalman filter optimal state estimation for forecast initialization.

Acknowledgments. The authors thank Dr. Bassam Bamieh for helpful discussions. This work was supported by NSF ATM-9623539 and by ONR N00014-99-1-0018.

APPENDIX A

Calculation of the Truncation Error

We obtain the H_∞ norm of the error of the approximation as in Glover (1984). The system state in transformed coordinates, x, is related to the system state in physical coordinates, ψ, through the transformation $x = \mathbf{T}\psi$. The dynamical operator $\tilde{\mathbf{A}}$ in transformed coordinates is related to the operator in the original coordinates, $\mathbf{A}$ by (17). The k-dimensional reduced-order dynamical operator is $\tilde{\mathbf{A}}_k$, which is equal to the first k rows and columns of $\tilde{\mathbf{A}}$. In these coordinates the k-dimensional vector x_k satisfies the transformed reduced-order equation

$$\frac{dx_k}{dt} = \tilde{\mathbf{A}}_k x_k + \mathbf{T}_{kn} u(t), \tag{A.1}$$

where $u(t)$ is an n-dimensional forcing in physical coordinates and $\mathbf{T}_{kn}$ is the first k rows of the transformation matrix $\mathbf{T}$. The full system is

$$\frac{dx}{dt} = \tilde{\mathbf{A}}x + \mathbf{T}u(t). \tag{A.2}$$

Consider now the compound system

$$\frac{dy}{dt} = \mathbf{E}y + \mathbf{F}u(t), \tag{A.3}$$

where $y(t)$ is the compound state $[x_k, x]^{\mathrm{T}}$ (T denotes transposition), $\mathbf{F} = [\mathbf{T}_{kn}, \mathbf{T}]^{\mathrm{T}}$, and $\mathbf{E}$ is the compound operator that evolves concurrently the approximate and full system

$$\mathbf{E} = \begin{pmatrix} \tilde{\mathbf{A}}_k & 0 \\ 0 & \tilde{\mathbf{A}} \end{pmatrix}. \tag{A.4}$$

The frequency response of the compound system is $\hat{y}(\omega) = \mathbf{R}_E(\omega)\mathbf{F}\hat{u}(\omega)$ where $\mathbf{R}_E(\omega)$ is the resolvent: $\mathbf{R}_E(\omega) = (i\omega\mathbf{I} - \mathbf{E})^{-1}$.

The Fourier amplitude of the error in physical coordinates is

$$\hat{e}(\omega) = \mathbf{T}_{nk}^{-1}\hat{x}_k(\omega) - \mathbf{T}^{-1}\hat{x}(\omega), \tag{A.5}$$

where $\mathbf{T}_{nk}^{-1}$ is the matrix consisting of the first k columns of $\mathbf{T}^{-1}$. The error can be written in terms of the resolvent of the compound system as

$$\begin{aligned} \hat{e}(\omega) &= [\mathbf{T}_{nk}^{-1}, -\mathbf{T}^{-1}]\hat{y}(\omega) \\ &= [\mathbf{T}_{nk}^{-1}, -\mathbf{T}^{-1}]\mathbf{R}_E(\omega)[\mathbf{T}_{kn}, \mathbf{T}]^{\mathrm{T}}\hat{u}(\omega), \end{aligned} \tag{A.6}$$

The H_∞ norm of the truncation error, denoted $\|\mathbf{A} - \mathbf{A}_k\|_\infty$, is the maximum over all real frequencies of the response of the compound system (A.6):

$$\|\mathbf{A} - \mathbf{A}_k\|_\infty \equiv \max_{\omega \in \Re} \|[\mathbf{T}_{nk}^{-1}, -\mathbf{T}^{-1}]\mathbf{R}_E(\omega)[\mathbf{T}_{kn}, \mathbf{T}]^{\mathrm{T}}\|_2, \tag{A.7}$$

where the subscript, 2, denotes the L_2 norm of the matrix, which is equal to its largest singular value. The error is the largest difference in the response of the full and approximate system that can be obtained by a sinusoidal forcing.

APPENDIX B

The Singular Values of Operators M and N

Operators $\mathbf{M}$ and $\mathbf{N}$ are defined in (40) and (41), respectively. To determine their singular values we need the adjoint operators $\mathbf{M}^\dagger$ and $\mathbf{N}^\dagger$. The adjoint of $\mathbf{M}$ is obtained by noting that

$$\begin{aligned} (\psi(t), \mathbf{M}\psi_0) &= \int_0^\infty \psi(t)^\dagger e^{\mathbf{A}t}\psi_0 \, dt = \left[\int_0^\infty e^{\mathbf{A}t}\psi(t) \, dt\right]^\dagger \psi_0 \\ &= [\mathbf{M}^\dagger\psi(t), \psi_0]. \end{aligned} \tag{B.1}$$

Hence

$$\mathbf{M}^\dagger\psi(t) = \int_0^\infty e^{\mathbf{A}^\dagger t}\psi(t) \, dt. \tag{B.2}$$

Therefore for any ψ_0 the product

$$\mathbf{M}^\dagger\mathbf{M}\psi_0 = \left(\int_0^\infty e^{\mathbf{A}^\dagger t}e^{\mathbf{A}t} \, dt\right)\psi_0 = \mathbf{Q}\psi_0, \tag{B.3}$$

and $\mathbf{Q}$ can be decomposed as

$$\mathbf{Q} = \mathbf{M}^\dagger\mathbf{M}, \tag{B.4}$$

and the singular values of $\mathbf{M}$ are revealed to be the square roots of the eigenvalues (also singular values) of $\mathbf{Q}$.

