Section 3

VALIDATION OF EISPACK

As part of its original development, EISPACK was subjected to thorough testing designed to exercise every statement in each subroutine and path in the package. Before a version of EISPACK is declared certified, however, it must also be validated at one or more test sites not involved in the adaptation. This validation has two important goals. One is to insure that the adaptation has been carried out successfully and that EISPACK performs as anticipated on that computer system. The second has been of major importance in the development of EISPACK; it is to provide feedback on the ease of use of EISPACK and its documentation. The procedures included in and the philosophy behind this validation are discussed in further detail in ([3],[5]).

Thus, validation at test sites includes two kinds of testing. The first employs test drivers and a collection of 82 matrices supplied with EISPACK; its purpose is to verify that the subroutines have been adapted to the machine and system properly, and to obtain data which can later be used as a check at other locations having the same computer system. These test drivers and matrices are included on the tape from the distribution center as described in Section 5.

The second kind of testing is the more informal; it seeks to employ EISPACK in actual scientific computations to measure the ease of use of the routines and their documentation. By its very nature, of course, this informal testing is not repeatable.

In order to simplify the communication of the results of the formal testing, a measure of performance for EISPACK, based on the backward error analysis of Wilkinson and Reinsch ([1],[2]), was defined. It is computed as:

$$\mu = \max_{1 \le i \le N} \frac{||Az_i - \lambda_i z_i||}{10 \cdot N \cdot \varepsilon \cdot ||A|| \cdot ||z_i||}$$

where each pair λ_i and z_i is an eigenvalue and corresponding eigenvector computed by EISPACK for the matrix A of order N, and where ε is the precision of arithmetic on the test machine (called MACHEP in the subroutines).

The factor 10 in the formula for μ was chosen empirically to obtain the following criterion for the performance of EISPACK:

If μ is less than 1, EISPACK has performed satisfactorily (as well as can be expected according to the backward error analysis for the particular precision of arithmetic). If μ is greater than 100, the performance is deemed poor. Finally, if μ is between 1 and 100, the performance is progressively marginal, and further investigation might be in order to verify that no error has occurred.

This measure of performance is also useful as a check on the correct installation of EISPACK; the residuals $Az_i - \lambda_i z_i$ are very sensitive to small perturbations in λ_i and z_i, and hence so is μ. μ may thus reflect small changes in the performance of EISPACK caused either by changes in the subroutines themselves or by the hardware, operating system, compiler, or library with which they are used. Hence it provides an excellent "quick check" of the correct installation of EISPACK on a hardware-software system. If the values of μ obtained from the test cases exactly duplicate those obtained at the test site, it is virtually certain that EISPACK has been correctly installed. If, however, the values of μ differ, an error may have occurred in the transmission or implementation of EISPACK, or the hardware-software system may differ from that on which it was tested; such differences should probably be investigated.

The collection of 82 test matrices has been accumulated from various sources and provides a wide spectrum of test cases from trivial to pathological. Most values of the performance index μ for these test cases were less than 1 for the systems on which EISPACK has been certified. The weaker performances were generally observed in the paths where some eigenvectors were computed; that is, where the inverse iteration subroutines (INVIT, CINVIT, TINVIT) were called.