Q-45

AN ALMOST-PERIODIC FOURIER TRANSFORM FOR USE WITH HARMONIC BALANCE

Gregory B. Sorkin†, Kenneth S. Kundert‡, and Alberto Sangiovanni-Vincentelli
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley, CA 94720

ABSTRACT

Harmonic balance is a powerful technique for the simulation of mildly nonlinear microwave circuits. This technique has had limited application for the analysis of almost-periodic circuits, such as mixers, due to the difficulties of transforming waveforms from the time domain to the frequency domain and vice versa. In this paper, a new Fourier transform algorithm for almost-periodic functions (APFT) is developed that is both efficient and accurate. Unlike previous attempts to solve this problem, the new algorithm does not constrain the input frequencies and uses the theoretically minimum number of time points.

1: INTRODUCTION

Harmonic balance [kundert86] is a powerful technique for the simulation of nonlinear microwave circuits. Harmonic balance requires the evaluation of the device equations in the frequency domain for linear elements and in the time domain for nonlinear elements. The discrete Fourier transform (DFT) is used to provide the conversion between the two domains when the signals are periodic. Currently there is no satisfactory way to analyze nonlinear circuits such as mixers that have two or more input signals with arbitrary input frequency and power, and hence have signals that are nonperiodic.

Signals in mixers are made up of several sinusoids at possibly nonharmonically related frequencies, and hence are almost periodic [bohr47]. This paper introduces an accurate and efficient algorithm, the APFT, for computing the forward and inverse Fourier transforms for almost-periodic functions. Unlike previous methods, the APFT does not constrain the input frequencies and uses the theoretical minimum number of time points. Also presented will be the implementation of a harmonic balance-based circuit simulator that uses the new Fourier transform for almost-periodic functions.

1.1: Definitions

A waveform x is *periodic* with *period* T if $x(t) = x(t+T)$ for all t. A waveform is *almost periodic* if it can be uniformly approximated by the sum of at most a countable number of sinusoids [hale80]. Thus, waveforms of the form

$$x(t) = \sum_{\omega_k \in \Lambda} (X_k^C \cos\omega_k t + X_k^S \sin\omega_k t), \tag{1}$$

where $\Lambda = \{\omega_0, \omega_1, \omega_2, \cdots\}$ and $X_k^C, X_k^S \in \mathbf{R}$, are almost periodic. We use $AP(\Lambda)$ to denote the set of all real-valued almost-periodic functions over the set of frequencies Λ. If Λ is finite with n elements, it is denoted Λ_n. When working with mixers, Λ consists of the sum and difference frequencies of usually two independent fundamental frequencies λ_1, λ_2 and their harmonics, though the transform that we present is not limited to this special case. For a mixer, Λ has the form

$$\Lambda(\lambda_1, \lambda_2) = \{\omega \mid \omega = k\lambda_1 + l\lambda_2;\ k,l \in \{\cdots, -1, 0, 1, \cdots\}\}. \tag{2}$$

If all frequencies $\omega_k \in \Lambda$ are distinct (i.e. $\omega_k \neq \omega_l$ for all $k \neq l$), then there exists an invertible linear operator $\mathbf{F}$ that maps x to X, where

$$X = [X_0^C, X_0^S, X_1^C, X_1^S, \cdots];$$

this operator is the Fourier transform.

Let x be a real-valued almost-periodic function, $x \in AP(\Lambda)$. For the problem of transforming between x and X to be numerically tractable, it is necessary to make Λ finite. In the case of a mixer, where $\Lambda = \Lambda(\lambda_1, \lambda_2)$, we assume that the Fourier coefficients of frequencies beyond a given order H are negligible, so that the truncated set of frequencies becomes

$$\Lambda_K(\lambda_1, \lambda_2) = \{\omega \mid \omega = k\lambda_1 + l\lambda_2;\ |k| + |l| \leq H,\ l \geq 0,\ l \neq 0 \text{ if } k < 0\}$$

The truncated set has $K = H^2 + H + 1$ frequencies. While the examples we present take this form, again, the APFT algorithm is not limited to this case.

