

FOURIER SERIES FOR ACCURATE, STABLE, REDUCED-ORDER MODELS IN LARGE-SCALE LINEAR APPLICATIONS*

K. WILLCOX[†] AND A. MEGRETSKI[‡]

Abstract. A new method, Fourier model reduction, for obtaining stable, accurate, low-order models of very large linear systems is presented. The technique draws on traditional control and dynamical system concepts and utilizes them in a way which is computationally very efficient. Discrete-time Fourier coefficients of the large system are calculated and used to construct a reduced-order model that preserves stability properties of the original system. Many coefficients can be calculated, which results in a very accurate representation of the system dynamics, but only a single factorization of the large system is required. The resulting system can be further reduced using explicit formulae for balanced truncation. The method is applied to two computational fluid dynamic systems, which model unsteady motion of a two-dimensional subsonic airfoil and unsteady flow in a supersonic diffuser. In both cases, the new method is found to work extremely well. Results are compared to models developed using the proper orthogonal decomposition and Arnoldi method. In comparison with these widely used techniques, the new method is computationally more efficient, preserves the stability of the original system, uses both input and output information, and, for smooth transfer functions, is valid over a wide range of frequencies.

Key words. model reduction, computational fluid dynamics, Fourier series

AMS subject classifications. 37N10, 42A16, 76M25

DOI. 10.1137/S1064827502418768

1. Introduction. Model reduction is a powerful tool which has been applied throughout many different disciplines, including controls, fluid dynamics, structural dynamics, and circuit design. In many situations, high-order, complicated numerical models accurately represent the problem at hand but are unsuitable for the desired application, for instance, for coupling between disciplines, as in aeroelastic analyses, or for control design. The goal of model reduction is to develop a model with a low number of states, which captures the system dynamics accurately over a range of frequencies and forcing inputs. For very large systems, the cost of reduction may also be an issue.

Many effective reduction techniques have been developed in a controls context. An optimal reduced model is one that minimizes the H-infinity norm of the difference between the reduced and original system transfer functions; however, no polynomial-time algorithm is known to achieve this goal. Algorithms such as optimal Hankel model reduction [1, 3, 12] and balanced truncation [16] have been widely used throughout the controls community to generate suboptimal reduced models with strong guarantees of quality. These algorithms can be performed in polynomial time; however, the computational requirements make them impractical for application to large systems such as those encountered in computational fluid dynamic (CFD) applications. In this case, system orders exceed 10^4 and the computation of grammians is impractical. For this reason, many of the control-based reduction concepts have not been transferred to other disciplines. Several methods have been developed for computing

*Received by the editors December 2, 2002; accepted for publication (in revised form) April 1, 2004; published electronically January 20, 2005.

<http://www.siam.org/journals/sisc/26-3/41876.html>

[†]Department of Aeronautics and Astronautics, Room 37-447, MIT, Cambridge, MA 02139 (kwillcox@mit.edu).

[‡]Department of Electrical Engineering and Computer Science, Room 32-D730, MIT, Cambridge, MA 02139 (ameg@mit.edu).

approximations to the grammians for large systems, including the approximate subspace iteration [2], least squares approximation [19], and Krylov subspace methods [9, 7]; however, these algorithms are complicated and computationally intensive.

Other techniques have been developed and widely used throughout the fluids community. The typical approach is to first determine an orthonormal basis, and to then project the large CFD system onto this reduced space to derive the reduced-order model. In particular, the proper orthogonal decomposition (POD), also known as Karhunen–Loève expansions [15], has been developed as a popular method of deriving basis vectors for high-order systems using the method of snapshots [21, 8]. This technique has been considered for a broad range of fluid dynamic applications, which are reviewed in Dowell and Hall [5]. Another class of reduction techniques that are based on matching moments of the system transfer function has been developed for analysis of large linear circuits. For instance, the Arnoldi algorithm can be used to generate vectors which form an orthonormal basis for the Krylov subspace, and has been applied to model reduction of RLC circuits [20] and turbomachinery aeroelastic applications [24].

While the POD and Arnoldi techniques have been shown to be effective for model reduction of large systems, they lack many of the desirable properties possessed by methods such as optimal Hankel model reduction. In particular, neither the POD nor the Arnoldi method is guaranteed to result in a stable reduced-order model derived from a stable CFD system. In practice, for CFD applications, these techniques sometimes yield unstable models despite the original system being stable. Moreover, the POD does not take account of system outputs when performing the reduction, and hence the reduced-order models produced may be inefficient. The POD has been suggested as a means to obtain an approximate balanced truncation for large systems using both inputs and outputs [13, 23]; however, the reduction approach is computationally expensive and offers no stability guarantees.

In this paper, a new technique for performing model reduction of very large systems is presented. This method draws on classical dynamical system and control theory concepts and applies them using an iterative procedure that is very efficient for large systems. The resulting reduced-order models are guaranteed to preserve the stability properties of the original system and have an associated error bound that depends on the smoothness of the original transfer function. The concept is similar to moment-matching methods, such as Arnoldi, except that the transfer function expansion is performed in the discrete frequency domain and a projection framework is not utilized. This idea was also explored in [4], where a Padé approximation is applied to a bilinear transformation of the transfer function. In [4], the Arnoldi or Lanczos algorithms are applied in the transformed domain; here, we take a different approach by computing the Fourier expansion of the discrete-frequency transfer function. Model reduction based on such a Fourier expansion was suggested in [6] and [25]. In [6] and [25], the Fourier coefficients of the expansion were calculated approximately using a fast Fourier transform (FFT) algorithm. The FFT approach requires frequency response data to be provided, making the algorithm prohibitively expensive for many large-scale applications.

