

Time-Mapped Harmonic Balance

Ognen J. Nastov
Motorola, Inc.
Austin, TX 78730
ojn@advtools.sps.mot.com

Jacob K. White
Massachusetts Institute of Technology
Cambridge, MA 02139
white@mit.edu

Abstract

Matrix-implicit Krylov-subspace methods have made it possible to efficiently compute the periodic steady-state of large circuits using either the time-domain shooting-Newton method or the frequency-domain harmonic balance method. However, the harmonic balance methods are not so efficient at computing steady-state solutions with rapid transitions, and the low-order integration methods typically used with shooting-Newton methods are not so efficient when high accuracy is required. In this paper we describe a Time-Mapped Harmonic Balance method (TMHB), a fast Krylov-subspace spectral method that overcomes the inefficiency of standard harmonic balance in the case of rapid transitions. TMHB features a non-uniform grid to resolve the sharp features in the signals. Results on several examples demonstrate that the TMHB method achieves several orders of magnitude improvement in accuracy compared to the standard harmonic balance method. The TMHB method is also several times faster than the standard harmonic balance method in reaching identical solution accuracy.

1 Introduction

The exploding demand for high performance wireless products has increased the need for more efficient and accurate simulation techniques for communication integrated circuits. Designers of such circuits are interested in some quantities which can be computed from small-signal analysis, but many, such as harmonic and intermodulation distortion, require the accurate computation of the circuit's steady-state. The two most commonly used approaches to computing a circuit's steady-state are the shooting-Newton method [1], and the Harmonic Balance (HB) method [5, 3]. Recent algorithmic developments, based on preconditioned matrix-implicit Krylov-subspace algorithms [4, 6, 8], have made these methods even more popular as now they can be used to easily analyze circuits with hundreds of devices.

The advantage of the shooting-Newton method is that it is a time domain method which can select time-points based on local error estimation. Therefore, shooting-Newton methods can easily handle circuits where the solution waveform has sharp transitions. The advantage of Harmonic Balance is that it is a spectrally accurate method, and therefore the solution converges exponentially fast with increasing harmonics. However, the effective time-steps used by the Harmonic Balance method are uniformly spaced, and this implies that the method requires a large number of harmonics

when the circuit solution contains very rapid transitions.

In this paper we describe a Time-Mapped Harmonic Balance method (TMHB), a fast Krylov-subspace spectral method utilizing a non-uniform grid to resolve the sharp features in the signals. At the core of the method are the grid selection strategies [7] and their use in construction of a time-map function specific to the simulated circuit. In the next section we overview the standard Harmonic Balance method. In Section 3 we detail the Time-Mapped Harmonic Balance algorithm. We derive the algorithm, give a Krylov-subspace based solution technique, describe the post-processing procedure used to obtain the actual Fourier coefficients from the TMHB solution, and detail the procedure used to construct the time-map function. In Section 4 we present results on several examples that demonstrate that the TMHB method achieves several orders of magnitude improvement in accuracy compared to the standard HB method. We also show that the TMHB method is several times faster than the standard HB method in reaching identical solution accuracy. Finally, conclusions are given in Section 5.

2 Standard Harmonic Balance

Consider a circuit described with N nonlinear differential equations:

$$\dot{q}(v(t)) + i(v(t)) + u(t) = 0 \quad (1)$$

where $v(t) \in \mathcal{R}^N$ is the vector of node voltages, $q(v(t)) \in \mathcal{R}^N$ the vector of node charges (or fluxes), $i(v(t)) \in \mathcal{R}^N$ the vector of resistive node currents, and $u(t) \in \mathcal{R}^N$ the vector of input sources.

Let the circuit be driven by a single periodic excitation input source with period T . Finding the periodic steady-state solution of this circuit consists of computing the N steady-state waveforms $v(t)$ on the solution domain $t \in [0, T]$. The periodic steady-state solution of (1) satisfies the two-point constraint:

$$v(T) = v(0) . \quad (2)$$

In the standard HB method, the solution waveforms are approximated with truncated Fourier series:

$$v(t) = \sum_{k=-K}^{k=K} V_k e^{j2\pi k f t} \quad (3)$$

with K the number of harmonics considered in the truncation. The method solves for the Fourier coefficients V_k . The approximation (3), in conjunction with the N circuit equations (1), results in the residual function:

$$f(V, t) \equiv \sum_{k=-K}^K j2\pi k f Q_k e^{j2\pi k f t} + i \left(\sum_{k=-K}^K V_k e^{j2\pi k f t} \right) + u(t) \quad (4)$$

which is to be minimized on $[0, T]$.

