ARTICLE IN PRESS

Journal of Computational Physics xxx (2011) xxx–xxx

Contents lists available at SciVerse ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit

Gary J. Chandler [a], Matthew P. Juniper [a,*], Joseph W. Nichols [b,c], Peter J. Schmid [b]

[a] Engineering Department, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK
[b] LadHyX, Ecole Polytechnique, 91128 Palaiseau, France
[c] Center for Turbulence Research, Stanford University, CA 94305, USA

ARTICLE INFO

Article history:
Received 17 June 2011
Received in revised form 8 November 2011
Accepted 9 November 2011
Available online xxxx

Keywords:
Hydrodynamic stability
Adjoint
Low Mach number
Non-normality
Global modes

ABSTRACT

This paper describes a derivation of the adjoint low Mach number equations and their implementation and validation within a global mode solver. The advantage of using the low Mach number equations and their adjoints is that they are appropriate for flows with variable density, such as flames, but do not require resolution of acoustic waves. Two versions of the adjoint are implemented and assessed: a discrete-adjoint and a continuous-adjoint. The most unstable global mode calculated with the discrete-adjoint has exactly the same eigenvalue as the corresponding direct global mode but contains numerical artifacts near the inlet. The most unstable global mode calculated with the continuous-adjoint has no numerical artifacts but a slightly different eigenvalue. The eigenvalues converge, however, as the time-step reduces. Apart from the numerical artifacts, the mode shapes are very similar, which supports the expectation that they are otherwise equivalent. The continuous-adjoint requires less resolution and usually converges more quickly than the discrete-adjoint but is more challenging to implement. Finally, the direct and adjoint global modes are combined in order to calculate the wavemaker region of a low density jet.

1. Introduction

This paper concerns the calculation of the hydrodynamic stability of variable density flows. It is the first of a series taken from Ref. [1] concerning the hydrodynamic stability of flames. The stability is calculated by linearizing the Navier–Stokes equations around a base flow and then obtaining an operator that describes the evolution of small perturbations to the flow. The eigenvalues of this operator give the frequency and growth rate of the perturbations. The corresponding eigenfunctions give their corresponding shapes [2]. Eigenmodes can be calculated in one dimension [3] (Section 25–31) in two dimensions [4] and in three dimensions [5]. In this paper, we will call these the *direct* eigenmodes.

Having calculated the direct eigenmodes of a flow, we wish to know their sensitivity to small perturbations in the direct linear operator. First, however, we need to define what we mean by sensitivity. This requires a global measure of the size of the small perturbations. This is typically taken to be the perturbation kinetic energy integrated over the whole domain, which is proportional to the inner product $\langle \mathbf{u}, \mathbf{u} \rangle \equiv \int_V \mathbf{u}^* \cdot \mathbf{u} \, dV$. One type of sensitivity is the change that a perturbation makes to this measure. This sensitivity is given by the eigenmodes of the corresponding adjoint operator, where the adjoint is defined with respect to the above inner product [6] (Section 3.3.2). (Other inner products and corresponding adjoint operators are discussed in depth by [7].) Another type of sensitivity is the change that a perturbation makes to an eigenvalue. This is a robust quantity, in

* Corresponding author.
E-mail address: mpj1001@cam.ac.uk (M.P. Juniper).

0021-9991/$ - see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2011.11.013

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

that it is not affected by the above choice of measure [8] (Section 2.2). The adjoint eigenfunctions are biorthogonal to the direct eigenfunctions [9,10]. In other words, for each direct eigenfunction, there is only one non-orthogonal adjoint eigenfunction. The sensitivities are calculated by various combinations of the direct eigenfunction with its corresponding adjoint eigenfunction [11].

If an operator, L, is normal, it commutes with its corresponding adjoint operator, L^+, (i.e. $L^+L = LL^+$) and, for each eigenmode, the adjoint eigenfunction is identical to the direct eigenfunction. This means that a given eigenmode is most sensitive to changes of the direct linear operator in the region in which its eigenfunction has highest amplitude. If an operator is non-normal, it does not commute with its corresponding adjoint operator and, for each eigenmode, the adjoint eigenfunction differs from the direct eigenfunction. This means that a given eigenmode is most sensitive to changes of the stability operator in a different region. For example, flows with advection are non-normal [8], and their direct eigenfunctions have higher amplitude downstream, whereas their adjoint eigenfunctions have higher amplitude upstream.

Adjoint eigenmodes can be calculated for systems with many degrees of freedom. This has led to their extensive use in receptivity analyses [10], structural sensitivity analyses [4,11], aerodynamic design optimization of steady flow [12–15], optimal flow control [16,17], as well as routines to calculate initial conditions that maximize the energy growth of linear systems [17–20] and nonlinear systems [21].

The hydrodynamic stability of many flows is strongly affected by the density profile [22,23]. The adjoint operator has been calculated for the uniform density linearized Navier–Stokes (LN–S) equations [4,11] but these are not appropriate for flows with variable composition and density, such as flames. The adjoint operator has also been calculated for the fully compressible LN–S equations [24,25]. The eigenmodes of these equations contain acoustic modes, however, which require high temporal resolution, are strongly affected by boundary conditions, and in many cases are not relevant to the system stability. For this reason, we will use the low Mach number (LMN) equations [26,27]. These equations permit density variations due to temperature or species composition but not due to compressibility. (For simplicity, only variations due to temperature are presented in this paper because variations due to composition are dealt with using the same technique.) The LMN equations have been used to simulate self-sustained oscillations in variable-density round jets [28,29], unsteady combustion in pre-mixed flames [30], turbulent reactive plumes [31], and self-sustained oscillations in variable-density lifted diffusion flames [28,32,33].

Adjoint techniques have been applied to steady low Mach number flow problems [34], where low Mach number preconditioning is used to accelerate convergence. This technique, however, is not suitable for time-dependent flows, such as the ones considered here. For time-dependent low Mach number flows with density variation it is advantageous to use the linearized LMN equations and an adjoint operator of these equations does not seem to have been calculated before. This process is technically more difficult than for the uniform density or fully compressible LN–S equations because of the way that the density is treated. In Section 2, a review is given of the derivation of the direct low Mach number governing equations. In Section 3, the adjoint equations are derived. In Section 4, the numerical procedures are described; one for the direct equations and two different techniques for the adjoint equations. In Section 6, the numerical procedures are validated against each other and against previous results where possible. In Section 7, the direct and adjoint global modes for a hot jet at $Re = 1000$ are presented and the two adjoint techniques are compared against each other. Finally, the sensitivity of the hot jet to spatially-localized feedback is calculated in the same fashion as described in [11]. The structural sensitivity described in [11] was for incompressible flow behind a cylinder – we extend this to incorporate variable density effects, which are important to the case of the hot jet. The main subject of this paper is construction and validation of the numerical procedure. Subsequent papers will present more of the physical results from [1].

2. The direct governing equations

In this paper, the direct and adjoint global modes are found using the low Mach number (LMN) formulation of the Navier–Stokes equations [26,29]. Starting from the non-dimensionalized fully-compressible Navier–Stokes equations, each variable is expressed in a form similar to $p = p^{(0)} + \gamma Ma^2 p^{(1)} + \cdots$, where p is the non-dimensional pressure. The Mach number, Ma, and the ratio of specific heats, γ, are defined in Table 1. In the limit of low Mach number, all terms of order Ma^2 disappear except for the pressure term in the momentum equation, which contains the factor $1/\gamma Ma^2$. This term implies that the LMN momentum equation is balanced by $\nabla p^{(1)}$, and that $\nabla p^{(0)} = \mathbf{0}$.

The domain is cylindrical and uses a collocated grid with cylindrical polar coordinates (r, θ, x). The inlet (the jet exit plane) is located at the axial location $x = 0$ and the outlet is at $x = X_{max}$, with uniformly spaced grid points in the axial direction. The lateral boundaries are at a radial distance $r = R_{max}$ from the centreline, with skewed grid spacing in the radial direction. For the results in this paper, only axisymmetric flows are considered, but the analysis can be extended to fully 3D regimes. The full derivation of the nonlinear equations along with a full description of the computational domain and non-dimensionalizations are given in Ref. [29]. The definitions of all non-dimensionalized variables and parameters are given in Table 1. The superscript (0) has been dropped on all non-dimensional variables except for p. Dimensional variables are denoted by $(\sim)$, ambient conditions are denoted by subscript (0) and conditions for the jet at entry are denoted by subscript (j). The non-dimensional nonlinear low-Mach-number equations are

$$\frac{D\rho}{Dt} + \rho(\nabla \cdot \boldsymbol{u}) = 0, \tag{1a}$$

$$\frac{\partial(\rho \boldsymbol{u})}{\partial t} + \nabla p^{(1)} - \nabla \cdot \left(\frac{1}{S_1 Re} \boldsymbol{\tau} - \rho \boldsymbol{u}\boldsymbol{u} \right) + Ri(1-\rho)\hat{\boldsymbol{g}} = \mathbf{0}, \tag{1b}$$

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

Table 1
Non-dimensionalized variables and parameters.

Reference quantities	
Ambient density (kg m^{-3})	$\tilde{\rho}_0$
Ambient temperature (K)	$\tilde{T}_0$
Jet density (kg m^{-3})	$\tilde{\rho}_j$
Jet temperature (K)	$\tilde{T}_j$
Jet diameter (m)	$\tilde{d}_j$
l^2-norm of jet velocity (m s^{-1})	$\|\tilde{\boldsymbol{u}}_j\| \equiv \tilde{u}_j$
Dynamic viscosity (N m^{-2} s)	μ
Thermal conductivity (W m^{-1} K^{-1})	λ
Gas constant (J kg^{-1} K^{-1})	R
Specific heat capacity at constant pressure (J kg^{-1} K^{-1})	c_p
Ratio of specific heats	γ
Acceleration due to gravity (m s^{-2})	$g \equiv \|\boldsymbol{g}\|$
Non-dimensional variables	
Velocity	$\boldsymbol{u} \equiv \tilde{\boldsymbol{u}}/\tilde{u}_j$
Density	$\rho \equiv \tilde{\rho}/\tilde{\rho}_0$
Temperature	$T \equiv \left(\tilde{T} - \tilde{T}_0\right)/\left(\tilde{T}_j - \tilde{T}_0\right)$
Pressure	$p \equiv \tilde{p}/\left(\tilde{\rho}_0 R\tilde{T}_0\right) \approx p^{(0)} + \gamma Ma^2 p^{(1)}$
Time	$t \equiv \tilde{t}\,\tilde{u}_j/\tilde{d}_j$
Spatial derivatives	$\boldsymbol{\nabla} \equiv \tilde{d}_j\tilde{\boldsymbol{\nabla}}$
Non-dimensional parameters:	
Mach number	$Ma \equiv \tilde{u}_j/\left(\gamma R\tilde{T}_0\right)^{1/2}$
Reynolds number	$Re \equiv \tilde{\rho}_j\tilde{u}_j\tilde{d}_j/\mu$
Prandtl number	$Pr \equiv \mu c_p/\lambda$
Richardson number	$Ri \equiv g\tilde{d}_j/\tilde{u}_j^2$
Density ratio	$S_1 \equiv \tilde{\rho}_0/\tilde{\rho}_j$

$$\rho\frac{\mathrm{D}T}{\mathrm{D}t} - \frac{1}{S_1 RePr}\nabla^2 T = 0, \tag{1c}$$

$$\rho((S_1 - 1)T + 1) = p^{(0)}, \tag{1d}$$

$$\boldsymbol{\nabla}p^{(0)} = \boldsymbol{0}, \tag{1e}$$

$$\boldsymbol{\tau} \equiv \left[\boldsymbol{\nabla}\boldsymbol{u} + (\boldsymbol{\nabla}\boldsymbol{u})^T\right] - \frac{2}{3}(\boldsymbol{\nabla}\cdot\boldsymbol{u})\boldsymbol{I},$$

where the identity matrix, $\boldsymbol{I}$, has the same matrix dimensions as $\boldsymbol{\nabla}\boldsymbol{u}$ and $\boldsymbol{\tau}$ is the viscous stress tensor. The vector $\hat{\boldsymbol{g}}$ is a unit vector in the direction of $\boldsymbol{g}$, the acceleration due to gravity. The pressures $p^{(0)}$ and $p^{(1)}$ are known as the thermodynamic pressure and hydrodynamic pressure respectively because $p^{(0)}$ is determined by the temperature in the state Eq. (1d) and $p^{(1)}$ is determined by the hydrodynamic forces in the momentum Eq. (1b).

