

Accurate Fourier Analysis for Circuit Simulators

Kenneth S. Kundert — Cadence Design Systems

Abstract

A new approach to Fourier analysis within the context of circuit simulation is presented that is considerably more accurate and flexible than the traditional SPICE approach. It is based on the direct computation of the Fourier integral rather than on the discrete Fourier transform and so it is not subject to aliasing. It can be used to accurately compute a small number of Fourier coefficients for broad-spectrum signals such as those generated by mixers, $\Sigma\Delta$ and pulse-width modulators, DACs, and SC-filters. In addition, techniques that reduce errors and provide the ability to resolve harmonics 120 dB – 140 dB below the carrier are presented.

1 Introduction

It is widely accepted that with care the Fourier analysis available in SPICE can resolve harmonics that are 60 dB – 80 dB below the carrier. A well known approach to providing a more accurate Fourier analysis is to uniformly sample the waveforms while forcing SPICE to compute the solution at each of the sample points and using the Discrete Fourier Transform (DFT) to compute the Fourier coefficients. This differs from the approach used in SPICE2 where the uniform sampling of the waveform is performed by interpolating.

A new approach to Fourier analysis is presented that computes the Fourier coefficients using the Fourier integral rather than the DFT. This approach has two important benefits. First, it uses the unequally spaced time-points naturally generated by the simulator. As a result, fewer time-points are used, making the simulation less expensive (assuming time required for computing the circuit response dominates over the time required for Fourier analysis). Second, this approach does not suffer from aliasing. It accurately computes the Fourier coefficients for a few harmonics, even though there is significant energy in the harmonics that are not being computed. This second advantage is an important one, because it allows the Fourier analysis to be performed on signals it would not normally be suited for, such as the waveforms generated by $\Sigma\Delta$ and pulse-width modulators, mixers, DACs and SC-filters.

The second advantage also results in the new method

being less sensitive to the error that results from the signal not being precisely periodic. This common situation results when analyzing signals generated by a circuit that has not completely settled. In this case, the signal at the beginning and the end of the period do not match. Fourier analysis assumes periodic signals, and treats a nonperiodic signal as a periodic signal with a discontinuous jump at the end of the analysis interval. Because it is discontinuous, the jump is an error with a very wide bandwidth. Both the Fourier integral and the DFT are sensitive to the portion of the error from the jump at the frequencies of the coefficients being computed. However, the DFT is also sensitive to the error at frequencies beyond the Nyquist limit (half the sampling frequency) because of aliasing. The wide bandwidth of the discontinuous jump results in a large amount of additional error in the coefficients computed by the DFT.

The computational complexity of the FFT is in the order of $N \log_2 N$ whereas the complexity of the Fourier integral is in the order of KN where N is the number samples used and K is the desired number of Fourier coefficients. The Fourier integral will be computationally less expensive if $K < \log_2 N$. The Fourier integral only requires storage for the K coefficients being computed, not for all of the N sample points. Thus, it requires less memory than the FFT if $K < N$.

2 Previous Approaches

SPICE2g6 computes the Fourier coefficients of a waveform by linearly interpolating to a uniformly-spaced grid and computing the DFT. The errors in this method spring from the interpolation and aliasing in the DFT. The `.FOURIER` command is intended to be used to compute the distortion of low-distortion sinusoids. Thus, aliasing is not typically an issue (though nonperiodic waveforms due to incomplete settling are). The DFT is applied to the interpolated signal, which differs from the true signal by an amount that is related to the curvature of the signal and the time between the interpolation point and the closest time-points computed by the simulator. The greater the curvature and the farther the interpolation point is from the time-points, the greater the error. In SPICE there are several reasons why interpolation can lead to unacceptable error:

3.1.1

25

1. In Fourier analysis, users are interested in resolving very small signals. Thus, even small interpolation errors are a problem.
2. SPICE's Fourier analyzer does not bound the time-step during the analysis period in order to assure reasonable accuracy.
3. SPICE controls the time-step to assure that second-order interpolation is accurate on capacitor charge waveforms and inductor flux waveforms, not to assure accurate Fourier analysis.
 - (a) The Fourier analysis is using first-order interpolation, which requires a smaller time-step to achieve the same accuracy as the second-order interpolation used by the simulator.
 - (b) Fourier analysis is performed on node voltages and voltage source currents rather than on capacitor charges and inductor fluxes. Thus, the simulator is not choosing the time-steps to control error in the signals being Fourier analyzed.
 - (c) If the circuit only contains very small capacitors, `CHGTOL` effectively loosens the local-truncation error criterion, allowing the simulator to take large time-steps even though the Fourier analyzer requires smaller steps.

