ON PARALLELIZABLE EIGENSOLVERS

Louis Auslander

Center for Large Scale Computation
City University of New York
25 W. 43rd St.
New York, NY 10036

and

Anna Tsao

Supercomputing Research Center
17100 Science Drive
Bowie, MD 20715-4300
email: tsao@super.org

Abstract. In this paper, we show how the intimate relationship between a certain matrix algebra and the ring of complex polynomials suggests a strategy for devising new parallel algorithms for the eigenproblem with a two-fold potential for parallelism. Specifically, we explore the algorithmic implications of the fact that an ability to compute nontrivial invariant subspaces of a matrix M results in an ability to perform block triangularization of M into two independent subproblems. We further propose a strategy for devising algorithms to compute nontrivial invariant subspaces of M which rely almost exclusively on two computational primitives: matrix multiplication and solving systems of linear equations. The future success of numerical computation must certainly depend on the availability of algorithms to perform these primitives fast on parallel architectures. The recursive application of the block triangularization scheme on the subproblems generated leads naturally to a divide and conquer strategy.

Appeared in Advances in Applied Mathematics, vol. 13, pp. 253-261, 1992.

Key words and phrases. invariant subspaces, eigenvalues, parallel algorithms.

1. Introduction

The numerical computation of all the eigenvalues and eigenvectors of a matrix arises naturally in many important applications. Large problems, which were once intractable, are now considered approachable with the advent of parallel computers. However, effective use of the computational power of these new machines will only result through improved algorithms, more suited to the architectural features of parallel computers.

In this paper, we show how the intimate relationship between a certain matrix algebra and the ring of complex polynomials suggests a strategy for devising new parallel algorithms for the eigenproblem. The ideas presented are an extension of the algorithm described by Auslander and Tsao in [1]. Specifically, the ideas presented in this paper are motivated by the algorithmic implications of the fact that an ability to compute nontrivial invariant subspaces of a matrix M results in an ability to perform block triangularization of M into two independent subproblems. We shall propose a strategy for devising algorithms to compute nontrivial invariant subspaces of M. The recursive application of such a block triangularization algorithm on the subproblems generated leads naturally to a divide and conquer strategy.

We shall demonstrate in Section 2 that invariant subspaces of a square matrix M can be computed if an algorithm for constructing rank deficient matrices in the algebra generated by M can be devised. Since, in fact, the algebra generated by M consists of all polynomials in M, our strategy for devising algorithms to compute these rank deficient matrices will rest on devising algorithms to construct appropriately chosen polynomials. We shall further demonstrate that any practical realization of the proposed strategy for computing invariant subspaces will lead to an algorithm with a two-fold potential for parallelism. First, any practical algorithm for computing the desired matrices will rely almost exclusively on two computational primitives: matrix multiplication and solving systems of linear equations. The future success of numerical computation must certainly depend on the ability to perform these two primitives fast on parallel architectures. The second notable attribute of the algorithm derives from the divide and conquer resulting from block triangularization. Practical considerations in devising numerical algorithms to implement the strategy described are discussed in Section 3.

We first introduce some notation which will be used throughout the paper. Matrices will be represented by capital letters. The sets of natural, real and complex numbers shall be denoted by $\mathbb{N}$, $\mathbb{R}$, and $\mathbb{C}$, respectively. Let $n \in \mathbb{N}$. Denote n-dimensional complex Euclidean space by $\mathbb{C}^n$, the vector space of $n \times n$ matrices over $\mathbb{C}$ by $\mathbb{C}^{n\times n}$, the multiplicative identity in $\mathbb{C}^{n\times n}$ by I_n, and the algebra of polynomials over $\mathbb{C}$ by $\mathbb{C}[x]$. We use analogous notation over $\mathbb{R}$. Let $M \in \mathbb{C}^{n\times n}$. The null space and range of M will be denoted by $\mathcal{N}(M)$ and $\mathcal{R}(M)$, respectively. We denote the spectrum and spectral radius of M by $\lambda(M)$ and $\rho(M)$; respectively. If $\mathcal{B}$ is a basis for $\mathbb{C}^n$, $[M]_{\mathcal{B}}$ will denote the matrix of M with respect to $\mathcal{B}$. Let W be an M-invariant subspace of $\mathbb{C}^n$. We denote by $M\big|_W$ the restriction of M to W.

