

A Variable Interval Width Quadrature Technique Based on Romberg's Method*

EDMUND K. MILLER[†]

MBAssociates, San Ramon, California 94583

Received May 26, 1969

A numerical integration or quadrature technique (based on Richardson extrapolation to the trapezoidal rule as formulated by Romberg) was developed and applied to evaluate some Fourier integrals which occur in the theory of an infinite cylindrical antenna. The computational method is based on an abscissa spacing which varies as powers of two and which is determined by the relative convergence of two successive approximations to the integral over a given interval. This results in a minimum-to-maximum abscissa spacing ratio as small as 10^{-6} . Numerical examples are given for integration contours following the real axis and deforming into the complex plane.

INTRODUCTION

The solutions to many physical problems may be expressed as integrals resulting from integral transformations, integration over source distributions, etc. It is frequently impossible to evaluate such integrals analytically unless some restrictive approximations can be made so that numerical integration must then be used. The purpose of this paper is to describe a powerful numerical integration (or quadrature) technique which is useful for a wide range of commonly encountered integrals.

The numerical approach was developed specifically in order to find the admittance of an infinite, cylindrical, perfectly-conducting dipole antenna. This problem arises in connection with the use of an antenna in the ionosphere as a diagnostic probe. The antenna is excited at a circumferential gap of nonzero thickness and is immersed in a lossy plasma medium that is, in the general case, both compressible and anisotropic.

The axial surface current on the antenna is readily expressed as a Fourier integral

* The research reported in this paper was supported by the National Aeronautics and Space Administration under Headquarters Contract No. NASr-54(05).

[†] MBAssociates, Post Office Box 196, San Ramon, California 94583. Formerly with High Altitude Engineering Laboratory, University of Michigan, Ann Arbor, Michigan.

over the transform variable β which, when evaluated at the edge of the exciting gap and for unit excitation voltage, yields the infinite antenna admittance. The details of the formulation are given by Miller [4]. Analytic evaluation of this Fourier integral for the antenna admittance is not possible (except on an approximate basis for the free-space case), so a numerical approach must be used. Since the integrand function is time-consuming to calculate and may peak sharply at some value of β , while at the same time the integration range of β goes from zero to infinity, the quadrature technique used must optimize (as nearly as possible) the number of integrand evaluations required to invert the Fourier integral.

It should be mentioned that the Romberg variable interval width quadrature technique, outlined below, is a special case of adaptive Simpson quadrature, which is described in the book by Davis and Rabinowitz [2] and the recent article by Lyness [3]. The technique considered here differs from usual adaptive Simpson quadrature in that it uses interval bisection rather than trisection, and it applies what appears to be a more efficient convergence testing criterion. In addition, the method may further reduce computation time by deforming the integration contour into the complex plane, as illustrated below for the infinite antenna problem. Besides the infinite antenna case, some numerical examples are considered.

NUMERICAL DEVELOPMENT AND EXAMPLES

1. Infinite Antenna Problem

The infinite antenna current $I(z)$ may be written in terms of its Fourier-transformed counterpart as

$$I(z) = \int_0^{\infty} i(\beta) \cos(\beta z) d\beta, \quad (1)$$

where z is the axial distance along the cylinder, measured from the center of the exciting gap of thickness δ .

Consider curve I_1 in Fig. 1. The real part of $i(\beta)$, $i_R(\beta)$, is shown as a function of β_R (contour C_1) for the case of the isotropic incompressible plasma and the following conditions: a frequency f of 2 MHz; an electron plasma frequency f_p of 1.5 MHz; an electron collision frequency ν of 10^4 sec^{-1} ; an electron temperature T of 0°K ; an antenna radius c of 1 cm; and $\delta = 0.1 \text{ cm}$. The sheathless case ($X = 0$) is shown, where X denotes the thickness in electron Debye lengths of a concentric free-space layer, or vacuum sheath, which is used in the analysis to approximate the actual inhomogeneous ion sheath.

