Slide 1

Methods for Large Scale Eigenvalue Problems

$$\mathbf{Ax} = \mathbf{Bx}\lambda$$

D.C. Sorensen

Rice University
Houston, TX

Euro - Par 2000
Munich

Slide 2

Outline – Methods for Large Scale Eigenvalue Problems

Implicitly Restarted Arnoldi Methods

Krylov Subspace Projection
Lanczos/Arnoldi k-Step Factorization
Implicit Restarting

Software
ARPACK
P_ARPACK

Applications
CFD
Regularization
Model Reduction

Slide 3

ARPACK Users' Guide

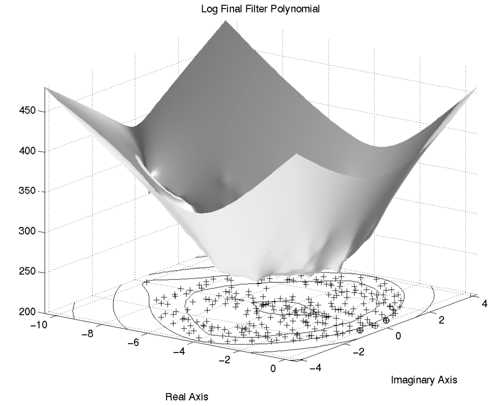

R.B. Lehoucq, D.C. Sorensen, C. Yang

Slide 4

Implicitly Restarted Arnoldi Method

$$\boxed{\mathbf{Ax} = \mathbf{Bx}\lambda}$$

A - may be nonsymmetric

Area Still Needs Research!

Arnoldi / Lanczos

Krylov Subspace Projection

Implicit Restart

Arnoldi $+$ Implicit QR

Slide 5

Implicitly Restarted Arnoldi Method

Clean Solution

Loss of Orthogonality	$\Longleftrightarrow$	Spurious Eigenvalues

Storage - Cost

Accuracy

Slide 6

ARPACK Applications

Computational Fluid Dynamics - Navier Stokes Eqn.
Stability/Bifurcation Analysis –
1M Degrees of Freedom - Sandia, NM

Design Dialectric Wave Guide - Maxwell Eqn.
Laser Diode Design – Lucent Tech.
Structural Analysis
Linear and Nonlinear vibration analysis
100K - 1M Degrees of Freedom - Sandia NM

Slide 7

APPLICATIONS contd.

Computational Chem - Schroedinger Eqn
- Reactive Scattering
- Nano-Scale Technology
- Polymer Analysis - Oak Ridge
- Quantum Chemistry – Self Consistent Field Calculations
- Large Dense Symmetric Eigenproblems
- Want 30-50% of Spectrum

Computational Biology (SVD)
- Image Enhancement - 3D image reconstruction
- Molecular Dynamics – Anlysis of Protein

Model Reduction – Dynamical Systems and Control

Slide 8

The Krylov Subspace

$$\mathcal{K}_k(\mathbf{A}, \mathbf{v}_1) = Span\{\mathbf{v}_1, \mathbf{A}\mathbf{v}_1, \mathbf{A}^2\mathbf{v}_1, \ldots, \mathbf{A}^{k-1}\mathbf{v}_1\}$$

Extract more eigen-information from the Power Sequence

Orthogonal Projection

Gram-Schmidt on Power Sequence gives O-N basis $\mathbf{V}$ for $\mathcal{K}_k(\mathbf{A}, \mathbf{v}_1)$

