

SANDIA REPORT

SAND99-1593
Unlimited Release
Printed August 1999

Large-Scale Eigenvalue Calculations for Stability Analysis of Steady Flows on Massively Parallel Computers

RECEIVED

AUG 30 1999

OSTI

Richard B. Lehoucq and Andrew G. Salinger

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of
Energy under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Prices available from (703) 605-6000
Web site: <http://www.ntis.gov/ordering.htm>

Available to the public from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A03
Microfiche copy: A01



DISCLAIMER

**Portions of this document may be illegible
in electronic image products. Images are
produced from the best available original
document.**

SAND99-1593
Unlimited Release
Printed August 1999

Large-Scale Eigenvalue Calculations for Stability Analysis of Steady Flows on Massively Parallel Computers

Richard B. Lehoucq
Applied Mathematics Department

Andrew G. Salinger
Parallel Computational Sciences Department

Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-1110

Abstract

We present an approach for determining the linear stability of steady states of PDEs on massively parallel computers. Linearizing the transient behavior around a steady state leads to a generalized eigenvalue problem. The eigenvalues with largest real part are calculated using Arnoldi's iteration driven by a novel implementation of the Cayley transformation to recast the problem as an ordinary eigenvalue problem. The Cayley transformation requires the solution of a linear system at each Arnoldi iteration, which must be done iteratively for the algorithm to scale with problem size. A representative model problem of 3D incompressible flow and heat transfer in a rotating disk reactor is used to analyze the effect of algorithmic parameters on the performance of the eigenvalue algorithm. Successful calculations of leading eigenvalues for matrix systems of order up to 4 million were performed, identifying the critical Grashof number for a Hopf bifurcation.

1 Introduction

Computing a numerical solution to the discretized Navier-Stokes equations for system sizes of $\mathcal{O}(10^4 - 10^5)$ is commonplace. Massively parallel computers are demonstrating the ability to simulate fluid flow where the systems size is $\mathcal{O}(10^6 - 10^7)$. However, the linear hydrodynamic stability of flows for these latter systems sizes is typically not computed.

A standard approach used to determine the stability of the solution is via a numerical time evolution of the discretized equations starting from a steady state. Another approach is to linearize the equations about the steady state and then compute several eigenvalues of a large-scale generalized eigenvalue problem. This latter approach is the subject of the study.

Our interest is in characterizing the stability of complex three dimensional systems with coupled fluid flow, heat transfer, and mass transfer. Hence we are interested in the numerical solution of large scale generalized eigenvalue problems for system sizes of $n = \mathcal{O}(10^6)$. Because the solution of generalized eigenvalue problem necessarily involves solving linear systems, sparse direct methods are not a viable alternative. They possibly require $\mathcal{O}(n^2)$ operations plus a prohibitive amount of memory. Instead, this paper considers the use of iterative methods for the linear solves on massively parallel machines. Along with a scalable eigensolver, such an approach allows stability analysis to be performed on large systems arising from 3D models.

We present a large scale eigenvalue algorithm that allows us to determine the linear stability of a representative problem of 3D incompressible flow of heat transfer in a rotating disk reactor. While the steady flow for this application is axisymmetric and can be computed with a 2D model, the stability of the flow to 3D disturbances is needed to confidently use the results to design reactors. We carefully discuss the influence of the various algorithmic parameters on the performance of the stability analysis. Successful calculations were performed on this problem where the order of the matrix eigenvalue problem was 4 million. Our algorithm identified a critical Grashof number for a Hopf bifurcation, above which the reactor exhibits undesirable flow behavior.

Our paper is organized as follows. Section 2 discusses the flow of heat transfer in a rotating disk reactor and the computation of the steady state solution. Section 3 formulates the eigenvalue problem used to compute the stability of the steady state and describes in detail our numerical scheme

for solving the large scale generalized eigenvalue problem using the Cayley transformation. Section 4 discusses in detail some of the issues related to using an iterative linear solver on parallel computers as part of the eigenvalue calculation. Section 5 applies the linear stability analysis capability to determine the critical Grashof number for a Hopf bifurcation. We summarize our findings along with concluding remarks in section 6.

2 Steady Flow Problem

Much of the behavior in a nonlinear system can be uncovered by tracking steady state solutions as they evolve with changes in system parameters. In this section, we will describe our representative flow problem and our calculation of the steady state solution. That we have computed a solution to the steady state equations does not indicate whether the solution is stable or unstable motivates the development of a linear stability analysis capability for large-scale flow problems.

The rotating disk reactor (RDR) is a common system for growing high quality thin films via chemical vapor deposition. A top view and cross sectional view of the reactor configuration are shown in Figure 1. The reactor consists of an outer cylindrical can and a smaller cylinder can inside, which is rotating and heated on top. On this heated disk, the deposition occurs via surface reactions. Plug flow enters the top circular area, passes over the heated, rotating disk, and through then through annular region before leaving the computational domain. Under certain conditions, the flow in the reactor is well represented by the von Karman similarity solution for flow over an infinite rotating disk [1], leading to very desirable growth conditions.

The rotating disk reactor is known from experiments and calculations to exhibit flow instabilities. This includes the formation of stable yet undesirable recirculation cells [1, 2, 3] as well as unsteady flows [4]. The steady state behavior of this system can be uncovered by bifurcation analysis of steady state flows [2, 5]. While these solutions are axisymmetric and require just 2D calculations, the stability analysis must be able to detect instabilities to non-axisymmetric states and so we have calculated the steady flow using a full 3D model.