The minimization of (4) is typically carried out by enforcing $f(V, t) = 0$ on the uniform grid of M collocation (interpolation) time-points where $M = 2K + 1$. This standard HB method is more accurately referred to as pseudospectral Harmonic Balance [2].

Permission to make digital/hardcopy of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copying is by permission of ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC 99, New Orleans, Louisiana
(c) 1999 ACM 1-58113-109-7/99/06..\$5.00

3 Time-Mapped Harmonic Balance

In contrast to standard Harmonic Balance, Time-Mapped Harmonic Balance (TMHB) utilizes a non-uniform grid of time-points. The non-uniform grid is selected such that it has increased resolution in the high-gradient regions of the solution waveforms, i.e. it resolves the sharp waveform features in order to obtain greater solution accuracy. We now introduce the notion of pseudo time, where pseudo time \hat{t} is related to real time via the time-map function λ such that $t = \lambda(\hat{t})$, $\lambda(0) = 0$, and $\lambda(T) = T$. The time-map function maps a uniform grid of pseudo time-points into the non-uniform grid of real time-points.

The time-map function $\lambda(\cdot)$ is constructed in two phases. The first phase is computing a grid of non-uniform time-points. These non-uniform time-points are determined by examining the results from solving the periodic steady-state problem using a shooting-Newton time-domain method with a low-order integration scheme. In the second phase, the non-uniform grid is spectrally interpolated to yield $\lambda(\cdot)$. The details of this construction are given in Subsection 3.3.

To derive the Time-Mapped Harmonic Balance (TMHB) method, consider that

$$\frac{d}{dt} = \frac{1}{\lambda'(\hat{t})} \frac{d}{d\hat{t}}. \quad (5)$$

Replacing the time-derivative in (1) with (5) yields

$$\frac{1}{\lambda'(\hat{t})} \frac{d}{d\hat{t}} q(v(\lambda(\hat{t})) + i(v(\lambda(\hat{t}))) + u(\lambda(\hat{t}))) = 0, \quad (6)$$

and the two-point constraint becomes

$$v(\lambda(T)) = v(\lambda(0)). \quad (7)$$

The solution waveforms in TMHB are approximated with truncated pseudo Fourier series:

$$v(t) = v(\lambda(\hat{t})) = \sum_{k=-K}^{k=K} \hat{V}_k e^{j2\pi k f \hat{t}} \quad (8)$$

where \hat{V}_k are the pseudo Fourier coefficients of the solution waveforms. Equations (6) and (8) yield the residual function

$$\hat{f}(\hat{V}, \hat{t}) \equiv \frac{1}{\lambda'(\hat{t})} \sum_{k=-K}^K j2\pi k f \hat{Q}_k e^{j2\pi k f \hat{t}} + i \left(\sum_{k=-K}^K \hat{V}_k e^{j2\pi k f \hat{t}} \right) + u(\hat{t}) \quad (9)$$

which is to be minimized on $[0, T]$. The minimization is carried out by a collocation method, enforcing $\hat{f}(\hat{V}, \hat{t}) = 0$ on the uniform pseudo grid of collocation points.

The non-uniform grid in real time in effect ‘‘stretches’’ out those regions of the solution waveforms with sharp features. As a result, the TMHB solution $v(t)$ in real time is the smoother waveform $v(\lambda(\hat{t}))$ when viewed in pseudo time, as illustrated in Figure 1. Since the waveform is smoother in pseudo time, its features are more easily resolved with an M -point uniform pseudo grid, compared to resolving the original fast varying waveform in real time with an M -point uniform real time grid in the HB method. Thus one expects better accuracy from the TMHB method.

The rapid transitions in the solution waveforms are better approximated with the pseudo Fourier series (8), whose building blocks are complex exponential basis functions with smoothly varying frequencies. The greater accuracy of the TMHB method stems from the smaller global truncation error of the pseudo Fourier series for the (smoother) solution waveform in pseudo time, compared to the global truncation error of the standard HB Fourier series approximation of the solution waveform in real time.




Figure 1: The smoothing effect of the non-uniform grid of TMHB: (A) v_{COIL} in real time; (B) time-map function; (C) v_{COIL} in pseudo time.