Arnoldi/Lanczos process solves

$$min\|\mathbf{AV} - \mathbf{VH}\|_2$$

Slide 9

The Arnoldi Factorization

$$\mathbf{AV} = \mathbf{VH} + \mathbf{f}\mathbf{e}_k^T$$

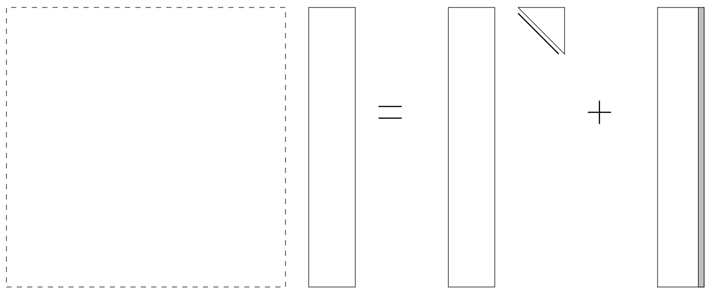

Slide 10

$$\mathbf{AV} = \mathbf{VH} + \mathbf{f}\mathbf{e}_k^T$$

Recursion	**Parallel**
(1) $\beta \leftarrow \|\mathbf{f}\|\,;\;\; \mathbf{v} \leftarrow \mathbf{f}\cdot\frac{1}{\beta};$	
(2) $\mathbf{H} \leftarrow \begin{bmatrix} \mathbf{H} \\ \beta\mathbf{e}_k^T \end{bmatrix};\;\; \mathbf{V} \leftarrow [\mathbf{V}, \mathbf{v}];$	
(3) $\mathbf{w} \leftarrow \mathbf{A}\mathbf{v}\,;$	} User supplied AV
(4) $\mathbf{h} \leftarrow \mathbf{V}^T\mathbf{w};$	} DGEMV dense matrix-vector product
(5) $\mathbf{f} \leftarrow \mathbf{w} - \mathbf{V}\mathbf{h};$	
(6) $\mathbf{H} \leftarrow [\mathbf{H}, \mathbf{h}]$	

Slide 11

The DGKS Correction

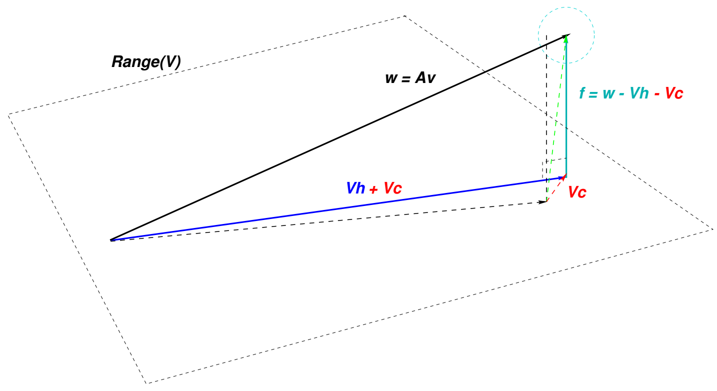

Slide 12

Approximate Eigenvalues and Vectors

$$\mathbf{AV} = \mathbf{VH} + \mathbf{f}\mathbf{e}_k^T$$

If $\quad \mathbf{H}\mathbf{y} = \mathbf{y}\theta\ , \quad \mathbf{x} = \mathbf{V}\mathbf{y}$

Then $\quad \|\mathbf{A}\mathbf{x} - \mathbf{x}\theta\| = \|\mathbf{f}\| \cdot |\mathbf{e}_k^T y|$

