A REVIEW OF THE LANCZOS METHOD APPLIED TO GENERALISED NON-SYMMETRIC EIGENPROBLEMS

SHARON SLADE

NUMERICAL ANALYSIS REPORT 18/86

DEPARTMENT OF MATHEMATICS

P.O. Box 220

University of Reading

Whiteknights

Reading RG6 2AX

Abstract

The Lanczos method is derived and analysed for application to the non-symmetric generalised eigenproblem $A\underline{x} = \lambda E\underline{x}$, where E may be singular. An identification test which effectively removes 'bad' eigenvalue approximations without knowledge of the true eigenvalues and eigenvectors of the system is also considered and both the method and the test applied to several test problems, including a matrix representing a 4-machine system employed by the South of Scotland Electricity Board.

Contents

1 INTRODUCTION

Problems requiring the computation of the eigenvalues of large matrices are now very commonplace and occur in such fields as transient power stability, vibration problems and the numerical solution of p.d.e.'s. Whilst efficient numerical methods are abundant for computing the eigenvalues of sparse, real, symmetric matrices, procedures which deal with non-symmetric matrices are rather less in evidence. However, since, in practical terms, most system matrices are non-symmetric, it is unrealistic not to attempt to create comparable techniques for this more general case. Such techniques as do exist, for example, those decribed in Stewart and Jennings (1981) and Saad (1980), are useful, but offer only a limited amount of information, the former being a simultaneous iteration procedure which calculates a few dominant eigenvalues, the latter based on Arnoldi's method, employing Hessenberg matrices to approximate the given matrix. In their favour, however, is their property, shared with Cullum and Willoughby's Lanczos technique (1986), which will be discussed further, that the original matrix is not modified.

In the first section, the Lanczos procedure with respect to both the symmetric and non-symmetric cases will be derived and analysed.

In the second section, a statement of Cullum and Willoughby's algorithm for the non-symmetric case will be given, together with those modifications which have been included plus some details omitted from their report (1986).

In the third section, numerical results as have been obtained to date will be given and compared with those of Cullum and Willoughby and in the final section, conclusions will be drawn and comparisons made between the effectiveness of the algorithm outlined in this report and other available numerical procedures.

2 DERIVATION AND ANALYSIS OF THE LANCZOS PROCEDURE

In this section, we derive the basic algorithm which describes the Lanczos process and consider its application to the problem $A\underline{x} = \lambda \underline{x}$, say, where A is real, nxn and symmetric. The motivation for the technique comes from the property that the eigensystem of a similarity transformation of a general matrix A is related in a straightforward manner to the eigensystem of A itself. The advantage of this is that A may be reduced to a more simple form from which the calculation of the eigenvalues and eigenvectors is substantially more efficient.

Associated with A is the characteristic polynomial

$$X(\lambda) = \lambda^n - \xi_{n-1} \lambda^{n-1} - \ldots - \xi_0 = 0$$

By the Cayley-Hamilton theorem

$$A^n = \xi_{n-1} A^{n-1} + \ldots + \xi_0 I \qquad (2.1)$$

Suppose we choose $\underline{r}$, an arbitrary vector which lies in the space of all the eigenvectors of A. Multiplying (2.1) by $\underline{r}$, we have

$$A^n \underline{r} = \xi_{n-1} A^{n-1} \underline{r} + \ldots + \xi_0 \underline{r}$$

i.e.

$$A^n \underline{r} = [\underline{r}, A\underline{r}, \ldots, A^{n-1}\underline{r}]\underline{\xi} \qquad (2.2)$$

where $\underline{\xi} = [\, \xi_0, \xi_1, \ldots, \xi_{n-1}\,]^T$.

Consider now the Krylov sequence

$$\underline{r}_0 = \underline{r}$$

$$\underline{r}_{j+1} = A\, \underline{r}_j \qquad j=0,\ldots,n-1$$

so that $\underline{r}_j = A^j\, \underline{r}_0$

Rewrite (2.2) as $[\underline{r}_0, \underline{r}_1, \ldots, \underline{r}_{n-1}]\underline{\xi} = \underline{r}_n$ or

$$B\underline{\xi} = \underline{r}_n . \qquad (2.3)$$

Therefore, the coefficients of the characteristic equation satisfy n linear equations in n unknowns. Multiplying (2.3) by B^T gives the normal equations

$$B^T B\underline{\xi} = B^T \underline{r}_n$$

or

$$M\,\underline{\xi} = \underline{K}$$

where

$$\begin{aligned} m_{i+1\ j+1} &= \underline{r}_i{}^T \underline{r}_j \\ &= (A^i \underline{r}_0)^T A^j \underline{r}_0 \\ &= \underline{r}_0{}^T A^{i+j} \underline{r}_0 \end{aligned} \qquad i,j=0,1,\ldots,n-1$$

and

$$K_{i+1} = \underline{r}_i{}^T \underline{r}_n \qquad i=0,1,\ldots,n-1$$

Note that the element m_{ij} depends only on the sum of the indices, so write

$$m_{i+1\ j+1} = \mu_{i+j} \qquad i+j=0,\ldots,2n-1$$

and

$$M = \begin{pmatrix} \mu_0 & \mu_1 & \cdot & \cdot & \mu_{n-1} \\ \mu_1 & & \cdot & \cdot & \\ & & & & \\ \cdot & & & & \cdot \\ & & & & \\ \mu_{n-1} & \cdot & \cdot & \cdot & \mu_{2n-2} \end{pmatrix} .$$

Lanczos gave an efficient algorithm for solving $M\underline{\xi} = \underline{K}$. Once $\underline{\xi}$ is determined, it is necessary to find the roots of the characteristic equation.

However, the numerical algorithm is not stable, i.e. a small perturbation in M greatly affects the vector $\underline{\xi}$ and, furthermore, the Krylov vectors $\underline{r}$, $A\underline{r}$, etc soon become parallel to the eigenvector associated with the dominant eigenvalue.

Assume an arbitrary starting vector $\underline{r}_1$ and some set of n independent vectors $\underline{x}_i$ and form the set of modifed Krylov vectors, $\underline{r}_2, \underline{r}_3, \ldots, \underline{r}_{n+1}$ defined by the relations

$$k_{i+1}\underline{r}_{i+1} = A\underline{r}_i - \sum_{j=1}^{i} h_{ji}\underline{r}_j . \qquad (2.4)$$

The k_{i+1} are scalars and usually employed as normalising factors. The h_{ji} are also scalars and are calculated so that $\underline{r}_{i+1}$ is orthogonal to $\underline{x}_1, \ldots, \underline{x}_1$.

Equations (2.4) may be rewritten as

$$A\,[\,\underline{r}_1, \ldots, \underline{r}_n\,] = [\,\underline{r}_1, \ldots, \underline{r}_n\,] \begin{bmatrix} h_{11} & h_{12} & \cdots & h_{1n} \\ k_2 & h_{22} & \cdots & h_{2n} \\ & & & \\ 0 & & k_n & h_{nn} \end{bmatrix}$$

or $A\,R = R\,H$

where H is upper Hessenberg.

However, the above process may break down if $\underline{r}_{i+1}$ is null. In this case we employ a new vector $\underline{v}_{i+1}$, orthogonal to $\underline{x}_1, \ldots, \underline{x}_i$. Thus we may write

$$\underline{v}_1 = \underline{r}_1 \;,\; \underline{r}_{i+1} = A\,\underline{v}_i - \sum_{j=1}^{i} h_{ji}\underline{v}_j \;,\; k_{i+1}\underline{v}_{i+1} = \underline{r}_{i+1} \;,\; \underline{r}_{i+1} \neq 0$$

$\underline{v}_{i+1}$ an arbitrary non-null vector orthogonal to $\underline{x}_1, \ldots, \underline{x}_i$, $\underline{r}_{i+1} = 0$. (2.5)

Choices of the initial set $X = \{\underline{x}_1, \ldots, \underline{x}_n\}$ obviously affect the resultant method. In order to create the procedure known as the Lanczos method, the $\underline{x}_i$'s are derived simultaneously with the $\underline{r}_i$'s and $\underline{v}_i$'s and are derived from A^T in the same way that the $\underline{v}_i$'s are derived from A. Such vectors will be labelled $\underline{t}_i$ and $\underline{w}_i$ so that we have

$$k^*_{i+1}\underline{v}_{i+1} = \underline{r}_{i+1} = A\underline{v}_i - \sum_{j=1}^{i} h^*_{ji}\underline{v}_j$$

$$k_{i+1}\underline{w}_{i+1} = \underline{t}_{i+1} = A^T\underline{w}_i - \sum_{j=1}^{i} h_{ji}\,\underline{w}_j \qquad (2.6)$$

where h_{ji} and h^*_{ji} are calculated so that $\underline{r}_{i+1}$ is orthogonal to $\underline{w}_1,\ldots,\underline{w}_i$ and $\underline{t}_{i+1}$ orthogonal to $\underline{v}_1,\ldots,\underline{v}_i$ respectively. Again we have

$$A\,V = V\,H \qquad \text{and} \qquad W^TV = L$$

$$A^TW = W\,H^* \qquad \text{and} \qquad V^TW = L^* \qquad (2.7)$$

where H and H^* are upper Hessenberg and L and L^* are lower triangular. From (2.7) we have $L = (L^*)^T \Rightarrow$ $L = L^*$ are diagonal

$= D$, say.

Then

$$H = V^{-1}A\,V = D^{-1}(H^*)^TD \qquad (2.8)$$

Since the LHS of (2.8) is upper Hessenberg and the RHS lower Hessenberg, it follows that both must be tri-diagonal so that

$$h_{ji} = h^*_{ji} = 0 \quad j=1,\ldots,i-2$$

$\Rightarrow$ if $A\underline{v}_i$ is orthogonalised with respect to $\underline{w}_{i-1}$ and $\underline{w}_i$ then it is also orthogonalised with respect to w_j $j=1,\ldots,i-2$. Similarly for A^Tw_i.

Furthermore, from (2.8), it follows that

$$h_{ii} = h^*_{ii}$$

and

$$h_{ii+1}k_{i+1} = h^*_{ii+1}k^*_{i+1}$$

so that, by choosing $k_j, k^*_j = 1$, $j=1,\ldots,n$,

$$h_{ii} = h^*_{ii} \Rightarrow \quad H = H^*$$

Rewriting (2.6) to encompass the above gives

$$\beta_{i+1}\underline{v}_{i+1} = \underline{r}_{i+1} = A\underline{v}_i - \alpha_i\underline{v}_i - \gamma_i\underline{v}_{i-1}$$

$$\gamma_{i+1}\underline{w}_{i+1} = \underline{t}_{i+1} = A^T\underline{w}_i - \alpha_i\underline{w}_i - \beta_i\underline{w}_{i-1} \qquad (2.10)$$

$$i=1,\ldots,n$$

where orthogonality decrees that

$$\alpha_i = \frac{\underline{w}_i A \underline{v}_i}{\underline{w}^T_i \underline{v}_i} \tag{2.11}$$

$$\beta_i = \frac{\underline{v}_{i-1} A^T \underline{w}_i}{\underline{v}^T_{i-1} \underline{w}_{i-1}} \tag{2.12}$$

$$\gamma_i = \frac{\underline{w}_{i-1}{}^T A \underline{v}_i}{\underline{w}^T_{i-1} \underline{v}_{i-1}} \tag{2.13}$$

Equations (2.10)-(2.13) define the basic Lanczos process.

The values assigned to α_i, β_i and γ_i decree that the Lanczos vectors $\{V_j\} = \{\underline{v}_1, \ldots, \underline{v}_j\}$ and $\{W_j\} = \{\underline{w}_1, \ldots, \underline{w}_j\}$ are real biorthogonal, i.e. $V_j{}^T W_j = I_j$. The real biorthogonal projections of $A\underline{v}_i$ and $A^T\underline{w}_i$ onto the vectors $\underline{w}_i$ and $\underline{v}_i$ are $\alpha_i \underline{v}_i$ and $\alpha_i \underline{w}_i$ respectively. Similarly, $\gamma_i \underline{v}_{i-1}$ and $\beta_i \underline{w}_{i-1}$ are the real biorthogonal projections of $A\underline{v}_i$ and $A^T\underline{w}_i$ onto $\underline{w}_{i-1}$ and $\underline{v}_{i-1}$ respectively.

So, having eliminated the majority of h_i's in the matrix H and introduced a new notation, the Lanczos matrix, now denoted by T_n (where n is the order of the matrix) is given by

$$T_n = \begin{vmatrix} \alpha_1 & \beta_2 & & & \\ \gamma_2 & \alpha_2 & \beta_3 & & 0 \\ & & & & \\ & & & & \\ 0 & & \gamma_{n-1} & \alpha_{n-1} & \beta_n \\ & & & \gamma_n & \alpha_n \end{vmatrix}$$

Clearly, not all of the values of α_i, β_i and γ_i need be calculated. The algorithm may be stopped at any point m, say, and the Lanczos matrix, T_m, constructed. Since the two sets of vectors V_j and W_j are biorthogonal, the Lanczos matrices, T_j, j=1,.. are the orthogonal

projections of A onto the subspaces V_j, i.e. $T_j = W_j^T A V_j$.

Note that breakdown can occur if the product $\gamma_{i+1}\beta_{i+1} = \underline{r}^T_{i+1}\underline{t}_{i+1}$ is zero. However, this is by no means unwelcome, since it implies the presence of an exact invariant subspace and, thus, that the eigenvalues yielded by the method to this point are, theoretically, correct. Parlett and Taylor (1984) have shown, though, that this happens only rarely and for very special matrices and choices of starting vectors. Cullum and Willoughby's adaptation of this method, which encompasses the non-symmetric case, entails the formulation of a Lanczos matrix whose resultant form is symmetric and, typically, complex. They achieve this by eliminating the variable γ_i by setting $\beta_i = \gamma_i = \sqrt{(\underline{r}_i^T \underline{t}_i)}$. In addition, they produce an equivalent rewritten form for α_i so that we have

$$\begin{aligned}
\beta_{i+1}\underline{v}_{i+1} &= A\underline{v}_i - \alpha_i\underline{v}_i - \beta_i\underline{v}_{i-1} \equiv \underline{r}_{i+1} \\
\beta_{i+1}\underline{w}_{i+1} &= A^T\underline{w}_i - \alpha_i\underline{w}_i - \beta_i\underline{w}_{i-1} \equiv \underline{t}_{i+1} \\
\alpha_i &\equiv (\alpha_i^v + \alpha_i^w)/2 \\
\beta^2_{i+1} &\equiv (\underline{r}^T_{i+1}\underline{t}^T_{i+1}) \\
\alpha^v_i \equiv \underline{w}^T_i(A\underline{v}_i - \beta_i\underline{v}_{i-1}) &\text{ and } \alpha^w_i \equiv \underline{v}^T_i(A^T\underline{w}_i - \beta_i\underline{w}_{i-1})
\end{aligned} \tag{2.14}$$

2.1 CONVERGENCE

A great deal of theory has been produced concerning the convergence of the Lanczos method upon application to a real, symmetric matrix A. Cullum and Willoughby make a statement in their report (1986) which summarizes the effectiveness of the Lanczos method and label it the Lanczos Phenomenon, it being:

Given a real, symmetric matrix A, use of the Lanczos method to generate an mxm matrix T_m will, for large enough m, cause every distinct eigenvalue of A to appear amongst the eigenvalues of T_m.

In theory then, the implication is that if we continue with the computation, at some point we will be able to extract all of the eigenvalues that we require. In practice, however, losses in biorthogonality somewhat nullify the validity of this assumption and 'extra' eigenvalues appear. These take the form of either duplicates of acceptable eigenvalues or bad approximations, labelled as 'spurious'.