In this paper, an algorithm is proposed that allows the Fourier coefficients of a discrete-time transfer function to be computed exactly. The algorithm uses an efficient iterative procedure, thus allowing its application to large-scale problems. A subset of the Fourier coefficients is used to construct a discrete-time (DT) reduced model. This model can then be converted to a continuous-time state-space model, or the method can be augmented with a further reduction step via balanced truncation. Due to the

particular structure of the DT model, the formulae for balanced truncation can be derived explicitly and applied efficiently.

The organization of the paper is as follows. In the following section, the dynamical system arising from implementation of CFD is briefly described. The Fourier model reduction (FMR) technique is then presented and compared with the Arnoldi and POD approaches. Examples are considered in the context of two different CFD applications, both of which require fluid models of low order. The first example investigates the unsteady motion of a two-dimensional airfoil, which forms the fluid component of an aeroelastic analysis. The second example derives a model of the flow dynamics of a supersonic inlet that will be used to derive active control strategies. Finally, conclusions are presented.

2. CFD model. Consider a general linearized CFD model, which can be written as

$$(2.1) \quad G: \quad E \frac{d}{dt} x = Ax + Bu, \quad y = Cx + Du,$$

where $x(t) \in \mathbf{R}^n$ contains the n unknown perturbation flow quantities at each point in the computational grid. For example, for two-dimensional, compressible, inviscid flow, which is governed by the Euler equations, the unknowns at each grid point are the perturbations in flow density, Cartesian momentum components, and flow energy. The vectors $u(t)$ and $y(t)$ in (2.1) contain the system inputs and outputs, respectively. The definition of inputs and outputs will depend upon the problem at hand. In aeroelastic analysis of a wing, inputs consist of wing motion, while outputs of interest are the forces and moments generated. For control purposes, the output might monitor a flow condition at a particular location which varies in response to a disturbance in the incoming flow.


The linearization matrices A , B , C , D , and E in (2.1) are evaluated at steady-state conditions. The descriptor matrix E is included for generality and may contain some zero rows, which arise from implementation of flow boundary conditions. On solid walls, a condition is imposed on the flow velocity, while at farfield boundaries certain flow parameters are specified, depending on the nature of the boundary (inflow/outflow) and the local flow conditions (subsonic/supersonic). Although these prescribed quantities could be condensed out of (2.1) to obtain a smaller state-space system, such a manipulation is often complicated and can destroy the sparsity of the system. The more general form of the system is therefore considered.

The system (2.1) is efficient for time computations since a time discretization, such as backward Euler, can be applied and the resulting large $n \times n$ system matrix factored just once for a time-dependent calculation. However, the order of the system is still prohibitively high for many applications, such as aeroelasticity and active flow control. In the next section, we present an efficient method with quality guarantees to reduce the size of the system while retaining an accurate representation of important flow dynamics.

3. Fourier series model reduction. We consider the task of finding a low-order, stable, continuous-time, linear-time invariant (LTI), state-space model

$$(3.1) \quad \hat{G}: \quad \frac{d}{dt} \hat{x}(t) = \hat{A}\hat{x}(t) + \hat{B}u(t), \quad \hat{y}(t) = \hat{C}\hat{x}(t) + \hat{D}u(t)$$

which approximates well the given stable model (2.1). We consider first the case of a single-input, single-output (SISO) system, and then the extension to a multiple-input, multiple-output (MIMO) problem. Typically, A and E in (2.1) are sparse,

FIG. 3.1. Comparing G and \hat{G} .

square matrices of very large dimension $n > 10^4$, and the desired order k of \hat{G} is less than 50.

The quality of \hat{G} as an approximation of G is defined as the H-infinity norm of the difference between their transfer functions:

$$(3.2) \quad \|\hat{G} - G\|_\infty = \sup_{\omega \in \mathbf{R}} |\hat{G}(j\omega) - G(j\omega)|,$$

which in turn equals the square root of the maximal energy

$$(3.3) \quad \|\hat{y} - y\|_2^2 = \int_t |\hat{y}(t) - y(t)|^2 dt$$

of the difference $e = \hat{y} - y$, which can be generated when testing both G and \hat{G} with an arbitrary unit energy input u as shown in Figure 3.1. With this measure of model reduction error, if a good reduced model \hat{G} is found, G can be represented, for design or analysis purposes, as a series connection (i.e., a sum) of \hat{G} and a small “uncertain” error system $\Delta = G - \hat{G}$, and the standard results from robustness analysis can be applied to predict the effect of replacing G with \hat{G} in even larger scale systems.

No efficient (polynomial-time) solution is known for the problem of minimizing $\|\hat{G} - G\|_\infty$ subject to the order and stability constraints imposed on \hat{G} . There are polynomial-time algorithms (optimal Hankel model reduction and balanced truncation) which produce suboptimal reduced models with strong guarantees of quality. However, the computational burden of these methods is still quite heavy, which makes their direct application impractical for $n > 10^4$.

In this paper, a very low complexity algorithm is introduced, which allows one to find an “intermediate” approximation of G by a model \bar{G} of order in the range of hundreds. While \bar{G} is not an optimal reduced model of G , it satisfies an attractive guaranteed H-infinity quality bound. After \bar{G} is found, the second round of more demanding model reduction, such as balanced truncation, is applied to \bar{G} to produce a high-quality, low-order, reduced model \hat{G} .

3.1. Fourier series of DT systems. Consider the full DT LTI system model g defined by the difference equations

$$(3.4) \quad g: \quad x(t+1) = ax(t) + bu(t), \quad y(t) = cx(t) + du(t),$$

where a, b, c, d are given matrices of coefficients, $x(t) \in \mathbf{R}^n$ is the system state, and $u(t), y(t)$ are scalar input and output. It will be assumed that g is stable, i.e., $\rho(a) < 1$, where $\rho(M)$ denotes the *spectral radius* of M , defined as the maximal absolute value of its eigenvalues.