Provides approximate Eigenpairs $(\mathbf{x}, \theta)$

Ritz vectors and *Ritz values.*

Error Estimate $\quad \|\mathbf{f}\| \cdot |\mathbf{e}_k^T y|$

Slide 13

Arnoldi – Ritz Values

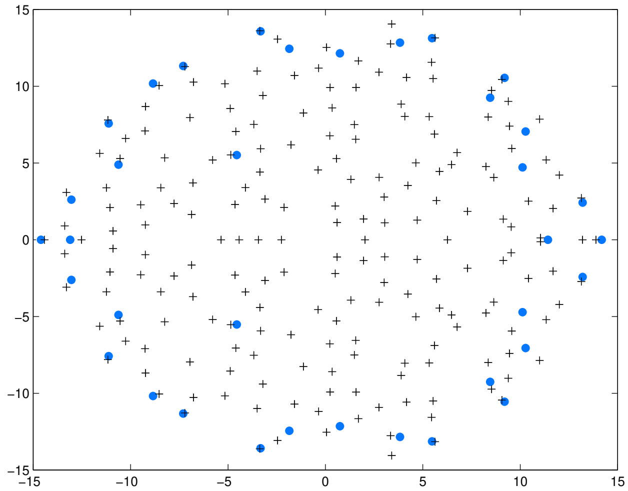

Slide 14

Arnoldi – Converging Ritz Values

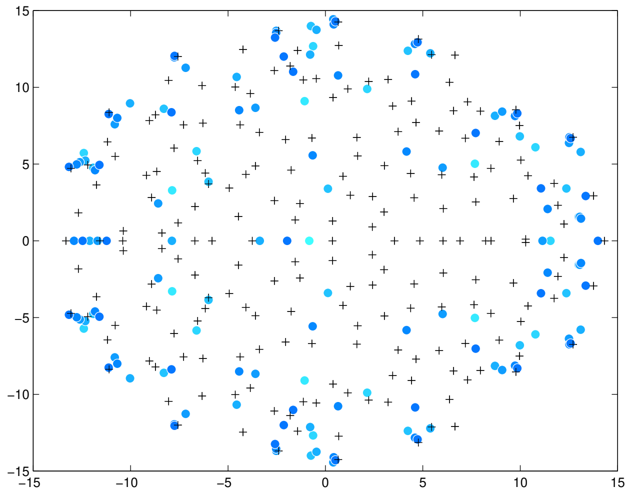

Slide 15

Traditional Methods Let k Grow ... Until

Problems:

Large k before "Wanted" λ appear

Cost of Orthogonalization is $O(nk^2)$

Cost to Examine Eigenvalues of $\mathbf{H}$ is $O(k^3)$, ($O(k^2)$ for symm.)

Cost of Storage is $O(nk)$

Slide 16

Main Idea – Restart

To Get $f(\mathbf{v}_1) = 0$

Force $\mathbf{v}_1 = \mathbf{q}_1\gamma_1 + \cdots + \mathbf{q}_k\gamma_k$, $\mathbf{A}\mathbf{q}_j = \mathbf{q}\lambda_j$

$$\mathbf{v}_1 \leftarrow \psi(\mathbf{A})\mathbf{v}_1$$

Slide 17

Main Idea – Restart

$$\mathbf{v}_1 \leftarrow \psi(\mathbf{A})\mathbf{v}_1 = \underbrace{\sum_{j=1}^{k} \mathbf{q}_j \psi(\lambda_j)\gamma_j}_{\text{Wanted}} + \underbrace{\sum_{j=k+1}^{n} \mathbf{q}_j \psi(\lambda_j)\gamma_j}_{\text{Unwanted}}$$

$\psi(\lambda) \rightsquigarrow$

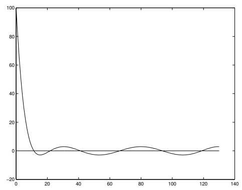

Slide 18

Arnoldi + Implicitly Shifted Q-R

<u>Update the Arnoldi Factorization</u>

$$\mathbf{AVQ} - \mathbf{VQ}(\mathbf{Q}^T\mathbf{HQ}) = \underline{\mathbf{f}\mathbf{e}_{k+p}^T\mathbf{Q}}$$