Cullum and Willoughby have developed some kind of identification test which examines each eigenvalue produced by the Lanczos matrix and sorts the good from the bad. Justifications, which will follow, have been made for the application of this test to the real, symmetric case, but there are none for the case when A is non-symmetric, although the test appears to work equally favourably in practice. In both cases orthogonality is lost as convergence of the eigenvalues of T_m to A occurs as m increases.

In order to provide a basis for the justification of the test, a series of definitions and lemmas follow:

Definition 2.1

Define a submatrix T_{ij} of the Lanczos matrix T_m where $1 \leq i \leq j \leq m$ as

$$T_{ij} = \begin{bmatrix} \alpha_i & \beta_{i+1} & & & 0 \\ \gamma_{i+1} & \alpha_{i+1} & \beta_{i+2} & & \\ & & & & \\ & & & & \beta_j \\ 0 & & & \gamma_j & \alpha_j \end{bmatrix}$$

In addition, define $\delta_k \equiv \gamma_k \beta_k \qquad 2 \leq k \leq m$

and write $T_k \equiv T_{1k}$ and $\hat{T}_k \equiv T_{km} \qquad 1 \leq k \leq m$.

Definition 2.2

Define the characteristic polynomial of the submatrices T_k and $\hat{T}_k$, $k = 1,2,\ldots,m$ of any tri-diagonal T_m to be

$$a_k(\mu) \equiv \det(\mu I - T_k) \text{ and } \hat{a}_k(\mu) \equiv \det(\mu I - \hat{T}_k)$$

respectively.

Lemma 2.1

For any μ, a simple eigenvalue of an mxm symmetric matrix T_m, the components of its right eigenvector may be given by

$$[u(k)]^2 = \frac{a_{k-1}(\mu)\hat{a}_{k+1}(\mu)}{a'_m(\mu)} \qquad k = 1,2,\ldots,m$$

where $\hat{a}_{m+1}(\mu) = 1$ and $a_0(\mu) = 1$.

Lemma 2.2

Let μ be an eigenvalue of T_m, a tri-diagonal matrix. By definition, $a_m(\mu) = 0$. Then

$$\hat{a}_2(\mu)a_{m-1}(\mu) = \prod_{k=2}^{m} \delta_k$$

where $\delta_k = \gamma_k \beta_k$.

Proofs for the above two lemmas may be found in Cullum and Willoughby (1985).

From Lemma 2.1, it can be seen that the mth component of u, the right eigenvector of a simple eigenvector μ of T_m is given by

$$[u(m)]^2 = \frac{a_{m-1}(\mu)}{a'_m(\mu)}$$

Suppose for this pair $\{\mu,\underline{u}\}$ we define the corresponding Ritz value and Ritz vectors $\{\mu,V_m\underline{u},W_m\underline{u}\}$ for A. Then, from the matrix formulation of equations (2.14), i.e.

$$AV_m = V_m T_m + \beta_{m+1} v_{m+1} \underline{e}_m^T \tag{2.15}$$

where $\underline{e}_m$ is the mth unit vector , we can form an expression for the residual norm which employs these Ritz vectors by multiplying (2.15) by $\underline{u}$, it being:

$$\frac{\|AV_m\underline{u} - \mu V_m\underline{u}\|}{\|V_m\underline{u}\|} = \frac{|\beta_{m+1}\, u(m)|\, \|v_{m+1}\|}{\|V_m\underline{u}\|} + E_m \tag{2.16}$$

In order to validate the identification test, it must first be shown that the residual norm (2.16) and that of the corresponding Ritz vector $W_m\underline{u}$ are in some sense sensible measures of the accuracy of the computed Ritz value μ.

2.2 Symmetric case

For the real, symmetric case, results are well established for the convergence of the Ritz value μ to the actual eigenvalue λ, the best known of which is:

Given any scalar μ and non-zero vector $\underline{x}$, $\exists$ an eigenvalue λ of A such that

$$|\lambda - \mu| \leq \frac{\|A\underline{x} - \mu\underline{x}\|}{\|\underline{x}\|} \qquad (2.17)$$

Clearly a small residual => a small error in the eigenvalue approximation. Furthermore, a perturbation of ε, say, in A will produce a perturbation no greater than ε in the approximation. In addition, if A is real and symmetric, $\|\underline{v}_{m+1}\| = 1$ in (2.16) and Paige (1976) has shown that the errors E_m are bounded by $\|A\|$ and the machine epsilon. Moreover, correct implementation of the method leads to values of $\|V_m \underline{u}\|$ which are not 'small'. The implication then is that the degree of approximation of a Ritz value and vector $(\mu,\underline{u})$ obtained from T_m to those of A can be estimated from $|\beta_{m+1} u(m)|$, a value easily obtained.

2.3 **Non-symmetric case**

Unfortunately, this linear relationship between a perturbation in the original matrix and a corresponding perturbation in the approximation exists only for the symmetric case. Results hold for diagonalisable matrices (Bauer and Fike (1960)), but more general results may be found in Kahan et al (1982).

Recall from the properties of the derived Lanczos matrix that

$$T_j = W_j^T A V_j$$

and let λ be a typical eigenvalue of T_j. Then

$$T_j \underline{p} = \underline{p}\theta, \quad \underline{q}^T T_j = \theta \underline{q}^T, \quad \underline{q}^T \underline{p} = 1$$

and suppress the dependence of θ, $\underline{p}$ and $\underline{q}$ on j.

Now V_j and W_j have the properties that

$$V_j^T W_j = I_j \tag{2.18}$$

$$AV_j - V_j T_j = (0,0,\ldots,0,\beta_{j+1}\underline{v}_{j+1}) \tag{2.19}$$

$$W_j^T A - T_j W_j^T = (0,0,\ldots,0,\gamma_{j+1}\underline{w}_{j+1}) \tag{2.20}$$

Multiply (2.19) and (2.20) by $\underline{p}$ and $\underline{q}^T$ and introduce ξ_j and ω_j, the last elements of $\underline{p}$ and $\underline{q}^T$ respectively, such that

$$AV_j\underline{p} - V_j\underline{p}\theta = \beta_{j+1}\xi_j\underline{v}_{j+1}$$

$$\underline{q}^T W_j^T A - \theta\underline{q}^T W_j = \gamma_{j+1}\omega_j \underline{w}_{j+1}^T$$

and the approximate eigenvalues are $\underline{x} = V_j\underline{p}$, $\underline{y}^T = \underline{q}^T W_j^T$

Corollary The closest matrix to A with $(\theta, \underline{x}, \underline{y}^T)$ as an eigenelement is A - E for an E satisfying

$$\|E\| = \max\left\{ \frac{|\beta_{j+1}\xi_j|\,\|\underline{v}_{j+1}\|}{\|\underline{x}\|} , \frac{|\gamma_{j+1}\omega_j|\,\|\underline{w}_{j+1}^T\|}{\|\underline{y}^T\|} \right\}$$

Kahan et al have shown that E is the rank 2 matrix

$$E = \frac{(\beta_{j+1}\xi_j)\,\underline{v}_{j+1}\underline{x}^T}{\|\underline{x}\|^2} + \frac{(\gamma_{j+1}\underline{w}_j)\,\underline{y}\underline{w}_{j+1}^T}{\|\underline{y}\|^2}$$

The object of interest, $|\lambda - \theta|$, is unknown, but when $\|E\|$ is small enough then

$$|\lambda - \theta| \leq \text{cond}(\theta).\|E\| + O(\|E\|^2)$$

Suppose we regard A as a perturbation of A - E, then we may take

$$\text{cond}(\theta) = \text{cond}(\theta, A - E)$$

$$= \frac{\|\underline{x}\|\|\underline{y}^T\|}{\|\underline{y}^T\underline{x}\|} = \|\underline{x}\|\|\underline{y}^T\|$$

Consequently, the Lanczos algorithm should be allowed to run until the perturbation $\|E\|$ is small and $\text{cond}(\theta).\|E\|$ satisfies the given tolerance. For $j<<n$, the cost of calculating an eigenvalue of T_j compares favourably with that of executing a step of the Lanczos process. In principle $\underline{x}$ and $\underline{y}^T$ are also computable, but this involves

the computation and formulation of V_j and W_j at a cost of 2jn operations, probably more than that of a step of Lanczos. Instead, Kahan et al suggest that the values $\|V_j\|_F^2$ and $\|W_j^T\|_F^2$ be saved in the fast store, where

$$\|V_j\|_F^2 = \sum_{i=1}^{j} \|\underline{v}_i\|^2 , \quad \|W_j^T\| = \sum_{i=1}^{j} \|\underline{w}_{iT}\|^2 ,$$

values which may easily be updated.

A more accessible bound then is given by

$$\text{cond}(\theta) = \|\underline{x}\|\|\underline{y}^T\| \leq \|V_j\|_F \|W_j^T\|_F \|\underline{p}\|\|\underline{q}\|$$

2.4 **Validation of the identification test**

From (2.16) and Lemma 2.1, we can see that

$$\frac{\|AV_m\underline{u} - \mu V_m\underline{u}\|}{\|V_m\underline{u}\|} = \frac{\|\beta_{m+1}\underline{v}_{m+1}\|[a_{m-1}(\mu)\hat{a}_2(\mu)]^{0.5}}{\|V_m\underline{u}\|\,[a'_m(\mu)\hat{a}_2(\mu)]^{0.5}} + E_m \qquad (2.21)$$

Employing Lemma 2.2 and adopting the procedure of Cullum and Willoughby in their non-symmetric method that $\beta_k = \gamma_k$, we have

$$\frac{\|AV_m\underline{u} - \mu V_m\underline{u}\|}{\|V_m\underline{u}\|} = \frac{|\Pi_{k=2}^{m+1}\beta_k|\,\|\underline{v}_{m+1}\|}{\|V_m\underline{u}\|[a'_m(\mu)\hat{a}_2(\mu)]^{0.5}} + E_m \qquad (2.22)$$

where E_m is the roundoff error which, in practice, appears to remain controlled and 'small', although theory is lacking here.

Cullum and Willoughby's identification test rests on this 'fact' and, in addition, requires that the Ritz vectors are not small and that the norms $\|\underline{v}_{m+1}\|$ and $\|\underline{w}_{m+1}\|$ are not big.These requirements when substituted into (2.22) imply that the size of the residual norm with respect to the eigenvalue μ is dependent on the $\hat{a}_2(\mu)$ term, so that small $\hat{a}_2(\mu)$ => a large residual term, which in turn implies a spurious eigenvalue. Hence, any eigenvalue which is both a member of T_m and its

submatrix $\hat{T}_2$ may be dismissed as having been produced by losses in biorthogonality. On a more positive note, numerically multiple eigenvalues may be accepted as 'good' approximations, since the implication here is that duplicates arise from eigenvalue approximations of matrices which are themselves good approximations of the original matrix A. This then is the justification and statement of Cullum and Willoughby's identification test.

3 STATEMENT OF THE ALGORITHM

The problem under consideration is the generalised eigenproblem, $A\underline{x}=\lambda E\underline{x}$, where both A and E may be sparse. Both matrices are generated and stored by the facility MATLAB and accessed by call statements in the program NSL7.

The majority of numerical methods for eigenvalue problems have the properties that they first yield the dominant eigenvalue and that the computation may be stopped at a point when all the eigenvalues of interest have emerged. In the same way, we shall exploit the Lanczos method by adapting the problem somewhat so that we may centre the computation around a particular range of eigenvalues.

We consider then the shifted inverse problem

$$(E^{-1}A - \delta I)^{-1}\underline{x} = \mu\underline{x}, \qquad \mu = 1/(\lambda - \delta)$$

which, in theory, will compute the eigenvalues μ which are closest to the shift δ. However,we will refrain from computing the inverse of the adjusted matrix directly, working instead with its LU decomposition and employing forward and backward substitution whenever the formation of a matrix-vector product is required.i.e. we consider

$$\begin{aligned} \underline{z} &= (E^{-1}A - \delta I)^{-1}\underline{v} \\ LU &= (E^{-1}A - \delta I) \\ L\underline{y} &= \underline{v} \\ U\underline{z} &= \underline{y} \end{aligned} \qquad (3.1)$$

and solve for $\underline{z}$. Ideally, we would like to avoid computing the inverse of the matrix E and deal with the problem in the following form:

$$\underline{z} = (A - \delta E)^{-1}E\underline{v}$$

$$LU = A - \delta E$$
$$L\underline{y} = E\underline{v} \qquad (3.3)$$
$$U\underline{z} = \underline{y}$$

Part of the computation requires that we take the product of the transpose of the matrix in question with a vector. This leads to the following formulation:

$$\underline{z} = E^T[(A - \delta E)^T]^{-1}\underline{v}$$
$$LU = A - \delta E$$
$$U^T\underline{w} = \underline{v} \qquad (3.3)$$
$$L^T\underline{y} = \underline{w}$$
$$E^T\underline{y} = \underline{z}$$

In order to implement the algorithm given by equations (2.14), two starting vectors $\underline{v}_1$ and $\underline{w}_1$ are required (we assume here that $\beta_0=0$). The vectors must satisfy the condition that $\|\underline{v}_1\| = \|\underline{w}_1\| = 1$ and further restrictions are made by imposing the condition that $\underline{v}_1 = \underline{w}_1$. There is little literature concerning the choices of these starting vectors. Cullum and Willoughby imply that, so long as the above two conditions are satisfied, the choices are arbitrary, whilst other authors (Parlett and Taylor (1984), Wilkinson (1965)) imply that the choices of the starting vectors play some small part in the resultant success of the method. The Power Method suggested itself as a sensible means of generating the vectors, and a few iterations of this were employed, the number of iterations being dependent on the size n of the matrix.

So, the algorithm for the Lanczos Method is given as:

Step 1

Read matrices A, E and order n.

Input the size m of the final Lanczos matrix T_m and the shift (complex), δ, required.

Generate starting vectors $\underline{v}_1$, $\underline{w}_1$ using Power Method and normalise.

Assume working matrix to be in the form $A1 = E^{-1}A - \delta I$.

N.B. For any following matrix-vector products of the form $(A1)^{-1}\ \underline{x} = \underline{y}$, $\underline{y}$ is found by solving the system $\underline{x} = A1\underline{y} = LU\underline{y}$ by forward and backward substitution.

Also note: Since only the previous, current and new values of each vector $\underline{v}$ and $\underline{w}$ are employed at any time, throughout the computation the vectors $\underline{v}_{i-1}$, $\underline{v}_i$ and $\underline{v}_{i+1}$ are stored in $\underline{v1}$, $\underline{v2}$ and $\underline{v3}$ respectively (similarly for $\underline{w}$'s) and are updated at the end of every step.

