

Analysing Forced Oscillators with Multiple Time Scales

Onuttom Narayan Jaijeet Roychowdhury
 UC Santa Cruz Bell Labs, Murray Hill
 USA USA

Abstract

We present a novel formulation, called the WaMPDE, for solving systems with forced autonomous components. An important feature of the WaMPDE is its ability to capture frequency modulation (FM) in a natural and compact manner. This is made possible by a key new concept: that of *warped time*, related to normal time through separate time scales. Using warped time, we obtain a completely general formulation that captures complex dynamics in autonomous nonlinear systems of arbitrary size or complexity. We present computationally efficient numerical methods for solving large practical problems using the WaMPDE. Our approach explicitly calculates a time-varying local frequency that matches intuitive expectations. Applications to voltage-controlled oscillators demonstrate speedups of two orders of magnitude.

1 Introduction

Oscillatory behaviour is crucial to the operation of many electronic systems, such as voltage-controlled oscillators (VCOs), phase-locked loops (PLLs), frequency dividers, $\Sigma\Delta$ modulators, etc.. It is, however, difficult to predict the response of such systems in a satisfactory and reliable manner. In this paper, we present the WaMPDE (Warped Multirate Partial Differential Equation), a new approach for analysing forced and unforced oscillatory systems. The WaMPDE provides a unified framework for treating phenomena like quasiperiodicity (particularly frequency modulation (FM)), mode-locking and period multiplication. (Conventional methods, discussed in Section 2, are typically error-prone and computationally intensive for oscillators in general, and especially for forced ones exhibiting FM-quasiperiodicity.) The key to the WaMPDE is a compact representation of FM signals using functions of several time variables, some of which are “warped”, *i.e.*, stretched or squeezed by different amounts at different times in order to make the density of the signal undulations uniform. The variation of this stretching is much slower than the undulations themselves, hence a multiple time approach is used to separate the time scales. A particularly important feature of the WaMPDE is that, unlike previous methods, it automatically and explicitly determines the local frequency as it changes with time. Our approach also eliminates the problem of growing phase error that limits previous numerical techniques for oscillators.

Numerical computations for the WaMPDE can be performed using time-domain or frequency-domain methods, or combinations. In particular, existing codes for previous methods like the MPDE and harmonic balance (see Section 2) can be modified easily to perform WaMPDE-based calculations. The use of iterative linear techniques (e.g., [Saa96, RLF98]) enables large systems to be handled efficiently.

The remainder of this paper is organized as follows. Section 2 contains a brief review of previous work. Section 3 is a tutorial-style exposition of the main concepts of the WaMPDE formulation, the mathematical details of which are presented in Section 4. In Section 5, the new methods are applied to practical VCO circuits and compared against existing techniques.

2 Previous Work

Most previous analyses of oscillators have typically apply purely linear concepts (e.g., [Ven82, Got97]) to obtain simple design formulae. Nonlinear analytical studies have largely been of polynomially-perturbed linear oscillators (e.g., [Far94]). For real oscillators, numerical simulation has been the predominant means of predicting detailed responses. A fundamental problem, however, is the intrinsic phase-instability of oscillators, leading to unbounded increase in phase error during simulation. Boundary-value methods like shooting (e.g., [NB95, TKW95]) and harmonic balance (e.g., [NB95, MFR95]) can be applied to unforced oscillators in steady-state, but not to forced oscillators with FM-quasiperiodic responses, which require an impractically large number of time-steps or variables (see Section 3). In prac-




Figure 1: Example 2-tone quasi-periodic signal $y(t)$




Figure 2: Corresponding 2-periodic bivariate form $\hat{y}(t_1, t_2)$

tice, the separation of the time scales is often reduced artificially to make the problem tractable. As illustrated in Section 5, such ad-hoc approaches can lead to qualitatively misleading results.