3.1 Matrix-Implicit Krylov-Subspace Approach

Equation (9) is now rewritten in the frequency domain yielding NM nonlinear algebraic equations

$$\hat{F}(\hat{V}) \equiv \Gamma \Lambda \Gamma^{-1} \Omega \Gamma q(\Gamma^{-1} \hat{V}) + \Gamma i(\Gamma^{-1} \hat{V}) + \Gamma u = 0 \quad (10)$$

where Ω is the diagonal frequency-domain differentiation matrix

$$\Omega = \begin{bmatrix} j2\pi(-K)fI_N & & & & \\ & j2\pi(-K+1)fI_N & & & \\ & & \ddots & \ddots & \\ & & & \ddots & \\ & & & & j2\pi KfI_N \end{bmatrix}, \quad (11)$$

Λ is the diagonal matrix

$$\Lambda \equiv \begin{bmatrix} \frac{1}{\lambda'(\hat{t}_1)} I_N & & & & \\ & \frac{1}{\lambda'(\hat{t}_2)} I_N & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \frac{1}{\lambda'(\hat{t}_M)} I_N \end{bmatrix}, \quad (12)$$

and I_N is the identity matrix of size N . The matrices Γ and Γ^{-1} are DFT matrices that perform the conversions from pseudo time to frequency and vice-versa

$$v = \Gamma^{-1} \hat{V}, \quad \Gamma^{-1} = \begin{bmatrix} e^{j2\pi(-K)f\hat{t}_1} I_N & \dots & e^{j2\pi K f \hat{t}_1} I_N \\ \vdots & & \vdots \\ e^{j2\pi(-K)f\hat{t}_M} I_N & \dots & e^{j2\pi K f \hat{t}_M} I_N \end{bmatrix}. \quad (13)$$

Since the pseudo grid is uniform, the DFT can be carried out in $O(NM \log M)$ operations using the FFT just as in the standard HB.

Applying the Newton method to (10) results in the iteration

$$J^{(l)} (\hat{V}^{(l+1)} - \hat{V}^{(l)}) \equiv (\Gamma \Lambda \Gamma^{-1} \Omega \Gamma C^{(l)} \Gamma^{-1} + \Gamma G^{(l)} \Gamma^{-1}) (\hat{V}^{(l+1)} - \hat{V}^{(l)}) = -\hat{F}(\hat{V}^{(l)}) \quad (14)$$

where l is the Newton iteration index. The block-diagonal matrices

C and G are

$$C \equiv \begin{bmatrix} C_1 & & & \\ & C_2 & & \\ & & \ddots & \\ & & & C_M \end{bmatrix}, \quad G \equiv \begin{bmatrix} G_1 & & & \\ & G_2 & & \\ & & \ddots & \\ & & & G_M \end{bmatrix} \quad (15)$$

where $C_m = \frac{dq(v(\lambda(\hat{t}_m)))}{dv} = \frac{dq(v(t_m))}{dv}$ and $G_m = \frac{di(v(\lambda(\hat{t}_m)))}{dv} = \frac{di(v(t_m))}{dv}$, and can therefore be evaluated in real time on the non-uniform grid of real time-points t_m .

The Newton iteration (14) is a linear problem. Explicitly forming and factoring the dense TMHB Jacobian J is very expensive, $O(NM^3)$. As in standard HB, a preconditioned iterative linear solver such as GMRES can be used to reduce the complexity to $O(NM^2)$. Further reductions in complexity are obtained by implicitly forming the GMRES matrix-vector product by sequential evaluation using FFTs, to $O(NM \log M)$. The diagonal blocks of the Jacobian work well as a standard preconditioner in most circuit examples. Therefore the complexity of TMHB is the same $O(NM \log M)$ as the state-of-the-art matrix-implicit Krylov-subspace standard Harmonic Balance [4, 6, 8].