The transfer function

$$(3.5) \quad g(z) = d + c(zI - a)^{-1}b$$

has the Fourier decomposition

$$(3.6) \quad g(z) = \sum_{k=0}^{\infty} g_k z^{-k},$$

where

$$(3.7) \quad g_0 = d, \quad g_k = ca^{k-1}b \quad (k = 1, 2, \dots).$$

The Fourier expansion converges exponentially for $|z| > \rho(a)$. Note that the first m Fourier coefficients g_k are easy to calculate using the “cheap” iterative process

$$(3.8) \quad g_k = ch_{k-1}, \quad h_k = ah_{k-1} \quad (k = 1, \dots, m), \quad \text{where } h_0 = b,$$

which is, in theory, “stable” since $\rho(a) < 1$. In practice, if a has eigenvalues that lie very close to the unit circle, numerical resolution may be lost as the iterations proceed.

Let \hat{g}_m denote the m th-order approximation of g based on the Fourier series expansion

$$(3.9) \quad \hat{g}_m(z) = \sum_{k=0}^m g_k z^{-k}.$$

The following simple result provides an estimate of the approximation error:

$$(3.10) \quad \|g - \hat{g}_m\|_{\infty} = \max_{|z|=1} |g(z) - \hat{g}_m(z)|;$$

this result ties it to the smoothness of G as follows.

THEOREM 3.1. *For $q = 1, 2, \dots$*

$$(3.11) \quad \|g - \hat{g}_m\|_{\infty}^2 \leq \frac{m^{1-2q}}{2\pi(2q-1)} \int_{-\pi}^{\pi} |g^{(q)}(e^{j\tau})|^2 d\tau,$$

where $g^{(q)}$ is the q th derivative of g with respect to τ .

Proof.

$$\begin{aligned} |g(z) - \hat{g}_m(z)|^2 &= \left| \sum_{k=m+1}^{\infty} g_k z^{-k} \right|^2 \\ &= \left| \sum_{k=m+1}^{\infty} g_k k^q z^{-k} k^{-q} \right|^2. \end{aligned}$$

Using the Cauchy inequality and considering $z = e^{j\tau}$,

$$\begin{aligned} |g(e^{j\tau}) - \hat{g}_m(e^{j\tau})|^2 &\leq \left(\sum_{k=m+1}^{\infty} |g_k|^2 k^{2q} \right) \left(\sum_{k=m+1}^{\infty} k^{-2q} \right) \\ &\leq \left(\sum_{k=0}^{\infty} |g_k|^2 k^{2q} \right) \left(\int_m^{\infty} \frac{dx}{x^{2q}} \right). \end{aligned}$$

Using the fact that

$$\left(\frac{\partial}{\partial \tau}\right)^q g(e^{j\tau}) = \sum_{k=0}^{\infty} (-jk)^q g_k e^{-jk\tau}$$

and applying Parseval's theorem, we obtain

$$|g(e^{j\tau}) - \hat{g}_m(e^{j\tau})|^2 \leq \frac{m^{1-2q}}{2\pi(2q-1)} \int_{-\pi}^{\pi} |g^{(q)}(e^{j\tau})|^2 d\tau.$$

3.2. Fourier series of continuous-time systems. Consider the full continuous-time LTI system model G defined by the system (2.1), where $u(t), y(t)$ are scalar input and output. If E is invertible, then $x = x(t)$ is the system state. It will be assumed that G is stable, i.e., that all roots of the characteristic equation $\det(sE - A) = 0$ have negative real part, and that $C(sE - A)^{-1}B$ remains bounded as $s \rightarrow \infty$.

Let $\omega_0 > 0$ be a fixed positive real number. The transfer function

$$(3.12) \quad G(s) = D + C(sE - A)^{-1}B$$

has the Fourier decomposition

$$(3.13) \quad G(s) = \sum_{k=0}^{\infty} G_k \left(\frac{s - \omega_0}{s + \omega_0}\right)^k.$$

This decomposition was suggested in [6]; however, no means was provided to calculate the Fourier coefficients G_k . Instead, an approximate FFT algorithm was used. Here, an efficient iterative procedure is proposed to directly calculate the Fourier coefficients as follows.

Consider the identity

$$G(s) = g(z) = d + c(zI - a)^{-1}b \quad \text{for } z = \frac{s + \omega_0}{s - \omega_0},$$

which allows one to apply the observations and theorem from the previous subsection to this case. Note that by comparing (3.6) and (3.13), it can be seen that $G_k = g_k$. The Fourier coefficients are therefore given by the formulae

$$(3.14) \quad G_0 = d, \quad G_k = ca^{k-1}b \quad (k = 1, 2, \dots),$$

$$(3.15) \quad d = D + C(\omega_0 E - A)^{-1}B,$$

$$(3.16) \quad a = (\omega_0 E + A)(\omega_0 E - A)^{-1},$$

$$(3.17) \quad c = 2\omega_0 C(\omega_0 E - A)^{-1},$$

$$(3.18) \quad b = -E(\omega_0 E - A)^{-1}B,$$

which are relatively straightforward to obtain using algebraic manipulations. An outline of the derivation is as follows.

Consider manipulation of the term $(sE - A)^{-1}$ in (3.12). Using the relationship between continuous and discrete frequency, we obtain

$$(3.19) \quad \begin{aligned} (sE - A)^{-1} &= \left(\omega_0 \frac{(z+1)}{(z-1)} E - A\right)^{-1} \\ &= \left(1 - \frac{1}{z}\right) (\omega_0 E - A)^{-1} \left[I - \frac{1}{z} (\omega_0 E + A) (A - \omega_0 E)^{-1}\right]^{-1}. \end{aligned}$$

The last term in the above equation can be expanded using the well-known Maclaurin series

$$(3.20) \quad \left[I - \frac{1}{z} \bar{a} \right]^{-1} = I + \frac{\bar{a}}{z} + \frac{\bar{a}^2}{z^2} + \frac{\bar{a}^3}{z^3} + \dots,$$

where $\bar{a} = (\omega_0 E + A)(A - \omega_0 E)^{-1}$. The next step is to substitute the expansion (3.20) into (3.19), multiply out the terms, and use the fact that

$$(3.21) \quad (\omega_0 E - A)^{-1} (\bar{a} - I) = -2\omega_0 (\omega_0 E - A)^{-1} E (\omega_0 E - A)^{-1}.$$

Finally, incorporating B , C , and D from (3.12) yields the final result given in (3.13) through (3.18).