To provide the most general results, we study just the fluid flow and heat transfer model (no mass transfer or reactions) and look at dimensionless




Figure 1: Top view and cross section of rotating disk reactor for chemical vapor deposition reactions. The surface elements shown correspond to a 94656 element mesh of hexahedrons and 500215 unknowns.

Table 1: The governing PDE's for incompressible flow with heat transfer are shown in dimensionless form, including the Navier-Stokes equations with the Boussinesq approximation, the continuity equation, and a heat balance.

Momentum	$\frac{dv}{dt} + v \cdot \nabla v = -\nabla P + \nabla^2 v + GrTe_z$
Total Mass	$\nabla \cdot v = 0$
Thermal Energy	$\frac{dT}{dt} + v \cdot \nabla T = \frac{1}{Pr} \nabla^2 T$

numbers, which are based on the assumptions of constant properties and the Boussinesq approximation for buoyancy. For the calculations in this paper, we have fixed the design parameters as shown in the figure, with $L/R = 1.0$, $W/R = 1.2$, and $H/R = 1.0$. The operating parameters in the model that are also fixed for these calculations include the Rotational Reynolds number, $Re_{rot} = (\Omega R^2)/\nu = 83.77$ and the Prandlt number $Pr = \nu/\alpha = 1.0$ where Ω is the rotation rate, ν is the kinematic viscosity, and α is the thermal diffusivity. The Reynolds number at the inlet is fixed at the matching condition, which is the flow rate that would be drawn by an infinite disk rotating at Re_{rot} . The asymptotic value for the inlet velocity [4] of $V = 0.884\sqrt{\Omega\nu}$ leads to a inlet Reynolds number of $Re = (2RV)/\nu = 16.18$. The final parameter is the Grashof number, measuring the relative strength of buoyancy forces to viscous forces. This parameter is varied at the end of the paper, but for most calculations is held constant at $Gr = (g\beta TR^3)/\nu^2 = 15000$, where g is the magnitude of gravity, β is the thermal expansion coefficient, and T is the temperature difference between the heated disk and the inlet (with the outer walls also being held at the inlet temperature).

The steady state Navier-Stokes equations with the Boussinesq approximation are solved along with the continuity equation for incompressible flows. In addition, a heat equation with convective and conduction terms is solved. The equations are shown in Table 1 and include the time dependent terms, which are important for the formulation of the stability (eigenvalue) calculation.

The computational domain is discretized using a mesh of 94656 hexahedral elements, which corresponds to 100043 nodes. The circular area is paved with an unstructured mesh, as can be seen from a top view in Figure 1(a),

while the axial direction is structured, as seen in the cross-sectional view in Figure 1(b). The mesh is partitioned with the Chaco graph partitioning package [6], using a multi-level method and Kernighan-Lin refinement. The mesh was partitioned into the same number of sub-domains as the number of processors for the run, which was 250 for most of the calculations described below. The calculations were performed on the Sandia-Intel Tflops Computer [7].

The MPSalsa, massively parallel, unstructured grid, reacting flow code is used to solve for the steady state solution [8, 9]. This code has been successfully used to analyze flows and deposition profiles in chemical vapor deposition reactors [10, 11]. MPSalsa uses a Galerkin/Least-Squares method [12] (GLS) to discretize the time-dependent Navier-Stokes equations of incompressible fluid flow. This formulation includes a pressure stabilization term so that the velocity components, temperature, and pressure fields can all be represented with the same trilinear basis functions.

This discretization leads to a system of 500215 unknowns. The equations are solved using a fully coupled Newton's method [13] with an analytically calculated Jacobian matrix. Within each iteration of Newton's method, the finite element residuals and Jacobian matrix are assembled in 2.0 seconds when run on 250 processors. The linear solve is performed with the Aztec package [14] using a GMRES iteration without restarts, row-sum matrix scaling, and an ILU preconditioner with up to 7 levels of fill in. An average GMRES solve required 80 iterations and 30 seconds (including the time to construct the preconditioner) to reach a drop in the scaled residual of 0.007. The steady state solution at $Gr = 15000$ was reached from a trivial initial guess in 7 minutes using 2 consecutive steady state solves for increasing Gr .

A visualization of the steady state flow is shown in Figure 2. Several streamlines are shown entering the top of the reactor, spiraling over the disk, and exiting through the annular region. This calculation does not give any information on the stability of the steady state solution to small perturbations.

For the remainder the article, unless otherwise stated, all numerical experiments on linear stability analysis algorithms are undertaken about the above steady state calculation.




Figure 2: Visualization of the three-dimensional steady state flow solution. Streamlines enter the top, pass over the heated disk, and leave through the annular region.

3 Stability Analysis Calculation

Suppose that $\mathbf{f}(\mathbf{y}, \dot{\mathbf{y}}, \mathbf{p}) = \mathbf{0}$ represents the discretized non-linear system of time-dependent Navier-Stokes equations modeling incompressible fluid flow coupled with heat transport. If we linearize the equation

$$\mathbf{f}(\mathbf{y}, \dot{\mathbf{y}}, \mathbf{p}) = \mathbf{0}$$

about the steady state $(\mathbf{y}_0, \mathbf{p}_0)$ to small perturbations $e^{\lambda t} \mathbf{z}$, we obtain the generalized eigenvalue problem $\mathbf{J}\mathbf{z} = \lambda \mathbf{B}\mathbf{z}$ where the Jacobian and mass matrices are $\mathbf{J} = \mathbf{f}_y(\mathbf{y}_0, \mathbf{0}, \mathbf{p}_0)$ and $\mathbf{B} = -\mathbf{f}_{\dot{y}}(\mathbf{y}_0, \mathbf{0}, \mathbf{p}_0)$ respectively. We denote the order of the matrices \mathbf{J} and \mathbf{B} by n . Because we use a GLS discretization scheme, the generalized eigenvalue problem can be written as