Step 2

$$\text{Set } \alpha_1 = (\underline{w}_1^T A1^{-1}\underline{v}_{-1} + \underline{v}_1^T(A1^{-1})^T\underline{w}_1)/2$$

$$\beta_2 = [(A1^{-1}\underline{v}_1 - \alpha_1\underline{v}_{-1})^T((A1^{-1})^T\underline{w}_1 - \alpha_1\underline{w}_{-1})]^{0.5}$$

Step 3

For i=2,...,m

$$\beta_{i+1}\underline{v}_{i+1} = A1^{-1}\underline{v}_i - \alpha_i\underline{v}_i - \beta_i\underline{v}_{i-1} \equiv \underline{r}_{i+1}$$

$$\beta_{i+1}\underline{w}_{i+1} = (A1^T)^{-1}\underline{w}_i - \alpha_i\underline{w}_i - \beta_i\underline{w}_{i-1} \equiv \underline{t}_{i+1}$$

$$\alpha_i = (\alpha_i^v + \alpha_i^w)/2$$

$$\beta_{i+1}^2 = \underline{r}_{i+1}^T\underline{t}_{i+1}$$

$$\alpha_i^v = \underline{w}_i^T(A1^{-1}\underline{v}_i - \beta_i\underline{v}_{i-1})$$

$$\alpha_i^w = \underline{v}_i^T((A1^T)^{-1}\underline{w}_i - \beta_i\underline{w}_{i-1})$$

Step 4

If $\beta_{i+1}=0$, breakdown occurs. Form matrix T_i as in Step 5 and compute eigenvalues to date as in Step 6. Restart process by implementing the last step of the Gram-Schmidt process to generate two new starting vectors $\underline{v}_1$ and $\underline{w}_1$, orthogonal to all previous vectors $\underline{w}_j$ and $\underline{v}_j$, j=1,...,i-1, respectively. (Unfortunately, this entails storing all vectors as they are generated, although they could be placed onto some sort of external storage device such as disc where they may be retrieved only as required.)

Step 5

Form Lanczos matrix, $T_m = \begin{bmatrix} \alpha_1 & \beta_2 & & & 0 \\ \beta_2 & \alpha_2 & \beta_3 & & \\ & & & & \\ & & & & \beta_m \\ 0 & & & \beta_m & \alpha_m \end{bmatrix}$

Step 6

Solve for eigenvalues $\mu = 1/(\lambda - \delta)$ and re-evaluate true eigenvalues, $\lambda = \delta + 1/\mu$.

4 NUMERICAL RESULTS

Numerical results obtained by the Lanczos method differ vastly in their accuracy and are largely dependent on the condition number of the shifted matrix, it being a measure of the matrix's sensitivity to perturbations. A condition number equal to, or just greater than, 1 indicates that a small change in the components of that matrix will not greatly alter the results obtained whilst a 'large' condition number implies that that matrix is extremely sensitive to such changes.

The matrices initially examined were all constructed via the MATLAB facility and typically had classical structures, i.e. of symmetric or Hessenberg type, although non-symmetric matrices were, of course, also examined, these being the main source of interest. Of these cases, various combinations were employed, differing both m (the final dimension of the Lanczos matrix) and the value of the shift.

The first group of test results were achieved from matrices specially generated to display certain characteristics, for example, distinct eigenvalues, multiple eigenvalues and defective eigenvectors.

As one would expect, matrices with low condition numbers gave reliable results, often picking up every eigenvalue to machine accuracy. In those cases where a shift was chosen to coincide with a known eigenvalue, the matrix predictably became very ill-conditioned and consequently the results were poor.

The second group of test problems involved a non-symmetric matrix, A, having eigenvalues spread over both the positive and negative planes. Various shifts were chosen, these taking values at eigenvalues, below and above eigenvalues and zero. Results were generally poor, but were considerably more promising in the case where the shift was chosen to be slightly greater than the dominant eigenvalue (see figures 1 and 2 at end of section).The problem here, of course, is that prior knowledge of the dominant eigenvalue would seem to be required to guarantee good approximations.

In the next section of problems, the matrix E was incorporated so that the problem became a generalised eigenproblem. Here, A was chosen to be symmetric and E to be symmetric positive-definite. For the first study with this particular combination of matrices, no shift was incorporated, but the value of m (determining the size of the Lanczos matrix T_m) was reduced each time. The results obtained are represented at the end of this section in figures 3-6. In the second study, m was consistently set as the size of A and E while the shift values were altered and finally, these values were retained while m was given a common, low value.

The results all displayed a common characteristic. A shift chosen to be slightly greater than the dominant positive eigenvalue produced very good results and multiple approximations, it having the effect of lowering the 'new' matrix condition number. A shift chosen slightly below the dominant eigenvalue behaved less well, whilst a shift chosen just above the next dominant eigenvalue produced good approximations to that and the next few eigenvalues, but did not manage to pick up the true dominant member of the set. Similar results were obtained for shifts chosen further down the range. A noticeable difference on the

accuracies obtained could be seen between the second group where A was non-symmetric and the third where A was symmetric, although in both cases similar matrix condition numbers were recorded with the symmetric case being a great deal more accurate. The most encouraging results were seen in the third section where m was repeatedly reduced. With m = n, the dimension of A, approximately 56% of distinct eigenvalues were recovered with a further 22% of duplicates. Reducing the Lanczos matrix to 80% of its original size caused 62% of all possible distinct eigenvalues (i.e. 80% of those of the original matrix) to be picked up with a further 12% of duplicates. With a Lanczos matrix 60% of its original size, 50% of all possible distinct values were obtained with 10% of duplicates. A value of m equal to 40% of n led to 80% of all possible recoveries whilst m equal to 20% of n gave a 70% success rate. In this case then, reducing the size of m did not impair the ability of the method to pick up those eigenvalues lying closest to the shift. Each reduction of m merely caused the Lanczos matrix to 'drop' those eigenvalues furthest from, and consequently of less interest than, the shift.

In considering a particular state matrix of the SSEB system, the method was used in both the shifted,inverse form and in its most basic form. A full computation of the eigenvalues employing the latter produced extreme discrepancies between the actual eigenvalues of the system and those given by Lanczos. However, after application of Cullum and Willoughby's identification test (1986), the most inaccurate of these were eliminated and the remainder proved to be reasonable approximations (see figures 7 and 8). It was unfortunate that these approximations did not, in some sense, lie 'outside' the true eigenvalues, for then they might have represented a reasonable bound for the true solution. Interestingly, reduction of the size of

the Lanczos matrix resulted in removal of the same spurious eigenvalues with no further reduction by the identification test (figures 9 and 10). Application of the shifted, inverse form produced similar results to the reduced Lanczos method (figure 11).

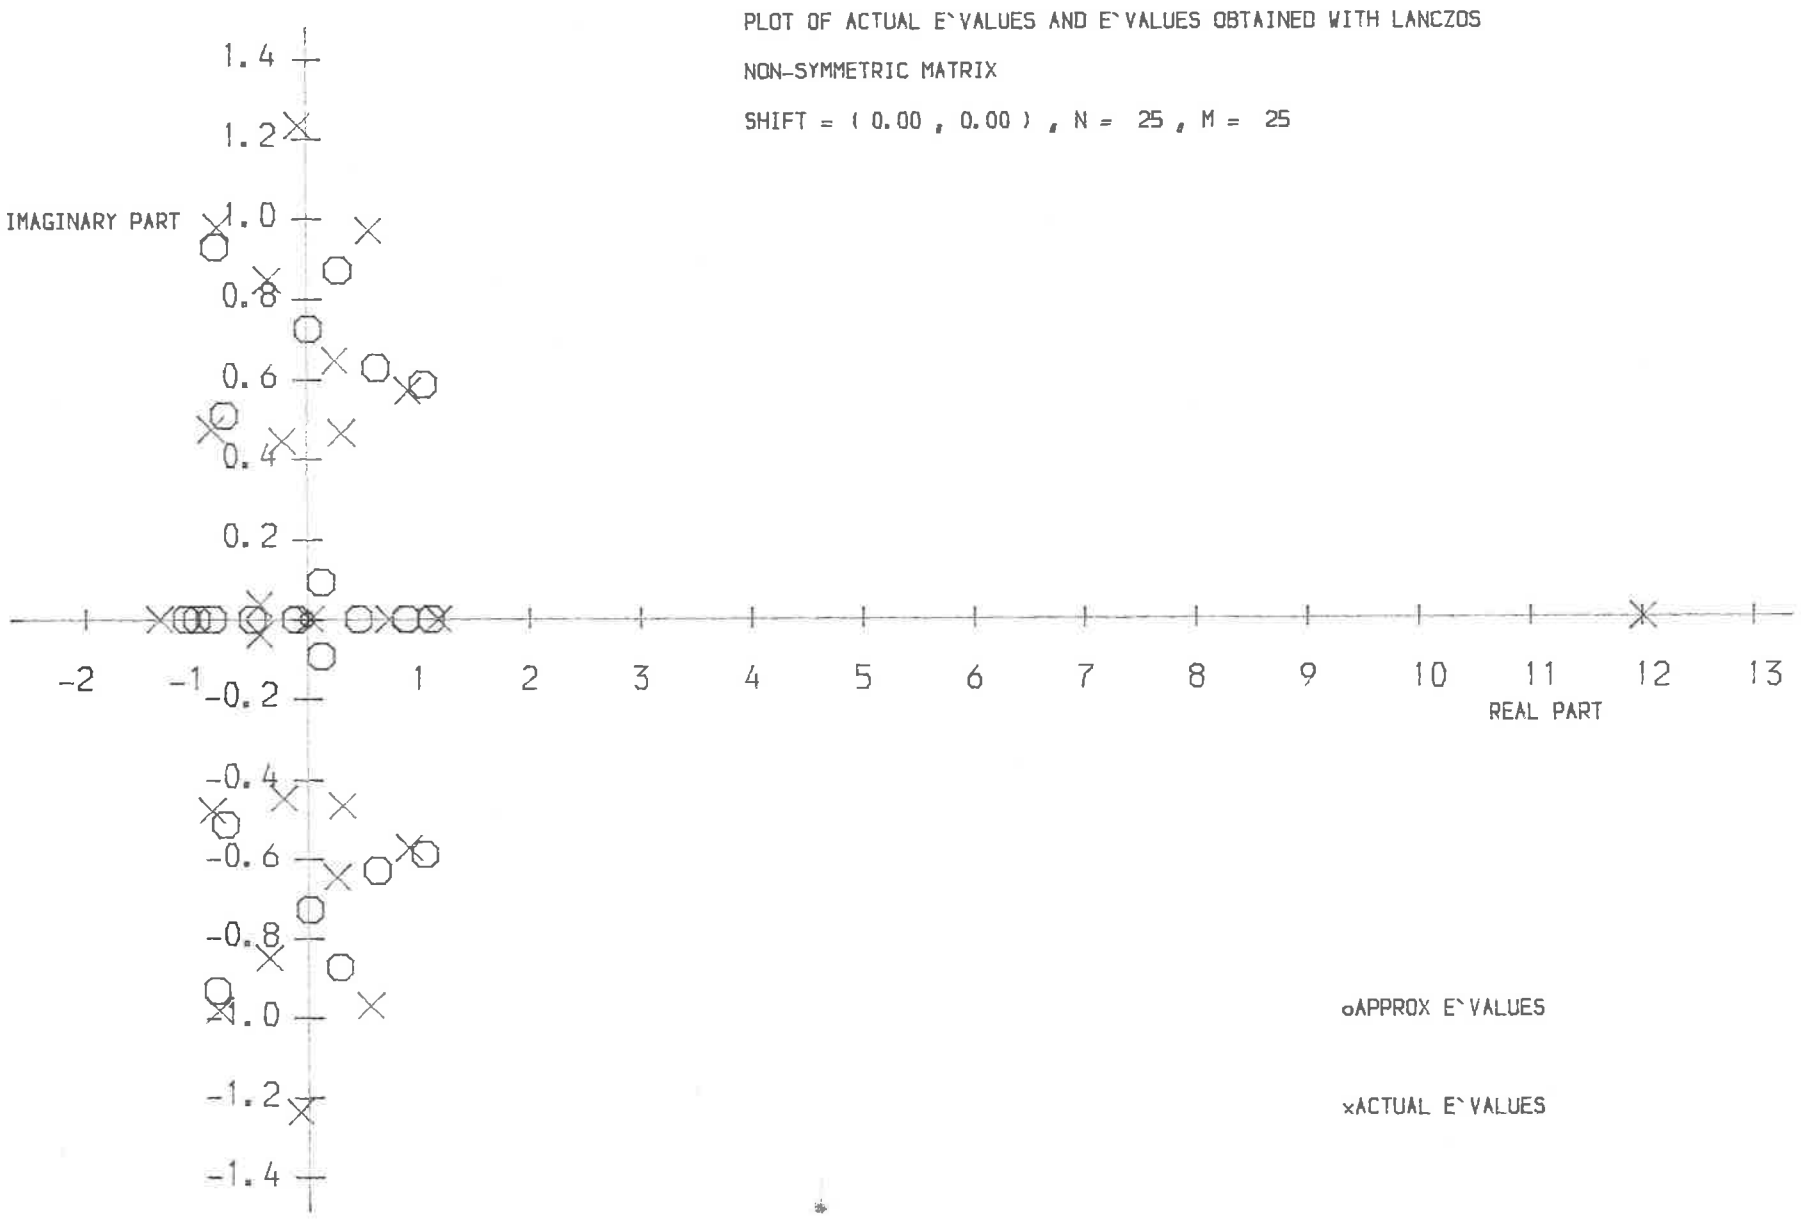
PLOT OF ACTUAL E`VALUES AND E`VALUES OBTAINED WITH LANCZOS
NON-SYMMETRIC MATRIX
SHIFT = (0.00 , 0.00) , N = 25 , M = 25
IMAGINARY PART
1.4
1.2
1.0
0.8
0.6
0.4
0.2
-0.2
-0.4
-0.6
-0.8
-1.0
-1.2
-1.4
-2
-1
1
2
3
4
5
6
7
8
9
10
11
12
13
REAL PART
oAPPROX E`VALUES
xACTUAL E`VALUES

