Technical Report CS74003-R

ON MAKING BAIRSTOW'S METHOD WORK

David A. Ault

Department of Computer Science

College of Arts and Sciences

Virginia Polytechnic Institute & State University

Blacksburg, Virginia 24061

March 1974

Key words and phrases: polynomial roots, Bairstow's method, approximation of polynomial roots, monitoring convergence, multiple roots.

each part of the program. Finally, the use of a language which is available on most computers is desirable. By avoiding special features of a language provided by one vendor and by restricting oneself to a common subset of a language, for example, Standard [12] or Basic [1] FORTRAN, the software becomes portable, and a potential user is not condemned to rewrite the program to be able to use it on a machine with a different compiler. These practices were followed in creating this program and the details are in [2]. The above comments apply to writing any software, not just numerical applications. The subject of this paper lies in the area between numerical analysis, i.e., the mathematical formulation of algorithms, and the computer programming techniques mentioned above.

There are several considerations that are not included in the above areas. One is the selection of the mathematical formulas which will minimize the possible generation of roundoff error. Another is the determination of when a number is zero, except for propagated roundoff error. Three others are the subject of this paper. That is, we consider several problems which are more closely related to the problem whose solution is sought and the numerical method used in the solution. The solution of polynomial equations by a Newton-type method needs a very accurate approximation to the desired root to assure convergence. If the approximation does not meet the strict criterion, then the iterations may actually diverge from the root. Moreover, it has been shown by Gabler [8] that if a polynomial has a single real root, the remaining being complex, then the iterations of Bairstow's method may remain away from a solution indefinitely. Further, if a solution is a multiple root of the polynomial, then it is also a root of the derivative of the polynomial and so the denominator in Newton's method approaches zero as the root is found. This magnifies the effect of the roundoff error in the numerator and limits the number of significant digits in the answer. These are the problems we will consider in the following paragraphs.

II. Obtaining Initial Starting Values

A. APPROXIMATION OF ONE ROOT

We would like to narrow the region over which to search for the existence of a real or complex root of the polynomial. There exist involved methods which will locate polynomial roots [11,pp. 355-359], but a relatively simple method that involved as few polynomial evaluations as possible and gave a reasonable guess as to the existence and location of a root was preferred. We were led to our choice of an algorithm by a discussion in Hamming [10, p. 105] of transformations on a polynomial under which the solutions are invariant. Assume $P(z)$ is a polynomial of degree n of the form

$$P(z) = \alpha_n z^n + \alpha_{n-1} z^{n-1} + \ldots + \alpha_1 z + \alpha_0,$$

where $z = x + yi$, x and y are real and $i = \sqrt{-1}$. The transformations are, for $c \neq 0$,

1. $P(z)$ into $cP(z)$,
2. $P(z)$ into $P(cz)$,

3. $P(z)$ into $z^nP(1/z)$.

The third transformation is used to divide the complex plane into two regions. The transformation is made by reversing the order of the coefficients of the polynomial. The roots of $z^nP(1/z)$ are the reciprocals of the roots of $P(z)$. Thus, given $r > 0$, a root of $P(z)$ lies in the region

$$\{z: |z| \leq r\} \tag{1}$$

or its reciprocal lies in the region

$$\{z: 1/|z| < 1/r\}. \tag{2}$$

The number r is chosen to be the absolute value of the geometric mean of the polynomial roots, that is, $r = |\alpha_0/\alpha_n|^{1/n}$. The author has shown in [2] that the geometric mean can be considered as the average distance of the roots from zero, whether the roots be real or complex. It is expected that, on the average, half of the roots will lie in the region (1) and half will lie in the region (2). Since complex roots occur in complex conjugate pairs, we need only consider that portion of each of the above regions for which the imaginary part of the complex number z is nonnegative Each of these two regions may be divided in half by considering the half where the real part of z is nonnegative and the half where it is negative as separate regions.