Other sources of error can also limit the resolution of the Fourier analysis. For example, discontinuities in the model equations due to incomplete models, errors, or bypass algorithms generate small jumps in the waveforms being analyzed. In addition, the small errors that result from incompletely converging Newton's method also contributes to reducing the resolution of the Fourier analysis.

The interpolation error is eliminated by forcing SPICE to place a time-point everywhere a sample is needed for the Fourier analysis. This dramatically improves the resolution of the Fourier analysis at the expense of requiring the simulator to evaluate the circuit at additional time-points. If the waveform being analyzed is a low-distortion sinusoid, the additional computational burden is negligible because the DFT will not require many samples to accurately compute the Fourier coefficients. However, if the signal has a broad spectrum because it has very fast transitions between two or more equilibrium points, then the DFT's need for small time-steps during the transitions, and its requirement that the samples be equally spaced, will require the simulator to take unnecessarily small time-steps during the flat portions of the waveforms between the transitions.

3 The Fourier Integral

For a periodic waveform x of period T , the Fourier coefficients are computed with the following formulas,

$$a_0 = \frac{1}{T} \int_t^{t+T} x(\tau) d\tau \quad (1)$$

$$a_k = \frac{2}{T} \int_t^{t+T} x(\tau) \cos\left(\frac{2\pi k\tau}{T}\right) d\tau \quad (2)$$

$$b_k = \frac{2}{T} \int_t^{t+T} x(\tau) \sin\left(\frac{2\pi k\tau}{T}\right) d\tau \quad (3)$$

where $k = 1, 2, \dots$. The derivation that follows uses (2) to demonstrate the calculations. Similar steps should be taken with (1) and (3) to compute a_0 and b_k .

The circuit simulator discretizes time and solves the system of equations that describe the circuit at $N + 1$ time-point, t_0, t_1, \dots, t_N . Equation (2) is rewritten as a sum of integrals over each time step,

$$a_k = \frac{2}{T} \sum_{n=1}^N \int_{t_{n-1}}^{t_n} x(\tau) \cos(k\omega\tau) d\tau \quad (4)$$

where $\omega = \frac{2\pi}{T}$ and $t_N - t_0 = T$.

Between time-points, x can be approximated with a low-order polynomial. This is very natural because the circuit simulator makes a very similar approximation in order to perform numerical integration. During time-step n ,

$$x(\tau) \approx \sum_{m=0}^M c_{mn} \tau^m \quad (5)$$

for $t_{n-1} \leq \tau \leq t_n$, where M is the order of the approximating polynomial. Substituting into (4) gives

$$a_k \approx \frac{2}{T} \sum_{n=1}^N \int_{t_{n-1}}^{t_n} \sum_{m=0}^M c_{mn} \tau^m \cos(k\omega\tau) d\tau. \quad (6)$$

Exchange the order of integration and summation,

$$a_k \approx \frac{2}{T} \sum_{n=1}^N \sum_{m=0}^M c_{mn} \int_{t_{n-1}}^{t_n} \tau^m \cos(k\omega\tau) d\tau. \quad (7)$$

Finally,

$$a_k \approx \frac{2}{T} \sum_{n=1}^N \sum_{m=0}^M c_{mn} \alpha_{mn} \quad (8)$$

where

$$\alpha_{mn} = \int_{t_{n-1}}^{t_n} \tau^m \cos(k\omega\tau) d\tau. \quad (9)$$

This integral is in a simple form that can be evaluated analytically using integration by parts and induction.

$$\alpha_{mn} = \sum_{i=0}^m \frac{m! \tau^i}{i! (k\omega)^{m-i+1}} \sin\left(k\omega\tau + \frac{\pi}{2}(m-i)\right) \Big|_{\tau=t_{n-1}}^{t_n} \quad (10)$$

where $0! \equiv 1$.