FIGURE 1

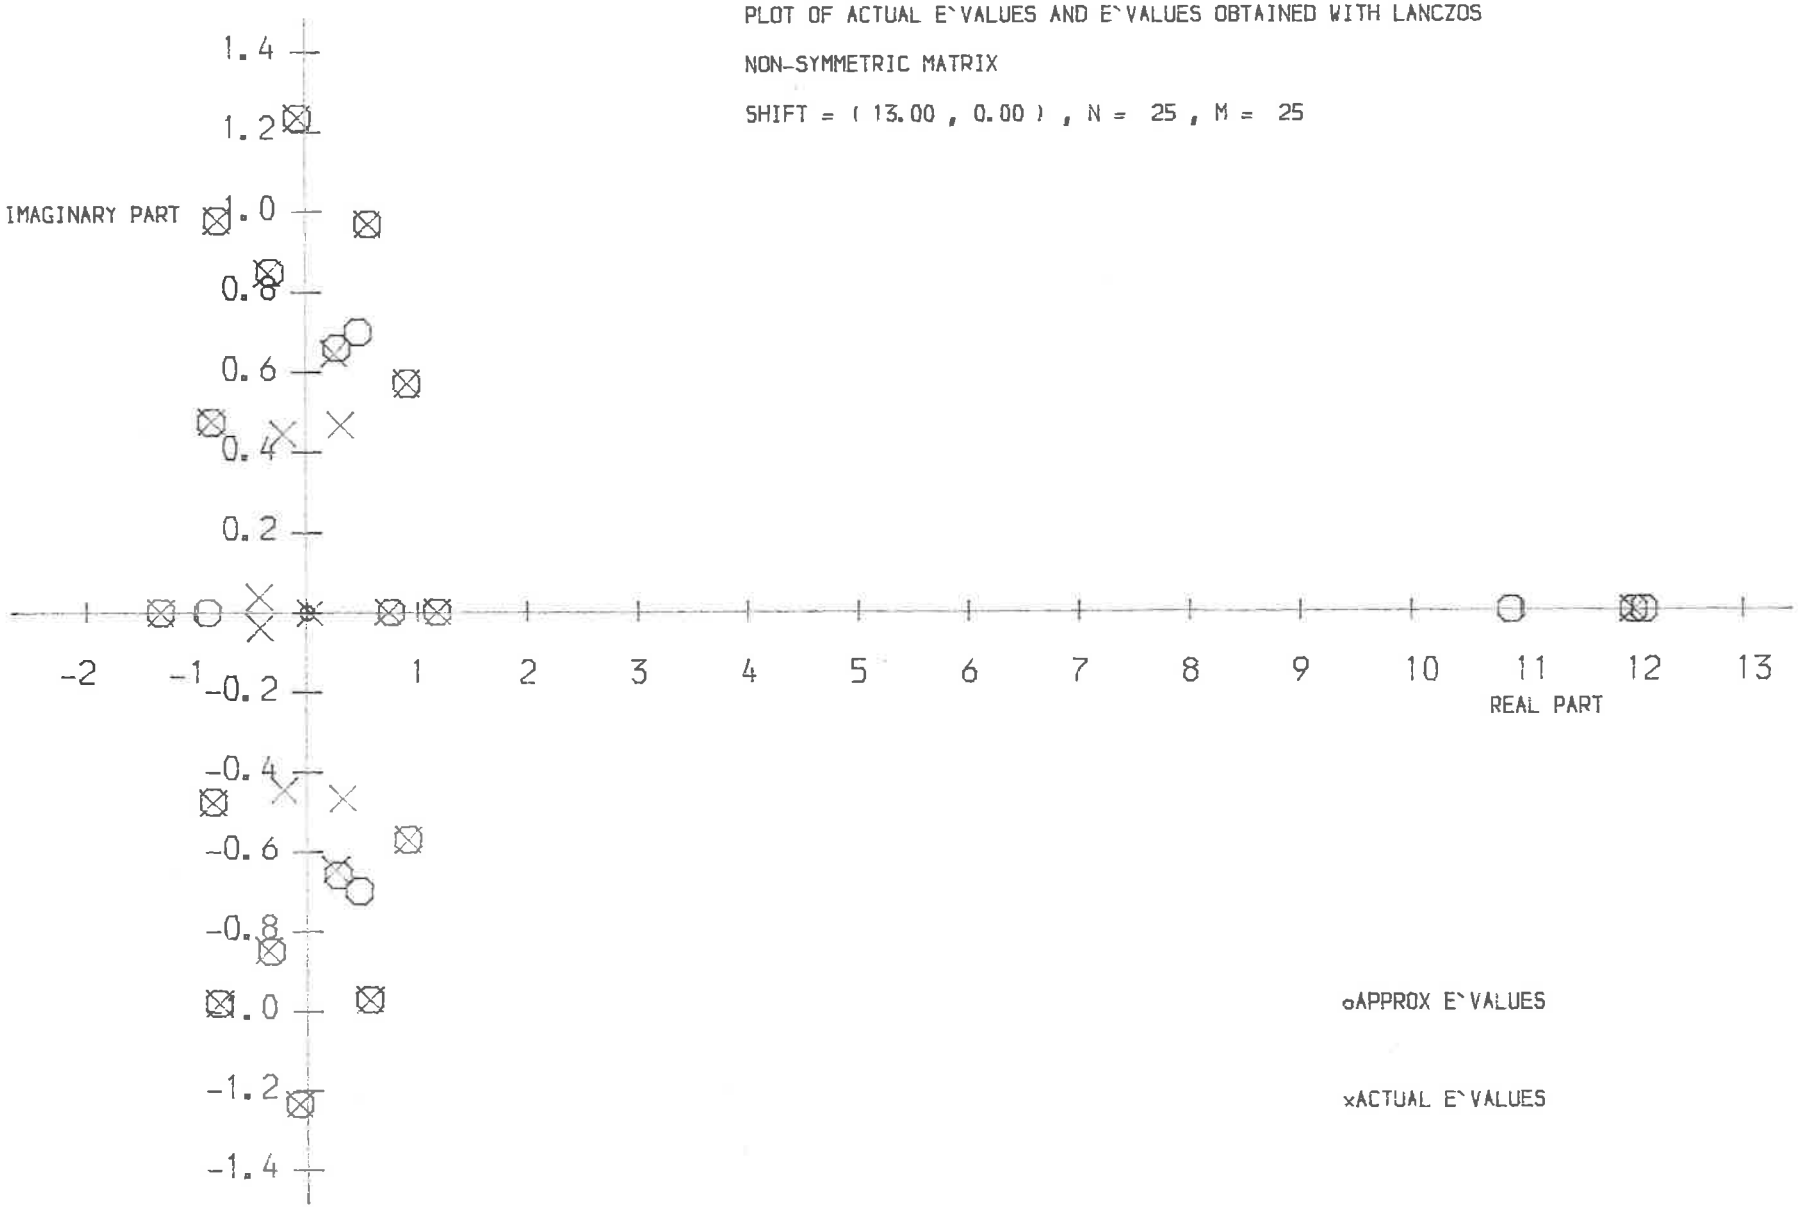
PLOT OF ACTUAL E`VALUES AND E`VALUES OBTAINED WITH LANCZOS
NON-SYMMETRIC MATRIX
SHIFT = (13.00 , 0.00) , N = 25 , M = 25
IMAGINARY PART
1.4
1.2
1.0
0.8
0.6
0.4
0.2
-0.2
-0.4
-0.6
-0.8
-1.0
-1.2
-1.4
-2
-1
1
2
3
4
5
6
7
8
9
10
11
12
13
REAL PART
oAPPROX E`VALUES
xACTUAL E`VALUES

FIGURE 2

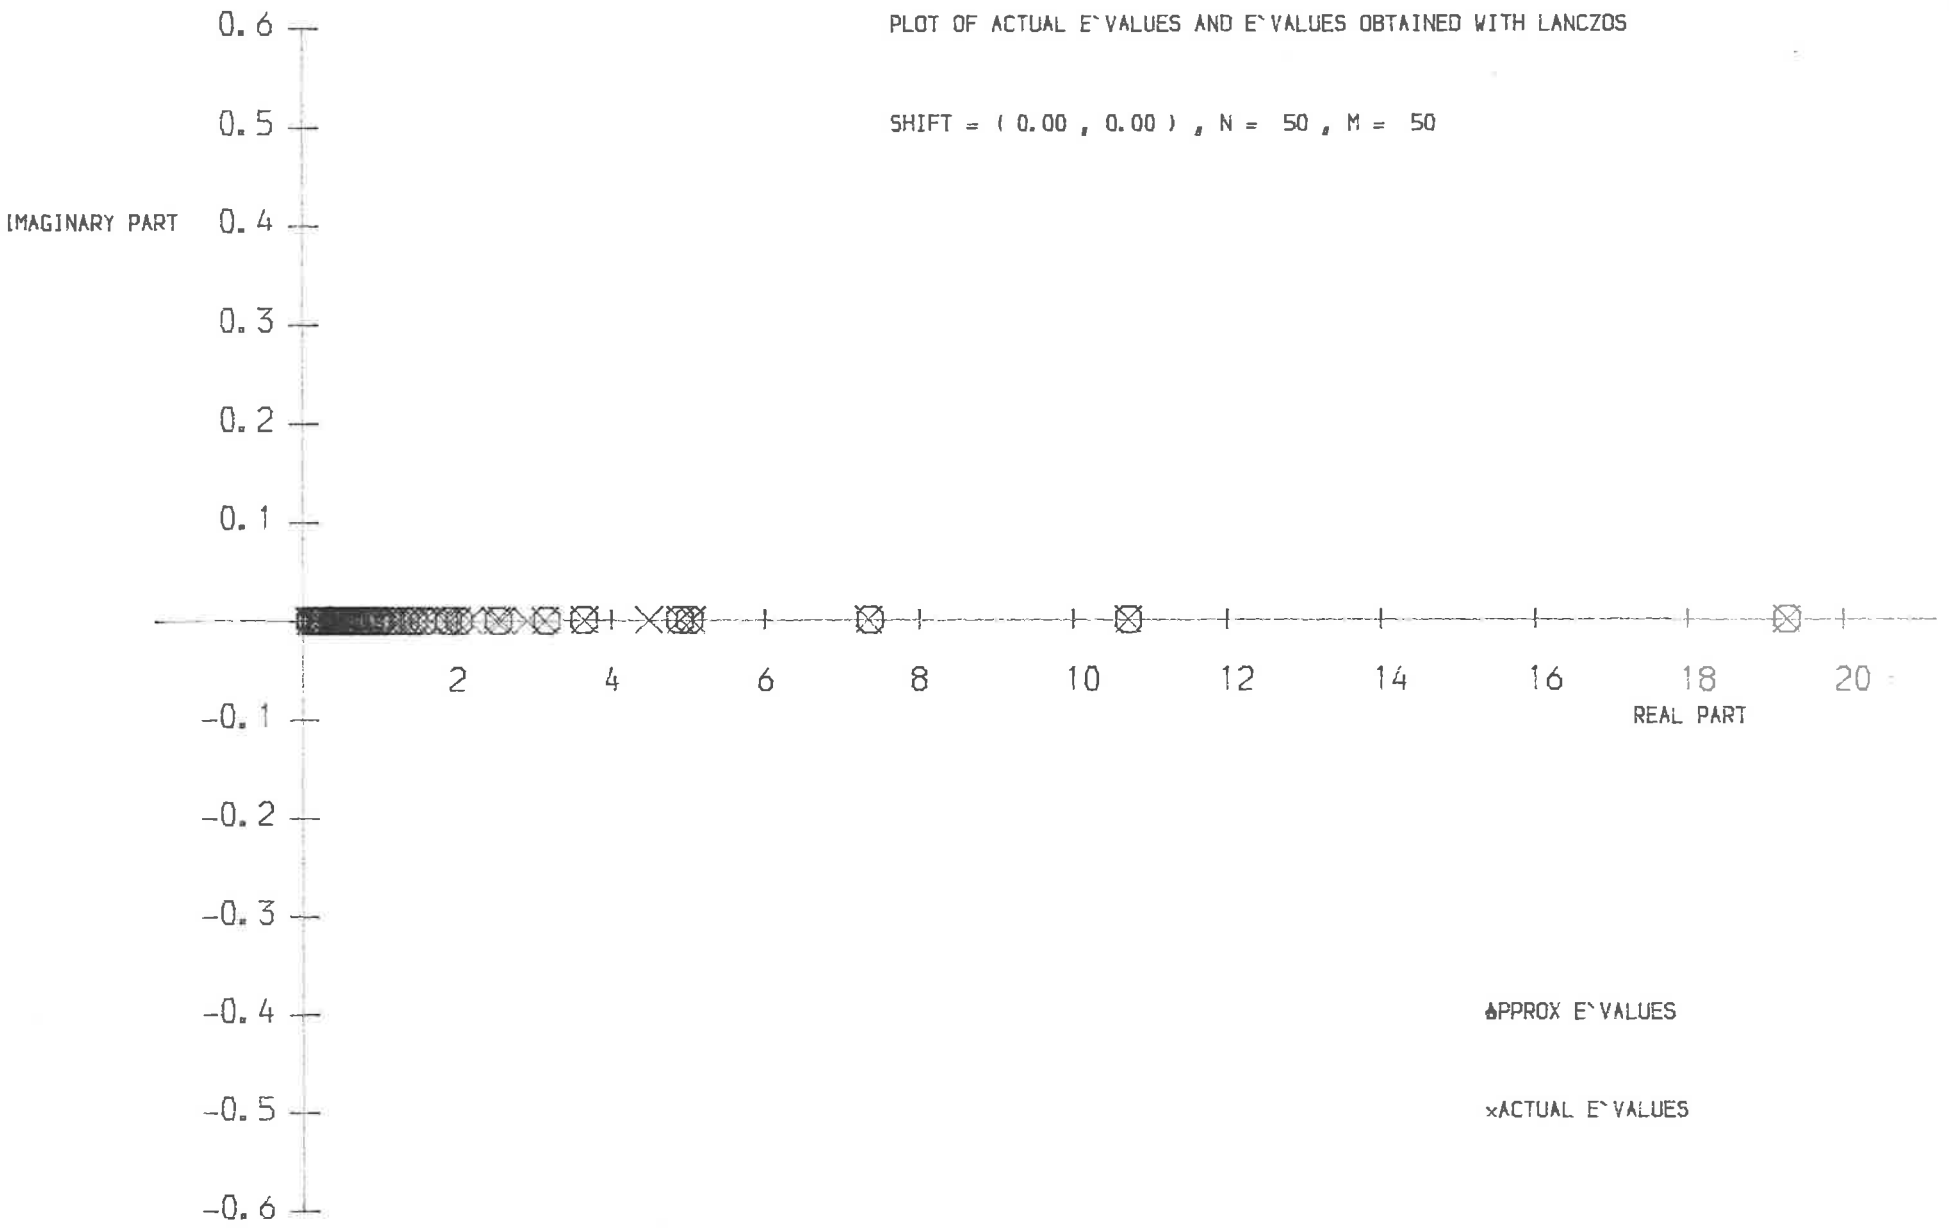
PLOT OF ACTUAL E`VALUES AND E`VALUES OBTAINED WITH LANCZOS
SHIFT = (0.00 , 0.00) , N = 50 , M = 50
IMAGINARY PART
0.6
0.5
0.4
0.3
0.2
0.1
-0.1
-0.2
-0.3
-0.4
-0.5
-0.6
2
4
6
8
10
12
14
16
18
20
REAL PART
APPROX E`VALUES
ACTUAL E`VALUES