$$\mathbf{AV}_+ - \mathbf{V}_+\mathbf{H}_+ = \mathbf{f}\mathbf{q}_{k+p}^T$$

Do Double Implicit Shift for Complex Conjugate Pair

Result After p - Shifts

$$\mathbf{v}_1^+ = \mathbf{V}_+\mathbf{e}_1 = \psi(\mathbf{A})\mathbf{v}_1$$

$$\psi(\lambda) = \tau \prod_{j=1}^{p} (\lambda - \mu_j)$$

Slide 19

Implicitly Restarted Arnoldi

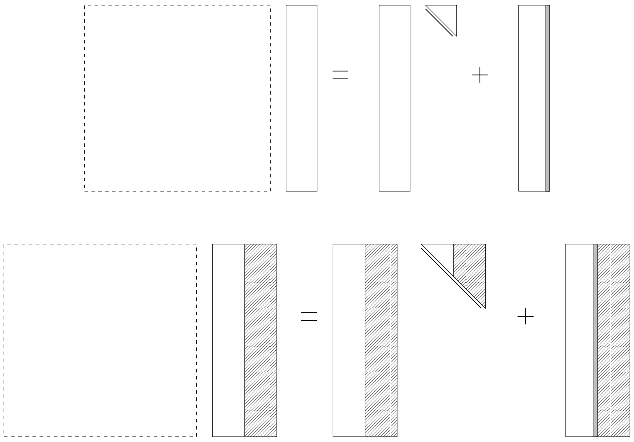

Slide 20

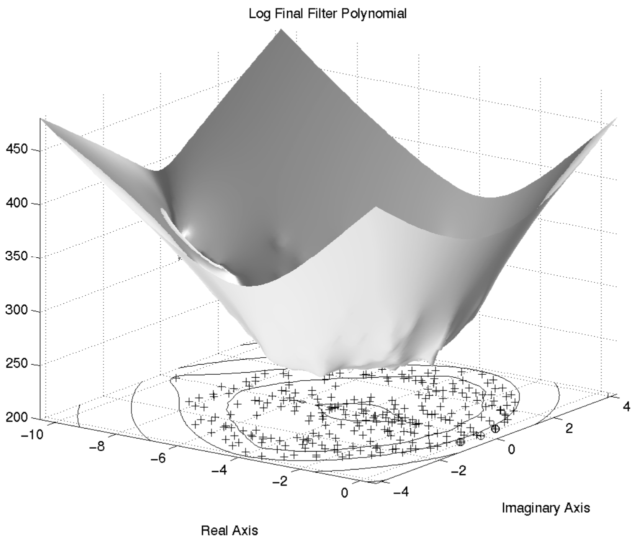

IRAM – Converging Ritz Values

Slide 21

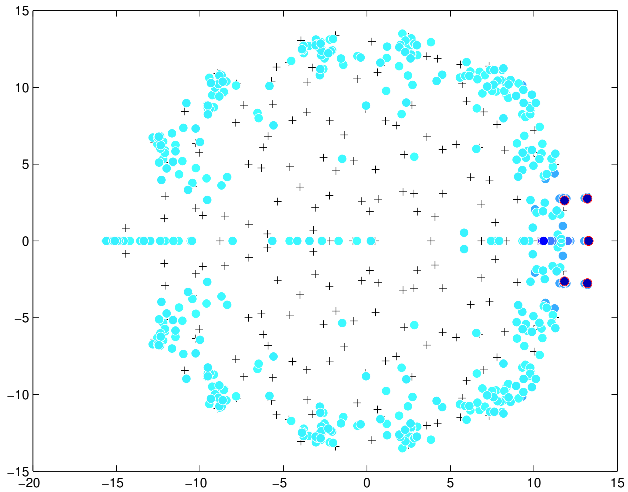

Spectral Transformation – Shift and Invert

Slide 22

$$\mathbf{A}\mathbf{q} = \mathbf{q}\lambda \iff (\mathbf{A} - \sigma\mathbf{I})^{-1}\mathbf{q} = \mathbf{q}(\frac{1}{\lambda - \sigma})$$

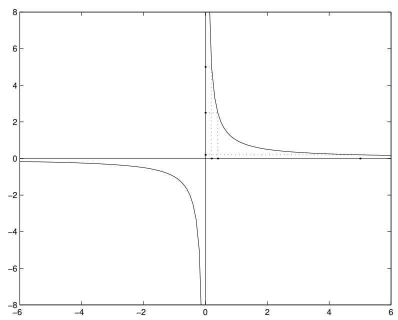

Eigenvalues near σ $\quad\rightarrow$ Extremes of $(\mathbf{A} - \sigma\mathbf{I})^{-1}$

