

A preconditioned Krylov technique for global hydrodynamic stability analysis of large-scale compressible flows

Christoph J. Mack^{a,b,*}, Peter J. Schmid^a

^a Laboratoire d'Hydrodynamique (LadHyX), CNRS-École Polytechnique, F-91128 Palaiseau, France

^b Department of Numerical Mathematics, Universität der Bundeswehr (UniBw), D-85577 Munich, Germany

ARTICLE INFO

Article history:

Received 19 February 2009

Received in revised form 15 September 2009

Accepted 20 September 2009

Available online 8 October 2009

Keywords:

Hydrodynamic stability

Compressible flow

Krylov methods

Cayley transformation

Jacobian-free framework

Direct numerical simulation

ABSTRACT

The combination of iterative Krylov-based eigenvalue algorithms and direct numerical simulations (DNS) has proven itself an effective and robust tool in solving complex global stability problems of compressible flows. A Cayley transformation is required to add flexibility to our stability solver and to allow access to specific parts of the full global spectrum which would be out of reach without such a transformation. In order to robustify the overall global stability solver an efficient ILU-based preconditioner has been implemented. With this Cayley-transformed DNS-based Krylov method two flow cases were successfully investigated: (i) a compressible mixing layer, a rather simple but well-known problem, which served as a test case and (ii) a supersonic flow about a swept parabolic body, a challenging large-scale flow configuration.

© 2009 Elsevier Inc. All rights reserved.

1. Introduction

Linear hydrodynamic stability analysis plays a central role in identifying the dynamic behavior of infinitesimal perturbations superimposed on a steady base flow. It is a crucial component for understanding the underlying mechanisms in a large variety of fluid-dynamical applications. A sound understanding of the prevailing instability mechanisms for general shear layers is, in turn, required to optimize and manipulate the inherent flow properties.

A classical tool to study the temporal instability of such flows is given by *local* stability theory, which in general relies on the existence of two homogeneous and one inhomogeneous spatial coordinate direction. This approach dates back nearly a hundred years and leads to an eigenvalue problems of moderate size which can be solved by standard direct techniques. The assumption of two homogeneous directions, however, restricts local stability theory to flows with simple geometries and simple flow physics. More complex and technologically relevant flow situations with several inhomogeneous directions and/or complicated flow physics such as supersonic flow about blunt bodies, are beyond its reach. Instead, this type of flow situations requires a *global* rather than a local approach.

Over the past decades direct numerical simulations (DNS) based on high-order spatial discretization schemes have established themselves as a widely used tool to study complex flow problems. They aim at capturing all relevant physical features of the flow by spatially resolving all dynamic scales; thus, no modeling efforts are required. The range of applications of state-of-the-art direct numerical simulations is truly impressive in scope and complexity, and it is the aim of global hydrodynamic stability analysis – and the objective of this article – to harness these strengths and capabilities.

* Corresponding author. Address: Laboratoire d'Hydrodynamique (LadHyX), CNRS-École Polytechnique, F-91128 Palaiseau, France. Tel.: +33 1 69 33 52 85; fax: +33 1 69 33 52 92.

E-mail address: cmack@ladhyx.polytechnique.fr (C.J. Mack).

Even though the global stability problem can be mathematically formulated, the lack of multiple homogeneous directions yields a linear stability matrix whose sheer size makes a direct solution prohibitively expensive [46]. Limited computational resources call for iterative eigenvalue methods (see, e.g., [11] for applications of iterative techniques in fluid mechanics) that extract stability information from the linearized flow in an approximate manner. Among these iterative solution techniques Krylov subspace methods [44] are particularly popular for fluid-dynamical applications. Even though an explicit formulation and storage of the stability matrix is a feasible alternative when combined with parallelization efforts [18], we will focus on a Jacobian- or matrix-free environment where the necessary information for the global stability problem is directly extracted from direct numerical simulations.

Compressible global stability problems have only recently been tackled (see, e.g., [9,10] for a study of the onset of transonic shock buffeting). Large-scale complex flow problems featuring multiple (temporal/spatial) scales and multi-physics (shear and compressibility effects, acoustic waves, etc.) exhibit a complicated spectrum which requires special means to extract the relevant flow behavior. In addition, an erratic convergence history of standard iterative techniques is observed which calls for additional physics-based measures, such as spectral transformations and preconditioning, to improve their convergence towards specific global modes. Spectral transformations deform the complex eigenvalue plane in order to make specific parts of the global spectrum accessible to iterative eigenvalue methods (an approach of this type has been proposed by [32]). However, these transformations come at the expense of solving a large-scale linear system which, in accordance with the above procedures, has to be accomplished using a preconditioned iterative method based on a matrix-free approach.

Krylov subspace methods for hydrodynamic stability analysis of the *incompressible* Navier–Stokes equations were first introduced by [11]. Their investigation of moderately complex flow situations allowed a simpler Jacobian-based implementation and did not require any type of spectral transformation. In the following years similar techniques have been compiled into the open-source package ARPACK [24]. A comparative study of transformed Krylov subspace techniques applied to problems from computational fluid dynamics can be found in, e.g., [32,50,52], and further large-scale stability calculations have been performed by [22,8]. The above studies, however, take advantage of the explicit presence of the linear stability matrix. More recently, Arbenz et al. [2] compared eigensolvers for large-scale three-dimensional modal analysis using AMG-preconditioned iterative techniques; all matrices are semi-positive definite and are, again, available explicitly. A state-of-the-art review of Krylov subspace techniques applied to a wide range of fluid flows of aerodynamical interest is given by Theofilis [46] with special emphasis on the global linear stability of non-parallel and three-dimensional flow configurations. Based on his encompassing study we conclude that further progress in the field of global stability analysis has to involve a purely iterative and matrix-free approach which provides the starting point of this article.

Preconditioning techniques (see, e.g., [19,33] for a recent overview) help improve the convergence of the linear system solver. In a survey article, Benzi [5] provides a detailed overview of recent preconditioning strategies for large linear systems mainly focusing on incomplete factorization techniques (ILU) and sparse approximate inverses (SPAI). A general discussion on algebraic multigrid (AMG) as a preconditioner can be found in [49], and a specific application to AMG-accelerated BiCG-Stab is given in [14]. While these references concentrate on preconditioning strategies for linear systems arising in Newton–Krylov methods, the literature on transformation and preconditioning techniques for large-scale eigenvalue problems is relatively sparse.

In this article we combine preconditioned Krylov-based techniques and direct numerical simulations (DNS) to obtain a robust DNS-based Jacobian-free global stability solver for compressible flows. In this manner, our contribution represents, on one hand, a generalization of the work of [11] towards compressible flows and, on the other hand, provides an extension of current tools for the global stability analysis of non-parallel and three-dimensional flows, as alluded to by [46]. It builds on previous global stability studies of compressible flow such as the work of Theofilis and colonius [47] who investigate the behavior of cavity flows for a wide range of Mach numbers, the study of Robinet [36] who treats the interaction of a shock with a laminar boundary layer within a global stability framework and the investigations of Crouch et al. [9,10] who treat the onset of transonic shock buffeting as a global stability problem. Even though these studies are based on compressible governing equations they do not follow the approach taken in this article where an iterative stability solver is coupled to a direct numerical simulation code, thus performing a DNS-based (matrix-free) global stability analysis.

In a historical and methodological context the linear stability of compressible flows can be studied by the following approaches: (a) the direct solution of a one-dimensional local stability problem, (b) the long-time integration of a direct numerical simulation (DNS) starting with small-amplitude perturbations, (c) the direct or iterative solution of the global linear stability matrix, and (d) the preconditioned iterative solution of a global linear stability problem. In this article we develop the methodology for the latter approach (d) based on direct numerical simulations. In Section 2 we present our test case, a compressible mixing layer, formulate the governing equations, briefly describe the linear stability theory for approaches (a), (c) and (d) and display our reference spectrum. This is followed by the description of the DNS-based global stability solver for approach (c) (Section 3) and our Cayley-transformed version (Section 4). We conclude with results in Section 5.

2. Compressible mixing layer: reference spectrum and DNS

It will be useful and instructive to demonstrate the global stability analysis of complex flows with the help of a somewhat simplified flow configuration that possesses all the relevant physical and numerical features of the full problem while still

allowing a comparison with reference solutions or analytical results. This way, we can design and assess the important components of the DNS-based global stability solver and gain valuable insight and experience in using the governing parameters to influence convergence properties and solution quality.

In view of the desired fields of application we consider a fully compressible flow that is dominated by advective, diffusive and wave-propagation phenomena. These features will be reflected in the associated global spectrum in form of shear modes, acoustic modes and combinations thereof which displays the typical features of many high-speed flows arising in aeronautical applications. The chosen reference flow configuration, however, should be sufficiently simple to still allow a solution by direct means; this point will particularly aid in the quality assessment of each approximate step taken in the overall global stability algorithm.

A flow configuration that satisfies most, if not all, of the above described prerequisites is the compressible mixing layer. A hyperbolic-tangent base velocity profile is assumed, and the flow field is subsequently linearized about this analytic base state resulting in the linearized compressible Navier–Stokes equations governing the perturbation field. The stability of this flow has been studied extensively, and the two landmark papers on this subject have been published by [30,7].

2.1. Governing equations

The dynamics of the mixing layer is governed by the compressible Navier–Stokes equations which have been formulated, using Cartesian tensor notation, for the pressure p , the velocities (u, v, w) and the entropy s . We define a viscous length scale δ (shear-layer thickness), a Reynolds number Re , a Mach number Ma and a Prandtl number Pr as

$$\delta = \frac{2u_\infty}{\omega_\delta}, \quad Re = \frac{u_\infty \delta}{\nu}, \quad Ma = \frac{u_\infty}{c_\infty}, \quad Pr = \frac{C_p \mu}{k}, \quad (1)$$

where $\omega_\delta = (du/dy)_{y=0}$ and the subscript ∞ stands for the freestream quantities. The remaining variables ν, μ, c_∞, C_p and k denote, respectively, the kinematic and dynamic viscosity, the speed of sound, the specific heat at constant pressure and the thermal conductivity. Using these expressions, the compressible Navier–Stokes equations can be rewritten in the following non-dimensionalized form:

$$\frac{\partial p}{\partial t} + u_j \frac{\partial p}{\partial x_j} = -\gamma^2 p \frac{\partial u_j}{\partial x_j} + (\gamma - 1) \left(\frac{\gamma Ma^2}{Re} \Phi + \frac{\gamma}{(\gamma - 1) Re Pr} \frac{\partial}{\partial x_j} \left(k \frac{\partial T}{\partial x_j} \right) \right), \quad (2a)$$

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} = -\frac{1}{\gamma Ma^2 \rho} \frac{\partial p}{\partial x_i} + \frac{1}{Re \rho} \frac{\partial}{\partial x_j} \left(\mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) - \frac{2}{3} \mu \frac{\partial u_k}{\partial x_k} \delta_{ij} \right), \quad (2b)$$

$$\frac{\partial s}{\partial t} + u_j \frac{\partial s}{\partial x_j} = \frac{\gamma Ma^2}{Re} \Phi + \frac{\gamma}{(\gamma - 1) Re Pr} \frac{\partial}{\partial x_j} \left(k \frac{\partial T}{\partial x_j} \right) \quad (2c)$$

with

$$\Phi \equiv \frac{1}{2} \mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right)^2 - \frac{2}{3} \mu \left(\frac{\partial u_k}{\partial x_k} \right)^2.$$

Herein, a freestream entropy $s_\infty = R_c$ was assumed, where R_c denotes the gas constant. The variables γ, ρ and T represent the specific heat ratio, the density and the temperature, respectively.