1.2: The Fourier Transform

By considering only a finite number of frequencies, it is possible to sample a waveform at a finite number of time points and calculate its Fourier coefficients. Since the spaces involved are now finite dimensional, the first representation theorem of linear algebra shows that the discrete Fourier transform (DFT) and its inverse (IDFT) can be viewed as matrices acting on the vectors of samples and coefficients, respectively. That is,

$$\begin{bmatrix} 1 & \cos\omega_1 t_1 & \sin\omega_1 t_1 & \cdots & \cos\omega_{K-1} t_1 & \sin\omega_{K-1} t_1 \\ 1 & \cos\omega_1 t_2 & \sin\omega_1 t_2 & \cdots & \cos\omega_{K-1} t_2 & \sin\omega_{K-1} t_2 \\ 1 & \cos\omega_1 t_3 & \sin\omega_1 t_3 & \cdots & \cos\omega_{K-1} t_3 & \sin\omega_{K-1} t_3 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 1 & \cos\omega_1 t_{2K-1} & \sin\omega_1 t_{2K-1} & \cdots & \cos\omega_{K-1} t_{2K-1} & \sin\omega_{K-1} t_{2K-1} \end{bmatrix} \begin{bmatrix} X_0 \\ X_1^C \\ X_1^S \\ \vdots \\ X_{K-1}^C \\ X_{K-1}^S \end{bmatrix} = \begin{bmatrix} x(t_1) \\ x(t_2) \\ x(t_3) \\ \vdots \\ x(t_{2K-1}) \end{bmatrix}. \tag{3}$$

If the frequencies ω_k are distinct, this system is invertible (for almost all choices of time points), and can be written in compact form as $\Gamma^{-1}X = x$. Inverting Γ^{-1} gives $\Gamma x = X$, and so Γ and Γ^{-1} are a discrete Fourier transform pair.

The DFT is a special case of (3) in which $\omega_k = k\omega$ and $t_m = mT/(2K-1)$, i.e. when the frequencies are all multiples of a single fundamental and the time points are chosen equally spaced within its period. DFT and its inverse, IDFT, have the desirable property of being well conditioned, which is to say that very little error is generated when transforming between x and X. From the matrix viewpoint, the high accuracy of DFT corresponds to the fact that the rows of Γ^{-1} are orthogonal. (We will say a bit more about this later.) Unfortunately, DFT and IDFT are defined only for periodic signals.

Given a finite set Λ_K of distinct frequencies ω_k, and a set of time points, we will say that Γ and Γ^{-1}, are one implementation of the almost-periodic Fourier transform for $AP(\Lambda_K)$. To perform either the forward (using Γ) or inverse (using Γ^{-1}) transform requires just a matrix multiply, or $(2K-1)^2$ operations, once Γ and Γ^{-1} are known, which is the same number of operations required by the DFT.

The difficulty is that if the time points are not chosen carefully, Γ^{-1} can be very ill-conditioned. A particularly bad strategy for choosing time points when signals are not periodic seems to be that of making them equally spaced. Unlike the periodic case, it is in general impossible to choose a set of time points over which the sampled sinusoids at frequencies in Λ_K are orthogonal. In fact, it is common for evenly sampled sinusoids at two or more frequencies to be nearly linearly dependent, which causes the severe ill-conditioning problems encountered in practice. The main contribution of this paper is the development of an algorithm for choosing time points which gives a well-conditioned system. We will present previous work with the same goal and then present the APFT algorithm.

1.3: Previous Work

Ushida and Chua [ushida84] use equally spaced points, but avoid the ill-conditioning problem by using extra time points. In doing so, the matrix Γ^{-1} becomes a tall rectangular matrix. To make the system square again, both sides of (3) are multiplied by $(\Gamma^{-1})^T$, which results in

† Also with IBM T. J. Watson Research Center, Yorktown Heights, NY 10598.

‡ Also with Hewlett-Packard, Santa Rosa, CA 95404.

0149-645X/87/0000-0717$01.00 © 1987 IEEE

$$(\Gamma^{-1})^T\Gamma^{-1}X = (\Gamma^{-1})^T x.$$

Thus (3) is converted into a least squares problem that is solved in the traditional manner using the normal equation. Unfortunately, the normal equation is notoriously ill-conditioned and so a new ill-conditioning problem may be introduced.