FIGURE 3

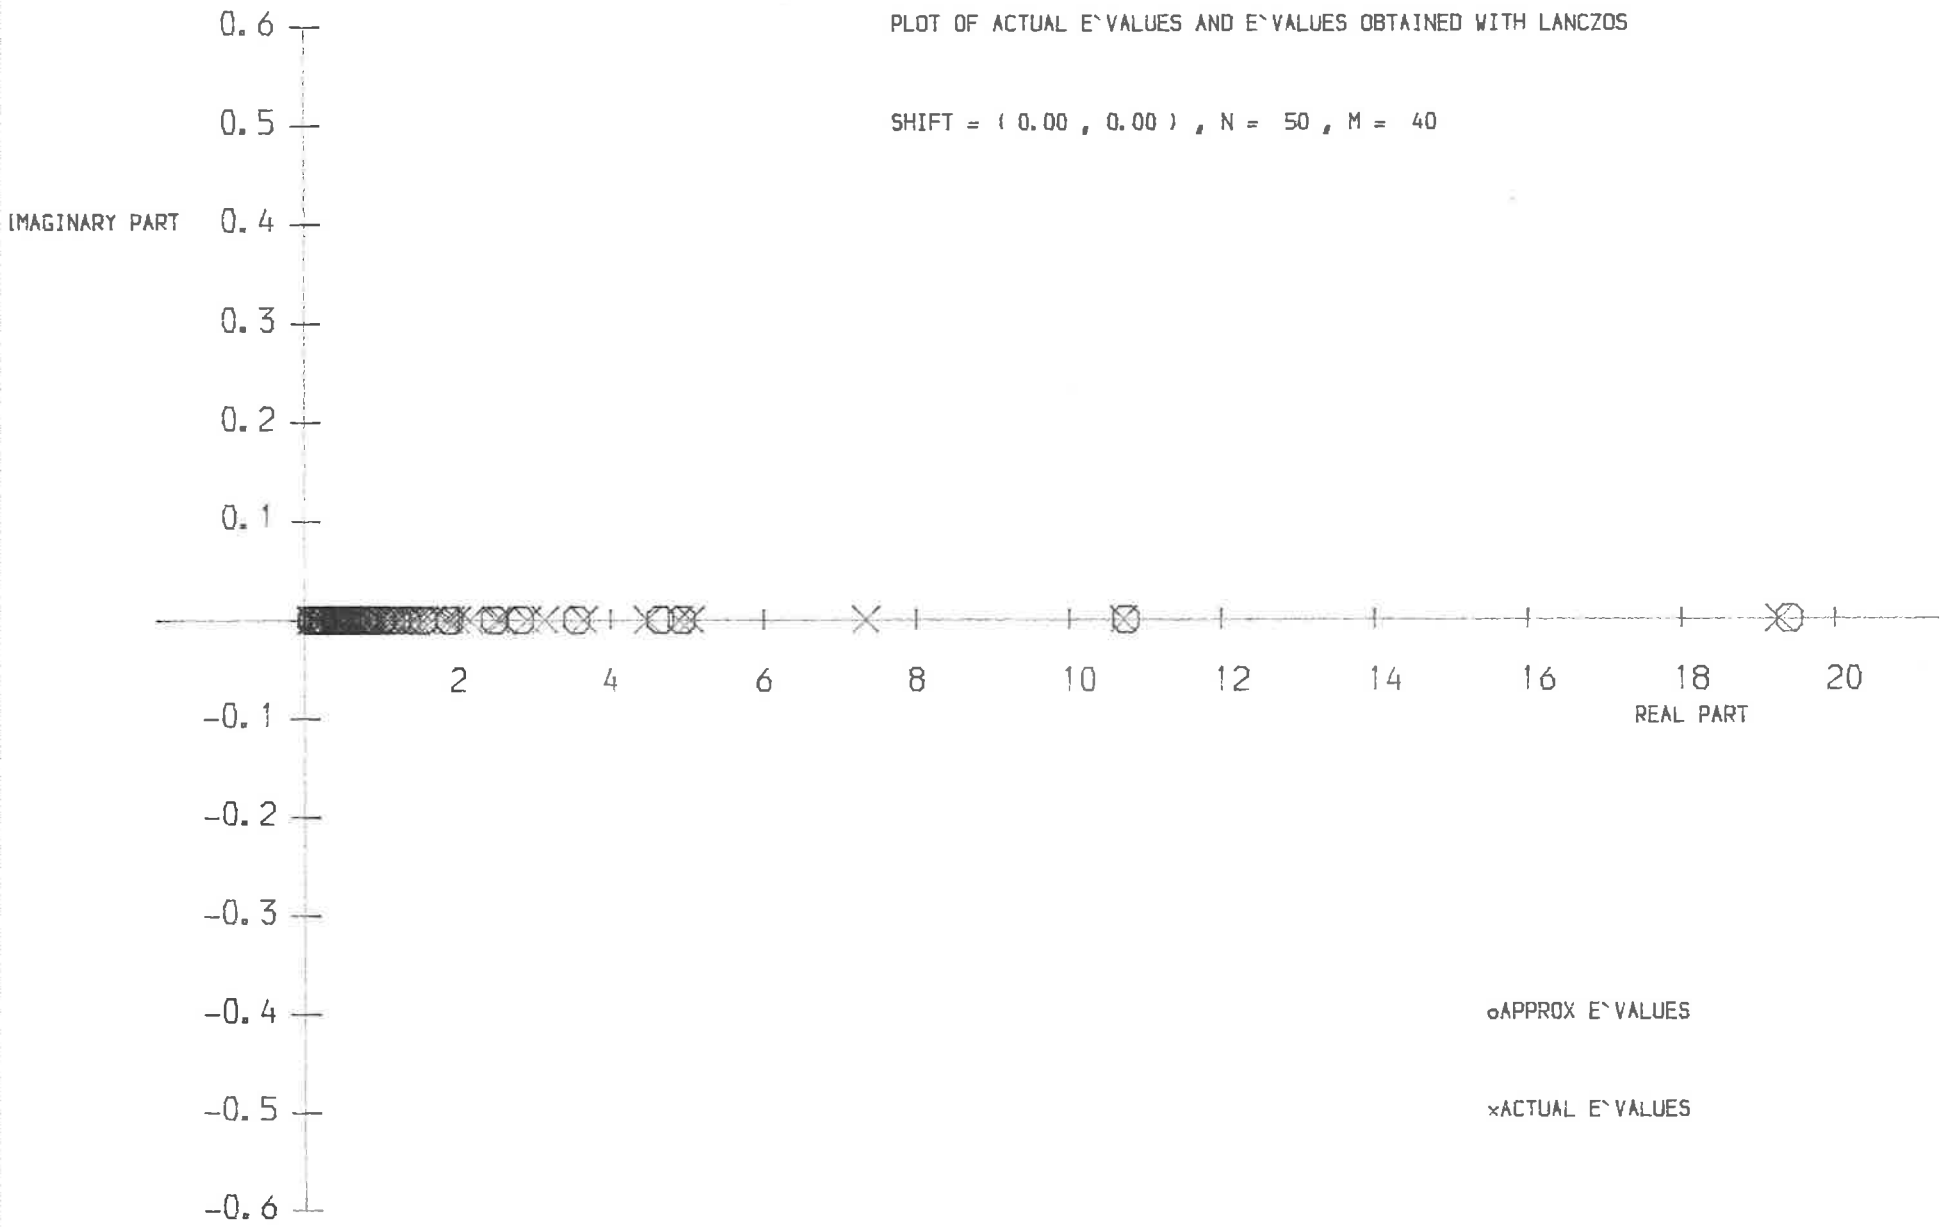
PLOT OF ACTUAL E`VALUES AND E`VALUES OBTAINED WITH LANCZOS
SHIFT = (0.00 , 0.00) , N = 50 , M = 40
IMAGINARY PART
0.6
0.5
0.4
0.3
0.2
0.1
-0.1
-0.2
-0.3
-0.4
-0.5
-0.6
2
4
6
8
10
12
14
16
18
20
REAL PART
oAPPROX E`VALUES
xACTUAL E`VALUES

FIGURE 4

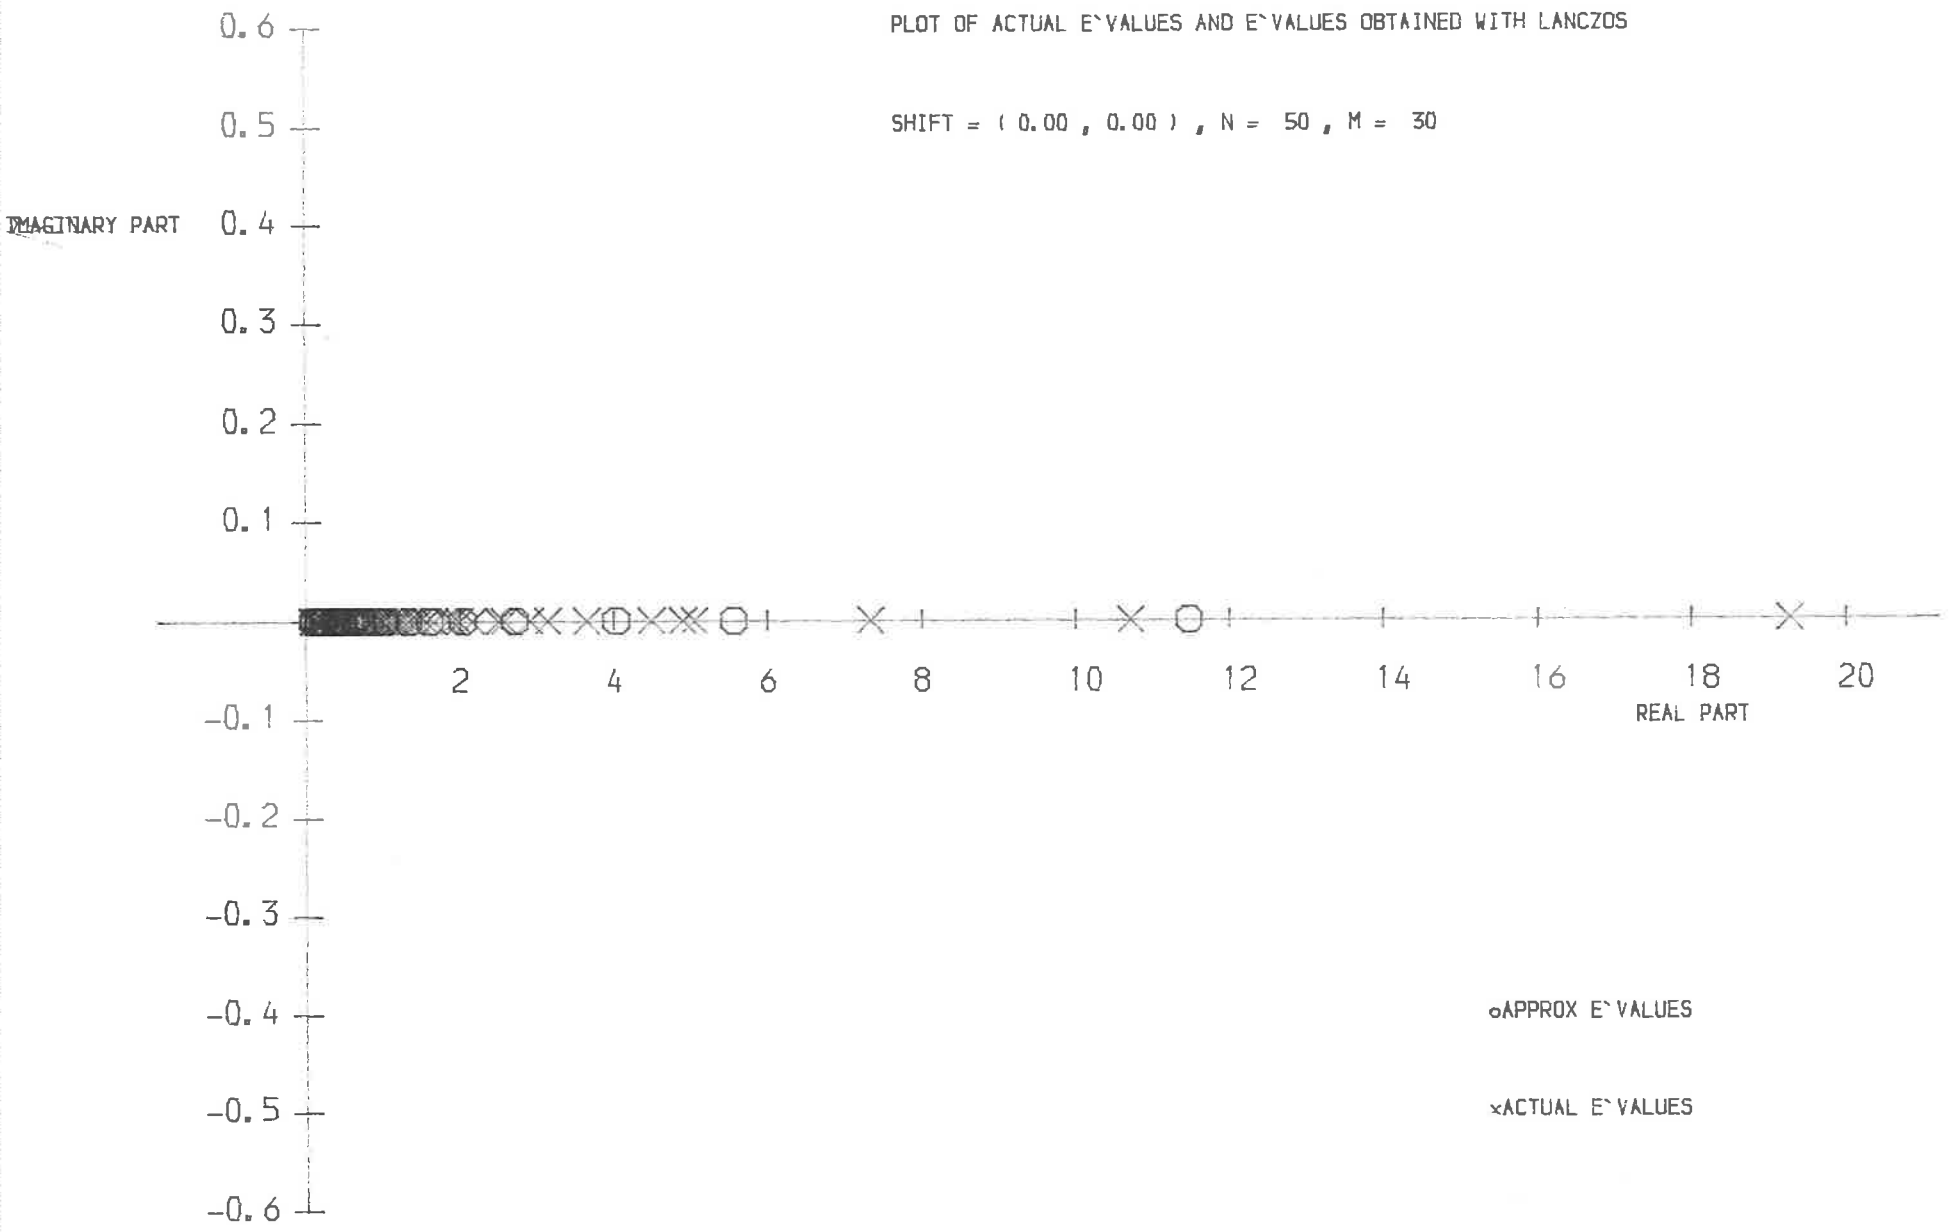
PLOT OF ACTUAL E`VALUES AND E`VALUES OBTAINED WITH LANCZOS
SHIFT = (0.00 , 0.00) , N = 50 , M = 30
IMAGINARY PART
0.6
0.5
0.4
0.3
0.2
0.1
-0.1
-0.2
-0.3
-0.4
-0.5
-0.6
2
4
6
8
10
12
14
16
18
20
REAL PART
oAPPROX E`VALUES
xACTUAL E`VALUES