The warped-time approach presented in this paper is a generalization of a recent multi-time approach (the Multirate Partial Differential Equation (MPDE) [BWLBG96, Roy97, Roy98]) for non-autonomous systems with widely separated time scales. Earlier efforts to generalize the MPDE to autonomous systems [BL98] used non-rectangular boundaries to capture frequency variation. It has been shown [Roy98], however, that this approach is limited to oscillations that eventually become periodic, and cannot, for instance, accommodate FM-quasiperiodicity.

3 Essential concepts

In this section, we introduce several concepts at the core of this work. We first review why it is advantageous to use two or more time scales for analysing quasiperiodic signals, using amplitude-modulated (AM) signals for illustration. Then we show that although frequency-modulated (FM) signals can be quasiperiodic, the multi-time approaches that work for AM do not confer the same advantages. Next, we introduce the concept of warped time and show how it can be used to remedy the situation for FM. Finally, we outline the basic features of the WaMPDE.

Consider the waveform $y(t)$ shown in Figure 1, a simple two-tone quasiperiodic signal given by:

$$y(t) = \sin\left(\frac{2\pi}{T_1}t\right) \sin\left(\frac{2\pi}{T_2}t\right), \quad T_1 = 0.02s, \quad T_2 = 1s \quad (1)$$

The two tones are at frequencies $f_1 = \frac{1}{T_1} = 50\text{Hz}$ and $f_2 = \frac{1}{T_2} = 1\text{Hz}$, *i.e.*, there are 50 fast-varying sinusoids of period $T_1 = 0.02s$ modulated by a slowly-varying sinusoid of period $T_2 = 1s$. When such signals result from differential-algebraic equation (DAE) systems being solved by numerical integration (*i.e.*, transient simulation), the time-




Figure 3: \hat{x}_1 : simplistic bivariate representation of FM signal

steps taken need to be spaced closely enough that each rapid undulation of $b(t)$ is sampled accurately. If each fast sinusoid is sampled at n points, the total number of time-steps needed for one period of the slow modulation is $n\frac{T_2}{T_1}$. To generate Figure 1, 15 points were used per sinusoid, hence the total number of samples was 750. This number can be much larger in applications where the rates are more widely separated, e.g., separation factors of 1000 or more are common in electronic circuits.

Now consider a multivariate representation of $y(t)$, obtained as follows: for the ‘fast-varying’ parts of $y(t)$, t is replaced by a new variable t_1 ; for the ‘slowly-varying’ parts, by t_2 . The resulting function, now of two variables, is denoted by $\hat{y}(t_1, t_2)$:

$$\hat{y}(t_1, t_2) = \sin\left(\frac{2\pi}{T_1}t_1\right) \sin\left(\frac{2\pi}{T_2}t_2\right) \quad (2)$$

Note that $\hat{y}(t_1, t_2)$ is periodic with respect to both t_1 and t_2 , i.e., $\hat{y}(t_1 + T_1, t_2 + T_2) = \hat{y}(t_1, t_2)$. The plot of $\hat{y}(t_1, t_2)$ on the rectangle $0 \leq t_1 \leq T_1, 0 \leq t_2 \leq T_2$ is shown in Figure 2. Because \hat{y} is bi-periodic, this plot repeats over the rest of the t_1 - t_2 plane. Note also that $\hat{y}(t_1, t_2)$ does not have many undulations, unlike $y(t)$ in Figure 1. Hence it can be represented by relatively few points, which, moreover, do not depend on the relative values of T_1 and T_2 , unlike Figure 1. Figure 2 was plotted with 225 samples on a uniform 15×15 grid – three times fewer than for Figure 1. This saving increases with increasing separation of the periods T_1 and T_2 .

Note further that it is easy to recover $y(t)$ from $\hat{y}(t_1, t_2)$, simply by setting $t_1 = t_2 = t$, and using the fact that \hat{y} is bi-periodic. Given any value of t , the arguments to \hat{y} are given by $t_i = t \bmod T_i$.

