Simulation of High-Q Oscillators

M. Gourary, S. Ulyanov, M. Zharov, S. Rusakov
NIISAPRAN, Russian Academy of Sciences, Moscow
K. K. Gullapalli, B. J. Mulvaney
Motorola Inc., Austin, Texas
{kiran, mulvaney, @adttx.sps.mot.com}

Abstract

We present a new technique, based on a continuation method, for oscillator analysis using harmonic balance. With the use of Krylov subspace iterative linear solvers, harmonic balance has become a very powerful method for the analysis of general nonlinear circuits in the frequency domain. However, application of the harmonic balance method to the oscillator problem has been difficult due to the very small region of convergence. The main contribution of this paper is a robust and efficient continuation method that overcomes this problem.

1 Introduction

Oscillators are used to generate the signals that are used as the time or phase references and are essential components in almost any electronic system. The design of an oscillator, a nonlinear circuit, requires a large signal steady state analysis. Using traditional time domain integration techniques for simulating these circuits can be very difficult, especially for high-Q crystal oscillators. Without special intervention, such simulations settle to the trivial DC solution. If oscillations do start, due to the high-Q, many cycles have to be simulated before the steady state is obtained. In addition, the truncation error tolerances need to be extremely stringent. The accuracy requirements and the high-Q mean that, in some cases, the simulation can take a few days. Even with more accurate integration schemes [1], time domain simulation of oscillators is inefficient.

Since many periodic waveforms typically consist of a fundamental frequency and a few harmonics, it is more natural to represent the waveforms with a truncated Fourier series. Contrast this with use of piece–wise polynomials used in time domain simulations. Using the Fourier representation converts the differential–alegbraic equations of the circuit equations to a set of nonlinear algebraic equations which connect the Fourier coefficients of the signals. The harmonic balance method achieves the simultaneous solution of these equations employing Newton–Raphson based methods. A very important advantage of harmonic balance is that the truncation error inherent in time domain analysis is absent when the waveforms are band limited. This makes harmonic balance the method of choice for a large class of important circuits in the wireless communications area. Also, the transient response of the circuit is completely ignored by the method, greatly speeding up the analysis of the steady state response.

Harmonic balance has been applied to oscillator simulation by adding the unknown frequency of oscillation to the set of state variables. However, oscillator simulation using harmonic balance has proven to be difficult due to the small region of convergence and the existence of the degenerate solution (the DC solution is a valid steady state). This is due to nature of the oscillator circuit itself, often requiring an extremely good guess for the frequency of oscillation and the oscillator waveforms to aid convergence. There have been many methods to improve the convergence of harmonic balance when applied to oscillator circuits [2], [3], [4]. However, these techniques are not always suitable. In this paper we propose a new continuation method that significantly improves our ability to simulate oscillators using harmonic balance.

We first present an overview of the theory of harmonic balance, including recent advances in the application of Krylov methods. We then discuss the oscillator problem and present a new continuation method. We conclude the paper with some examples that demonstrate the robustness and efficiency of the new method.

2 State of the art of Harmonic Balance

Harmonic balance is well established as a simulation technique for nonlinear circuits driven by one or more periodic inputs [5], [6], [7]. Harmonic balance exploits the fact that the periodic or quasi-periodic waveforms in many circuits are most compactly described in terms of their Fourier coefficients. If the period T of a waveform x contains M uniformly spaced samples, then,

$$x_l = \sum_{k=0}^{M-1} X_k e^{2\pi\imath\frac{kl}{M}} \quad l = 0, \ldots, M-1 \tag{1}$$

The goal of harmonic balance is to find the Fourier coefficients, X_k, of all the circuit waveforms. These are related by Kirchoff's current law (KCL), formulated in the frequency domain. KCL states that the algebraic sum of currents entering each node of a circuit must be zero.

Using the modified nodal analysis [8], the circuit can be described by a system of algebraic–differential equations:

$$f(x,t) + \frac{d}{dt}q(x,t) = 0 \tag{2}$$

where f is a vector of the conductive contributions to the circuit equations and q is the vector of charges and fluxes. Each node contributes on equation arising from KCL. In addition, there are

branch equations for the inductors and voltage sources. f and q are functions of x, which in addition to the node voltages v, can consist of inductor and independent voltage source currents.

Discretizing $\frac{d}{dt}q(x,t)$ yields the following nonlinear system of equations:

$$h = F + DQ = 0 \tag{3}$$

where D is a block matrix with diagonal blocks and represents the finite difference approximation of the differentiation operator, and F and Q are vectors of length NM, with N the number of nodes and M the number of samples of each waveform:

$$\begin{array}{l} F = [f_1^0 \cdots f_N^0 \cdots f_1^{M-1} \cdots f_N^{M-1}] \\ Q = [q_1^0 \cdots q_N^0 \cdots q_1^{M-1} \cdots q_N^{M-1}] \end{array} \tag{4}$$

To obtain the KCL equations in the frequency domain, we apply the discrete Fourier transform to Eq. 3 and obtain:

$$H = \Gamma F + \Gamma D \Gamma^{-1} \Gamma Q = 0 \tag{5}$$

where Γ is a block matrix of diagonal blocks and represents the discrete Fourier transform. Since differentiation in time domain transforms to multiplication by $\imath\omega$ in the frequency domain, we have

$$D = \Gamma^{-1} \imath\Omega\Gamma \tag{6}$$

where

$$\Omega = \begin{bmatrix} \omega_0 I & & \\ & \ddots & \\ & & \omega_{M-1} I \end{bmatrix} \tag{7}$$

and

$$\omega_i = \begin{cases} \frac{2\pi i}{T} & \text{if } 0 \le i \le \frac{M-1}{2} \\ \frac{2\pi(i-M)}{T} & \text{if } \frac{M-1}{2} < i < M \end{cases} \tag{8}$$

I is the $N \times N$ identity matrix. D is a dense block matrix with diagonal blocks, which exactly differentiates waveforms represented by truncated Fourier series.

The Jacobian, J, required for the solution of Eq. 5 using Newton–Raphson methods is

$$\Gamma \begin{bmatrix} G_0 & & \\ & \ddots & \\ & & G_{M-1} \end{bmatrix} \Gamma^{-1} + \imath\Omega\Gamma \begin{bmatrix} C_0 & & \\ & \ddots & \\ & & C_{M-1} \end{bmatrix} \Gamma^{-1} \tag{9}$$

or

$$J = \Gamma G \Gamma^{-1} + \imath\Omega\Gamma C \Gamma^{-1}$$

where G and C are the Jacobians of f and q respectively. G_i and C_i are $N \times N$ matrices.

