

PERSONAL GLIMPSE

How I Was Led to the Frequency Approach

R. W. Hamming

Department of Computer Science, Naval Postgraduate School, Monterey, California 93943

Probably the first necessary step in adopting the frequency approach to numerical analysis is to realize that there are serious defects in the polynomial approach. I will give a few examples of these, more or less in the order that they came to my attention.

In 1945 while in the computer group at Los Alamos I discovered that the classical Simpson's iterated formula for numerical integration makes no statistical sense. If you think of the integration of a function $y(x)$ from 0 to 1 as computing the average value of the function in that interval, then the iterated formula is

$$[y(0) + 4y(h) + 2y(2h) + 4y(3h) + 2y(4h) + \cdots + y(2nh)]/3n$$

(where n is the number of double intervals used). But to believe that a sample of the integrand at one point around the middle of the range is twice as important as the adjacent values is nonsense! Thus, although the formula may be impeccably right in a mathematical sense, it fails to pass the reasonable statistical test of common sense.

In Scarborough's classic book (of many editions) he derived the error term for each polynomial interpolation formula as if it were different in each case, yet it was obvious to me that from the same sample points you get a unique interpolation formula (though the form of writing may involve different function values and assorted differences), and hence the error term depends only on the sample points used and not on the form in which it is written (of course within round-off errors). Therefore I was led to think in terms of classes of formulas rather than in terms of the isolated ones which filled the texts at that time. Furthermore it was clear to me that what we were computing at Los Alamos had very little to do with interpolation.

The next event that I recall occurred when I was teaching a night class in numerical methods. (How better to learn the topic than to teach it?) I was ending the lecture with the problem of finding the derivative to a set of data and I said, following the text book, "You find the interpolating polynomial, differentiate it, and evaluate the derivative at a sample point." Even as I spoke I realized the utter stupidity of what I was saying—at a sample point the interpolating polynomial is almost certainly crossing the function, and a worse point to evaluate a derivative could hardly be found! The bell saved me, but during the hour's drive home I had time to meditate on the mathematically correct but, nevertheless, idiotic result. It also became clear to me in the same course that the Weierstrass theorem on the approximation of functions by polynomials both was irrelevant and produced unbelievably high degree polynomials. Thus I finally realized that many of the classical numerical methods are plain silly.

It was somewhere around this time that I found that the (ideal) mathematical result that the powers of x are linearly independent over any interval (and over $n + 1$ points up to the n th power of x), while true, is false in the presence of "noise" (round-off errors), and it fails for even moderate degree polynomials. You can find this in Lanczos' *Applied Analysis*, if not for yourself (as I did). In a sense Hilbert's matrix underlines this fact clearly. This is the reason, I believe, that practical numerical methods for integrating differential equations seldom go above fifth order.

My next adventure came when I was using Milne's method for integrating ordinary differential equations (which I inherited and took over without benefit of careful thought) and found that it was unstable in some situations! A very careful reading of Milne's two books showed that he never even hinted that his

method was unstable, although he observed that a certain method was unstable. I was in a desperate situation from which I could not withdraw without a great loss of face. How about using the Adams–Bashforth method? Well, I was not going to do that until I understood things a lot better than I did, and that meant, at least to me, that I had to create a general theory in which both methods occurred as special cases.

It was at about this time that I developed what I call *the direct method*. You write down the formula (a linear form) with m unknown coefficients using the samples of the function you are prepared to use. Then you make the formula exact for $1, x, \dots, x^{m-1}$, and since it is linear it will be exactly true for any polynomial of degree $(m - 1)$. If there is an interpolating polynomial, then the resulting formula will be the same as if you found the interpolating polynomial and did the corresponding operation on it. For example, to derive Simpson’s formula we are given the three function values $y(-1), y(0)$, and $y(1)$. Thus we write

$$\int_{-1}^1 f(x) dx = ay(-1) + by(0) + cy(1)$$

and require that this be exactly true for the functions $1, x$, and x^2 . We get:

Function	LHS	RHS
if $y = 1$:	2	$= a + b + c$
if $y = x$:	0	$= -a + c$
if $y = x^2$:	$2/3$	$= a + c$

Solving these equations we get Simpson’s coefficients directly and with a lot less effort than with the classical method of first finding the interpolating polynomial. To repeat the argument, if the formula is exact for any polynomial of a given degree, then it must be true for the particular polynomials which are the powers of x . And if it is true for the individual powers, it must be true for any linear combination of them. It was only much later that I realized the lack of perfect equivalence of the two methods; there might be no interpolating polynomial but there might still be a corresponding formula!