3.3. Reduced model construction. To construct an m th-order reduced model, one first calculates the Fourier coefficients, g_0, g_1, \dots, g_m . The DT reduced model is then given by

$$(3.22) \quad \begin{aligned} \hat{g} : \quad \hat{x}[t + 1] &= \hat{a}\hat{x}[t] + \hat{b}u[t], \\ \hat{y}[t] &= \hat{c}\hat{x}[t] + \hat{d}u[t], \end{aligned}$$

where

$$(3.23) \quad \hat{a} = \begin{bmatrix} 0 & 0 & 0 & 0 & \dots \\ 1 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ 0 & 0 & 1 & 0 & \dots \\ 0 & 0 & 0 & \ddots & \dots \end{bmatrix}, \quad \hat{b} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \end{bmatrix},$$

$$\hat{c} = [g_1 \quad g_2 \quad \dots \quad g_m], \quad \hat{d} = g_0.$$

It should be noted that if the original system g is stable, then the reduced system \hat{g} is also guaranteed to be stable. This is due to the orthogonality properties of the Fourier expansion; taking a truncated number of Fourier terms, as in (3.9), automatically results in a stable approximation of a stable system.

Different alternatives could be chosen for the DT system representation. The controller canonical form above was selected, as in [6], for the purpose of an efficient second step of reduction. An effective approach is to use the efficient iterative procedure to calculate several hundred coefficients, resulting in an intermediate reduced model of the form (3.22). A second reduction step using balanced truncation can now be performed easily, since the expressions for the grammians are known explicitly. For the DT reduced model (3.22), the controllability matrix is the identity matrix and the observability matrix is the Hankel matrix that has \hat{c} as its first row. The balancing vectors can therefore be obtained by computing the singular vectors of the m th-order Hankel matrix

$$(3.24) \quad \Gamma = \begin{bmatrix} g_1 & g_2 & g_3 & \dots & g_{m-1} & g_m \\ g_2 & g_3 & g_4 & \dots & g_m & 0 \\ g_3 & g_4 & g_5 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ g_{m-1} & g_m & 0 & \dots & 0 & 0 \\ g_m & 0 & 0 & \dots & 0 & 0 \end{bmatrix}.$$

The Hankel singular values, σ_i , $i = 1, 2, \dots, m$, of the intermediate reduced system are given by the singular values of Γ .

The final reduced system, \tilde{g} , is computed as

$$(3.25) \quad \tilde{a} = V_k' \hat{a} V_k, \quad \tilde{b} = V_k' \hat{b}, \quad \tilde{c} = \hat{c} V_k, \quad \tilde{d} = \hat{d},$$

where V_k is a matrix whose columns are the first k singular vectors of Γ , and k is chosen according to the distribution of Hankel singular values. Specifically, it can be shown that the following error bound holds:

$$(3.26) \quad \|\hat{G} - \tilde{G}\|_\infty \leq 2 \sum_{i=k+1}^m \sigma_i,$$

where \hat{G} is the transfer function of the intermediate system obtained using m Fourier coefficients, and \tilde{G} is the final reduced model derived by retaining k states in the balanced truncation step. It should be noted that the application of optimal Hankel model reduction to the intermediate reduced model would yield a bound on the error lower than that in (3.26); however, balanced truncation can be applied very efficiently since the grammians are known explicitly. This also avoids the explicit construction of the intermediate system, as can be seen in the following algorithm.

3.4. Model reduction algorithm. The FMR algorithm is summarized in the following steps:

1. Choose a value of ω_0 . The value of ω_0 should reflect the frequency range of interest. The nominal value is unity; however, if the response at high frequencies is of interest, a higher value of ω_0 should be chosen. One can visualize the transformation from continuous to discrete time as a mapping of the imaginary axis in the s -plane to the unit circle in the z -plane. The value of ω_0 then describes the compression of frequencies around the unit circle.
2. Calculate $m + 1$ Fourier coefficients using (3.14)–(3.18). Using the iterative procedure, any number of coefficients can be calculated with a single n th-order matrix factorization.
3. Using (3.24), calculate the m th-order Hankel matrix. Calculate its singular values and singular vectors.
4. Using (3.25), construct a k th-order DT system, \tilde{g} . The value of k is chosen according to the distribution of Hankel singular values of the intermediate system.
5. Convert the k th-order, DT reduced model to a continuous-time model using the relationships

$$(3.27) \quad \hat{A} = \omega_0 (\hat{a} - I)^{-1} (\hat{a} + I),$$

$$(3.28) \quad \hat{B} = 2\omega_0 (\hat{a} - I)^{-1} \hat{b},$$

$$(3.29) \quad \hat{C} = -\hat{c} (\hat{a} - I)^{-1},$$

$$(3.30) \quad \hat{D} = \hat{d} - \hat{c} (\hat{a} - I)^{-1} \hat{b}.$$

The error bounds corresponding to the above algorithm are given in (3.11) and (3.26) for the first and second stages of the reduction, respectively. The Hankel singular values of the intermediate system provide a straightforward, quantitative means to choose k , the size of the final reduced model; however, the error bound given

in (3.11) is not readily computable. One way to determine the appropriate number of Fourier coefficients, m , is to monitor the magnitudes of the coefficients, g_k , which tend to decrease quickly with increasing index k . A tolerance on the coefficient magnitude can thus be used to set a stopping criterion for the iterative procedure. Note that the exact distribution of coefficients will depend upon the transfer function under consideration. For near resonant systems, the magnitudes of the Fourier coefficients will not decrease quickly, and many coefficients will be needed to gain an accurate representation. For such systems, other algorithms, such as Arnoldi-based reduction, may yield better results.