Gilmore [gilmore86] samples the waveform using several small sets of equally spaced time points. The DFT is then applied to each individually, however these sets have too few time points to prevent aliasing in the computed spectra. The aliasing is eliminated by taking an appropriate linear combination of the computed spectra. Since the DFT is used, the method is constrained to periodic signals, though it can be much more efficient that the standard DFT on sparse spectra. Its numerical stability is unknown.

2: THEORETICAL BACKGROUND

2.1: Condition Number and Orthonormality

It is now necessary to discuss the conditioning of a system of equations, a concept alluded to earlier. A good general discussion can be found in [golub83]. From our perspective, the condition number of a matrix is important because it is a measure of how much errors are amplified during the course of solving a matrix equation. For example, in the presence of no other errors, the uncertainty in the result is given by the condition number times the roundoff error. Formally, the condition number of a matrix Γ is defined as $\kappa(\Gamma) = \|\Gamma\| \, \|\Gamma^{-1}\|$.

The problem of ill-conditioning in (3) can be visualized by considering each equation as defining a hyperplane in the Euclidian space $\mathbf{R}^{2K-1}$. Let $\rho_i \in \mathbf{R}^{2K-1}$ be such that ρ_i^T is the ith row in Γ^{-1}, then the ith hyperplane is defined by $\rho_i^T X = x(t_i)$. Thus, ρ_i is a vector orthogonal to the hyperplane. The solution to (3) is the intersection of all the hyperplanes. If the system is degenerate because two or more planes are coincident, then the intersection is not a single point and the system of equations has an infinite number of solutions. If there are no coincident hyperplanes, but two or more of the planes are nearly parallel, then a unique solution exists, however very high precision arithmetic will be needed to find it accurately.

Because a system is degenerate if and only if the defining matrix has zero determinant, it is natural to try to relate the determinant to the condition number. In fact, though, there is no direct connection between determinant and condition number [golub83]. But a matrix is also degenerate if and only if there is a linear dependence among its row vectors, and so it is also natural to suppose that a matrix will have small (good) condition number if its rows are relatively orthogonal (and thus "far" from being linearly dependent). We now prove a statement along these lines.

Consider an invertible $n \times n$ matrix A. Suppose that the rows a_i of A, regarded as vectors, are nearly orthonormal. In particular, suppose that each vector has unit length and that the orthogonal component of each vector a_i with respect to the space S_i spanned by the others is at least $\alpha \le 1$ (it would be exactly 1 if the vectors were precisely orthonormal).

Now, when forming the product $A^{-1}A = I$, each row of A^{-1} can be thought of as determining a linear combination of the rows of A, which yields a row in the identity matrix — a vector of length 1. Suppose that the ith element in a row of A^{-1} had absolute value $r > 1/\alpha$. Then the component of the resulting linear combination in the direction orthogonal to S_i is determined solely by ra_i, and will have magnitude greater than $r\alpha > 1$. Since the linear combination was to be a vector of unit length, this is a contradiction. Thus no element of any row of A^{-1}, and thus no element of A^{-1}, has absolute value greater than $1/\alpha$.

From this it follows that since $A \in \mathbf{R}^{n \times n}$, $\|A^{-1}\|_\infty$ (the l_∞ norm of A^{-1}) is no more than n/α. Also, the fact that A was to have row vectors of norm 1 means that $\|A\|_\infty$ is no more than n, so that $\kappa(A) \le n^2/\alpha$.

In short, the near-orthonormality of a matrix places an upper bound on its condition number.

2.2: Condition Number and Time Point Selection

Given a finite set of frequencies Λ_K, any set of $N = 2K - 1$ time points yields a Γ^{-1} whose row vectors (consisting of a single 1 and a set of sine-cosine pairs) have norms between K and $1+\sqrt{2}(K-1)$, which is adequate to satisfy the essence of the demand above that the row vectors be of unit length. If we could find a set of time points producing row vectors that were also nearly orthogonal, according to the definition above, we would be assured of having a well-conditioned Γ.