The TMHB yields the pseudo Fourier coefficients \hat{V} of the voltage waveform solutions that can be related to the real Fourier coefficients V . Note that if time-domain waveforms are desired, due to (8), an inverse FFT readily yields the voltage waveforms at the non-uniform grid of real time-points.

$$v(t) \equiv v(\lambda(\hat{t})) = \Gamma^{-1} \hat{V}. \quad (16)$$

3.2 Computing the Real Time Fourier Coefficients

To compute the real time Fourier coefficients V , we use the following “unmap” procedure. We introduce a non-uniform oversampled grid in pseudo time such that $\lambda(\cdot)$ maps this oversampled grid in pseudo time to a uniform oversampled grid in real time. Since $\hat{t} = \lambda^{-1}(t)$, (8) can be rewritten as

$$v(t) = \sum_{k=-K}^{k=K} \hat{V}_k e^{j2\pi k f \lambda^{-1}(t)}. \quad (17)$$

The summation in (17) is then evaluated to give the solution waveforms at the oversampled uniform grid in real time. Note that the summation cannot be carried out by an inverse FFT since the pseudo time-points $\lambda^{-1}(t_m)$ form a non-uniform grid. Finally, since the $v(t)$'s are now known on a uniform grid in real time, we can use the FFT to compute the real Fourier coefficients V .

Note that this procedure actually yields more than $M = 2K + 1$ Fourier coefficients. The additional Fourier coefficients represent the higher frequencies captured by the non-uniform grid in TMHB. These coefficients are shown to match the Fourier coefficients of the “exact” solution quite well (Figure 2). Without oversampling, these coefficients would be zero and some of the additional accuracy obtained by the TMHB method would be lost.

In effect the M pseudo Fourier coefficients \hat{V} “pack” high frequency information content, and in order to preserve it, we must carry the “unmap” procedure utilizing the oversampling frequencies above K .

The rate of oversampling is determined by the Nyquist frequency $f_\sigma = \frac{1}{2h_{min}}$ corresponding to the smallest spacing h_{min} in the non-uniform grid in real time, where from

$$K_\sigma = \left\lceil \frac{f_\sigma}{f} \right\rceil = \left\lceil \frac{1}{2h_{min}f} \right\rceil \quad (18)$$

where K_σ is the number of oversampling harmonics.

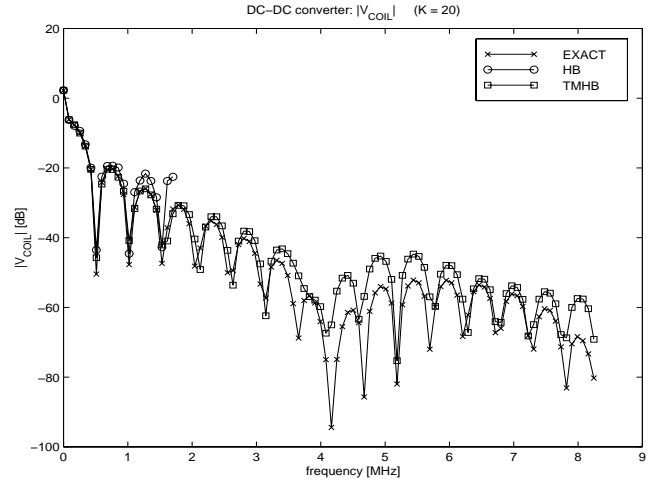


Figure 2: TMHB matching of high-frequency coefficients.

The unmap procedure described above is in essence an oversampled Fourier interpolation of the solution waveforms $v(t)$. This interpolation uses the discrete waveform values of $v(t)$ at the non-uniform grid in real time to generate the discrete values of $v(t)$ at the oversampled uniform grid in real time. It is crucial to use a spectrally accurate oversampled interpolation in order to preserve the accuracy of the solution. Local interpolation schemes (linear or quadratic) are not suitable for this task as they would introduce errors that are larger than the errors from the Fourier approximation of the solution.

3.3 The Time-Map Function

The first step in determining the time-map function λ is to determine a set of S non-uniform real time time-points. The success of the TMHB method is crucially dependent on this time-point selection [7], and the strategies used require an initial guess for the solution waveforms. In particular, an approximate solution is computed using a shooting-Newton method [5] with a low-order time integration scheme. The S non-uniform time-points for the TMHB method are then selected based on balancing two criteria: using small time-steps in the fast-varying regions of the approximate solution waveforms, and insuring that the time-steps do not change too rapidly. Although using a shooting-Newton method to compute the approximate solution is expensive, the cost is kept low by loosening the convergence tolerance. In addition, this shooting-Newton solution is useful as an initial guess for the TMHB.