The physical properties μ, λ, R, c_p and γ described in Table 1, are assumed to be uniform and constant and can therefore be used directly in the non-dimensional variables and parameters. The Prandtl number Pr expresses a non-dimensional ratio of momentum diffusivity to thermal diffusivity. The parameter S_1 is the ratio of ambient density to jet density, which implies that $S_1 Re$ forms a Reynolds number based on the ambient density instead of the jet density. The Richardson number, Ri expresses a ratio of gravitational potential energy to kinetic energy. For the results in this paper, it is set to zero, but we carry it through the derivation as it may be useful in some cases, such as vertically orientated diffusion flames.

The dimensional equation of state, $\tilde{p} = \tilde{\rho}R\tilde{T}$, implies that the ambient pressure $\tilde{p}_0 = \tilde{\rho}_0 R\tilde{T}_0$. Together with the definition of p in Table 1, this implies that the non-dimensional ambient pressure $p_0 = \tilde{p}_0/\tilde{\rho}_0 R\tilde{T}_0 = 1$. Eq. (1e) implies that $p^{(0)}$ is uniform in space and, assuming that the boundary conditions on $p^{(0)}$ are time independent, $p^{(0)}$ is also constant in time. Under the assumption of low Mach number, $p^{(0)} \approx p$ and we conclude therefore that $p^{(0)} \approx 1$.

The base flow is a steady solution to (1) and satisfies

$$\boldsymbol{\nabla}\cdot\bar{\boldsymbol{m}} = 0, \tag{2a}$$

$$\boldsymbol{\nabla}\bar{p}^{(1)} - \boldsymbol{\nabla}\cdot\left(\frac{1}{S_1 Re}\bar{\boldsymbol{\tau}} - \bar{\rho}\bar{\boldsymbol{u}}\bar{\boldsymbol{u}}\right) + Ri(1-\bar{\rho})\hat{\boldsymbol{g}} = \boldsymbol{0}, \tag{2b}$$

$$\bar{\boldsymbol{m}}\cdot\boldsymbol{\nabla}\overline{T} - \frac{1}{S_1 RePr}\nabla^2\overline{T} = 0, \tag{2c}$$

$$\bar{\rho}((S_1 - 1)\overline{T} + 1) = 1, \tag{2d}$$

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

$$\bar{\boldsymbol{\tau}} \equiv \left[\nabla\bar{\boldsymbol{u}} + (\nabla\bar{\boldsymbol{u}})^T\right] - \frac{2}{3}(\nabla\cdot\bar{\boldsymbol{u}})\boldsymbol{I},$$
$$\bar{\boldsymbol{m}} \equiv \bar{\rho}\bar{\boldsymbol{u}},$$

where $p^{(0)}$ has been replaced by the value 1 in (2d). The nonlinear Eq. (1) are linearized about the base flow to form:

$$\frac{\partial\rho'}{\partial t} + \nabla\cdot\boldsymbol{m}' = 0, \tag{3a}$$
$$\frac{\partial\boldsymbol{m}'}{\partial t} + \nabla\cdot(\bar{\rho}\bar{\boldsymbol{u}}\boldsymbol{u}' + \bar{\rho}\boldsymbol{u}'\bar{\boldsymbol{u}} + \rho'\bar{\boldsymbol{u}}\bar{\boldsymbol{u}}) + \nabla p' - \frac{1}{S_1 Re}\nabla\cdot\boldsymbol{\tau}' - Ri\rho'\hat{\boldsymbol{g}} = \boldsymbol{0}, \tag{3b}$$
$$\bar{\rho}\frac{\partial T'}{\partial t} + \bar{\rho}\bar{\boldsymbol{u}}\cdot\nabla T' + \boldsymbol{m}'\cdot\nabla\overline{T} - \frac{1}{S_1 RePr}\nabla^2 T' = 0, \tag{3c}$$
$$\frac{\rho'}{\bar{\rho}} + \frac{(S_1-1)T'}{((S_1-1)\overline{T}+1)} = 0, \tag{3d}$$
$$\boldsymbol{\tau}' \equiv \left[\nabla\boldsymbol{u}' + (\nabla\boldsymbol{u}')^T\right] - \frac{2}{3}(\nabla\cdot\boldsymbol{u}')\boldsymbol{I}, \tag{3e}$$
$$\boldsymbol{m}' \equiv \bar{\rho}\boldsymbol{u}' + \bar{\boldsymbol{u}}\rho', \tag{3f}$$

where $\boldsymbol{m}'$ is the linearized momentum perturbation and p' is the linear perturbation of $p^{(1)}$. For the rest of this paper, the primes on the direct linear perturbation variables are dropped and p refers to the linear perturbation of the hydrodynamic pressure $p^{(1)}$.

3. The adjoint governing equations

The system of linearized perturbation equations, (3), can be expressed in terms of the direct linear operator, L:

$$\frac{\partial\boldsymbol{q}}{\partial t} - \mathrm{L}\boldsymbol{q} = \boldsymbol{0}, \tag{4}$$

$$\boldsymbol{q} \equiv \begin{bmatrix}\boldsymbol{m}\\ T\end{bmatrix},$$

where $\boldsymbol{m}$ and T are functions of the axial and radial coordinates, x and r, and time t. The state vector $\boldsymbol{q}$ does not contain p and ρ because these can be derived from $\boldsymbol{m}$ and T using (3).

Our main interest is calculating the global modes of (4) and their corresponding frequencies and growth rates. We are therefore interested in non-trivial solutions of (4) of the form

$$\boldsymbol{q}(x,r,t) = \hat{\boldsymbol{q}}(x,r)\exp(\sigma t). \tag{5}$$

Substituting (5) into (4), we can write the new system of equations as

$$\sigma\hat{\boldsymbol{q}} - \mathrm{L}\hat{\boldsymbol{q}} = \boldsymbol{0}. \tag{6}$$

To form the adjoint equations and find the adjoint global modes, an inner product must be defined:

$$\langle\boldsymbol{q}_1,\boldsymbol{q}_2\rangle \equiv \frac{1}{V}\int_V \boldsymbol{q}_1^H\boldsymbol{q}_2\,dV, \tag{7}$$

where $\boldsymbol{q}_1^H$ is the Hermitian of $\boldsymbol{q}_1$. In this notation, $\boldsymbol{q}_1^H\boldsymbol{q}_2 \equiv \sum_i q_{1i}^* q_{2i}$, where superscript $(*)$ denotes the complex conjugate.

The adjoint eigenmodes themselves depend on this choice of norm but, when they are recombined with the direct modes to give the sensitivity of the eigenvalue to changes in the operator L, and other robust measures, the effect of the norm cancels out. Using the inner product, Eq. (6) is premultiplied by an arbitrary vector, $\hat{\boldsymbol{q}}^+$, to give:

$$\langle\hat{\boldsymbol{q}}^+,\sigma\hat{\boldsymbol{q}}\rangle - \langle\hat{\boldsymbol{q}}^+,\mathrm{L}\hat{\boldsymbol{q}}\rangle = 0. \tag{8}$$

The vector $\hat{\boldsymbol{q}}^+$ will soon be identified with the adjoint global mode corresponding to $\hat{\boldsymbol{q}}$. The adjoint equations can be derived from the continuous direct Eq. (8) and then discretized or can be derived directly from the discretized direct equations. These two paths, although subtly different, are theoretically equivalent. They are known as 'discretization of the adjoint system' (DA) versus 'adjoint of the discrete system' (AD), [35], 'finite difference of adjoint' (FDA) versus 'adjoint of finite difference' (AFD), [36], and 'optimize then discretize' (OTD) versus 'discretize then optimize' (DTO), [16]. Both paths are reviewed by Ref. [15] in the context of design methods for steady flows. Some of the programming benefits of the AD approach in the context of aircraft design are described in Ref. [37], as well as software that can calculate the discrete-adjoint operator automatically.

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

3.1. Continuous-adjoint approach

3.1.1. General formulation

The continuous-adjoint operator of L is labelled L^+ and is defined implicitly by

$$\langle L^+\hat{\boldsymbol{q}}^+, \hat{\boldsymbol{q}}\rangle \equiv \langle \hat{\boldsymbol{q}}^+, L\hat{\boldsymbol{q}}\rangle. \tag{9}$$

L^+ can be found by rearranging (8) with integration by parts to give:

$$\langle \sigma^*\hat{\boldsymbol{q}}^+, \hat{\boldsymbol{q}}\rangle - \langle L^+\hat{\boldsymbol{q}}^+, \hat{\boldsymbol{q}}\rangle = b. \tag{10}$$

By selecting appropriate boundary conditions, the boundary term, b, can be set to zero, which means that (9) is satisfied. At this point $\hat{\boldsymbol{q}}^+$ can still be an arbitrary vector but, in order for (10) to be satisfied for arbitrary $\hat{\boldsymbol{q}}$, $\hat{\boldsymbol{q}}^+$ must satisfy

$$\sigma^*\hat{\boldsymbol{q}}^+ - L^+\hat{\boldsymbol{q}}^+ = \boldsymbol{0}. \tag{11}$$

A vector $\hat{\boldsymbol{q}}^+$ satisfying (11) is a global mode of the continuous-adjoint operator L^+ with a corresponding eigenvalue equal to σ^*. Global modes of L^+ will be referred to hereafter as continuous-adjoint global modes and are time-independent solutions of the system

$$-\frac{\partial \boldsymbol{q}^+}{\partial t} - L^+\boldsymbol{q}^+ = \boldsymbol{0}, \tag{12}$$

$$\text{where} \quad \boldsymbol{q}^+ \equiv \boldsymbol{q}^+(x, r, t) \equiv \hat{\boldsymbol{q}}^+(x, r)\exp(-\sigma^* t),$$

$$\text{and} \quad \boldsymbol{q}^+ \equiv \begin{bmatrix} \boldsymbol{m}^+ \\ T^+ \end{bmatrix}.$$

The minus sign in front of the time derivative term in (12) appears due to the integration by parts of the time derivative term in (4). Remembering that adjoint simulations are run backwards in time, t could be replaced with $-t^+$, which would put the adjoint system of (11) and (12) in a similar form to the direct system of (4)–(6).