Because each node in a circuit is typically connected only to a few other nodes, and all the Fourier coefficients are coupled, the Jacobian is a block dense matrix with the blocks themselves being sparse. The degree to which any two frequencies are coupled is determined by the nonlinearity of the circuit.

For robust convergence of the Newton method, it is necessary to invert the exact Jacobian. Applying direct linear solvers is prohibitive as the storage required is $O(NM^2)$ and the factoring cost is $O(N^\alpha M^3)$, $\alpha = 1.5 - 2$. This limited the broader application of harmonic balance for many years. In recent years, Krylov subspace based iterative methods have dramatically improved the situation. These methods solve $Ax = b$ by repeatedly performing matrix-vector multiplications involving A.

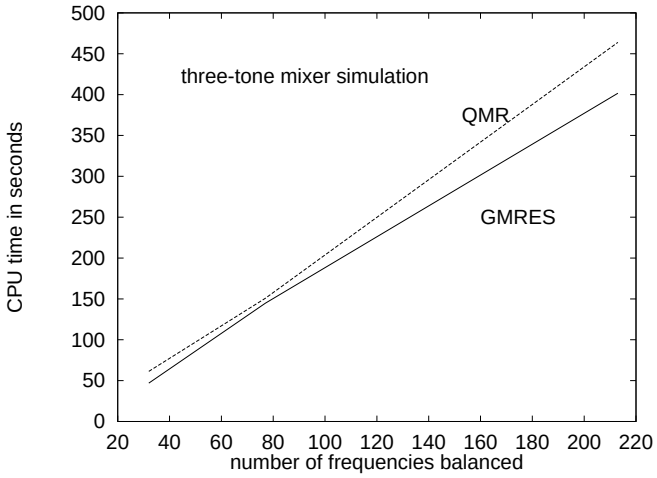

Figure 1: By using the FFT and the Jacobian decomposition in (10) in matrix–vector products, the cost of simulation grows almost linearly with number of frequencies. We compare both the QMR and GMRES methods

2.1 Iterative Linear Solvers

The computational cost of harmonic balance is dominated by the solution of

$$J\Delta X = -H$$

at each Newton iterate, where H is given in Equation 5. A typical RFIC receiver front-end can require the solution of $1,000,000$ equations. Direct methods have been adapted with much success to large, sparse linear systems that arise in circuit analysis. However, for the large dense system of equations that arise in harmonic balance, direct solution methods require excessive storage and computation and their use is prohibitive.

Krylov–subspace methods [9], when combined with suitable preconditioning, are powerful algorithms for the iterative solution of large, sparse linear systems. One of the features of Krylov-subspace methods is that they solve $Ax = b$ using the matrix A only in the form of matrix vector products, and thus they readily exploit the sparsity and structure of the linear problem.

Two Krylov methods exploited for harmonic balance are the quasi–minimal residual (QMR) [10], [11] and generalized minimal residual (GMRES) methods [12], [13].

2.2 Applying Krylov-subspace methods to Harmonic Balance

If the Jacobian in harmonic balance is explicitly formed by carrying out the product

$$J = \Gamma G \Gamma^{-1} + \imath\Omega\Gamma C \Gamma^{-1}, \tag{10}$$

the cost of the matrix vector products required in GMRES is $O(NM^2)$. The product JX can be performed more efficiently, in $O(NMlogM)$ operations, if we use the decomposed Jacobian and execute the multiplication from right to left:

$$JX = \Gamma(G(\Gamma^{-1}X)) + \imath\Omega\Gamma(C(\Gamma^{-1}X)) = \Gamma(Gx) + \imath\Omega\Gamma(Cx) \tag{11}$$

This exploits the FFT and the sparsity of the block diagonal G and C matrices. Figure 1 shows the cost of a three–tone mixer simulation as a function of the number of frequencies in the Fourier representation. Notice that the cost grows only slightly superlinearly with the number of frequencies. This circuit had 558 resistors, 216

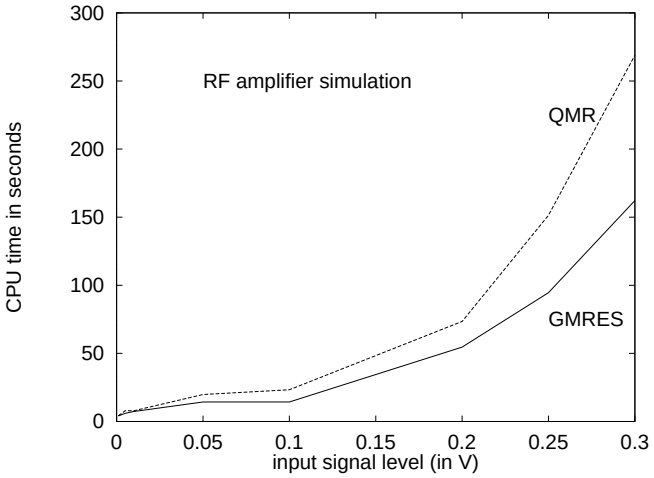

Figure 2: The block diagonal preconditioner (12) becomes increasingly ineffective as the circuit is driven with larger signals. Notice that GMRES is better than QMR

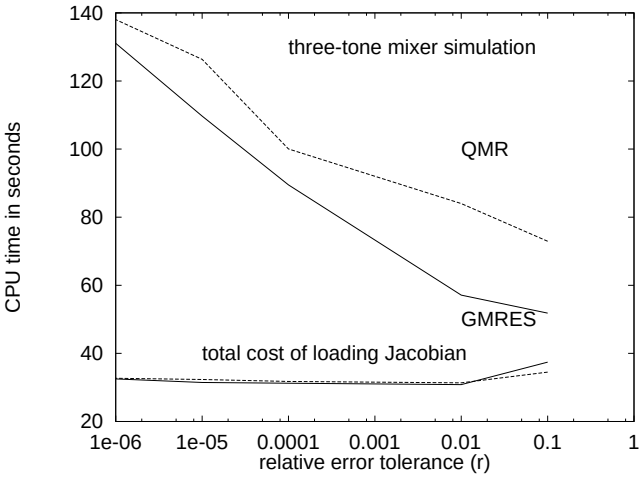

Figure 3: The cost of simulation as a function of the relative error tolerance in the linear iterative solver. The almost constant total cost of loading the Jacobian, proportional to the number of Newton iterations, shows that convergence of Newton's method is unaffected.

capacitors, 1 inductor, 4 diodes and 172 bipolar transistors, all together generating 900 nodes. Simulating this kind of circuit using direct sparse matrix techniques is not practical.