$$\begin{pmatrix} \mathbf{L} & -\mathbf{C} \\ -\mathbf{C}^T + \mathbf{G} & \mathbf{K} \end{pmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \begin{pmatrix} \mathbf{M} & \mathbf{0} \\ \mathbf{N} & \mathbf{0} \end{pmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} \lambda \quad (1)$$

where \mathbf{u} is the vector of fluid velocity components and temperature unknowns, \mathbf{p} is the pressure, \mathbf{M} is the symmetric positive definite matrix of

the overlaps of the finite element basis functions, \mathbf{N} is an up-winded mass matrix, \mathbf{L} is the sum of the discretized diffusion, nonlinear convection and any possible reaction operators, \mathbf{C} is the discrete gradient, \mathbf{C}^T is the discrete divergence operator, and \mathbf{G} and \mathbf{K} (pressure Laplacian) are stabilization terms arising from the GLS.

The steady state is stable if $\text{Real}(\lambda) < 0$ for all the eigenvalues of (1). Hence, computing approximations to the right-most eigenvalues determines the stability of the steady state.

3.1 Formulation of the Eigenvalue Problem

To compute the right-most eigenvalues, a shift-invert spectral transformation [15] is typically used to transform (1) into the standard eigenvalue problem

$$\mathbf{T}_s \mathbf{z} = (\mathbf{J} - \sigma \mathbf{B})^{-1} \mathbf{B} \mathbf{z} = \gamma \mathbf{z}, \quad \gamma = \frac{1}{\lambda - \sigma}. \quad (2)$$

The above formulation maps the infinite eigenvalues of (1) (arising from singular \mathbf{B}) to zero. By selecting the pole σ near the imaginary axis, the right-most eigenvalues are mapped by \mathbf{T}_s into those of largest magnitude. However, because \mathbf{J} and \mathbf{B} are real matrices, we only allow a real σ to keep the computation in real arithmetic. Although a natural choice is to select a zero pole, the resulting transformation might miss a Hopf bifurcation (complex conjugate pair of eigenvalues that cross into the right half of the complex plane). This occurs, for instance, when the distance to the Hopf bifurcation is greater than the distance to other (perhaps stable) eigenvalues of (1). The paper [16] discusses these issues in some detail.

The computational burden is in solving the linear set of equations with coefficient matrix $(\mathbf{J} - \sigma \mathbf{B})^{-1} \mathbf{B}$. The standard approach is to use a sparse direct solver to factor $\mathbf{J} - \sigma \mathbf{B}$ and then solve linear sets of equations. Although this transformation maps the eigenvalues near the pole to those of largest magnitude, the transformation also maps the eigenvalues far from the pole to zero. Hence, the spectral condition number (the ratio of the largest to smallest, in magnitude, eigenvalues) of \mathbf{T}_s can be quite large. The resulting linear systems will be difficult to solve because the rate of convergence with an iterative method [17, 18] depends strongly upon the spectral condition number.

A better conditioned linear set of equations is achieved when using a generalized Cayley [15] transformation

$$\mathbf{T}_c \mathbf{z} = (\mathbf{J} - \sigma \mathbf{B})^{-1} (\mathbf{J} - \mu \mathbf{B}) \mathbf{z} = \gamma \mathbf{z}, \quad \gamma = \frac{\lambda - \mu}{\lambda - \sigma}. \quad (3)$$

We call μ the zero of the Cayley transform. In contrast to shift-invert transform, the Cayley transform maps any eigenvalues of (1) far from the pole close to one. If we are able to select a pole σ that is to the right of all the eigenvalues (1) and choose $\mu > \sigma$, then the smallest eigenvalue of \mathbf{T}_c is no smaller than one (in magnitude). Moreover, by judiciously choosing the pole, we can approximately bound the largest eigenvalue of \mathbf{T}_c (in magnitude) resulting in a small (say order ten) spectral condition number.

The last two paragraphs describe a delicate balancing act. On the one hand, the ability to compute the rightmost eigenvalues (λ 's) requires that the Cayley transformation map these values to γ 's that are the largest (in magnitude). Such a situation allows the eigensolver to perform well. On the other hand, the iterative solver used to solve the linear systems arising from the Cayley transformation is negatively impacted if the ratio of $\max(|\gamma|)$ to $\min(|\gamma|)$ (the spectral condition number of \mathbf{T}_c) is large.

We remark that although the Cayley and shift-invert spectral transformation both involve $(\mathbf{J} - \sigma \mathbf{B})^{-1}$, the system of linear equations solved by each transformation is distinct. Given a vector \mathbf{x} , the Cayley system requires the solution of

$$(\mathbf{J} - \sigma \mathbf{B}) \mathbf{v} = (\mathbf{J} - \mu \mathbf{B}) \mathbf{x} \quad (4)$$

so that $\mathbf{v} = \mathbf{T}_c \mathbf{x}$. Instead, the shift-invert system solves $(\mathbf{J} - \sigma \mathbf{B}) \mathbf{v} = \mathbf{B} \mathbf{x}$. That the spectral condition number of the Cayley system can be tightly bounded (via a careful choice of σ and μ) implies that the Cayley system results in a better conditioned set of linear equations.