3.1.2. Specific formulation

The time derivative terms and continuous-adjoint operator form a set of adjoint equations, which are created by successive integration by parts of the set of direct Eq. (3). The exact form of the adjoint equations depends on the exact form of the direct equations. The aim of this section is to find a set of adjoint equations that closely mimics the set of direct equations so that they both can be solved using a similar algorithm. The incompressible and fully-compressible direct Navier–Stokes equations are naturally in a suitable form. The low-Mach-number Navier–Stokes equations, however, are not.

A naive approach would be to apply integration by parts to (3). The $\partial\rho/\partial t$ term in (3a), however, leads to adjoint equations that require $\partial p^+/\partial t$ to be approximated from p^+ at the current and previous time steps. This approximation is not present in the direct algorithm and leads to large numerical errors. The direct and adjoint low-Mach-number equations closely mimic one another when the direct equations take the form

$$(S_1 - 1)\left(\bar{m}_i\frac{\partial \widehat{T}}{\partial x_i} - \frac{1}{S_1 RePr}\frac{\partial^2 \widehat{T}}{\partial x_i^2}\right) + \frac{\partial}{\partial x_i}\left(\frac{\hat{m}_i}{\bar{\rho}}\right) = 0, \tag{13a}$$

$$\sigma\hat{m}_i + \frac{\partial}{\partial x_j}\left(\frac{\bar{m}_j\hat{m}_i}{\bar{\rho}} + \frac{\hat{m}_j\bar{m}_i}{\bar{\rho}} - \frac{\hat{\rho}\bar{m}_j\bar{m}_i}{\bar{\rho}^2}\right) + \frac{\partial \hat{p}}{\partial x_i} - \frac{1}{S_1 Re}\left(\frac{\partial^2}{\partial x_j^2}\left(\frac{\hat{m}_i}{\bar{\rho}} - \frac{\hat{\rho}\bar{m}_i}{\bar{\rho}^2}\right) + \frac{1}{3}\frac{\partial^2}{\partial x_j \partial x_i}\left(\frac{\hat{m}_j}{\bar{\rho}} - \frac{\hat{\rho}\bar{m}_j}{\bar{\rho}^2}\right)\right) - Ri\hat{\rho}\hat{g}_i = 0, \tag{13b}$$

$$\sigma\widehat{T} + \frac{\bar{m}_i}{\bar{\rho}}\frac{\partial \widehat{T}}{\partial x_i} + \frac{\hat{m}_i}{\bar{\rho}}\frac{\partial \overline{T}}{\partial x_i} - \frac{1}{S_1 RePr}\frac{1}{\bar{\rho}}\frac{\partial^2 \widehat{T}}{\partial x_i^2} = 0, \tag{13c}$$

$$\frac{\hat{\rho}}{\bar{\rho}^2} + (S_1 - 1)\widehat{T} = 0. \tag{13d}$$

To obtain (13) from (3) a number of steps have been taken: (5) is substituted into (3); $\boldsymbol{u}$ and $\boldsymbol{\tau}$ are removed by using (3f) and (3e); (2d) is substituted into (3d) to remove $\overline{T}$; (3a), (3c) and (3d) are divided through by $\bar{\rho}$; (3d) is differentiated with respect to time and substituted into (3c), which is then substituted into (3a); and the terms that contained $\hat{m}_i$ in (13a) have been combined using (2d) to form:

$$(S_1 - 1)\hat{m}_i\frac{\partial \overline{T}}{\partial x_i} + \frac{1}{\bar{\rho}}\frac{\partial \hat{m}_i}{\partial x_i} = \hat{m}_i\frac{\partial}{\partial x_i}\left((S_1 - 1)\overline{T} + 1\right) + \left((S_1 - 1)\overline{T} + 1\right)\frac{\partial \hat{m}_i}{\partial x_i} = \frac{\partial}{\partial x_i}\left(\frac{\hat{m}_i}{\bar{\rho}}\right). \tag{14}$$

If the adjoint variable multiplying (13a) is labelled $\hat{p}^+$, then forming the adjoint equations with the terms on the left hand side in (14) would result in both $\hat{p}^+$ and $\nabla\hat{p}^+$ appearing in the adjoint momentum equation. This would make the formation

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

and solution of a Poisson equation more difficult. The rearrangement in (14) allows the adjoint momentum equation to be formed with just $\nabla\hat{p}^+$ and is therefore necessary for the adjoint Poisson equation to be solved.

Two new adjoint variables are required to premultiply (13a) and (13d). The labelling of these variables is unimportant but it is helpful to think of the variable premultiplying (13a) as the adjoint pressure and the variable premultiplying (13d) as the adjoint density. The adjoint momentum and temperature naturally premultiply (13b) and (13c) respectively. Through successive integration by parts, the adjoint equations corresponding to (13) are

$$\frac{\partial \hat{m}_i^+}{\partial x_i} = 0, \tag{15a}$$

$$\sigma^* \hat{m}_i^+ - \frac{\bar{m}_j}{\bar{\rho}}\left(\frac{\partial \hat{m}_i^+}{\partial x_j} + \frac{\partial \hat{m}_j^+}{\partial x_i}\right) - \frac{1}{S_1 Re\bar{\rho}}\left(\frac{\partial^2 \hat{m}_i^+}{\partial x_j^2} + \frac{1}{3}\frac{\partial^2 \hat{m}_j^+}{\partial x_j \partial x_i}\right) - \frac{1}{\bar{\rho}}\frac{\partial \hat{p}^+}{\partial x_i} + \frac{\widehat{T}^+}{\bar{\rho}}\frac{\partial \overline{T}}{\partial x_i} = 0, \tag{15b}$$

$$\sigma^* \widehat{T}^+ - \bar{m}_i \frac{\partial}{\partial x_i}\left(\frac{\widehat{T}^+}{\bar{\rho}} + (S_1 - 1)\hat{p}^+\right) - \frac{1}{S_1 RePr}\frac{\partial^2}{\partial x_i^2}\left(\frac{\widehat{T}^+}{\bar{\rho}} + (S_1 - 1)\hat{p}^+\right) + (S_1 - 1)\hat{\rho}^+ = 0, \tag{15c}$$

$$\bar{m}_i \bar{m}_j \frac{\partial \hat{m}_i^+}{\partial x_j} + Ri\bar{\rho}^2 \hat{m}_i^+ \hat{g}_i + \frac{\bar{m}_i}{S_1 Re}\left(\frac{\partial^2 \hat{m}_i^+}{\partial x_j^2} + \frac{1}{3}\frac{\partial^2 \hat{m}_j^+}{\partial x_j \partial x_i}\right) + \hat{\rho}^+ = 0. \tag{15d}$$

Eq. (15) shows that $\hat{p}^+$ and $\hat{\rho}^+$ can be derived from $\hat{\boldsymbol{m}}^+$ and $\widehat{T}^+$ in the same way as $\hat{p}$ and $\hat{\rho}$ can be derived from $\hat{\boldsymbol{m}}$ and $\widehat{T}$ using (13).

In forming (15c), the base flow condition (2a) has been used to shift $\bar{m}_i$ outside the derivative. This ensures that the adjoint pressure appears only inside a derivative, which reduces numerical error because the solution to the adjoint pressure Poisson equation is accurate only up to an arbitrary constant. Alternatively, $(\widehat{T}/\bar{\rho})\nabla\cdot\bar{\boldsymbol{m}}$ could have been added to the left-hand side of (13c). The base flow condition (2a) ensures this extra term is approximately zero, but now integration by parts gives (15c) directly without requiring any further use of (2a). In practice, the forms in (13c) and (15c) give the smallest discrepancy between the direct and adjoint global mode frequencies and have therefore been used for this work.

3.2. Discrete-adjoint approach

Eqs. (4), (6) and (8) are discretized to give:

$$\frac{\partial \boldsymbol{q}}{\partial t} - \boldsymbol{A}\boldsymbol{q} = \boldsymbol{0}, \tag{16}$$

$$\sigma\hat{\boldsymbol{q}} - \boldsymbol{A}\hat{\boldsymbol{q}} = \boldsymbol{0}, \tag{17}$$

$$(\hat{\boldsymbol{q}}^+)^H(\boldsymbol{D}\sigma\hat{\boldsymbol{q}}) - (\hat{\boldsymbol{q}}^+)^H(\boldsymbol{D}\boldsymbol{A}\hat{\boldsymbol{q}}) = \boldsymbol{0}, \tag{18}$$

where the state vector, $\boldsymbol{q}$, and global mode vector, $\hat{\boldsymbol{q}}$, have changed from continuous functions, $\boldsymbol{q}(x,r,t)$ and $\hat{\boldsymbol{q}}(x,r)$, to discrete arrays of vectors located at grid points that span the domain. Alternatively the continuous scalar fields that form the continuous vectors can be discretized to form separate discrete scalar fields, which are then combined to form discrete vectors of arrays. The matrix $\boldsymbol{A}$ is the discretized version of the direct linear operator L. The diagonal volume distribution matrix $\boldsymbol{D}$ represents the volumes associated with each grid point as a fraction of the total volume of the domain and is required so that the terms in (18) are equivalent to the inner products in (8). The product $\boldsymbol{q}_1^H\boldsymbol{D}\boldsymbol{q}_2$ represents a discrete volume summation in the same way that the inner product (7) represents a volume integral.

To find the discrete-adjoint global mode, each of the terms in (18) is rearranged so that the discrete direct global mode alone follows the Hermitian operator:

$$\left(\sigma^*\boldsymbol{D}^H\hat{\boldsymbol{q}}^+\right)^H(\hat{\boldsymbol{q}}) - \left(\boldsymbol{A}^H\boldsymbol{D}^H\hat{\boldsymbol{q}}^+\right)^H(\hat{\boldsymbol{q}}) = \boldsymbol{0}, \tag{19}$$

$$\Rightarrow \quad \text{for arbitrary } \hat{\boldsymbol{q}} : \sigma^*\boldsymbol{D}^H\hat{\boldsymbol{q}}^+ - \boldsymbol{A}^H\boldsymbol{D}^H\hat{\boldsymbol{q}}^+ = \boldsymbol{0} \tag{20}$$

$$\Rightarrow \sigma^*\hat{\boldsymbol{q}}^+ - (\boldsymbol{D}^H)^{-1}\boldsymbol{A}^H\boldsymbol{D}^H\hat{\boldsymbol{q}}^+ = \boldsymbol{0} \tag{21}$$

$$\Rightarrow \sigma^*\hat{\boldsymbol{q}}^+ - \boldsymbol{D}^{-1}\boldsymbol{A}^H\boldsymbol{D}\hat{\boldsymbol{q}}^+ = \boldsymbol{0}, \tag{22}$$

where the final step arises because $\boldsymbol{D}$ is a real diagonal matrix and therefore $\boldsymbol{D}^H = \boldsymbol{D}$. By comparing (20) to (11) it is clear that the discretized adjoint operator, $\boldsymbol{A}^+$ is given by:

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

$$\boldsymbol{A}^{+} \equiv \boldsymbol{D}^{-1}\boldsymbol{A}^{H}\boldsymbol{D} \tag{23}$$

and that the discrete-adjoint global mode, labelled $\hat{\boldsymbol{q}}^{\oplus}$, is equivalent to the continuous-adjoint global mode, $\hat{\boldsymbol{q}}^{+}$. By establishing a rigorous relationship between the discrete- and continuous-adjoint global modes and eigenvalues, the two approaches can be seen to be equivalent.