We now have four regions, each of which may contain one or more roots. To achieve an approximation to one of the roots, we select one of the four regions and approximate the absolute value of $P(z)$ over the region by a bivariate interpolating polynomial of the form

$$f(x,y) = a_0 + a_1x + a_2y + a_3xy + a_4x^2 + a_5y^2.$$

To record in the program that we have chosen the region where the real part of z is negative, the geometric mean r is made negative. Thus, if $s = |r|$, then $f(x,y)$ is chosen to be the bivariate polynomial which interpolates to $|P(z)|$ at the points in the complex plane 0, si/2, si, r/2, r and r + si. This leads to the system of equations: $P(0) = a_0$, and

$$b_1 = P(si/2) - a_0 = a_2s/2 + a_5s^2/4 \tag{3}$$

$$b_2 = P(si) - a_0 = a_2s + a_5s^2 \tag{4}$$

$$b_3 = P(r/2) - a_0 = a_1r/2 + a_4r^2/4 \tag{5}$$

$$b_4 = P(r) - a_0 = a_1r + a_4r^2 \tag{6}$$

$$b_5 = P(r+si) - a_0 = a_1r + a_2s + a_3rs + a_4r^2 + a_5s^2.$$

Equations (3) and (4) can be used to solve for a_2 and a_5:

$$a_2 = (4b_1 - b_2)/s = (4b_1 - b_2)/|r|,$$

$$a_5 = (2b_2 - 4b_1)/s^2 = (2b_2 - 4b_1)/r^2.$$

Similarly, equations (5) and (6) can be used to find a_1 and a_4:

$$a_1 = (4b_3 - b_4)/r \text{ and } a_4 = (2b_4 - 4b_3)/r^2.$$

Using these values for a_1, a_2, a_4 and a_5, we can solve for a_3:

$$a_3 = (b_5 - b_4 - b_2)/(r\cdot s) = (b_5 - (b_4 + b_2))/(r\cdot|r|).$$

The choice of the points of interpolation was made with two objectives in mind. First, choosing three of the six points along the real axis should aid in the better approximation of real roots. This is an important special case because a small imaginary part in the initial approximation of a single real root can lead to a poor approximation to a quadratic factor of P(z). Second, the choice of the points along the real axis and along the imaginary axis results in a system of linear equations which can be solved algebraically and the coefficients of the interpolating polynomial can be calculated in a few arithmetic operations. This avoids the need for another routine to solve a system of linear equations.

Once the coefficients of f(x,y) have been determined, the minimum is sought by first solving for the point at which both its partial derivatives are zero. This point is easily determined by setting the partial derivatives of f(x,y) equal to zero and solving the resulting system of two linear equations in two unknowns. If this solution does not lie in the chosen region, then the minimum of f(x,y) over the region is its minimum value along the boundary of the region. The minimum along each edge of the boundary may be found by restricting f(x,y) to that edge and locating the point at which the derivative of this restricted function is zero.

Since the bivariate polynomial f(x,y) approximates the absolute value of P(z) over a region, it should follow the general shape of $|P(z)|$ and a minimum of f(x,y) over that region should occur at a point near a zero of P(z). This does not say that f(x,y) will be zero at this point; it may be positive or negative and not even close to zero. This is not important to us. We only use the point at which f(x,y) achieves a minimum to approximate a zero of P(z). The value of P(z) at this approximation is tested and it is this value that should be close to zero. The reader may have noticed that the size of this region will depend upon the geometric mean r of the roots, which in turn depends upon the magnitude of the roots. One might expect that the smaller the region, the better the interpolating polynomial would approximate the behavior of $|P(z)|$ and hence the better the approximation to a root. This indeed has been our experience and so the test for the closeness to zero of $|P(z)|$ at the approximate root is in proportion to the width of the region. If $|P(z)|$ at the approximate root is not close to zero, another region is chosen.

B. CALCULATION OF A SECOND ROOT

We assume that one approximate root has been found. If this root has either a real or an imaginary part which is close to zero, then we would prefer that the approximate root be real or pure imaginary, so the "near zero" part is set to zero. If the resulting number is complex, then its complex conjugate is selected as the second root.