I used this method to obtain the family of integration formulas for differential equations. I wrote down the form, using the data I planned to use, then made the formula exact for the corresponding powers of x , but saved a couple of coefficients for stability and round-off properties that I deemed to be important. Thus I was led to two parameter families of predictor and corrector formulas for integrating ordinary differential equations, and could study the matter of sta-

bility and round-off of the whole family. I could then better understand what I was doing when I picked a particular method for solving ordinary differential equations.

Now it should be clear, and I must emphasize this point, that in use one formula differs from another only by the numerical values of the coefficients used and does not affect the amount of arithmetic the machine must do, save when you can get 0 coefficients and therefore save one or more multiplications and additions. Thus I found suitable integration formulas and emerged with my personal pride and a published paper.

This is perhaps the point at which it is necessary to explain equally spaced sampling and the technical words “band-limited functions.” We almost always use equally spaced samples, and it is convenient to think of the sampling rate as the unit of time—thus problems sampling in microseconds and problems sampling in years are brought to a common basis. Now the sampling theorem states that you need two samples for the highest frequency present, and this is often known as the *Nyquist rate*. Thus the highest frequency is one-half a cycle per unit of time. The simple reason behind this is the pair of trigonometric identities

$$\sin \pi(1 + a)t - \sin \pi(1 - a)t = 2 \sin \pi t \cos \pi a t$$

$$\cos \pi(1 + a)t - \cos \pi(1 - a)t = -2 \sin \pi t \sin \pi a t.$$

At the unit sampling rate, $\sin \pi n = 0$ and the higher frequency $(1 + a)$ has the same sample values as the lower frequency $(1 - a)$ —the two frequencies cannot be distinguished from each other! Each frequency above the rate of $\frac{1}{2}$ per cycle is *aliased* into a lower frequency! *This aliasing is fundamental*; it is due to the sampling and cannot be ignored. As a result we have to deal only with frequencies f in the band $-\frac{1}{2}$ to $\frac{1}{2}$ —there are no other frequencies present after the sampling is done. The polynomial approach has no such clarity in explaining the effects of the sampling rate. Even unsampled original data are generally approximately band limited in the frequencies that have large amplitudes; for example, hi-fi stereo sound systems have a limited band in which they reproduce the sound reliably.

My real introduction to the frequency approach began when I was asked to solve on a computer what amounted to a 28th-order system of ordinary differential equations. I could estimate the cost of the programming (in those days in absolute binary!) and knew how many trajectories they wanted and the corresponding real times of them. Now the error terms in the polynomial approach depended on the

size of the 5th derivatives, and who could possibly guess at them! But a mistake in the estimate of a factor of 32 could double the machine costs (which were in the thousands of dollars and hence large for those days). Fortunately for me, I knew of Shannon's sampling theorem, but this theorem involves samples from minus to plus infinity! Having at best data on only one side, it seemed reasonable to expect that maybe 4 samples would be about right, and also having only a short run of data maybe another factor of 2 would be needed—hence perhaps a sampling rate of 8 per the highest frequency present in the function might be reasonable. I pointed this out to the men who wanted the solution and the need for them to commit themselves to the highest frequencies present in the solution I was to find. Meanwhile I turned to the relay computer we had at that time to test this conjecture. I found experimentally, and then theoretically, that around 7 samples was the minimum, and by 10 samples I was more than safe for the accuracy with which they knew any of their constants. They came back with the remark, "We will worry about anything above 10 cycles per second, and you worry about those below." So I went ahead, and as soon as possible tested the program at 1/60 of a second, getting hash more or less, and at 1/120 and getting perfect results. So we next ran, and continued to run, at 1/70 of a second. Had I not been able to get the solutions, I at least could have said to them, "You have higher frequencies than you thought and have just spent \$8000 in programming costs to find this out."

It was not long after this that I heard that a West Coast subcontractor was having trouble simulating the launch of a guided missile and was using a step size of 1/1000 to 1/10,000 of a second. I immediately asserted that they had a programming error and that the error was of the order of 300 or more. Well, this debugging of a program across the country was correct! It converted any lingering doubts I had that for problems in which the frequency was of physical significance the frequency approach was greatly superior to the classical polynomial approach. One simply asks what the error in the formula is for each of the frequencies in the band of interest and then adjusts the unknown coefficients so that the error term is, generally, Chebyshev in form (in the band of course).

