JOURNAL OF COMPUTATIONAL PHYSICS 4, 19–29 (1969)

Numerical Integration of Analytic Functions[1]

CHARLES SCHWARTZ

Department of Physics, University of California, Berkeley, California 94720

Received August 23, 1968

ABSTRACT

A method of numerical integration is presented which is simple, versatile and accurate. It is particularly valuable for integrals with complicated end point singularities.

The method we shall discuss starts with a consideration of the Euler–Maclaurin sum formula:

$$\sum_{n=0}^{N-1} \delta F(n\delta) = \int_0^{\Delta} F(x)\,dx - \frac{\delta}{2}[F(\Delta) - F(0)] + \sum_{l=1}^{\infty} \frac{(-1)^{l+1} B_l}{(2l)!} \delta^{2l} \left(\frac{d}{dx}\right)^{2l-1} F \Big|_{x=0}^{x=\Delta} \tag{1}$$

where $\Delta = N\delta$ and B_l are the Bernoulli numbers. If the term with 1/2 the end points is included with the sum we get the trapezoidal rule, which leaves an error of order δ^2; and if we add two such series at alternate points so as to cancel this δ^2 error we get Simpson's rule, with error of order δ^4, and so forth. We start with the observation that if the function F and its first p derivatives vanish at the endpoints of the integration, then the simple sum (1) differs from the integral in order δ^{p+2}. The general theme of the method we propose is to change variables under the integration sign to obtain rapid decrease of the error with δ while using this simplest summation formula. We are here reminded that the Euler–Maclaurin formula is probably an asymptotic formula and its predictions may not always be reliable (imagine a function with essential singularities at both end points, such as $e^{-1/x(1-x)}$, or a function periodic over the interval such that all the correction terms in (1) cancel out); but one may nevertheless use it as a qualitative indication of the sort of convergence, as δ goes to zero, to be expected.

[1] Research supported in part by the Air Force Office of Scientific Research, Office of Aerospace Research, United States Air Force, under Grant AF-AFOSR-130-66.

To make our analysis specific we shall now consider the integral

$$I = \int_{-\infty}^{\infty} dx\, F(x) \tag{2}$$

and its approximation by the sum

$$\sum = \sum_{n=-\infty}^{+\infty} \delta F(n\delta). \tag{3}$$

We assume that the function F goes to zero at the end points fast enough that both integral and sum converge, and we assume that F is an analytic function on the real line where we are integrating and also within the strip $|\operatorname{Im}(x)| \leqslant w$ in the complex plane. We now construct the integral representation

$$\sum = \oint \frac{dx\, F(x)}{e^{2\pi i x/\delta} - 1} \tag{4}$$

where the closed contour is the following.

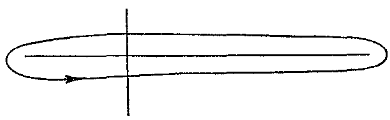

Now we separate this into two integrals

$$\sum = \sum_+ + \sum_-$$

$$\sum_- = \int_{-\infty - i\epsilon}^{+\infty - i\epsilon} \frac{dx\, F(x)}{e^{2\pi i x/\delta} - 1} \tag{5}$$

$$\sum_+ = -\int_{-\infty + i\epsilon}^{+\infty + i\epsilon} \frac{dx\, F(x)}{e^{2\pi i x/\delta} - 1} = \int_{-\infty + i\epsilon}^{+\infty + i\epsilon} dx\, F(x) \left[1 + \frac{1}{e^{-2\pi i x/\delta} - 1}\right]$$

and we can relate the sum to the integral

$$\sum = I + \mathscr{E} \tag{6}$$

with the error given by the integrals

$$\mathscr{E} = \int_{-\infty - i\epsilon}^{+\infty - i\epsilon} \frac{dx\, F(x)}{e^{2\pi i x/\delta} - 1} + \int_{-\infty + i\epsilon}^{+\infty + i\epsilon} \frac{dx\, F(x)}{e^{-2\pi i x/\delta} - 1}. \tag{7}$$

Now using the assumed properties of $F(x)$ we move the first integral down a distance w in the complex plane, and move the second one up by w (The number w has a rather evanescent quality[2] since one can easily find examples where it might appear the contour could be moved infinitely far away. We are simply taking as an example the case where the integrand has some singularity at a distance w from the real axis. For a different example see the latter part of Appendix A.)