Given the S non-uniform real time time-points, the next step is the construction of the time-map function $t = \lambda(\hat{t})$ that relates the uniform grid in pseudo time to the non-uniform grid in real time. In order to preserve the spectral accuracy of the TMHB method, the time-map function must be smooth, and we must be able to compute its first derivative with spectral accuracy or better as it is used in (9). Furthermore, to ensure strict monotonicity of the non-uniform grid of real time-points, the time-map function must be strictly monotonic, i.e. $\lambda'(\hat{t}) > 0$ for all $\hat{t} \in [0, T]$. Finally, for unmap purposes, we also need to be able to compute $\lambda^{-1}(t)$.

We first represent $\lambda(\hat{t})$ as a sum of a linear part and a T-periodic part $\lambda_\phi(\hat{t})$

$$t = \lambda(\hat{t}) = \hat{t} + \lambda_\phi(\hat{t}). \quad (19)$$

The periodic part $\lambda_\phi(\hat{t})$ is chosen to be a Fourier polynomial interpolant $\phi(\hat{t})$ of order S such that the interpolatory condition

$$t_s = \hat{t}_s + \phi(\hat{t}_s) \quad (20)$$

is exactly satisfied at the points (\hat{t}_s, t_s) where t_s are the S non-uniform real-time time-points, and \hat{t}_s are S uniform pseudo time-points. The interpolant $\phi(\hat{t})$ is the truncated Fourier series

$$\phi(\hat{t}) = \sum_{k=-J}^J \Phi_k e^{j2\pi k f \hat{t}} \quad (21)$$

where $2J + 1 = S$. The coefficients Φ_k can be computed with an inverse FFT of size S

$$\begin{bmatrix} \Phi_{-J} \\ \dots \\ \Phi_J \end{bmatrix} = \Gamma^{-1} \begin{bmatrix} t_1 - \hat{t}_1 \\ \dots \\ t_S - \hat{t}_S \end{bmatrix}. \quad (22)$$

Thus the time-map function is constructed as:

$$\lambda(\hat{t}) = \hat{t} + \sum_{k=-J}^J \Phi_k e^{j2\pi k f \hat{t}} \quad (23)$$

and this approximation exactly passes through the points (\hat{t}_s, t_s) .

The first derivative of the time-map function is

$$\lambda'(\hat{t}) = 1 + \sum_{k=-J}^J j2\pi k f \Phi_k e^{j2\pi k f \hat{t}} \quad (24)$$

and is exact.

The $\lambda(\cdot)$ function (23) and its first derivative (24) are now evaluated at M uniform pseudo time-points to yield the M -point non-uniform grid in real time and the matrix of time-map derivatives Λ .

Due to the Fourier nature of the representation (23), $\lambda(\hat{t})$ may exhibit high frequency oscillations and violate the monotonicity requirement. In practice, for the grids selected, if S is sufficiently large, this violation rarely happens, and can be resolved by damping the oscillations with an exponential filter μ_k , yielding a filtered construction

$$\lambda_\mu(\hat{t}) = \hat{t} + \sum_{k=-J}^J \mu_k \Phi_k e^{j2\pi k f \hat{t}} \quad (25)$$

where $\mu_k = e^{-\delta(\frac{k}{S})^\gamma}$ and δ and γ are filter parameters. Note that the filtered approximation no longer passes through the points (\hat{t}_s, t_s) . In addition, the filtered approximation can introduce an offset τ such that $\lambda(0) = \tau$ and $\lambda(T) = T + \tau$. While this offset causes no problems to the TMHB method, excessive filtering can deteriorate the quality of the approximation.

The values of $\lambda^{-1}(t)$ at the oversampled uniform times t_m is required in order to compute the actual Fourier coefficients. This is accomplished by applying Newton's method to the nonlinear equation $\lambda(\hat{t}_m) - t_m = 0$ and solving for \hat{t}_m at each time point t_m .

4 Results

In this Section we compare the performance of the TMHB method with standard state-of-the-art matrix-implicit Krylov-subspace Harmonic Balance [4, 6, 8]. Both the standard HB and TMHB methods were implemented in Mica, Motorola's SPICE-like circuit simulator.

The best candidates for the TMHB method are circuits whose solution waveforms undergo rapid transitions. Many highly nonlinear circuits will exhibit such waveforms. For these circuits the pseudo Fourier series solution representations of the TMHB method will be much more efficient than the standard Fourier series used in the standard HB method.