3.1.2

4 Interpolation Error

Computing the Fourier series using the Fourier integral not only avoids error due to aliasing, it is also less sensitive to the error that results from interpolation. The DFT is equivalent to applying the trapezoidal rule with a uniformly spaced grid to the Fourier integral. Thus, to compute a particular coefficient, the input signal is first multiplied by sinusoid, the product is interpolated to the uniform grid and then integrated. The new approach first interpolates the signal alone to the grid. Once interpolated, the remaining two steps are performed analytically. Interpolating the signal alone introduces considerably less error than does interpolating the product of the signal and the sinusoid, particularly for the high frequency sinusoids used to compute the upper harmonics. In addition, the simulator is controlling the time-step so that interpolating the signal is accurate, not the product of the signal and some sinusoid.

While use of the Fourier integral reduces the effect of interpolation error, it does not completely eliminate it. It remains very important to reduce the error due to interpolation in order to provide sufficient resolution.

5 Implementation in Spectre

Computation of Fourier coefficients using the Fourier integral was implemented in the Spectre circuit simulator. Spectre is different from SPICE in a key way that helps to improve the resolution of the Fourier analysis. Rather than choosing the time-step to control the error in polynomial approximation of the *charge* waveforms on capacitors only (neglecting inductors), Spectre chooses the time-step to control error in the polynomial approximations of the *voltage* waveforms on capacitors. With the addition of the Fourier analyzer, Spectre also chooses the time-step to control error in the polynomial approximations of the voltage waveforms at the input of the Fourier analyzer. As a result, simply tightening RELTOL directly acts to improve the accuracy of the Fourier analysis.

Several additional methods are applied simultaneously in Spectre to reduce error from interpolation.

1. Higher order interpolation methods are used. Typically, second-order interpolation is used, though in special circumstances first order interpolation is used on at most a few time-points.
2. The Fourier analyzer forces Spectre to place a time-point at both the beginning and the end of the Fourier analysis period (using Spectre’s break-points).
3. The Fourier analyzer limits the maximum size of the time-step during the Fourier analysis (only) to be no

harmonic	SPICE2g6	Spectre	Harm. Bal.
fund	875 mV	1 V	1 V
2 nd	-40.5 dB	-98.6 dB	-98.5 dB
3 rd	-43.2 dB	-116 dB	-120 dB
4 th	-24.3 dB	-125 dB	-169 dB
5 th	-33.3 dB	-124 dB	-143 dB
6 th	-33.3 dB	-125 dB	-196 dB
7 th	-24.3 dB	-125 dB	-165 dB
THD	9.55 %	0.0012 %	0.0012 %

Table 1: Results from $\mu a741$ using default tolerances and time-steps.

greater than one-tenth of the period of the largest harmonic being computed.

Finally, Spectre does not use bypass algorithms.

6 Examples

6.1 Distortion of a Feedback Amplifier

Fourier analysis was applied to a low distortion feedback amplifier. The amplifier is a $\mu a741$ op-amp in unity-gain configuration. The input frequency is low at 1 kHz and the input amplitude is small at 1 Vp, so the distortion is very low. The distortion was computed from 9 harmonics using the Fourier analysis in SPICE2g6 and in Spectre on a 10ms transient analysis. Nothing was done in either simulation to improve the results (default values were used for both RELTOL and the maximum time-step). Harmonic balance was also run on the circuit as a control (periodic harmonic balance has the property of becoming extremely accurate as signals approach being pure sinusoids). The results, given in Table 1, show that because SPICE does not control the time-step to accurately compute Fourier coefficients, its results are often inaccurate. Spectre with RELTOL = 10^{-3} resolves signals down to about -120 dB. This “noise” floor results from errors in the waveform as computed by the simulator rather than from errors in the Fourier calculations. This is confirmed by running Spectre with RELTOL = 10^{-6} , as shown in Table 2. In this case, with Spectre the number and placement of the time-points during the Fourier analysis interval did not change, and so the accuracy of the Fourier analysis is unchanged. However, the tightened convergence criteria results in the waveform being computed more accurately and the “noise” floor dropping to -160 dB.