$$\mathscr{E} = \int_{-\infty}^{\infty} \frac{dx\, F(x - iw)}{e^{A} e^{2\pi i x/\delta} - 1} + \int_{-\infty}^{\infty} \frac{dx\, F(x + iw)}{e^{A} e^{-2\pi i x/\delta} - 1} \tag{8}$$

where

$$A = \frac{2\pi w}{\delta} \tag{9}$$

Since we are interested in the error as δ goes to zero, we have e^{A} as a very large number and so we write

$$\mathscr{E} \sim e^{-A} \left[\int_{-\infty}^{\infty} dx\, F(x - iw)\, e^{-2\pi i x/\delta} + \int_{-\infty}^{\infty} dx\, F(x + iw)\, e^{2\pi i x/\delta} \right] \tag{10}$$

$$= o(e^{-2\pi w/\delta}) \tag{11}$$

This result does indeed show that the error decreases faster than any power of δ, and suggests that we may have here a very powerful method for numerical integration. (Halving the interval *doubles* the number of correct figures in the answer.) The practical goal is to get small error with a modest number N of points contributing to the sum $\sum$. We would terminate the sum (3) when the neglected contributions were smaller than the desired accuracy, but whether this means a small or large number of terms depends on the function F and how rapidly it dies as its argument goes to infinity.

Let us take an example: $F(x) = (e^{x} + e^{-x})^{-1}$. We have a closest singularity at $x = \pm i\pi/2$, so $w = \pi/2$. To get an error of only $\mathscr{E}$ we would terminate the sum at $x = \pm \ln \mathscr{E}^{-1}$. Thus we have the total number of points used is

$$N = \frac{2}{\delta} \ln \mathscr{E}^{-1}$$

and the error is given by

$$\mathscr{E} \approx e^{-\pi^2/\delta} = e^{-\pi^2 N/2 \ln \mathscr{E}^{-1}}$$

Solving we get

$$[\ln \mathscr{E}^{-1}]^2 = \frac{\pi^2 N}{2}$$

[2] The author is grateful to a reviewer for pointing this out.

and finally

$$\mathscr{E} \approx e^{-\pi\sqrt{N/2}} \tag{12}$$

This is a rapid convergence rate: for $N = 10$ the error is about 10^{-3}, and for $N = 100$, about 10^{-10}. (We shall later see how to make this even better.)

By contrast consider the example $F(x) = (1 + x^2)^{-1}$. We have $w = 1$, but we must terminate at $x = \pm\mathscr{E}^{-1}$ to leave errors of order $\mathscr{E}$ in the summation. Thus

$$N = \frac{2}{\delta}\,\mathscr{E}^{-1}$$

$$\mathscr{E} \approx e^{-2\pi/\delta} = e^{-\pi N\mathscr{E}}$$

or

$$\mathscr{E} \approx \frac{\ln \mathscr{E}^{-1}}{\pi N}. \tag{13}$$

Here we see that the convergence is worse than $1/N$, and this is extremely bad. The technique we shall use to improve this is a change of variable in the integral which will make the resulting integrand decay more rapidly at $\pm\infty$. We shall use the exponential function for these changes, not because it is unique as a function, but because it is readily evaluated by computing machines. Thus with

$$x = e^y - e^{-y} \tag{14}$$

we get

$$I = \int_{-\infty}^{\infty} dx\, F(x) = \int_{-\infty}^{\infty} dy\, G(y)$$

where

$$G(y) = (e^y + e^{-y})\, F\,(e^y - e^{-y}). \tag{15}$$

Then if F decays as a power

$$F(x) \xrightarrow[x\to\infty]{} x^{-p}, \tag{16}$$

G decays exponentially

$$G(y) \xrightarrow[y\to\infty]{} e^{-(p-1)y}$$

and we would expect a faster convergence of the error as a function of the number of points N. Of course we must also ask how this change of variables has moved the singularities around before we can be sure that we have really gained something. For this last example where the poles were at $x = \pm i$, these are transformed to a string of poles on the imaginary axis, the closest at a distance $w = \pi/6$; therefore we expect a good formula, something like (12). If we use this transformation on the first problem we reach the conclusion that the rate of error decrease is improved from something like $\exp(-N^{1/2})$ to something like $\exp(-N/\ln N)$. As one repeats this game there are the two competing effects of the faster decrease of the function, which means we need fewer points, and the moving of the singularities closer to the real axis, which degrades the error formula. Just where the optimum lies is too complicated to say in general, but it is easy enough to find out by experimenting on a computer with any given problem.