3.2 Solution of the Eigenvalue Problem

We employ an implicitly restarted Arnoldi method (IRAM) as implemented in the parallel implementation [19] of the ARPACK [20] to compute eigenvalues and eigenvectors of the generalized eigenvalue problem. We have slightly modified the P_ARPACK subroutines pdnaupd and pdneupd to implement the Cayley transformation. We refer the reader to [20] for full details about the software and underlying algorithm.

Start P_ARPACK with the vector $v = J^{-1}Bx$ where x is random vector. Select a pole and zero for T_c .

1. Compute m iterations of Arnoldi's method with T_c using the starting vector v . Compute m eigenvalues (the θ 's approximating the γ 's) of the order m upper Hessenberg matrix constructed by P_ARPACK.
2. Map the θ 's to $\hat{\lambda}$'s (approximations to the λ 's) via the inverse Cayley transformation.
3. Exit if the k rightmost $\hat{\lambda}$'s satisfy the user specified tolerance.
4. Implicitly restart Arnoldi's method resulting in an updated starting vector v .
5. Update σ and μ using the current approximate eigenvalues.

Figure 3: Computing the leading eigenvalues of $Jz = Mz\lambda$ using the Cayley transformation and IRAM.

Figure 3 lists the scheme used for computing several (say k) right-most eigenvalues of (1). A few remarks are in order. The starting vector is chosen so that it does not contain any components [21] in the null-space of B . For all the eigenvalue problems solved, the value of $m = 24$ was used. At step 2, the eigenvalues (θ 's approximating γ 's) of the m by m Hessenberg matrix are mapped back to the system defined by (1) via the inverse Cayley transformation resulting in approximations $\hat{\lambda}$'s. The eigensolver is terminated when these k rightmost approximate eigenvalues satisfy the user specified tolerance. The check that must be satisfied is $\|J\hat{z} - \hat{\lambda}B\hat{z}\|/\|B\hat{z}\|$ where \hat{z} is the associated approximate eigenvector. By implicitly restarting the Arnoldi iteration, we compute a new starting vector for the subsequent run (step 1). Implicitly restarting is an efficient and stable manner to restart Arnoldi's method so that storage requirements remain fixed for the computation. Finally, at step 5, the new pole and zero for the next Cayley transformation are updated so that the spectral condition number of T_c is of order ten.

We remark that there are two iterations—an outer and an inner iteration. The outer iteration is Step 1 of the algorithm listed in Figure 3. During each of these outer iterations, there is an inner iteration used to solve the linear set of equations (4) arising from applying T_c . We use a GMRES iteration for solving this linear set of equations. The next section discuss details associated with the inner iteration.

The two parameters σ and μ in the Cayley transformation give a considerable amount of flexibility over what eigenvalues will be located by the Arnoldi's method, how accurately they will be calculated, and how expensive the calculation will be. The major consideration is the size of $|\gamma|$ for the eigenvalues λ of interest. Eigenvalues λ that are mapped to large $|\gamma|$ will emerge and quickly be approximated by Arnoldi's method. We present results that quantify the various trade-offs in picking these parameters while preserving the spectral condition number of T_c .

A good choice for these parameters is for the right-most eigenvalues of interest, λ_i for $i = 1 : k$, to have real parts that satisfy $2\sigma - \mu < \text{Real}(\lambda_i) < \sigma < \mu$. This implies that these λ_i are mapped so that $|\gamma(\lambda_i)| \geq 2$ as long as $|\text{Imag}(\lambda_i)|$ is not large compared to $\sigma - \text{Real}(\lambda_i)$.

To illustrate how the Cayley transformation maps eigenvalues of the system, we plot the magnitude of Cayley transformation in Figure 4. This figure shows how λ is mapped to $|\gamma|$ for fixed values of $\sigma = 20$ and $\mu = 80$ and four imaginary portions of λ . Note that as the real part of λ decreases, $|\gamma|$ approaches one. The σ and μ values in this plot map the real eigenvalues in the range of $-40 < \text{Real}(\lambda) < 20$ to magnitudes in the Cayley transformed system of $|\gamma| > 2$, which is sufficiently well separated from the many eigenvalues near $|\gamma| = 1$ for the eigensolver. For any real eigenvalues satisfying $\text{Real}(\lambda) < -40$, the Cayley transformation maps these eigenvalues so that $1 < |\gamma(\lambda)| < 2$. Hence, Arnoldi's method will provide the best approximations to the eigenvalues satisfying $2\sigma - \mu < \text{Real}(\lambda_i) < \sigma$.

Figure 4 also indicates that eigenvalues with large imaginary parts are mapped to small $|\gamma|$. Therefore it is difficult to compute approximations to eigenvalues with large imaginary parts. For instance an eigenvalue at $0 \pm 50i$ might not be located if there are many eigenvalues with large $|\gamma|$, such as near $-5 \pm 0i$. This problem is resolved by increasing the Arnoldi space m needed by P_ARPACK.

An appropriate choice of σ , μ , and the size of the Arnoldi space m is therefore a tradeoff between two factors: selecting m large enough so that the right-most eigenvalues λ are reliably computed by the eigensolver and avoiding large values of $|\gamma|$ so that the resulting linear systems can be efficiently solved with preconditioned iterative methods.




Figure 4: Plot of the transformation of the eigenvalues in the physical system to those in the Cayley transformed system. The magnitude of the transformed eigenvalue is plotted against the real part of λ for three different imaginary contributions to λ .