The most striking feature exhibited by I_1 is the sharp peak in $i_R(\beta)$ at a β value of approximately 0.03 m^{-1} . This peak corresponds to the wave number of a plane electromagnetic (EM) wave propagating in the plasma medium. The current will

propagate down the infinite antenna with a wave number corresponding to the EM wave, since the peak near $\beta = 0.03 \text{ m}^{-1}$ acts nearly as a delta function in the Fourier equation [Eq. (1)]. The peak in $i_R(\beta)$ is related to the source-free characteristic waves which may propagate on an infinite cylinder.




FIG. 1. The real part of the transformed current $i_R(\beta)$ as a function of β along contour C_1 for the infinite cylindrical antenna in an isotropic plasma medium.

Efficient and accurate numerical evaluation of Eq. (1) requires a sophisticated numerical quadrature technique, in view of the integrand behavior (Fig. 1) and the fact that each $i(\beta)$ value may need about 1 sec or more of computer time (IBM 7090). A variable abscissa spacing (or interval width) must be used, since it is not economical to employ, over the entire β axis, the abscissa point density required to integrate the peak accurately. In addition, it is highly desirable to use a quadrature technique based on equally spaced abscissa points over each integration interval, since any increase of order, iteration, or decrease of interval size, should make use of previously calculated integrand values.

2. ROMBERG QUADRATURE

The Newton-Cotes class of quadrature formulas is based on equally-spaced abscissa points. Special cases of this method are the well-known trapezoidal rule

for the order one, and the parabolic rule for the order two. An extension of the trapezoidal rule by Romberg, involving Richardson extrapolation, leads to a particularly powerful quadrature technique that, with some modification in application, is quite successful in integrating functions of the type shown in Fig. 1.

Richardson extrapolation involves a linear combination of two computed answers, F_1 and F_2 , for the integral of $f(x)$ over the integration range (a, b) (Ralston [7]; Bauer *et al.*, [1]). Let

$$F = F_1 + \sum_{i=K}^{\infty} a_i h_1^i, \quad (2a)$$

$$F = F_2 + \sum_{i=K}^{\infty} a_i h_2^i, \quad (2b)$$

where F denotes the correct answer, h_1 and h_2 are the intervals used, the a_i are constants which may depend on $f(x)$ but not on h , and the sums represent the errors in the calculation (which are taken to be of order K). Multiplying Eq. (2a) by k_2^K , Eq. (2b) by h_1^K , subtracting, and setting $h_1 = 2h_2 = 2h$, we obtain the following

$$F = \frac{1}{2^K - 1} (2^K F_2 - F_1) + \sum_{i=K+1}^{\infty} \left(\frac{2^K - 2^i}{2^K - 1} \right) a_i h^i. \quad (3)$$

It may be seen that the leading error term in Eq. (3) is of order $K + 1$. Thus, the error in the new computed value in Eq. (3) has been reduced without evaluating the integrand at more abscissa points than are required to obtain F_2 . In addition, all the abscissa points required to find F_1 are also used to calculate F_2 . This satisfies our requirement that any change in the quadrature scheme should make use of previously calculated integrand values. Richard extrapolation may be applied to any computation based on a fixed spacing h , but, for reasons to be given below, the trapezoidal rule appears best fitted for that role.

Let the trapezoidal rule answer for 2^k subintervals on the integration range of (a, b) be denoted by $T_{0,k}$, which may be expressed as

$$T_{0,k} = h_k [(f_0/2) + f_1 + f_2 + \cdots + f_{2^k-1} + (f_{2^k}/2)], \quad (4)$$

where

$$h_k = (b - a)/2^k$$

and

$$f_i = f(a + ih_k).$$

them. However, bracketing a sharp peak in $f(x)$ by (a, b) may actually result in oscillating $T_{k,0}$ Romberg answers which converge more slowly to F than do the $T_{0,k}$ trapezoidal rule answers. To make the Romberg technique more efficient and to more nearly optimize the number of integrand evaluations required, the integration range (a, b) should be divided into smaller intervals. This will give a variable abscissa spacing suited to the integrand behavior and will avoid bracketing a sharply varying $f(x)$ by a too-wide integration interval. Assuming that $f(x)$ is continuous in (a, b) , a sequence of intervals (a_i, b_i) which spans the range (a, b) can be chosen which ensures the desired smoothness of $f(x)$ over each interval (a_i, b_i) . The power of the Romberg technique is partly due to the method it provides for choosing the (a_i, b_i) . One way in which this may be done in integrating a curve of the type shown in Fig. 1 is as follows.