We have up to now assumed the range $(-\infty, +\infty)$ for the given integral. If we are given something else, we will start off with an exponential transformation.
If

$$0 \leqslant x < \infty; \qquad x = e^y \qquad -\infty < y < \infty \tag{17}$$

If

$$0 \leqslant x \leqslant 1; \qquad x = \frac{1}{1 + e^{-y}} \qquad -\infty < y < \infty \tag{18}$$

Furthermore suppose that the function F has singularities at the end points of the integral (see also Appendix B); this will usually not cause any difficulty since after our variable changes we put an essential singularity at each end point, and the original singularity is usually smothered. If the function F has singularities within the original range of integration, we would just break it up into several separate integrals with singularities only at the end points and proceed as above.

We will now show several examples. The manner of studying error vs. N is slightly different from that described above. The desired accuracy in the sum $\sum$ is specified first, and then the summation is repeated several times with diminishing spacing δ.

#1. $\int_0^1 dx$ We first made the transformation (18) and then repeated the transformation (14) M times. The sum was stopped when contributions were less than 10^{-20}. Figure 1 shows a logarithmic plot of the absolute magnitude of the error as a function of N, the number of points used, for $M = 0, 1, 2, 3$. $M = 1$ appears to be best. Also shown on the same figure is the integral $\int_0^1 dx \ln 1/x$ treated the same way, and the results are also very good. The calculation was repeated using a scaling factor in the transformation (14), i.e.

$$x = C(e^y - e^{-y}), \tag{19}$$

and $C = 1$ appeared to be about optimum. (We have not tried to analyse this.)

#2. $\int_0^\infty dx\, e^{-x}$ Here we start with (17) and then proceed as with #1. The results in Figure 2 show again excellent results, with larger M best, but $M = 1$ still very good.

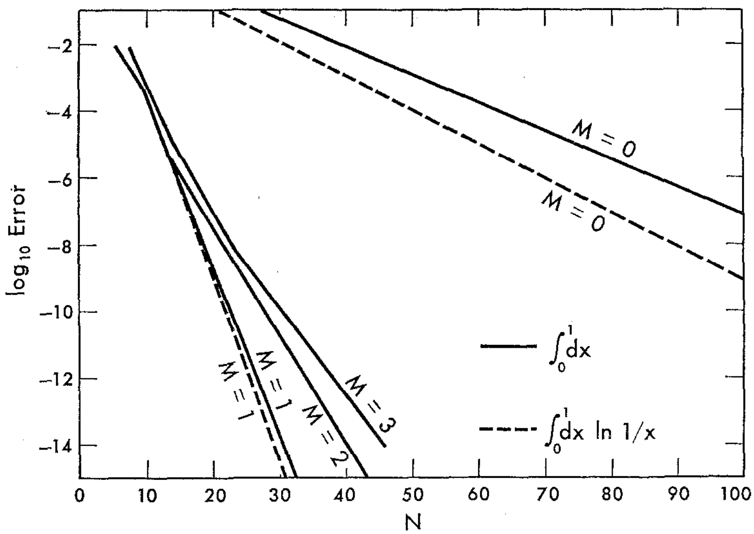

FIG. 1. Error vs. N by the present method for the examples #1.

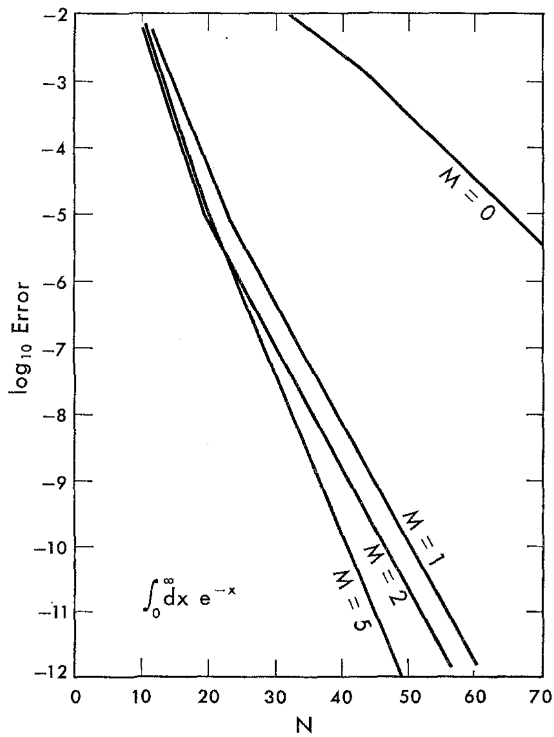

FIG. 2. Error vs. N by the present method for example #2.