4. Numerical procedure

A nonlinear axisymmetric DNS code is used to obtain a steady solution of the base flow equations. If the flow is unstable, the base flow may be calculated by adding selective frequency damping (SFD), [38], to the nonlinear Eq. (1). At a fixed point, the effect of the SFD terms is nil, but they allow an unstable fixed point to be reached numerically by artificially damping unstable frequencies along the convergence path. In our calculations, simple time-stepping is used until the SFD system converges to a steady solution that satisfies (2).

The code uses axisymmetric cylindrical polar coordinates with sixth-order compact finite difference schemes [39] to compute spatial derivatives in the axial and radial directions. Values along the centreline of the domain are calculated with asymptotic equations [40] that differ slightly from the regular equations. A 4th-order Runga–Kutta time marching scheme is used, involving an explicit version of the projection method used by [41], which requires the solution of a pressure Poisson equation. The Poisson solver uses direct cosine transforms (DCTs) along each line of grid points in the axial direction to decouple the system into separate wave numbers. A standard complex fast Fourier transform (FFT) algorithm is used with pre- and postprocessing routines that allow the efficient computation of two real DCTs for each complex FFT [42–44]. The forward DCTs are applied first to the right hand side of the Poisson equation. An efficient algorithm (based on the Thomas algorithm) is then applied to solve the decoupled system of one-dimensional equations. Finally the inverse DCTs are applied to give the pressure. Further details of the code can be found in [28]. This section contains a summary of the relevant parts of the direct and adjoint solution algorithms.

4.1. Computing global modes

The eigenvalues and global modes are found using ARPACK, which uses the implicitly restarted Arnoldi method [45]. An introduction to these methods is given by Ref. [2]. ARPACK calculates the eigenvalues and eigenvectors of a virtual time-stepper matrix that represents the evolution of the flow from time t to time $t + N\Delta t$, where Δt is the time advanced by a single application of the time-stepping algorithm and N is the number of applications. If an explicit single-step time discretization is used, such as a Runge–Kutta method, (16) and its discrete-adjoint equivalent can be expressed as

$$\boldsymbol{q}(t + N\Delta t) = \boldsymbol{M}^{N}\boldsymbol{q}(t), \tag{24}$$

$$\boldsymbol{q}^{\oplus}(t - N\Delta t) = (\boldsymbol{M}^{H})^{N}\boldsymbol{q}^{\oplus}(t), \tag{25}$$

where $\boldsymbol{M}$ is the matrix exponential of $\boldsymbol{A}\Delta t$, $\boldsymbol{M}^{H}$ is the matrix exponential of $\boldsymbol{A}^{H}\Delta t$ and superscript (N) denotes the matrix raised to the power of N.

The continuous-adjoint described by (12) is discretized in space and time to give:

$$-\frac{\partial \boldsymbol{q}^{+}}{\partial t} - \boldsymbol{A}^{+}\boldsymbol{q}^{+} = \boldsymbol{0}, \tag{26}$$

$$\text{and} \quad \boldsymbol{q}^{+}(t - N\Delta t) = \boldsymbol{M}^{+N}\boldsymbol{q}^{+}(t), \tag{27}$$

where $\boldsymbol{M}^{+}$ is the matrix exponential of $\boldsymbol{A}^{+}\Delta t$.

The eigenmodes and eigenvalues calculated by ARPACK correspond to the matrices $\boldsymbol{M}^{N}$, $(\boldsymbol{M}^{H})^{N}$ and $\boldsymbol{M}^{+N}$. The eigenmodes are the same as those of the matrices $\boldsymbol{A}$, $\boldsymbol{A}^{H}$ and $\boldsymbol{A}^{+}$, but the eigenvalues depend on $N\Delta t$. From (17), σ is an eigenvalue of $\boldsymbol{A}$ and if λ is the corresponding eigenvalue of $\boldsymbol{M}^{N}$ then

$$\lambda \equiv \exp(\sigma N\Delta t), \tag{28}$$

$$\Rightarrow \sigma \equiv \frac{1}{N\Delta t}(\ln(|\lambda|) + \mathrm{i}\arg(\lambda)), \tag{29}$$

where $||$ is the absolute value of a complex number.

A two-sided Arnoldi algorithm, which gives both the direct and adjoint eigenmodes, has been proposed by [46]. The two-sided algorithm uses a combination of the approximate left eigenvectors of the direct solution as the initial state for the adjoint calculation. While this improves the speed of convergence of the adjoint algorithm, the use of parallel computing allows the direct and adjoint calculations to be run side by side and is therefore quicker overall. [11] use a two-sided inverse iteration algorithm to efficiently calculate the discrete-adjoint modes at the same time as the direct modes.

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

4.2. Direct linear algorithm

Converting (13) into time-dependent form and discretizing in time with a 1st-order explicit Euler scheme for illustration, the direct system can be written as

$$\nabla \cdot \boldsymbol{m}_{n+1} - k\,\mathrm{h}(\boldsymbol{m}_{n+1}, T_{n+1}) = 0, \tag{30a}$$

$$\frac{\boldsymbol{m}_{n+1} - \boldsymbol{m}_n}{\Delta t} + \nabla p_* - \mathrm{f}(\boldsymbol{m}_n, \rho_n) = \boldsymbol{0}, \tag{30b}$$

$$\frac{T_{n+1} - T_n}{\Delta t} - \mathrm{h}(\boldsymbol{m}_n, T_n) = 0, \tag{30c}$$

$$\rho_n + kT_n = 0, \tag{30d}$$

where f and h are functions representing the remaining terms in the equations and k is a constant scalar field that is dependent upon the base flow. Subscript (n) denotes the value of the variable at time step n and subscript $(*)$ denotes an intermediate value.

The variable ρ_n can be eliminated using (30d). The variable p_*, however, requires the formation of a Poisson equation. Taking the divergence of (30b) and rearranging gives

$$\nabla^2 p_* = \frac{1}{\Delta t}\left(\nabla \cdot [\boldsymbol{m}_n + \Delta t\,\mathrm{f}(\boldsymbol{m}_n, \rho_n)] - \nabla \cdot \boldsymbol{m}_{n+1}\right). \tag{31}$$

The variables T_{n+1} from (30c) and $\boldsymbol{m}_n$ are substituted into (30a) to give an approximation for $\nabla \cdot \boldsymbol{m}_{n+1}$

$$\nabla \cdot \boldsymbol{m}_{n+1} \approx k\,\mathrm{h}(\boldsymbol{m}_n, T_{n+1}),$$

$$\nabla \cdot \boldsymbol{m}_{n+1} \approx (S_1 - 1)\bar{\rho}\left(-\boldsymbol{m}_n \cdot \nabla\overline{T} - \bar{\boldsymbol{m}} \cdot \nabla T_{n+1} + \frac{1}{S_1 RePr}\nabla^2 T_{n+1}\right). \tag{32}$$

The Poisson equation is then solved for p_* and (30b) can then be used to find $\boldsymbol{m}_{n+1}$. The term inside the square brackets in (31) is the auxiliary momentum

$$\boldsymbol{m}_* = [\boldsymbol{m}_n + \Delta t\,\mathrm{f}(\boldsymbol{m}_n, \rho_n)], \tag{33}$$

$$\Rightarrow \boldsymbol{m}_{n+1} = \boldsymbol{m}_* - \Delta t \nabla p_*. \tag{34}$$

To improve temporal accuracy a 4th-order Runge–Kutta scheme is used, which is formed from the 1st-order Euler scheme shown above.

4.3. Continuous-adjoint algorithm

Converting (15) into time-dependent form and discretizing in time with a 1st-order explicit Euler scheme, the continuous-adjoint system can be written as

$$\nabla \cdot \boldsymbol{m}^+_{n+1} = 0, \tag{35a}$$

$$-\left(\frac{\boldsymbol{m}^+_{n+1} - \boldsymbol{m}^+_n}{\Delta t}\right) - \frac{1}{\bar{\rho}}\nabla p^+_* + \mathrm{f}^+(\boldsymbol{m}^+_n, T^+_n) = \boldsymbol{0}, \tag{35b}$$

$$-\frac{T^+_{n+1} - T^+_n}{\Delta t} + \mathrm{h}^+(\rho^+_n, T^+_n, p^+_n) = 0, \tag{35c}$$

$$\rho^+_n + \mathrm{k}^+(\boldsymbol{m}^+_n) = 0, \tag{35d}$$

where f^+, h^+ and k^+ are functions representing the remaining terms in (15b), (15c) and (15d) respectively.

The derived variable ρ^+_n can be eliminated by substituting (35d) into (35c). The derived variable p^+_*, however, requires the formation of a Poisson equation. Taking the divergence of (35b) and rearranging gives

$$\nabla^2 p^+_* = \frac{1}{\Delta t}\left(\nabla \cdot \left(\bar{\rho}\left[\boldsymbol{m}^+_n + \Delta t\,\mathrm{f}^+(\boldsymbol{m}^+_n, T^+_n)\right]\right) - \boldsymbol{m}^+_{n+1} \cdot \nabla\bar{\rho} - \bar{\rho}\nabla \cdot \boldsymbol{m}^+_{n+1}\right). \tag{36}$$

Eq. (35a) is used to remove the $\nabla \cdot \boldsymbol{m}^+_{n+1}$ term from (36). The remaining $\boldsymbol{m}^+_{n+1}$ is approximated by $\boldsymbol{m}^+_n$ in a similar way to the way in which $\boldsymbol{m}_{n+1}$ is approximated by $\boldsymbol{m}_n$ in (32). The adjoint Poisson equation is then solved for p^+_*. The adjoint pressure is then substituted into (35) to find $\boldsymbol{m}^+_{n+1}$ and T^+_{n+1}. The term inside the square brackets in (36) is the adjoint auxiliary momentum field

$$\boldsymbol{m}^+_* = \left[\boldsymbol{m}^+_n + \Delta t\,\mathrm{f}^+(\boldsymbol{m}^+_n, T^+_n)\right], \tag{37}$$

$$\Rightarrow \boldsymbol{m}^+_{n+1} = \boldsymbol{m}^+_* - \frac{\Delta t}{\bar{\rho}}\nabla p^+_*. \tag{38}$$

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

In order for the diffusive term in the adjoint momentum equation to be well-behaved it is necessary to use a negative time step, which is achieved by changing the sign in front of terms containing Δt in the above equations. Again, the 4th-order Runge–Kutta scheme is used to improve temporal accuracy.

4.4. Discrete-adjoint algorithm

To form a discrete-adjoint algorithm, the direct algorithm needs to be represented as one matrix (although this matrix is never actually formed) and the discrete-adjoint is the conjugate transpose of this matrix. First, all the steps in the direct solution algorithm are broken down and represented as simple matrix–vector products. At the start of the time step, only the variables contained in the state vector exist. At points in the time step algorithm, derived variables are introduced and the vector grows. When the derived variables are no longer needed they are removed and the vector shrinks, eventually back to the size at which the time step started. For example, the state equation, (3d), is represented as a matrix with 4 rows and 3 columns that multiplies the state vector containing three scalar fields: m_x, m_r and T. The top 3 rows of this matrix form a 3×3 identity matrix. The bottom row calculates ρ from a combination of m_x, $m_r$and T and this forms the fourth field in a new vector containing m_x, m_r, T and ρ. Once ρ is no longer needed, a matrix with 3 rows and 4 columns is used to convert back to a three-field vector.