3. Variation of the Interval Width

Initially an arbitrary interval size in β is used to begin the calculation at $\beta = 0$, and the quantities $T_{0,0}$, $T_{0,1}$, and $T_{1,0}$ are computed. This operation requires three abscissa points. A comparison of $T_{0,1}$ and $T_{1,0}$ determines whether the desired accuracy (or convergence) has been attained; we require that

$$\left| \frac{(T_{1,0} - T_{0,1})}{T_{1,0}} \right| \leq E_C, \quad (7)$$

where the maximum acceptable convergence error is E_C . If this inequality, or convergence test, is not satisfied, $T_{0,2}$, $T_{1,1}$, and $T_{2,0}$ are then computed. This requires two more abscissa points midway between the points already calculated. The quantities $T_{1,1}$ and $T_{2,0}$ are now compared for convergence, and if it is still not attained, the original interval is divided in two. The new T -matrix elements are computed for the new intervals from the previously evaluated integrand values (which have been stored in the computer). This process of halving the interval is continued until the convergence test is satisfied, using no less than three nor more than five integrand values per interval.

On the other hand, if the convergence test is satisfied on a given interval, the succeeding interval is doubled in size if the resulting endpoints are consistent with the abscissa points of previously calculated integrand values. Suppose, for example, the sequence of calculated abscissa points is as shown in Fig. 2, with abscissa points 1 and 4 the end points of the first interval on which the convergence criterion is satisfied. Then the second interval used will be that defined by abscissa points 4 and 3, and which has the same width on the first interval. If convergence occurs on this interval as well, then the third interval to be used will have end points 3 and 2, and is then twice the width of intervals 1 and 2. Convergence on interval 3 would result in a fourth interval with twice the width of interval 3, etc. It should be

observed that with the bisectioning scheme used to decrease the interval size, the stored integrand values are always encountered in three point intervals which progressively double in size, with the exception of the first interval encountered after the first successful convergence test, which is of the same size as the first interval.




FIG. 2. Typical abscissa point spacing sequence which may result using interval bisection and Romberg variable interval width quadrature.

Note that one of the advantages of the Romberg technique is the convergence test provided by comparison of elements in the T array. This test does not need additional integrand evaluations, as do other rules (trapezoidal, Simpson's, Gaussian, etc.), which require an increase of quadrature order to provide two answers for a convergence test. The two row elements $T_{m-1,1}$ and $T_{m,0}$, rather than the two Romberg answers $T_{m-1,0}$ and $T_{m,0}$, are used for the convergence test, since the row elements have been found by experience to generally provide an earlier and more sensitive indication of convergence. A demonstration of this is provided by the T array given in Ralston ([7], p. 125) for the example

$$\int_1^3 dx/x.$$

At least three points per interval are used ($k = 1$) since this is the minimum number required for Romberg integration. A maximum of five points ($k = 2$) is allowed, since halving the interval would require these abscissa points to check the convergence on each of the two new intervals, while the use of more than five per interval would generally result in a greater point density than necessary over part of the interval. In particular, when a sharp peak in $f(x)$ is bracketed by the interval, $T_{k,0}$ may converge more slowly than $T_{0,k}$ (when $k \gg 2$) and the convergence test may then become unreliable. This restriction to a maximum of five abscissa points per interval is perhaps the most important feature of the technique; its influence on the accuracy and efficiency of the variable interval width Romberg method cannot be over emphasized. Application of this technique to curves of the kind shown in Fig. 1 consistently sets the minimum interval size (and maximum abscissa point density) at the peak of the curve, which results in minimum-to-maximum abscissa spacings of 10^{-8} or smaller.