2. A divide and conquer strategy for the complete eigenproblem

Throughout the remainder of the paper, we shall let M denote a matrix in $\mathbb{C}^{n\times n}$ having spectrum $\lambda(M) = \{\lambda_1, \ldots, \lambda_k\}$, where $2 \le k \le n$. Let the minimal polynomial of M be given by

$$m(x) = \prod_{i=1}^{k} f_i^{r_i}(x), \tag{1}$$

where $f_i(x) = (x - \lambda_i)$ and $d = \deg m(x) = \sum_{i=1}^{k} r_i$. By the Primary Decomposition Theorem for matrices [3, p. 220], the invariant subspaces $\mathcal{N}\big(f_i^{r_i}(M)\big)$, $i = 1, \ldots, k$; comprise a direct sum decomposition of $\mathbb{C}^n$. A complete eigensolution for M consists of finding $\lambda(M)$ and the subspaces $\mathcal{N}\big(f_i^{r_i}(M)\big)$, $i = 1, \ldots, k$. The set $\{I_n, M, M^2, \cdots, M^{d-1}\}$ generates a d-dimensional commutative matrix algebra with identity over $\mathbb{C}$, $\mathfrak{A}(M)$. We shall show how the algebraic properties of $\mathfrak{A}(M)$ suggest a method of computing the invariant subspaces of M. We give a self-contained treatment of the relevant results for the sake of clarity.

Consider the ring homomorphism $h : \mathbb{C}[x] \to \mathfrak{A}(M)$, defined in the following way: for any polynomial $f(x) = \sum_{i=0}^{N} c_i x^i \in \mathbb{C}[x]$,

$$h(f) = f(M) = \sum_{i=0}^{N} c_i M^i,$$

where $M^0 = I_n$. Exploitation of the fact that matrices in $\mathfrak{A}(M)$ are completely determined by polynomials in $\mathbb{C}[x]$ is the central theme of our proposed strategy for devising new eigensolvers. The correspondence between the operations of addition and multiplication in $\mathbb{C}[x]$ and $\mathfrak{A}(M)$ means that computation of matrices in $\mathfrak{A}(M)$ consists simply of polynomial arithmetic for matrices.

Suppose that $a \in \mathbb{C}[x]$ has the property that $a(M)$ has nonzero rank $r < n$. Then both $\mathcal{N}\big(a(M)\big)$ and $\mathcal{R}\big(a(M)\big)$ are invariant subspaces of M. Hence, invariant subspaces of M can be found by producing such a polynomial a. We say that a (or $a(M)$) is a *rank-r invariant subspace annihilator*. Let W denote a nontrivial invariant subspace of M and $\mathcal{B}_1$ be a basis for W. Then $\mathcal{B}_1$ can be extended by a set of vectors $\mathcal{B}_2$ to form a basis for $\mathbb{C}^n$. Denote by U_1 and U_2 the matrices having as columns the elements of $\mathcal{B}_1$ and $\mathcal{B}_2$, respectively. Since W is an invariant subspace of M,
the change of basis for M with respect to $\mathcal{B}_1 \cup \mathcal{B}_2$ is block triangular, i.e.,

$$[M]_{\mathcal{B}_1 \cup \mathcal{B}_2} = [\,U_1 \quad U_2\,]^{-1} M\,[\,U_1 \quad U_2\,] = \begin{bmatrix} M_1 & M_2 \\ 0 & M_3 \end{bmatrix},$$

where $M_1 = M\big|_W$ and the matrices M_1 and M_3 have dimensions r and $n-r$, respectively. Then

$$\lambda(M) = \lambda(M_1) \cup \lambda(M_3).$$

Continuing recursively in this manner, we can clearly produce a sequence of similarity transformations which block triangularize the matrix M. The process will conclude when each of the blocks has only a single, possibly repeated, eigenvalue. Note that the similarity transformations can easily be guaranteed to be unitary, which may be desirable for numerical reasons.

Assuming the existence of an algorithm for computing nontrivial invariant subspace annihilators, the divide and conquer strategy suggested above can be summarized as follows:

Invariant subspace decomposition algorithm (ISDA).

Suppose that M has at least two distinct eigenvalues.

I. *Compute $a(M)$, where $a \in \mathbb{C}[x]$ is a nontrivial invariant subspace annihilator.*
II. *Compute a basis $\mathcal{B}_1$ for $\mathcal{N}(a(M))$ (or $\mathcal{R}(a(M))$) and extend $\mathcal{B}_1$ to a basis $\mathcal{B}$ for $\mathbb{C}^n$.*
III. *Compute the block triangular matrix $[M]_{\mathcal{B}}$.*
IV. *Update the similarity transformations leading to the triangular decomposition, if eigenvectors are desired.*

The entire process is repeated for both subproblems. We continue recursively until each subproblem has only one eigenvalue, possibly repeated. Nontrivial subproblems resulting from defective Jordan blocks might then need to be dealt with via other means. Note that step IV need not be completed before proceeding to the next stage of the divide and conquer. The invariant subspace decomposition algorithm reduces a problem of size n into two smaller subproblems of sizes r and $(n-r)$. Hence i iterations of the ISDA reduces the original problem to 2^i subproblems. Ideally then, r should be as close to $\frac{n}{2}$ as possible.