The non-dimensionalized base flow in the streamwise x -direction is given by

$$u_0(x, y, z) = \tanh(y) \quad -\infty < y < +\infty, \quad (3)$$

where the subscript 0 stands for base flow quantities; y and z represent the normal and the spanwise coordinate direction, respectively (see Fig. 1). The base pressure $p_0(x, y, z) = 1$ and the base entropy $s_0(x, y, z) = 1$ are assumed to be constant.

The compressible Navier–Stokes equations, the equation of state, Fourier's law for the thermal conductivity and Sutherland's law, where ambient conditions are used as a reference state, for the viscosity fully describe the flow. For all simulations shown in this article, we consider the motion of a compressible fluid modeled as a perfect gas with constant specific heat ratio $\gamma = 1.4$ and constant Prandtl number $Pr = 0.71$.

2.2. Linear stability analysis

In a next step towards a linear stability analysis we assume a three-dimensional small-amplitude perturbation field $\epsilon \phi' = \epsilon(p', u', v', w', s')^T$ superimposed on a steady base flow ϕ_0 .

$$\phi(x, y, z, t) = \phi_0(x, y, z) + \epsilon \phi'(x, y, z, t) \quad (4)$$

From this the linearized Navier–Stokes equations for the perturbations

$$\frac{\partial \phi'}{\partial t} = \mathcal{L}(\phi_0) \phi' \quad (5)$$




Fig. 1. Sketch of a mixing layer showing the base velocity profile (in blue), some of the relevant flow parameters and the Cartesian coordinate system (in red). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

are obtained, where $\mathcal{J}(\phi_0)$ denotes the linear stability operator. The compressible mixing layer (see Fig. 1) is then assumed to be homogeneous in the x - and the z -direction, and periodic boundary conditions are applied in these directions. In the remaining inhomogeneous y -direction, all disturbance quantities are assumed to decay exponentially in the freestream.

2.2.1. Local stability approach

Referring back to approach (a), mentioned at the end of the introduction, and taking advantage of the separability of the governing equations in the two homogeneous coordinate directions, the disturbance field $\phi'(x, y, z, t)$ can be taken of the following traveling-wave form

$$\phi'(x, y, z, t) = \tilde{\phi}(y) e^{i(\alpha x + \beta z - \omega t)}, \quad (6)$$

where $\tilde{\phi}(y)$ denotes the complex amplitude, and α and β are the real wavenumbers of the perturbation in the streamwise and the spanwise direction, respectively. The temporal long-term evolution of this type of disturbance is characterized by ω whose real part describes the frequency ω_r and whose imaginary part the corresponding growth rate ω_i . Upon substitution of the above expression (6) into the perturbation equations (5) we obtain a *one-dimensional* eigenvalue problem $\omega \tilde{\phi} = \mathcal{J}(\phi_0) \tilde{\phi}$ for ω and $\tilde{\phi}$. Once the dependence on the inhomogeneous y -direction is eliminated by a sixth-order compact finite difference scheme [25] using n_y grid points, we arrive at an eigenvalue problem for the linear stability matrix, i.e., the $5n_y \times 5n_y$ Jacobian matrix, which can be solved numerically by direct means for each wavenumber pair (α, β) and for each value of the remaining parameters.

2.2.2. Global stability approach

Many flow problems are characterized by a complex geometry or complex flow physics which no longer permits an assumption of the form (6), e.g., owing to several inhomogeneous coordinate directions. Supersonic flow about a swept parabolic body, which will be considered later in this article (see Section 5.2), represents an example of this type. In this case, a *global* rather than a *local* stability approach has to be considered which forms the basis for approach (c) and (d) as mentioned at the end of the introduction.

For the global approach the disturbance field $\phi'(x, y, z, t)$ is assumed to satisfy

$$\phi'(x, y, z, t) = \tilde{\phi}(x, y) e^{i(\beta z - \omega t)}, \quad (7)$$

where, as before, $\tilde{\phi}(x, y)$ denotes the complex amplitude and β the real spanwise wavenumber of the perturbation. This formulation is referred to as *BiGlobal* in [46]. The long-term temporal stability of the perturbation is given by the global eigenvalue ω . Under these assumptions and an appropriate discretization of the x - and y -dependence, the *two-dimensional* (global) stability problem can formally be written as

$$\omega \tilde{\phi}(\mathbf{x}, \mathbf{y}) = \mathbf{J}(\phi_0) \tilde{\phi}(\mathbf{x}, \mathbf{y}), \quad (8)$$

where $\mathbf{J}(\phi_0)$ represents the $n \times n$ linear stability matrix (the Jacobian), i.e., the discretized Navier–Stokes equations linearized about the base state ϕ_0 , with $n = 5n_x n_y$ as the dimension of this (complex-valued) eigenvalue problem; n_x and n_y denote the number of grid points in the x - and y -direction, respectively.

For linear stability matrices of moderate size, e.g., a dimension $n \sim \mathcal{O}(10^5)$, this eigenvalue problem can still be solved directly (see, e.g., [46] for a discussion on the solution of global eigenvalue problems). For more complex flow configurations which require a large domain as well as a high spatial resolution, however, the direct solution of the (global) eigenvalue

problem (8) is prohibitively expensive. Instead, iterative solution techniques have to be employed to extract pertinent stability information.

2.3. Long-time integration of the initial value problem

An alternative approach to address the stability of complex flows is given by the long-time solution of a linearized initial value problem

$$\frac{\partial \phi}{\partial t} = \mathcal{F}(\underbrace{\phi_0 + \epsilon \phi^i}_{\phi}), \quad (9)$$

where \mathcal{F} represents the right-hand side of the nonlinear Navier–Stokes equations. The solution of (9) could be obtained from direct numerical simulations (see, e.g., [46,20] for applications in fluid dynamics). This technique corresponds to approach (b) as mentioned at the end of the introduction. Starting with an arbitrary initial condition of small amplitude $\epsilon \phi^i$ (such that nonlinear effects can be neglected) superimposed on a steady base flow ϕ_0 , the solution of the initial value problem converges towards the least-stable global mode as time progresses. The corresponding global eigenvalue can be computed from this global mode in a straightforward manner.

2.4. Reference spectrum

A typical spectrum of the compressible mixing layer is displayed in Fig. 2(a), and the values of the four eigenvalues depicted in circles are shown in Table 1. This multi-branch spectrum was obtained using the local stability approach (see Section 2.2.1), and it reveals characteristic features such as (i) acoustic modes with very small damping but rather high phase

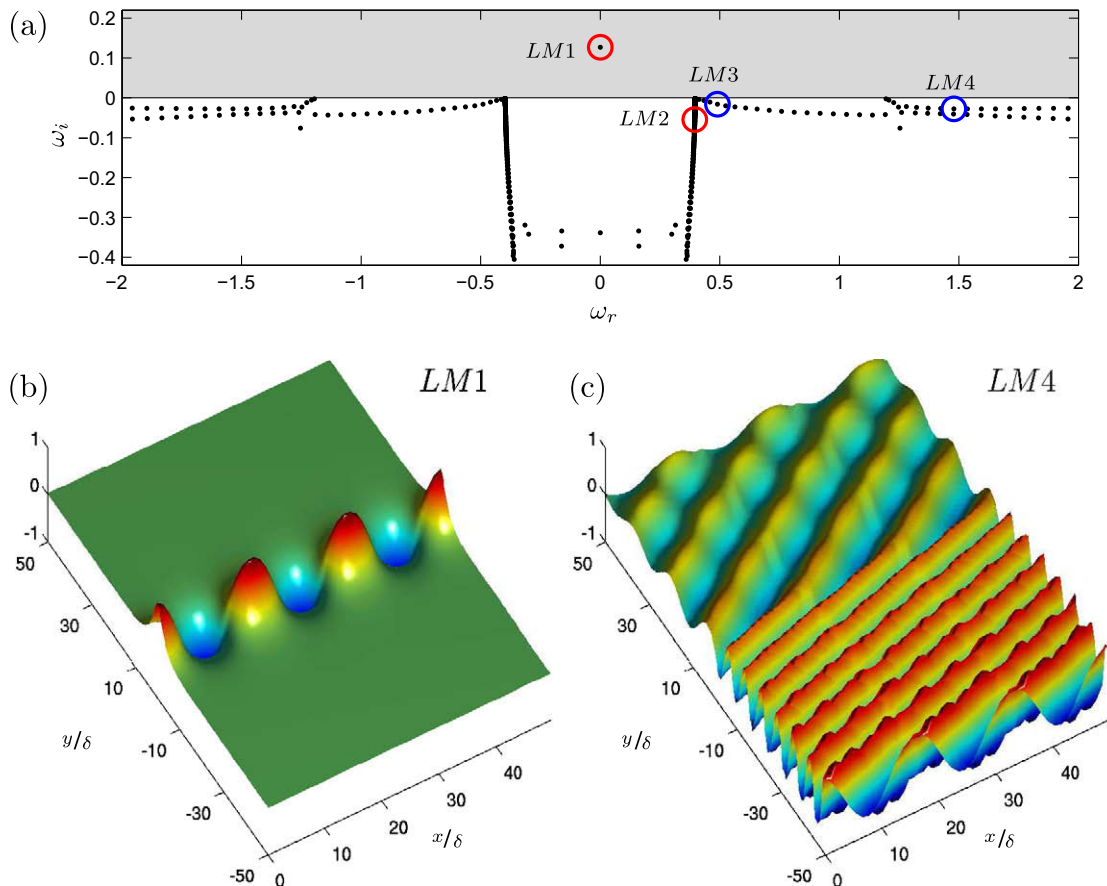


Fig. 2. (a) Directly computed reference spectrum of a compressible mixing layer for $Re = 1000, Ma = 0.5, \alpha = 0.397$ and $\beta = 0$ (see Table 2, ConfigIV); $n_y = 201$ grid points were used to resolve the inhomogeneous y -direction; unstable half-plane in grey. The spatial structure of two (local) modes is visualized by the (normalized) disturbance pressure using three periods in the periodic x -direction: (b) unstable (discrete) shear mode (LM1), (c) faster-moving acoustic mode (LM4) from the acoustic branch.