FIGURE 5

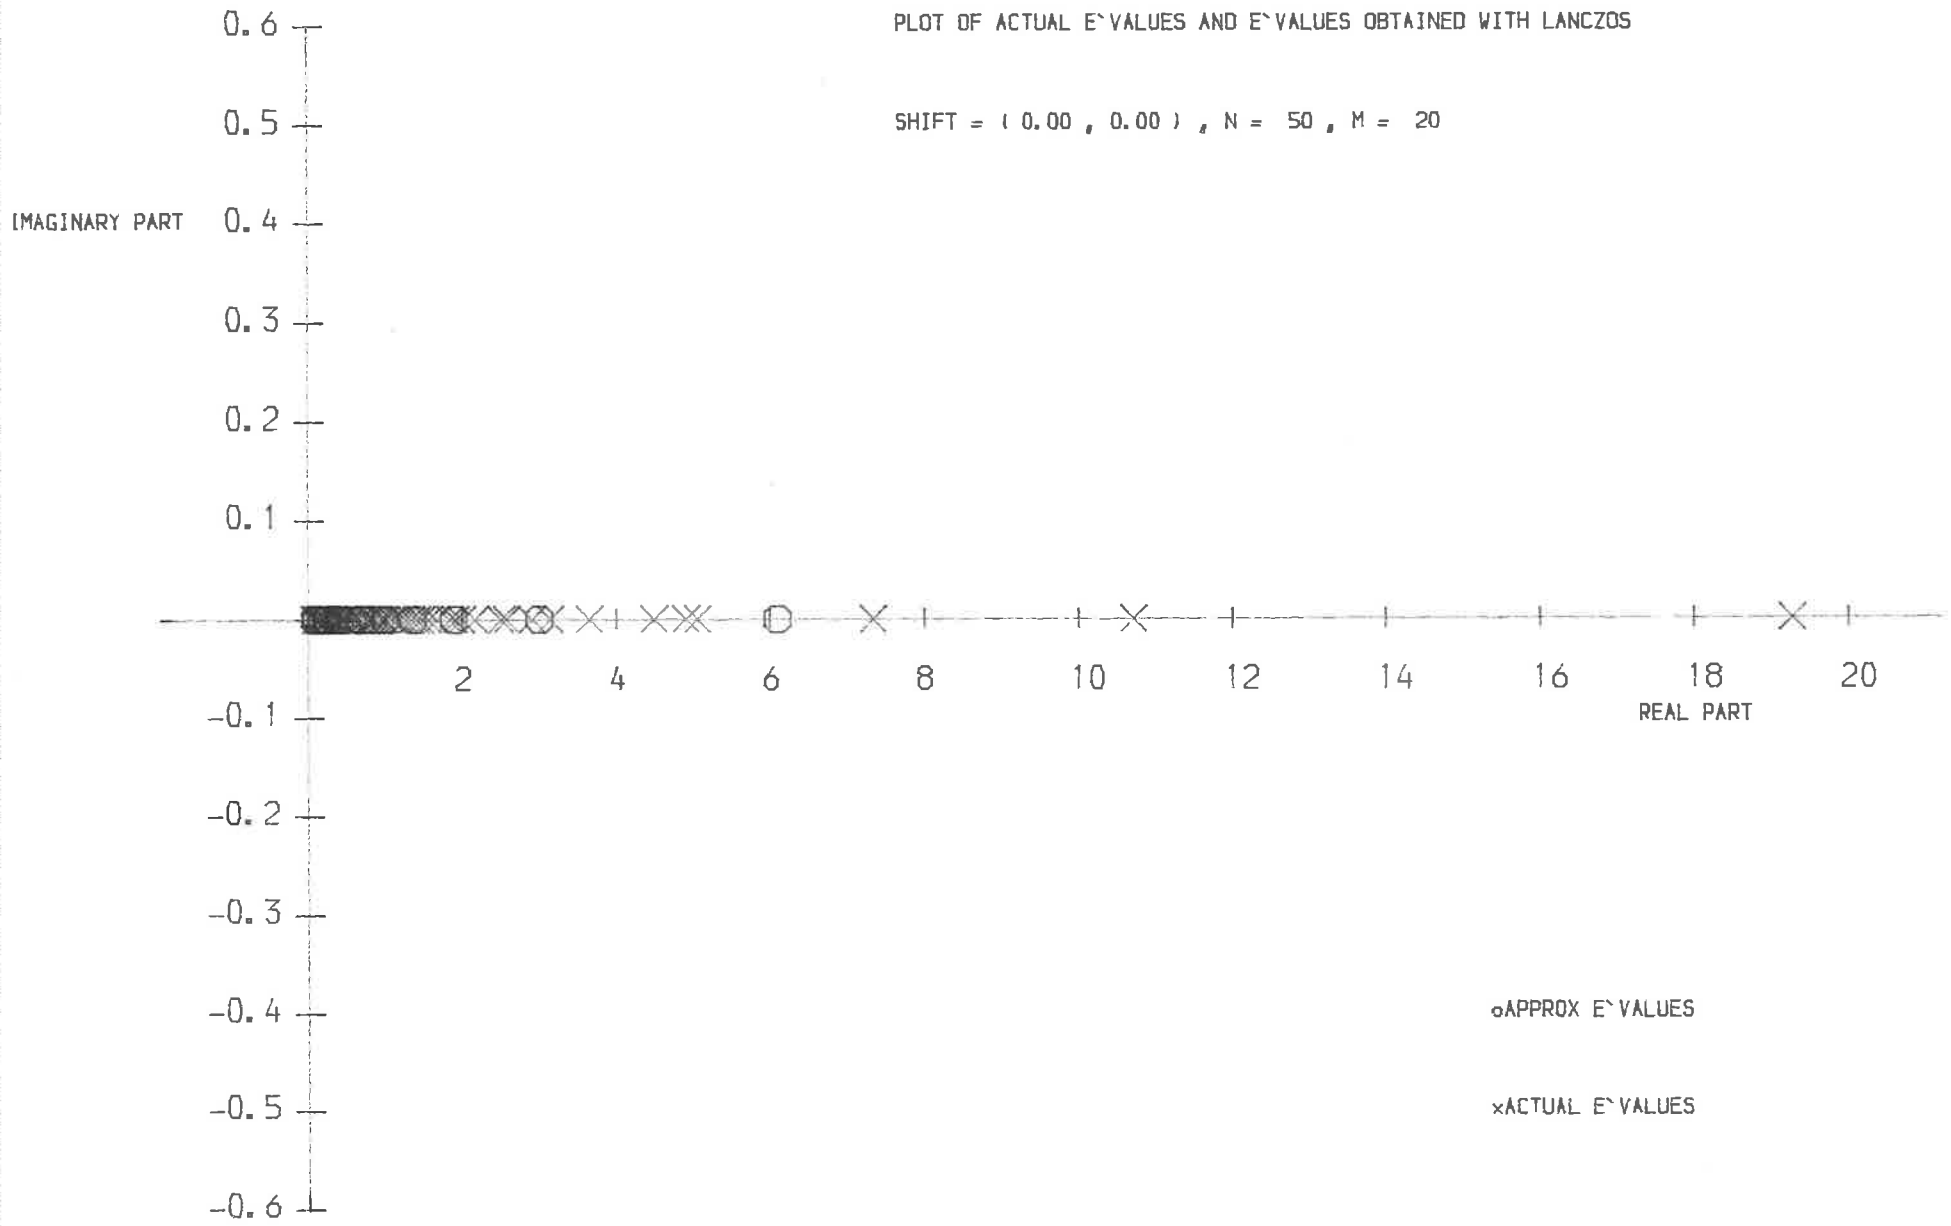
PLOT OF ACTUAL E`VALUES AND E`VALUES OBTAINED WITH LANCZOS
SHIFT = (0.00 , 0.00) , N = 50 , M = 20
IMAGINARY PART
0.6
0.5
0.4
0.3
0.2
0.1
-0.1
-0.2
-0.3
-0.4
-0.5
-0.6
2
4
6
8
10
12
14
16
18
20
REAL PART
oAPPROX E`VALUES
xACTUAL E`VALUES

FIGURE 6

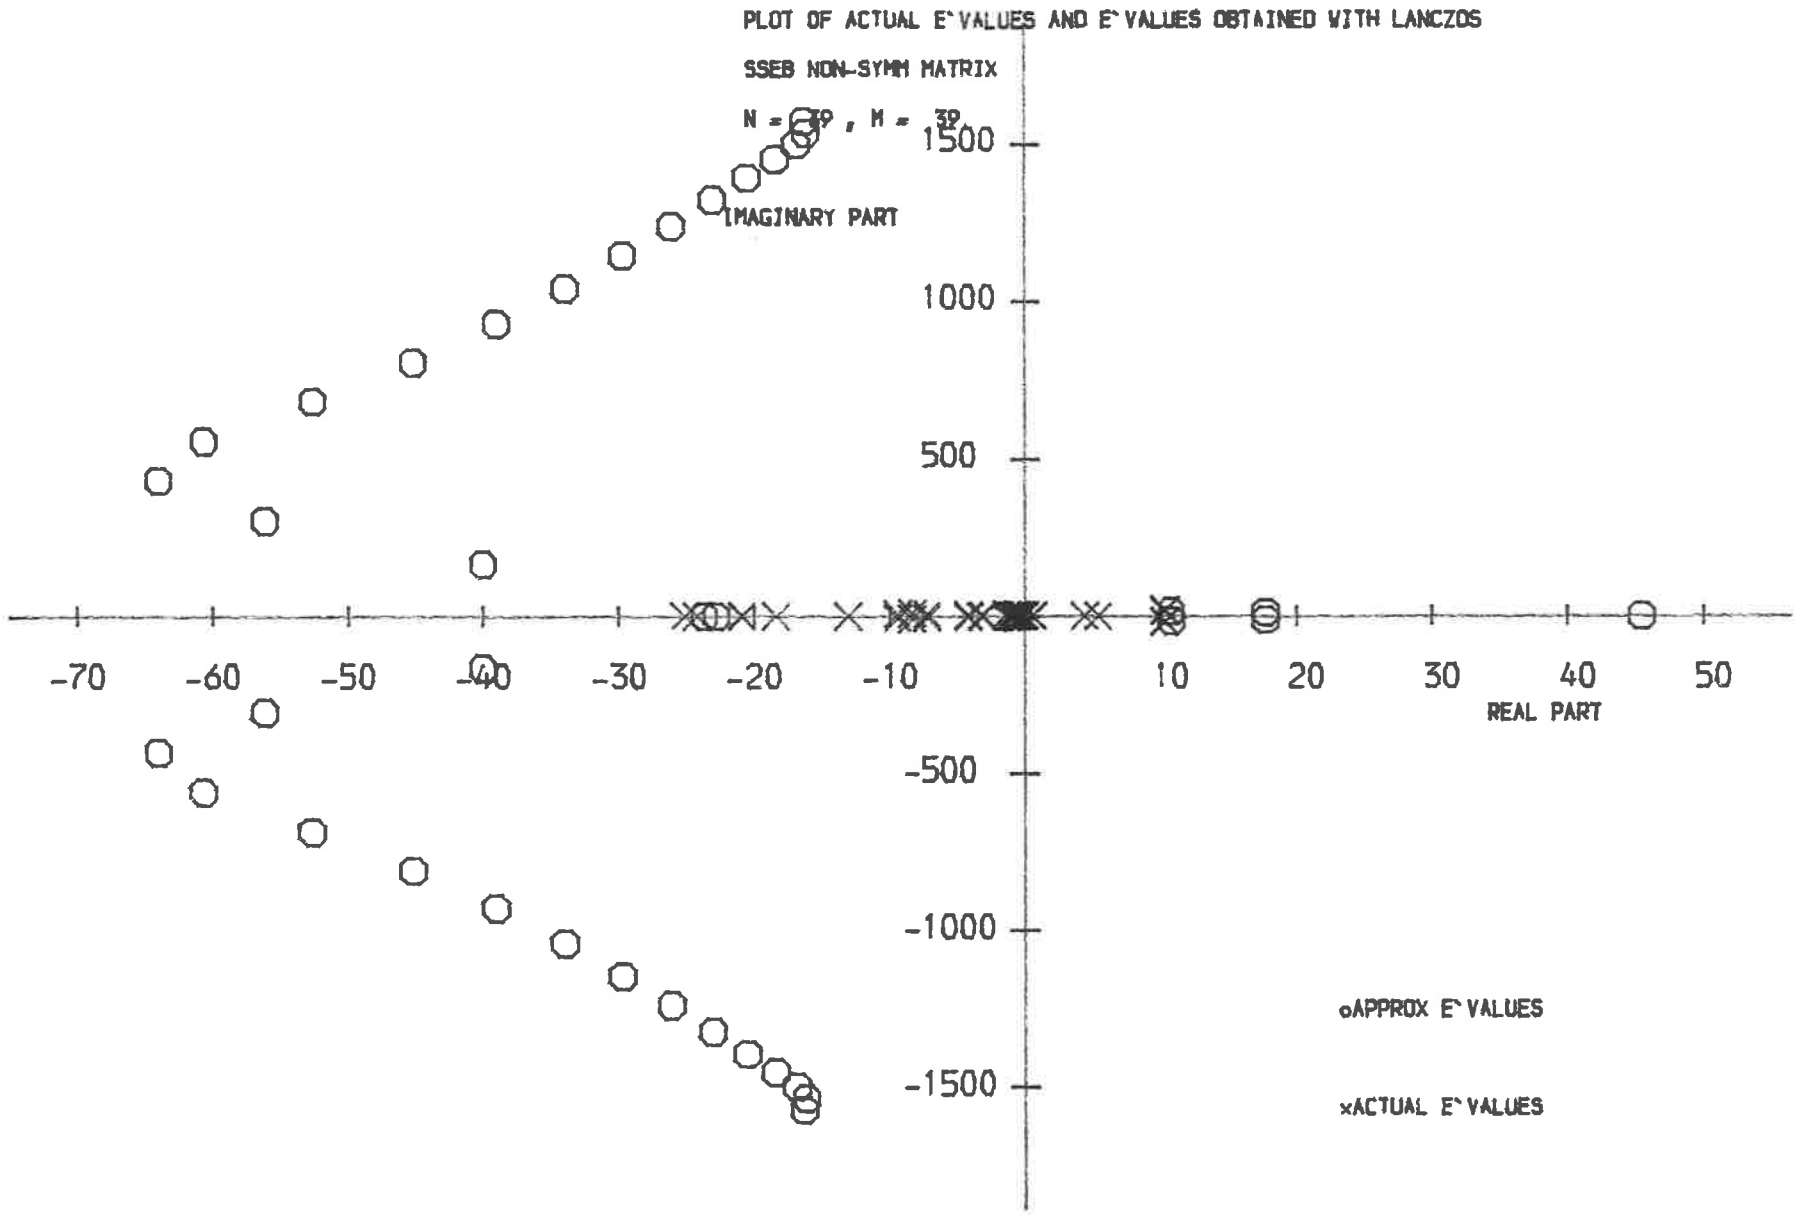
PLOT OF ACTUAL E`VALUES AND E`VALUES OBTAINED WITH LANCZOS
SSEB NON-SYMM MATRIX
N = 39 , M = 39
IMAGINARY PART
1500
1000
500
-500
-1000
-1500
-70
-60
-50
-40
-30
-20
-10
10
20
30
40
50
REAL PART
oAPPROX E`VALUES
xACTUAL E`VALUES