But the relation between the time points and the orthogonality of the resultant row vectors is clearly rather involved; finding a set of times which define nearly-orthogonal row vectors seems to be quite difficult. One approach might be to write down *a priori* a set of orthogonal vectors and then look for time points that generate vectors close to these prespecified ones; this is equivalent to defining the approximate phases of each sine wave and looking for a time where every wave will be in the appropriate phase. This in turn can be thought of as a set of approximate equalities modulo 2π, but it is far from clear under what circumstances a solution exists or how to go about finding it.

Another approach might be to choose time points equally spaced within a time interval larger than the period corresponding to the smallest nonzero frequency in Λ_K. As we show later however, experimentally this method of time-point selection gives the worst results of any method we tried.

2.3: Condition Number and Aliasing

As mentioned previously, the condition number provides a measure of how much error is amplified during a calculation. Roundoff is one source of error in the transform, but there is another that is normally much larger, and that is error due to aliasing. Aliasing occurs when we truncate Λ to make it finite. The Fourier coefficients of the frequencies omitted from Λ are presumably small but may not be exactly 0, and thus these frequencies make a contribution to the vector x of samples; this contribution was unaccounted for in the calculation of Γ^{-1} and Γ. Because of this, the computation of X will be in error.

Fortunately, this error can be bounded. Suppose that the overlooked sinusoids contributed some "error" δx to the observed sample vector $x+\delta x$. From this we have calculated the Fourier coefficients $X+\delta X$ by the formula

$$X+\delta X = \Gamma(x+\delta x).$$

By construction we know that in the case where there is no such aliasing contribution, $X = \Gamma x$. Thus, $\delta X = \Gamma \delta x$, and $\|\delta X\| \le \|\Gamma\| \, \|\delta x\|$. But by definition $\kappa = \|\Gamma\| \, \|\Gamma^{-1}\|$, and as shown previously, for the l_∞ norm, $\|\Gamma^{-1}\| \ge K$. So $\|\Gamma\| \le \kappa/K$, and

$$\|\delta X\| \le \frac{\kappa}{K} \|\delta x\|.$$

That is, κ/K represents the degree to which the coefficients of aliasing frequencies are amplified in their contribution to the calculated Fourier coefficients. Thus, it is very important when aliasing is present to select a set of time points so that $\kappa(\Gamma^{-1})$ is small.

3: THE APFT ALGORITHM

3.1: Time Point Selection

The conception of the orthogonal-selection algorithm came from the some of theoretical ideas discussed above.

First, we thought that if selecting evenly-spaced time points was likely to yield row vectors particularly close to being linearly dependent, we might be better off selecting time points randomly from a time interval larger than the period corresponding to the smallest nonzero frequency in Λ_K. (We chose an interval equal to three times this period.) Such a choice is particularly attractive given the complexity of the relationship between the time points and the orthogonality of the row vectors; making any more intelligent choice of time points seem quite difficult.

Second, we realized that in essence the problem in recovering X from x is that the linear system may be close to being underdetermined, in a numerical sense. So adding additional equations should increase the accuracy of the calculation of X. In fact, if more than N time points are chosen, Γ^{-1} becomes a tall rectangular matrix, and its pseudo-inverse Γ is a wide rectangular matrix satisfying $X = \Gamma x$.

Oversampling with twice as many randomly-selected time points as theoretically necessary proves to be successful: it yields a very well conditioned system. However, when using the transform in the context of harmonic balance, all the nonlinear devices must be evaluated at each time point, which is an expensive operation because of the complexity of the nonlinear device models. Thus oversampling is a costly remedy. However, it made it clear that the rows of the tall Γ^{-1} matrix spanned the space well (in a numerical sense), and we wondered whether some subset of them might not suffice to do the same job.

The *orthogonal selection* algorithm solves just this problem; from a Γ^{-1} whose dimension is $2N$ by N, where $N = 2K - 1$, it selects a set of just N rows, thus requiring no extra time samples. Because of the expense of sampling, the fact that only the minimum number of time points is used, and not 1.5 to 2 times as many as required by the other methods, is one of the significant advantages of the APFT algorithm.

