

Parallel Numerical Algorithms

Chapter 12 – Eigenvalue Problems

Prof. Michael T. Heath

Department of Computer Science
University of Illinois at Urbana-Champaign

CSE 512 / CS 554



Outline

- 1 Basics
- 2 Power Iteration
- 3 QR Iteration
- 4 Krylov Methods
- 5 Other Methods



Eigenvalues and Eigenvectors

- Given $n \times n$ matrix A , find scalar λ and nonzero vector x such that

$$Ax = \lambda x$$

- λ is *eigenvalue* and x is corresponding *eigenvector*
- A always has n eigenvalues, but they may be neither real nor distinct
- May need to compute only one or few eigenvalues, or all n eigenvalues
- May or may not need corresponding eigenvectors



Problem Transformations

- **Shift**: for scalar σ , eigenvalues of $\mathbf{A} - \sigma\mathbf{I}$ are eigenvalues of \mathbf{A} shifted by σ , $\lambda_i - \sigma$
- **Inversion**: for nonsingular \mathbf{A} , eigenvalues of \mathbf{A}^{-1} are reciprocals of eigenvalues of \mathbf{A} , $1/\lambda_i$
- **Powers**: for integer $k > 0$, eigenvalues of \mathbf{A}^k are k th powers of eigenvalues of \mathbf{A} , λ_i^k
- **Polynomial**: for polynomial $p(t)$, eigenvalues of $p(\mathbf{A})$ are values of p evaluated at eigenvalues of \mathbf{A} , $p(\lambda_i)$



Similarity Transformations

- B is *similar* to A if there is nonsingular T such that

$$B = T^{-1}AT$$

- Then

$$By = \lambda y \Rightarrow T^{-1}ATy = \lambda y \Rightarrow A(Ty) = \lambda(Ty)$$

so A and B have same eigenvalues, and if y is eigenvector of B , then $x = Ty$ is eigenvector of A

- Similarity transformations preserve eigenvalues, and eigenvectors are easily recovered



Similarity Transformations

Forms attainable by similarity transformation

A	T	B
distinct eigenvalues	nonsingular	diagonal
real symmetric	orthogonal	real diagonal
complex Hermitian	unitary	real diagonal
normal	unitary	diagonal
arbitrary real	orthogonal	real block triangular (real Schur)
arbitrary	unitary	upper triangular (Schur)
arbitrary	nonsingular	almost diagonal (Jordan)



Preliminary Reduction

- Eigenvalues easier to compute if matrix first reduced to simpler form by similarity transformation
- Diagonal or triangular most desirable, but cannot always be reached in finite number of steps
- Preliminary reduction usually to tridiagonal form (for symmetric matrix) or Hessenberg form (for nonsymmetric matrix)
- Preliminary reduction usually done by orthogonal transformations, using algorithms similar to QR factorization, but transformations must be applied from both sides to maintain similarity



Parallel Algorithms for Eigenvalues

- Algorithms for computing eigenvalues and eigenvectors employ basic operations such as
 - vector updates ($saxpy$)
 - inner products
 - matrix-vector and matrix-matrix multiplication
 - solution of triangular systems
 - orthogonal (QR) factorization
- In many cases, parallel implementations will be based on parallel algorithms we have already seen for these basic operations, although there will sometimes be new sources of parallelism



Power Iteration

- Simplest method for computing one eigenvalue-eigenvector pair is *power iteration*, which in effect takes successively higher powers of matrix times initial starting vector

- $x_0 =$ arbitrary nonzero vector

for $k = 1, 2, \dots$

$$y_k = Ax_{k-1}$$

$$x_k = y_k / \|y_k\|_\infty$$

end

- If A has unique eigenvalue λ_1 of maximum modulus, then power iteration converges to eigenvector corresponding to dominant eigenvalue