The almost linear–time cost of harmonic balance shown in Figure 1 also depends on the convergence of GMRES in a few iterations. To improve the convergence rate, GMRES is applied to a preconditioned version of J:

$$JP^{-1}Y = -H \quad PX = Y$$

where the preconditioner P should be an easily invertible approximation of J. If the circuit deviates from linearity only mildly, we can construct a good preconditioner by ignoring all frequency coupling. This yields the block diagonal matrix

$$P = \begin{bmatrix} \overline{G}_0 + \imath\omega_0\overline{C}_0 & & \\ & \ddots & \\ & & \overline{G}_0 + \imath\omega_{M-1}\overline{C}_0 \end{bmatrix} \tag{12}$$

where the elements of $\overline{G}_0$ and $\overline{C}_0$ are the time averages of the G and C matrices respectively. This preconditioner has proven to be very effective. For very nonlinear circuits, however, the preconditioner is not very effective as shown in Figure 2 where we compare the simulation cost of RF amplifier as the input signal to the circuit is increased, driving the circuit into increasingly nonlinear operation.

A better preconditioner can be constructed by pruning the Jacobian not all the way to a block–diagonal matrix, but to optimally include the extent of frequency coupling. The cost of applying the preconditioner increases with the extent of frequency coupling retained. But this can be more than offset by the reduction in the number of GMRES iterations required.

The efficiency of harmonic balance also depends on the error tolerances specified for GMRES. In Figure 3 we compare the simulation time as r in $\|Ax_n - b\| \leq r\|b\|$ is varied. The convergence of the Newton iterates is not effected by r but the cost of GMRES can vary significantly.

3 The Oscillator problem

The periodic steady state of an oscillator is determined not by external sources, but by the circuit itself, so the period is an additional unknown. With a small extension the harmonic balance method discussed in the preceding sections can be applied to the analysis of oscillators. One approach is to add the unknown frequency to the harmonic balance equations of the circuit [14]. An additional equation fixing the phase of the fundamental component of one waveform is added to choose one of many equivalent solutions.

$$H(X,\omega) = F(X) + \imath\Omega(\omega)Q(X) = 0 \tag{13}$$

$$\Im(X_n^1) = 0 \tag{14}$$

This set of equations can be solved using Newton–Raphson based methods where the Jacobian is

$$J_{osc} = \begin{bmatrix} J & \frac{\partial H}{\partial \omega} \\ (e_n^{1,im})^T & 0 \end{bmatrix} \tag{15}$$

where

$$\frac{\partial H}{\partial \omega} = \imath\frac{d\Omega}{d\omega}Q$$

J is the Jacobian in Eq. 9 and $e_n^{1,im}$ is the unit vector that selects the imaginary part of fundamental component of the n^{th} waveform.

However, without an extremely good initial guess for both the oscillation frequency and the node voltages, the method either tends to converge to the trivial DC solution, or fails to converge at all. Oscillators are nonlinear and very frequency selective and Newton's method can fail because the unknown frequency is adjusted at every iteration along with the node voltages. The node voltages at the intermediate iterates are such that KCL is not satisfied, and at some iterates, are quite unphysical. Updating the frequency based on such node voltages can lead to divergence, or make convergence difficult.

3.1 Review of past work

The oscillator problem consists of determining a good starting point for the node voltages and frequency of oscillation, and avoiding the DC solution. In this section we describe the techniques used to estimate the frequency of oscillation and past attempts to improve the convergence of harmonic balance in oscillator analysis.

The robustness and convergence rate of harmonic balance depends greatly on the initial guess for the frequency of oscillation.

circuit	*guess*		*actual*
	kurokawa	pole–zero	
1	10.0173MHz	10.01966MHz	10.01280MHz
2	2.23014MHz	2.267678MHz	2.209776MHz
3	723.7238KHz	732.8542KHz	730.2530KHz
4	159.155KHz	159.1544KHz	159.1549KHz
5	3.93757MHz	3.799555MHz	3.379978MHz
6	8.76376MHz	8.394621MHz	7.724068MHz
7	260.991KHz	260.5295KHz	273.3039KHz
8	1.21937KHz	1.188461KHz	1.213109KHz

Table 1: The difference between the large signal oscillation frequency and that predicted by linear circuit analysis for various oscillators.

This is most readily obtained by studying the circuit stability around DC. For some frequency, a circuit designed to be an oscillator will amplify small signals. As these signals grow, the nonlinearity of the circuit will limit the amplitude of oscillations and steady state is reached. One linear analysis technique is to study the Nyquist locus [15, 2]. A suitable value for the unknown frequency satisfies the following

$$\Im[\Delta(\omega_0)] = 0 \quad (16)$$
$$\Re[\Delta(\omega_0)] < 0 \quad (17)$$

where $\Delta = det[Y(\omega)]$ is the determinant of the admittance matrix, computed at the DC operating point. Another criteria, the Kurokawa condition [16] can also be used to estimate the frequency:

$$\Im[Z(\omega_0)] = 0 \quad (18)$$
$$\Re[Z(\omega_0)] < 0 \quad (19)$$

where Z is the driving point impedance at some port in the circuit. Another method is to obtain the poles, λ_i, of the circuit. Poles that lie in the right-half-plane indicate the frequencies at which small signal perturbations to the DC state will grow. So, a good starting frequency is

$$\omega_0 = \Im(\lambda) \quad (20)$$
$$\Re(\lambda) \geq 0 \quad (21)$$

With the recent advances in Krylov methods, pole–zero analysis can be performed very efficiently even for very large circuits.