The scalar fields are discretized to form a single state vector containing scalar values. This is achieved by consecutively storing values in the axial direction at each discrete radial location in turn. The middle values for each discretized scalar field correspond therefore to the centreline of the domain. The middle rows of the corresponding matrix represent the centreline equations and the other rows represent the regular equations. From this starting point every action that alters the state vector is represented as a matrix–vector product, which is then transposed. A few key highlights are discussed below with full details provided in [1].

The Poisson solver requires the transpose of forward and inverse DCTs. Conveniently the inverse DCT is similar to the transpose of the forward DCT and the forward DCT is similar to the transpose of the inverse DCT. In both cases, only the value corresponding to the constant coefficient requires altering, which can be easily done in the pre- and postprocessing routines.

A matrix representation of the 4th-order Runge–Kutta scheme is formed by expanding the state vector at each of the four steps. At each intermediate Runge–Kutta step the new vector is added to the bottom of the state vector. The last matrix–vector product is then

$$[\boldsymbol{q}_{n+1}] = [0\ \boldsymbol{I}/6\ \boldsymbol{I}/3\ \boldsymbol{I}/3\ \boldsymbol{I}/6] \begin{bmatrix} \boldsymbol{q}_n \\ \boldsymbol{q}_n^1 \\ \boldsymbol{q}_n^2 \\ \boldsymbol{q}_n^3 \\ \boldsymbol{q}_n^4 \end{bmatrix}. \tag{39}$$

With the Runge–Kutta scheme represented in this way the separate matrices can be transposed and reordered accordingly to provide the exact transpose of the whole 4th-order Runga–Kutta algorithm.

5. Boundary conditions

One of the most challenging aspects of developing an adjoint solver is the correct treatment of boundary conditions. Similar to the direct problem, this depends on the specific flow configuration. In this section, therefore, we carefully describe boundary conditions developed for the base flow, direct perturbations, and continuous-adjoint perturbations in the case of a hot jet. The discrete-adjoint algorithm in some sense circumvents these difficulties by using the exact matrix transpose, but as we will see introduces numerical artifacts near the boundaries.

The boundary conditions used for the base flow, direct perturbations, and continuous-adjoint perturbations are summarized in Table 2. The discrete-adjoint boundary conditions are the exact matrix transpose of the discretized direct boundary conditions.

Before discussing the boundary conditions for the different simulations, some general points can be made. Although setting boundary conditions on momentum and temperature provides enough information to describe a unique solution [47], the efficient solution of the pressure Poisson equation in the projection method requires boundary conditions for the hydrodynamic pressure and auxiliary momentum fields, which must be consistent with the system of equations and the boundary conditions already set.

On the boundaries, the projection steps ((34) for the direct case) imply that the auxiliary momentum is equal to the actual momentum plus a multiple of $\mathbf{V}p_*$. It is convenient therefore to set $\mathbf{V}p_* = 0$, which allows the boundary conditions on the auxiliary momentum to be the same as those for the actual momentum. To set $\mathbf{V}p_* = 0$, it is necessary and sufficient to set the spatial derivatives tangential to the boundary ($\partial p_*/\partial\tau$) and normal to the boundary ($\partial p_*/\partial n$) to zero, where the spatial co-ordinates τ and n are tangential and normal to the boundary surface respectively. For the axisymmetric case with no swirl and $Re \gg 1$, $\partial p_*/\partial\tau \approx 0$ on all boundaries. A high Reynolds number is required so that viscous corrections near the inlet, which cause a radial pressure gradient at inlet, are small. The solution of the pressure Poisson equation however defines the pressure only to within a constant and the overall pressure level can therefore drift if a Dirichlet condition is not set

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

Table 2
Boundary conditions. D → Dirichlet, h-D → homogeneous Dirichlet, h-N → homogeneous Neumann, C → convective, v-T-F → viscous traction free.

Field	Base flow			Direct linear			Continuous-adjoint		
	$x = 0$	X_{max}	R_{max}	$x = 0$	X_{max}	R_{max}	$x = 0$	X_{max}	R_{max}
m_x	D	C	v-T-F	h-D	C	h-D	h-D	h-D	h-D
m_r	h-N	C	v-T-F	h-D	C	h-D	h-D	h-D	h-D
T	D	C	h-D	h-D	C	h-D	h-D	h-D	h-D
p	h-N	h-D	h-D	h-N	h-D	h-D	h-D	h-N	h-D
m_{x*}	D	C	v-T-F	h-D	C	h-D	h-D	h-D	h-D
m_{r*}	h-N	C	v-T-F	h-D	C	h-D	h-D	h-D	h-D

on part of the pressure boundary. Far from the inlet, where the flow is varying slowly, $\partial p_*/\partial n \approx 0$, allowing a homogeneous-Dirichlet condition to be set explicitly on the lateral and outlet boundaries. $\partial p_*/\partial n = 0$ can then be set as the pressure boundary condition at the inlet. Setting a homogeneous-Neumann pressure condition at the inlet and a homogeneous-Dirichlet condition at the outlet helps generate a pressure build up at start up, which pushes any transients towards the outlet and out of the domain.

5.1. Base flow boundary conditions

The boundary conditions for the base flow are the same as those used in the nonlinear simulations by [29]. At inlet these are Dirichlet for axial momentum and temperature, with top-hat inlet profiles formed from Michalke's profile number two [48]:

$$y(0,r) \equiv \frac{1}{2}\left(1+\tanh\left(\frac{1}{4}\frac{\tilde{d}_j}{2\theta}\left(\frac{1}{2r}-2r\right)\right)\right), \tag{40}$$

$$T(0,r) = y(0,r), \tag{41}$$

$$\rho(0,r) = \frac{1}{(S_1-1)T(0,r)+1}, \tag{42}$$

$$m_x(0,r) = \rho(0,r)(u_c + (1-u_c)y(0,r)), \tag{43}$$

where θ is the momentum thickness, u_c is the non-dimensional coflow velocity and the non-dimensional jet diameter is 1. A homogeneous-Neumann condition is used for the radial momentum. A small positive coflow velocity helps reduce the accumulation of numerical errors by slowly advecting them out of the domain. With $u_c \ll 1$, however, the inlet conditions approximate a jet exiting from a hole in a flat wall. The lateral boundaries allow fluid entrainment through a viscous traction-free momentum boundary condition and a homogeneous-Dirichlet temperature condition. The outlet boundary is a convection boundary condition for all fields [49].

5.2. Direct linear boundary conditions

The base flow at the inlet and lateral boundaries is assumed to be inwards and a homogeneous-Dirichlet boundary condition is set for the momentum and temperature perturbation fields.

For shorter domains, a convection outlet boundary condition is necessary for all perturbation fields to allow disturbances to flow out of the domain with minimal reflections. For long domains, a homogeneous-Dirichlet condition can be used if perturbations have reached a small enough amplitude at the exit. If this is the case, a convection condition produces almost identical results, which is why it has been used for all simulations.

5.3. Continuous-adjoint boundary conditions

The boundary terms, which arise from the integration by parts in the derivation of the adjoint equations, are represented by b in Eq. (10). Assuming $Re \gg 1$, the boundary terms containing a factor of $1/Re$ can be neglected. The remaining boundary terms forming b are

$$\frac{1}{V}\int_V \frac{\partial}{\partial x_i}\left(\frac{\hat{p}^{+*}\hat{m}_i}{\bar{\rho}}\right) + (S_1-1)\frac{\partial}{\partial x_i}\left(\hat{p}^{+*}\hat{T}\bar{m}_i\right) + \frac{\partial}{\partial x_i}\left(\hat{m}_i^{+*}\hat{p}\right) + \frac{\partial}{\partial x_i}\left(\frac{\hat{T}^{+*}\hat{T}\bar{m}_i}{\bar{\rho}}\right)\cdots + \frac{\partial}{\partial x_i}\left(\frac{\bar{m}_i\hat{m}_j\hat{m}_j^{+*}}{\bar{\rho}}\right)$$
$$+\frac{\partial}{\partial x_i}\left(\frac{\hat{m}_i\bar{m}_j\hat{m}_j^{+*}}{\bar{\rho}}\right) + (S_1-1)\frac{\partial}{\partial x_i}\left(\hat{T}\bar{m}_i\bar{m}_j\hat{m}_j^{+*}\right)dV = 0. \tag{44}$$

Using the divergence theorem, the volume integral in (44) can be transformed into a surface integral

$$\frac{1}{V}\int_S \left(\frac{\hat{p}^{+*}\hat{m}_i}{\bar{\rho}}\right) + (S_1-1)\left(\hat{p}^{+*}\widehat{\bar{T}}\bar{m}_i\right) + \left(\hat{m}_i^{+*}\hat{p}\right) + \left(\frac{\widehat{\bar{T}}^{+*}\widehat{\bar{T}}\bar{m}_i}{\bar{\rho}}\right)\cdots + \left(\frac{\bar{m}_i\hat{m}_j\hat{m}_j^{+*}}{\bar{\rho}}\right) + \left(\frac{\hat{m}_i\bar{m}_j\hat{m}_j^{+*}}{\bar{\rho}}\right) + (S_1-1)\left(\widehat{\bar{T}}\bar{m}_i\bar{m}_j\hat{m}_j^{+*}\right)n_i dS = 0, \tag{45}$$

where n_i is the outward pointing unit normal vector of the surface S.

The flow direction for the adjoint is reversed. The inlet is at X_{max} and the outlet is at $x = 0$. The boundary conditions for the adjoint pressure are therefore also reversed to push any transients out of the domain. p_*^+ is therefore set to zero on $x = 0$ and on $r = R_{max}$. For the axisymmetric case with no swirl, $\partial p_*^+/\partial\tau \approx 0$ on all boundaries, implying $p_*^+ \approx 0$ on X_{max}. $\partial p_*^+/\partial n$ is set to zero on X_{max} and it has been assumed that $\partial p_*^+/\partial n \approx 0$ on R_{max} and on $x = 0$ in order to set $\nabla p_*^+ = 0$ on all of the boundaries.

On the lateral boundaries, the homogeneous-Dirichlet conditions imposed on all the direct perturbation variables ensure the boundary terms in (45) are zero. A convenient choice is to set the same conditions on the adjoint variables. At $x = X_{max}$, a homogeneous-Dirichlet condition on $\boldsymbol{m}^+$ and T^+, together with $p_*^+ \approx 0$, ensures the boundary terms in (45) are approximately zero. At $x = 0$ a homogeneous-Dirichlet condition set on $\boldsymbol{m}^+$ together with the homogeneous-Dirichlet conditions on $\boldsymbol{m}$ and T ensure all the boundary terms in (45) are zero. It is also convenient to set a homogeneous-Dirichlet condition on T^+ at $x = 0$.

Unlike at the direct outlet at X_{max}, the adjoint perturbations are large near $x = 0$ and this causes reflections at the homogeneous-Dirichlet boundary condition. If the domain is too short, the reflections set up standing waves between the inlet and outlet. This type of behaviour has also been observed in spatial mixing layers by [50], who suggested that the streamwise boundaries of the finite domain triggered the global resonances.