Four strongly nonlinear circuits were simulated with the HB and TMHB methods: a diode rectifier powered with a 50Hz sine

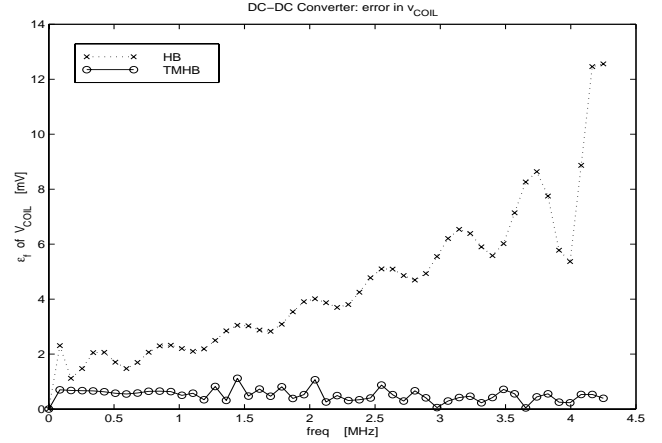


Figure 3: DC-DC converter circuit: error in the computed Fourier coefficients of v_{COIL} for $K = 50$, in dB.




Figure 4: DC-DC converter circuit, v_{COIL} computed with: (A) standard HB; (B) TMHB, at same number of harmonics $K = 50$.

input, a DC-DC converter with a 85kHz sine input, a BiCMOS switching mixer with a 1.8GHz square wave LO, and a BiCMOS IF preamplifier circuit driven into distortion with a 0.1V 110MHz sine input. Both the standard HB and TMHB methods in all runs used the same shooting-Newton solution guess. The Fourier coefficients V_k^* of the "exact" solution were computed using a standard HB method with a very large number of harmonics.

The four circuits were first simulated with both the standard HB and TMHB methods at a fixed number of harmonics. A plot of the frequency-domain pointwise error $\epsilon_f(kf) = |V_k - V_k^*|$ in dB in each computed Fourier coefficient V_k of the computed voltage v_{COIL} versus frequency in the DC-DC converter is shown in Figure 3 (the number of harmonics was $K = 50$). v_{COIL} was chosen because it is the signal with sharpest features in the circuit. The plot illustrates that the TMHB method computes each individual harmonic much more accurately than the standard HB. A plot of the computed v_{COIL} waveforms with HB and TMHB at $K = 10$ is given in Figure 4 illustrating the smaller TMHB error in the time domain. Similar results were observed for the computed waveforms in the other three circuits as well.

Next, the four circuits were repeatedly simulated with the standard HB and TMHB methods using increasing numbers of harmonics. Figures 5, 6, 7, and 8 show the L_∞ norm of the frequency-domain pointwise error ϵ_f in dB, for the computed Fourier coeffi-

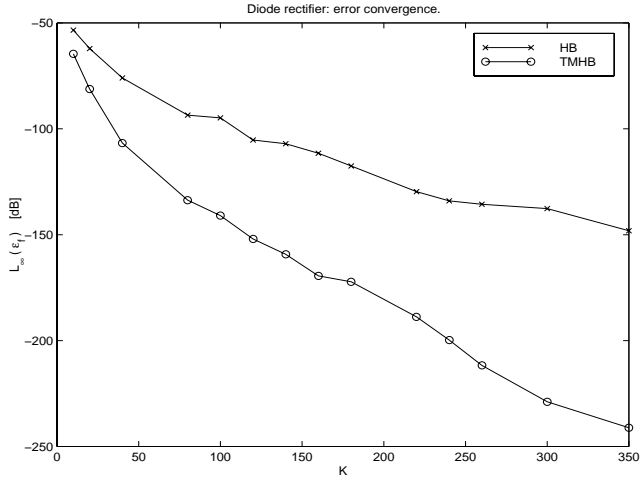


Figure 5: Diode rectifier circuit: L_∞ of the error in i_{V1N} , in dB.




Figure 7: Switching mixer: L_∞ of the error in i_{V31} , in dB.