It is very important to understand that the actual frequency of oscillation can be different from the frequency obtained using linear analysis techniques. Table 1 shows the difference between the oscillation frequency and the frequencies predicted by linear analysis. The difference is because the capacitances and conductances of semiconductor devices are nonlinear functions of the terminal voltages. While the differences between the estimated and actual frequencies is not large in most of the circuits, they are large when compared to the domain of convergence of the solution.

To determine the initial values for the node voltages, Rizzoli et al. [2] hold one Fourier component of one waveform fixed, while minimizing $H(X, \omega_T)$ using the quasi–Newton method. But the minimization problem can be as hard as the original problem of solving the nonlinear equations.

Recently, Ngoya et al. [4] presented a two-tier approach to improve oscillator simulation. They use the concept of a probe to create a circuit–probe combination for which KCL can be satisfied for any frequency. The probe is a voltage source at a specified frequency and an open circuit at all other frequencies. In terms of the circuit–probe combination, the oscillation condition is met when the probe current is zero for a finite probe voltage. Then all the circuit equations are satisfied without the probe being a part of the circuit.

In this approach oscillator simulation is formulated as a two–level problem. In the upper level, the following two–dimensional nonlinear problem:

$$\Re(I_{probe}(V_{probe}, \omega_{osc})) = 0 \quad (22)$$
$$\Im(I_{probe}(V_{probe}, \omega_{osc})) = 0 \quad (23)$$

is solved using a Newton method. V_{probe} is the probe voltage and I_{probe} is the probe current. To select one of the many equivalent solutions, the phase of V_{probe} is fixed, usually at zero. In the lower level, given the probe magnitude and frequency specified by the upper level iteration, the probe current is obtained using standard harmonic balance. Notice that the frequency is updated only in the outer level iterations, restricting the frequency iterates to follow only those paths along which the voltages are such that KCL is satisfied for the circuit–probe combination. This is a much more stable process.

The above method converts the analysis of an autonomous circuit into an analysis of a series of closely related nonautonomous circuits, easily handled using standard harmonic balance. We have found the two–level method to be more robust than the direct solution of Equations (13) and (14). However, the upper level Newton iteration requires the computation of the exact Jacobian

$$J_{probe} = \begin{bmatrix} \frac{\partial \Re[I_{probe}]}{\partial V_{probe}} & \frac{\partial \Re[I_{probe}]}{\partial \omega} \\ \frac{\partial \Im[I_{probe}]}{\partial V_{probe}} & \frac{\partial \Im[I_{probe}]}{\partial \omega} \end{bmatrix} \quad (24)$$

for the upper level 2×2 system. This Jacobian can be constructed by solving:

$$J\frac{\partial X}{\partial V_{probe}} = -\frac{\partial H}{\partial V_{probe}} \quad (25)$$
$$J\frac{\partial X}{\partial \omega} = -\frac{\partial H}{\partial \omega} \quad (26)$$

For each outer–level Newton iteration, this adds two extra linear system solves. Further, these linear systems must be solved very accurately since the solution is used in constructing the outer–level Jacobian. If the two–tier Newton scheme is applied directly, the extra cost of constructing the upper–level Jacobian is not significant because the outer–level Newton converges in a few iterations and the total number of linear solves is much greater than the number of linear solves needed for the sensitivity analysis. However, it is not very efficient for use in a curve tracing or continuation algorithm.

The two–tier method requires a guess for the amplitude of the fundamental component of the probe waveform. Ngoya et al. select this so that it minimizes the probe current at the initial estimate for the frequency. If the initial estimate is exactly the oscillation frequency, then the probe current vanishes when the probe amplitude is the first harmonic of the free running oscillator waveform at the probe terminals. If the frequency differs from the oscillation frequency then the probe current will exhibit a minima. The idea is that this can serve as a good starting point for the Newton iterations.

However, the starting point selected using this scheme may not belong to the domain of attraction of the solution. An even more serious limitation of the method is that it fails when the initial guess for the frequency is not close enough to the actual oscillation frequency. In some circuits, the probe current will exhibit a minimum only if the frequency is extremely close to the true oscillation frequency, as illustrated in Figure 4.

This is an important limitation since the frequency of oscillation is estimated using linear analysis techniques, and the minima we are searching for occurs far from the DC solution. Of course we can search for the minima of $|I_{probe}(V_{probe})|$ at several frequencies

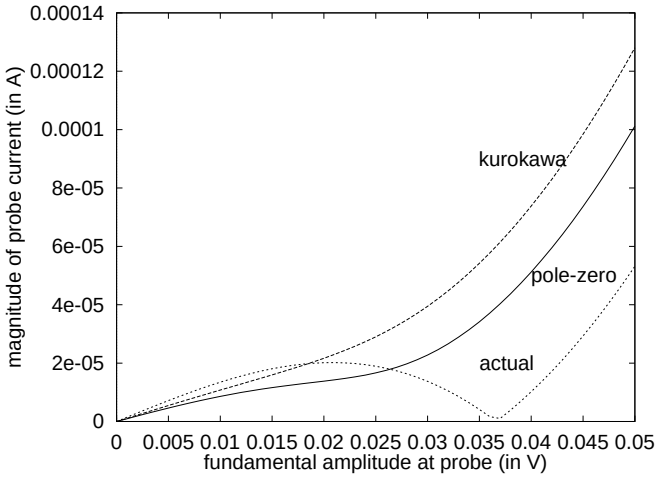

Figure 4: $|I_{probe}(V_{probe})|$ at $f = 8.763758$MHz estimated by *kurokawa's* condition, at $f = 8.394621$MHz estimated by *pole–zero* analysis and at the *actual* oscillation frequency $f = 7.724068$MHz. The scheme described by Ngoya et al. clearly fails in this case.