The orthogonal selection algorithm is a variation of the Gram-Schmidt orthogonalization procedure [dahlquist74]. Its input is the matrix formed by randomly choosing twice as many time points as necessary and forming the corresponding row vectors, ρ_i. Initially, these vectors all have the same Euclidean length (i.e., l_2 norm). One of these vectors, say ρ_1, is chosen arbitrarily. Any component in the direction of ρ_1 is removed from the remaining vectors using

$$\rho_i \leftarrow \rho_i - \frac{\rho_1^T \rho_i}{\rho_1^T \rho_1} \rho_1 \qquad i = 2, \cdots, 2N. \tag{4}$$

The vectors that remain are now orthogonal to ρ_1. Since the vectors were initially the same length, the largest remaining vector was originally most orthogonal to ρ_1. It is chosen to play the role of ρ_1 for the next iteration of the algorithm. This process repeats until the required N vectors have been chosen. The time points that correspond to these vectors are the time points used to form Γ^{-1}. This algorithm is detailed below.

APFT Orthogonal Selection Algorithm

Given:

$\Lambda_K = \{0, \omega_1, \omega_2, \cdots, \omega_{K-1}\}$, the set of frequencies.

Task:

To find a set of $N = 2K-1$ time points that results in a well-conditioned Γ^{-1}.

Algorithm:

$\omega_{min} \leftarrow \min(\{|\omega_k| : 1 \le k < K\})$

for ($i \leftarrow 1, \cdots, 2N$)

{ **random()** *is a function that returns numbers uniformly distributed between 0 and 1.*

$t_i \leftarrow \frac{6\pi}{\omega_{min}}$ **random()**

$\rho_i^{(1)} \leftarrow [\,1 \;\; \cos\omega_1 t_i \;\; \sin\omega_1 t_i \;\; \cdots \;\; \cos\omega_{K-1} t_i \;\; \sin\omega_{K-1} t_i\,]^T$

}

for ($j \leftarrow 1, \cdots, N$)

{ **argmax()** *is a function that returns the index of the largest member of a set*

$k =$ **argmax**$(\{\|\rho_i^{(j)}\| : j \le i \le 2N\})$

swap$(\rho_j^{(j)}, \rho_k^{(j)})$

swap(t_j, t_k)

for ($i \leftarrow j, \cdots, 2N$)

$$\rho_i^{(j+1)} \leftarrow \rho_i^{(j)} - \frac{\rho_j^{(j)T} \rho_i^{(j)}}{\rho_j^{(j)T} \rho_j^{(j)}} \rho_j^{(j)}$$

}

Results:

The set $\{t_i : 1 \le i \le N\}$ contains the desired time points.

Once the time points are selected, Γ^{-1} is constructed with the rows $\rho_i^{(1)}$. It is easy to verify that the time points were well-chosen either by calculating the condition number $\kappa = \|\Gamma\| \, \|\Gamma^{-1}\|$ or by computing $\varepsilon = \|\Gamma^{-1}\Gamma - I\|$ where I is the identity matrix; both are excellent measures of the numerical stability of the transform.

3.2: Constructing the Transform Matrix

There is another problem which we have heretofore ignored. The arguments to the sine and cosine functions in (3) are potentially very large, which results in excessive roundoff error. For example, assume $\lambda_1 = 2\pi 10^9$ and $\lambda_2 = 2\pi(10^9+\sqrt{2})$. Then $\omega_{min} = 2\pi\sqrt{2}$ and so the time points will fall in the range of 0 to $3\sqrt{1/2}$ seconds. Thus, $\omega_j t_i$ could be as large as 10^{11}, which causes two problems. First, on most computer systems, the trigonometry routines are not designed to handle such large arguments and will return meaningless results. This problem is easily avoided by subtracting from the argument as many multiples of 2π as possible without making it negative. The second problem is more troublesome. The approximately 10^{10} multiples of 2π in the argument have no effect on the result except to reduce its accuracy by about 10 digits. Since the $\omega_j t_i$ product must be formed (and so truncated to a finite number of digits by the computer) before the multiples of 2π can be removed, the important digits are lost and cannot be reclaimed. While this error cannot be eliminated, it can be controlled by exploiting the way the frequency set Λ_K is constructed. From (2), the product $\omega_j t_i$ can be written

$$\omega_j t_i = k(\lambda_1 t_i) + l(\lambda_2 t_i).$$

Let $$\psi_{im} = \mathbf{fract}\left[\frac{\lambda_m t_i}{2\pi}\right] \qquad m = 1, 2; \quad 1 \le i \le N \tag{5}$$