Table 1

Values of the four eigenvalues $\omega = \omega_r + i\omega_i$ marked by circles in Fig. 2(a); they belong to the unstable shear mode (LM1), a representative weakly-damped mode from the continuous shear-layer branch (LM2) and a slow-moving (LM3) as well as a faster-moving mode (LM4) from the acoustic branch.

Mode	LM1	LM2	LM3	LM4
ω_r	0	0.395973	0.459261	1.478955
ω_i	0.127146	-0.040255	-0.010807	-0.027631

velocities (which reflects the fact that sound waves propagate quickly and without significant attenuation in their amplitude), (ii) shear-layer modes displaying small phase velocities (which reflects the fact that shear instabilities propagate with the local base velocity which is rather small inside the shear layer), and (iii) a continuous spectrum of modes that describes the perturbation dynamics in the freestream and ensures the completeness of the eigenfunction basis [40]. Instabilities (with $\omega_i > 0$) are observed only on the shear-layer branch. The spatial structure of two representative eigenfunctions, one from the shear-layer branch (LM1) and one from the acoustic branch (LM4) are shown in Fig. 2(b) and (c).

This spectrum which displays many features of more complex spectra for compressible flows will serve as a reference spectrum on which we will develop, test and validate iterative solution techniques for the global hydrodynamic stability problem.

2.5. Direct numerical simulations and validation of the code

The direct numerical simulations (DNS) are performed on a non-uniformly distributed grid using n_x and n_y grid points in the x - and y -direction with a clustering of the n_y grid points near $y = 0$ to better resolve the shear layer. The governing equations are solved using a characteristic-type formulation [42] and discretized employing fifth- and sixth-order compact difference schemes for the inviscid and viscous terms, respectively [1,25]. For the initial value problem the temporal discretization is accomplished by a fourth-order Runge–Kutta scheme (see [26] for details).

In what follows, all numerical investigations will focus on four selected flow configurations (Config I–IV) that represent a range of physical and numerical features. The governing parameters for these examples are given in Table 2.

To validate the code direct solutions of the one-dimensional eigenvalue problem (see Section 2.2.1) as well as solutions of the two-dimensional initial value problem (9) via direct numerical simulations are performed. In Table 3 these results are compared with Blumen's findings for selected parameter combinations (Config I–III), and the results for the parameter choice of our reference case (Config IV) are shown too. In this manner, the Jacobian (used as a preconditioner) and the DNS-code are validated. In addition, results from the solution of the two-dimensional (global) eigenvalue problem (8) obtained from a DNS-based iterative stability solver (without transformation) are included as well. The used Krylov subspace method that generated the latter results is described in Section 3, its spectrally (Cayley-) transformed version is developed in Section 4 and evaluated using the given reference spectrum (see Fig. 2).

3. Iterative stability analysis

The global linear stability analysis of complex fluid flows leads to a non-Hermitian eigenvalue problem (8) whose solution requires iterative algorithms such as Krylov subspace methods. In many fluid-dynamical applications only a few eigenvalues

Table 2

Governing parameters of the four selected flow configurations (Config I–IV) which are defined by the shear-layer thickness δ , the Reynolds number Re , the Mach number Ma as well as the disturbance wavenumbers α and β .

Config #	δ	Re	Ma	α	β
I	0.1	$\rightarrow \infty$	0.1	0.433	0
II	0.1	$\rightarrow \infty$	0.5	0.397	0
III	0.1	$\rightarrow \infty$	0.9	0.208	0
IV	0.1	1000	0.5	0.397	0

Table 3

Growth rate ω_i of the unstable shear mode as obtained by using the approaches (a)–(c) described at the end of the introduction; a resolution of $n_x = 32$ and $n_y = 201$ grid points is used; Blumen's [7] results are also included. For Config IV the unstable shear mode could not be computed via approach (b).

Config #	[7]	approach (a) matrix-based	approach (b) DNS-based	approach (c) DNS-based
I	0.187	0.187521	0.1875	0.187521
II	0.141	0.141161	0.1412	0.141167
III	0.055	0.054731	0.0547	0.054723
IV	–	0.127146	–	0.127155

are needed to answer questions of interest, a requirement that also favors iterative techniques. In what follows, we adopt a linear algebra notation with the matrix $\mathbf{A} \equiv \mathbf{J}(\phi_0)$, the vector $\mathbf{x} \equiv \tilde{\phi}$ and the eigenvalue $\lambda \equiv \omega$, yielding a standard eigenvalue problem given as $\lambda \mathbf{x} = \mathbf{A} \mathbf{x}$.

3.1. Krylov subspace methods

A common class of iterative eigenvalue algorithms is based on the premise that the full stability problem can be projected onto a lower m -dimensional vector space given by an m -dimensional Krylov subspace sequence

$$\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1) = \text{span} \{ \mathbf{v}_1, \mathbf{A} \mathbf{v}_1, \mathbf{A}^2 \mathbf{v}_1, \dots, \mathbf{A}^{m-1} \mathbf{v}_1 \}. \tag{10}$$

This sequence consists of repeated applications of a matrix \mathbf{A} to an initial vector \mathbf{v}_1 . The spectrum of the subsequent projected system then approximates the spectrum of the full stability problem given by \mathbf{A} .

Among the general class of Krylov subspace methods, we choose the *Implicitly Restarted Arnoldi Method* (IRAM) proposed by [43]. This method is briefly described as follows (for a more complete discussion the reader is referred to the cited author): The Arnoldi method constructs an orthonormal basis $\mathbf{V}_m = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m]$ of the Krylov subspace \mathcal{K}_m which is then used to decompose a matrix \mathbf{A} in the following way:

$$\mathbf{A} \mathbf{V}_m = \mathbf{V}_m \mathbf{H}_m + \mathbf{f}_m \mathbf{e}_m^T. \tag{11}$$

\mathbf{H}_m denotes an m -dimensional upper Hessenberg matrix (with $m \ll n$), \mathbf{f}_m is the residual vector orthogonal to the basis \mathbf{V}_m , and \mathbf{e}_m represents a unit-vector in the m th component. Multiplying both sides of (11) from the left by \mathbf{V}_m^* and using the fact that \mathbf{V}_m is unitary, we obtain

$$\mathbf{V}_m^* \mathbf{A} \mathbf{V}_m = \mathbf{H}_m, \tag{12}$$

where the superscript $*$ denotes the Hermitian conjugate. The eigenvalues $\{\theta_j\}$ of the Hessenberg matrix \mathbf{H}_m , the so-called Ritz values, are approximations of the eigenvalues $\{\lambda_j\}$ of the matrix \mathbf{A} , and the associated eigenvectors $\tilde{\mathbf{x}}_j$ of \mathbf{A} , the so-called Ritz vectors, can be calculated using the orthonormal basis \mathbf{V}_m as

$$\tilde{\mathbf{x}}_j = \mathbf{V}_m \mathbf{y}_j, \tag{13}$$

where \mathbf{y}_j denotes the eigenvector of \mathbf{H}_m associated with the eigenvalue θ_j . In general, some of the Ritz pairs $(\tilde{\mathbf{x}}_j, \theta_j)$ closely approximate the eigenpairs $(\mathbf{x}_j, \lambda_j)$ of \mathbf{A} , and the quality of this approximation usually improves as the dimension m of the Krylov subspace sequence \mathcal{K}_m increases. In practice, however, the dimension of this subspace is limited by memory restrictions, and its ortho-normalization is progressively affected by numerical errors as m increases. For this reason, the Arnoldi factorization (11) needs to be periodically restarted with a new starting vector \mathbf{v}^+ .

Sorensen's implicit restarting strategy [43] computes this new starting vector \mathbf{v}^+ by a polynomial approximation of Krylov vectors that damps $p = m - k$ undesired Ritz pairs, where k denotes the number of desired Ritz pairs. Lehoucq and Scott [23] and Morgan [31] studied the issue of restarting and compared implicit restarting with other schemes; furthermore, Morgan [31] and Sorensen [44] concluded that using implicit restarting and applying exact shifts in connection with the Arnoldi method is optimal. For details, including a discussion on the convergence behavior, we refer the reader to the above-mentioned literature.

For the sake of completeness, a further class of Krylov subspace methods known as subspace iteration techniques (see, e.g., [44]) is worth mentioning. As an example, Heeg and Geurts [16] successfully applied these techniques in their studies on spatial instabilities of the incompressible attachment-line flow.

3.2. Jacobian-free framework

The form of the Krylov subspace sequence (10) indicates that the Jacobian matrix $\mathbf{A} \equiv \mathbf{J}(\phi_0)$ does not need to be formed explicitly; rather, only matrix-vector products are necessary to perform the Arnoldi decomposition. These matrix-vector products are readily obtained from direct numerical simulations (DNS) via

$$\mathbf{A} \mathbf{v}_i \approx \frac{\mathbf{F}(\phi_0 + \epsilon \mathbf{v}_i) - \mathbf{F}(\phi_0)}{\epsilon} = \left. \frac{\partial \mathbf{F}(\phi)}{\partial \phi} \right|_{\phi=\phi_0} \mathbf{v}_i + \mathcal{O}(\epsilon) \quad \text{with } i = 1, 2, \dots, m - 1, \tag{14}$$

where ϵ is a user-specified parameter, ϕ_0 and $\mathbf{v}_i \equiv \phi'$ denote, as before, the base flow and a disturbance field, respectively, and \mathbf{F} represents the discretized right-hand side of the nonlinear Navier–Stokes equations. This first-order finite difference approximation of the Jacobian matrix $\mathbf{J}(\phi_0)$ allows a Jacobian-free framework where right-hand side evaluations from direct numerical simulations provide the input for the iterative stability solver. Consequently, an explicit linearization of the governing Eqs. (2a–c) is no longer required. A matrix-free approach reduces memory requirements considerably and removes the problem of explicitly forming and storing the high-order Jacobian matrix. This advantage significantly simplifies the overall global stability method.

The choice of ϵ is, however, not obvious: if ϵ is too large, the derivative will be poorly approximated and if ϵ is too small, the result will be contaminated by roundoff errors. A widely used choice represents

$$\epsilon = \frac{\|\phi_0\|_2}{\|\mathbf{v}_i\|_2} \epsilon_0, \quad (15)$$

where ϵ_0 is a small parameter which is typically chosen somewhat larger than the square root of machine epsilon. For a discussion on common choices for ϵ as well as higher-order approximations for Eq. (14) the reader is referred to [19] and the references therein.

It should be mentioned though, that this matrix-free formulation further introduces a considerable amount of flexibility in forming the linear stability matrix. The call to the right-hand side \mathbf{F} in Eq. (14), which in our case consists of a direct numerical simulation code for the compressible Navier–Stokes equations, can easily be substituted by other numerical discretizations of the corresponding governing equations. In particular, the entire hierarchy of common models for high-Reynolds number flows in complex geometries (such as, e.g., large-eddy simulations (LES), detached-eddy simulations (DES), vortex-particle methods and even commercial codes) can be treated within the same framework. Stability results using this Jacobian-free framework can thus be obtained for any flow whose main features can be captured to a sufficient degree of accuracy by appropriate numerical simulations.