If a convection condition is used for $\hat{m}_i^+$ at $x = 0$ instead of a homogeneous-Dirichlet condition, the boundary term $(\hat{m}_i^{+*}\hat{p})$ is no longer precisely zero. It is, however, approximately zero because $\hat{p}$ is approximately zero at $x = 0$. There is also the question of which is physically more correct. The jet is modelled as exiting through a hole in a wall. A wall suggests that no perturbations, adjoint or direct, should pass out through the boundary at the jet exit plane, implying that a homogeneous-Dirichlet condition for the adjoint momentum is appropriate. Inside the jet, however, perturbations could travel upstream of the jet exit plane, implying that a convection condition for the adjoint momentum could be more appropriate. In practice we found that the homogeneous-Dirichlet condition produces continuous-adjoint eigenvalues slightly closer to the direct eigenvalues and so this is the condition that has been used for this work.

It is difficult to compare the continuous-adjoint boundary conditions with the discrete-adjoint boundary conditions because the algorithms in which they are applied are different. In both cases, an approximation of the true adjoint boundary condition is satisfied at the end of the time step. The continuous-adjoint algorithm enforces the boundary conditions directly after the state vector is altered, whereas the discrete-adjoint algorithm applies the transpose of the direct boundary conditions at the start of the time step and then updates the state vector.

6. Numerical validation

To check the direct linear code, the boundary between absolute and convective instability is calculated as S_1 is varied and the results are compared to those of Ref. [51]. The response to an impulse is calculated in a long domain with a parallel base flow and the energy at the site of the impulse is monitored. The flow parameters of $Re = 1000$ and $R/\theta = 20$ from Ref. [51] were used with a non-dimensional domain size of 6.0×84.5 (diameter $\times$ length) with 285×7271 grid points and time step $\Delta t = 0.003$. The energy at the site of the impulse was found to decay for density ratios $S_1 \leqslant 1.61$ and grow for ratios above this. For this flow configuration, the absolute/convective instability boundary in Fig. 9 of [51] lies at $S_1 \approx 1.54$(corresponding to $S \approx 0.65$ in their figure). Given the difficulties in using a time marching code to obtain absolute/convective boundaries for parallel flow, this agreement is better than expected.

The discrete-adjoint and continuous-adjoint time-marching algorithms are checked against the direct algorithm at each timestep by calculating $\boldsymbol{q}^{\oplus}(t_1)\cdot\boldsymbol{M}^N\boldsymbol{q}(t_2) - \boldsymbol{q}(t_2)\cdot(\boldsymbol{M}^H)^N\boldsymbol{q}^{\oplus}(t_1)$, which should be zero. For the discrete-adjoint code, this quantity is zero to machine precision, demonstrating that the discrete-adjoint algorithm is the exact adjoint of the direct governing equations and numerical scheme. For the continuous-adjoint code, this quantity is not quite zero because the numerical scheme of the continuous-adjoint algorithm has different truncation errors to that of the direct algorithm.

The global modes of all algorithms are then calculated with ARPACK for the base flow described in the next section. The discrete-adjoint and continuous-adjoint eigenvalues are compared against the direct eigenvalues. The discrete-adjoint eigenvalues match the direct eigenvalues to six decimal places, which is the precision specified for the eigenvalue calculation in ARPACK. The continuous-adjoint eigenvalues do not quite match the direct eigenvalues because of the different numerical schemes. Table 3 compares the direct and adjoint eigenvalues of the most unstable eigenmode as the temporal resolution is increased. The results in Tables 3 and 4 were calculated using the mid-res setup from Section 7. The direct and adjoint eigenvalues should be complex conjugates of each other. The real component is the growth rate and the imaginary component is $2\pi St$, where $St \equiv \tilde{f}\tilde{d}_j/\tilde{u}_j$ where $\tilde{f}$ is the frequency of the global mode in Hertz. This base flow is only just globally unstable, which is why the changes in the growth rate seem large when measured from zero. The absolute and relative discrepancies between the direct and adjoint eigenvalues are more informative, however, and the final column shows that these are $\sim$1% and reduce as the temporal resolution increases.

We also checked that the direct and adjoint modes are bi-orthogonal. The discrete-adjoint modes are bi-orthogonal to the direct modes to the precision specified in ARPACK, as expected. We evaluated the value of the inner product (7) between the

Table 3
Eigenvalues of the direct, σ_d, and continuous-adjoint, σ_{ca}, calculations, which would be complex conjugates of each other for a perfect scheme. The absolute discrepancy, $\mathrm{abs}(\sigma_d - \sigma_{ca}^*)$ and relative discrepancy $\mathrm{abs}(\sigma_d - \sigma_{ca}^*)/\mathrm{abs}(\sigma_d)$ decrease as the timestep decreases.

δt	Eigenvalue		Discrepancy	
	Direct	Continuous-adjoint	Absolute	Relative
0.0100	0.0450 − 1.0946i	0.0632 + 1.0950i	0.0183	0.0167
0.0050	0.0496 − 1.0946i	0.0650 + 1.0957i	0.0155	0.0141
0.0025	0.0519 − 1.0947i	0.0651 + 1.0960i	0.0132	0.0121

Table 4
The inner product of the first adjoint eigenfunction with the first, second, third and fourth direct eigenfunctions. This shows that the eigenfunctions are nearly bi-orthogonal and become more bi-orthogonal as the timestep decreases.

δt	Direct mode 1	Direct mode 2	Direct mode 3	Direct mode 4
0.0100	1	2.85×10^{-3}	4.44×10^{-5}	1.61×10^{-6}
0.0050	1	2.36×10^{-3}	4.20×10^{-5}	1.29×10^{-6}
0.0025	1	2.21×10^{-3}	3.82×10^{-5}	1.10×10^{-6}

Table 5
Eigenvalues of the direct, continuous-adjoint and discrete adjoint calculations at three different spatial resolutions.

Case	Direct	Continuous-adjoint	Discrete-adjoint
Low-res	0.0613 − 1.0646i	0.0769 + 1.0645i	0.0613 + 1.0647i
Mid-res	0.0450 − 1.0946i	0.0632 + 1.0950i	0.0450 + 1.0946i
High-res	0.0423 − 1.1044i	0.0605 + 1.1056i	0.0423 + 1.1044i

direct and continuous-adjoint eigenmodes. These are shown in Table 4 for the first adjoint mode and the first four direct modes. The modes are nearly bi-orthogonal and become more so as the resolution increases.

The discrepancies given in Tables 3 and 4 and also in Table 5 in the next section do not reduce to zero, because there are sources of error other than the one that is being investigated. For example a coarse grid will give a large discrepancy even with a small time step. Another source of error is the assumption that $Re \gg 1$, which is used to set some boundary terms to zero.

We also compared the eigenvalues of the direct global mode to the non-reacting nonlinear global mode frequencies and local analysis in the numerical study of Ref. [32]. The non-reacting frequency calculated in this study is within 4% of the frequencies predicted by the local analysis of Ref. [32]. These values are approximately 20% lower than the frequency of the nonlinear global mode, but Ref. [32] points out that this in agreement with the findings of Ref. [52].

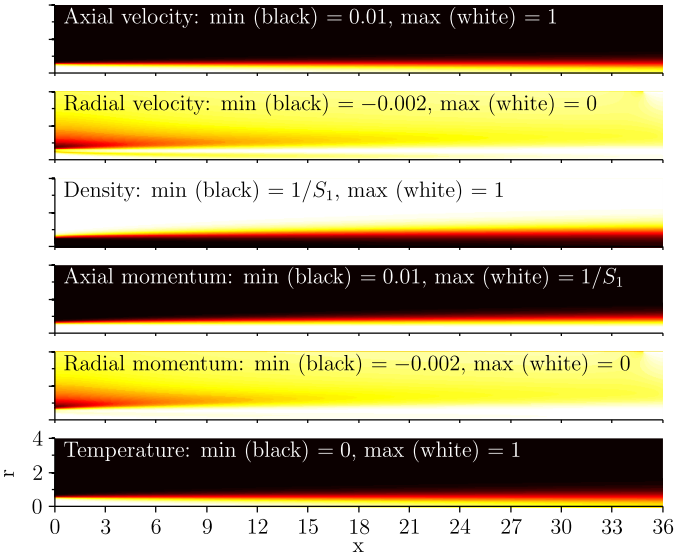

Fig. 1. Non-dimensional base flow at high resolution.

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

Figs. 3 and 4, and Figs. 5 and 6, which will be described in the next section, show that the continuous-adjoint and discrete-adjoint eigenfunctions agree well. From all of these tests, we conclude that the discrete and continuous adjoint algorithms are working correctly.

7. Results

All the results in this paper are at Reynolds number 1000 (see Table 1). This is large enough to satisfy $Re \gg 1$ but small enough to allow sufficient spreading of the shear layer over the length of the domain so that the local profiles are stable at exit. The Prandtl number is 0.7. The inlet profile has a shear layer thickness parameter $\tilde{d}_j/(2\theta) = 12.5$ and coflow $u_c = 0.01$. The Richardson number is zero and the density ratio $S_1 = 7.0$. The non-dimensional axial domain length is 36.0 and the non-dimensional domain diameter is 8.0. This setup represents a confinement ratio (domain diameter to jet diameter) of 8, which is large enough for the effects of confinement to be negligible [53].

Three grid sizes are used: 127×1027 (low-res), 181×1449 (mid-res) and 255×2049 (high-res), corresponding to radial $\times$ axial grid points ($N_r \times N_x$). The corresponding time steps are $\Delta t = 0.00707$ for the high-res simulations and $\Delta t = 0.01$ for the others. The grid is uniform in the axial direction and skewed in the radial direction to give higher resolution near the shear layers and boundaries.

The base flow used in this paper, Fig. 1, is a hot jet discharging into a cold reservoir of identical fluid. It is steady but unstable. The top three frames show the axial velocity, radial velocity and density. The bottom three frames show the axial momentum, radial momentum and temperature, which are the quantities used in the calculations.

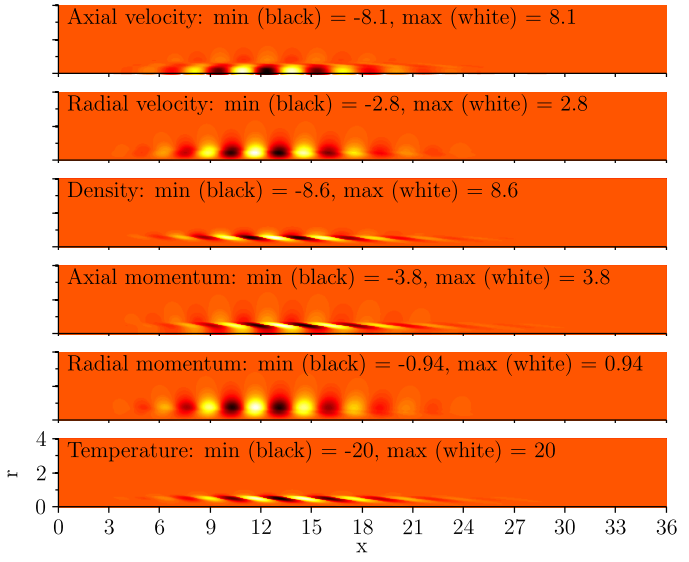

Fig. 2. Real part of the non-dimensional most-unstable direct global mode at high resolution.

