Efficient AC and Noise Analysis of Two-Tone RF Circuits

Ricardo Telichevesky and Ken Kundert
Cadence Design Systems, San Jose California

Jacob White
Massachusetts Institute of Technology, Cambridge, Massachusetts

Abstract— **In this paper we present a preconditioned recycled Krylov-subspace method to accelerate a recently developed approach for ac and noise analysis of linear periodically-varying communication circuits. Examples are given to show that the combined method can be used to analyze switching filter frequency response, mixer 1/f noise frequency translation, and amplifier intermodulation distortion. In addition, it is shown that for large circuits the preconditioned recycled Krylov-subspace method is up to forty times faster than the standard optimized direct methods.**

I. Introduction

The intensifying demand for very high performance portable communication systems has greatly expanded the need for simulation algorithms that can be used to efficiently and accurately analyze frequency response and noise of RF communication circuits such as mixers, switched-capacitor filters and narrow-band amplifiers. Such simulation problems seem deceptively simple because the quantities of interest, noise spectral density and small-signal frequency response, are of sufficiently small amplitude that they can be computed using linearization. However, communication circuits like mixers are usually driven by an additional large amplitude periodically time-varying source, and this implies that the linearization used for analysis must be time-varying.

Consider a specific example, a mixer being used to perform frequency down-conversion. For this example, one of the input signals will be a large signal local oscillator running at a fixed frequency, and the other input could be an RF carrier. Even if the RF carrier is a very small amplitude sine wave, the mixer will produce outputs at frequency-translated harmonics of the large signal local oscillator. It is only by including the time-variation of the linearized system that one could derive a frequency translation result.

Algorithms have recently been developed for the type of linear time-varying analysis needed to determine noise and frequency response in periodically driven RF circuits [5], [7], [8]. Unfortunately, the computational cost of the direct factorization techniques typically used in these algorithms grows so fast with circuit size that the methods, although effective, have been restricted to problems with only a few hundred nodes. In particular, the direct approach involves factoring a dense $N \times N$ system that is also expensive to construct.

In this paper, we describe a preconditioned recycled Krylov subspace method with a semi-implicit matrix representation to solve the system of equations generated by linear time-varying analysis algorithms. Krylov-subspace methods have previously been used to accelerate the solution of systems of equations generated by problems associated with steady-state analysis. Preconditioned iterative techniques have been applied to harmonic balance methods [2], [6] and to the problem of periodic steady-state analysis [12] but not, to our knowledge, to the problem of periodic time-varying linear analysis. In the next section we describe the periodic time-varying linear analysis algorithm in more detail, and in Section III. we present how to precondition the problem so that the Krylov subspace method converges rapidly. Also in Section III. we present a new Krylov subspace recycling algorithm which makes simulation of frequency sweeps particularly efficient. The recycling method is unlike standard multiple right-hand side methods because it adapts to a changing system matrix. In Section IV. we present results on several examples to demonstrate that the recycled preconditioned Krylov-subspace method is faster than direct techniques for large problems, and is nearly forty times faster on a thousand node industry-supplied RF mixer example. Finally, conclusions are given in Section V.

II. Background

In this section we briefly review the system of equations generated by the linear time-varying analysis approach in [7], though we refer the reader to [8] for the adjoint formulation used for noise analysis. We then show that the computational cost of solving the generated system with standard, but carefully applied, direct techniques grows

33rd Design Automation Conference ®

like N^3, limiting the method to relatively small systems. In the next section, we will describe our new recycled Krylov-subspace approach to solving these equations that grows nearly linearly with problem size.

A. The Linear Periodically-Varying Approach

Consider a circuit whose input is the sum of two periodic signals, $u_L(t) + u_s(t)$, where $u_L(t)$ is an arbitrary periodic waveform with period T_L and $u_s(t)$ is a sinusoidal waveform of frequency f_s whose amplitude is small. Using modified nodal analysis [3], the differential equations for the circuit can be written in the form

$$f(v(t), t) = i(v(t)) + \dot{q}(v(t)) + u_L(t) + u_s(t) = 0, \quad (1)$$

where $u_L(t), u_s(t) \in \mathcal{R}^N$ are the vector of large and small signal input sources, $v(t) \in \mathcal{R}^N$ is the vector of node voltages, and $i(v(t)), q(v(t)) \in \mathcal{R}^N$ are the vectors of resistive node currents and node charges or fluxes respectively.