While Richardson extrapolation could be applied to other quadrature rules to begin with (for example, the parabolic rule), use of the trapezoidal rule allows a convergence test involving $T_{0,1}$ and $T_{1,0}$ to be made with the fewest abscissa points. Richardson extrapolation as applied to the parabolic rule as a special case for example leads to the quantity $T_{2,k}$.

It is worthwhile, before concluding this discussion on the interval width variation, to mention a potential limitation of the convergence criterion established by Eq. (7). A smoothly varying function whose sampled values within the interval under consideration happen to fall near a straight line may satisfy Eq. (7) while producing a $T_{0,1}$ value which is grossly in error.

Increasing the minimum number of abscissa points used per interval could alleviate this possibility to some extent, though not eliminate it since $T_{m-1,1}$ and $T_{m,0}$ will be close in value whenever the $2^m + 1$ abscissa values chosen lie near a polynomial curve of degree $2m - 1$. Application of the technique to arbitrary integrals should, therefore, keep in mind this potential accuracy limitation. The author has, however, not knowingly encountered any difficulties from this false convergence in using the method.

4. The Error Test

A key to the accurate and efficient evaluation of an integral, using the above method, is specification of the acceptable convergence error, E_C , which will give the overall computational accuracy, E , desired. The value used for E_C should be a function both of the integrand variation and the interval width to maintain a nearly uniform interval-to-interval absolute accuracy. Intuitively, it seems inefficient for two adjacent intervals to possess absolute errors which differ by one or two orders of magnitude. This might conceivably happen if the relative accuracies in each interval were the same.

The quadrature summing process, adding together the T values for all the intervals, may be viewed as the addition of a column of numbers. Each number in this column should be sufficiently accurate to achieve the desired accuracy in the sum with a minimum of integrand evaluations. Choices which might be used to determine E_C for this purpose are:

- (1) E_C const,
- (2) E_C scaled to the largest T value in the column, and
- (3) E_C scaled to the sum of T values preceding the current T value.

Of these options, method (2) is least conservative while still achieving an overall error of about E_C .

If the largest T values were known before the calculation, use of (2) would give T values which have their first uncertain digit in the same location relative to the

decimal point. However, the largest T value is not known until the computation is finished. Thus, before applying method (2), E_C must be scaled according to the maximum T value encountered up to that point in the calculation. Consequently, the error test using the E_C determined by (2) is somewhat more conservative in the calculation before the maximum T value has been encountered.

Method (2) determines E_C from

$$\begin{aligned} E_C(n) &= E | T_m / T^{(n)} | \\ &= ET_R. \end{aligned} \quad (8)$$

The superscript n denotes the n -th interval where the Romberg answer is $T^{(n)}$, T_m is the maximum T value encountered to that point, and E is the desired overall relative error. However, the $E_C(n)$ of Eq. (8) can become larger than one. In this case, the error test becomes meaningless. This, and the fact that it may be desirable to keep $E_C(n)$ equal to E until T_m exceeds $T^{(n)}$ by a specified margin [to obtain a somewhat more conservative $E_C(n)$] suggests an alternative to Eq. (8):

$$E_C(n) = E \left[\frac{1 + AT_R}{1 + BT_R} \right], \quad (9)$$

where $B \leq A$ and $A \leq 1$. We see that $E_C(n)$ remains about equal to E until $T_R \sim A^{-1}$, then increases monotonically as T_R increases until $T_R \sim B^{-1}$; beyond this point, $E_C(n)$ is relatively constant with a value of about EA/B .

The use of Eq. (9) to determine E_C has been successful in obtaining a final relative accuracy of E_R close to E for a wide variety of integrals, including the examples given below. Generally, the results are rather insensitive to the value of A , but somewhat more dependent on B . The usual values employed in the calculations presented below have been $A = 10^{-2}$ and $B = 10 \times E$ or A , whichever is smaller. Note that the use of a variable E_C value is mandatory for the evaluation of integrals having sharp cusps, such as example (c) below. A constant E_C value, in such cases, may result in a relative computation error which decreases by several orders of magnitude with a small decrease of E below a critical value. Because of the halving and doubling characteristics of the Romberg variable interval width method, if E_C were not varied there would be a range of E_R which could not be obtained regardless of the specified value of E .