Power Iteration

- Convergence rate of power iteration depends on ratio $|\lambda_2/\lambda_1|$, where λ_2 is eigenvalue having second largest modulus
- It may be possible to choose shift, $\mathbf{A} - \sigma\mathbf{I}$, so that ratio is more favorable and yields more rapid convergence
- Shift must then be added to result to obtain eigenvalue of original matrix



Parallel Power Iteration

- Power iteration requires repeated matrix-vector products, which are easily implemented in parallel for dense or sparse matrix, as we have seen
- Additional communication may be required for normalization, shifts, convergence test, etc.
- Though easily parallelized, power iteration is often too slow to be useful in this form



Inverse Iteration

- *Inverse iteration* is power iteration applied to A^{-1} , which converges to eigenvector corresponding to dominant eigenvalue of A^{-1} , which is reciprocal of *smallest* eigenvalue of A
- Inverse of A is not computed explicitly, but only factorization of A (and only once) to solve system of linear equations at each iteration
- $x_0 =$ arbitrary nonzero vector
for $k = 1, 2, \dots$
 Solve $Ay_k = x_{k-1}$ for y_k
 $x_k = y_k / \|y_k\|_\infty$
end



Inverse Iteration

- Shifting strategy can greatly accelerate convergence
- Inverse iteration is especially useful for computing eigenvector corresponding to approximate eigenvalue, since it converges rapidly when approximate eigenvalue is used as shift
- Inverse iteration also useful for computing eigenvalue closest to any given value β , since if β is used as shift, then desired eigenvalue corresponds to smallest eigenvalue of shifted matrix



Parallel Inverse Iteration

- Inverse iteration requires initial factorization of matrix A and solution of triangular systems at each iteration, so it appears to be much less amenable to efficient parallel implementation than power iteration
- However, inverse iteration is often used to compute eigenvector in situations where
 - approximate eigenvalue is already known, so using it as shift yields very rapid convergence
 - matrix has previously been reduced to simpler form (e.g., tridiagonal) for which linear system is easy to solve



Simultaneous Iteration

- To compute many eigenvalue-eigenvector pairs, could apply power iteration to several starting vectors simultaneously, giving *simultaneous iteration*

$\mathbf{X}_0 =$ arbitrary $n \times q$ matrix of rank q

for $k = 1, 2, \dots$

$$\mathbf{X}_k = \mathbf{A}\mathbf{X}_{k-1}$$

end

- $\text{span}(\mathbf{X}_k)$ converges to invariant subspace determined by q largest eigenvalues of \mathbf{A} , provided $|\lambda_q| > |\lambda_{q+1}|$
- Normalization is needed at each iteration, and columns of \mathbf{X}_k become increasingly ill-conditioned basis for $\text{span}(\mathbf{X}_k)$



Orthogonal Iteration

- Both issues addressed by computing QR factorization at each iteration, giving *orthogonal iteration*

$\mathbf{X}_0 =$ arbitrary $n \times q$ matrix of rank q

for $k = 1, 2, \dots$

 Compute reduced QR factorization

$$\hat{\mathbf{Q}}_k \mathbf{R}_k = \mathbf{X}_{k-1}$$

$$\mathbf{X}_k = \mathbf{A} \hat{\mathbf{Q}}_k$$

end

- Converges to block triangular form, and blocks are triangular where moduli of consecutive eigenvalues are distinct



Parallel Orthogonal Iteration

- Orthogonal iteration requires matrix-matrix multiplication and QR factorization at each iteration, both of which we know how to implement in parallel with reasonable efficiency