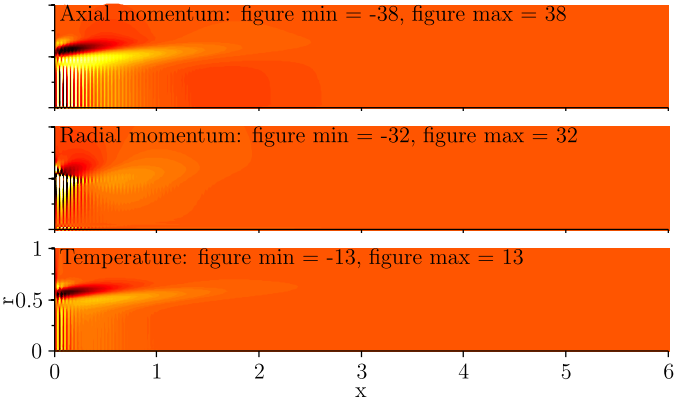

Fig. 3. Real part of the non-dimensional most-unstable discrete-adjoint global mode at high resolution (volume distribution corrected). Black regions correspond to minimum values, white regions correspond to maximum values.

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

The eigenvalues for the direct, continuous adjoint and discrete adjoint global modes at all three resolutions are given in Table 5. The difference between the eigenvalues for the mid-res and high-res cases is smaller than the difference between the low-res and mid-res case, which shows that the eigenvalues are converging with increasing spatial resolution.

The real part of the most-unstable direct, discrete-adjoint, and continuous-adjoint global modes at high-resolution are shown in Figs. 2–4 respectively. In this paper, if the minimum/maximum values of the figures are stated as 'min'/'max', they also refer to the minimum/maximum values of the data, whereas if they are stated as 'figure min'/'figure max', they refer to the figure values only and differ from the actual minimum/maximum values of the data. In the modal figures, the minimum/maximum values have been set using the value with larger magnitude so that the figure is centred on zero. Only a portion of the domain close to the inlet is shown in Figs. 3 and 4 because the low-amplitude downstream structure of the adjoint modes does not show up with the contours used in these figures. The modes have been normalized so that $\langle \hat{\boldsymbol{q}}, \hat{\boldsymbol{q}} \rangle = \langle \hat{\boldsymbol{q}}^{+}, \hat{\boldsymbol{q}}^{+} \rangle = \langle \hat{\boldsymbol{q}}^{\oplus}, \hat{\boldsymbol{q}}^{\oplus} \rangle = 1$.

The direct global modes for axial and radial velocity have the cat's eye pattern that is characteristic of the Kelvin–Helmholtz instability [3] (Section 22) [54] but with additional fluctuations in density. The global mode has maximum amplitude at $x \approx 12$ and an analysis similar to that of [11] suggests that it is caused by a wavemaker region located near the upstream edge of the global mode.

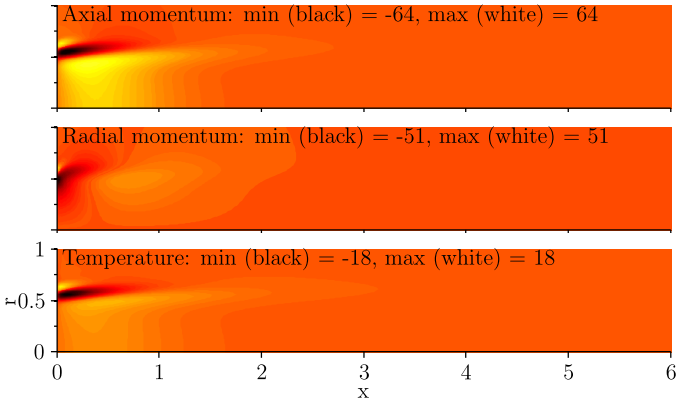

Fig. 4. Real part of the non-dimensional most-unstable continuous-adjoint global mode at high resolution.

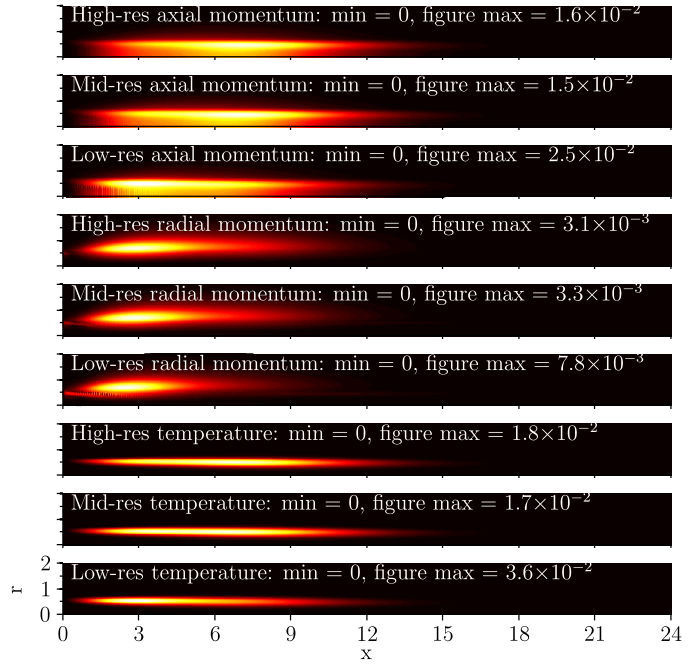

Fig. 5. Absolute value of the non-dimensional discrete-adjoint global mode multiplied by the direct global mode at every grid point (volume distribution corrected). Black regions correspond to minimum values, white regions correspond to maximum values.

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

Figs. 3 and 4 show that the adjoint mode is concentrated around the shear layer at the inlet to the domain, as is expected. The discrete-adjoint mode (Fig. 3) has a similar structure to the continuous-adjoint mode (4) but is obscured by a numerical artifact close to the inlet. We have seen this nonphysical behaviour near the boundaries in other discrete-adjoint systems and it shows the sensitivity to a numerical instability rather than a physical instability. The same effect has been seen by Ref. [55] and [36] when using the discrete-adjoint approach, with Ref. [55] attributing the effect to the discrete adjoint problem being an inconsistent discretization of the continuous adjoint problem. Nevertheless, the strong similarity between the discrete-adjoint and continuous-adjoint modes away from the boundaries is encouraging because they were calculated via two different approaches.

Structural sensitivity maps are formed from various combinations of the direct and adjoint global modes [11]. Each combination corresponds to a different sensitivity. Before presenting the receptivity to spatially-localized feedback, which is the equivalent to Fig. 17 of [11], it is useful to examine a simple combination of each direct global mode with its corresponding adjoint global mode and to see how it changes with grid resolution. The absolute value of the product between the direct and adjoint global modes at every point in the domain for axial momentum, radial momentum and temperature are shown in Fig. 5 for the discrete-adjoint and Fig. 6 for the continuous-adjoint.

The discrepancy between the mid-res and high-res results for both the discrete-adjoint (Fig. 5) and continuous-adjoint (Fig. 6) is small, suggesting that the high-res results have more than enough spatial resolution. There is a slight difference, however, between the exact locations of the high amplitude region given by the continuous-adjoint and the discrete-adjoint calculations. It is likely that this difference is due to the temporal resolution because we know from Table 3 that the eigenvalues converge closer still as the temporal resolution increases.

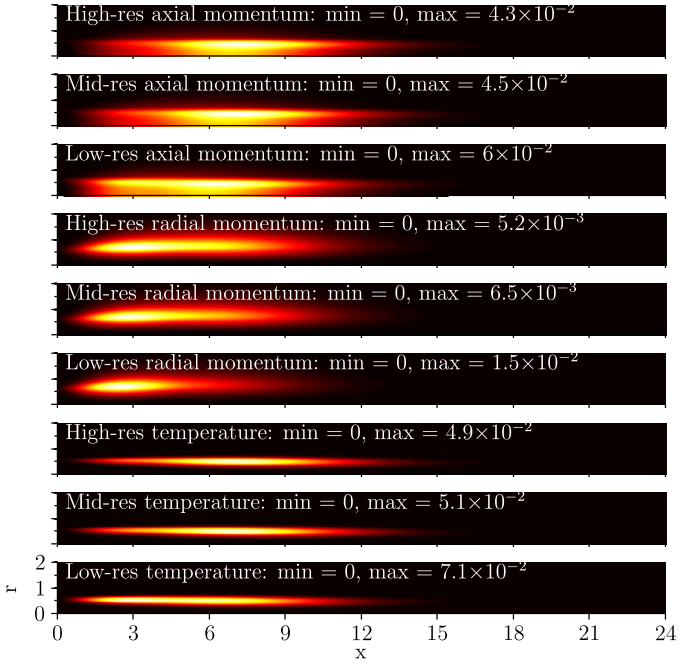

Fig. 6. Absolute value of the non-dimensional continuous-adjoint global mode multiplied by the direct global mode at every grid point. Black regions correspond to minimum values, white regions correspond to maximum values.

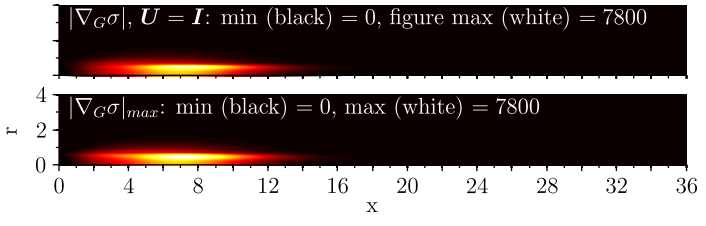

Fig. 7. Sensitivity ($\nabla_G \sigma$) of the most unstable eigenvalue (σ) to an increase in the feedback (G) between velocity and force: (a) when the force is aligned and in phase with the velocity; (b) when the feedback is in the most influential direction, as in Fig. 17 of Ref. [11].

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

Although the purpose of this paper is to describe and validate the numerical method, it is worth including a physical result. Considering a feedback mechanism in which the fluctuating velocity at a given point in the flow causes a fluctuating force at that point (such as that caused by placing a small object in the flow), Fig. 7 shows the sensitivity of the most unstable eigenvalue to the strength of that feedback. Fig. 7(a) shows the sensitivity if the feedback is in the same direction as the velocity. Fig. 7(b) shows the sensitivity if the feedback is in the direction that has the most influence. The latter is the equivalent to Fig. 17 of Ref. [11] and shows the wavemaker region of the flow, which is centred on $x \approx 7$. More results can be found in [1] and will be reported in future papers.

8. Conclusions

The first goal of this paper was to describe the calculation of adjoint global modes for variable density shear flows using a low-Mach-number formulation of the Navier–Stokes equations. Due to the treatment of density, the formation of the adjoint equations for the low-Mach-number case is more complex than for the incompressible and fully-compressible cases. By careful rearrangement of the linear equations, a set of adjoint equations has been derived that can be solved with a similar algorithm to that for the direct equations. This is known as the continuous-adjoint algorithm. A complimentary discrete-adjoint algorithm was constructed, using a matrix representation that expands and contracts the state vector. This method provides an efficient framework with which to calculate the matrix transpose for a wide range of algorithmic structures, whilst maintaining programming efficiency.