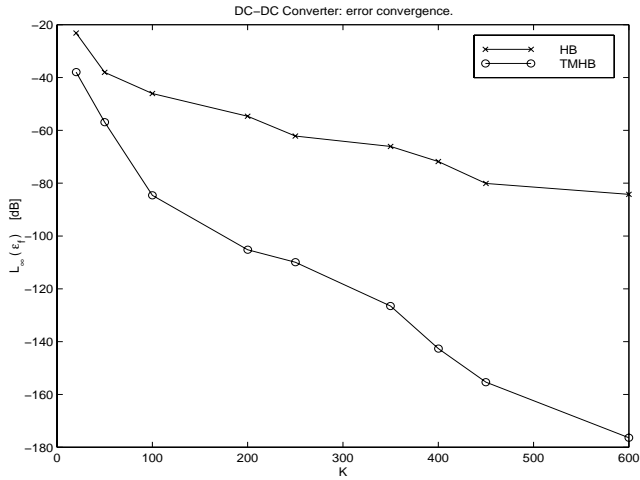


Figure 6: DC-DC converter circuit: L_∞ of the error in v_{COIL} , in dB.

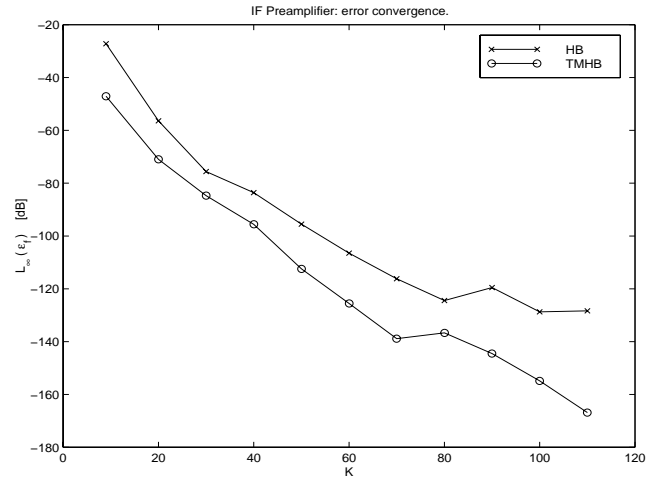


Figure 8: IF preamplifier: L_∞ of the error in v_{OUTP} , in dB.

coefficients of a selected voltage or current waveform versus the number of harmonics K . The plots show orders of magnitude improvements in the accuracy of the TMHB solution compared to the standard HB solution. For example, at $K = 200$ the TMHB solution of v_{COIL} in the DC-DC converter is about 100dB (5 orders of magnitude) more accurate than the standard HB solution. The L_∞ error as well as the errors in each individual harmonic for the remaining waveforms show the same superior error convergence properties.

4.1 Runtime Efficiency and Storage Requirements of TMHB

A logical way to measure the runtime efficiency of the TMHB method is to compare standard HB and TMHB runs achieving similar accuracies. Table 1 summarizes these findings. The results were obtained on a Sun Ultra-2 300MHz workstation. The accuracy ϵ_f was the L_∞ norm of the frequency domain pointwise error in the computed Fourier coefficients for the waveforms used for the error convergence profile plots.

The total CPU times include the time spent in the non-uniform grid selection, as well as a complete unmap of all solution waveforms in the circuit. The complete unmap of all solution waveforms

is in general unnecessary in practice as only a few waveforms are of interest. A partial unmap of only the few needed waveforms can generate significant total CPU time savings for larger circuits with hundreds of waveforms.

Table 1 shows significant CPU runtime speedups for three of the four simulated circuits. For both the diode rectifier and the IF preamplifier, a speedup of 1.6 is achieved. For the DC-DC converter the speedup is a factor of 6.

The total CPU times T for the HB and TMHB methods in reaching a specific accuracy ϵ_f in v_{COIL} from the DC-DC converter circuit are shown in Figure 9 (A). The accuracy measure was again the L_∞ norm of the frequency domain pointwise error in the computed Fourier coefficients. For less stringent accuracies, the total CPU times for the TMHB method are comparable to the HB CPU times due to the TMHB overhead in the non-uniform grid selection and waveform unmap. The situation is drastically different for accuracies better than -50dB: the TMHB becomes up to several times faster than the HB method. In addition the speedup factor grows with increases in required accuracies.

The memory storage requirements for the TMHB method are the same as for the standard HB method, growing linearly with the number of harmonics K due to the storage of the Krylov subspace vectors in the GMRES linear solver. Since the TMHB method can

Circuit	N	ϵ_f	Standard HB					TMHB				
			K	T	T_L	I_L	I_N	K	T	T_L	I_L	I_N
Diode Rectifier	6	-200	650	43.2	33.0	273	16	240	27.2	6.84	187	14
DC-DC Converter	9	-100	1000	1080	1053	2487	14	180	177	156	2112	12
Switching Mixer	105	-130	150	67.3	21.8	37	8	45	62.8	13.1	73	9
IF Preamplifier	289	-155	170	1065	861	417	18	90	662	514	441	17

Table 1: Comparison of the standard HB and TMHB methods at same achieved solution accuracy. N is the number of equations for the circuit and ϵ_f is the achieved accuracy in dB. K is number of harmonics, T is total CPU time, T_L is linear solve time, I_L is number of GMRES iterations, I_N is number of Newton iterations. All times are in seconds.