#3. $\int_0^1 dx/((x - .5)^2 + a^2)$ Treated as #1 for two values of a, shown in Figure 3. As a gets smaller the results get poorer.

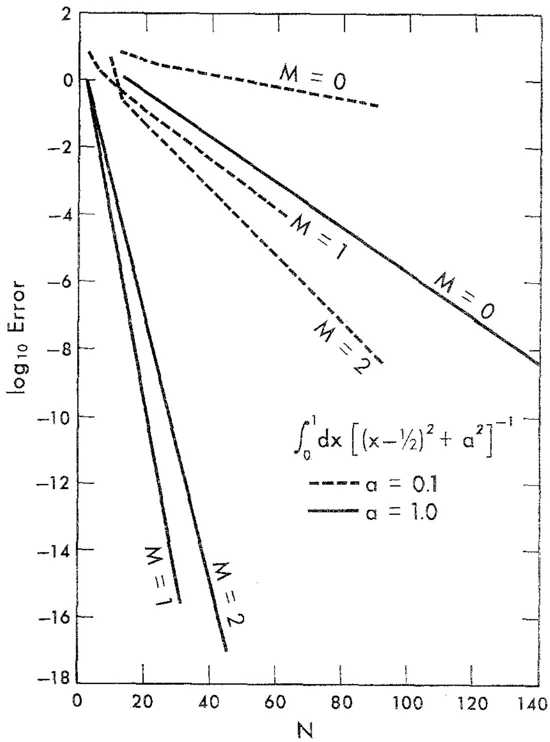

FIG. 3. Error vs. N by the present method for example #3 at $a = 0.1$ and $a = 1.0$.

Discussion

It is commonly said that the most efficient method for numerical integration is some form of Gaussian quadrature. We wish to make the general claim that the method described in this paper is competitive, and perhaps even superior because in addition to yielding rapid convergence it is simple to use and is very versatile. For any given integral one can construct a sequence of Gaussian points and weights which will be the best, but it may be a nuisance to get these from tables; and if the integrand does not have the analytic properties assumed for the Gaussian formula then much of the power of that method is lost. The present method on the other hand is quite insensitive to end point behavior and is very simple to program. (This method was found while trying to do a set of integrals involving products of several Bessel functions such that the power series expansion about one end point had various powers of logs and thus no Gaussian formulas were available.)

In all these methods what really controls the rate of convergence is the nearness of singularities in the complex plane. The real artistry of numerical integration lies in learning to make changes of the variable such that these singularities are mapped farther away, and this must be studied separately for every problem. For example suppose we had the integral $\int_{-\infty}^{\infty} dx f(x)$ where the function f was singular at $x = \pm ia$, then for small a we would get slow convergence of our method as well as of the other classical methods. We can map this singularity off to infinity; for example by

$$z = \frac{x}{\sqrt{x^2 + a^2}}; \qquad x = \frac{az}{\sqrt{1 - z^2}} \tag{20}$$

which gives us the new representation for the integral

$$\int_{-1}^{1} \frac{1}{a} \frac{dz}{\sqrt{1 - z^2}} f(x). \tag{21}$$

We have introduced new singularities at the end points, but these can be well handled either by our method or by the appropriate Gaussian method. Of course if the function f has other singularities we may bring them in closer by this transformation and possibly lose more than we gain. But then we must study the function closer and be more clever with the transformations we need to apply.