4. Convergence acceleration and control

The need for convergence acceleration arises from the fact that high resolution simulations of complex fluid flow physics (with, e.g., the coexistence of shear and acoustic modes as illustrated in our reference spectrum, see Fig. 2(a)) lead to an unpredictable and erratic convergence behavior of the simple Arnoldi method (without transformation). In addition, any thorough investigation of complex fluid flow behavior requires us to focus on specific parts of the global spectrum. For example, to investigate the acoustic near- and far-field as to its structure and directivity it is necessary to extract global modes from the acoustic branch. To direct the convergence of the Arnoldi method towards these modes a transformation of the complex eigenvalue plane can be used. In short, both convergence acceleration and convergence control are desired for an effective DNS-based global stability solver. It must be stressed, however, that any technique employed to accelerate or control the convergence of the global stability solver has to be of an iterative and matrix-free nature to preserve the applicability of the method to a wide range of complex fluid-dynamical problems.

4.1. Inexact Cayley transformation

To accelerate and control the convergence behavior of the global stability algorithm the *Cayley transformation* [15] is applied. This transformation consists of a two-parameter conformal mapping of the complex plane and, for generalized eigenvalue problems of the form $\lambda \mathbf{B} \mathbf{x} = \mathbf{A} \mathbf{x}$, is defined as

$$\mathbf{T}_C(\sigma, \mu) \equiv (\mathbf{A} - \sigma \mathbf{B})^{-1} (\mathbf{A} - \mu \mathbf{B}), \quad (16)$$

where σ and μ are the mapping parameters. For standard eigenvalue problems, as it is the case for our problem, we have $\mathbf{B} = \mathbf{I}$ where \mathbf{I} denotes the identity matrix. This transformation represents a more general mapping than the more commonly applied shift-invert technique (see, e.g., [12,46]).

The complex parameter σ acts as a shift parameter, and eigenvalues close to it are mapped far into the right-half plane for $\text{Imag}\{\lambda\} < \text{Imag}\{\sigma\} < \text{Imag}\{\mu\}$ while eigenvalues far from it are mapped close to one (see Fig. 3). The second complex parameter μ introduces an additional stretching-and-rotation effect on the transformed spectrum. Its major role, however, consists in controlling the condition number of the linear transformation. For this reason, the Cayley transformation in general yields a better-conditioned linear system than the shift-invert transformation (see [22]), an important advantage for its iterative solution. In addition, [21] report “the superior numerical performance of a Cayley transformation over that of a shift-invert transformation within an Arnoldi method when using an iterative linear solver”.

The eigenvalues λ of (\mathbf{A}, \mathbf{B}) are then recovered from the eigenvalues ξ of the transformed problem via

$$\lambda = \frac{\sigma \xi - \mu}{\xi - 1}, \quad (17)$$

while the eigenvectors \mathbf{x} are not affected by the transformation.

Fig. 3 demonstrates the Cayley transformation for our reference spectrum (see Fig. 2(a)). Two parameter settings for σ and μ are displayed: the first parameter combination focuses on the unstable shear mode *LM1* (see Table 1) whereas the second parameter combination aims at extracting a specific mode *LM4* from the acoustic branch. The first configuration illustrates the possibility of convergence acceleration of the Arnoldi method, while the second configuration demonstrates the Cayley transformation as a convergence control tool. The general mapping between the complex λ - and the complex ξ -plane is further visualized by the dashed Cartesian grid in Fig. 3(a) and its mapped counterparts in Fig. 3(b) and (c). For an overview of available acceleration techniques we refer the reader to [3].

The Cayley transformation (16) requires the solution of the following generally non-Hermitian linear system

$$(\mathbf{A} - \sigma \mathbf{I}) \mathbf{v}_{j+1} = (\mathbf{A} - \mu \mathbf{I}) \mathbf{v}_j \quad (18)$$

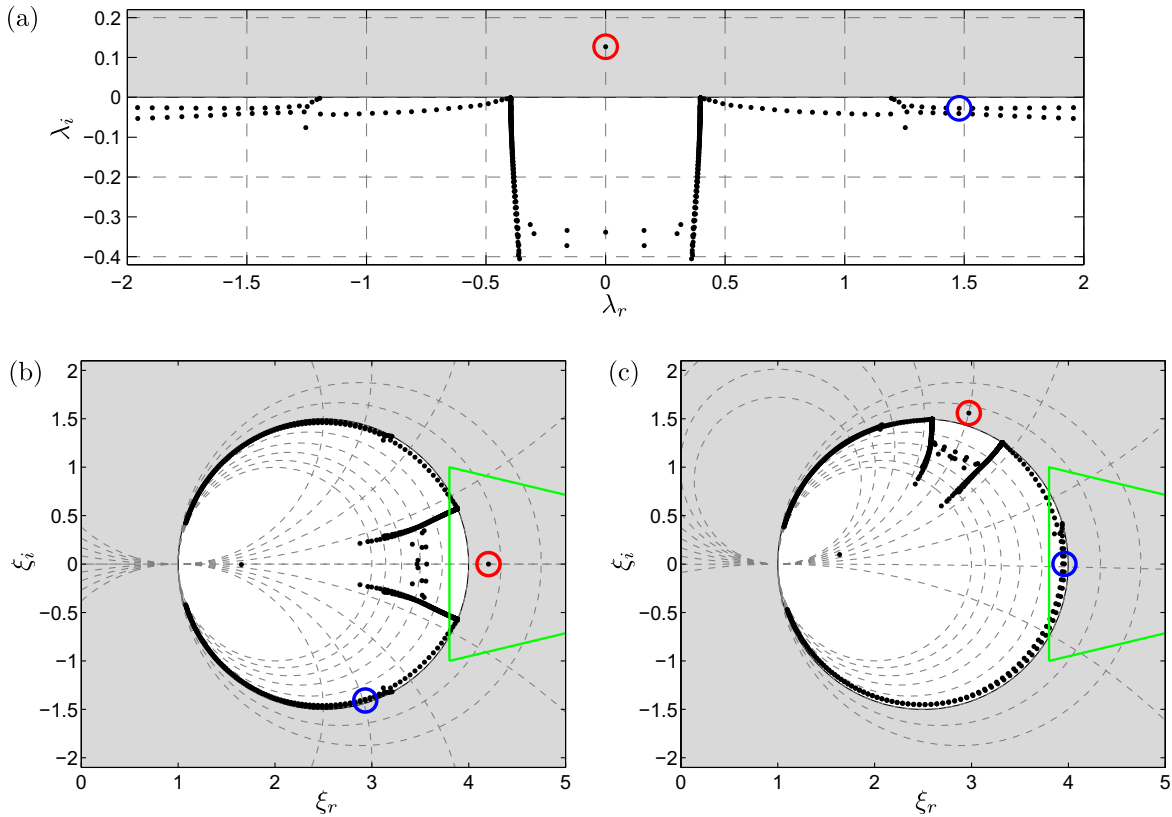


Fig. 3. (a) Relabeled reference spectrum ($\lambda \equiv \omega$) as presented in Fig. 2(a); the two eigenvalues marked by circles belong to ω_{LM1} (in red) and ω_{LM4} (in blue); (b) and (c) display Cayley-transformed spectra for the mapping parameters $\sigma = \text{Real}\{\omega_{LM1}\} + 2i$ and $\mu = \text{Real}\{\omega_{LM1}\} + 8i$ and $\sigma = \text{Real}\{\omega_{LM4}\} + 2i$ and $\mu = \text{Real}\{\omega_{LM4}\} + 8i$, respectively. A dashed Cartesian grid as well as the unstable half-plane (in grey) and their mapped counterparts are also shown. The green line indicates the region of convergence of the Cayley-transformed Arnoldi method. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

for each outer step of the Arnoldi method to construct the $(j + 1)$ th Krylov vector in (10). This is accomplished by using a Krylov-based iterative linear solver. From the commonly applied techniques of Generalized Minimum Residual (GMRES) method [38], stabilized Bi-Conjugate Gradient Iteration (BiCGStab) [51] and transpose-free quasi minimum residual (TFQMR) method [13], we choose the BiCGStab algorithm since its three-term recurrence relation results in low memory requirements. The solution of the linear system (18) by an iterative method can only be approximate; as a consequence, the Cayley transformation (16) is necessarily inexact (see, e.g., [29] for a discussion on inexact transformations).

The necessity of an iterative solution technique for solving (18), however, restricts the choice of Cayley parameters, and the resulting linear system yields solutions only if the shift parameter σ is chosen sufficiently far from an eigenvalue, thus avoiding an ill-conditioned matrix. Considerations like this can be ignored when a direct inversion is attempted. For a discussion of the choice of the Cayley parameters the reader is referred to, e.g., [8] and the references therein.

4.2. Preconditioning

An efficient iterative solution of (18) requires a reliable and robust preconditioning technique. This has also been stated by Benzi [5] who considers preconditioning as the “most critical ingredient in the development of efficient solvers for challenging problems in scientific computation”. For this reason, there exists a large body of literature on preconditioning strategies, and the reader is referred to, e.g., [5,35,33] for an overview.

State-of-the-art preconditioning techniques require in general a preconditioning matrix in explicit form, as reported by [19]. These authors also discuss Jacobian-free preconditioning strategies for linear systems and conclude that “the only iterative method that can be implemented in a fashion that is literally matrix-free is a Krylov method”.

In our global stability algorithm we maintain a Jacobian-free implementation via direct numerical simulations but assume the preconditioning matrix in explicit form. Applying this (shifted) preconditioning matrix $\mathbf{P}_\sigma = \mathbf{P} - \sigma \mathbf{I}$ from the right, the modified expression of the finite difference approximation (14) reads

$$\underbrace{\mathbf{A} \mathbf{P}_\sigma^{-1} \mathbf{p}}_{\mathbf{p}} \approx \frac{\mathbf{F}(\phi_0 + \epsilon \mathbf{P}_\sigma^{-1} \mathbf{p}) - \mathbf{F}(\phi_0)}{\epsilon}, \quad (19)$$

where \mathbf{p} denotes an iteration vector in the BiCGStab algorithm (see [51] for details). This expression has to be evaluated each time a matrix-vector product is required in the inner iteration, the solution of Eq. (18) of the Krylov-based linear solver. The outer Arnoldi iteration is not affected by \mathbf{P}_σ .