If $u_s(t)$ is assumed to be zero, then the periodic steady solution, $v_L(t)$ is the solution of (1) which also satisfies the two-point constraint

$$v(T_L) - v(0) = 0. \quad (2)$$

Now assume $u_s(t)$ is not zero, but is small. We can consider the new solution to be a perturbation on $v_L(t)$, as in

$$\frac{d}{dt} q(v_L(t)+v_s(t)) + i(v_L(t)+v_s(t)) + u_L(t) + u_s(t) = 0, \quad (3)$$

where $v_s(t)$ is the difference between the exact solution to (1) and the solution computed assuming $u_s(t) = 0$.

Linearizing about $v_L(t)$, which is accurate only if $u_s(t)$ is small, yields a time-varying linear system of the form

$$\frac{d}{dt}\left(\frac{dq(v_L(t))}{dv_L} v_s(t)\right) + \frac{di(v_L(t))}{dv_L} v_s(t) + u_s(t) = 0. \quad (4)$$

for $v_s(t)$. Here, $v_s(t)$ can be interpreted as the small signal response to $u_s(t)$. In addition, from the theory of periodically time-varying systems, it is known that if $u_s(t) = U e^{j2\pi f_s t}$, then the steady-state response is given by

$$v_s(t) = \sum_{k=-\infty}^{\infty} V_k \cdot e^{j2\pi(f_s + k/T_L)t} \quad (5)$$

from which it follows that

$$v_s(t + T_L) = v_s(t)\, e^{j2\pi f_s T_L}. \quad (6)$$

Equation (6) in the periodically time-varying linear steady-state problem is analogous to (2) in the standard steady-state problem. Note that (6) also implies that the entire small-signal steady-state response of the periodic time-varying system is determined by the behavior of $v_s(t)$ on any interval of length T_L.

Now consider solving (4) using the backward-Euler discretization,

$$\left(\frac{1}{h_j}\frac{dq(v_L(t_j))}{dv_L} + \frac{di(v_L(t_j))}{dv_L}\right)\tilde{v}_s(t_j) - \frac{1}{h_j}\frac{dq(v_L(t_{j-1}))}{dv_L}\tilde{v}_s(t_{j-1}) + u_s(t_j) = 0. \quad (7)$$

where $\tilde{v}_s(t_j)$ is the approximation to $v_s(t)$ at the j^{th} backward-Euler timestep, and $h_j = t_j - t_{j-1}$. Note, (7) represents a sequence of M linear systems, where M is the number of timesteps.

Combining (6) with (7), and assuming $t_M = T_L$, yields complex linear system of $M \times N$ equations for the M N-length vectors of v_s at each of M timesteps. In particular,

$$\begin{bmatrix} \frac{C_1}{h_1} + G_1 & & & -\frac{C_M}{h_1}\cdot\alpha(f_s) \\ -\frac{C_1}{h_2} & \frac{C_2}{h_2} + G_2 & & \\ & \ddots & \ddots & \\ & & -\frac{C_{M-1}}{h_M} & \frac{C_M}{h_M} + G_M \end{bmatrix} \begin{bmatrix} \tilde{v}_s(t_1) \\ \tilde{v}_s(t_2) \\ \vdots \\ \tilde{v}_s(t_M) \end{bmatrix} = \begin{bmatrix} -u_s(t_1) \\ -u_s(t_2) \\ \vdots \\ -u_s(t_M) \end{bmatrix} \quad (8)$$

where $\alpha(f_s) \equiv e^{-j2\pi f_s T_L}$, C_j denotes $dq(v_L(t_j))/dv_L$, and G_j denotes $di(v_L(t_j))/dv_L$.

B. Standard Direct Solution Method

Although (8) can be solved using sparse matrix techniques, a more efficient approach is typically applied, one that exploits the fact that the matrix is mostly block lower triangular and is typically solved for a sweep of frequencies. To describe this standard approach, we consider solving (8) and let L be the lower triangular portion of the matrix in (8),

$$L \equiv \begin{bmatrix} \frac{C_1}{h_1} + G_1 & & & \\ -\frac{C_1}{h_2} & \frac{C_2}{h_2} + G_2 & & \\ & \ddots & \ddots & \\ & & -\frac{C_{M-1}}{h_M} & \frac{C_M}{h_M} + G_M \end{bmatrix} \quad (9)$$

and define B as

$$B \equiv \begin{bmatrix} 0 & \dots & 0 & -\frac{C_M}{h_1} \\ 0 & \ddots & 0 & 0 \\ \vdots & & \ddots & \vdots \\ 0 & \dots & 0 & 0 \end{bmatrix}. \quad (10)$$

Using the above notation, (8) becomes

$$(L + \alpha(f_s)B)\ \tilde{v}_s = -u_s(f_s). \quad (11)$$

As L is lower triangular, its inverse is easily applied by factoring the diagonal blocks and back-solving. Formally, the result can be written as

$$\left(I + \alpha(f_s)L^{-1}B\right)\ \tilde{v}_s = -L^{-1}u_s(f_s), \quad (12)$$

though L^{-1} need not be explicitly computed.