3.5. MIMO systems. Although the above algorithm was derived for SISO systems, it can be easily extended to the MIMO case. In fact, the steps above are essentially unchanged. Consider the general case of p inputs and q outputs. For step 2, one can still use (3.14)–(3.18) to calculate the Fourier coefficients, although now each g_k is a $q \times p$ matrix. The iterative procedure is still very efficient, since consideration of multiple inputs and outputs requires only additional system solves.

In step 3, one again uses (3.24) to construct the Hankel matrix. The size of this matrix will now be $mq \times mp$. The balanced truncation in step 4 proceeds as before, noting only that each scalar entry of \hat{a} and \hat{b} in (3.23) is replaced with a block entry (zero matrix or identity matrix, as appropriate) of size $p \times p$. Finally, the transformations in step 5 are valid for the MIMO case.

In the results section, two examples will be presented that demonstrate the effectiveness of this approach. We first compare the advantages of FMR with other commonly used techniques.

3.6. Comparison with alternative reduction methods. The POD is the most widely used reduction method in the fluid dynamics community. The basic idea is to collect a set of snapshots, which are solutions of the linearized CFD model at selected time instants or frequencies. These snapshots are then combined to form an efficient basis, which is optimal in the sense that it minimizes the error between the exact and projected CFD data. While the method results in reduced-order models which accurately reproduce those dynamics included in the sampling process, the models are not valid outside the sampled range. If snapshots are obtained in the time domain, it can be difficult to choose an appropriate forcing function, and typically many snapshots are required to obtain accurate results. Using the frequency domain to compute snapshots is often more convenient; however, each frequency sampled requires an n th-order matrix factorization. Moreover, the reduced-order models obtained via the POD offer no guarantee of stability.

The Arnoldi method has also been used for model reduction of large CFD systems. The goal of moment-matching techniques, such as the Arnoldi method, is to determine a reduced-order model by matching moments (or Taylor series coefficients) of the high-order system transfer function. The Arnoldi method has an advantage over the POD in that a sequence of basis vectors can be generated in the frequency domain with just a single n th-order matrix factorization, and the approach of generating an intermediate model which can be subsequently further reduced via Hankel model reduction or balanced truncation also sometimes works effectively for this technique. A two-stage reduction process using Arnoldi followed by balanced truncation has been used by a number of authors for control and circuit applications [6, 11, 18]. The Arnoldi method can be applied using system inputs and/or outputs [17]; however, the resulting reduced-order models are not guaranteed to be stable. Arnoldi-based reduction does preserve the definiteness of the system, which for many applications is

TABLE 3.1
Comparison of reduction techniques for CFD systems.

| Method | Reduction cost | Frequency range | Stability preserved | Reduction considers |
|--------------------|---------------------------------|-------------------|---------------------|---------------------|
| Time POD | $O(1)$ | Selected | No | Inputs |
| Frequency POD | $O(\text{No. of snapshots})$ | Selected | No | Inputs |
| Arnoldi | $O(1)$ | Restricted | No | Inputs, Outputs |
| Multipoint Arnoldi | $O(\text{No. of freq. points})$ | Wide | No | Inputs, Outputs |
| FMR | $O(1)$ | Problem dependent | Yes | Inputs, Outputs |

sufficient. In particular, for circuit applications, reduced models obtained using the Arnoldi method preserve the passivity of the system. For CFD applications, passivity is not important and preservation of definiteness offers no guarantees for stability. In practice, unstable models are often generated if a large number of Arnoldi basis vectors is used for the first reduction step.

Since the Arnoldi basis vectors are derived from an expansion about zero frequency, for some applications the resulting reduced-order model can be large if frequencies far away from zero are of interest. The multiple frequency point Arnoldi method attempts to address this issue by considering transfer function expansions about multiple frequency points. In this way, accurate models can be derived over a specified frequency range. The reduction cost in this case is proportional to the number of frequency points considered. By generating multiple Arnoldi vectors at each frequency point, the samples can be placed further apart without loss in accuracy. The multipoint Arnoldi method therefore provides a way to trade between the low reduction cost of Arnoldi and the frequency span of POD.

Table 3.1 compares the attributes of each of these reduction techniques with the new FMR approach described in this paper. Reduction cost refers to the number of n th-order factorizations required. This is the dominating factor in reduction of large CFD systems, although one should also consider the number of system solves required. For time-domain POD, many system solves are required, covering the full snapshot simulation. For frequency-domain POD, one complex solve is required for each frequency chosen. Arnoldi and multiple-point Arnoldi require one solve per Arnoldi vector generated, while FMR requires one solve per Fourier coefficient. The exact computational cost depends on the problem at hand, but in general terms the cost of time-domain POD, Arnoldi, and FMR can be classed as “low,” that of multiple-point Arnoldi as “medium,” and that of frequency-domain POD as “high.”