Preconditioning the inner iterations raises two important questions related to the choice of \mathbf{P} and the manner of solving the required linear system $\mathbf{P}_\sigma \hat{\mathbf{p}} = \mathbf{p}$. For \mathbf{P} we choose a second-order finite difference approximation of the Jacobian matrix. Owing to the sparsity of \mathbf{P} we can take advantage of efficient incomplete decomposition techniques. The degree of “incompleteness” is given by the chosen sparsity structure of the decomposition. For general matrices, the LU-decomposition results in upper/lower triangular matrices that are dense. Incomplete decompositions, on the other hand, yield matrices that have a characteristic sparsity pattern and can be inverted efficiently by standard algorithms. In our case we choose an incomplete LU-decomposition, i.e., the dual truncation technique ILUT(p, τ) in which dropping during the factorization is based on two user-specified parameters: the fill level p and the drop tolerance τ [37]. This strategy was successfully applied by [35] to precondition complex-valued matrices, and as a dropping rule for a given fill level maximally p super-diagonal and p sub-diagonal elements are kept in each row of LU.

Denoting our low-order approximation of the Jacobian matrix by \mathbf{P}_{low} we can recast Eq. (19) as follows

$$\mathbf{A}_{high} \mathbf{P}_{low, \sigma}^{-1} \mathbf{p} \approx \frac{\mathbf{F}_{high}(\phi_0 + \epsilon \mathbf{P}_{low, \sigma}^{-1} \mathbf{p}) - \mathbf{F}_{high}(\phi_0)}{\epsilon}, \quad (20)$$

which illustrates the combination of high-order Jacobian evaluation with low-order preconditioning, denoted by the subscripts *high* and *low*, respectively. For the sake of simplicity, these subscripts are omitted for \mathbf{A} and \mathbf{P} in the following.

Ideal preconditioning would result in eigenvalues of $\mathbf{A}_\sigma \mathbf{P}_\sigma^{-1} = (\mathbf{A} - \sigma \mathbf{I})(\mathbf{P} - \sigma \mathbf{I})^{-1}$ at one. In practice, however, one has to be content with a clustering of the eigenvalues of $\mathbf{A}_\sigma \mathbf{P}_\sigma^{-1}$ about one. The quality of a preconditioner can thus be measured by the distance of these eigenvalues from one but also by their distance from the origin which is necessary to avoid ill-conditioning. These eigenvalues depend on four factors: (i) the discretization in $\mathbf{A} \equiv \mathbf{J}(\phi_0)$, (ii) the choice of \mathbf{P} (discretization, formulation, etc.), (iii) the technique employed to efficiently invert \mathbf{P}_σ and (iv) the choice of the shift parameter σ .

4.3. Proposed global stability algorithm (PCIRAM)

In summary, our proposed DNS-based global stability solver consists of the five steps (S1–S5) shown below. This solver requires the user to specify several parameters which are related to the implicitly restarted Arnoldi method (m, k, tol_A), the Jacobian-free implementation (ϵ_0), the Cayley transformation (σ, μ), the iterative linear solver (tol_B), and the ILUT-preconditioner (p, τ). Additionally, the starting vectors \mathbf{v}_1 and $\mathbf{v}_{j+1,0}$ for the Arnoldi method and the iterative linear solver are required.

- S1. Compute the base flow ϕ_0 (if not available in analytic form)
- S2. Compute and save $\mathbf{F}_{high}(\phi_0)$ (required for the Jacobian-free framework)
 - ⇒ Call right-hand-side in DNS to compute $\mathbf{F}_{high}(\phi_0)$
- S3. Setup the preconditioning matrix \mathbf{P} by a low-order approximation of the high-order Jacobian matrix \mathbf{A} , and compute the incomplete LU-decomposition of $\mathbf{P}_\sigma = \mathbf{P} - \sigma \mathbf{I}$.
- S4. Perform outer Arnoldi iterations to solve the (global) eigenvalue problem

$$\lambda \mathbf{x} = \mathbf{A} \mathbf{x}$$

- (A) Choose initial condition \mathbf{v}_1
- (B) Iterate *until convergence*: for $j = 1, 2, \dots$
 - B.1 Apply Cayley transformation and compute \mathbf{v}_{j+1}

$$(\mathbf{A} - \sigma \mathbf{I}) \mathbf{v}_{j+1} = \underbrace{(\mathbf{A} - \mu \mathbf{I}) \mathbf{v}_j}_{\mathbf{b}}$$

- (a) Obtain the matrix-vector product in \mathbf{b} using

$$\mathbf{A} \mathbf{v}_j \approx \frac{\mathbf{F}_{high}(\phi_0 + \epsilon \mathbf{v}_j) - \mathbf{F}_{high}(\phi_0)}{\epsilon}$$

- ⇒ Call right-hand side in DNS to compute $\mathbf{F}_{high}(\phi_0 + \epsilon \mathbf{v}_j)$

- (b) Perform inner iterations of ILU-preconditioned BiCGStab to solve

$$(\mathbf{A} - \sigma \mathbf{I}) \mathbf{P}_\sigma^{-1} \mathbf{P}_\sigma \mathbf{v}_{j+1} = \mathbf{b}$$

- (b.1) Choose initial condition $\mathbf{v}_{j+1,0}$
- (b.2) Iterate *until convergence*: for $i = 1, 2, \dots$
 - (a) Incomplete solution of $\mathbf{P}_\sigma \hat{\mathbf{p}} = \mathbf{p}$
 - (b) Obtain the matrix-vector product via

$$\mathbf{A}\hat{\mathbf{p}} \approx \frac{\mathbf{F}_{high}(\phi_0 + \epsilon\hat{\mathbf{p}}) - \mathbf{F}_{high}(\phi_0)}{\epsilon}$$

\Rightarrow Call right-hand side in DNS to compute $\mathbf{F}_{high}(\phi_0 + \epsilon\hat{\mathbf{p}})$

- (c) Check for convergence of $\mathbf{v}_{j+1,i}$ via

$$\|\mathbf{r}_i\|/\|\mathbf{b}\| \leq tol_B,$$

where \mathbf{r}_i denotes the current residual error and tol_B is a user-specified tolerance parameter.

B.2 Check for convergence of the desired Ritz pairs $(\tilde{\mathbf{x}}_j, \theta_j)$ using the Ritz estimate

$$|\beta_m \mathbf{e}_m^T \mathbf{y}_j| \leq \max(\epsilon_M \|\mathbf{H}_m\|, tol_A \cdot |\theta_j|),$$

where $\beta_m = \|\mathbf{f}_m\|$, ϵ_M stands for machine epsilon and tol_A denotes a user-specified tolerance parameter.

S5. Finally, recover eigenvalues λ of \mathbf{A} via

$$\lambda = \frac{\sigma\theta - \mu}{\theta - 1},$$

while the eigenvectors \mathbf{x} follow from $\tilde{\mathbf{x}}$.

5. Results

After having established and analyzed the components of an iterative global stability solver based on direct numerical simulations (DNS), we now demonstrate its effectiveness in extracting information of the perturbation dynamics on two examples. The first example continues our test case of the compressible mixing layer introduced in Section 2. The second example concerns supersonic flow about a swept parabolic body where the global treatment of the associated stability problem will yield new physical results and provide a significant numerical challenge to our global stability algorithm.

All numerical simulations shown in this article have been performed on an SGI Altix 4700 with a clock rate of 1.6 GHz. For all cases considered in the next sections, the initial condition \mathbf{v}_1 for the implicitly restarted Arnoldi method (IRAM) has been taken as a field of randomly distributed values, localized in space in order to satisfy the appropriate boundary conditions.

5.1. Example 1: compressible mixing layer

The compressible mixing layer, introduced earlier, represents a generic flow configuration that can be observed, at least locally, in many technological and industrial applications. In our investigation we will focus on the unstable shear mode (LM1) and a specific mode from the acoustic branch (LM4) as depicted in Fig. 2 as well as on some selected parameter settings (see Table 2 for details); we will discuss flexibility, accuracy, robustness and efficiency of our global stability solver employing three methods: (i) the (simple) implicitly restarted Arnoldi method (IRAM), (ii) its Cayley-transformed but unpreconditioned version (abbreviated as CIRAM) and (iii) its Cayley-transformed and ILU-preconditioned variant (abbreviated as PCIRAM). As parameters for the Arnoldi method we set $m = 80$, $k = 32$ and $tol_A = 10^{-4}$, and we choose $\epsilon_0 = 10^{-8}$ for the Jacobian-free implementation.

5.1.1. Increasing flexibility: the Cayley-transformed IRAM

We start by studying the Cayley-transformed implicitly restarted Arnoldi method (PCIRAM) as our proposed global stability algorithm (see Section 4.3) and demonstrate its flexibility in computing specified regions of the global spectrum (related to shear modes or acoustic modes) by adjusting the Cayley parameters. For this purpose, the two configurations displayed in Fig. 3(b) and (c), the first for the Cayley parameters $\sigma = \text{Real}\{\omega_{LM1}\} + 5i$ and $\mu = \text{Real}\{\omega_{LM1}\} + 20i$ and the second applying $\sigma = \text{Real}\{\omega_{LM4}\} + 5i$ and $\mu = \text{Real}\{\omega_{LM4}\} + 20i$ have been chosen, and the global spectra (a) and associated global modes (b) and (c) are visualized by the disturbance pressure $\tilde{p}(x, y)$ in Fig. 4. Guess values for the Cayley parameters have to be taken when no prior information or estimates of the flow behavior are available; an iterative adapting of these parameters is conceivable. The real part of each (normalized) mode is plotted using three periods in the periodic x -direction. A clear distinction between global modes with support in the areas of highest shear (GM1) and acoustic modes with non-