QR Iteration

- If we take $X_0 = I$, then orthogonal iteration should produce *all* eigenvalues and eigenvectors of A
- Orthogonal iteration can be reorganized to avoid explicit formation and factorization of matrices X_k
- Instead, sequence of unitarily similar matrices is generated by computing QR factorization at each iteration and then forming reverse product, giving *QR iteration*

$$A_0 = A$$

for $k = 1, 2, \dots$

 Compute QR factorization

$$Q_k R_k = A_{k-1}$$

$$A_k = R_k Q_k$$

end



QR Iteration

- In simple form just given, each iteration of QR method requires $\Theta(n^3)$ work
- Work per iteration is reduced to $\Theta(n^2)$ if matrix is in Hessenberg form, or $\Theta(n)$ if symmetric matrix is in tridiagonal form
- Preliminary reduction is usually done by Householder or Givens transformations
- In addition, number of iterations required is reduced by preliminary reduction of matrix
- Convergence rate also enhanced by judicious choice of shifts



Parallel QR Iteration

- Preliminary reduction can be implemented efficiently in parallel, using algorithms analogous to parallel QR factorization for dense matrix
- But subsequent QR iteration for reduced matrix is inherently essentially serial, and yields little parallel speedup for this portion of algorithm
- This may not be of great concern if iterative phase is relatively small portion of total time, but it does limit efficiency and scalability



Krylov Subspace Methods

- *Krylov subspace* methods reduce matrix to Hessenberg (or tridiagonal) form using only matrix-vector multiplication
- For arbitrary starting vector x_0 , if

$$\mathbf{K}_k = [\mathbf{x}_0 \quad \mathbf{A}\mathbf{x}_0 \quad \cdots \quad \mathbf{A}^{k-1}\mathbf{x}_0]$$

then

$$\mathbf{K}_n^{-1}\mathbf{A}\mathbf{K}_n = \mathbf{C}_n$$

where \mathbf{C}_n is upper Hessenberg (in fact, companion matrix)



Krylov Subspace Methods

- To obtain better conditioned basis for $\text{span}(\mathbf{K}_n)$, compute QR factorization

$$\mathbf{Q}_n \mathbf{R}_n = \mathbf{K}_n$$

so that

$$\mathbf{Q}_n^H \mathbf{A} \mathbf{Q}_n = \mathbf{R}_n \mathbf{C}_n \mathbf{R}_n^{-1} \equiv \mathbf{H}$$

with \mathbf{H} upper Hessenberg



Krylov Subspace Methods

- Equating k th columns on each side of equation $AQ_n = Q_nH$ yields recurrence

$$Aq_k = h_{1k}q_1 + \cdots + h_{kk}q_k + h_{k+1,k}q_{k+1}$$

relating q_{k+1} to preceding vectors q_1, \dots, q_k

- Premultiplying by q_j^H and using orthonormality,

$$h_{jk} = q_j^H Aq_k, \quad j = 1, \dots, k$$

- These relationships yield *Arnoldi iteration*, which produces upper Hessenberg matrix column by column using only matrix-vector multiplication by A and inner products of vectors



Arnoldi Iteration

$\mathbf{x}_0 =$ arbitrary nonzero starting vector

$\mathbf{q}_1 = \mathbf{x}_0 / \|\mathbf{x}_0\|_2$

for $k = 1, 2, \dots$

$\mathbf{u}_k = \mathbf{A}\mathbf{q}_k$

for $j = 1$ **to** k

$h_{jk} = \mathbf{q}_j^H \mathbf{u}_k$

$\mathbf{u}_k = \mathbf{u}_k - h_{jk}\mathbf{q}_j$

end

$h_{k+1,k} = \|\mathbf{u}_k\|_2$

if $h_{k+1,k} = 0$ **then** stop

$\mathbf{q}_{k+1} = \mathbf{u}_k / h_{k+1,k}$

end



Arnoldi Iteration

- If

$$Q_k = [q_1 \quad \cdots \quad q_k],$$

then

$$H_k = Q_k^H A Q_k$$

is upper Hessenberg matrix

- Eigenvalues of H_k , called *Ritz values*, are approximate eigenvalues of A , and *Ritz vectors* given by $Q_k y$, where y is eigenvector of H_k , are corresponding approximate eigenvectors of A
- Eigenvalues of H_k must be computed by another method, such as QR iteration, but this is easier problem if $k \ll n$