If the resulting number is real and nonzero, then we need another real root to be used to calculate an initial approximation to a quadratic factor of P(z). It is possible that there is only one real root. If this is the case, it will be detected in the Bairstow iteration and the real root determined by the method described in section III B. Assume that another real root does exist. Then the magnitude of each real root contributes to the magnitude of the geometric mean. Thus, the size of the second root may be approximated by dividing the square of the geometric mean by the magnitude of the known approximation. Note that r is the geometric mean of these two roots. The problem is to determine the sign of the approximation to the second root. Since the magnitude of the first approximate root is not greater than r, the second root will be at least as big as r, i.e., it is a larger root. The value of α_{n-1}/α_n is the negative of the sum of the real parts of all of the roots of P(z). If it is negative, the larger roots have a positive real part; if it is positive, the larger roots have a negative real part. So we choose the sign of the second approximation to be the opposite of the sign of α_{n-1}/α_n. If α_{n-1} is equal to zero, then the roots are either paired with real parts about the same size, only opposite in sign, or there are several roots with small real parts of one sign and another root with a large real part of the opposite sign. In either of the last two cases we want to choose an approximation of opposite sign and on the other side of the geometric mean from the first root. The two complex or two real roots are used to find the coefficients p and q of an approximate quadratic factor, $z^2 - pz - q$, of P(z). If the roots are x + yi and x - yi, then p = 2x and $q = -(x^2 + y^2)$. If the roots are x_1 and x_2, then $p = x_1 + x_2$ and $q = -x_1x_2$.

If the first approximate root is zero, then it is assumed that there are one or more roots near zero (all of the zero roots were eliminated early in the algorithm) and the coefficients p and q of $z^2 - pz - q$ are chosen to be zero.

C. EXAMPLES OF STARTING APPROXIMATIONS

To illustrate how the algorithms work to find initial starting approximations, we present two examples. Although all of the real arithmetic in the program is done in double precision (16 decimal digits on an IBM 360 or 370 series computer), our illustrations will be presented in at most three decimal digits.

In the first example, $P(z)$ is the polynomial

$$P(z) = z^4 + 2z^3 + 6z^2 + 8z + 8.$$

The initial approximation is the pure imaginary complex number $z = 1.7i$. Since it is complex, its complex conjugate is chosen as the other approximate root and the coefficients of the approximate quadratic factor of $P(z)$ are $p = 0$ and $q = -2.9$. The Bairstow iteration scheme converges to $p = 0$ and $q = -4$ (accurate to 15 decimal digits) in five iterations. This quadratic expression factors into the two complex roots $\pm 2i$.

In the second example, $P(z)$ is the polynomial

$$P(z) = z^4 - 6z^3 - 39z^2 - 16z + 60.$$

The initial approximation is $z = .75$. Since the geometric mean for this polynomial is 2.6, the magnitude of the second approximate root is $(2.6)^2/.75 = 9.0$. Since -6 is negative, the sign of the second root is made positive. The coefficients of an approximate quadratic factor of $P(z)$ are $p = 9.8$ and $q = 6.8$. The Bairstow iteration scheme converges to $p = 11.$ and $q = 10.$ in six iterations. This quadratic expression factors into the two real roots 1.0 and 10.

III. Monitoring Convergence

A. DIRECTING THE ITERATION TOWARD THE ROOT

Even if a "good" approximation is made to a quadratic factor of P(z), the Bairstow iteration scheme may step away from the solution (and away from any solution). It may even step towards the solution for a few iterations and then suddenly diverge. To keep the method on a convergent path, the magnitude of the polynomial is measured at each new approximation to p and q. If z_1 and z_2 are the two roots of $z^2 - pz - q$, then the measure is $|P(z_1) \cdot P(z_2)|$, see Hamming [10, p. 110] or Ault [2]. Beginning with the first approximation to p and q, the magnitude of the polynomial is kept for each successful iteration towards a solution. When the next iteration in calculated, the magnitude of P(z) is calculated at this new approximation. As long as the sequence of magnitudes is nonincreasing, the most recent one is saved and the Bairstow iteration scheme is allowed to continue. If the magnitude of P(z) suddenly increases, then it is assumed that the sequence is diverging. Two things are done to help the iteration scheme find a better path to a solution. First, the next approximation is taken to be half the previous step size from the previous approximation to p and q (this is the last accepted approximation). Second, the allowable maximum magnitude is doubled. The maximum allowable magnitude is increased in this manner to prevent the iteration from becoming stuck in one spot and failing to continue toward the solution.

There are several reasons why an iteration could step a large distance away from the solution. Loss of significant digits in subtractive cancellation could produce a meaningless denominator in the calculation of the step size for p and for q in Bairstow's method. The path of the iterations may land upon a point on the surface generated by $P(z)$ which is near a saddle point. This may cause the denominator in the calculation of the step size for p and for q to be small and produce an unusually large step. Third, the iterations may land near a line in the p,q-plane generated by a real root of $P(z)$. This topic will be discussed in more detail in part B of this section. This may cause a large step, usually along a line "parallel" to this line.