4 Preconditioned Iterative Linear Solves

The computationally intensive part of the eigenvalue calculation is the linear solve with \mathbf{T}_c (inner iteration) that occurs during each outer iteration of Arnoldi's method. Since we are targeting very large problems and algorithms that scale to thousands of processors, we are limited to preconditioned iterative linear solves of distributed matrices. In this section we first discuss the tolerances used for the linear solver and eigensolver, the details associated with our use of Aztec [14] and the outcome of a mesh resolution study.

4.1 Error Tolerances

Figure 5 plots the residuals associated with three rightmost eigenvalues and eigenvectors versus the convergence tolerance used for the linear solver. The eigensolver residual error is defined as

$$\frac{\|\mathbf{J}\hat{\mathbf{z}} - \hat{\lambda}\mathbf{B}\hat{\mathbf{z}}\|}{\|\mathbf{B}\hat{\mathbf{z}}\|} \quad (5)$$

where $\hat{\lambda}$ and \hat{z} are the computed eigenvalue and eigenvector approximations. The residual contains the normalization with $B\hat{z}$ because P_ARPACK normalizes $\|\hat{z}\| = 1$ and so (5) is independent of the scaling of the data. The linear solver uses the criterion

$$\frac{\|b - Ax_j\|}{\|Bv\|} < \eta \quad (6)$$

where $A = J - \sigma B$ and $b = (J - \mu B)v$ from (4), and η is a tolerance parameter that must be chosen. Here, v is the distributed unit vector provided by P_ARPACK that is to be transformed via T_c during the i -th ($1 \leq i \leq m$) outer iteration and x_j is the approximate solution after j GMRES iterations (the inner iteration).

In the experiment shown in Figure 5 we show the influence of η on the eigensolver residual error 5. The residual error of the rightmost eigenvalue pair (denoted by the solid line) stops decreasing $\eta \approx 10^{-3}$. The residual error of the next two eigenvalues stop decreasing when $\eta \approx 10^{-6}$. Driving the residual errors lower would require a larger Arnoldi space m or a different choices of σ and μ . For the rest of the calculations in this section the linear solver tolerance was fixed at $\eta = 10^{-3}$.

A series of calculations are presented in Figure 6 to illustrate the trade-offs in choosing σ . The right-most eigenvalue of the steady state calculation has real part equal to 0.3 and we set $\mu = 80$. As σ is increased from 1 to 70, the maximum $|\gamma|$ decreases (recall that $\max(|\gamma|)$ is approximately equal to the spectral condition number of T_c). This decrease is seen to correspond directly to the decrease in the CPU time and memory requirements for the linear solve, as measured by solution time and the average number of GMRES iterations needed for a solve. However, as σ is increased so do the residuals (5). For this problem, a choice of $\sigma = 20$ provides a balance between efficiency and accuracy. The trends seen as a function of σ point to a remedy for systems where the preconditioned linear solver is not able to reach the specified tolerance: increase σ and μ until the linear problem can be solved and then increase the number of outer iterations needed by the eigensolver (and therefore the number of linear solves required) until (5) is sufficiently small for the rightmost eigenvalues.




Figure 5: The error in the eigenvalue calculation for the three rightmost eigenvalues is shown to be a function of the acceptance criterion of the iterative linear solver. The eigensolver residual error and linear solver tolerance is given in (5) and (6).




Figure 6: The model eigenvalue calculation is repeated for several values of the σ parameter in the Cayley transformation. The effect of σ on (a) the magnitude of the three largest complex θ pairs, (b) the residual errors (5) associated with the three largest complex θ pairs in the transformed system, (c) the time for the calculation, and (d) the average number of GMRES iterations needed for each of the 24 linear solve are shown.

4.2 Using Aztec

Another issue associated with preconditioned iterative solvers is the robustness of the algorithm in reaching a specified tolerance. This includes the access to, and selection of, an appropriate preconditioner and solution algorithm. For the calculations in this paper, the Aztec linear solver library was used ¹. An ILUT preconditioner with considerable fill-in was selected so that the preconditioner required almost 4 times more memory than the matrix itself. The preconditioner is computed only one time, and reused for each iteration of Arnoldi's method needed by the eigensolver.

Since Figure 6(d) shows that a few hundred GMRES iterations are possibly needed during each outer iteration, the numerical stability of the GMRES implementation becomes critical. Originally, a classical Gram-Schmidt scheme was used for the orthogonalization, but the lack of numerical stability prevented the GMRES algorithm from reaching the required tolerance. Two alternative orthogonalization schemes were used successfully: two-step classical Gram-Schmidt (CGS) and a modified Gram-Schmidt (MGS). The CGS method uses two steps of orthogonalization (the second step is the correction for the possible loss of orthogonality of the Arnoldi basis vectors) but the number of global communication points remains fixed (at two) independent of the GMRES iteration. On the other hand, the MGS scheme requires i communications to orthogonalize i vectors (at GMRES iteration i) but no additional floating point operations (flops). Both schemes reached the specified tolerance and provided identical results in terms of the number of GMRES iterations needed.