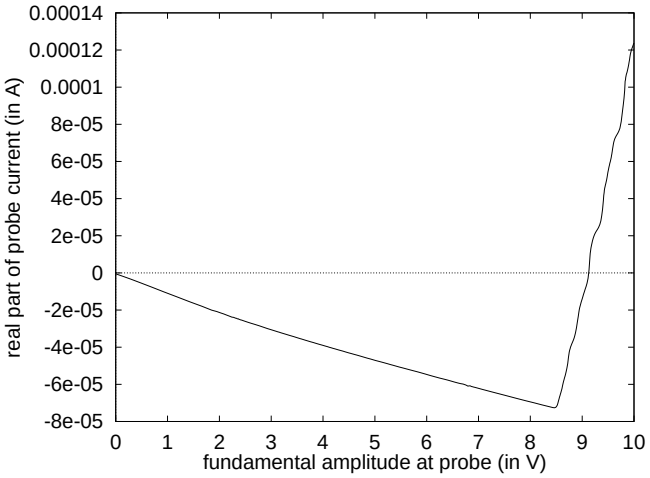

Figure 5: $\Re[I_{probe}(V_{probe})]$ vs V_{probe}. The oscillator solution is bracketed when $\Re[I_{probe}(V_{probe})]$ changes sign. The frequency is continually changing so that (V_{probe}, f) lies on $\Im[I_{probe}] = 0$.

in the neighbourhood of the frequency given by (17-21). However this search can be expensive since it involves repeatedly solving many harmonic balance problems for each trial frequency.

Before leaving this section, we note that in order to avoid converging to the DC solution, many authors [3], [2], [4] suggest solving the modified system:

$$\frac{H(X, \omega_{osc})}{\|X_{ac}\|} = 0 \tag{27}$$

where the $X_{ac} = X - X_{dc}$. This may not always work since it can still exhibit a local minimum at $X_{ac} = 0$ to which Newton's method can easily be attracted. Even otherwise, this function may not offer any advantages in terms of convergence of Newton's method.

Clearly, we need a method for oscillator simulation that does not require a good initial guess for the amplitude of oscillation.

3.2 A New Continuation Method

As shown in the previous section, the method of Ngoya et al. [4] is limited because $I_{probe}(V_{probe})$ may not have a minima if the initial guess for frequency is not close enough to the oscillation frequency. To develop a better method to generate a starting point for oscillator analysis, let us consider the process of oscillation build-up. When the oscillation is small amplitude and Kurokawa's condition is satisfied, $\Re[I_{probe}] < 0$. As the oscillation amplitude grows, the active devices and the positive feedback responsible for the negative input resistance begin to saturate, and pull the poles of the circuit toward the imaginary axis. Ultimately, the amplitude stabilizes and the probe current is zero. If the amplitude grows beyond this, the active devices respond so that the oscillations are damped. In other words, as the amplitude grows, the port resistance becomes positive and so does the probe current. As the oscillation amplitude changes, the frequency also changes.

This leads us to a natural continuation strategy that will identify or bracket the solution we are seeking. We retain the idea of a probe. But, instead of keeping the frequency fixed while searching for a good starting point for the probe amplitude, we fix the phase of I_{probe} to zero, and let the frequency vary. We then trace the curve

$$\Im[I_{probe}(V_{probe}, \omega)] = 0 \tag{28}$$

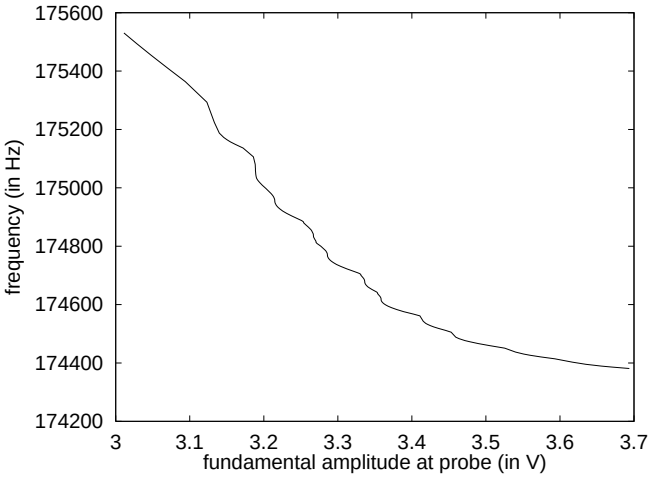

Figure 6: Curve $\Im[I_{probe}(V_{probe}, f = \frac{\omega}{2\pi})] = 0$ showing many turning points.

starting with a small signal for V_{probe}. As the probe amplitude increases, we let the frequency adjust to satisfy $\Im[I_{probe}] = 0$.

As the curve (28) is traced, a change in sign of $\Re[I_{probe}]$ signals that we have bracketed the solution for $\Re[I_{probe}] = 0$. Figure 5 shows $\Re[I_{probe}]$ as a function of V_{probe} for a typical oscillator.

This method is robust because the continuation method is started with a small signal excitation at the probe. So the frequency of oscillation estimated by linear analysis methods is valid. In addition, this method is also suitable for obtaining all steady state solutions when multiple solutions exist.

In tracing the curve (28), we can treat V_{probe} as a continuation parameter. Because the curve (28) can be very complex, as shown in Figure 6, with many turning points, we use a curve tracing algorithm [17], [18]. We use the backward–difference formulas as predictors, with automatic step–length and order control to efficiently trace (28). At each step along the curve we solve

$$H = F(X, V_{probe}) + \imath\Omega(\omega)Q(X, V_{probe}) = 0 \tag{29}$$

$$\Im[I^1_{probe}] = 0 \tag{30}$$

$$(V_{probe} - C_V)^2 + (\omega - C_\omega)^2 = \delta^2 \tag{31}$$

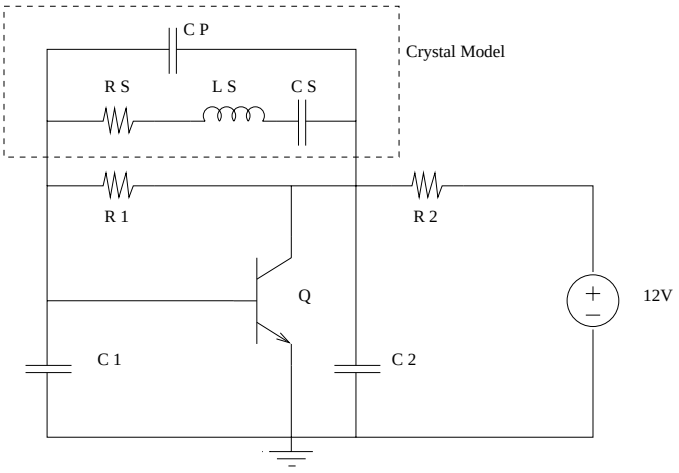

Figure 7: Schematic of a crystal oscillator.