and $$\phi_{im} = 2\pi(k\psi_{i1} + l\psi_{i2}). \tag{6}$$

Now $\phi_{im} = \omega_m t_i - 2\pi n$, where n is some integer and $|\phi_{im}| \le 2\pi(|k|+|l|)$. Since k and l are small integers, ϕ_{im} is an appropriate argument to trigonometry routines on all computers. The dominant source of roundoff error results from multiplying $\lambda_m/2\pi$ by t_i, because the product is formed before the **fract**() operator (which removes any integer portion and leaves only the fractional part) is applied to remove the troublesome multiples of 2π. By using (5) and (6), the roundoff error can be viewed as resulting from roundoff error in λ_1, λ_2 and t_i. Since the t_i are chosen randomly, their roundoff errors are of no concern. The amount of roundoff in λ_1 and λ_2 is predictable and so a warning can be given if is becomes significant in comparison to ω_{min}.

3.3: APFT Algorithm Results

The APFT orthogonal selection algorithm requires on the order of M^2N operations, where M is the number of time point candidates used and N is the number of Fourier coefficients. Since we have used $M = 2N$, the asymptotic complexity of the algorithm is the same as that of the matrix inversion needed to compute Γ.

We note that while the initialization of the APFT (that is the time point selection, formation of Γ^{-1} and the inversion of Γ^{-1} to find Γ) requires on the order of N^3 operations, the actual forward and inverse transform requires N^2 operations, the same as the DFT. Thus the expensive part of the APFT is performed only once per set of frequencies; after this initial overhead has been paid, the APFT is as efficient as the DFT.

To show the numerical stability of our method, we will compare the condition number of Γ^{-1} when time points are 1) evenly spaced, 2) randomly spaced, and 3) determined by the orthogonal selection algorithm. The condition number, κ, is proportional to the errors in computing the inverse. In our experiments, $\varepsilon \approx 10^{-16}\kappa$. Bear in mind that even the DFT, which is theoretically the best conditioned algorithm for the simpler periodic case, has a condition number $\kappa \approx N$, so the best we could hope for is linear growth of the condition number with the number Fourier coefficients. Observe that as shown by the results given in Figure 1, the condition number from orthogonal selection is in fact experimentally observed to grow linearly with K. That of random selection appears to grow quadratically, and that of evenly spaced to grow exponentially.

The case chosen for comparison was with fundamentals $\lambda_1 = 2\pi 10^9$ and $\lambda_2 = 2\pi(10^9 + \sqrt{2})$. Thus, the fundamentals differ by only 1 part in 10^9; also, because the fundamentals are incommensurable, the signal is not periodic. Comparisons of the condition numbers are shown in Figure 1 with the order H varying between 1 and 10. To smooth the wide variation seen in the results for the case of randomly selected time points, each condition number plotted is the geometric mean of 10 trials. Similarly, because different intervals give widely varying results for evenly-spaced points, those condition numbers are geometrically averaged over 10 intervals ranging from 1.5 to 4.5 times $2\pi/\omega_{min}$. Results obtained from orthogonal selection are so consistent that no averaging was needed, as evidenced by the smoothness of that curve. Graphing the condition number clearly shows that both randomly chosen and equally spaced samples will have accuracy problems when the number of frequencies is large. Orthogonal selection from $2N$ randomly selected time points, the algorithm we have described, always results in a reasonable condition number. The table below gives a summary of information on the APFT with the orthogonal selection algorithm. Execution times were measured on a VAX 8650 running ULTRIX 2.0.

Recall that coefficients of frequencies not in Λ_K can be amplified by up to κ/K in the process of aliasing. For order $H = 10$, this amplification factor equals approximately 10^8 for evenly spaced points, 2000 for randomly spaced points, and 10 for points chosen using orthogonal selection. Thus, even if the coefficients of neglected frequencies are small, for evenly and randomly spaced points the error δX due to aliasing may be so large as to dominate over the desired coefficients X.