Being at Bell Telephone Laboratories helped, as the frequency approach arises naturally both in communication problems and in feedback control problems using the Nyquist criterion of stability. But being ignorant of the current development of recursive digital filters, I went ahead on my own—and it is probably lucky that I did. The classical filter theory insists on "bounded input, bounded output," but it is immediately evident that if you are to integrate a constant

(hence bounded) function, then the solution must rise linearly. Stability meant to me no "exponential growth." Later, when integrating trajectories into the moon, I tried to get the corresponding second-order differential equation methods and I found immediately that I would have a double zero on the boundary between stability and instability. Reflection showed that this had to be true when there is no drag to produce first derivatives in the formula. Now, of course, the error could rise quadratically! Hardly stable by their condition. But thinking things through for myself, I have come to the conclusion that the classical definitions from the analog past are simply inappropriate for the digital world. I have been unable to convince the classicists that their criterion of stability prevents them from accurately simulating even simple integration, let alone more interesting and important problems.

The objection so often raised that the corresponding computing will be much more than the computing for the polynomial approach has already been forestalled above—no matter which functions you use to make the formula exact, or at least close to exact in a Chebyshev sense, the data assumed will determine the amount of arithmetic needed. Of course the derivation of the formula (which must be done only once in your entire career) may be harder in the frequency approach than in the polynomial approach, but is that a matter of any importance? Hardly!

The main gains in the frequency approach are (when it is appropriate): (1) from the physical understanding of the problem you can accurately estimate the spacing to use, and hence the amount of machine time needed, and (2) from the results of the computations you can often give physical insight. It is perhaps the latter that is the most valuable asset of the frequency approach in the long run. The troubles that can arise in the polynomial case often depend on the position of singularities in the corresponding complex plane (which are seldom known in practice), while for the frequency approach the troubles are visible on the real line (where the data occur).

I must now digress and discuss "band-limited functions" a bit more. The basic mathematical tool is the Fourier transform of the time function $g(t)$,

$$g(t) = \int_{-\infty}^{\infty} G(f)e^{2\pi ift} df,$$

where clearly $G(f)$ is the amount of the corresponding frequency f in the signal $f(t)$. It is a nice fact that the inversion formula is almost perfectly symmetric,

$$G(f) = \int_{-\infty}^{\infty} g(t)e^{-2\pi ift} dt.$$

Note that the formulas we are using are in *rotations* f , and not in radians.

The effect of sampling at unit spacing is to alias all higher frequencies into the range (band) $-\frac{1}{2} < f < \frac{1}{2}$. Thus the effective range of integration is similarly reduced in the transform in f .

Let us illustrate the frequency approach method with the same Simpson's integration formula, written in the form

$$y_{n+1} = y_{n+1} + ay'_{n+1} + by'_n + ay'_{n-1},$$

where $y(x)$ is the integral and $y'(x)$ the integrand. It is natural to want a constant input ($y' = c$) to give exactly the correct answer. Hence putting $y(x) = x$ we get the equation

$$1 = -1 + a + b + a$$

or

$$2 = 2a + b$$

as one condition on the coefficients. To get the other condition we try computing the error $R(f)$

$$R(f) = \text{LHS} - \text{RHS}$$

for the frequency function $e^{2\pi if t}$. Thus we will have the error curve as a function of f .

We want this error curve to have the shape of a Chebyshev polynomial so that the maximum error for any frequency in the interval will be a minimum. It is well known that expansions in Chebyshev polynomials converge much more rapidly than the corresponding truncated Taylor series—indeed the last coefficient a_n of the Taylor series is divided by 2^{n-1} to get the corresponding coefficient of the Chebyshev expansion. We therefore need to write the exponential as a series in Chebyshev polynomials using the following easily derived identity

$$e^{izx} = J_0(z) + 2 \sum_{m=1}^{\infty} i^m J_m(z) T_m(x),$$

where $J_m(z)$ are the Bessel functions of order m and $T_m(x)$ are the Chebyshev polynomials of order m .

We need to identify the variables in the identity with those of our problem. Let us suppose, as in the case of Tick's formula, that we want the Chebyshev criterion to hold only for some fraction $k < 1$ of the Nyquist interval. The reason for this is that for several reasons we oversample and take more than the minimum number of samples for the highest frequencies we think are there. First, the actual function may

not have a sharp cutoff at the highest frequencies, but rather trails off toward zero. If we sampled at the lowest rate possible, then the sampling would alias the weak higher frequencies on top of those we were concerned with. Second, by taking many more samples than we need, we can get some statistical averaging to lower the random effects of the measurements. Thus we are often interested mainly in frequencies up to some fraction k of the Nyquist frequency $f = \frac{1}{2}$.