In Table 3.1, frequency range refers to the range of validity of the resulting reduced-order models. This range is selected a priori for the POD and multipoint Arnoldi approaches. The validity of the Arnoldi-based reduced-order model is restricted to frequencies close to the expansion points; by choosing these expansion points appropriately, good results can be obtained, especially for near resonant systems (see, for example, [10]). The frequency range of validity for the FMR algorithm depends on the problem at hand. More specifically, as can be seen in the error bound in (3.11), the accuracy is dependent on the smoothness of the original transfer function. For relatively smooth transfer functions, such as those encountered in many CFD applications, a Fourier-based reduced model will often produce good approximation over a wide frequency range. For near resonant systems with a sharp peak in the transfer function, many Fourier coefficients will be required to obtain an accurate model. If the response contains a number of resonance peaks, the number of Fourier

coefficients required is likely to be prohibitive, and FMR is not recommended as a reduction algorithm. Similarly, if the original transfer function contains eigenvalues that are very close to the imaginary axis, numerical resolution may be lost in the power iteration (3.8) and only a few Fourier coefficients may be calculated accurately. Arnoldi-based techniques can avoid this problem using orthogonalization to yield a sound numerical implementation.

Finally, we note that none of the alternative methods are guaranteed to produce a stable reduced-order model. For certain applications, Arnoldi-based reduction methods are guaranteed to preserve the passivity of the system. FMR offers no theoretical guarantees of passivity preservation. In practice, for many applications where passivity is not an issue, the POD and Arnoldi approaches can often generate unstable models even though the original system is stable. For these applications, FMR provides an alternative approach that may yield more accurate reduced models. Two test cases will now be presented that demonstrate the new methodology. Results will also be shown to compare FMR against Arnoldi and POD.

4. Results. Two CFD applications will be considered. Each has very different flow dynamics and uses a different CFD formulation; however, FMR will be shown to work very effectively for both.

4.1. Subsonic airfoil. The first example is a two-dimensional NACA 0012 airfoil operating in unsteady plunging motion with a steady-state Mach number of 0.755. The flow is assumed to be inviscid, so the governing equations are the linearized Euler equations, which have four unknowns per grid point. A finite volume CFD formulation is used [22] with a CFD mesh containing 3482 grid points, which corresponds to a total of $n = 13,928$ unknowns in the linear state-space system. The input to this system is a rigid plunging motion (vertical motion of the airfoil), while the output of interest is the lift force generated. This input and output are typical for an aeroelastic analysis, which typically would also include pitching (angular) motion as an input and airfoil pitching moment as an output.

Fourier coefficients were generated using (3.14)–(3.18) for several different values of ω_0 . Figure 4.1 shows the transfer function of the CFD system compared with the transfer functions of the resulting reduced-order models. The size of each of these models is $m = 200$, which represents two orders of magnitude reduction from the CFD. However, as can be seen in Figure 4.1, the transfer functions match very closely over a large frequency range. Since the original CFD model is stable, these reduced-order models are also all guaranteed to be stable.

A model of size $m = 200$ is still relatively large for aeroelastic applications; however, it can be used as an intermediate step for further reduction. Balanced truncation was applied to the 200th-order DT model for $\omega_0 = 1$, and the first thirty Hankel singular values of the system are shown in Figure 4.2. These data indicate that a further reduction to five states can be achieved with virtually no loss in accuracy. While applying optimal Hankel model reduction to the intermediate reduced model would yield a lower bound on the error, balanced truncation can be applied very efficiently to the DT system since the grammians are known explicitly. The transfer function of a five-state model is shown in Figure 4.3 and can be seen to match the CFD data extremely well.

In order to compare the performance of the new method, reduced-order models for this problem were computed using the POD and Arnoldi methods. POD snapshots were obtained by causing the airfoil to plunge in sinusoidal motion at selected frequencies. Frequencies were selected at 0.1 increments from $\omega = 0$ to $\omega = 2.0$,




FIG. 4.1. Transfer function from plunging motion to lift force for subsonic airfoil. Results from CFD model ($n = 13,928$) are compared to reduced-order models derived using 200 Fourier coefficients with $\omega_0 = 1, 5, 10$.




FIG. 4.2. Hankel singular values for subsonic airfoil case. Values are calculated using a 200th-order model derived using FMR with $\omega_0 = 1$.

requiring 21 complex factorizations and solves. From these snapshots a set of POD basis vectors was obtained. Arnoldi vectors were generated about $\omega = 0$. Using the same approach as for the new method, 200 Arnoldi vectors were generated (with the cost of a single n th-order factorization and 200 solves), and balanced truncation was applied to the resulting 200th-order system. Figure 4.4 shows the transfer functions of fifth-order reduced-order models constructed using POD, Arnoldi, and the new FMR approach. It can be seen that even though it is much more expensive to compute than the other methods, POD has the worst performance. In particular, the figure




FIG. 4.3. Transfer function from plunging motion to lift force for subsonic airfoil. Results from CFD model ($n = 13,928$) are compared to a fifth-order reduced-order model derived from the $m = 200$, $\omega_0 = 1$ FMR model via balanced truncation.




FIG. 4.4. Transfer function from plunging motion to lift force for subsonic airfoil. Results from CFD model ($n = 13,928$) are compared to fifth-order reduced-order models derived using FMR, Arnoldi and POD. Both the Arnoldi and FMR models have been created using balanced truncation from an $m = 200$ system.

shows that, as expected, the POD-based reduced-order model has a large error outside the frequency range included in the snapshots. For this example, the Arnoldi-based reduced-order model shows similar accuracy to the new approach.

4.2. Supersonic diffuser. For the second example, we consider unsteady flow through a supersonic diffuser as shown in Figure 4.5. The diffuser operates at a nominal Mach number of 2.2; however, it is subject to perturbations in the incoming




FIG. 4.5. Mach contours for steady flow through supersonic diffuser. Steady-state inflow Mach number is 2.2.

flow, which may be due (for example) to atmospheric variations. In nominal operation, there is a strong shock downstream of the diffuser throat, as can be seen from the Mach contours plotted in Figure 4.5. Incoming disturbances can cause the shock to move forward towards the throat. When the shock sits at the throat, the inlet is unstable, since any disturbance that moves the shock slightly upstream will cause it to move forward rapidly, leading to unstart of the inlet. This is extremely undesirable, since unstart results in a large loss of thrust. In order to prevent unstart from occurring, one option is to actively control the position of the shock. This control may be effected through flow bleeding upstream of the diffuser throat. In order to derive effective active control strategies, it is imperative to have low-order models which accurately capture the relevant dynamics.