There are a couple of details worthy of note. A temptation might be to compute several matrices Γ^{-1} by random time-point selection, and then use the selection that gave the best condition number. This works surprisingly well, but (for, say, 10 trials) it is still 10 to 100 times worse

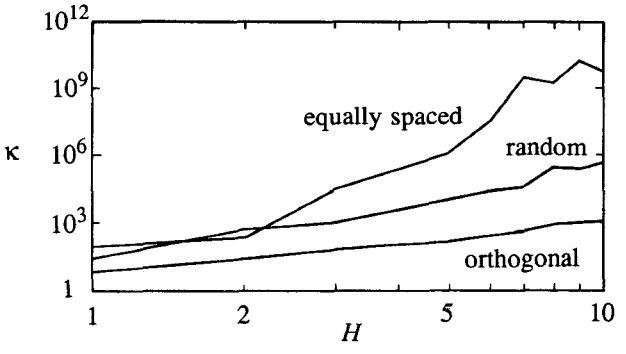

Figure 1 : Condition number versus order H for time points chosen evenly spaced, randomly, and using the orthogonal selection algorithm.

APFT Summary						
H	K	N	κ	ε	t_{init}	$t_{transform}$
1	3	5	6	2.8×10^{-17}	17 ms	0
2	7	13	24	8.3×10^{-17}	67 ms	0.3 ms
3	13	25	64	1.1×10^{-16}	280 ms	1.7 ms
4	21	41	113	1.6×10^{-16}	1.1 s	3.6 ms
5	31	61	143	1.1×10^{-16}	3.3 s	8.5 ms
6	43	85	270	2.3×10^{-16}	8.6 s	17 ms
7	57	113	420	2.9×10^{-16}	20 s	30 ms
8	73	145	790	3.3×10^{-16}	41 s	49 ms
9	91	181	950	4.8×10^{-16}	79 s	77 ms
10	111	221	1200	4.6×10^{-16}	142 s	116 ms

than orthogonal selection, with no justifying increase in efficiency. Also, the difference between the algorithms' results again increases rapidly with the order of harmonics. Unsurprisingly, the spread of the condition numbers obtained from orthogonal selection applied to different sets of random times is much smaller than that of random selection; thus its results are not only better but also more predictable.

Finally, the function taking a time point to a row vector of Γ^{-1} is complex; for example, it is unclear how to invert it or how to characterize its range. Modeling it as random, that is as producing vectors uniformly distributed in an N-ball, gives interesting theoretical results, including expectations of the condition numbers obtained from randomly selected time points and from time points chosen by the orthogonal selection algorithm. We are currently working to develop this theoretical model, and we hope to be able to report on it in the near future.

4: HARMONIC BALANCE

Kundert and Sangiovanni [kundert86] show that if a circuit consists of nonlinear resistors and capacitors as well as independent current sources and arbitrary linear devices, and if nodal analysis is used to formulate the circuit equations, then the time domain circuit equation

$$f(v,t) = i(v(t)) + \dot{q}(v(t)) + \int_{-\infty}^{t} y(t-\tau)v(t)dt + u(t) = 0 \qquad (7)$$

can be written in the frequency domain as

$$F(V) = I(V) + \mathbf{j}\Omega Q(V) + YV + U = 0 \qquad (8)$$

where $\mathbf{F}v=V$ represents the node voltages, $\mathbf{F}i(v)=I(V)$ represents the current through the nonlinear resistors, $\mathbf{F}q(v)=Q(V)$ represents the charge through the nonlinear capacitors, y is the impulse response of the linear portion of the circuit and Y is its phasor equivalent, $\mathbf{F}u=U$ represents the current from the input sources, and $\mathbf{F}f(v)=F(V)$ represents the sum of the current entering the nodes. For (8) to be valid, it is necessary to assume that u is periodic, that v is a locally asymptotically stable periodic solution, and that y is causal and asymptotically stable. Assume for simplicity that (7) describes a one node circuit, and so is a scalar equation.