Slide 23

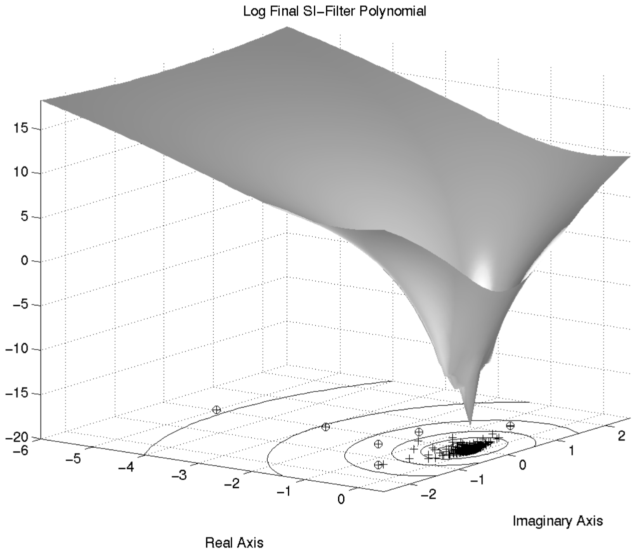

Slide 24

ARPACK Software

$$\mathbf{Ax} = \mathbf{Bx}\lambda$$

Slide 25

ARPACK and P_ARPACK

R. Lehoucq, D.C. Sorensen, P.A. Vu, C. Yang
K. Maschhoff - Parallel

ARPACK is a package of Fortran77 subroutines which implement the Implicitly Restarted Arnoldi Method for solving large sparse eigenvalue problems.

- Used to compute a Few (k) Eigenvalues - Specified Properties
- Symmetric, Non-symmetric, Standard or Generalized
- Fixed Storage Requirement O(nk)
- Shift Invert Mode, Band Matrix, and SVD Templates
- Reverse Communication Interface
- Users Guide – SIAM

Slide 26

ARPACK and P_ARPACK

Builds on Existing Lin. Alg. Software and MP Interfaces

BLAS , LAPACK, BLACS, MPI

Distributed on Request

http://www.caam.rice.edu/software/ARPACK

Also available from Netlib@ornl.gov under ScaLAPACK

Slide 27

Reverse Communication

```
10  continue
    call snaupd (ido, bmat, n, which, nev, tol,
   &             resid, ncv, v, ldv, iparam, ipntr,
   &             workd, workl, lworkl, info)

    if (ido .eq. newprod) then

        call matvec ('A', n, workd(ipntr(1)),
   &                           workd(ipntr(2)))
     else
        return
     endif
     go to 10
```

- Avoid Passing **A** Through Calling Sequence or COMMON

Slide 28

Shift-Invert Modes $\mathbf{Ax} = \mathbf{Bx}\lambda$

- Regular Mode

$$\mathbf{B}^{-1}\mathbf{Ax} = \mathbf{x}\lambda$$

- Inverse Mode

$$\mathbf{A}^{-1}\mathbf{Bx} = \mathbf{x}\frac{1}{\lambda}$$

- Shift Inverse Mode

$$(\mathbf{A} - \sigma\mathbf{B})^{-1}\mathbf{Bx} = \mathbf{x}\nu \qquad \lambda = \sigma + 1/\nu$$

$$< \mathbf{x}, \mathbf{y} >= \mathbf{x}^T\mathbf{By}$$

Slide 29

CFD – Linear Stability Analysis - Fluid Between Rotating Plates

P. Lequere and O. Daube LIMSI-CNRS, France

First, the flow is symmetric with respect to the mid plane
Fork bifurcation occurs at Re=110 $\Rightarrow$ assymetric state.

Supercitical Hopf bifurcation at Re $\approx$ 2200.

ARPACK computes 6 most unstable eigenvalues and vectors
Sequence of highly non-symmetric problems of 60,000 variables
Typical Problem – 229 matrix-vector products
$\|Ax - x\lambda\| \approx 10^{-6}$

Slide 30

```
=============================================
= Nonsymmetric implicit Arnoldi update code =
= Version Number:  2.1                      =
= Version Date:    8/31/95                  =
=============================================
= Summary of timing statistics              =
=============================================

Total number update iterations            =    44
Total number of OP*x operations           =   229
Total time in user OP*x operation         =   262.41
Total time in snaup2 routine              =   267.31
Total time in reorthogonalization phase   =     0.27
Total time in applying the shifts         =     0.77
```

Slide 31

Semiconductor Laser Design

L. Cowsar   ATT Bell Labs. (Now Lucent)

The eigenvalue problems arise in computing the lasing frequencies of semiconductor lasers

ARPACK computes the first 3 or 4 eigenvalues and vectors
Problems involving 500,000 to 1,000,000 unknowns.