There are several transfer functions of interest in this problem. The shock position will be controlled by monitoring the average Mach number at the diffuser throat. The reduced-order model must capture the dynamics of this output in response to two inputs: the incoming flow disturbance and the bleed actuation. In addition, total pressure measurements at the diffuser wall are used for sensing. The response of this output to the two inputs must also be captured.

The CFD formulation for this problem is described fully in Lassaux [14]. Again, we consider the linearized two-dimensional Euler equations. The CFD model has 3078 grid points and 11,730 unknowns.

Figure 4.6 shows the transfer function between bleed actuation and average throat Mach number. Plotted are the results from the CFD code and FMR with $m = 200$ and $\omega_0 = 1, 5, 10$. The frequency axis is nondimensionalized by $f_0 = a_0/h$, where a_0 is the freestream speed of sound and h is the height of the diffuser. As the figure shows, the reduced-order models capture the dynamics very accurately over a wide range of frequencies. Typical disturbances might be on the range $[0, 2f_0]$; however, this model also captures higher frequencies should they occur. At higher frequencies, a slight error can be observed for the $\omega_0 = 1$ model; however, $\omega_0 = 5, 10$ yield very accurate results. This discrepancy highlights the importance of selecting a frequency scaling factor appropriately. Note that a frequency $\omega_0 = 1$ corresponds to $f/f_0 = 1/2\pi$. The first 200 Fourier coefficients for the $\omega_0 = 5$ case are plotted in Figure 4.7 and can be seen to reduce in magnitude quite quickly.

While $m = 200$ represents a significant reduction in order from the CFD, a model of this size is too large for derivation of active control strategies. Once again, we apply balanced truncation to further reduce the size of the model. The first thirty Hankel singular values are plotted in Figure 4.8 for the $\omega_0 = 5$ case. Figure 4.8 indicates that at least ten states should be retained to achieve an accurate representation. The transfer functions of reduced-order models with five and ten states are compared against CFD in Figure 4.9. While the model with five states has some error, with just ten states the results are almost indistinguishable.




FIG. 4.6. Transfer function from bleed actuation to average throat Mach number for supersonic diffuser. Results from CFD model ($n = 11,730$) are compared to reduced-order models derived using 200 Fourier coefficients with $\omega_0 = 1, 5, 10$.




FIG. 4.7. First 200 Fourier coefficients for supersonic diffuser transfer function from bleed actuation to average throat Mach number. $\omega_0 = 5$.

The results obtained using the combined FMR/balanced truncation reduction approach are compared to Arnoldi models in Figure 4.10. For this case, the two-step reduction procedure could not be used with the Arnoldi method, since including more than fifty basis vectors in the first reduced model led to an unstable system. Instead, the reduced-order models were calculated directly by projection onto the first k Arnoldi basis vectors. It can be seen in Figure 4.10 that with ten states, FMR is indistinguishable from the CFD results, while the Arnoldi model has considerable error for $f/f_0 > 1.3$. If the number of states is increased to thirty, the fit at lower frequencies




FIG. 4.8. Hankel singular values for supersonic diffuser case. Values are calculated using a 200th-order model derived using FMR with $\omega_0 = 5$.




FIG. 4.9. Transfer function from bleed actuation to average throat Mach number for supersonic diffuser. Results from CFD model ($n = 11,730$) are compared to reduced-order models with 5 and 10 states derived from the $m = 200$ FMR model via balanced truncation.

improves; however, some large oscillations are observed for high frequencies. If the number of basis vectors is further increased, the behavior at high frequencies becomes increasingly poor, until eventually an unstable system is obtained. This example highlights one of the shortcomings of the Arnoldi model reduction approach.

FMR is also applied to the transfer function between an incoming density perturbation and the average Mach number at the diffuser throat. This transfer function represents the dynamics of the disturbance to be controlled and is shown in Figure 4.11. As the figure shows, the dynamics contain a delay and are thus more




FIG. 4.10. Transfer function from bleed actuation to average throat Mach number for supersonic diffuser. Results from CFD model ($n = 11,730$) are compared to an FMR/BT model with 10 states and Arnoldi models with 10 and 30 states.




FIG. 4.11. Transfer function from incoming density perturbation to average throat Mach number for supersonic diffuser. Results from CFD model ($n = 11,730$) are compared to 200th-order FMR models with $\omega_0 = 5, 10$. The $\omega_0 = 10$ model is further reduced to $k = 30$ via balanced truncation.

difficult for the reduced-order model to approximate. Results are shown for FMR with $m = 200$ and $\omega_0 = 5, 10$. With $\omega_0 = 5$, the model has significant error for frequencies above $f/f_0 = 2$. Choosing a higher value of ω_0 improves the fit, although some discrepancy remains. These higher frequencies are unlikely to occur in typical atmospheric disturbances; however, if they are thought to be important, the model could be further improved either by evaluating more Fourier coefficients or by choosing a higher value of ω_0 . The $\omega_0 = 10$ model is further reduced via balanced truncation to a system with thirty states without a noticeable loss in accuracy.

5. Conclusions. Fourier model reduction is a new method for model reduction of very large linear systems. The method yields accurate, guaranteed stable reduced models, which can be derived using an efficient iterative procedure. An effective use of the method is to derive an intermediate DT reduced system, with more states than desired but which captures the relevant dynamics very accurately, and then to apply a second model reduction technique. Balanced truncation can be applied very efficiently to further reduce the system, since the system grammians are known explicitly. Efficient calculation of the intermediate system is a critical step in enabling the application of control-based theory to very large systems.