As an example of the manner in which this technique works, consider the polynomial

$$P(z) = z^7 - 3z^6 - 2z^5 + 20z^4 - 56z^3 + 68z^2 - 44z + 12.$$

The first approximation to a quadratic factor was $z^2 - z + .6$. After five iterations, the Bairstow iteration scheme converged to $z^2 - 1.2z + .7$. The reduced polynomial would be of degree 5. The first approximation yields a real root .75. The second real root is 2.75 and so the approximate quadratic factor is $z^2 - 3.5z + 2$. Even though the approximate real roots are close to the actual real roots (they are 0.64 and 2.84) the Bairstow iteration immediately begins to diverge. In figure 1, the initial starting value is labeled 0 and the first iteration is labeled 1. The step size is halved three times, each time beginning at the original approximation, until the iteration labeled 4 begins a new sequence 4,5,6,7,8,9, which converges to the solution. The points labeled 7 and 8 are not shown because of the large scale of the figure.

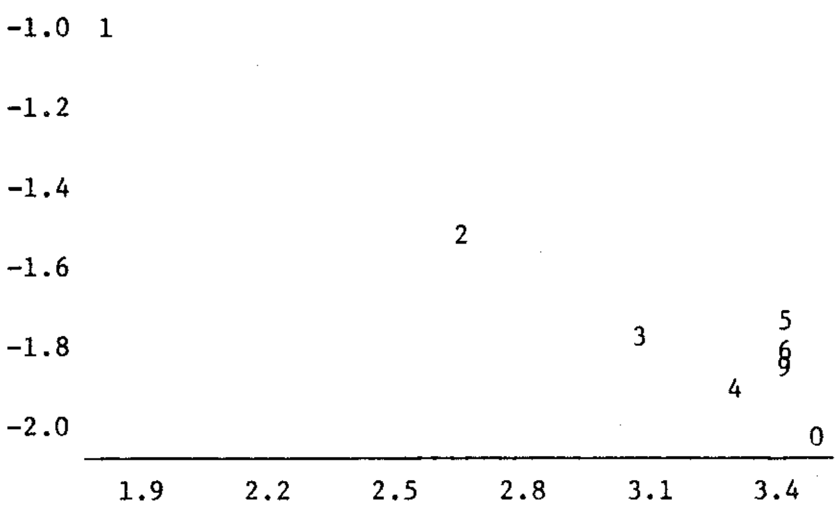

Figure 1

B. THE PROBLEM OF A SINGLE REAL ROOT

Even if the iteration steps are not diverging, they may never converge to a quadratic factor of $P(z)$. Let (p_i, q_i), $i = 0,1,2,\dots,$ denote the sequence of approximations to the coefficients of a quadratic factor $z^2 - pz - q$ of $P(z)$. It has been shown by Gabler [8] that if $P(z)$ has a real root $z = k$ and if for some index n,

$$k^2 - p_n k - q_n = 0,$$

then for the following index (and hence for all further indices)

$$k^2 - p_{n+1} k - q_{n+1} = 0.$$

Geometrically, this result states that if some iteration lands on the line

$$\{(p,q): k^2 - pk - q = 0\} \qquad (7)$$

in the p,q-plane, then all of the remaining iterations will be confined to that line. It has been observed that the iterations will "parallel" this line if they land close to, but not necessarily on, the line. The consequences of this theorem to our algorithm depend on the number of real roots of P(z).

If there is more than one real root, there is a line in the p,q-plane determined by each distinct root. If the root is a multiple root, then the line (7) at the point $p = 2k$ and $q = -k^2$ yields a quadratic factor of P(z). This is because $(z - k)^2 = z^2 - 2kz + k^2$. The lines given by (7) for any two distinct real roots intersect at a unique point. Suppose k_1 and k_2 are distinct real roots of P(z), then the point of intersection of the two lines determined by k_1 and k_2 is $p = k_1 + k_2$ and $q = -k_1k_2$. In each of these cases, we want the sequence of iterations to follow along one of the lines until it arrives at a solution. This will produce two real roots.