Note that if the integrand under consideration oscillates in sign, it is better to sum the negative and positive parts separately. The final answer is the difference between these two sums. Since the error of each is well-determined, the final error can also be estimated accurately. However, because errors are additive, the overall answer will be less accurate than either sum, thus, requiring a smaller E value than would otherwise be necessary. If the positive and negative sums are of the same magnitude, the final answer may be unreliable even with a smaller value of E .

The quadrature method outlined above is applicable to both finite and infinite (but convergent) range integrals, over any contour in the complex plane. When applied to an infinite range integral, it is necessary to estimate the truncation error which arises from termination of the integration range at a finite value. This can be done by comparing an analytically-integrated, large-argument approximation of the integrand, with the quadrature answer to determine the point at which the truncation error is acceptably small. Contour integration in the complex plane requires treating a complex integrand, where the real and imaginary parts are subjected to the convergence and truncation error tests separately. (Of course this may also happen on the real axis.)

Integrand functions, which have an oscillatory or sharply peaked behavior for which analytically integrable approximations are available, are frequently encountered. In this case, the numerical integration can be carried out more efficiently by integrating the difference between the actual and approximate integrand and adding to this answer the analytic integral of the approximation. This technique is successfully employed in certain electromagnetic problems (Miller and Burke [5]), and is also frequently used to sum infinite series.

Finally, we may mention that the quadrature routine described above was initially programmed to save all calculated integrand values. If, for example, the original or starting integration interval is bisected a number of times, many of the integrand values calculated would not be required to perform the quadrature process over the smaller intervals which result. These integrand values were stored sequentially in the computer, and then recalled as the integration encountered their abscissas. This was an economical procedure since the storage and recall of these values took much less computer time than their calculation. In handling integrals whose integrands are not especially time-consuming to obtain, it may be more efficient to discard the unused values, and recalculate them as required. An option to follow either procedure would be a desirable feature of a general-purpose library routine.

An example of the integrand value sequence and numbering system used was presented in Fig. 2. The stored integrand values are retrieved via an index vector which successively locates them according to their original order of calculation as the abscissa is monotonically increased. The end points and widths of the intervals are also similarly recorded, so that a minimum of reordering is required with each interval bisection. A computer listing for the variable interval width algorithm is available from the author.

5. Numerical Examples

The integral [Eq. (1)] of the curve I_1 (Fig. 1) was evaluated by the above scheme. It required a total of 105 integral evaluations to obtain an infinite antenna conductance value of 7.073×10^{-4} mho with a specified error E of 5×10^{-3} and

an estimated truncation error (the error which results from truncating the infinite integration interval at a finite value of the upper limit) of less than one unit in the fourth significant figure. To check this result (as well as to evaluate the feasibility of avoiding the near singularity exhibited by curve I_1), a second value for the conductance was calculated using the deformed contours C_2 and C_3 shown in Fig. 3. This resulted in curves I_2 and I_3 , respectively. A conductance value in this case was obtained from integrating I_2 and I_3 as 7.065×10^{-4} mho, with the same truncation and convergence errors as before, using a total of 45 abscissa points. The difference between these two answers is well within the specified error, and a considerable decrease in the number of points was required to perform the integration, using the deformed contour.




FIG. 3. The real part of the transformed current $i_R(\beta)$ as a function of β along contours C_2 and C_3 for the infinite cylindrical antenna in an isotropic plasma medium.

The spacing between successive integrand evaluations and the corresponding $T_{K,0}$ values as a function of the abscissa value β for the integrand I_1 is shown in Fig. 4. As a result of the sharp decrease in abscissa spacing in the vicinity of the peak, the integrand evaluations cluster about the peak in I_1 . This demonstrates that the technique concentrates the integrand evaluations in regions where the integrand is changing most rapidly. In addition, the maximum-to-minimum abscissa spacing ratio varies in the vicinity of the integrand peak by more than three orders of magnitude. This ratio increases to about 10^8 when the larger abscissa values

(near the truncation point at approximately 1000 m^{-1}) are considered. Obviously, the abscissa spacing required to integrate the peak of I_1 accurately could not be used to cover the entire range of integration, so a variable abscissa spacing is necessary. The contribution to the integral from each interval is relatively constant over the peak in I_1 , since the interval width decreases in proportion to the increase in the integrand. This is advantageous from a convergence error viewpoint because it results in the addition of a sequence of nearly equal numbers of comparable accuracy.