The above discussion has illustrated two important features: 1. the bivariate form can require far fewer points to represent numerically than the original quasiperiodic signal, yet 2. it contains all the information needed to recover the original signal completely. These concepts are the key to the MPDE approach [BWLBG96, Roy97, Roy98] for analysing non-autonomous systems. We refer the reader to, e.g., [Roy98, Roy97] for further details of the non-autonomous case.

When the DAEs are autonomous, frequency-modulation (FM) can be generated. FM cannot, in general, be represented compactly as in Figure 2. We illustrate the difficulty with an example. Consider the following prototypical FM signal $x(t)$ with instantaneous frequency $f(t)$:

$$x(t) = \cos(2\pi f_0 t + k \cos(2\pi f_2 t)), \quad f(t) = f_0 - k f_2 \sin(2\pi f_2 t) \quad (3)$$

Following the same approach as for (1), a bivariate form can be defined to be

$$\hat{x}_1(t_1, t_2) = \cos(2\pi f_0 t_1 + k \cos(2\pi f_2 t_2)), \quad \text{with } x(t) = \hat{x}_1(t, t). \quad (4)$$

Note that \hat{x}_1 is periodic in t_1 and t_2 , hence $x(t)$ is quasiperiodic with frequencies f_0 and f_2 . Unfortunately, $\hat{x}_1(t_1, t_2)$, illustrated in Figure 3, is not a simple surface with only a few undulations like Figure 2. When $k \gg 2\pi$, i.e., $k \approx 2\pi m$ for some large integer m , then $\hat{x}_1(t_1, t_2)$ will undergo about m oscillations as a function of t_2 over one period T_2 . In practice, k is often of the order of $\frac{f_0}{f_2} \gg 2\pi$, hence this number of undulations can be very large. Therefore it becomes difficult to represent \hat{x}_1 efficiently by sampling on a two-dimensional grid. It is also clear,

from Figure 3, that representing (3) in the frequency domain will require a large number of Fourier coefficients to capture the undulations.

The WaMPDE approach of this work resolves this problem by bending the path along which $y(t)$ is evaluated away from the diagonal line $t_1 = t_2 = t$. Since along the bent path $t_2 = t$, but t_1 is no longer equal to t , we refer to t_1 as a warped time-scale. As mentioned in Section 1, this effectively results in stretching and squeezing the time axis differently at different times to even out the period of the fast undulations.

We illustrate warping by returning to (3). Consider the following new multivariate representation $\hat{x}_2(\tau_1, \tau_2)$, together with the warping function $\phi(\tau_2)$:

$$\hat{x}_2(\tau_1, \tau_2) = \cos(2\pi\tau_1), \quad \phi(\tau_2) = f_0\tau_2 + \frac{k}{2\pi} \cos(2\pi f_2\tau_2). \quad (5)$$

We now retrieve our one-dimensional FM signal (i.e., (3)) as

$$x(t) = \hat{x}_2(\phi(t), t). \quad (6)$$

Note that both \hat{x}_2 and ϕ , given in (5), can be easily represented with relatively few samples, unlike \hat{x}_1 in (4). Note further that $\phi(t)$ is the sum of a linearly increasing term and a periodic term, hence its derivative is periodic. This periodic derivative is equal to the instantaneous frequency, given in (3), of $x(t)$. For general FM signals (possibly with non-sinusoidal waveforms and varying amplitudes), this derivative, still well-defined, is the local frequency of the signal.

To find an efficient bivariate representation, a crucial step in our approach is to avoid specifying the function $\phi(t)$ a priori, but to impose a smooth ‘‘phase’’ condition instead on the bivariate function, and use this to calculate ϕ . The phase condition can, for instance, require that the phase of the τ_1 -variation of the function should vary only slowly (or not at all) as τ_2 is changed. Alternatively, a time-domain condition on the bivariate function (or a derivative) can be specified.