If P(z) has a single real root k, then the iterations along (7) will continue indefinitely because complex roots are found in complex conjugate pairs and $z - k$ is always a factor of $z^2 - pz - q$ whenever $k^2 - pk - q = 0$. Suppose the iterations land on the line (7) corresponding to the single real root k. For any quadratic expression $z^2 - pz - q$, there exists a polynomial Q(z) of degree $n - 2$ and real numbers β_1 and β_0 such that

$$P(z) = Q(z)(z^2 - pz - q) + \beta_1 z + \beta_0.$$

If (p,q) lies on the line (7), then

$$0 = P(k) = Q(k)\cdot 0 + \beta_1 k + \beta_0.$$

This equation may be solved for k to find $k = -\beta_0/\beta_1$. If the same answer occurs for several successive iterations then this number provides a good approximation to be used in the Birge-Vieta method [6, p. 34] to find the real root. The polynomial may then be deflated and the search continued for the other roots.

The above test for a real root will work even when the iterations are successfully converging on a quadratic factor corresponding to two real roots. To prevent this test from stopping convergence, the magnitude of the polynomial is considered too. If it is growing small, then the Bairstow iteration is allowed to continue even though a real root could be isolated.

It is possible that the elimination of a real root is best, even if there are other real roots; this occurs in the following example. It demonstrates how the single real root is detected. Let P(z) be the polynomial

$$P(z) = z^3 - 3z^2 - 60z - 100.$$

The roots of P(z) are -5, -2 and 10. The first acceptable approximate root was $z = -2.1$. P(z) is a deflated polynomial from a polynomial of degree five with a geometric mean of 2.9. Thus, the second approximate root is $|(2.9)^2/(-2.1)| = 4.0$. The sign of both roots is correct and the approximation -2.1 is close to the real root -2.0, but the magnitude of the second approximate root is too small. The sequence of iterations beginning with $p = 1.9$ and $q = 8.4$ is given in table 1. This shows how the Bairstow iterations began to slowly converge, then jumped far away from the solution along the line in the p,q-plane corresponding to the root $z = -2$. This was detected by using the sequence of approximations to the real root which can be found in column 4 of table 1. The real root was determined in five iterations of the Birge-Vieta method and the polynomial was deflated.

iteration number	p	q	approximate real root
0	1.90	8.40	-2.05
1	.88	4.16	-1.89
2	1.32	6.24	-1.97
3	37.66	75.46	-1.90

TABLE 1. The sequence of Bairstow iterations and the approximations to a real root of the polynomial $P(z) = z^3 - 3z^2 - 60z - 100$.

IV. Coping with Multiple Roots

If after a fixed number of steps, the Bairstow iteration fails to converge to a quadratic factor of $P(z)$ and if a single real root has not been detected, then the magnitude of the polynomial is tested at the most recent approximate roots. If it is essentially zero, then one or more of these roots may be a multiple root. The fact that the denominator in the calculation of the correction factors to p and q approaches zero as p and q approach the coefficients of a quadratic factor corresponding to a multiple root introduces roundoff error into the calculation. This prevents convergence to more than a limited number of significant digits. If $P(z)$ has a multiple root m, either real or complex, then m will be a root of the derivative of $P(z)$. If for some integer $j > 0$, $P^{(i)}(m) = 0$, $i = 1,\ldots,j-1$, and $P^{(j)}(m) \neq 0$, then m is a root of multiplicity j and

$P^{(j-1)}(z)$ may be used to find the root without the problems of roundoff error due to a near zero denominator in the calculation of the correction factors to p and q. When the root is found to the number of significant digits desired, we use the root to deflate the polynomial P(z) j times. A check is made at each step of the deflation to guarantee that m is a zero of the remaining polynomial.

To implement the scheme to determine multiplicity, the coefficients of P(z) are transferred to another array and the derivative operations are performed on the copied coefficients. The coefficients of P(z) are changed only if a multiple root of P(z) is successfully determined. This scheme is so straight-forward that one wonders if there is not some unseen problems with it. The alternative is to use repeated application of Horner's method to evaluate the polynomial and as many derivatives as are needed [3]. This would have to be done at each iteration to find the multiple root. We see two objections to this method. First, if the multiplicity of the root is large, this method appears to require many more arithmetic operations. This takes time and possibly introduces more roundoff error. Second, the type of arithmetic operation that occurs with more frequency is addition. This introduces the possibility of a loss of significant digits, due to the addition of nearly equal numbers which are opposite in sign, and the propagation of relative error is often more severe under addition than under multiplication. This can be seen by studying the analysis of the propagation of relative error using process graphs as presented by Dorn and McCracken [6, pp. 80-94].