To evaluate the nonlinear devices in (8) it is necessary to convert the node voltage spectrum V into the waveform v and evaluate the nonlinear devices in the time domain. The response is then converted back into the frequency domain. Thus, $I(V)=\mathbf{F}i(\mathbf{F}^{-1}V)$, and similarly $Q(V)=\mathbf{F}q(\mathbf{F}^{-1}V)$. Now that we have developed the APFT, it can be used with (8) to allow harmonic balance to be applied to almost-periodic systems. Assume that $v,u\in AP(\Lambda_K)$ and that a set of time points $\{t_0, t_1, \cdots, t_{2K-1}\}$ has been chosen so that Γ^{-1} is nonsingular. Then $I(V)=\Gamma i(\Gamma^{-1}V)$ and $Q(V)=\Gamma q(\Gamma^{-1}V)$.

Applying Newton-Raphson to solve (8) results in the iteration

$$J(V^{(j)})(V^{(j+1)} - V^{(j)}) = -F(V^{(j)}) \qquad (9)$$

where

$$J(V) = \frac{\partial F}{\partial V} = \frac{\partial I}{\partial V} + \mathbf{j}\Omega\frac{\partial Q}{\partial V} + Y.$$

The derivation of $\partial I/\partial V$ follows with help from the chain rule.

$$I(V) = \Gamma i(\Gamma^{-1}V)$$

$$\partial I(V) = \Gamma\frac{\partial i(\Gamma^{-1}V)}{\partial v}\Gamma^{-1}\partial V$$

$$\frac{\partial I(V)}{\partial V} = \Gamma\frac{\partial i(\Gamma^{-1}V)}{\partial v}\Gamma^{-1}$$

The derivation of $\partial Q/\partial V$ is identical. Now everything needed to evaluate (9) is available. If the sequence generated by (9) converges, its limit point is the desired solution to (8).

4.1: Harmonic Balance Results

The APFT algorithm has been integrated into *Harmonica*, our harmonic balance based circuit simulator. *Harmonica* was then used to simulate a GaAs FET [statz87] double-balanced mixer with a 50 mV 5 GHz RF input signal and a 500 mV 5.001 GHz LO input signal. The output is at 1 MHz and passes through a high Q 1 MHz bandpass lattice filter. The circuit consists of 6 GaAs FETs and 27 nodes and was simulated with order $H = 5$, which corresponds to 31 frequencies. *Harmonica* required 4.5 megabytes of physical memory and 470 seconds on a VAX 8650 to complete the simulation, though this should improve as the implementation is polished. The circuit, with the center frequency of the output filter adjusted accordingly, was also simulated with the LO frequency set as close as 1Hz away from the 5 GHz RF with no apparent change in accuracy. Note that the combination of the widely separated frequencies and the high Q output filter make it prohibitively expensive to find the steady state response of this circuit with a time domain simulator.

5: CONCLUSION

A new almost-periodic Fourier transform was presented that is both efficient and accurate. This transform was combined with harmonic balance to allow circuits with widely separated frequencies to be accurately simulated. Work is continuing on the APFT and its application in harmonic balance to further increase the efficiency the algorithm and to explore its error mechanisms.

REFERENCES

[bohr47] Harald Bohr, *Almost Periodic Functions*, Chelsea, 1947.

[dahlquist74] Germund Dahlquist and Åke Björck, *Numerical Methods*, Prentice-Hall, 1974.

[gilmore86] Rowan Gilmore, "Nonlinear circuit design using the modified harmonic balance algorithm," *IEEE Transactions on Microwave Theory and Techniques*, vol. MTT-34, no. 12, December 1986, pp. 1294-1307.

[golub83] Gene H. Golub and Charles F. Van Loan, *Matrix Computations*, The Johns Hopkins University Press, 1983.

[hale80] Jack K. Hale, *Ordinary Differential Equations*, Krieger, 1980.

[kundert86] Kenneth S. Kundert and Alberto Sangiovanni-Vincentelli, "Simulation of nonlinear circuits in the frequency domain," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. CAD-5, October 1986, pp. 521-535.

[statz87] Hermann Statz, Paul Newman, Irl W. Smith, Robert A. Pucel and Hermann A. Haus, "GaAs FET device and circuit simulation in SPICE," *IEEE Transactions on Electron Devices*, vol. ED-34, no. 2, 1987, pp. 160-169.

[ushida84] A. Ushida and L.O. Chua, "Frequency-domain analysis of nonlinear circuits driven by multi-tone signals," *IEEE Transactions on Circuits and Systems*, vol. CAS-31, no. 9, September 1984, pp. 766-778.