The above discussion has summarized our basic strategy for representing FM efficiently; it now remains to concretize these notions in the framework of an arbitrary dynamical system defined by DAEs. This is accomplished in the following section by the WaMPDE, which is a partial differential equation similar to the MPDE, but with a multiplicative factor of $\frac{\partial \phi}{\partial t}$ modifying one of the differential terms. By solving the WaMPDE together with the phase condition mentioned above, compact representations of the solutions of autonomous systems can be found by efficient numerical methods.

4 The Warped Multirate Partial Differential Equation (WaMPDE)

We consider a nonlinear system modelled using vector differential-algebraic equations (DAEs), a description adequate for circuits [CL75] and many other applications:

$$\frac{d}{dt} q(x(t)) + f(x(t)) = b(t) \quad (7)$$

In the circuit context, $x(t)$ is a vector of node voltages and branch currents; $q()$ and $f()$ are nonlinear functions describing the charge/flux and resistive terms, respectively. $b(t)$ is a vector forcing term consisting of inputs, usually independent voltage or current sources.

We now define the two-dimensional WaMPDE to be¹:

$$\omega(\tau_2) \frac{\partial q(\hat{x})}{\partial \tau_1} + \frac{\partial q(\hat{x})}{\partial \tau_2} + f(\hat{x}(\tau_1, \tau_2)) = b(\tau_2). \quad (8)$$

τ_1 is a warped time scale while τ_2 is unwarped. The warped time variable has an associated frequency function $\omega(\tau_2)$, which depends on the unwarped time variable. \hat{x} is bi-variate, i.e., a function of the two time variables.

The utility of (8) lies in its special relationship with (7). Consider any solution \hat{x} of (8), together with the relations

$$x(t) = \hat{x}(\phi(t), t), \quad \phi(t) = \int_0^t \omega(\tau_2) d\tau_2. \quad (9)$$

¹Far more general forms for the WaMPDE have been developed [NR98], of which (8) is a special case that suffices for the present exposition.

With these definitions, it can be shown by substitution that $x(t)$ above satisfies (7).

Next, we describe how (8) can be solved to determine $\hat{x}(\tau_1, \tau_2)$ and $\omega(\tau_2)$. We first assume that $\hat{x}(\tau_1, \tau_2)$ is periodic in τ_1 with period 1:

$$\hat{x}(\tau_1, \tau_2) = \sum_{i=-\infty}^{\infty} \hat{X}_i(\tau_2) e^{ji\tau_1} \quad (10)$$

We note that if $\hat{x}(\tau_1, \tau_2)$ satisfies (8), then so does for $\hat{x}(\tau_1 + \Delta, \tau_2)$, for any $\Delta \in \mathcal{R}$ – this is simply because (8) is autonomous in the τ_1 time scale. We remove this ambiguity in the same way as for unforced autonomous systems, *i.e.*, by fixing the phase of (say) the k^{th} variable to some value,² *e.g.*, 0. This is the phase constraint mentioned in Section 3.

We expand (8) in one-dimensional Fourier series in τ_1 , and also include the phase constraint, to obtain:

$$\sum_{i=-\infty}^{\infty} \left(\frac{\partial \hat{Q}_i(\tau_2)}{\partial \tau_2} + j\omega(\tau_2)\hat{Q}_i(\tau_2) + \hat{F}_i(\tau_2) \right) e^{ji\tau_1} = b(\tau_2) \quad (11)$$

$$\Im \left\{ \hat{X}_l^k(\tau_2) \right\} = 0 \quad (12)$$

$\hat{Q}_i(\tau_2)$ and $\hat{F}_i(\tau_2)$ are the Fourier coefficients of $q(\hat{x}(\tau_1, \tau_2))$ and $f(\hat{x}(\tau_1, \tau_2))$, respectively. k and l are fixed integers; $\hat{X}_l^k(\tau_2)$ denotes the l^{th} Fourier coefficient of the k^{th} element of \hat{x} .