The second goal was to compare the continuous-adjoint and discrete-adjoint approaches. The most unstable global mode calculated with the discrete-adjoint approach has exactly the same eigenvalue as the corresponding direct global mode but contains numerical artifacts near the inlet. These artifacts reduced with increased resolution but did not disappear completely in the high resolution case. The most unstable global mode calculated with the continuous-adjoint approach has an eigenvalue that differs by 1.5% from that of the corresponding direct global mode but does not contain numerical artifacts. The difference reduces as the timestep reduces. Apart from the numerical artifacts, the mode shapes are very similar, which supports the expectation that they are otherwise equivalent. The continuous-adjoint requires less resolution and usually converges more quickly than the discrete-adjoint but is more challenging to implement.

Finally, for a hot jet at $Re = 1000$, the sensitivity to spatially-localized feedback was calculated. This shows the wavemaker region of the flow, which is the region that has most influence on the eigenvalue of the first global mode. The advantage of using the low Mach number equations and their adjoints is that they are appropriate for flows with variable composition and density but do not require resolution of acoustic waves. More results will be reported in future papers, including the sensitivity analysis of lifted jet diffusion flames.

Acknowledgements

The authors thank Colm Caulfield and Flavio Giannetti for helpful discussions during development of the code. This work was supported by EPSRC and Rolls Royce under grant CASE/CNA/04/80 and was performed using the Darwin Supercomputer of the University of Cambridge High Performance Computing Service (http://www.hpc.cam.ac.uk/), provided by Dell Inc. using Strategic Research Infrastructure Funding from the Higher Education Funding Council for England.

References

[1] G.J. Chandler, Sensitivity analysis of low density jets and flames, Ph.D thesis, University of Cambridge, 2010.
[2] P.J. Schmid, Nonmodal stability theory, Ann. Rev. Fluid Mech. 39 (2007) 129–162.
[3] P. Drazin, W. Reid, Hydrodynamic Stability, Cambridge University Press, 1981.
[4] D.C. Hill, A theoretical approach for analyzing the re-stabilization of wakes, AIAA Paper (1992). 92-0067.
[5] S. Bagheri, P. Schlatter, P. Schmid, D. Henningson, Global stability of a jet in crossflow, J. Fluid Mech. 642 (2009) 33–44.
[6] P.J. Schmid, D.S. Henningson, Stability and transition of shear flows, Springer, New York, 2001.
[7] B. Protas, T.R. Bewley, G. Hagen, A computational framework for the regularization of adjoint analysis in multiscale PDE systems, J. Comput. Phys. 195 (2004) 49–89.
[8] J.-M. Chomaz, Global instabilities in spatially developing flows: non-normality and nonlinearity, Ann. Rev. Fluid Mech. 37 (2005) 357–392.
[9] H. Salwen, C.E. Grosch, The continuous spectrum of the Orr–Sommerfeld equation. Part 2. Eigenfunction expansions, J. Fluid Mech. 104 (1981) 445–465.
[10] D.C. Hill, Adjoint systems and their role in the receptivity problem for boundary layers, J. Fluid Mech. 292 (1995) 183–204.
[11] F. Giannetti, P. Luchini, Structural sensitivity of the first instability of the cylinder wake, J. Fluid Mech. 581 (2007) 167–197.
[12] O. Pironneau, On optimum design in fluid mechanics, J. Fluid Mech. 64 (1974) 97–110.
[13] A. Jameson, Aerodynamic design via control theory, J. Sci. Comput. 3 (3) (1988) 233–260.
[14] A. Jameson, Re-engineering the design process through computation, J. Aircraft 36 (1) (1999) 36–50.
[15] J.C. Newman, A.C. Taylor, R.W. Barnwell, P.A. Newman, G.J.-W. Hou, Overview of sensitivity analysis and shape optimization for complex aerodynamic configurations, J. Aircraft 36 (1) (1999) 87–96.
[16] T.R. Bewley, Flow control: new challenges for a new renaissance, Prog. Aerosp. Sci. 37 (1) (2001) 21–58.
[17] A. Guegan, P. Schmid, P. Huerre, Optimal energy growth and optimal control in swept Hiemenz flow, J. Fluid Mech. 566 (2006) 11–45.
[18] P. Andersson, M. Berggren, D.S. Henningson, Optimal disturbances and bypass transition in boundary layers, Phys. Fluids 11 (1) (1999) 134–150.
[19] P. Luchini, Reynolds-number-independent instability of the boundary layer over a flat surface: optimal perturbations, J. Fluid Mech. 404 (2000) 289–309.
[20] P. Corbett, A. Bottaro, Optimal linear growth in swept boundary layers, J. Fluid Mech. 435 (2001) 1–23.
[21] M.P. Juniper, Triggering in the horizontal Rijke tube: non-normality, transient growth and bypass transition, J. Fluid Mech. 667 (2011) 272–308.
[22] P.A. Monkewitz, K.D. Sohn, Absolute instability in hot jets, AIAA J. 26 (8) (1988) 911–916.

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013

[23] A. Lingens, K. Neemann, J. Meyer, M. Schreiber, Instability of diffusion flames, Proc. Combust. Inst. 26 (1996) 1053–1061.
[24] L.I. Cerviño, T.R. Bewley, J.B. Freund, S.K. Lele, Perturbation and adjoint analyses of flow-acoustic interactions in an unsteady 2D jet, CTR Proc. Summ. Prog. (2002) 27–39.
[25] P. Meliga, A theoretical approach for the onset and control of unsteadiness in compressible afterbody flows, Ph.D. thesis, École Polytechnique, 2008.
[26] P.A. McMurtry, W.-H. Jou, J.J. Riley, R.W. Metcalfe, Direct numerical simulations of mixing layers with heat release, AIAA J. 24 (1986) 962.
[27] P.A. McMurtry, Direct numerical simulations of a reacting mixing layer with chemical heat release, Ph.D. thesis, University of Washington, 1987.
[28] J.W. Nichols, Simulation and stability analysis of jet diffusion flames, Ph.D. thesis, University of Washington, 2005.
[29] J.W. Nichols, P.J. Schmid, J.J. Riley, Self-sustained oscillations in variable-density round jets, J. Fluid Mech. 582 (2007) 341–376.
[30] H.N. Najm, P.S. Wyckoff, O.M. Knio, A semi-implicit numerical scheme for reacting flow, J. Comput. Phys. 143 (1998) 381–402.
[31] A.W. Cook, J.J. Riley, Direct numerical simulation of a turbulent reactive plume on a parallel computer, J. Comput. Phys. 129 (1996) 263–283.
[32] J.W. Nichols, P.J. Schmid, The effect of a lifted flame on the stability of round fuel jets, J. Fluid Mech. 609 (2008) 275–284.
[33] J.W. Nichols, J.-M. Chomaz, P.J. Schmid, Twisted absolute instability in lifted flames, Phys. Fluids 21 (2) (2009). 015110–8.
[34] V.G. Asouti, A.S. Zymaris, D.I. Papadimitriou, K.C. Giannakoglou, Continuous and discrete adjoint approaches for aerodynamic shape optimization with low Mach number preconditioning, Int. J. Numer. Methods Fluids 57 (2008) 1485–1504.
[35] C.R. Vogel, J.G. Wade, Analysis of costate discretizations in parameter estimation for linear evolution equations, SIAM J. Contr. Optim. 33 (1) (1995) 227–254.
[36] Z. Sirkes, T. Tziperman, Finite difference of adjoint or adjoint of finite difference?, Mon Weath. Rev. 125 (7) (1997) 3373–3378.
[37] M.B. Giles, M.C. Duta, J.-D. Müller, N.A. Pierce, Algorithm developments for discrete adjoint methods, AIAA J. 41 (2) (2003) 198–205.
[38] E. Åkervik, L. Brandt, D.S. Henningson, J. Hœpffner, O. Marxen, P. Schlatter, Steady solutions of the Navier–Stokes equations by selective frequency damping, Phys. Fluids 18 (6) (2006). 068102–1–068102–4.
[39] S.K. Lele, Compact finite difference schemes with spectral-like resolution, J. Comput. Phys. 103 (1992) 16–42.
[40] G.S. Constantinescu, S.K. Lele, A highly accurate technique for the treatment of flow equations at the polar axis in cylindrical coordinates using series expansions, J. Comput. Phys. 183 (2002) 165–186.
[41] A.J. Chorin, Numerical solution of the Navier–Stokes equations, Math. Comput. 22 (104) (1968) 745–762.
[42] J.W. Cooley, P.A.W. Lewis, P.D. Welch, The fast Fourier transform algorithm: programming considerations in the calculation of sine, cosine and Laplace transforms, J. Sound Vibr. 12 (1970) 315–337.
[43] P.N. Swarztrauber, The methods of cyclic reduction, Fourier analysis and the FACR algorithm for the discrete solution of Poisson's equation on a rectangle, SIAM Rev. 19 (3) (1977) 490–501.
[44] U. Schumann, R.A. Sweet, Fast Fourier transforms for direct solution of Poisson's equation with staggered boundary conditions, J. Comput. Phys. 75 (1988) 123–137.
[45] R.B. Lehoucq, D.C. Sorensen, C. Yang, ARPACK Users' Guide, SIAM, 1998.
[46] A. Ruhe, The two-sided Arnoldi algorithm for non-symmetric eigenvalue problems, in: B. Kågström, A. Ruhe (Eds.), Matrix pencils, Lecture Notes in Mathematics, vol. 973, Springer, Berlin, Heidelberg, 1983, pp. 104–120.
[47] R.L. Sani, J. Shen, O. Pironneau, P.M. Gresho, Pressure boundary condition for the time-dependent incompressible Navier–Stokes equations, Int. J. Numer. Methods Fluids 50 (2006) 673–682.
[48] A. Michalke, Survey on jet instability theory, Prog. Aerosp. Sci. 21 (1984) 159–199.
[49] P.M. Gresho, Some interesting issues in incompressible fluid dynamics, both in the continuum and in numerical simulation, in: J.W. Hutchinson, T.Y. Wu (Eds.), Advances in Applied Mechanics, Vol. 28, Academic Press, 1991, pp. 45–140.
[50] J.C. Buell, P. Huerre, Inflow/outflow boundary conditions and global dynamics of spatial mixing layers, CTR Proc. Summ. Prog. (1988) 19–27.
[51] L. Lesshafft, P. Huerre, Linear impulse response in hot round jets, Phys. Fluids 19 (2) (2007). 024102–1–024102–11.
[52] L. Lesshafft, P. Huerre, P. Sagaut, M. Terracol, Nonlinear global modes in hot jets, J. Fluid Mech. 554 (2006) 393–409.
[53] M.P. Juniper, The effect of confinement on the stability of non-swirling round jet/wake flows, J. Fluid Mech. 605 (2008) 227–252.
[54] O. Tammisola, F. Lundell, P. Schlatter, A. Wehrfritz, L.D. Söderberg, Global linear and nonlinear stability of viscous confined plane wakes with co-flow, J. Fluid Mech. 675 (2011) 397–434.
[55] R. Hartmann, Adjoint consistency analysis of discontinuous Galerkin discretizations, SIAM J. Num. Anal. 45 (6) (2007) 2671–2696.

Please cite this article in press as: G.J. Chandler et al., Adjoint algorithms for the Navier–Stokes equations in the low Mach number limit, J. Comput. Phys. (2011), doi:10.1016/j.jcp.2011.11.013