There was a significant difference in the scalability of the two algorithms as the number of processors was changed. The time required to perform a single linear solve (using a pre-calculated preconditioner) was recorded with the number of processors being varied from 100 to 1000. The message of this calculation is clear when presenting the total CPU time (calculated as the wall clock time multiplied by the number of processors) as shown in Figure 7. When the problem was run on 100 processors, the extra communications required by MGS were slightly less expensive than the extra flops required by the CGS algorithm. However, as the number of processors is increased, it is seen that the communication time in MGS starts to dominate, while the total CGS time remains relatively flat. At 1000 processors, the CGS routine

¹<http://www.cs.sandia.gov/CRF/aztec1.html>




Figure 7: The parallel efficiency of two stable orthogonalization schemes in the GMRES algorithm is compared. The extra communications of the Modified Gram-Schmidt (MGS) approach scale poorly with number of processors, while the extra operations of the 2-step Classical Gram-Schmidt (CGS) approach scale well.

requires only about a quarter of the time of the MGS method.

The results in Figure 7 show that the two-step classical Gram-Schmidt (CGS) scheme scales much better than the modified Gram-Schmidt (MGS) scheme. It should be pointed out that the inter-processor communication rate of the Sandia-Intel Tflop computer is very fast compared to more loosely coupled parallel machines, where we would expect the difference to be more dramatic and the crossover point (at around 175 processors for this case) to occur at fewer processors.

4.3 Mesh Resolution

All the calculations previously discussed were carried out for a single finite element mesh corresponding to just over a half million unknowns (see section 2). While we have shown above that the eigenvalues are accurate for this given discretized system, a mesh resolution study verifies that these eigenvalues are good approximations to those of the continuous PDE model. The results of such a study are shown in Table 2. The six eigenvalues with largest

Table 2: Mesh resolution studies at $Gr = 15000$ on the six eigenvalues with largest real parts. The coarse mesh results would indicate a stable solution, but the finest mesh shows that two pairs of eigenvalues have positive real parts.

Number of Unknowns	First Eigenvalue	Second Eigenvalue	Third Eigenvalue
.25 Million	$-0.08 \pm 25.33i$	$-0.49 \pm 9.63i$	$-1.44 \pm 5.96i$
.50 Million	$0.35 \pm 25.16i$	$-0.05 \pm 9.50i$	$-1.13 \pm 5.91i$
1 Million	$0.57 \pm 25.06i$	$0.21 \pm 9.36i$	$-0.98 \pm 6.01i$
2 Million	$0.73 \pm 25.02i$	$0.39 \pm 9.31i$	$-0.85 \pm 6.03i$
4 Million	$0.84 \pm 24.94i$	$0.50 \pm 9.22i$	$-0.78 \pm 6.08i$
2D mesh	$1.06 \pm 24.69i$		
.50 Million			

real parts at $Gr = 15000$ are shown for five successively finer meshes, each doubling the number of unknowns of the previous. They range from 250 thousand to 4 million unknowns. A final calculation on a very fine 2D axisymmetric mesh of 0.5 million unknowns was used to verify the first eigen mode, which was determined to be axisymmetric through visualization of the eigenvectors.

The parameter value was chosen to be near a Hopf bifurcation. What we find is that while the coarsest mesh indicates a stable steady state, the second coarsest mesh (which is used in all other computations in this paper) shows one unstable eigenpair, and the three finest meshes predict that two eigenvalues are unstable. While the change in the eigenvalues with successive refinement is slowing, the values are still changing even when the number of unknowns increases from 2 to 4 million unknowns.

We recall that the accuracy of the nonlinear steady state calculation and of the linear solves within the eigensolver are additional sources of error that may be influencing the convergence rate of the mesh resolution study. These results imply that very fine meshes may be needed to pinpoint the exact parameter value of a Hopf bifurcation in 3D problems. However, the fact that the coarsest mesh is converging to the same physical modes implies that the system's behavior can be quickly explored with a relatively coarse mesh,

and the finer meshes are only needed to locate the parameter values to a higher degree of accuracy.

For the finest mesh, the steady state solution was reached from a trivial guess in three continuation steps in the Gr number, using 2.5 hours of CPU time on 1024 processors. The eigenvalues were calculated in under 5 hours, where each linear solve required about 12 minutes. The linear solves used the ILUT preconditioner with a fill-in factor of 8 and an average of 825 GMRES iterations.

5 Reactor Analysis

In this section we apply the linear stability analysis capability that has been presented above. Experiments have shown that the desirable non-recirculating flow in the rotating disk reactor can go unstable to periodic oscillations [4]. It is important during reactor design to be able to locate this instability. With that goal in mind, the steady-state solution branch was tracked using first-order continuation and the leading eigenvalues were calculated at each step. The calculations were performed on the standard mesh corresponding to half a million unknowns.

Figure 8 shows how the six eigenvalues that have largest real part at $Gr = 15000$ evolve from $Gr = 10000$ to $Gr = 16000$. By interpolating between the symbols to where the curves cross the imaginary axis, the first Hopf bifurcation is seen to occur near 14800, the second just above 15000 and a third near 15500. By including the trends seen in the mesh resolution study in Table 2, where the systems became less stable with more refined meshes, we can extrapolate that with a finer mesh the first Hopf bifurcation would fall in the range $Gr = 14000 - 14500$.