FIG. 4. The variation of the abscissa spacing and $T_{k,0}$ Value for integration of curve I_1 of Fig. 1 by the Romberg variable interval-width technique.

The advantage of the Romberg variable interval width technique is not restricted in application to sharply peaked functions of the type shown by curve I_1 . A comparison of relative accuracies obtained from Simpson's rule, Romberg quadrature using a fixed abscissa spacing over the integration interval, and the present technique, are shown as a function of the number of integrand evaluations in Fig. 5 for the three integrals below:

$$\int_0^1 (1 - 0.5x^4)^{-1} dx, \quad (\text{a})$$

$$\int_0^1 (1 + 100x^2)^{-1} dx, \quad (\text{b})$$

$$\int_0^1 |0.5 + x|^{1/2} dx. \quad (\text{c})$$

(The relative accuracy E_R is defined as the absolute value of the difference between the exact answer and the quadrature answer divided by the exact answer.) The superiority of the variable interval width method over the straightforward Romberg technique and Simpson's rule is clearly shown.

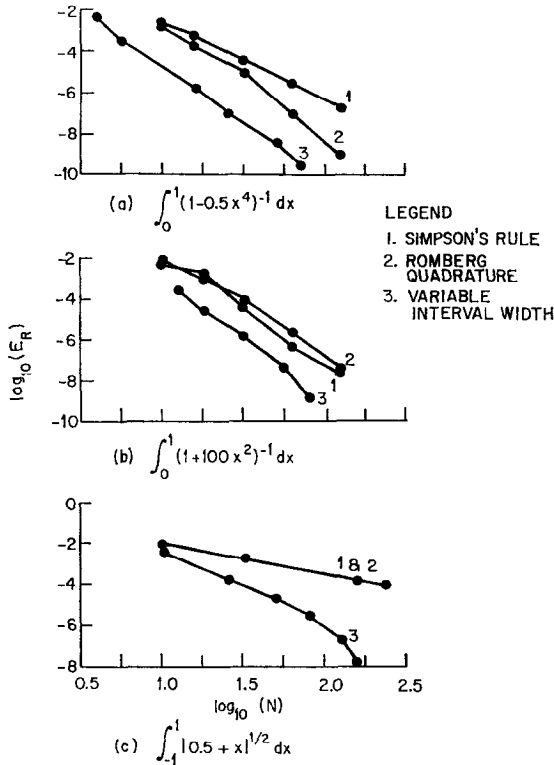


FIG. 5. Comparison of the relative accuracies E_R obtained for three integrals. The number of integrand evaluations is denoted by N .

These results were compared with the recent study presented by O'Hara and Smith [6], who considered various quadrature techniques, including Gaussian and Curtis-Clenshaw quadrature for three integrals similar to the examples given here. This comparison shows that the variable interval width Romberg technique is generally competitive with these more accurate methods for examples (a) and (b) and exceeds the most accurate method of those considered for example (c).

The importance of varying the interval width for a function which is not sharply peaked in the integration range is demonstrated by these examples. It would

seem to indicate that a comparison of numerical quadrature technique which does not allow for variable interval widths may be somewhat misleading as to the relative power of the various methods. The conclusion reached by O'Hara and Smith [6] that the Romberg method is inferior to Clenshaw-Curtis and Gaussian quadrature would seem to need reexamination. A particular feature of the variable interval width technique, it should be noted again, is that the abscissa points are equally spaced in the interval, and no integrand values need be discarded when the abscissa spacing is decreased. This is contrary to the case for Clenshaw-Curtis quadrature when the interval size is decreased, and the Gaussian quadrature, when the order or interval size is changed. The variable interval width technique may generally be expected to require fewer total integrand evaluations to achieve a desired accuracy. The accuracy obtained from Gaussian and Clenshaw-Curtis quadrature, however, may be greater than that provided by the Romberg variable interval width method when only those integrand evaluations required to calculate the final answer are counted.

