Stephen A. Dyer and Justin S. Dyer

Numerical Integration

Introduction

If we want to find the integral $y(nT)$ of a function (e.g., a signal) $x(t)$ over $[0,\, nT]$, we would solve

$$y(nT) = \int_0^{nT} x(t)\, dt \,. \tag{1}$$

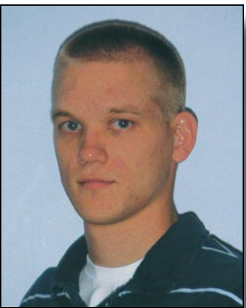

But what if the integration cannot be carried out exactly? Or what if $x(t)$ is known at only a finite number of points $x(0), x(T), x(2T), \ldots, x(nT)$ (i.e., $x(t)$ has been sampled and we have knowledge only of the values of the resulting sequence $x[k] = x(kT)$, $k = 0, \ldots, n$, where T is the sampling interval)? In situations such as these, we would use numerical integration (or, numerical quadrature) to estimate the number obtained if we could carry out (1).

You may recall from elementary calculus the notion of a *Riemann sum* to approximate the integral. One form of this approximation is

$$y(nT) \approx \sum_{k=1}^{n} x(kT)\Delta t = \sum_{k=1}^{n} x[k]T \,.$$

Indeed, if $x(t)$ is piecewise continuous and we fix the interval of integration and then let $n \to \infty$ (making T successively smaller in the process), we arrive at the Riemann integral. In effect, this is, in fact, how the Riemann integral is *defined*.

Integration is a smoothing operation, and numerical integration is, in its essence, a stable operation. Many formulas (or rules) exist. The most basic of these is the so-called *rectangular rule*, which effectively amounts to a Riemann sum with finite n (or, equivalently, fixed T). Perusal of any elementary text on numerical analysis would uncover several other classical formulas, such as the *trapezoidal rule* and *Simpson's rule*. Some other rules include the *midpoint rule*, the *corrected trapezoidal rule*, *Tick's rule*, *Simpson's three-eighths rule*, and *Bode's rule*. And the list goes on.

While some would argue that basic rules such as these are generally of only historical interest, with the possible exception of extended (or composite) versions of the simpler formulas, comparing various of the rules can help develop intuition. And, besides, there indeed remain situations, especially in various instrument designs, that beg for the implementation of some of these basic integration rules.

Numerical Integration as Recursive Digital Filtering

Although only a rare text on numerical analysis presents it as such, numerical integration can be modeled as a digital filtering operation. (Hamming's [1] is such a text, and we use his presentation and results as inspiration for what follows.)

For our purposes here, we will consider, along with true, or ideal, integration, the following rules for numerical integration: The rectangular rule, the trapezoidal rule, Simpson's rule, and a rule similar in form to Simpson's, called Tick's rule [1]. We will treat these numerical-integration rules as the recursive—or infinite-impulse-response (IIR)—digital filters they are, defining them by the respective difference equations that describe them.

Since they are linear, time-invariant digital filters, we can further describe each in terms of its transfer function, its impulse response, and its frequency response. (We will find each of these for the trapezoidal rule, and leave it to the reader to do similarly for the remaining numerical-integration rules.)

Once the frequency responses of the various integrators have been found, we can easily make comparisons, in the frequency domain, of their relative performance.

In all that follows, we will, with no loss of generality, normalize the sampling interval T to unity (i.e., 1 second, if the independent variable t represents *time*). Thus, the sampling frequency $f_s = 1/T$ is likewise unity (1 Hz), and the Nyquist frequency f_N is then 0.5 Hz.

Difference Equations for the Various Rules

The difference equations describing the four numerical-integration rules considered are as follows:

Rectangular Rule. With the rectangular rule, the estimate of the integral $y[n]$ at sample point n of the sampled function $x[n]$ depends only on the estimate of the integral at the previous sample and the value of the function at the current sample. Specifically, the difference equation (again, with the sampling interval taken to be $T = 1$) is

$$y[n] = y[n-1] + x[n] \,. \tag{2}$$

Trapezoidal Rule. The trapezoidal rule adds the area under the trapezoidal element established by $x[n-1]$ and $x[n]$

1094-6969/08/$25.00©2008IEEE

on the nth interval to the estimate $y[n-1]$ of the integral obtained up through the $(n-1)$th interval. The difference equation is

$$y[n] = y[n-1] + \frac{1}{2}x[n] + \frac{1}{2}x[n-1]\,. \tag{3}$$

Simpson's Rule. The trapezoidal rule is a two-point formula giving exact results when the function being integrated is a polynomial of degree one. Simpson's rule is a three-point formula giving exact results, it turns out, when integrating not only polynomials of degree 2 but also polynomials of degree 3. The difference equation is

$$y[n] = y[n-2] + \frac{1}{3}x[n] + \frac{4}{3}x[n-1] + \frac{1}{3}x[n-2]\,. \tag{4}$$

Tick's Rule. Another three-point formula, described in [1] and attributed to Leo Tick, is described by the difference equation

$$y[n] = y[n-2] + 0.3584x[n] + 1.2832x[n-1] + 0.3584x[n-2]\,. \tag{5}$$

This rule has a nice spectral property that we will note shortly.