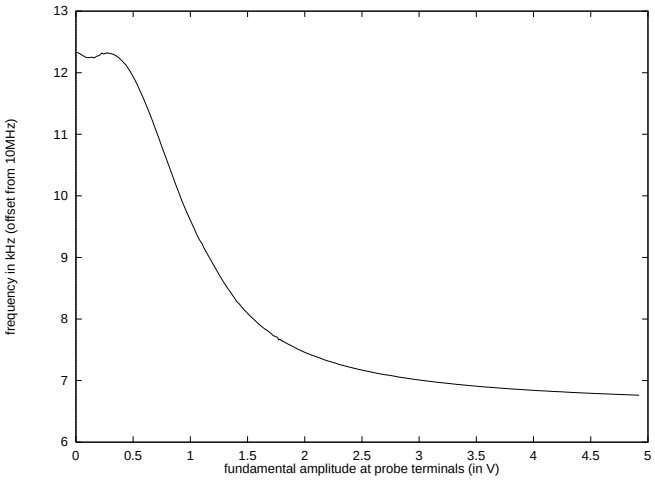

Figure 8: The curve $\Im[I_{probe}(V_{probe}, f)] = 0$ for the pierce oscillator.

where δ is the step size and (C_V, C_ω) is the previous point on the curve. V_{probe} is the value assigned to the real part of the probe voltage. The imaginary part of the probe voltage is set to zero, to pick on of many equivalent solutions. This set of equations can be solved using Newton–Raphson based methods where the Jacobian is

$$J_{trace} = \begin{bmatrix} J & \frac{\partial H}{\partial \omega} & \frac{\partial H}{\partial V_{probe}} \\ e^T & 0 & 0 \\ 0 & 2(\omega - C_\omega) & 2(V_{probe} - C_V) \end{bmatrix} \tag{32}$$

where

$$\frac{\partial H}{\partial V_{probe}} = \frac{\partial F}{\partial V_{probe}} + \imath\Omega \frac{\partial Q}{\partial V_{probe}}$$

and e is the unit vector that selects the imaginary part of fundamental component of the probe current which, in the MNA formulation, is part of X. A suitable preconditioner for (32) is

$$P_{trace} = \begin{bmatrix} P & \frac{\partial H}{\partial \omega} & \frac{\partial H}{\partial V_{probe}} \\ e^T & 0 & 0 \\ 0 & 2(\omega - C_\omega) & 2(V_{probe} - C_V) \end{bmatrix} \tag{33}$$

where P, given in Eq. 12, is the preconditioner used in standard harmonic balance.

The continuation method can also be formulated to trace the $\Re[I_{probe}(V_{probe})]$ characteristic as shown in Figure 5. We have also implemented an efficient method for this purpose by explicitly switching parameters between V_{probe} and I_{probe} at the turning points.

4 Examples

In this section we present a few examples of oscillator simulation. First, we consider the simulation of a Pierce oscillator. Then we simulate the characteristics of a 2 GHz VCO. We conclude this paper with an example of efficient oscillator startup simulation.

4.1 Pierce oscillator

Consider the Pierce oscillator shown in Figure 7. The element values for the Pierce oscillator are: R_1=100KΩ, R_2=2.2KΩ, C_1=100pF, C_2=100pF, C_P= 25pF, C_S=99.5fF, R_S = 6.4Ω and L_S=2.55mH. The bipolar transistor has $\beta = 500$. The circuit has poles at $(1.467978 \times 10^4, 6.295085 \times 10^7)$ and $(1.467978 \times 10^4, -6.295085 \times 10^7)$,

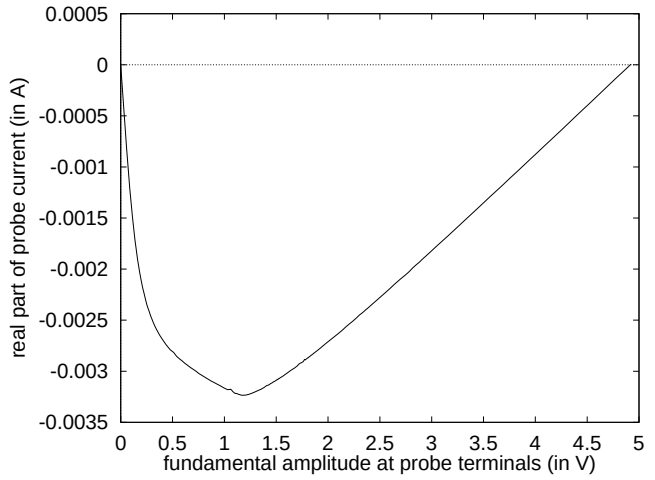

Figure 9: Plot of $\Re[I_{probe}(V_{probe})]$ for the pierce oscillator. The trace, including the final solution was completed in 15 seconds on a 200MHz ultraSPARC workstation, taking 150 steps.

that lie in the right half plane. So, we set $f = 10.01894$MHz as the starting point for the oscillation frequency. The crystal Q is approximately $\frac{\omega L_S}{R_S} = 2.5 \times 10^4$.

We introduce the probe across the collector and emitter terminals of the transistor. For this configuration, Figure 8 shows the curve $\Im[I_{probe}] = 0$, and Figure 9 shows the probe current as a function of probe voltage along this curve.

The Fourier expansion of the waveforms in this circuit was truncated at $M = 64$. The entire simulation took 15 seconds on a 200MHz ultraSPARC workstation. The frequency of oscillation is determined to be $f_{osc} = 10.00676335199140$MHz. The oscillator waveforms at the base and collector of the transistor are obtained using the inverse fast Fourier transform and shown in Figure 10.