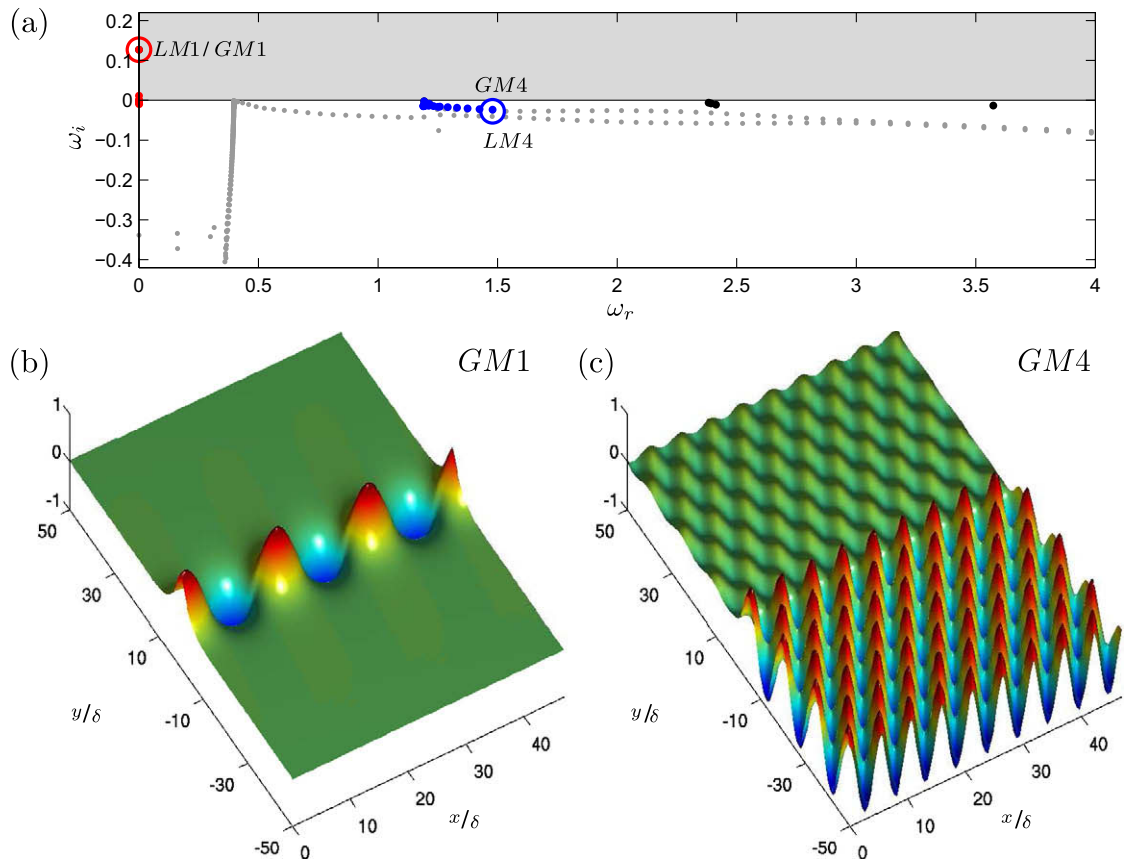


Fig. 4. (a) Two spectra (in red and blue) which have been computed via our DNS-based global stability solver PCIRAM using the Cayley transformations displayed in Fig. 3(b) and (c), respectively. The eigenvalues belonging to ω_{LM1} and ω_{LM4} (see Table 1) are indicated by circles; the directly obtained spectrum (in dark grey) as well as the unstable half-plane (in light grey) are also shown. The associated global modes $GM1$ (b) and $GM4$ (c) are visualized by the real part of the (normalized) pressure. Parameters given in Table 2 (Config IV) have been chosen, and 32×201 grid points have been used to resolve the x - and the inhomogeneous y -direction, respectively (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.).

zero amplitudes in the freestream ($GM4$) can be made. The successful extraction of these modes from the general spectrum demonstrates the increase in flexibility of the DNS-based global stability solver as the Cayley transformation is added. As expected, the two-dimensional global spectrum is far more complex due to a superposition of modes with multiple streamwise scales as a consequence of a discretization in the streamwise x -direction. Therefore, it should not come as a surprise that the Arnoldi method does not necessarily converge to the global acoustic mode shown in Fig. 2(c). Instead, the PCIRAM extracts the least-stable global modes with respect to the chosen Cayley parameters.

The possibility of the Cayley-transformed IRAM of exploring desired parts of the full global spectrum represents at the same time a significant drawback of the Arnoldi method without such a transformation which, applied to the same choice of flow parameters, converges to only the unstable mode and a random sample of other modes, such as fast-traveling acoustic modes (see Fig. 4(a) black dots). No influence over the convergence towards specific modes, however, can be exerted.

5.1.2. Increasing accuracy: the influence of the parameter ϵ_0

An important component of our global stability solver contains the replacement of the exact Jacobian matrix $\mathbf{J}(\phi_0)$ by a first-order finite-difference approximation (see Eq. (14)). This approximation crucially depends on a user-defined parameter ϵ_0 . In particular, we are interested in the influence of ϵ , computed via Eq. (15), on the accuracy of the Ritz values as well as the Ritz pairs.

For this analysis, we resort to the (simple) implicitly restarted Arnoldi method (IRAM) to delineate the effects of ϵ_0 from those introduced by the inexact Cayley transformation and by the preconditioner. The influence of ϵ_0 on our solution is measured by two quantities, namely, the relative error and the direct residual.

The first quantity, the relative error e_1 , is defined as

$$e_1(\epsilon_0) = \frac{|\omega_{2D,DNS}(\epsilon_0) - \omega_{1D}|}{|\omega_{1D}|}, \quad (21)$$




Fig. 5. (a) Relative error e_1 and (b) direct residual $\|\mathbf{r}_1\|$ as a function of the user-specified parameter ϵ_0 for four selected flow configurations (i)–(iv); see text for details. The (simple) implicitly restarted Arnoldi method (IRAM) was applied.

where $\omega_{2D,DNS} \equiv \theta_1$ denotes the least-stable Ritz value of the global DNS-based calculation and ω_{1D} represents the least-stable eigenvalue of the one-dimensional eigenvalue problem obtained by direct means (as shown in Table 3).

The second quantity, the direct residual $\|\mathbf{r}_1\|$ (see [24]), provides a measure of the accuracy of the computed least-stable Ritz pair $(\tilde{\mathbf{x}}_1, \theta_1)$; it is defined as

$$\|\mathbf{r}_1(\epsilon_0)\| = \frac{\|\mathbf{J}(\phi_0)\tilde{\mathbf{x}}_1 - \tilde{\mathbf{x}}_1\theta_1\|}{|\theta_1|}. \quad (22)$$

To cover a range of parameters, the following four flow configurations have been investigated (see Table 2 for details about Config I and III): (i) a low Mach-number case with a rather low resolution (Config I, $n_x = 16, n_y = 101$), (ii) the same low Mach-number case but with a higher resolution (Config I, $n_x = 32, n_y = 201$), (iii) a high-Mach number case with the previous high resolution (Config III, $n_x = 32, n_y = 201$) and, finally, (iv) a high-Mach number case with a low Reynolds number (Config III with $Re = 1000, n_x = 32, n_y = 201$). The results are summarized in Fig. 5 where the evolution of the relative error e_1 and the norm of the direct residual $\|\mathbf{r}_1\|$ are displayed versus ϵ_0 (with ϵ_0 ranging from 10^{-2} to 10^{-15} as a consequence of the double precision arithmetic).

In the case of the relative error e_1 (see Fig. 5(a)), the curves for each of the four cases show a similar shape: a rather small relative error plateau for a moderate range of ϵ_0 (10^{-6} to 10^{-11}) and a rapidly increasing relative error as we tend towards larger (10^{-2}) and smaller (10^{-15}) values of ϵ_0 . The range of values where e_1 is nearly independent of the parameter ϵ_0 , however, depends itself on the value of e_1 ; this means that the more accurate the least-stable Ritz value is determined – resulting in a smaller value of e_1 , – the more narrow is its range of independence from ϵ_0 .

Different results are found for the direct residual $\|\mathbf{r}_1\|$ (see Fig. 5(b)): instead of a plateau suggesting a range of optimal values for ϵ_0 , a distinct choice of $\epsilon_0 \approx 10^{-8}$ yields the lowest residual norm; for values larger or smaller than this critical value, the residual norm increases substantially.

It is important to keep in mind, though, that for non-Hermitian matrices $\mathbf{A} \equiv \mathbf{J}(\phi_0)$ a low direct residual does not necessarily imply an equivalent low error, and that the converged Ritz pair $(\tilde{\mathbf{x}}, \theta)$ may not represent an accurate approximation of the corresponding eigenpair $(\mathbf{x}, \lambda) \equiv (\phi, \omega)$. Direct residual and error are linked via the condition number of the system matrix \mathbf{A} , and an ill-conditioned system may yield inaccurate solutions. For precisely this reason do we observe a distinct plateau in the error (see Fig. 5(a)) but a lack thereof in the direct residual (see Fig. 5(b)). As far as the user is concerned the precise choice of ϵ_0 is not critical as long as it falls within the range of values defined by the plateau.

5.1.3. Increasing robustness: an analysis of the growth rate and the direct residual

The results of our investigation into robustness of the proposed methods are demonstrated and summarized in Table 4. In it we compare findings obtained by applying (i) the simple untransformed, (ii) the Cayley-transformed but unpreconditioned and (iii) the Cayley-transformed and ILUT-preconditioned Arnoldi method. Only the convergence to the least-stable shear mode ω_1 , with $\text{Real}\{\omega_1\} = 0$, and its dependence on a representative choice for the governing parameters are considered here. We observe that for the low-Mach number case (see Table 2, Config I) the Cayley-transformed but unpreconditioned Arnoldi method failed to produce accurate results (see CIRAM in Table 4). The ILUT-preconditioned version, on the other hand, converged towards the least-stable Ritz pair. With the exception of the failed computation employing the (unpreconditioned) CIRAM for the low-Mach number case (Config I) the accuracy of the results listed in Table 4 are satisfactory.

Table 4

Growth rates $\text{Imag}\{\omega_1\}$ and associated direct residuals $\|\mathbf{r}_1\|$ applying the simple (IRAM), the Cayley-transformed but unpreconditioned (CIRAM) and the Cayley-transformed and ILUT-preconditioned Arnoldi method (PCIRAM) for the four flow configurations defined in Table 2 and a resolution 32×201 . For the inexact Cayley-transformation we set $\sigma = 5i, \mu = 20i$ and $\text{tol}_B = 10^{-5}$ and as preconditioner we used $\text{ILUT}(10, 10^{-2})$.

Config #	IRAM		CIRAM		PCIRAM	
	$\text{Imag}\{\omega_1\}$	$\ \mathbf{r}_1\ $	$\text{Imag}\{\omega_1\}$	$\ \mathbf{r}_1\ $	$\text{Imag}\{\omega_1\}$	$\ \mathbf{r}_1\ $
I	0.187521	3.3e-5	(0.114675)	2.0e1	0.187520	7.5e-4
II	0.141167	5.1e-6	0.141167	4.4e-4	0.141165	6.3e-5
III	0.054723	2.6e-5	0.054723	1.1e-4	0.054723	5.4e-5
IV	0.127155	6.2e-6	0.127154	5.7e-5	0.127154	4.4e-5

5.1.4. Increasing efficiency: the performance of the proposed stability solver

The application of a preconditioner \mathbf{P} in the solution of Eq. (18) has proven imperative for the extraction of the least-stable mode for low-Mach numbers (see Table 4). Even for converging cases the preconditioner plays a secondary role as it dramatically improves the efficiency of the Cayley-transformed Arnoldi method (see Table 5). With this in mind, we compare the convergence behavior of our solver by applying various preconditioner matrices from the class of incomplete LU decomposition techniques denoted by $\text{ILUT}(p, \tau)$.