Figure 9: DC-DC Converter: (A) total CPU time T , (B) number of harmonics K for HB and TMHB to reach a specific solution accuracy ϵ_f in v_{COIL} .

achieve the same solution accuracy as the standard HB method with a smaller number of harmonics, it follows that significant memory savings can be achieved by using the TMHB method. In particular, from Table 1, we can measure the memory savings roughly as the ratio of the needed numbers of harmonics K for the standard HB and the TMHB method respectively. For example, the memory savings range from a factor of 1.9 for the IF preamplifier, to a factor of 5.5 for the DC-DC converter.

Figure 9 (B) shows the required numbers of harmonics K needed by the HB and the TMHB methods, versus the reached accuracy in the v_{COIL} waveform for the DC-DC converter circuit. Since the storage requirements are proportional to K the plot demonstrates that the TMHB method storage requirements at same solution accuracy are not only smaller than those of the HB, but also grow less rapidly for higher accuracy computations.

5 Conclusions

In this paper we described the Time-Mapped Harmonic Balance method (TMHB), a fast Krylov-subspace spectral method that overcomes the limitations of standard state-of-the-art Krylov-subspace harmonic balance method for circuits with rapid transitions. The non-uniform grid in the TMHB method resolves the sharp features in the signals. The computational results show that at same numbers of harmonics the TMHB method achieves up to five orders of magnitude improvement in accuracy. The TMHB method retains the same complexity as the standard HB method, is up to six

times faster than the standard HB method in reaching identical solution accuracy, and uses up to five times fewer harmonics and less computer memory. The TMHB runtime speedup factor and storage savings favorably increase for stricter accuracy requirements, making TMHB well suited for high accuracy simulations of large strongly nonlinear circuits with rapid transitions.

Acknowledgments

The authors would like to thank Kiran Gullapali, Jing Lee, and Brian Mulvaney for their valuable advice, help, and support during this research effort. This work was supported by grants from Motorola, Inc., and the MAFET Consortium.

References

- [1] Thomas J. Aprille and Timothy N. Trick. "Steady-State Analysis of Nonlinear Circuits with Periodic Inputs." *Proceedings of the IEEE*, Vol. 60, No. 1, pp. 108–114, January 1972.
- [2] C. Canuto, M.Y. Hussaini, A. Quarteroni, and T.A. Zang. *Spectral Methods in Fluid Dynamics*. Springer-Verlag, Berlin, New York, 1987.
- [3] Rowan Gilmore and Michael B. Steer. "Nonlinear Circuit Analysis Using the Method of Harmonic Balance - A Review of the Art. Part I - Introductory Concepts." *Int. J. on Microwave and Millimeter Wave Computer Aided Engineering*, Vol. 1, No. 1, 1991.
- [4] P. Heikkilä. *Object-Oriented Approach to Numerical Circuit Analysis*. Ph.D. dissertation, Helsinki University of Technology, January 1992.
- [5] Kenneth S. Kundert, Jacob K. White, and Alberto Sangiovanni-Vincentelli. *Steady-State Methods for Simulating Analog and Microwave Circuits*. Kluwer Academic Publishers, 1990.
- [6] R. Melville, P. Feldmann, and J. Roychowdhury. "Efficient Multi-Tone Distortion Analysis of Analog Integrated Circuits". *Proceedings of the Custom Integrated Circuits Conference*, May 1995.
- [7] Ognjen J. Nastov and Jacob K. White. "Grid Selection Strategies for the Time-Mapped Harmonic Balance Simulation of Circuits with Rapid Transitions." *Proceedings of the IEEE Custom Integrated Circuits Conference*, May 1999.
- [8] R. Telichevesky, K. Kundert, and J. White. "Efficient Steady-State Analysis Based on Matrix-Free Krylov-Subspace Methods". *Proceedings of the IEEE Design Automation Conference*, pp. 480–484, 1995.