The linear problem comes from a symmetric elliptic eigenvalue problem discretized by hybrid mixed finite element methods on a locally refined, solution adapted octree mesh.

Slide 32

P_ARPACK
Portable Parallel Implementation of ARPACK

for

Distributed Memory Parallel

Architectures

Kristi Maschhoff

Now at Tera Computers

Slide 33

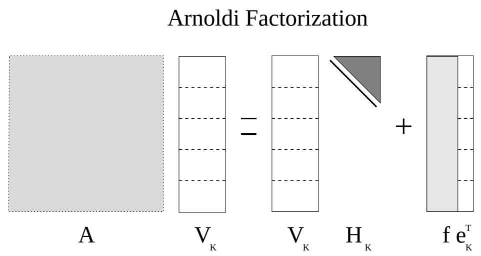

$\mathbf{H}_k$ replicated on every processor

$\mathbf{V}_k$ is distributed across a 1-D processor grid

$\mathbf{f}_k$ and workspace distributed accordingly

Slide 34

P_ARPACK - Platforms Used (So Far)

- SUN4 (SPARC 20)
  - MPI - MPICH 1.1.10
  - MPI-BLACS - MPICH 1.1.10
- SGI Power Challenge Cluster - Rice University
  - PVM-BLACS
  - MPI
- IBM SP2 - Maui HPCC
  - MPL-BLACS
  - MPI - MPICH 1.0.10
  - MPI-BLACS - MPICH 1.0.10

Slide 35

- CRAY T3D - Ohio Supercomputer Center
  - MPI
  - BLACS
- Intel Paragon - Sandia Nat. Labs
  - MPI
  - NX-BLACS
- ASCII Red - Sandia Nat. Labs
  - MPI

Slide 36

Parallel Performance - Non-symmetric Code

Nodes	Problem Size	Total Time (s)	Efficiency
1	100,000 * 1	40.53	
4	100,000 * 4	40.97	0.98
8	100,000 * 8	42.48	0.95
12	100,000 * 12	42.53	0.95
16	100,000 * 16	42.13	0.96
32	100,000 * 32	46.59	0.87
64	100,000 * 64	54.47	0.74
128	100,000 * 128	57.69	0.70

Internal Scalability of Parallel ARPACK

Parallel Performance - Non-symmetric Code

Slide 37

Platform: IBM SP2 - MPL-BLACS
Maui HPCC

No Communication for MV Product

Objective of Test:
Show that provided a user supplies an efficient Matrix-Vector product, the internal performance of ARPACK will not be a detriment to the overall performance and scalability.

Applications Using P_ARPACK

Slide 38

Iterative Determination of Vibrational Energy Levels of Four-Atom Molecules

Rich Lehoucq and Steve Gray (Argonne National Laboratory)

Slide 39

Three Experiments on the IBM SP at Argonne - HOCO

	Experiment 1	Experiment 2	Experiment 3
N	1,092,420	1,638,630	2,287,044
Nodes	54	54	54
ARPACK			
Total time	5.5 hours	7.9 hours	11 hours
NEV	41	52	52
MV Prod	733	903	894
Time used			
by MV prod	98%	97%	97%
NCV	71	81	81
Restarts	39	50	49
Tol	10^{-5}	10^{-7}	10^{-7}

P_ARPACK internal costs = 2% - 3% of the total computation time.

Slide 40

Structural Analysis – SANDIA

K. Alvin, M. Bhardwaj, D. Day, G. Reese, J. Red-Horse

Used P_ARPACK on ASCII Red Machine
9632 Xeon Pentium II 333 MHz processors.

They used 4054/2 (1 proc. + 1 co-proc for memory 256MB)
Peak rate: 2027 *333 Mflop/s
Kernal Runs at about 10% of peak.