6.2 Distortion of a $\Sigma\Delta$ Modulator

An important feature of the Fourier integral is that it is not subject to aliasing, allowing it to be applied to broad-

3.1.3

harmonic	SPICE2g6	Spectre	Harm. Bal.
fund	997 mV	1 V	1 V
2 nd	-55.3 dB	-98.5 dB	-98.5 dB
3 rd	-58.1 dB	-120 dB	-120 dB
4 th	-53.3 dB	-159 dB	-169 dB
5 th	-55.9 dB	-142 dB	-143 dB
6 th	-55.9 dB	-164 dB	-196 dB
7 th	-53.3 dB	-165 dB	-165 dB
THD	0.48 %	0.0012 %	0.0012 %

Table 2: Results from $\mu\alpha 741$ with $REL TOL = 10^{-6}$.




Figure 1: Waveform generated by $\Sigma\Delta$ converter driven with a 1V sine wave.

spectrum signals. This is illustrated by computing the first 7 harmonics of the signal shown in Figure 1. This signal was generated by driving a clocked first-order two-level $\Sigma\Delta$ converter with a 1V sine wave. The signal has a period of 1 s and the transitions last 1 μs .

The output of the $\Sigma\Delta$ converter is a long sequence of integers. The distortion of the converter is determined solely from the value and the order of the integers, not the details of the waveform. This rather special characteristic allows the use of the FFT, which is employed here to verify the accuracy of the Fourier integral approach. To compute the Fourier coefficients using the FFT one simply samples the output once per clock cycle for 2^K cycles. The magnitude of the lowest harmonics computed from the $\Sigma\Delta$ converter waveform by the Fourier integral and FFT are shown in Table 3.

6.3 Distortion Products of a Mixer

In the previous example, it was possible, and generally desirable, to use the FFT because the desired spectrum could be computed from a uniform sampling of the output with one sample per clock period (it was possible and desirable to ignore the details of the waveform over individual clock cycles). However, uniform sampling is impractical when trying to compute the distortion of a mixer. Consider the example of a mixer from a hypothetical satellite receiver for video signals. The RF input frequency is 2.7GHz and the IF output frequency is 500MHz. To deter-

harmonic	Fourier Integral	FFT
fund	1 V	1 V
2 nd	-60.5 dB	-63.3 dB
3 rd	-54.7 dB	-54.3 dB
4 th	-54.9 dB	-55.3 dB
5 th	-51.5 dB	-51.6 dB
6 th	-47.3 dB	-48.9 dB
7 th	-47.0 dB	-47.7 dB
THD	0.964 %	1.414 %

Table 3: The magnitude of the lowest harmonics computed from the $\Sigma\Delta$ converter waveform by the Fourier integral and FFT. The discrepancy in THD results from the Fourier integral using far fewer harmonics (7 versus 512) when computing the THD.

harmonic	$REL TOL = 10^{-3}$	$REL TOL = 10^{-6}$
495 th	-56.5 dB	-95.4 dB
496 th	-81.7 dB	-93.1 dB
497 th	-60.0 dB	-68.2 dB
498 th	-52.1 dB	-53.3 dB
499 th	-6.06 dB	-6.05 dB
500 th	56.2 mV	56.3 mV
501 th	-6.05 dB	-6.05 dB
502 th	-57.7 dB	-53.3 dB
503 th	-63.9 dB	-68.4 dB
504 th	-77.4 dB	-98.1 dB
505 th	-56.8 dB	-101 dB

Table 4: The magnitude of the harmonics needed to determine intermodulation distortion of a mixer calculated using the Fourier integral.

mine the distortion of the mixer, the input is modulated at 1MHz. The IF frequency is the 500th harmonic of the modulation frequency. Spectre is requested to compute the five harmonics above and below the IF. The computation is still efficient, because the Fourier coefficients are being computed for only 11 harmonics. Of course, the simulation is slow because small time steps are needed to accurately follow the high frequencies and because of the need to simulate for at least one period of the modulation frequency. However, it is still faster than if the FFT were used because the output waveform has sharp transitions that would require a very small sampling interval. The minimum step size needed during the simulation was six times smaller than the average step size, implying that to get comparable accuracy with the FFT will require six times as many time point and so will be six times as slow. \square

3.1.4