Example: The Trapezoidal Rule

We now find the transfer function, the impulse response, and the frequency response of the trapezoidal rule.

Transfer Function

The Z-transform (ZT) of (3) yields

$$Y(z) = z^{-1}Y(z) + \frac{1}{2}\left[X(z) + z^{-1}X(z)\right],$$

from which we can write the transfer function $H(z)$ as

$$H(z) = \frac{Y(z)}{X(z)} = 0.5\frac{z+1}{z-1}\,. \tag{6}$$

(For those of you who remember your linear-systems theory: Is this filter stable, marginally stable, etc.?)

Impulse Response

The impulse response $h[n]$ can be found as the inverse ZT (IZT) of the transfer function $H(z)$. Expanding $H(z)/z$ into partial fractions yields

$$H(z) = -0.5 + \frac{z}{z-1}\,, \tag{7}$$

from which we can find the IZT as

$$h[n] = \begin{cases} 0.5, & n = 0 \\ 1.0, & n > 0\,, \end{cases} \tag{8}$$

which is the impulse response associated with the trapezoidal rule.

Frequency Response

In obtaining the frequency response, we first recall the Laplace-transform variable (i.e., the complex-frequency variable) $s = \sigma + j\omega$ and the mapping to the z-plane, $z = e^{sT}$, where, again, T is the sampling interval.

To obtain the frequency response of a continuous-time (CT) system, we evaluate the transfer function H(s) at $s = j\omega$. For a discrete-time (DT) system, we then have

$$z = e^{sT}\Big|_{s=j\omega} = e^{j\omega T}\,.$$

So the frequency response of the DT system is

$$\begin{aligned} H(z)\Big|_{z=e^{j\omega T}} &= H(e^{j\omega T}) = \overline{H}(\omega) \\ &= H(e^{j2\pi fT}) = \overline{H}(f) \end{aligned} \tag{9}$$

Applying (9) to (6), with $T = 1$, we find the frequency response for the trapezoidal rule to be

$$\begin{aligned} \overline{H}(f) &= 0.5\frac{e^{j2\pi f}+1}{e^{j2\pi f}-1} \\ &= \frac{1}{j2}\frac{\cos\pi f}{\sin\pi f} \\ &= \frac{1}{j2}\cot\pi f\,. \end{aligned} \tag{10}$$

Relative Performance

How does a given numerical-integration rule act, at a particular frequency, in comparison with true integration? To answer that question, we simply need to take the ratio $\overline{H}(f)/\overline{H}_{\text{true}}(f)$ of the frequency response of that numerical integrator to that of a true integrator and plot the ratio over the Nyquist interval.

Recall that a true integrator has as its transfer function $H_{\text{true}}(s) = 1/s$. Its frequency response is then

$$\overline{H}_{\text{true}}(f) = \frac{1}{j2\pi f}\,. \tag{11}$$

For our example of the trapezoidal rule, the ratio is thus

$$\frac{\overline{H}(f)}{\overline{H}_{\text{true}}(f)} = \pi f\ \cot\pi f\,. \qquad \text{(trapezoidal rule)} \tag{12}$$

Equation (12) tells what the trapezoidal rule does at frequency f with respect to a true integrator—whether it accentuates or plays down that spectral component in the function whose integral is being estimated.

Another way of looking at the ratio $\overline{H}(f)/\overline{H}_{\text{true}}(f)$ is as the frequency response of a linear system that takes the ideally integrated version of the function $x(t)$ being considered and converts it to the corresponding numerically integrated version.

We note that if the desire is simply to obtain the ratio $\overline{H}(f)/\overline{H}_{\text{true}}(f)$ vs. f for a given numerical-integration rule, there is no particular need to obtain its frequency response $\overline{H}(f)$ in closed form if you have written—or otherwise have available to you—functions such as `ct_response()` and `dt_response()` [2] for evaluating frequency responses of CT and DT systems, given their transfer functions.

Another thing to point out with regard to investigating relative performance is that some of the numerical-integration rules have phase error with respect to true integration. (The rectangular rule is one such example.) Thus, when plotting the performance of various rules, it would be better to graph the *magnitude*—and, if necessary, the phase—of $\overline{H}(f)/\overline{H}_{\text{true}}(f)$ vs. f.