Thus we want the variable x in the Chebyshev polynomials to reach the value 1 at $k/2$, that is, $-1 < x < 1$ goes into $-k/2 < f < k/2$; hence

$$x = 2f/k.$$

But clearly

$$xz = 2\pi ft$$

so

$$z = \pi kt$$

and the fundamental identity becomes

$$e^{2\pi if t} = J_0(\pi kt) + 2i^m J_m(\pi kt) T_m(2f/k).$$

Putting this into our Simpson's formula we get the messy expression

$$\begin{aligned} J_0(\pi k) + 2 \sum_{m=1}^{\infty} i^m J_m(\pi k) T_m(x) \\ = J_0(-\pi k) + 2 \sum_{m=1}^{\infty} i^m J_m(-\pi k) T_m(x) \\ + a\pi k [J'_0(\pi k) + 2 \sum_{m=1}^{\infty} i^m J'_m(\pi k) T_m(x)] \\ + b\pi k [J'_0(0) + 2 \sum_{m=1}^{\infty} i^m J'_m(0) T_m(x)] \\ + a\pi k [J'_0(-\pi k) + 2i^m J'_0(-\pi k) T_m(x)]. \end{aligned}$$

Because of the oddness and evenness of the Bessel functions, the coefficients of $T_{2m}(x)$ all vanish, leaving only the odd indexed terms.

Setting the coefficient of $T_1(x)$ equal to zero (and dropping a $2i$ factor) will make the error term resemble the next Chebyshev polynomial $T_3(x)$. We get

$$\begin{aligned} J_1(\pi k) = J_1(-\pi k) \\ + \pi k [aJ'_1(\pi k) + bJ'_1(0) + aJ'_1(-\pi k)] \end{aligned}$$

or

$$2\pi kaJ'_1(\pi k) + \pi kbJ'_1(0) = 2J_1(\pi k).$$

Using $J_1'(0) = \frac{1}{2}$ and the earlier condition on the coefficients a and b , namely

$$2a + b = 2,$$

we get on eliminating b

$$a = \{2J_1(\pi k)/(\pi k) - 1\} / \{2J_1'(\pi k) - 1\},$$

from which the following table is easily computed.

k	a	b
0.0	0.33333	1.33333
0.1	0.33425	1.33150
0.2	0.33703	1.32594
0.3	0.34177	1.31647
0.4	0.34862	1.30276
0.5	0.35785	1.28431
0.6	0.36979	1.26043
0.7	0.38493	1.23014
0.8	0.40394	1.19213
0.9	0.42771	1.14457
1.0	0.45742	1.08496

These values do not give the exactly equal ripple error curves because we have taken only the next term in the series and neglected the still higher order terms.

We see that the case $k = 0$ is exactly the classical polynomial case of Simpson's formula. The particular case Tick worked out, using a computer to compute the error curves and leveling them by repeated trials,

was for $k = \frac{1}{2}$, and he got $a = 0.3584$, which differs slightly (0.00055) from ours as he exactly leveled the error curve and we neglected the small effects of higher degree terms. The case $k = 1$ gives just about twice the hamming window coefficients, 0.23, 0.54, 0.23, as another check.

Thus we see that the main difference in the formulas derived by the polynomial approach and by the frequency approach is that the coefficients differ somewhat, but the physical interpretations differ quite a bit. Both can be derived by the method of undetermined coefficients using the same data, but the frequency formulas are closely connected with the known physics when the method is appropriate.

I trust that this essay clearly shows the path by which I was fairly directly led, step by step, to the frequency approach (though of course I have omitted many details). I was simply following my motto:

The purpose of computing is insight, not numbers.

I have given only the highlights of the early development. For further details see:

1. Lanczos, C. *Applied Analysis*. Prentice-Hall, Englewood Cliffs, NJ, 1956;
2. Hamming, R. W. *Numerical Methods for Scientists and Engineers*. Dover, New York, 1972, 2nd ed.; 1986, reprint;
3. Hamming, R. W. *Digital Filters*. Prentice-Hall, Englewood Cliffs, NJ, 1989, 3rd ed.;

and of course the current literature which has greatly extended the development without altering the fundamentals.