(11) and (12) together form a DAE system which can be solved for isolated solutions. In practice, the Fourier series (11) can be truncated to $N_0 = 2M + 1$ terms with i restricted to $-M, \dots, M$. In this case, (11) and (12) lead to $N_0 + 1$ equations in the same number of unknown functions of τ_2 .

Applying periodic or initial boundary conditions to the DAE system (11) and (12) leads to quasiperiodic or envelope-modulated FM solutions, and also captures other interesting phenomena like mode locking and period multiplication. First, we consider periodic boundary conditions.

4.1 Quasiperiodic and envelope solutions

Assume $b(t)$ periodic with period T_2 or angular frequency $\omega_2 = \frac{2\pi}{T_2}$. Also assume that the solution of (8) is periodic in both arguments, *i.e.*, $\hat{x}(\tau_1, \tau_2)$ is $(1, T_2)$ -periodic and $\omega(\tau_2)$ is T_2 -periodic. $\omega(\tau_2)$ and $\phi(t)$ can then be written as:

$$\omega(\tau_2) = \omega_0 + p'(\tau_2), \quad \phi(t) = \omega_0 t + p(t) \quad (13)$$

where ω_0 is a constant, $p'(\cdot)$ is a zero-mean T_2 -periodic waveform, and its integral $p(t)$ is a T_2 -periodic function.

We motivate these assumptions by showing that such periodic forms for $\hat{x}(\cdot, \cdot)$ and $\omega(\cdot)$ capture FM- and AM-quasiperiodicity, mode-locking and period multiplication.

By expressing $\hat{x}(\tau_1, \tau_2)$ in Fourier series and substituting into (9), we obtain:

$$x(t) = \sum_{i,k=-\infty}^{\infty} \hat{X}_{i,k} e^{ji(\omega_0 t + p(t))} e^{jk\omega_2 t} \quad (14)$$

where the constants $\hat{X}_{i,k}$ are Fourier coefficients of \hat{x} .

Consider, for example, the term of (14) with $i = 1$ and $k = 0$:

$$\hat{X}_{1,0} e^{j(\omega_0 t + p(t))} = \hat{X}_{1,0} \cos(\omega_0 t + p(t)) + j \hat{X}_{1,0} \sin(\omega_0 t + p(t)) \quad (15)$$

When $\omega(t)$ is nontrivially T_2 -periodic, $p(t)$ is also nontrivially T_2 -periodic. (15) can then readily be recognized to be a frequency-modulated signal with instantaneous frequency $\omega(t)$. Hence the WaMPDE with periodic solutions can capture not only FM signals, but also the more general form of (14).

Next, we indicate how (11) and (12), with periodic boundary conditions, can be turned into a set of nonlinear equations for numerical

²or some slow function of τ_2 ; the selection of a slowly-varying phase condition is, in fact, the key to compact numerical representation of $\hat{x}(\cdot, \cdot)$.

solution³. (11) and (12) is discretized at N_1 points along the τ_2 axis, covering the interval $[0, T_1)$. The differentiation operator is replaced by a numerical differentiation formula (*e.g.*, Backward Euler or Trapezoidal), and when the periodic boundary condition $\hat{X}_i(0) = \hat{X}_i(T_1)$ is applied, a system of $N_1(N_0 + 1)$ nonlinear algebraic equations in $N_1(N_0 + 1)$ unknowns is obtained. This set of equations is solved with any numerical method for nonlinear equations, such as Newton-Raphson or continuation, to obtain the solution of the WaMPDE. Further, when iterative linear algebra and factored-matrix methods (*e.g.*, [Saa96, RLF98]) are employed, computation and memory requirements grow almost linearly with size, making calculations practical for even large systems.