- W76 - 3 Million variable FEM Structural Model
  200 smallest eigenvalues
- Used inverse mode with FETI (C. Farhat)
- ARPACK cost = 5% of total time

CVD Reactor Stability – SANDIA

R. Lehoucq, A. Salinger

Used P_ARPACK on Sandia-Intel Tflop
1024 processors .

Slide 41

- 4 Million variable FEM 3D Model of CVD Reactor
  Computed 6 eigenvalues of largest real part

- Used AZTEC to solve linear systems
  Cayley Transformation $(\mathbf{A} - \sigma_1 \mathbf{B})\mathbf{w} = (\mathbf{A} + \sigma_2 \mathbf{B})\mathbf{v}$
  ILUT preconditioner with GMRES
  2-step CGS required for consistent convergence.

CVD Reactor

Slide 42

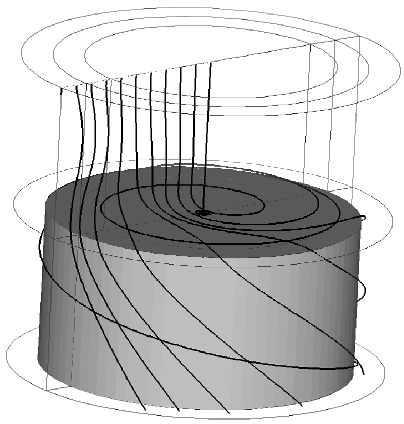

Slide 43

CVD Bifurcation Analysis

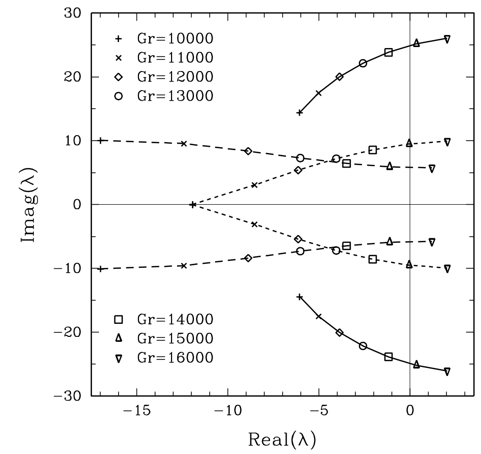

Figure 8: The tracking of the six largest eigenvalues as a function of parameter indicate a Hopf bifurcation near 14800.

Slide 44

2-step CGS vs MGS

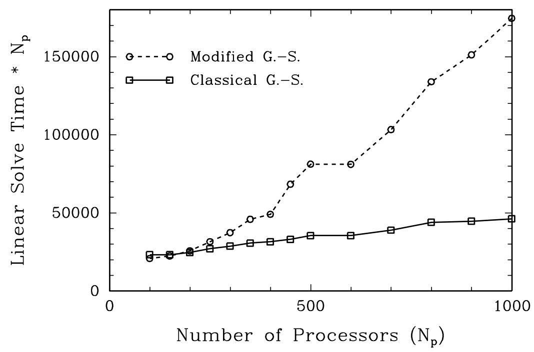

Figure 7: The parallel efficiency of two stable orthogonalization schemes in the GMRES algorithm is compared. The extra communications of the Modified Gram-Schmidt (MGS) approach scale poorly with number of processors, while the extra operations of the 2-step Classical Gram-Schmidt (CGS) approach scale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

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 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 cross-over point (at around 175 processors for this case) to occur at fewer processors.

Slide 45

Current and Future Research Directions

$$\boxed{\mathbf{Ax} = \mathbf{Bx}\lambda}$$

Pre-conditioning - Spectral Enhancement

Inexact Solves

Slide 46

Truncated RQ

Chao Yang and D.C.S.