One outstanding problem is the following. Consider the integral equation

$$f(x) = \int dy\, K(x, y) f(y) + g(x). \tag{22}$$

We wish to replace the integral by some discrete sum and then solve a finite algebraic (matrix) problem for the function f at a discrete set of mesh points. Now often the kernel $K(x, y)$ will have singularities at points y which depend on the value of x. Thus when we do our best to get an accurate integration formula —by making clever mappings in the variable y to move singularities far away— these will be different mappings for different x; and so we will not get a closed system of algebraic equations for the function f at fixed mesh points. A test case is the Schrödinger equation with a Yukawa potential, in momentum space:

$$\varphi(p) = \lambda \int_0^{\infty} q^2\, dq \frac{1}{p^2 + \kappa^2} \frac{1}{2pq} \ln \left[\frac{(p + q)^2 + 1}{(p - q)^2 + 1}\right] \varphi(q) \tag{23}$$

The author has not been able to do better than about one percent accuracy in the eigenvalue λ with a ten point mesh for the integral by any method. In contrast a variational calculation of this problem with ten trial functions can give λ accurate to about ten decimal places.

Summary

We recommend the following method for numerical integration of a given analytic function.

1) Use the transformation (17) or (18) to make the range of integration $(-\infty, \infty)$ if it is not so to start with.

2) Then make the further transformation (14) once.

3) Now evaluate the sum (3), truncating when contributions fall below the desired accuracy, and reduce the value of the spacing δ until the answer has the desired accuracy.

Appendix A

Another realization of the error formula (7) is obtained by expanding the denominators in geometric series,

$$\mathscr{E} = \sum_{l=1}^{\infty} \left[\tilde{F}\left(\frac{2\pi}{\delta} l\right) + \tilde{F}\left(-\frac{2\pi}{\delta} l\right)\right] \tag{A1}$$

involving the Fourier transforms

$$\tilde{F}(p) = \int_{-\infty}^{\infty} dx\, F(x)\, e^{ipx}. \tag{A2}$$

So the rapid decrease of the error as $\delta \to 0$ is given in terms of the rapid decrease of $\tilde{F}(p)$ as $p \to \infty$. The dominant term would usually be $l = 1$:

$$\mathscr{E} \approx \tilde{F}\left(\pm \frac{2\pi}{\delta}\right). \tag{A3}$$

Thus if $F(x)$ is analytic in the strip $|\operatorname{Im} x| < w$, we get as before

$$\mathscr{E} = o(e^{-(2\pi/\delta)w}). \tag{A4}$$

For another example:

$$F(x) = e^{-bx^2} \tag{A5}$$

we find

$$\mathscr{E} \approx e^{-\pi^2/b\delta^2}. \tag{A6}$$

Appendix B

The contour integral representation can be used to analyse the error in the finite sum (1)

$$\sum_{n=0}^{N-1} \delta F(n\delta) = \oint \frac{dx\, F(x)}{e^{2\pi i x/\delta} - 1} \tag{B1}$$

We start with the integral over the narrow contour:

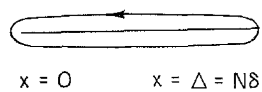

(B2)

and then open it up to:

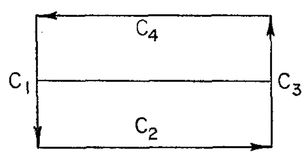

(B3)

after writing

$$\frac{1}{e^z - 1} = -1 - \frac{1}{e^{-z} - 1} \tag{B4}$$

for the parts in the upper half plane. Now we have

$$\sum_{n=0}^{N-1} \delta F(n\delta) = \int_0^{\Delta} F(x)\, dx + R_1 + R_2 + R_3 + R_4 . \tag{B5}$$

If the contour C_2 is a distance w from the real axis, we have

$$R_2 \sim o(e^{-2\pi w/\delta}) \tag{B6}$$

and similarly for R_4. The interesting terms are

$$R_1 = i \int_0^w \frac{d\xi [F(i\xi) - F(-i\xi)]}{e^{2\pi\xi/\delta} - 1} \tag{B7}$$

$$R_3 = -i \int_0^w \frac{d\xi [F(\Delta + i\xi) - F(\Delta - i\xi)]}{e^{2\pi\xi/\delta} - 1} . \tag{B8}$$

As $\delta \to 0$ the region of importance in these integrals is $\xi \gtrsim \delta$. If we can expand $F(x)$ in a Taylor series about $x = 0$ and about $x = \Delta$, then we recover the Euler Maclaurin series (1). However if F is not analytic at either endpoint, these formulas still let us estimate the error. For example, suppose that as $x \to 0$ F behaves as

$$x^p(\ln x)^q \tag{B9}$$

even for non-integral powers p and q. By scaling the integral (B7) we see that the error behaves as

$$\mathscr{E} \sim \delta^{p+1}(\ln \delta)^q. \tag{B10}$$