An additional advantage of the Romberg variable interval width technique is the conformity obtained between the acceptable convergence error and the actual integration error. The results of this technique for examples (a), (b), and (c) are presented in Table I. The convergence test used provides a reliable, generally con-

TABLE I
Actual Integration Accuracy Achieved for a Specified Convergence Error^a

Specified Error E	Relative Accuracy E_R Integrand		
	(a)	(b)	(c)
$A = 10^{-2}$ $B = \min[A, 10E]$	$\int_0^1 (-0.5 \times x^4)^{-1} dx$	$\int_0^1 (1 + 100x^2)^{-1} dx$	$\int_{-1}^1 0.5 + x ^{1/2} dx$
10^{-1}	7×10^{-3}	4×10^{-2}	4×10^{-3}
10^{-2}	6×10^{-4}	4×10^{-3}	4×10^{-3}
10^{-3}	5×10^{-4}	4×10^{-2}	4×10^{-3}
10^{-4}	3×10^{-5}	3×10^{-5}	2×10^{-4}
10^{-5}	1×10^{-6}	2×10^{-6}	2×10^{-5}
10^{-6}	1×10^{-7}	7×10^{-8}	3×10^{-6}
10^{-7}	3×10^{-9}	3×10^{-9}	4×10^{-7}
10^{-8}	9×10^{-11}	1×10^{-9}	2×10^{-8}

^a Using the Romberg variable interval width method for the integrals shown.

servative estimate of the actual error. This is a particularly attractive feature of the Romberg technique, and one which provides considerable confidence in the results obtained from it.

CONCLUSION

The Romberg integration technique with variable interval width provides an efficient and powerful quadrature method. The technique is a general-purpose routine which is especially useful for integrating functions which are time-consuming to calculate. It provides an accurate and sensitive test on the convergence accuracy at each stage in the calculation, without requiring redundant computations of the integrand function. Since it employs abscissa spacings based on powers of two, the technique may make use, in the final answer, of all integrand values calculated. The resultant variable interval width gives an abscissa point density which is tailored to and determined by the function being integrated. Because the method provides a reliable convergence test, it is well suited to be a library routine where the user need only specify, in addition to the function to be integrated, the integration endpoints and accuracy desired. These features, together with a contour integration capability and an integrand-value storage (if desired), provide a versatile, reliable, and powerful quadrature method for general-purpose computer usage.

ACKNOWLEDGMENT

The author wishes to acknowledge the contribution of Mr. Gerald J. Burke of MBAssociates in carrying out some of the numerical computations.

REFERENCES

1. F. L. BAUER, H. RUTISHAUSER, AND E. STIEFEL, "New Aspects in Numerical Quadrature," *Proc. Symp. Appl. Math.* 15-th *Amer. Math. Soc.*, Providence, R. I., pp. 199-218, 1963.
2. P. DAVIS AND P. RABINOWITZ, "Numerical Integration," Blaisdell Publishing Co., London, 1967.
3. J. N. LYNESS, Notes on the Adaptive Simpson Quadrature Routine, *J. Assoc. Comput. Mach.*, July, 1969.
4. E. K. MILLER, The admittance of an infinite cylindrical antenna in a Lossy, compressible, anisotropic plasma, *Canad. J. Phys.* **46** (1968), 1109-1118.
5. E. K. MILLER AND G. J. BURKE, Numerical Integration Methods *IEEE Trans. Antennas Propagat.* **AP-17**, No. 5, (1969), 669-672.
6. H. O'HARA AND F. J. SMITH, Error Estimation in the Clenshaw-Curtis Quadrature Formula, *Comput. J.*, **II**, No. 2, (1968), 213-219.
7. A. RALSTON, *A First Course in Numerical Analysis*, p. 21, McGraw-Hill, 1965.