We conclude this section by giving two examples of how this method worked. First, let P(z) be the polynomial

$$P(z) = z^4 - 5z^3 + 6z^2 + 4z - 8.$$

The initial approximations to roots of P(z) generated the coefficients $p = 3.4$ and $q = -3.0$ of an approximate quadratic factor of P(z). At the end of the maximum 20 iterations, $p = 3.999960$ and $q = -3.999920$. The magnitude of the polynomial at the roots of this quadratic factor is on the order of 10^{-26}, a number which is certainly near zero. It was determined that the root was real and its multiplicity was three. Three iterations with the second derivative of P(z) produced the root $z = 2.0$, correct to 16 decimal digits, and the polynomial was deflated for each root. The roots are 2, 2, 2 and -1.

Second, let P(z) be the polynomial

$$P(z) = z^5 + 3z^4 + 4z^3 - 4z - 4.$$

The initial approximations to the coefficients of a quadratic factor are $p = -2.5$ and $q = -1.9$. After the maximum 20 iterations, $p = -2.0000004$ and $q = -1.9999994$. The magnitude of the polynomial at the roots of this quadratic factor is on the order of 10^{-23}. It was determined that the roots were complex and of multiplicity two. Three iterations of Bairstow's method, using the first derivative of P(z) and the above values of p and q, produced the coefficients $p = -2.0$ and $q = -2.0$, correct to 16 decimal digits. The polynomial was deflated and the roots are $-1 \pm i$, $-1 \pm i$, and 1.

V. Conclusion

We conclude that Bairstow's method can be made to work successfully to find the roots of many polynomials. A relatively simple starting method may be used to make initial approximations to a quadratic factor of a given polynomial and Bairstow's method will often converge successfully. However, we have illustrated examples of and provided methods to correct cases where the Bairstow iteration scheme did not converge on its own. If we accept the assumption that we are to use Bairstow's method to find both real and complex,single and multiple roots, then each of these correction techniques, or similar alternatives, is necessary. Bairstow's method is attractive because of its fast rate of convergence as compared to other methods, but it is less popular because its convergence is unreliable. It is yet to be determined whether the additions described here will make it competitive with other methods. The author hopes that the corrective routines will add minimal execution time to the algorithm and that this method will prove to be faster than and at least as accurate as present methods.

REFERENCES

1. American National Standard Basic FORTRAN (ANSI X3.10-1966), American National Standards Institute, Inc., 1430 Broadway, New York, N.Y. 10018, 1966.

2. Ault, David A., A Total Algorithm for Polynomial Roots Based upon Bairstow's Method, Technical Report No. CS74002-R, Computer Science Department, Virginia Polytechnic Institute and State Univ., Blacksburg, Virginia, 1974.

3. Carrano, F. M., A Modified Bairstow Method for Multiple Zeros of a Polynomial, Mathematics of Computation 27 (1973), 781-792.

4. Conte, S. D., Elementary Numerical Analysis, McGraw-Hill Book Company, New York, 1965.

5. Dijkstra, E. W., Notes on Structured Programming, APIC Studies in Data Processing No. 8, Academic Press, New York, 1972, 1-82.

6. Dorn, William S. and Daniel D. McCracken, Numerical Methods with FORTRAN IV Case Studies, John Wiley and Sons, Inc., New York, 1972.

7. Foulk, C. R., Smooth Programming, First Computer Science Conference (Columbus, Ohio), Feb., 1973.

8. Gabler, W., Entry Areas for Pairs of Roots Using Bairstow's Method, Zeitschrift fur Angewandte Mathematik und Mechanik 49, 4 (1969), 249-250.

9. Gear, C. W., Experience and Problems with the Software for the Automatic Solution of Ordinary Differential Equations, Mathematical Software, J. R. Rice (ed.), Academic Press, New York, 1971, 211-227.

10. Hamming, Richard W., Numerical Methods for Scientists and Engineers, 2nd edition, McGraw-Hill Book Company, New York, 1973.

11. Ralston, Anthony, A First Course in Numerical Analysis, McGraw-Hill Book Company, New York, 1966.

12. USA Standard FORTRAN (USAS X3.9-1966), United States of America Standards Institute, 10 East 40th Street, New York, N. Y. 10016, 1966.