The bulk of the computation required by this algorithm consists of performing two fundamental primitives: matrix multiplication (steps I, III, and IV) and solving systems of linear equations (step II). Therefore, any speedup realized on a parallel architecture for either of these two primitives will result in a corresponding performance gain in any practical implementation.

The ideas presented do not constitute a *numerical* algorithm. In the next section, we shall discuss some practical issues related to the computation of these invariant subspace annihilators.

3. Some practical considerations

The triangularization strategy described in the previous section relies on exploitation of the fact that rank deficient matrices in $\mathfrak{A}(M)$ possess a structure which exposes information about the invariant subspaces of M. We restate here the lemma found in [3, p. 263].

Lemma 1. *Let M be a linear operator on $\mathbb{C}^n$ with minimal polynomial m given by (1) and let $\mathbb{C}^n = \mathcal{N}(f_1^{r_1}(M)) \oplus \cdots \oplus \mathcal{N}(f_k^{r_k}(M))$ be the primary decomposition for M. If W*

is a subspace of $\mathbb{C}^n$ *which is invariant under* M*, then*

$$W = \Big(W \cap \mathcal{N}\big(f_1^{r_1}(M)\big)\Big) \oplus \cdots \oplus \Big(W \cap \mathcal{N}\big(f_k^{r_k}(M)\big)\Big).$$

Suppose now that $a \in \mathbb{C}[x]$ has the property that

$$\gcd(a, m) = \prod_{i=1}^{k} f_i^{\nu_i}, \tag{2}$$

where $0 \leq \nu_i \leq r_i$, $i = 1, \ldots, k$. Then Lemma 1 implies that

$$\bigoplus_{\nu_i > 0} \mathcal{N}\big(f_i^{\nu_i}(M)\big) \subset \mathcal{N}\big(a(M)\big). \tag{3}$$

We also have

Lemma 2. *Let* $M \in \mathbb{C}^{n\times n}$ *have minimal polynomial* m *given by (1). Suppose that* $p \in \mathbb{C}[x]$ *has the property that*

$$\gcd(p, m) = \prod_{i\in\mathfrak{S}} f_i^{r_i},$$

where $\mathfrak{S} \subset \{1, \ldots, k\}$*. Then*

$$\mathcal{N}\big(p(M)\big) = \bigoplus_{i\in\mathfrak{S}} \mathcal{N}\big(f_i^{r_i}(M)\big).$$

Proof. Let $q \in \mathbb{C}[x]$ be defined by

$$q = \prod_{i\notin\mathfrak{S}} f_i^{r_i}.$$

Since p and q are relatively prime, there exist polynomials r and $s \in \mathbb{C}[x]$ such that

$$rp + sq = 1 \tag{4}$$

(see [4, Chapter II, §4]). Suppose that $\vec{x} \in \mathcal{N}\big(p(M)\big) \cap \mathcal{N}\big(q(M)\big)$. Then

$$\vec{x} = (rp + sq)(M)\vec{x} = r(M)p(M)\vec{x} + s(M)q(M)\vec{x} = \vec{0}.$$

Hence we have shown

$$\mathcal{N}\big(p(M)\big) \cap \mathcal{N}\big(q(M)\big) = \{\vec{0}\}. \tag{5}$$

For each $i \in \mathfrak{S}$, since $f_i^{r_i} \big| p$, $\mathcal{N}\big(f_i^{r_i}(M)\big) \subset \mathcal{N}\big(p(M)\big)$ and therefore

$$\bigoplus_{i\in\mathfrak{S}} \mathcal{N}\big(f_i^{r_i}(M)\big) \subset \mathcal{N}\big(p(M)\big).$$

In an analogous fashion, we have

$$\bigoplus_{i \notin \mathfrak{S}} \mathcal{N}\big(f_i^{r_i}(M)\big) \subset \mathcal{N}\big(q(M)\big).$$

By (5) and the fact that $\bigoplus_{i \in \mathfrak{S}} \mathcal{N}\big(f_i^{r_i}(M)\big)$ and $\bigoplus_{i \notin \mathfrak{S}} \mathcal{N}\big(f_i^{r_i}(M)\big)$ are complementary subspaces, the lemma is proved. $\square$