By applying initial conditions rather than periodic boundary conditions, (11) and (12) can be solved for aperiodic $(\{\hat{X}_i(\tau_2)\}, \omega(\tau_2))$. These envelope-modulated solutions can be useful for investigating transient behaviour in systems with FM. To obtain envelope solutions, (11) and (12) are solved by time-stepping in τ_2 using any DAE solution method, starting from (say) $\tau_2 = 0$. An initial condition $(\{\hat{X}_i(0)\}, \omega(0))$ is specified. For typical applications, a natural initial condition is the solution of (7) with no forcing, *i.e.*, with $b(t)$ constant. The procedure for discretizing of the WaMPDE for quasiperiodic or time-stepping solutions is similar to that for the MPDE; further details may be found in [Roy98].

5 Applications

A voltage-controlled oscillator (VCO) was simulated using the new WaMPDE-based numerical techniques. The oscillator consisted of an LC tank in parallel with a nonlinear resistor, whose resistance was negative in a region about zero and positive elsewhere. This led to a stable limit cycle. The capacitance was varied by adjusting the physical plate separation of a novel MEMS (Micro ElectroMechanical Structure) varactor with a separate control voltage.

The damping parameter of the mechanical structure was initially assumed small, corresponding to a near vacuum.




Figure 4: VCO: frequency modulation

An envelope simulation was conducted using purely time-domain numerical techniques for both τ_1 and τ_2 axes. The initial control voltage of 1.5V resulted in an initial frequency of about 0.75MHz; the controlling voltage was a sinusoid with time-period 30 times that of the unforced oscillator. Figure 4 shows the resulting change in local frequency, which varies by a factor of almost 3.

Figure 5 depicts the bivariate waveform of the capacitor voltage (*i.e.*, one entry of the vector $\hat{x}(\tau_1, \tau_2)$), with the warped τ_1 axis scaled to the oscillator's nominal time-period of 1 μ s. It is seen that the controlling voltage changes not only the local frequency, but also the amplitude and shape of the oscillator waveform.

The circuit was also simulated by traditional numerical ODE methods ("transient simulation"). The waveform from this simulation, together with the 1-dimensional waveform obtained by applying (9) to Figure 5, are shown in Figure 6. The match is so close that it is difficult to tell the two waveforms apart; however, the thickening of

³We outline a time-domain method for the τ_2 axis, leading to a mixed frequency-time method; purely time-domain or frequency-domain methods are equally straightforward.




Figure 5: VCO: bivariate representation of capacitor voltage the lines at about $60\mu\text{s}$ indicates a deviation of the transient result from the WaMPDE solution. Frequency modulation can be observed in the varying density of the undulations.

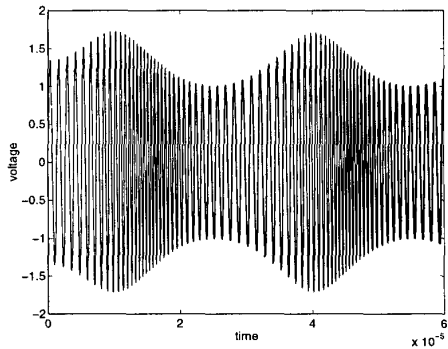


Figure 6: VCO: WaMPDE vs transient simulation

The VCO was simulated again after two modifications: the damping of the MEMS varactor was increased to correspond to an air-filled cavity, and the controlling voltage was varied much more slowly, *i.e.*, about 1000 times slower than the nominal period of the oscillator. The controlling voltage was sinusoidal with a period of 1ms. Figure 7 shows the new variation in frequency; note the settling behaviour and the smaller change in frequency, both due to the slow dynamics of the air-filled varactor.