FIGURE 7

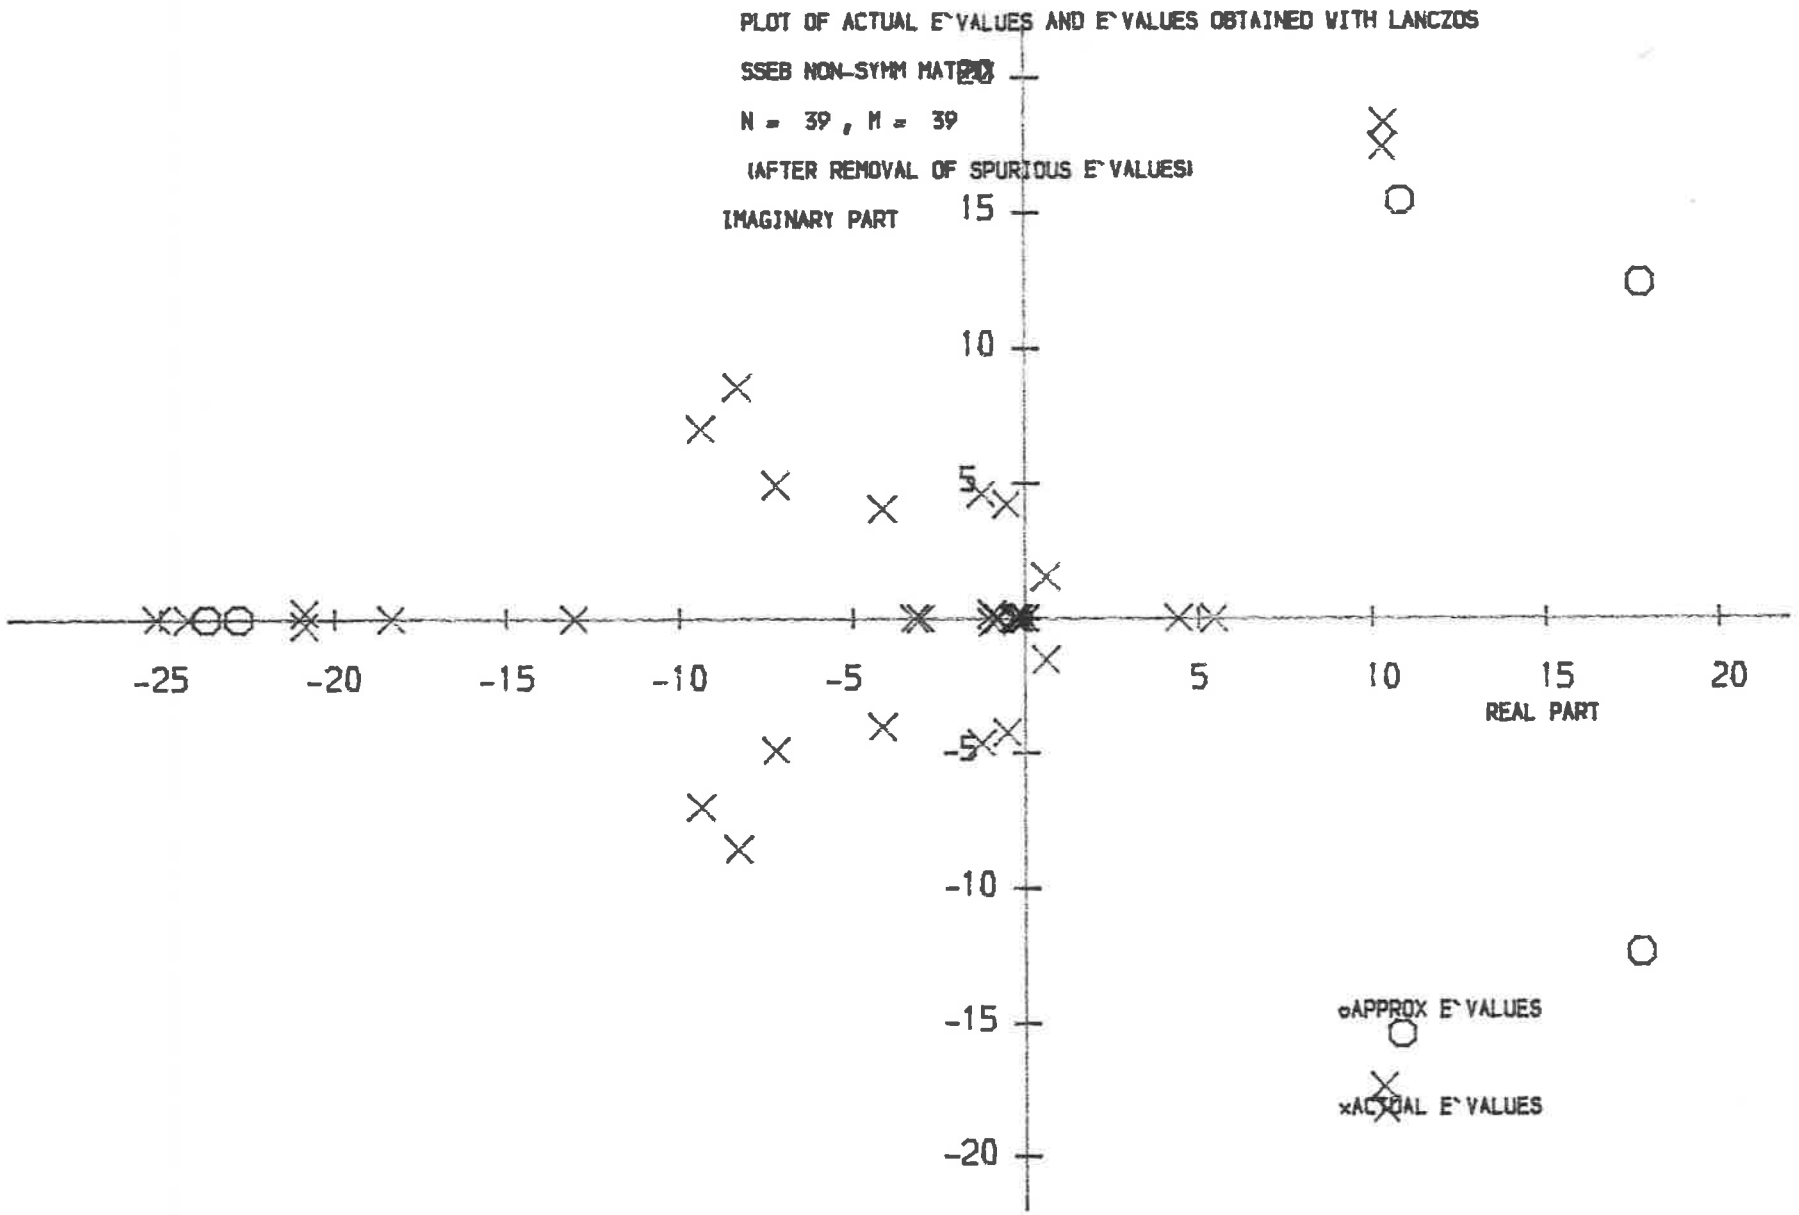
PLOT OF ACTUAL E`VALUES AND E`VALUES OBTAINED WITH LANCZOS
SSEB NON-SYMM MAT
N = 39 , M = 39
(AFTER REMOVAL OF SPURIOUS E`VALUES)
IMAGINARY PART
15
10
5
-5
-10
-15
-20
-25
-20
-15
-10
-5
5
10
15
20
REAL PART
oAPPROX E`VALUES
xACTUAL E`VALUES

FIGURE 8

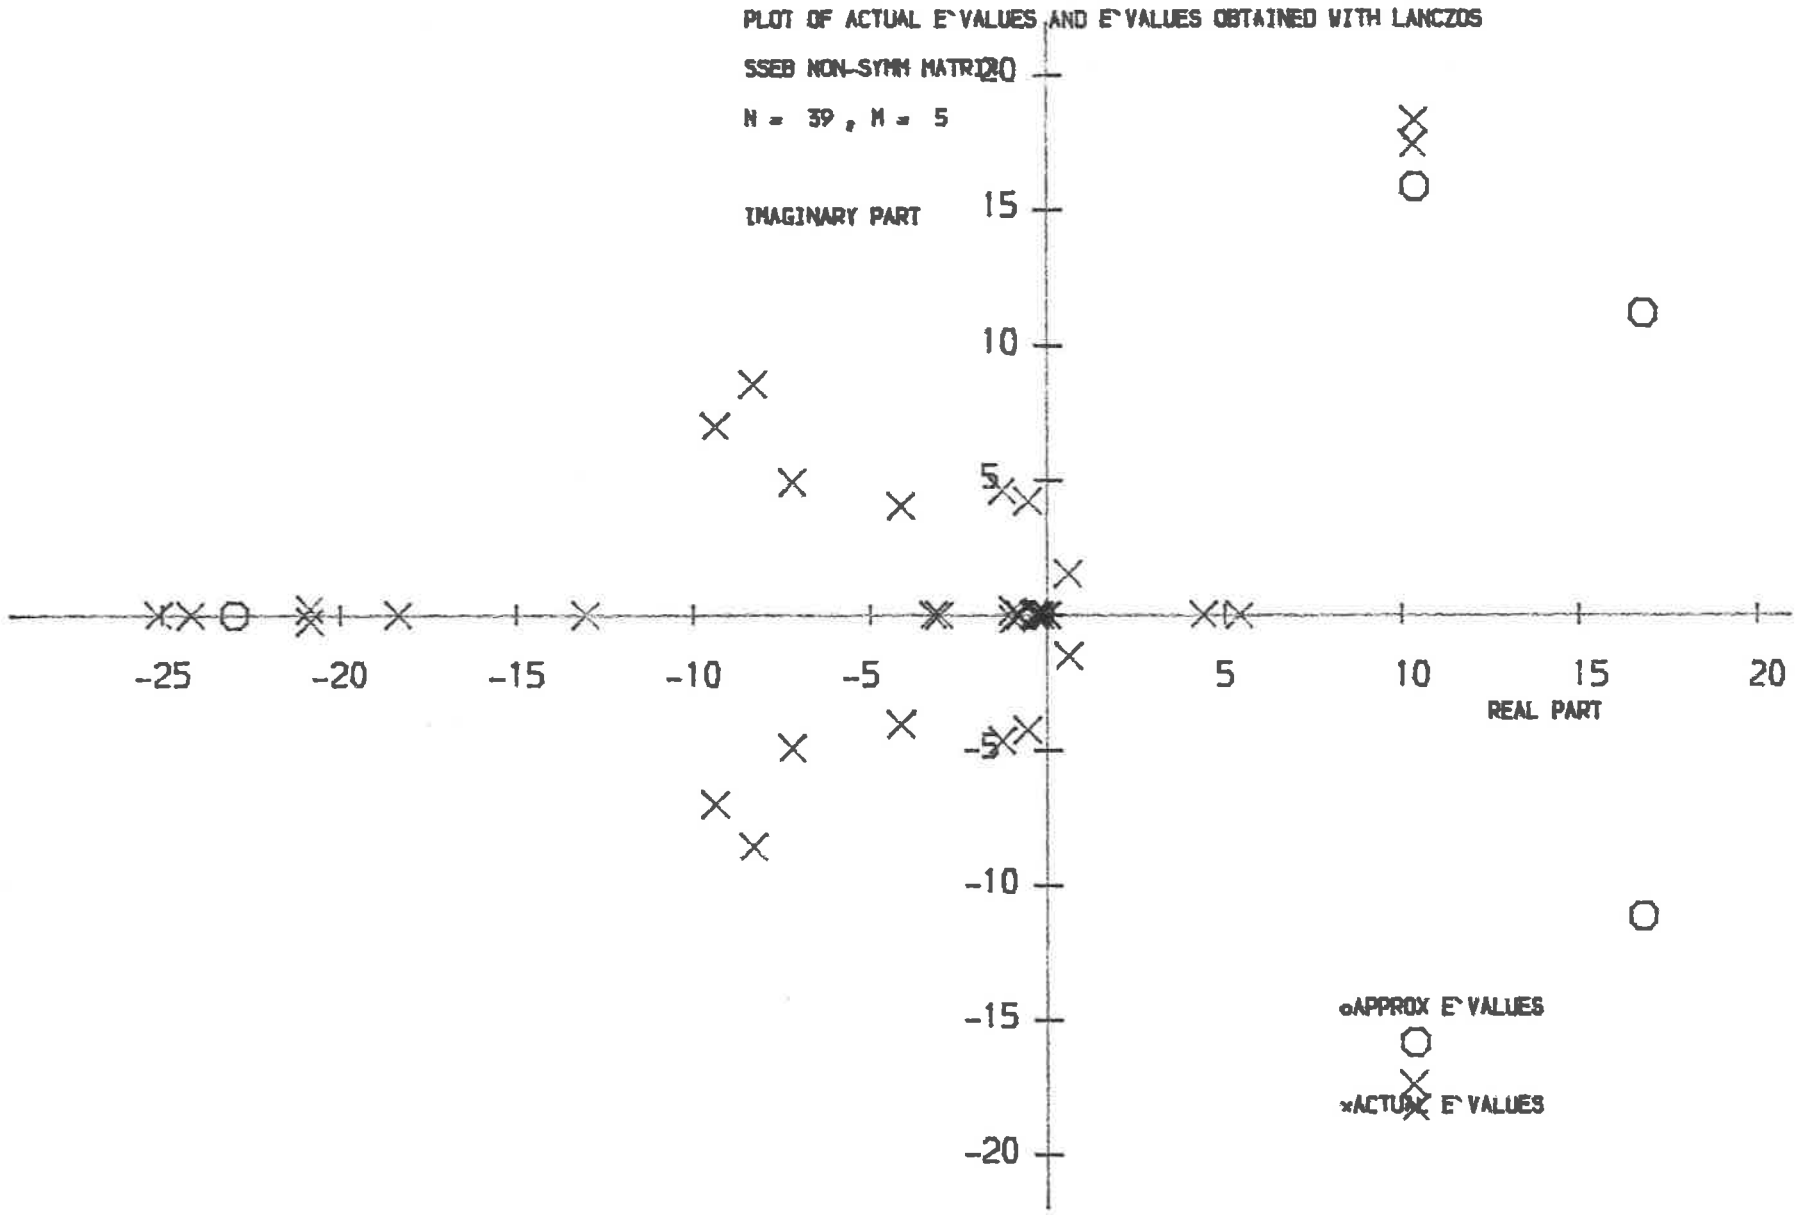
PLOT OF ACTUAL E`VALUES AND E`VALUES OBTAINED WITH LANCZOS
SSEB NON-SYMM MATRIX
N = 39 , M = 5
IMAGINARY PART
20
15
10
5
-5
-10
-15
-20
-25
-20
-15
-10
-5
5
10
15
20
REAL PART
oAPPROX E`VALUES
xACTUAL E`VALUES