The adjoint operator of $\mathbf{N}$ is found by noting that

$$\begin{aligned} (\psi_0, \mathbf{N}u) &= \psi_0 \int_{-\infty}^0 e^{-\mathbf{A}t}u(t) \, dt \\ &= \int_{-\infty}^0 (e^{-\mathbf{A}^\dagger t}\psi_0)^\dagger u(t) \, dt = (\mathbf{N}^\dagger\psi_0, u), \end{aligned} \tag{B.5}$$

and consequently

$$\mathbf{N}^\dagger\psi_0 = e^{-\mathbf{A}^\dagger t}\psi_0, \tag{B.6}$$

for $t < 0$. We then have that the product

$$\begin{aligned} \mathbf{N}\mathbf{N}^\dagger\psi_0 &= \left(\int_{-\infty}^0 e^{-\mathbf{A}t}e^{-\mathbf{A}^\dagger t} \, dt\right)\psi_0 \\ &= \left(\int_0^\infty e^{\mathbf{A}t}e^{\mathbf{A}^\dagger t} \, dt\right)\psi_0 = \mathbf{P}\psi_0, \end{aligned} \tag{B.7}$$

and the matrix $\mathbf{P}$ can be decomposed as

$$\mathbf{P} = \mathbf{N}\mathbf{N}^\dagger. \tag{B.8}$$

As a result the singular values of $\mathbf{N}$ are the square roots of the eigenvalues (also singular values) of $\mathbf{P}$.

REFERENCES

Adamjan, V. M., D. Z. Arov, and M. G. Krein, 1971: Analytic properties of Schmidt pairs for a Hankel operator and the generalized Schur–Takagi problem (in Russian). *Math. USSR Sb.,* **15,** 31–73. (Translation 1978: *Amer. Math. Soc. Transl.*, **111,** 133–175).

Drazin, P. G., and W. H. Reid, 1981: *Hydrodynamic Stability*. Cambridge University Press, 527 pp.

Eckart, C., and G. Young, 1936: The approximation of one matrix by another of lower rank. *Phychometrika,* **1,** 211–218.

Farrell, B. F., and P. J. Ioannou, 1993a: Stochastic forcing of the linearized Navier–Stokes equations. *Phys. Fluids A,* **5,** 2600–2609.

——, and ——, 1993b: Stochastic dynamics of baroclinic waves. *J. Atmos. Sci.,* **50,** 4044–4057.

——, and ——, 1994: Variance maintained by stochastic forcing of non-normal dynamical systems associated with linearly stable shear flows. *Phys. Rev. Lett.,* **72,** 1188–1191.

——, and ——, 1996: Generalized stability. Part I: Autonomous operators. *J. Atmos. Sci.,* **53,** 2025–2041.

——, and ——, 2001: State estimation using a reduced-order Kalman filter. *J. Atmos. Sci.,* in press.

Glover, K., 1984: An optimal Hankel-norm approximation of linear

multivariable systems and their L^∞-error bounds. *Int. J. Control,* **39,** 1115–1193.

Hasselmann, K., 1988: PIPs and POPs: The reduction of complex dynamical systems using principal interaction and oscillation patterns. *J. Geophys. Res.,* **93** (D9), 11 015–11 021.

Holmes, P. J., J. L. Lumley, G. Berkooz, J. C. Mattingly, and R. W. Wittenberg, 1997: Low-dimensional models of coherent structures in turbulence. *Phys. Rep.,* **287,** 337–384.

Horn, R. A., and C. R. Johnson, 1991: *Topics in Matrix Analysis.* Cambridge University Press, 607 pp.

Ince, E. L., 1956: *Ordinary Differential Equations.* Dover, 558 pp.

Ioannou, P. J., 1995: Nonnormality increases variance. *J. Atmos. Sci.,* **52,** 1155–1158.

Lanczos, C., 1950: An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Res. Nat. Bur. Stand.,* **45,** 255–282.

Luenberger, D. G., 1969: *Optimization by Vector Space Methods.* Wiley, 352 pp.

Mirsky, L., 1960: Symmetric gage functions and unitarily invariant norms. *Quart. J. Math.,* **11,** 50–59.

Moore, B. C., 1981: Principal component analysis in linear systems: Controllability, observability, and model reduction. *IEEE Trans. Autom. Control,* **AC–26,** 17–31.

Schmidt, E., 1907: Zur theorie der linearen und nichtlinearen integralgleichungen. I Teil. Entwicklung willkürlichen funktionen nach system vorgeschriebener. *Math. Ann.,* **63,** 433–476.

Stewart, G. W., and J.-G. Sun, 1990: *Matrix Perturbation Theory.* Academic Press, 365 pp.

Zhou, K., and J. C. Doyle, 1998: *Essentials of Robust Control.* Prentice Hall, 411 pp.