Figure 7: Modified VCO: frequency modulation

The new bivariate capacitor voltage waveform (not shown) was a sinusoid along the warped time scale τ_1 , but unlike Figure 5, varied very little along the forcing time-scale. This was corroborated by transient simulation, the full results of which are not depicted due to the density of the fast oscillations. A small section of the one-dimensional waveform, consisting of a few cycles around 0.3ms, is

shown in Figure 8. The one-dimensional WaMPDE output of (9) is compared against two runs of direct transient simulation, using 50 and 100 points per nominal oscillation period, respectively. It can be seen that even at an early stage of the simulation, direct transient simulation with 50 points per cycle builds up significant phase error. This is reduced considerably when 100 points are taken per cycle, but further along (not shown), the error accumulates again, reaching many multiples of 2π by the end of the simulation at 3ms. In contrast, the WaMPDE achieves much tighter control on phase because the phase condition (a time-domain equivalent of (12)) explicitly prevents build-up of error. To achieve accuracy comparable to the WaMPDE, transient simulation required 1000 points per nominal cycle, with a resulting speed disadvantage of two orders of magnitude.




Figure 8: Modified VCO: WaMPDE vs transient (a few cycles at 10% of the full run; phase errors from transient increase later)

6 Conclusion

We have presented a new, efficient, approach for analysing the dynamics of oscillatory systems. The approach uses multiple time scales and time warping functions to obtain a partial differential formulation (the WaMPDE) for autonomous dynamical systems. Solving the WaMPDE by efficient numerical methods enables us to predict complex phenomena, such as frequency modulation, in large autonomous systems quickly and accurately. We have extended the notion of instantaneous frequency to general settings and provided methods for calculating it explicitly. We have applied our methods to VCO circuits and shown that they have significant speed and accuracy advantages over previously existing techniques.

Acknowledgments

We thank Alper Demir, Hans-Georg Brachtendorf, Wim Sweldens and Anirvan Sengupta for discussions of the FM representation problem.

References

- [BL98] H.G. Brachtendorf and R. Laur. Transient Simulation of Oscillators. Technical Report ITD-98-34096K, Bell Laboratories, 1998.
- [BWLBG96] H.G. Brachtendorf, G. Welsch, R. Laur, and A. Bunse-Gerstner. Numerical steady state analysis of electronic circuits driven by multi-tone signals. *Electrical Engineering (Springer-Verlag)*, 79:103–112, 1996.
- [CL75] L.O. Chua and P.M. Lin. *Computer-aided analysis of electronic circuits : algorithms and computational techniques*. Prentice-Hall, Englewood Cliffs, N.J., 1975.
- [Far94] M. Farkas. *Periodic Motions*. Springer-Verlag, 1994.
- [Got97] I.M. Gottlieb. *Practical oscillator handbook*. Oxford, 1997.
- [MFR95] R.C. Melville, P. Feldmann, and J. Roychowdhury. Efficient multi-tone distortion analysis of analog integrated circuits. In *Proc. IEEE CICC*, pages 241–244, May 1995.
- [NB95] A. Nayfeh and B. Balachandran. *Applied Nonlinear Dynamics*. Wiley, 1995.
- [NR98] O. Narayan and J. Roychowdhury. General Partial Differential Forms for Autonomous Systems, 1998. Bell Laboratories Internal Document.
- [RLF98] J. Roychowdhury, D. Long, and P. Feldmann. Cyclostationary Noise Analysis of Large RF Circuits with Multitone Excitations. *IEEE J. Solid-State Ckts.*, 33(2):324–336, Mar 1998.
- [Roy97] J. Roychowdhury. Efficient Methods for Simulating Highly Nonlinear Multi-Rate Circuits. In *Proc. IEEE DAC*, 1997.
- [Roy98] J. Roychowdhury. Analysing Circuits with Widely-Separated Time Scales using Numerical PDE Methods. *IEEE Trans. Ckts. Syst. - I: Fund. Th. Appl.*, 1998. In press.
- [Saa96] Y. Saad. *Iterative methods for sparse linear systems*. PWS, Boston, 1996.
- [TKW95] R. Telichevsky, K. Kundert, and J. White. Efficient Steady-State Analysis based on Matrix-Free Krylov Subspace Methods. In *Proc. IEEE DAC*, pages 480–484, 1995.
- [Ven82] G.D. Vendelin. *Design of amplifiers and oscillators by the S-parameter method*. Wiley, 1982.