Arnoldi Iteration

- Arnoldi iteration expensive in work and storage because each new vector q_k must be orthogonalized against all previous columns of Q_k , which must be stored
- So Arnoldi process usually restarted periodically with carefully chosen starting vector
- Ritz values and vectors produced are often good approximations to eigenvalues and eigenvectors of A after relatively few iterations
- Work and storage costs drop dramatically if matrix symmetric or Hermitian, since recurrence then has only three terms and H_k is tridiagonal (so usually denoted T_k), yielding *Lanczos iteration*



Lanczos Iteration

$$\mathbf{q}_0 = \mathbf{0}$$

$$\beta_0 = 0$$

$\mathbf{x}_0 =$ arbitrary nonzero starting vector

$$\mathbf{q}_1 = \mathbf{x}_0 / \|\mathbf{x}_0\|_2$$

for $k = 1, 2, \dots$

$$\mathbf{u}_k = \mathbf{A}\mathbf{q}_k$$

$$\alpha_k = \mathbf{q}_k^H \mathbf{u}_k$$

$$\mathbf{u}_k = \mathbf{u}_k - \beta_{k-1} \mathbf{q}_{k-1} - \alpha_k \mathbf{q}_k$$

$$\beta_k = \|\mathbf{u}_k\|_2$$

if $\beta_k = 0$ **then stop**

$$\mathbf{q}_{k+1} = \mathbf{u}_k / \beta_k$$

end



Lanczos Iteration

- α_k and β_k are diagonal and subdiagonal entries of symmetric tridiagonal matrix \mathbf{T}_k
- As with Arnoldi, Lanczos iteration does not produce eigenvalues and eigenvectors directly, but only tridiagonal matrix \mathbf{T}_k , whose eigenvalues and eigenvectors must be computed by another method to obtain Ritz values and vectors
- If $\beta_k = 0$, then algorithm appears to break down, but in that case invariant subspace has already been identified (i.e., Ritz values and vectors are already exact at that point)



Lanczos Iteration

- In principle, if Lanczos algorithm is run until $k = n$, resulting tridiagonal matrix is orthogonally similar to A
- In practice, rounding error causes loss of orthogonality, invalidating this expectation
- Problem can be overcome by reorthogonalizing vectors as needed, but expense can be substantial
- Alternatively, can ignore problem, in which case algorithm still produces good eigenvalue approximations, but multiple copies of some eigenvalues may be generated



Krylov Subspace Methods

- Virtue of Arnoldi and Lanczos iterations is ability to produce good approximations to extreme eigenvalues for $k \ll n$
- Moreover, they require only one matrix-vector multiplication by A per step and little auxiliary storage, so are ideally suited to large sparse matrices
- If eigenvalues are needed in middle of spectrum, say near σ , then algorithm can be applied to matrix $(A - \sigma I)^{-1}$, assuming it is practical to solve systems of form $(A - \sigma I)x = y$



Parallel Krylov Subspace Methods

- Krylov subspace methods composed of
 - vector updates (`saxpy`)
 - inner products
 - matrix-vector multiplication
 - computing eigenvalues/eigenvectors of tridiagonal matrices
- Parallel implementation requires implementing each of these in parallel as before
- For early iterations, Hessenberg or tridiagonal matrices generated are too small to benefit from parallel implementation, but Ritz values and vectors need not be computed until later



Jacobi Method

- *Jacobi method* for symmetric matrix starts with $A_0 = A$ and computes sequence

$$A_{k+1} = J_k^T A_k J_k$$

where J_k is plane rotation that annihilates symmetric pair of off-diagonal entries in A_k

- Plane rotations are applied repeatedly from both sides in systematic sweeps through matrix until magnitudes of all off-diagonal entries are reduced below tolerance
- Resulting diagonal matrix is orthogonally similar to original matrix, so diagonal entries are eigenvalues, and eigenvectors given by product of plane rotations