Lemma 2 then implies that for a satisfying (2),

$$\mathcal{N}\big(a(M)\big) \subset \bigoplus_{\nu_i > 0} \mathcal{N}\big(f_i^{r_i}(M)\big). \tag{6}$$

Equation (3) together with (6) suggest a means of computing invariant subspace annihilators for M, namely, by constructing a polynomial which maps a substantial number of the eigenvalues of M to zero. Although these invariant subspace annihilators are easily computed in theory, the difficulty in practice will lie in devising an efficient method for computing them given little *a priori* spectral information. We indicate a possible strategy for devising an efficient algorithm for computing approximate invariant subspace annihilators. Let $\mathcal{D}$ be some bounded region in the complex plane containing $\lambda(M)$. Suppose $p \in \mathbb{C}[x]$ maps the the region $\mathcal{D}$ into the complex plane in such a way that approximately half the eigenvalues of M are mapped near 0 and the rest near the unit circle. Then p approximates an invariant subspace annihilator having the desired properties.

Lemma 2 shows that if an algorithm can be devised for computing the matrices $p(M)$ and $q(M)$, the algorithm in Section 2 can be extended into one for performing block diagonalization of M into the direct sum $M\big|_{W_{\mathfrak{S}}} \oplus M\big|_{W_{\widetilde{\mathfrak{S}}}}$, where

$$W_{\mathfrak{S}} = \bigoplus_{i \in \mathfrak{S}} \mathcal{N}\big(f_i^{r_i}(M)\big),$$

$$W_{\widetilde{\mathfrak{S}}} = \bigoplus_{i \notin \mathfrak{S}} \mathcal{N}\big(f_i^{r_i}(M)\big),$$

and $\mathfrak{S}$ is a nontrivial proper subset of $\{1, \ldots, k\}$. It is not necessary, however, to actually construct *both* p and q. Indeed, we have

LEMMA 3. *Let $M \in \mathbb{C}^{n \times n}$ have minimal polynomial m given by (1) and $\mathfrak{S}$ be a nontrivial subset of $\{1, \ldots, k\}$. Suppose that $p, q \in \mathbb{C}[x]$ satisfy*

$$\gcd(p, m) = \prod_{i \in \mathfrak{S}} f_i^{r_i} \qquad \textit{and} \qquad \gcd(q, m) = \prod_{i \notin \mathfrak{S}} f_i^{r_i}.$$

Then

$$\mathcal{R}\big(p(M)\big) = \mathcal{N}\big(q(M)\big).$$

Proof. Define q as in the proof of Lemma 2. Suppose $\vec{x} \in \mathcal{R}\big(p(M)\big)$. Then there exists a vector $\vec{y}$ such that $p(M)\vec{y} = \vec{x}$. Since $m \big| pq$ and m is the minimal polynomial of M,

$$q(M)\vec{x} = q(M)\big[p(M)\vec{y}\big] = \vec{0}.$$

Hence $\vec{x} \in \mathcal{N}\big(q(M)\big)$ and we have shown that

$$\mathcal{R}\big(p(M)\big) \subset \mathcal{N}\big(q(M)\big). \tag{7}$$

Using the proof of Lemma 2, we see that

$$n = \dim\big(\mathcal{N}\big(p(M)\big)\big) + \dim\big(\mathcal{N}\big(q(M)\big)\big).$$

It is well-known that

$$n = \dim\big(\mathcal{N}\big(p(M)\big)\big) + \dim\big(\mathcal{R}\big(p(M)\big)\big),$$

which in turn yields $\dim\big(\mathcal{N}\big(q(M)\big)\big) = \dim\big(\mathcal{R}\big(p(M)\big)$; together with (7), we then have that $\mathcal{R}\big(p(M)\big) = \mathcal{N}\big(q(M)\big)$. $\square$

In particular, Lemma 3 implies that $\mathcal{N}\big(p(M)\big)$ and $\mathcal{R}\big(p(M)\big)$ are complementary invariant subspaces of M. Lemma 3 clearly suggests a few alternative strategies for block diagonalizing M, all of which are equivalent mathematically. Practical considerations such as availability of appropriate algorithms or known properties of the matrices being considered may single out a clear choice for a given application.