QR vs **RQ**

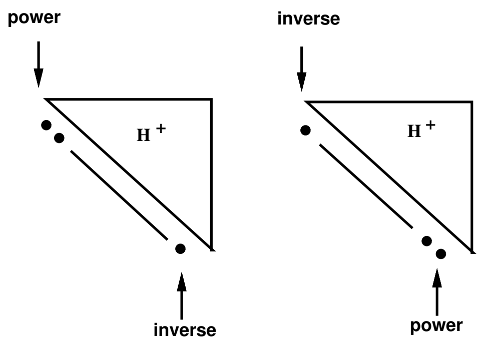

Slide 47

Truncating the RQ–iteration

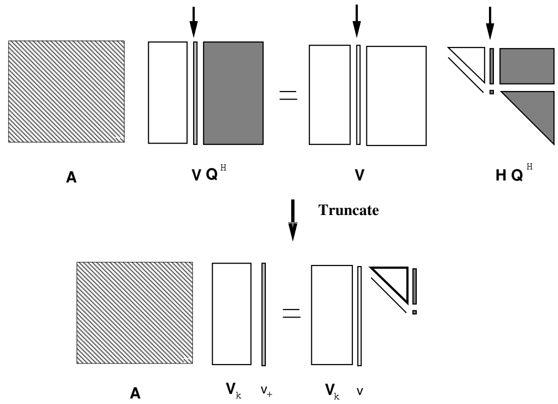

Slide 48

Effect of Pre-Conditioning

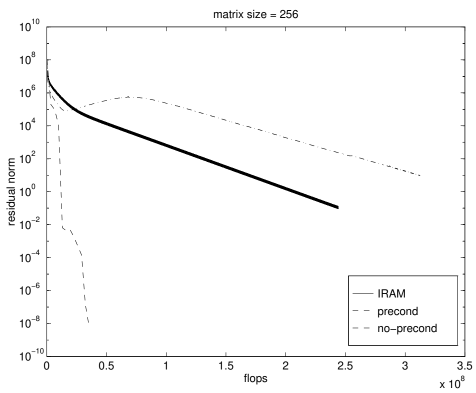

Jacobi - Davidson

Slide 49

- Sleijpen, Van der Vorst

1. $\mathbf{Hy} = \mathbf{y}\theta$; $\mathbf{x} = \mathbf{Vy}$;
   $\mathbf{r} \leftarrow \mathbf{Ax} - \mathbf{x}\theta$
2. "Solve"$\mathcal{P}(\mathbf{A} - \theta\mathbf{I})\mathcal{P}\mathbf{w} = \mathbf{r}$, $\mathcal{P} = (\mathbf{I} - \mathbf{xx}^T)$
   Pre-conditioned iterative method;
3. Orthogonalize $\mathbf{w}$ against $\mathbf{V}$
   to get new direction v;
4. $\mathbf{V} \leftarrow (\mathbf{V}, \mathbf{v})$ and $\mathbf{H} \leftarrow \mathbf{V}^T\mathbf{AV}$;

Rational Krylov

Slide 50

- A. Ruhe

1. $\mathbf{AVH} = \mathbf{VK}$;
2. Solve $(\mathbf{A} - \mu\mathbf{I})\mathbf{w} = \mathbf{Vt}$,
   Sparse Direct Method;
3. Orthogonalize w against $\mathbf{V}$
   get new direction $\mathbf{v}$,
   Arnoldi-like formula;
4. $\mathbf{V} \leftarrow (\mathbf{V}, \mathbf{v})$; $\mathbf{H} \leftarrow (\mathbf{H}, \mathbf{h})$; $\mathbf{K} \leftarrow (\mathbf{K}, \mathbf{k})$;

Slide 51

Extensions of IRAM

- Regularization of Discrete Ill-Posed Problems
  $min\|\mathbf{b} - \mathbf{Ax}\| \quad \text{s.t.} \quad \|\mathbf{x}\| \leq \Delta$
- Model Reduction of LTI state space systems
  $\dot{\mathbf{x}} = \mathbf{Ax} + \mathbf{Bu}, \quad \mathbf{y} = \mathbf{Cx}$
- Both use specialized eigenvalue calculations

Slide 52

Summary

- ARPACK – widely used eigenvalue software
  Many Applications
  http://www.caam.rice.edu/software/ARPACK
- Parallel computation effective on very large scale problems
  3-D stability analysis possible
- Implicit Restarting – Extend to new areas
  Examples: Regularization , Model Reduction