We also attempted to simulate this oscillator with the method described in [4]. However, we could not obtain a suitable starting point for the probe amplitude required by the method. At $f = 10.01894$MHz, $|I_{probe}(V_{probe})|$ does not exhibit any minimum. At the exact frequency of oscillation, as determined by our continuation method, there exists a minimum (which is also a zero). However, at a frequency $f = 1.01 f_{osc}$, which is only slightly away

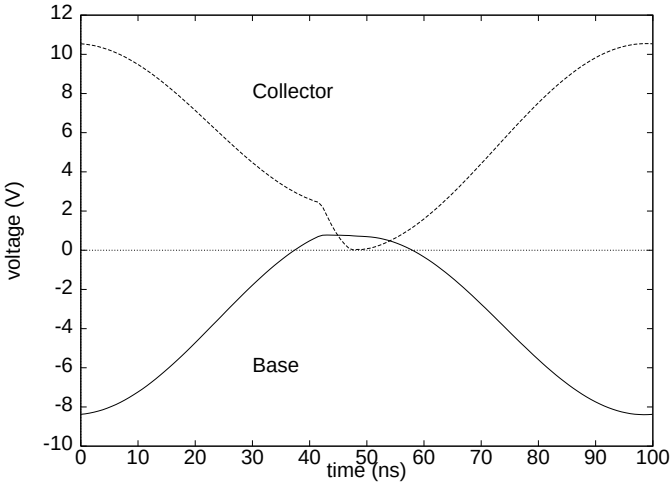

Figure 10: The base and emitter voltage waveforms for the Pierce oscillator.

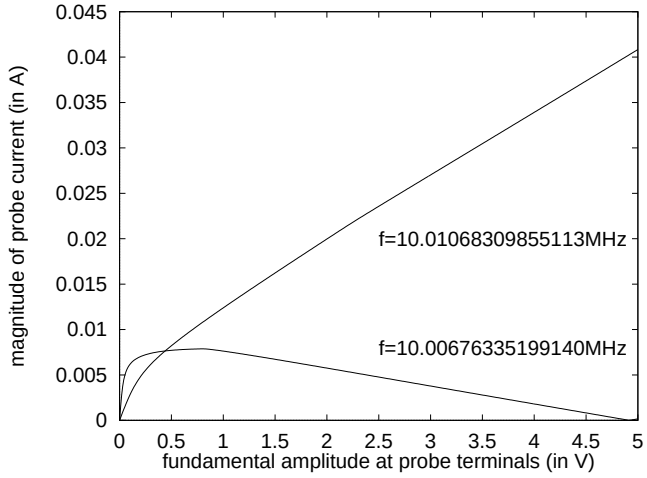

Figure 11: Magnitude of probe current as a function of probe amplitude at $f = f_{osc} = 10.00676335199140$MHz and at $f = 1.01 f_{osc} = 10.01068309855113$MHz. Because the frequency of oscillation is not known, the task of finding the minimum is frustrating.

from the exact oscillation frequency, this minimum disappears, as shown in Figure 11. Since it is impossible to guess the frequency of oscillation so accurately, the limitations of the method in [4] are clear.

Although the above circuit is simple, it illustrates the general difficulties in oscillator simulation and how our continuation method overcomes them. Our method was successfull in finding the solution in all circuits in our test-suite containing 15 oscillators. The largest of these contained 30 active devices. We found the 2-level direct newton method to be very sensitive to the supplied guess for the amplitude of oscillation. Contrast this with our method where the continuation is started with a small (or even zero) amplitude. This also makes our method faster compared to other methods because harmonic balance converges very quickly when the probe amplitude is small, and along the continuation path, we have a very good guess for the solution at each step. Each step along the continuation path takes only 2-3 iterations.

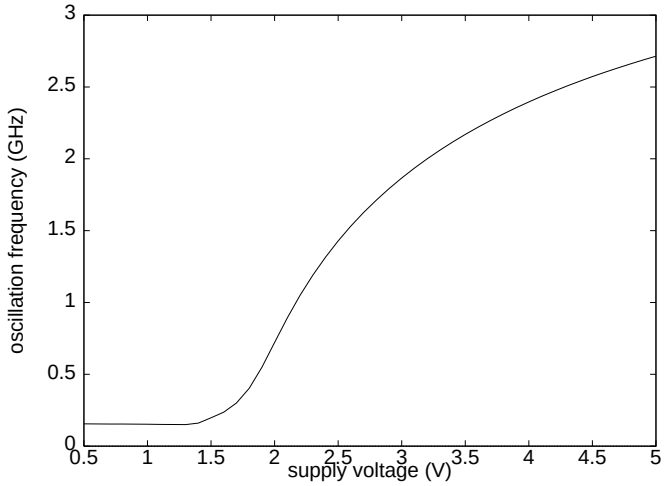

Figure 12: Dependence of the oscillation frequency on supply voltage.

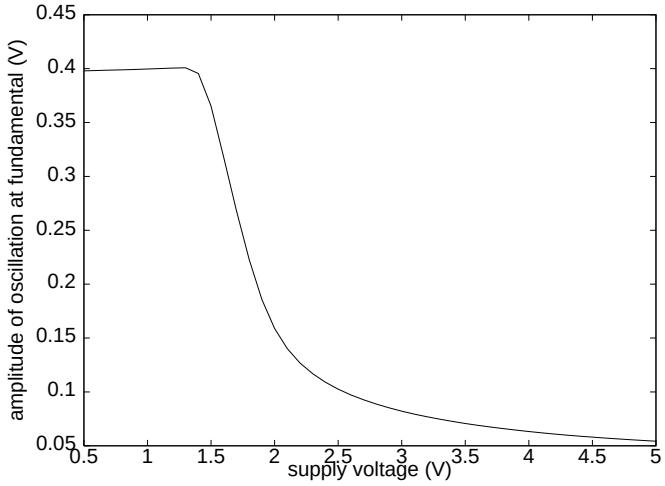

Figure 13: Oscillation amplitude at fundamental as a function of supply voltage.

4.2 VCO characteristics

Here, we simulate the characteristics of a VCO. It consists of 4 BJTs and a total of 30 nodes. Figure 12 shows how the frequency of the oscillator depends on the supply voltage. Figure 13 shows how the supply voltage effects the amplitude of oscillation at the fundamental. At low supply voltages, the oscillator waveform is almost a square wave, and at high supply voltages it is sinusoidal. The simulation, sweeping the supply voltage from 0.5V to 5.0V in steps of 0.1V took less than 10 minutes.