We briefly discuss the implications of the two most important practical requirements: numerical accuracy and parallel speed. The strategy described above is a mathematical outline for designing eigensolvers which presumes exact arithmetic. The success of any implementation of the ISDA given in Section 2 will be dependent on how accurately *approximate* invariant subspace annihilators can be computed using finite precision arithmetic. Furthermore, although we have shown that there are many invariant subspace annihilators for a given matrix M of order n, it is probable that any practical realization of the ISDA which presupposes no *a priori* spectral information must rely on constructing polynomials in M of extremely high degree having approximately $\frac{n}{2}$ eigenvalues near zero. The majority of the total computation time will be spent on the first "divide" and hence on matrix multiplication of matrices of order n. In order to be competitive with existing eigensolvers on parallel architectures, the number of such matrix multiplications cannot be very large.

It is interesting to consider the power method [2, 6] in the above context. The power method uses the approximate invariant subspace annihilator, x^k, where k is some suitably chosen large integer. Unfortunately, this choice of polynomial is unsatisfactory because of the use of finite precision arithmetic. In finite precision, the range or null space of a rank deficient matrix can only be computed accurately when the zero and nonzero singular values are well-separated, which in turn implies that the nonzero singular values must be all be of approximately the same magnitude. In the case of the power method, this happens in general only when k is so large that M^k has very small rank compared to that of M. Hence, high accuracy in the results when solving the complete eigenproblem using the power method exacts a high computational cost.

As illustrated by our discussion of the power method, the use of finite precision arithmetic imposes a critical constraint on the choice of approximate invariant subspace annihilator. This constraint suggests a possible approach to designing algorithms to approximate invariant subspace annihilators which we illustrate by considering the case where M has real eigenvalues. Without loss of generality, assume that $\lambda(M) \subset [0,1]$. Let $c > 0$ and suppose $a \in \mathbb{R}[x]$ maps approximately half the eigenvalues of M to $\{z : |z| < \delta_1\}$ and the remainder to $\{z : |z - c| < \delta_2\}$, where $\delta_1 > 0$ and $\delta_2 > 0$ are small relative to c. For well-conditioned matrices, this choice of a will separate the singular values as desired.

Consider, for example, the incomplete Beta functions, $B_j, j \in \mathbb{N}$, defined on $[0,1]$ by

$$B_j(x) = \frac{\displaystyle\int_0^x t^j(1-t)^j\,dt}{\displaystyle\int_0^1 t^j(1-t)^j\,dt} = \sum_{k=0}^{j} \binom{2j+1}{j-k}\binom{j+k}{k}(-1)^k x^{j+k+1}, \quad j \in \mathbb{N}.$$

For each j, B_j is a polynomial of degree $(2j+1)$ which increases on $[0,1]$, and has fixed points at 0, $\frac{1}{2}$, and 1. Let χ be the function defined on $[0,1]$ by

$$\chi(x) = \begin{cases} 0, & \text{if } 0 \le x < \frac{1}{2}; \\ \frac{1}{2}, & \text{if } x = \frac{1}{2}; \\ 1, & \text{if } \frac{1}{2} < x \le 1. \end{cases}$$

Then for $x \in [0,1]$,

$$\lim_{j\to\infty} B_j(x) = \chi(x). \tag{8}$$

Hence, for large j, B_j is a suitable candidate for a.

The rate of convergence in (8) is quite slow; hence, for an arbitrary matrix, computation of $B_j(M)$ for j large enough to approximate the desired invariant subspace annihilators would be computationally prohibitive. Lederman and Tsao [5] have observed that recursive applications of any fixed function in the family leads to a sequence of functions which converges pointwise to χ very quickly. This observation has allowed them to design and implement a new promising eigensolver for real diagonalizable matrices with real eigenvalues following the strategy outlined in this paper.

Although the methodology discussed in the previous paragraph suggests a number of analogous strategies for matrices with complex eigenvalues, finding appropriate algorithms for computing invariant subspace annihilators for general matrices remains an interesting open problem.

References

1. L. Auslander & A. Tsao, "On a divide and conquer algorithm for the eigenproblem via complementary invariant subspace decomposition", Supercomputing Research Center Technical Report SRC-89-003, 1989.
2. G. Golub & C. F. Van Loan, "Matrix Computations", Johns Hopkins University Press, Baltimore, 1989.
3. K. Hoffman & R. Kunze, "Linear Algebra", Prentice-Hall, New York, 1971.
4. S. Lang, "Algebra", Addison-Wesley, New York, 1971.
5. S. Lederman & A. Tsao, "A parallelizable eigensolver for real diagonalizable matrices with real eigenvalues", (in preparation).
6. J. H. Wilkinson, "The Algebraic Eigenvalue Problem", Oxford Science Publications, Oxford, 1965.