Results have been presented for two large dynamical systems, arising from implementation of CFD methods. While the flow dynamics of the two systems are very different—one is a subsonic external flow, the other a supersonic internal flow—the new method is shown to work extremely effectively in both cases. The cost to evaluate the reduced-order models is much lower than for the commonly used proper orthogonal decomposition. Moreover, the new method yields reduced-order models that produce accurate results over a wide frequency range for smooth transfer functions, account for system inputs and outputs, and preserve the stability properties of the original CFD system.

Acknowledgment. The authors would like to thank Professor Jacob White for helpful discussions regarding the development of this research.

REFERENCES

- [1] V. ADAMJAN, D. AROV, AND M. KREIN, *Analytic properties of Schmidt pairs for a Hankel operator and the generalized Schur-Takagi problem*, Math. USSR-Sb., 15 (1971), pp. 31–73.
- [2] M. BAKER, D. MINGORI, AND P. GOGGINS, *Approximate subspace iteration for constructing internally balanced reduced order models of unsteady aerodynamic systems*, in Proceedings of the 37th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference and Exhibit, Salt Lake City, UT, 1996.
- [3] M. BETTAYEB, L. SILVERMAN, AND M. SAFONOV, *Optimal approximation of continuous-time systems*, in Proceedings of the 19th IEEE Conference on Decision and Control, Albuquerque, NM, 1980.
- [4] C.-P. CHEN AND D. WONG, *Error bounded Padé approximation via bilinear conformal transformation*, in Proceedings of the 36th IEEE/ACM Design Automation Conference, 1999, pp. 7–12.
- [5] E. DOWELL AND K. HALL, *Modeling of fluid-structure interaction*, Ann. Rev. Fluid Mech., 33 (2001), pp. 445–490.
- [6] G. GU, P. KHARGONEKAR, AND E. LEE, *Approximation of infinite-dimensional systems*, IEEE Trans. Automat. Control, 34 (1989), pp. 610–618.
- [7] T. GUDMUNDSSON AND A. LAUB, *Approximate solution of large sparse Lyapunov equations*, IEEE Trans. Automat. Control, 39 (1994), pp. 1110–1114.
- [8] P. HOLMES, J. LUMLEY, AND G. BERKOOZ, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, Cambridge, UK, 1996.
- [9] I. M. JAIMOUKHA AND E. M. KASENALLY, *Krylov subspace methods for solving large Lyapunov equations*, SIAM J. Numer. Anal., 31 (1994), pp. 227–251.
- [10] M. KAMON, *Fast Parasitic Extraction and Simulation of Three-dimensional Interconnect via Quasistatic Analysis*, Ph.D. thesis, Department of Electrical Engineering & Computer Science, MIT, Cambridge, MA, 1988.
- [11] M. KAMON, F. WANG, AND J. WHITE, *Recent improvements to fast inductance extraction and simulation*, in Proceedings of the 7th Topical Meeting on Electrical Performance of Electronic Packaging, West Point, NY, 1998, pp. 281–284.
- [12] S.-Y. KUNG AND D. LIN, *Optimal Hankel-norm model reductions: Multivariable systems*, IEEE Trans. Automat. Control, 26 (1981), pp. 832–852.

- [13] S. LALL, J. E. MARSDEN, AND S. GLAVASKI, *A subspace approach to balanced truncation for model reduction of nonlinear control systems*, Internat. J. Robust Nonlinear Control, 12 (2002), pp. 519–535.
- [14] G. LASSAUX, *High-Fidelity Reduced-Order Aerodynamic Models: Application to Active Control of Engine Inlets*, master's thesis, Department of Aeronautics and Astronautics, MIT, Cambridge, MA, 2002.
- [15] M. LOÉVE, *Probability Theory*, D. Van Nostrand, New York, 1955.
- [16] B. MOORE, *Principal component analysis in linear systems: Controllability, observability, and model reduction*, IEEE Trans. Automat. Control, 26 (1981), pp. 17–31.
- [17] J. PHILLIPS, *Projection-based approaches for model reduction of weakly nonlinear, time-varying systems*, IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems, 22 (2003), pp. 171–187.
- [18] J. R. PHILLIPS, *Automated extraction of nonlinear circuits macromodels*, in Proceedings of the IEEE 2000 Custom Integrated Circuits Conference, 2000, pp. 451–454.
- [19] A. SCOTTEDWARD HODEL, *Least squares approximate solution of the Lyapunov equation*, in Proceedings of the 30th IEEE Conference on Decision and Control, 1991.
- [20] L. SILVEIRA, M. KAMON, I. ELFADEL, AND J. WHITE, *A coordinate-transformed Arnoldi algorithm for generating guaranteed stable reduced-order models of RLC circuits*, Comput. Methods Appl. Mech. Engrg., 169 (1999), pp. 377–389.
- [21] L. SIROVICH, *Turbulence and the dynamics of coherent structures. Part 1: Coherent structures*, Quart. Appl. Math., 45 (1987), pp. 561–571.
- [22] K. WILLCOX, *Reduced-Order Aerodynamic Models for Aeroelastic Control of Turbomachines*, Ph.D. thesis, Department of Aeronautics and Astronautics, MIT, Cambridge, MA, 2000.
- [23] K. WILLCOX AND J. PERAIRE, *Balanced model reduction via the proper orthogonal decomposition*, AIAA J., 40 (2002), pp. 2323–2330.
- [24] K. WILLCOX, J. PERAIRE, AND J. WHITE, *An Arnoldi approach for generation of reduced-order models for turbomachinery*, Comput. & Fluids, 31 (2002), pp. 369–389.
- [25] N. WU AND G. GU, *Discrete Fourier transform and H_∞ approximation*, IEEE Trans. Automat. Control, 35 (1990), pp. 1044–1046.