The results of our numerical experiments are shown in Fig. 6(a) displaying the relative residual $\|r\|/\|b\|$ as a function of the number of BiCGStab iterations. Best results are obtained – not surprisingly – by applying the complete LU-preconditioner, since no approximations (other than the low-accuracy discretization in \mathbf{P}) have been made, resulting in a true inverse of the linear system $\mathbf{P}_\sigma \hat{\mathbf{p}} = \mathbf{p}$ (see Section 4.2). To reach the same relative residual level of $\|r\|/\|b\| \leq \text{tol}_B = 10^{-5}$ the LU-preconditioned BiCGStab required 12 instead of 122 iterations (unpreconditioned). The impressive convergence acceleration, however, hides the fact that computing the full LU-decomposition as well as the solution of $\mathbf{P}_\sigma \hat{\mathbf{p}} = \mathbf{p}$ is, due to the number of non-zero entries nnz of \mathbf{P}_σ^{-1} (see Table 5), excessively costly for general matrices and thus unattractive for our application. Preconditioners based on the incomplete LU-decomposition also show a dramatic increase in convergence speed but do not incur the cost of a full LU-decomposition; $\text{ILUT}(10, 10^{-2})$ and $\text{ILUT}(10, 10^{-3})$ require 21 and 24 iterations, respectively, to reach a relative residual level of $\|r\|/\|b\| \leq \text{tol}_B = 10^{-5}$. Therefore, ILUT-preconditioners represent a class of effective and efficient convergence acceleration techniques.

The size of the fill level p , at least for values within a range that still optimizes memory requirements, does not substantially influence the convergence behavior. [35] report that “the rule of thumb is to take a large $[p]$ value, and use $[\tau]$ to control the amount of fill-in. This generally yields good results without compromising memory efficiency.” We have identified a fill level $p = 10$ as satisfying this requirement, and we found a drop tolerance $\tau = 0.01 - 0.005$ to be an optimal choice for $\text{ILUT}(10, \tau)$ in terms of cost-efficiency of our iterative linear solver. The reader is referred to the latter authors as well as to [6] for a discussion on tuning ILUT-type preconditioners and on enhancing their performance by using techniques such as reordering.

Results from numerical experiments for supersonic flow about a swept parabolic body (see Fig. 8) are presented in Fig. 6(b). As before, ILUT-based preconditioning is found to be capable of dramatically improving the speed of convergence




Fig. 6. Relative residual as a function of the number of BiCGStab iterations for unpreconditioned (noP stands for “no preconditioner \mathbf{P} ”) and selected preconditioned computations: (a) Config IV with a resolution of 32×201 grid points and (b) convergence results for a more challenging flow case given by supersonic flow about a swept parabolic body (see Section 5.2).

Table 5

Results from performance tests for the simple and the Cayley-preconditioned Arnoldi method showing the number of non-zero entries nnz of \mathbf{P}_σ^{-1} , the total number of required matrix-vector multiplications (# of matvec), the number of outer iterations of the Arnoldi method (# of outer iterations) and the CPU time; Config IV with a resolution of 32×201 grid points is used (see Table 2). For the inexact Cayley-transformation we choose $\sigma = 5i, \mu = 20i$ and $tol_B = 10^{-5}$.

Method	Ptype	nnz	# of matvec	# of outer iterations	CPU time (h)
IRAM	–	–	338245	13548	13.28
CIRAM	–	–	6252	211	6.65
PCIRAM	ILUT(10, 10^{-1})	159263	6909	218	4.10
PCIRAM	ILUT(10, 10^{-2})	219977	5777	172	3.58
PCIRAM	ILUT(10, 10^{-3})	331412	6286	192	3.89
PCIRAM	ILUK(0)	674532	5242	153	4.06
PCIRAM	ILUK(1)	1891106	5587	175	5.60

for BiCGStab. The results indicate, however, that this time a larger value of the fill level p is required to obtain a robust preconditioner; again, a drop tolerance $\tau = 0.01 - 0.005$ seems to be an appropriate choice.

In Table 5 we provide details on the performance tests for the (simple) implicitly restarted Arnoldi method (IRAM) as well as variants of the Cayley-transformed IRAM. This table contains the number of non-zero entries nnz of \mathbf{P}_σ^{-1} , the total number of matrix-vector multiplications (# of matvec), the number of outer iterations of the Arnoldi method (# of outer iterations) and the CPU time. It is found that ILU-based preconditioning techniques can be successfully applied to increase the efficiency of the global stability solver and that PCIRAM with ILUT-based preconditioning performs best in terms of CPU time.

5.1.5. Preconditioned spectra

Finally, to judge the effectiveness of the applied preconditioners we extract the product of \mathbf{P}_σ^{-1} and the high-order system matrix \mathbf{A} using the finite difference approximation

$$\mathbf{M}(:, i) = \mathbf{A}\mathbf{P}_\sigma^{-1}\mathbf{e}_i \approx \frac{\mathbf{F}_{high}(\phi_0 + \epsilon\mathbf{P}_\sigma^{-1}\mathbf{e}_i) - \mathbf{F}_{high}(\phi_0)}{\epsilon} \quad \text{with } i = 1, 2, \dots, n, \quad (23)$$

and compute the eigenvalues $\lambda_{\mathbf{M}}$ of $\mathbf{M}_\sigma = \mathbf{M} - \sigma\mathbf{P}_\sigma^{-1}$. The notation $\mathbf{M}(:, i)$ stands for the i th column of the matrix $\mathbf{M} \in \mathbb{C}^{n \times n}$, and the above expression is analogous to Eq. (20). The asymptotic convergence behavior of the preconditioned system can then be deduced from how closely the eigenvalues of \mathbf{M}_σ cluster about one – the ideal spectrum.

In Fig. 7 we present the spectra of \mathbf{M}_σ for four preconditioners employed in the previous subsection (see Table 5), where a parameter value of $\epsilon_0 = 10^{-8}$ and a shift parameter $\sigma = 5i$ have been used. The resolution was decreased to 8×201 grid points in order to be able to perform a (complete) eigenvalue decomposition. The spectrum based on a complete LU decomposition (in black) is used as a reference to assess the quality of four ILU-based preconditioners. It is found that decreasing the drop tolerance in ILUT(p, τ) from $\tau = 10^{-1}$ to 10^{-3} (see Fig. 7(a)–(c)) leads to more spectral clustering about one and, as a consequence, to an increase in convergence speed of the preconditioned BiCGStab (see Fig. 6(a)). The configuration ILUK(1) displays the best spectral properties (see Fig. 7(d)). This preconditioning technique comes, however, at the expense of more floating-point operations (represented by a larger computational time) due to a larger number of non-zero entries nnz of \mathbf{P}_σ^{-1} (see Table 5).

5.2. Example 2: supersonic flow about a swept parabolic body

We will now turn our attention to the second example which describes supersonic flow about a swept parabolic body. In this configuration the flow impinges through a bow shock onto the body forming a local stagnation flow near the attachment line which further downstream turns into a three-dimensional curved boundary-layer flow. It thus should not come as a surprise that this flow comprises a multitude of instability features that will also be reflected in the full global spectrum. There exists an abundance of literature, (e.g., [45,4,17,39,27]) that provides evidence for the following characteristic properties regarding the perturbation dynamics: (i) inside the boundary layer boundary-layer modes are present which can be divided into distinct structures near the stagnation line and a region further downstream; (ii) as a result of compressibility acoustic modes will appear; (iii) the interaction between the moving bow shock and the boundary layer might give rise to a special type of instabilities; (iv) finally, wave packet modes propagating near the edge of the boundary layer describe the convective nature of the flow. As we have previously seen for the compressible mixing layer, the mere existence of acoustic modes has put considerable strain on the global stability analysis which necessitated the use of additional tools such as a Cayley transformation and preconditioning to extract pertinent stability information from the direct numerical simulations (DNS). With its even more complex stability features, supersonic flow about a swept parabolic body requires to an even greater extent the incorporation of such tools into an effective, robust and efficient DNS-based global stability solver. Without these tools certain parts of the global spectrum will simply not be accessible, and a complete picture of the full perturbation dynamics of this flow would be out of reach.

Our direct numerical simulations are based on the same implementation as used before for the compressible mixing layer; in addition, a moving curvi-linear grid and shock-fitting techniques have been incorporated (for details the reader




Fig. 10. Spatial structure of a sample of associated global boundary-layer modes: (a,b) shape of the chordwise and the normal velocity component u and v , respectively, of a slow-moving boundary-layer mode ($B1$) in the $x-y$ -plane near the attachment line; (c) shape of two faster-moving modes ($B2, B3$) displayed by iso-surfaces of the normal velocity v .

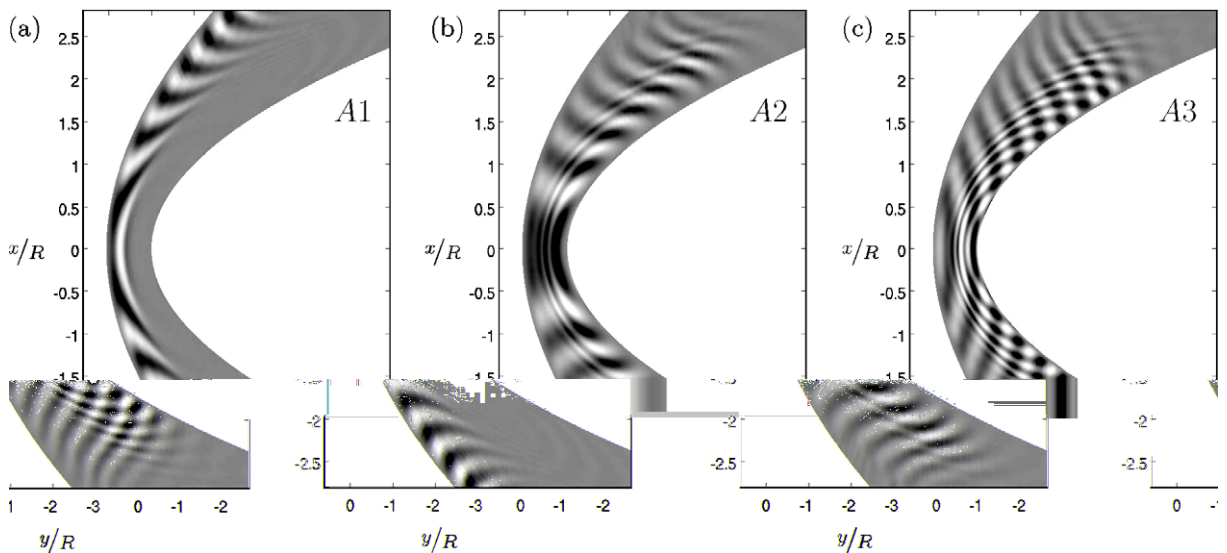


Fig. 11. Spatial structure of a sample of associated global acoustic modes visualized by the pressure p in the $x-y$ -plane: (a) unstable mode ($A1$) and (b) faster-moving stable modes ($A2, A3$).

is further mandatory to robustify the stability solver and to enhance its performance. Among the class of ILU-based techniques, ILUT was found to perform best.

For large-scale applications where a parallel computing approach is necessary or desirable the proposed global stability solver can be parallelized in a straightforward manner. For a successful application of a parallel version of the underlying DNS-code the reader is referred to [41]; for a parallel implementation of the employed implicitly restarted Arnoldi method (IRAM), the publically available software package PARPACK [28] provides parallelization details. Only the ILU-based preconditioning technique might be difficult to optimize for parallel environments.