The eigenvectors associated with these largest eigenvalues are the perturbations that will not get damped out if the Hopf bifurcation has occurred. Visualization of the eigenvectors gives information to the designers that could be used to modify the reactor design or operation to delay the onset of these unwanted instabilities. Since the instabilities involve oscillations between the real and imaginary parts of the eigenvector, each of which corresponds to a three-dimensional flow field, it was not possible to produce satisfactory still pictures for this publication. What the visualization found was that the first Hopf bifurcation is an axisymmetric state with a toroidal roll cell. The os-




Figure 8: The tracking of the six largest eigenvalues as a function of parameter indicate a Hopf bifurcation near 14800.

cillation is the roll cell being forced out by a counter-rotating roll cell. The second Hopf bifurcation breaks symmetry with a mode 1 instability, with a single large roll cell over the disk that rotates in time. The third Hopf bifurcation is a mode 2 symmetry breaking, where there is up-flow in two quadrants of the disk and down-flow in the two others. Again, this flow structure precesses around reactor in time. It is interesting that the modes 0, 1, and 2 symmetry breakings occur at nearly the same conditions. While this problem could have been solved using a two-dimensional model with an axisymmetric formulation and complex arithmetic for the non-axisymmetric modes, the methods were developed as a general 3D capability. And since Cartesian coordinates were used to model the system, the fact that the solution was axisymmetric did not simplify the calculations in any way.

6 Summary and Conclusions

A massively parallel code for calculating steady-states of incompressible and reacting flows has been linked to a library for calculating selected eigenvalues using P_ARPACK for the purpose of linear stability analysis. A novel implementation of the Cayley transform has been presented and analyzed for an example of 3D flow and heat transfer in a rotating disk CVD reactor. This implementation allows control over the spectral condition number of the linear system that must be solved during each step Arnoldi's iteration used by P_ARPACK making it particularly well suited for use with scalable iterative linear solvers.

By using sophisticated linear algebra algorithms and software for iterative solutions of large, sparse, distributed matrices, we were able to calculate several rightmost eigenvalues for linearized systems corresponding to 4 million unknowns and 530 million nonzero matrix entries on 1024 parallel processors.

The stability of the flow in the rotating disk reactor was analyzed as a function of the Grashof number, Gr . The desirable flow field was found to go unstable in the range of $Gr = 14000 - 14500$ after extrapolating the results to finer meshes. While this flow goes unstable to an axisymmetric mode, there are mode 1 and mode 2 instabilities that go unstable at slightly higher values of Gr .

We have shown that determining the linear stability of steady state solutions arising from discretization of 3D incompressible flow PDE's is possible.

We have also demonstrated the potential impact by locating a flow instability in an engineering system that can be used to interpret certain experimental results and guide the design of the next generation, scale-up reactors.

References

- [1] G. Evans and R. Greif. A numerical model of the flow and heat transfer in a rotating disk chemical vapor deposition reactor. *J. Heat Transfer ASME*, 109:928–935, 1987.
- [2] S. Kieda D.I. Fotiadis and K.F. Jensen. Transport phenomena in vertical reactors for metalorganic vapor phase epitaxy. *J. Crystal Growth*, 102:441–470, 1990.
- [3] C. Weber, C. van Opdorp, and M. de Keijser. Modeling of gas-flow patterns in a symmetric vertical vapor-phase-epitaxy reactor allowing asymmetric solutions. *J. Appl. Phys.*, 67:2109–2118, 1990.
- [4] W. G. Breiland and Greg H. Evans. Design and verification of nearly ideal flow and heat transfer in a rotating disk chemical vapor deposition reactor. *Journal of the Electrochemical Society*, 138:1806–1816, 1991.
- [5] L. A. Romero and A. G. Salinger. Stability analysis of flow in the rotating disk reactor. *in preparation*, 1999.
- [6] B. Hendrickson and R. Leland. The Chaco user's guide: Version 2.0. Technical Report SAND94-2692, Sandia National Labs, Albuquerque, NM, June 1995.
- [7] T. G. Mattson and G. Henry. An overview of the Intel TFLOPS supercomputer. *Intel Technology Journal*, 1, 1998.
- [8] J. N. Shadid, A. G. Salinger, R.C. Schmidt, T. M. Smith, S. A. Hutchinson, G. L. Hennigan, K. D. Devine, and H.K. Moffat. Mpsalsa version 1.5, a finite element computer program for reacting flow problems part 1: Theoretical development. Technical Report SAND98-2864, Sandia National Laboratories, Albuquerque, NM, 1999.

- [9] A. G. Salinger, K. D. Devine, G. L. Hennigan, H. K. Moffat, S. A. Hutchinson, and J. N. Shadid. Mpsalsa: a finite element computer program for reacting flow problems part 1 - user's guide. Technical Report SAND95-2331, Sandia National Laboratories, Albuquerque, NM, 1996.
- [10] A. G. Salinger, J. N. Shadid, S. A. Hutchinson, G. L. Hennigan, K. D. Devine, and H. K. Moffat. Massively parallel computation of 3d flow and reactions in chemical vapor deposition reactors. Technical Report SAND97-3092, Sandia National Laboratories, Albuquerque, NM, 1997.
- [11] A. G. Salinger, J. N. Shadid, S. A. Hutchinson, G. L. Hennigan, K. D. Devine, and H. K. Moffat. Analysis of gallium arsenide deposition in a horizontal chemical vapor deposition reactor using massively parallel computations. *J. Crystal Growth*, In press, 1999.
- [12] J. R. Hughes, L. P. Franca, and G. M. Hulbert. A new finite element formulation for computational fluid dynamics: VIII. the Galerkin/least-squares method for advective-diffusive equations. *Comp. Meth. App. Mech. and Eng.*, 73:173–189, 1989.
- [13] J.N. Shadid, R.S. Tuminaro, and H.F. Walker. An inexact Newton method for fully coupled solution of the Navier-Stokes equations with heat and mass transport. *Journal of Computational Physics*, 137:155–185, 1997.
- [14] S. A. Hutchinson, L. V. Prevost, R. S. Tuminaro, and J. N. Shadid. Aztec user's guide: Version 2.0. Technical report, Sandia National Laboratories, Albuquerque, NM, 1998.
- [15] K. Meerbergen, A. Spence, and D. Roose. Shift-invert and Cayley transforms for the detection of rightmost eigenvalues of nonsymmetric matrices. *BIT*, 34:409–423, 1994.
- [16] K. A. Cliffe, T. J. Garratt, and A. Spence. Eigenvalues of the discretized navier-stokes equation with application to the detection of Hopf bifurcations. *Advances in Computational Mathematics*, 1:337–356, 1993.
- [17] Y. Saad. *Iterative Methods for Sparse Linear Systems*. PWS, Boston, MA, 1996.