FIGURE 9

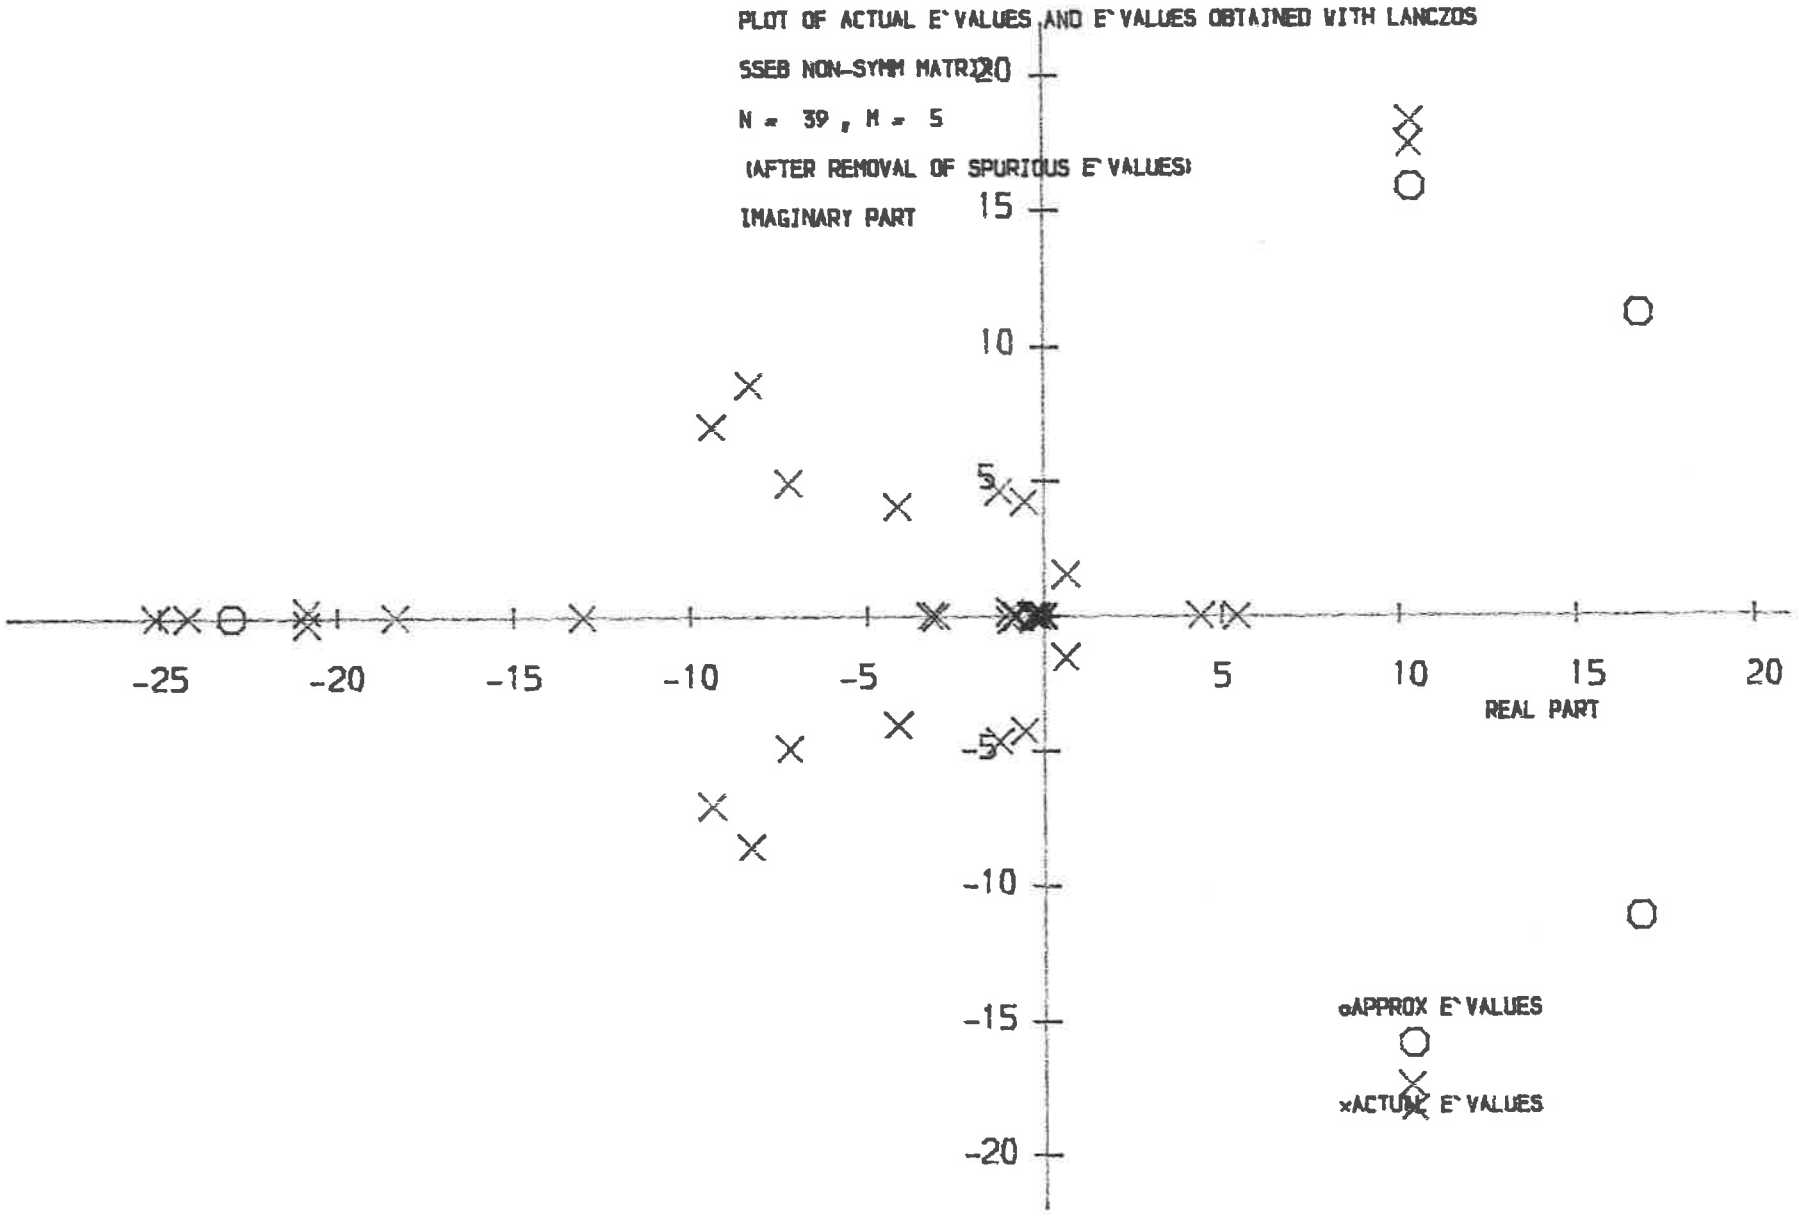
PLOT OF ACTUAL E`VALUES AND E`VALUES OBTAINED WITH LANCZOS
SSEB NON-SYMM MATRIX
N = 39 , M = 5
(AFTER REMOVAL OF SPURIOUS E`VALUES)
IMAGINARY PART
20
15
10
5
-5
-10
-15
-20
-25
-20
-15
-10
-5
5
10
15
20
REAL PART
oAPPROX E`VALUES
xACTUAL E`VALUES

FIGURE 10

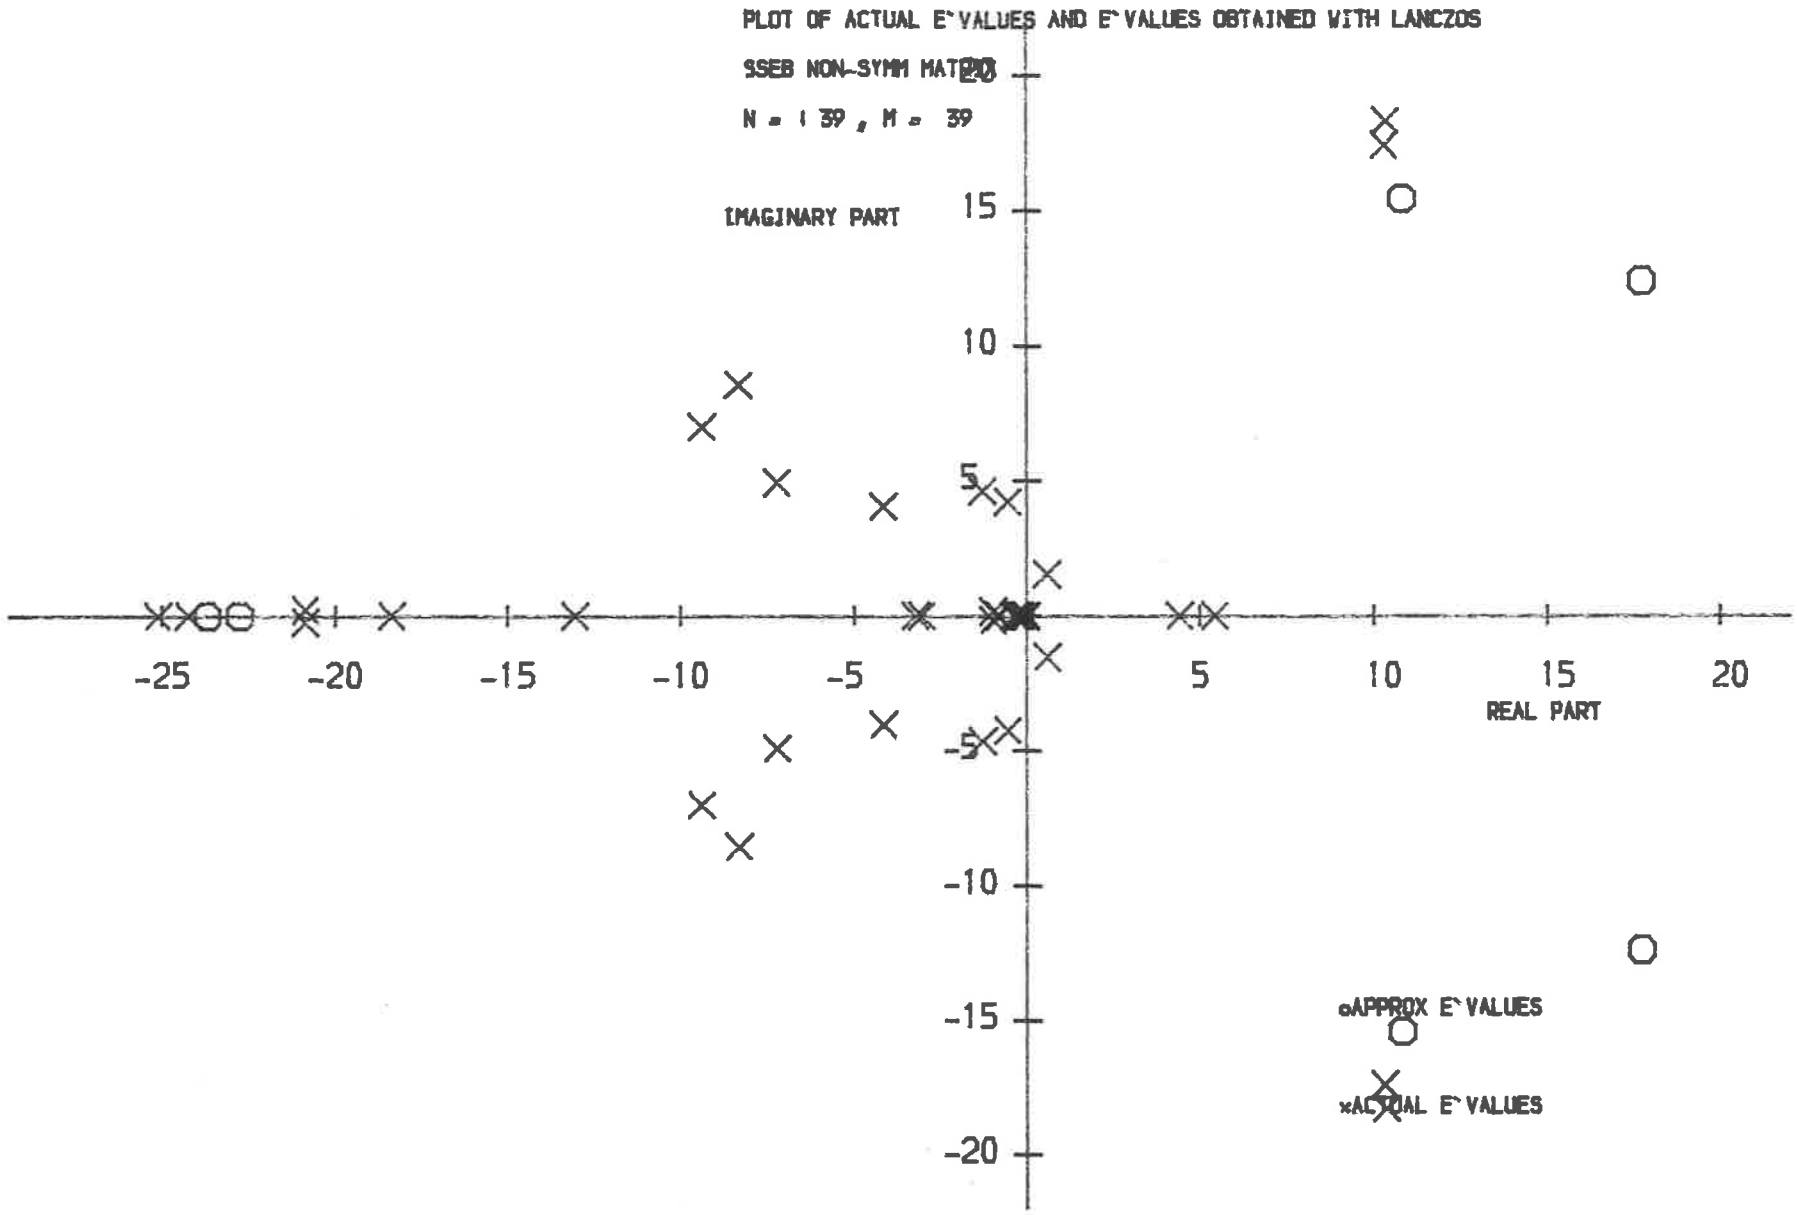
PLOT OF ACTUAL E`VALUES AND E`VALUES OBTAINED WITH LANCZOS
SSEB NON-SYMM MAT
N = 139 , M = 39
IMAGINARY PART
REAL PART
20
15
10
5
-5
-10
-15
-20
-25
-20
-15
-10
-5
5
10
15
20
oAPPROX E`VALUES
xACTUAL E`VALUES

FIGURE 11

5 CONCLUSIONS

Observation of the shifted inverse formulation points to an area of possible concern. Logically it would seem that a 'good' choice of δ, i.e. one that happens to coincide with, or nearly with, an actual eigenvalue of the system, would, by the very definition of the problem, cause the shifted, inverted matrix to be singular. A contradiction of sorts appears to arise. A choice of shift which should increase the rate of convergence would seem to promote the opposite. Fortunately, Wilkinson (1965, Ch.9, Sec.49) has shown that such a choice has only positive effects and any doubts regarding the production of singular matrices may be discounted.

Comparison with the QR method, it being the main alternative technique, shows that, while Lanczos inevitably loses some eigenvalues in deflating the original matrix to tri-diagonal form, its advantage lies in systems where the matrix is sparse and/or where only a few eigenvalues are desired. In such cases, the number of floating point operations required for the QR is $O(n^3)$ compared to the Lanczos process which employs $O(n)$ flops to obtain the tri-diagonal form, followed by a further $O(n)$ flops to compute the eigenvalues using a QR method which takes advantage of the tri-diagonal structure. Naturally, the workload is greatly reduced by the implementation of the shifted, inverse form, which 'relabels' the eigenvalue of interest as the dominant eigenvalue so that an 'incomplete' Lanczos matrix may be calculated, which, in theory, contains those eigenvalues of interest.

Results obtained by, amongst others, Cullum and Willoughby (1986) are very encouraging, but may display the method under favourable conditions. Results obtained to date by myself using this shifted, inverse form are less encouraging, in particular for the state matrix of the SSEB system which is very badly conditioned and may account for the lack of accuracy obtained. However, I would find it hard to maintain confidence in the method for matrices which are other than symmetric. Alternatives to the version defined by equations (2.14) are available and include a similar process which employs reorthogonalisation at every step, obviously advantageous in that accuracy is greatly enhanced, but clearly vastly more expensive. Block Lanczos methods are also in evidence which ultimately yield block tri-diagonal matrices and are especially useful for computing a few of the extreme eigenvalues of a sparse matrix.

ACKNOWLEDGEMENTS

I would like to acknowledge the financial support of the SSEB and the help and advice of Dr.N.K.Nichols with especial thanks for the time and encouragement offered by Dr.P.K.Sweby.

REFERENCES

Cullum J. and Willoughby R.A., Lanczos Methods for Large Symmetric Eigenvalue Computations, Birkhauser-Boston, January 1985.

Cullum J. and Willoughby R.A., A Practical Procedure for Computing Eigenvalues of Large Sparse Nonsymmetric Matrices, Large Scale Eigenvalue Problems, North-Holland Publ.,1986.

Golub G. and van Loan C., Sparse Matrix Computations

Kahan W., Parlett B.N. and Jiang E., Residual bounds on approximate eigensystems of nonnormal matrices, SIAM J. Numer. Anal. 19 pp.470-484, 1982.

Paige C.C., Error analysis of the Lanczos algorithm for tridiagonalizing a symmetric matrix, J. Inst. Math. Appl., 18, 341-349.

Parlett B.N. and Taylor D.R., A look-ahead Lanczos algorithm for unsymmetric matrices, Math. Comp (to appear), 1984.

Parlett B.N., Taylor D.R. and Lui Z-S., The look-ahead Lanczos algorithm for large unsymmetric eigenproblems, Proc. of INRIA 6th Int. Conf. on Comp. Methods in App. Sciences and Eng., Dec 12-16, 1983, Versailles, France.

Saad Y., Variations on Arnoldi's method for computing eigenelements of large unsymmetric matrices, Linear Algebra Appl., 34 pp.269-295, 1980.

Stewart W.J. and Jennings A., A Simultaneous iteration algorithm for real matrices, ACM Trans. Math. Software, 7(2) pp.184-198, 1981.
Wilkinson,J., The Algebraic Eigenvalue Problem, O.U.P.,1965.