Even though, the presented hydrodynamic global stability solver has been designed to treat complex stability problems it should be kept in mind that the interaction of multi-physical processes and numerical convergence behavior leads to a complicated dynamics which requires the careful adjustment of the governing parameters to obtain robust solutions. The appropriate choice of parameters has to be equally based on a physical understanding of the flow and a familiarity with the convergence behavior of the iterative methods. For this reason, each flow configuration under investigation has its inherent dynamical properties, and the question of proper parameter choices for a robust convergence of the global stability solver has to be answered anew.

Acknowledgments

Financial support from the Deutsche Forschungsgemeinschaft (DFG), the Studienstiftung des Deutschen Volkes, the Alexander-von-Humboldt Foundation and the ANR program “chaires d’excellence” is gratefully acknowledged. Olivier Marquet is thanked for his insightful comments on the manuscript.

References

- [1] N.A. Adams, K. Shariff, A high-resolution hybrid compact-ENO scheme for shock-turbulence interaction problems, *J. Comput. Phys.* 127 (1996) 27–51.
- [2] P. Arbenz, U.L. Hetmaniuk, R.B. Lehoucq, R.S. Tuminaro, A comparison of eigensolvers for large-scale 3D modal analysis using AMG-preconditioned iterative methods, *Int. J. Numer. Methods Eng.* 64 (2005) 204–236.
- [3] Z. Bai, Progress in the numerical solution of the nonsymmetric eigenvalue problem, *Numer. Lin. Alg. Appl.* 2 (3) (1992) 219–234.
- [4] P. Balakumar, M.R. Malik, Discrete modes and continuous spectra in supersonic boundary layers, *J. Fluid Mech.* 239 (1992) 631–656.
- [5] M. Benzi, Preconditioning techniques for large linear systems: a survey, *J. Comput. Phys.* 182 (2002) 418–477.
- [6] M. Benzi, J.C. Haws, M. Tuma, Preconditioning highly indefinite and nonsymmetric matrices, *SIAM J. Sci. Comput.* 22 (2000) 1333–1353.
- [7] W. Blumen, Shear layer instability of an inviscid compressible fluid, *J. Fluid Mech.* 40 (4) (1970) 769–781.
- [8] E.A. Burroughs, L.A. Romero, R.B. Lehoucq, A.G. Salinger, Linear stability of flow in a differentially heated cavity via large-scale eigenvalue calculations, *Int. J. Numer. Methods Heat Fluid Flow* 14 (2004) 803–822.
- [9] J.D. Crouch, A. Garbaruk, D. Magidov, Predicting the onset of flow unsteadiness based on global instability, *J. Comput. Phys.* 224 (2) (2007) 924–940.
- [10] J.D. Crouch, A. Garbaruk, D. Magidov, A. Travin, Origin of transonic buffet on aerofoils, *J. Fluid Mech.* 628 (2009) 357–369.
- [11] W.S. Edwards, L.S. Tuckerman, R.A. Friesner, D.C. Sorensen, Krylov methods for the incompressible Navier–Stokes equations, *J. Comput. Phys.* 110 (1994) 82–102.
- [12] T. Ericsson, A. Ruhe, The spectral transformation Lanczos method for the numerical solution of large sparse generalized symmetric eigenvalue problems, *Math. Comput.* 35 (152) (1980) 1251–1268.
- [13] R.W. Freund, A transpose-free quasi-minimal residual algorithm for non-Hermitian linear systems, *SIAM J. Sci. Comput.* 14 (2) (1993) 470–482.
- [14] T. Füllentuch, K. Stüben, Algebraic multigrid for selected PDE systems, in: *Elliptic and Parabolic Problems*, 2002, pp. 399–410.
- [15] T.J. Garratt, G. Moore, A. Spence, A generalised Cayley transform for the numerical detection of Hopf bifurcations in large systems, *Contributions in Numerical Mathematics*, vol. 2, World Sci. Ser. Appl. Anal. World Sci. Publishing, 1993, pp. 177–195.
- [16] R.S. Heeg, B.J. Geurts, Spatial instabilities of the incompressible attachment-line flow using sparse matrix Jacobi–Davidson techniques, *Appl. Sci. Res.* 59 (1998) 315–329.
- [17] R.D. Joslin, Direct simulation of evolution and control of three-dimensional instabilities in attachment-line boundary layers, *J. Fluid Mech.* 291 (1995) 369–392.
- [18] V. Kitsios, D. Rodriguez, V. Theofilis, A. Ooi, J. Soria, BiGlobal stability analysis in curvilinear coordinates of massively separated lifting bodies, *J. Comput. Phys.* 228 (19) (2009) 7181–7196.
- [19] D.A. Knoll, D.E. Keyes, Jacobian-free Newton–Krylov methods: a survey of approaches and applications, *J. Comput. Phys.* 193 (2) (2004) 357–397.
- [20] A. Le Duc, J. Sesterhenn, R. Friedrich, Instabilities in compressible attachment-line boundary layers, *Phys. Fluids* 18 (044102) (2006).
- [21] R.B. Lehoucq, K. Meerbergen, Using generalized Cayley transformations within an inexact rational Krylov sequence method, *SIAM J. Matrix Anal. Appl.* 20 (1) (1998) 131–148 (electronic).
- [22] R.B. Lehoucq, A.G. Salinger, Large-scale eigenvalue calculations for stability analysis of steady flows on massively parallel computers, *Int. J. Numer. Methods Fluids* 36 (2001) 309–327.
- [23] R.B. Lehoucq, J.A. Scott, An Evaluation of Software for Computing Eigenvalues of Sparse Nonsymmetric Matrices, Preprint MCS-P547-1195, Argonne National Laboratory, Argonne, IL, 1996.
- [24] R.B. Lehoucq, D.C. Sorensen, C. Yang, ARPACK Users’ Guide. Software, Environments, and Tools. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods, 1998.
- [25] S.K. Lele, Compact finite difference schemes with spectral-like resolution, *J. Comput. Phys.* 103 (1992) 16–42.
- [26] C.J. Mack, P.J. Schmid, Direct numerical simulations of hypersonic flow about a swept parabolic body. *Comput. Fluids* (2009), submitted for publication.
- [27] C.J. Mack, P.J. Schmid, J.L. Sesterhenn, Global stability of swept flow around a parabolic body: connecting attachment-line and crossflow modes, *J. Fluid Mech.* 611 (2008) 205–214.
- [28] K.J. Maschhoff, D.C. Sorensen, A portable implementation of ARPACK for distributed memory parallel architectures, in: *Copper Mountain Conference on Iterative Methods*, SIAM, 1996.
- [29] K. Meerbergen, D. Roose, The restarted Arnoldi method applied to iterative linear system solvers for the computation of rightmost eigenvalues, *SIAM J. Matrix Anal. Appl.* 18 (1) (1997) 1–20.
- [30] A. Michalke, On the inviscid instability of the hyperbolic-tangent velocity profile, *J. Fluid Mech.* 19 (1964) 543–556.
- [31] R.B. Morgan, On restarting the Arnoldi method for large non-symmetric eigenvalue problems, *Math. Comput.* 65 (215) (1996) 1213–1230.
- [32] M. Morzynski, K. Afanasiev, F. Thiele, Solution of the eigenvalue problems resulting from global non-parallel flow stability analysis, *Comput. Methods Appl. Mech. Eng.* 169 (1999) 161–176.
- [33] A. Nejat, C. Ollivier-Gooch, Effect of discretization order on preconditioning and convergence of a high-order unstructured Newton–GMRES solver for the Euler equations, *J. Comput. Phys.* 227 (2008) 2366–2386.
- [34] D. Obrist, P.J. Schmid, On the linear stability of swept attachment-line boundary layer flow. Part 1. Spectrum and asymptotic behaviour, *J. Fluid Mech.* 493 (2003) 1–29.
- [35] D. Osei-Kuffuor, Y. Saad, A Comparison of Preconditioners for Complex-valued Matrices, Technical Report, University of Minnesota Supercomputing Institute Research Report UMSI 2007/139, 2007.
- [36] J.-C. Robinet, Bifurcations in shock-wave/laminar-boundary layer interaction: global instability approach, *J. Fluid Mech.* 579 (2007) 85–112.
- [37] Y. Saad (Ed.), *Iterative Methods for Sparse Linear Systems*, second ed., Society for Industrial and Applied Mathematics, 2003.
- [38] Y. Saad, M.H. Schultz, GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* 7 (3) (1986) 856–869.
- [39] W.S. Saric, H.L. Reed, E.B. White, Stability and transition of three-dimensional boundary layers, *Annu. Rev. Fluid Mech.* 35 (2003) 413–440.

- [40] P.J. Schmid, D.S. Henningson, *Stability and transition in shear flows*, Applied Mathematical Sciences, Springer, 2001. vol. 142.
- [41] J. Schulze, P.J. Schmid, J.L. Sesterhenn, Exponential time integration using Krylov subspaces, *Int. J. Num. Methods Fluids* 60 (6) (2009) 591–609.
- [42] J.L. Sesterhenn, A characteristic-type formulation of the Navier–Stokes equations for high-order upwind schemes, *Comput. Fluids* 30 (2001) 37–67.
- [43] D.C. Sorensen, Implicit application of polynomial filters in a k-step Arnoldi method, *SIAM J. Matrix Anal. Appl.* 13 (1) (1992) 357–385.
- [44] D.C. Sorensen, Numerical methods for large eigenvalue problems, *Acta Numer.* 11 (2002) 519–584.
- [45] P.R. Spalart, Direct numerical study of leading-edge contamination. In: AGARD-CP-438, 1988, pp. 5/1–5/13.
- [46] V. Theofilis, Advances in global linear instability analysis of nonparallel and three-dimensional flows, *Prog. Aero. Sci.* 39 (2003) 249–315.
- [47] V. Theofilis, T. Colonius, T., 2004. Three-dimensional Instabilities of Compressible Flow Over Open Cavities: Direct Solution of the BiGlobal Eigenvalue Problem, AIAA Paper 2004–2544.
- [48] V. Theofilis, A. Fedorov, D. Obrist, U.C. Dallmann, The extended Görtler–Hämmerlin model for linear instability of three-dimensional incompressible swept attachment-line boundary layer flow, *J. Fluid Mech.* 487 (2003) 271–313.
- [49] U. Trottenberg, C.W. Oosterlee, A. Schüller, *Multigrid*, Academic Press, London, 2000.
- [50] L.S. Tuckerman, F. Bertagnolio, O. Daube, P. LeQuéré, D. Barkley, Stokes preconditioning for the inverse Arnoldi method, *Notes Numer. Fluid Mech.* 74 (2000) 241–255.
- [51] H.A. van der Vorst, Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of non-symmetric linear systems, *SIAM J. Sci. Statist. Comput.* 13 (1992) 631–644.
- [52] J. Zhang, Preconditioned Krylov subspace methods for solving nonsymmetric matrices from CFD applications, *Comput. Methods Appl. Mech. Eng.* 189 (2000) 825–840.