Figure 1 shows the magnitude of the ratio $\overline{H}(f)/\overline{H}_{\text{true}}(f)$ vs. f for the four numerical integrators considered, with the curve corresponding to ideal integration being included for reference.

Which Integrator to Choose?

Given the results presented, which numerical-integration rule would you choose? Clearly, from a spectral point of view, all of them should perform well if the energy in the function to be integrated were confined to frequencies well below the Nyquist (say, $0.05f_s$). Simpson's rule more closely matches true integration for a larger range of frequencies, and Tick's rule does well for an even broader range—from 0 to about $0.25f_s$. But, look what happens at higher frequencies. Both of these rules provide excessive amplification at frequencies approaching the Nyquist. If the signal being integrated contains significant high-frequency energy (or has high-frequency noise embedded in it), then the results of applying Simpson's rule or Tick's rule could be disappointing at the very least, and disastrous in many cases. (Apply Simpson's or Tick's rule to a signal consisting of a sinusoid at f_N. The samples in the sequence are proportional to [1, –1, 1, –1, 1, –1, …]. What happens? This is easy enough to do by hand.)

The rectangular rule similarly provides amplification over true integration for high frequencies, although it is relatively small—about 4 dB—at the Nyquist, whereas both Simpson's rule and Tick's rule have unbounded amplification at the Nyquist.

The trapezoidal rule, in contrast, provides increasing *attenuation* with increasing frequency, with the gain reaching zero at the Nyquist frequency. The effects of roundoff, high-frequency noise, etc., tend to be squashed by the trapezoidal rule, in contrast with the other rules considered herein.

The answer to which integration rule to use eventually gets back to how much error per step is tolerable. Clearly, the signal to be integrated requires considerable oversampling to achieve a reasonable small per-step error. Either expanding Figure 1 to show more detail over, say, $0 \leq f \leq 0.1f_s$, or tabulating the values of the various curves over that range will aid in selecting an appropriate sampling interval. For example, sampling at a rate of 10 times the highest frequency in the function to be integrated results in a per-step underestimate of less than 3% when the trapezoidal rule is used, an overestimation by less than 0.1% with Simpson's rule, and an underestimation by less than 0.4% with Tick's rule.

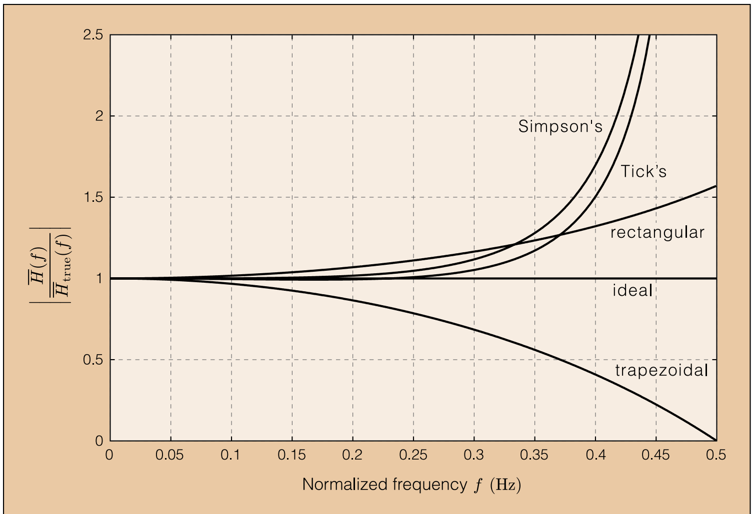

Fig. 1. Ratio $|\overline{H}(f)/\overline{H}_{\text{true}}(f)|$ of the magnitude response of the numerical integrator to that of an ideal (i.e., true) integrator for several numerical integrators.

References

[1] R. W. Hamming, *Numerical Methods for Scientists and Engineers*, 2nd ed. New York: McGraw-Hill, 1973.

[2] S. A. Dyer and J. S. Dyer, "System response and passband transformations," *IEEE Instrum. & Meas. Mag.*, vol. 2, no.3, pp. 43–46, Sep. 1999.

Stephen A. Dyer (s.dyer@ieee.org) has just completed his third term as President of the IEEE Instrumentation and Measurement Society. He is a Professor of Electrical and Computer Engineering at Kansas State University and is President of Silver Wolf Technologies, Inc., where he designs instrumentation. His edited book, *Wiley Survey of Instrumentation and Measurement*, published by Wiley–IEEE Press, is available as a 1,112-page hardbound volume and as an electronic book.

Justin S. Dyer (j.dyer@ieee.org) holds the B.S.E.E. and the M.S. in electrical engineering from Kansas State University. He recently served for two years as the Graduate Student Representative to the I&M Society's Administrative Committee and is currently a Ph.D. student in statistics at Stanford University.