Parallel Jacobi Method

- Jacobi method, though slower than QR iteration serially, parallelizes better
- Parallel implementation of Jacobi method performs many annihilations simultaneously, at locations chosen so that rotations do not interfere with each other (analogous to parallel Givens QR factorization)
- Computations are still rather fine grained, however, so effectiveness is limited on parallel computers with unfavorable ratio of communication to computation speed



Bisection or Spectrum-Slicing

- For real symmetric matrix, can determine how many eigenvalues are less than given real number σ
- By systematically choosing various values for σ (*slicing spectrum* at σ) and monitoring resulting count, any eigenvalue can be isolated as accurately as desired
- For example, by computing inertia using $A = LDL^T$ factorization of $A - \sigma I$ for various σ , individual eigenvalues can be isolated as accurately as desired using interval bisection technique



Sturm Sequence

- Another spectrum-slicing method for computing individual eigenvalues is based on *Sturm sequence* property of symmetric matrices
- Let $p_r(\sigma)$ denote determinant of leading principal minor of $A - \sigma I$ of order r
- Zeros of $p_r(\sigma)$ strictly separate those of $p_{r-1}(\sigma)$, and number of agreements in sign of successive members of sequence $p_r(\sigma)$, for $r = 1, \dots, n$, gives number of eigenvalues of A strictly greater than σ
- Determinants $p_r(\sigma)$ easy to compute if A transformed to tridiagonal form before applying Sturm sequence technique



Parallel Multisection

- Spectrum-slicing methods can be implemented in parallel by assigning disjoint intervals of real line to different tasks, and then each task carries out bisection process using serial spectrum-slicing method for its subinterval
- Both for efficiency of spectrum-slicing technique and for storage scalability, matrix is first reduced (in parallel) to tridiagonal form, so that each task can store entire (reduced) matrix
- Load imbalance is potential problem, since different subintervals may contain different numbers of eigenvalues



Parallel Multisection

- Clustering of eigenvalues cannot be anticipated in advance, so dynamic load balancing may be required to achieve reasonable efficiency
- Eigenvectors must still be determined, usually by inverse iteration
- Since each task holds entire (tridiagonal) matrix, each can in principle carry out inverse iteration for its eigenvalues without any communication among tasks
- Orthogonality of resulting eigenvectors cannot be guaranteed, however, and thus communication is required to orthogonalize them



Divide-and-Conquer Method

- Another method for computing eigenvalues and eigenvectors of real symmetric tridiagonal matrix is based on *divide-and-conquer*
- Express symmetric tridiagonal matrix T as

$$T = \begin{bmatrix} T_1 & O \\ O & T_2 \end{bmatrix} + \beta uu^T$$

- Can now compute eigenvalues and eigenvectors of smaller matrices T_1 and T_2
- To relate these back to eigenvalues and eigenvectors of original matrix requires solution of *secular equation*, which can be done reliably and efficiently



Divide-and-Conquer Method

- Applying this approach recursively yields divide-and-conquer algorithm that is naturally parallel
- Parallelism in solving secular equations grows as parallelism in processing independent tridiagonal matrices shrinks, and vice versa
- Algorithm is complicated to implement and difficult questions of numerical stability, eigenvector orthogonality, and load balancing must be addressed



References

- Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, eds., *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*, SIAM, 2000
- J. W. Demmel, M. T. Heath, and H. A. van der Vorst, Parallel numerical linear algebra, *Acta Numerica* 2:111-197, 1993
- P. J. Eberlein and H. Park, Efficient implementation of Jacobi algorithms and Jacobi sets on distributed memory architectures, *J. Parallel Distrib. Comput.*, 8:358-366, 1990
- G. W. Stewart, A Jacobi-like algorithm for computing the Schur decomposition of a non-Hermitian matrix, *SIAM J. Sci. Stat. Comput.* 6:853-864, 1985
- G. W. Stewart, A parallel implementation of the QR algorithm, *Parallel Comput.* 5:187-196, 1987