4.3 Simulating Oscillator Startup

In the introduction we mentioned the problem of simulating the long startup of a high-Q oscillator. Using a time-varying Fourier representation, we can write a differential equation for the evolution of the Fourier coefficients [19]. This allows us to follow the start up without having to take the extremely small time-steps necessary to capture the high frequency oscillations. The high frequency oscillations are simulated in the frequency domain using harmonic balance whereas the startup envelopes are simulated in the time domain. As an example, we simulate the startup of a crys-

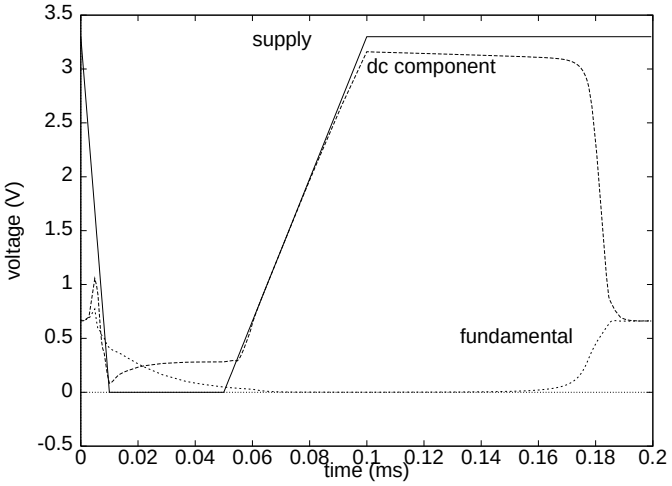

Figure 14: Fast simulation of startup transients using a mixed frequency-time method.

tal oscillator.

The circuit, a 20MHz oscillator, consisted of 30 MOSFETs and a total of 80 nodes. We used 4 harmonics in the Fourier representation of the oscillator waveforms and simulated the evolution of these harmonics as a function of time. The supply is initially on, then turned off, and then turned back on as shown in Figure 14. Also shown in the figure are the dc and fundamental components of the oscillator output. Notice the initial increase in the dc and fundamental components as the supply is turned off. This is due the active devices temporarily entering the high gain region of operation. Then, when the power supply is turned on, initially only the dc component of the oscillator waveform can respond to the increasing supply voltage. After some delay, the fundamental component starts to build-up. The entire simulation takes only 10 minutes compared to several hours using purely time domain simulation. All the features captured using the mixed frequency-time simulation were also observed in the purely time domain simulation.

References

[1] G. Denk, "Efficient simulation of a quartz oscillator circuit", *Zeitschrift fur Angewandte Mathematik und Mechanik*, vol. 76, no. suppl. 1, pp. 393–394, 1996.

[2] V. Rizzoli, A. Costanzo, and A. Neri, "Harmonic balance analysis of microwave oscillators with automatic suppression of degenerate solution", *Electronics Letters*, vol. 28, no. 3, pp. 256–257, Jan. 1992.

[3] C. R. Chang, M. B. Steer, S. Martin, and Jr. E. Reese, "Computer-aided analysis of free-running microwave oscillators", *IEEE Transactions on Microwave Theory and Techniques*, vol. 39, no. 10, pp. 1735–1745, Oct. 1991.

[4] E. Ngoya, A. Suárez, R. Sommet, and R. Quéré, "Steady state analysis of free or forced oscillators by harmonic balance and stability investigation of periodic and quasi–periodic regimes", *International Journal of Microwave and Millimeter-Wave Computer-Aided Engineering*, vol. 5, no. 3, pp. 210–233, Mar. 1995.

[5] K. S. Kundert and A. Sangiovanni-Vincentelli, "Simulation of nonlinear circuits in the frequency domain", *IEEE Transactions on Computer–Aided Design*, vol. CAD-5, pp. 521–535, Oct. 1986.

[6] V. Rizzoli and A. Neri, "State of the art and present trends in nonlinear microwave cad techniques", *IEEE Transcations on Microwave Theory and Techniques*, vol. 36, no. 2, pp. 343–364, Feb. 1988.

[7] R. J. Gilmore and M. B. Steer, "Nonlinear circuit analyis using the method of harmonic balance - a review of the art", *International Journal of Microwave and Millimeter-Wave Computer–Aided Engineering*, vol. 1, no. 4, pp. 159–180, Apr. 1991.

[8] J. Vlach and K. Singhal, *Computer Methods for Circuit Analysis and Design*, Van Nostrand Reinhold, 1983.

[9] Y. Saad, *Iterative Methods for Sparse Linear Systems*, PWS Publishing Company, 1996.

[10] H. G. Brachtendorf, G. Welsch, and R. Laur, "Fast simulation of the steady–state of circuits by the harmonic balance technique", *Proceedings of the International Symposium on Circuits and Systems*, pp. 1388–1391, 1995.

[11] R. C. Melville, P. Feldman, and J. Roychowdhury, "Efficent multi-tone distortion analysis of analog integrated circuits", *Proc. IEEE CICC*, May 1995.

[12] P. Heikkila, M. Valtonen, and T. Veijola, "Harmonic balance of nonlinear circuits with multitone excitation", *Proceedings of the 10th European Conference on Circuit Theory and Design*, vol. 2, pp. 802–811, 1991.

[13] M. Gourary, S. Rusakov, S. Ulyanov, M. Zharov, K. Gullapalli, and B. Mulvaney, "Iterative solution of linear systems in harmonic balance analysis", *IEEE MTT-S International Microwave Symposium Digest*, pp. 1507–1510, 1997.

[14] K. Kundert, J. White, and A. Sangiovanni-Vincentelli, *Steady–State Methods for Simulating Analog and Microwave Circuits*, Kluwer Academic Publishers, Boston, 1990.

[15] J. Paillot, J. Nallatamby, M. Hessane, R. Quere, M. Prigent, and J. Rousset, "A general program for steady state, stability, and fm noise analysis of microwave oscillators", *IEEE MTT-S International Microwave Symposium Digest*, pp. 1287–1290, 1990.

[16] K. Kurokawa, "Some basic characteristics fo broadband negative resistance oscillators", *Bell Syst. Tech. J.*, vol. 48, pp. 1937–1955, July 1969.

[17] A. Ushida and L. O. Chua, "Tracing solution curves of nonlinear equations with sharp turning points", *International Journal of Circuit Theory and Applications*, vol. 12, no. 1, pp. 1–21, Jan. 1984.

[18] K. Yamamura and T. Sekiguchi, "A modified spherical method for tracing solution curves", *IEICE Transactions: Fundamentals*, vol. E78-A, no. 9, pp. 1233–1238, Sept. 1995.

[19] J. Royychowdhury, "Efficient methods for simulating highly nonlinear multi-rate circuits", *Proc. Design Automation Conference*, pp. 269–274, 1997.