- [18] Anne Greenbaum. *Iterative Methods for Solving Linear Systems*. SIAM, Philadelphia, PA., 1997.
- [19] K. J. Maschhoff and D. C. Sorensen. P_ARPACK: An efficient portable large scale eigenvalue package for distributed memory parallel architectures. In Jerzy Wasniewski, Jack Dongarra, Kaj Madsen, and Dorte Olesen, editors, *Applied Parallel Computing in Industrial Problems and Optimization*, volume 1184 of *Lecture Notes in Computer Science*, Berlin, 1996. Springer-Verlag.
- [20] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK USERS GUIDE: Solution of Large Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*. SIAM, Philadelphia, PA, 1998.
- [21] Karl Meerbergen and Alastair Spence. Implicitly restarted Arnoldi with purification for the shift-invert transformation. *Mathematics of Computation*, 218:667–689, 1997.

Distribution

EXTERNAL DISTRIBUTION:

Valmor de Almeida
Oak Ridge national Laboratory
PO Box 2008, MS6224
Oak Ridge, TN 37831-6224

Rich A. Cairncross
Department of Chemical Engineering
Drexel University
Philadelphia, PA 19104

Dr. Kostas Christodoulou
Avery Dennison Corp.
2900 Bradley Street
Pasadena, CA 91107-1599

Professor Andrew Cliffe
AEA Technology
424.4 Harwell
Didcot, OX11 ORA
United Kingdom

Jeffrey J. Derby
Dept. of Chemical Eng. and Materials Science
University of Minnesota
421 Washington Ave. S.E.
Minneapolis, MN 55455

Dr. Ian Gates
Engineering Systems Technology Center
3M Center, Building 518-1-01
St. Paul, MN 55414

Dr. Paul Fischer
Building 221
Mathematics and Computer Science Division
9700 S. Cass Avenue
Argonne, IL 60439

Dan Hitchcock
US Department of Energy
SCS, ER-30 GTN
Washington, DC 20585

Fred Howes
US Department of Energy
OSC, ER-30, GTN
Washington, DC 20585

John Guckenheimer
Department of Mathematics
Malott Hall
Cornell University
Ithaca, NY, 14853-4201

Professor Herb Keller
Applied Mathematics 217-50
California Institute of Technology
Pasadena, CA 91125

Ioannis Kevrekidis
Department of Chemical Engineering
E-Quad
Princeton NJ, 08544

T.J. Mountziaris
Department of Chemical Engineering
303 Furnas Hall
State University of New York
Buffalo, NY 14260-4200

Professor Dirk Roose
K.U. Leuven
Department of Computer Science
Celestijnlaan 200A
3000 Heverlee
Belgium

Todd Salamon
Rm. 7D-212
Lucent Technologies
600 Mountain Avenue
Murray Hill, NJ 07974

Professor A. Spence
Department of Mathematical Sciences
University of Bath
Bath BA2 7AY
United Kingdom

Professor Simon Tavener
Math Department
Penn State University
University Park, PA 16802

Konstantinos Theodoropoulos
Department of Chemical Engineering
E-Quad
Princeton NJ, 08544

David West
DOW Chemical, Engineering Sciences
2301 N. Brazosport Blvd., B-1226
Freerport, TX 77541-3257

Dr. H.C. Yee
Senior Staff Scientist
NASA Ames Research Center
Mail Stop T27B-1
Moffett Field, CA 94035

INTERNAL DISTRIBUTION:

1	MS 1111	Sudip Dosanjh, 9221
1	MS 1111	John Shadid, 9221
10	MS 1111	Andrew Salinger, 9221
1	MS 1111	Roger Pawlowski, 9221
1	MS 1110	David Womble, 9222
1	MS 1110	Richard Lehoucq, 9222
1	MS 1110	Louis Romero, 9222
1	MS 1110	David Day, 9222
1	MS 1110	Beth Burroughs, 9222
1	MS 1110	Ray Tuminaro, 9222
1	MS 1110	Mike Heroux, 9222
1	MS 0819	Timothy Trucano, 9231
1	MS 0819	Allen Robinson, 9231
1	MS 0321	William Camp, 9200
1	MS 0601	Harry Moffat, 1126
1	MS 0826	Wahid Hermina, 9111
1	MS 0826	Randy Schunk, 9111
1	MS 0826	Matt Hopkins, 9111
1	MS 0826	Phil Sackinger, 9111
1	MS 0826	Ken Chen, 9111
1	MS 0834	Mario Martinez, 9112
1	MS 0834	Polly Hopkins, 9112
1	MS 0827	C. C. Wong, 9114
1	MS 9018	Central Technical Files, 8540-2
5	MS 0899	Technical Library, 4916
1	MS 0612	Review & Approval Desk, 4912
		For DOE/OSTI