Examining (12) reveals another important feature, the MN by MN matrix $L^{-1}B$ has nonzero entries *only in the last N columns*. Therefore, the matrix $(I + \alpha(f_s)L^{-1}B)$ has the form

$$\begin{bmatrix} I & \dots & 0 & \alpha(f_s)P_1 \\ 0 & I & 0 & \alpha(f_s)P_2 \\ \vdots & & \ddots & \vdots \\ 0 & \dots & 0 & I + \alpha(f_s)P_M \end{bmatrix}, \tag{13}$$

where $P_i \in \mathcal{R}^{N \times N}$ is the $((i-1)N)+1$ through iN rows of the last N columns of $L^{-1}B$.

Note, the structure of (13) is such that $\tilde{v}_s$ can be computed in stages. First, P_M is computed by forming the N products Be_i, where the e_i's are the first N unit vectors in $\mathcal{R}^{N \times M}$, followed by backsolving N times with L. Then, for each frequency point one must: form the right-hand side by backsolving with L, compute $\tilde{v}_s(t_M)$ by factoring $I + \alpha(f_s)P_M$, and finally compute the rest of $\tilde{v}_s$ from (8) using the known value of $\tilde{v}_s(t_M)$ and backsolving again with L. The total cost of the calculation is

$$(N + 2FP)(\text{backsolve}) + FP(\text{dense factorization}) \tag{14}$$

where FP is the number of frequency points. Since each backsolve with L costs *at least* $O(MN)$, and so the time to compute the N backsolves required to construct P_M is *at least* $O(MN^2)$ and the cost of factoring the dense P_M matrix is $O(N^3)$. This rapid growth of computation time with problem size will severely limit the size problem which can be attacked with linear time varying analysis. In the next section, we will describe a more efficient approach to solving (8) based on using recycled Krylov subspace based iterative methods.

III. Fast Frequency Sweeps using the Recycled Krylov Subspace Method

Krylov-subspace methods have proven to be effective for periodic steady-state problems [12], and in this section we show that the method is efficient for two-tone linear time-varying analysis as well. We start by describing how to apply the Krylov subspace methods to solving (8) by reinterpreting the use of L^{-1} in (12) as a preconditioner which also reduces the size of the system which must be solved iteratively. We then present a recycling optimization which improves the method's efficiency when performing frequency sweeps associated with noise and AC analysis.

A. Applying Preconditioned Krylov Subspace Methods

As described in the previous section, solving (8) with direct methods grows like N^3, and this can severely limit the size problems which can reasonably be solved. Instead, consider solving (8) with a Krylov-subspace based iterative algorithm like GCR [1], [10], given in very general form in Algorithm I.

Algorithm I
(*Krylov-subspace algorithm for solving* $Ax = b$.)

Guess at a solution, x^0.
Initialize the search direction $p^0 = b - Ax^0$.
Set $k = 1$.
do {
  Select $p^k \in \{p^0, Ap^0, A^2p^0, ..., A^{k-1}p^0\}$ in
    $x^k = x^{k-1} + p^k$
   to minimize $||r^k|| = ||b - Ax^k||$.

  If $||r^k|| <$ *tolerance*, return x^k as the solution.
  else Set $k = k + 1$.
}

We have taken a particularly stylized approach to describing the Krylov-subspace method in Algorithm I in order to focus on matrix-vector products, whose costs dominate in this particular problem. However, any practical implementation would not compute the Krylov subspace by directly computing $A^k p_0$, the vectors would too quickly line up with the dominant eigenspace of A. Instead, orthogonalization procedures are introduced both to insure that the Krylov subspace is accurately spanned and that the minimization is easy to perform.

Before applying a Krylov-subspace method to solving (11), consider preconditioning the problem by multiplying through by L^{-1} as in

$$(I + \alpha(f_s)L^{-1}B)\tilde{v_s} = -L^{-1}u_s(f_s). \tag{15}$$

Of course, one would not actually invert L, but simply factor it, so that the multiplication by the inverse could be accomplished by backward substitution. And note, being block lower bidiagonal implies L can be factored by just factoring the diagonal blocks. In our case, the diagonal blocks are just the $\frac{1}{h}C + G$ matrices which would be factored in a standard transient analysis.

Examining (15) reveals another important feature of the preconditioned problem. As noted in the previous section, the MN by MN matrix $H = L^{-1}B$ has nonzero entries only in the last N columns. This implies that (15) can be solved by solving for the last N entries in $\tilde{v_s}$. This last observation reduces the cost of applying the Krylov subspace method.

Finally, there is the question of how many iterations the Krylov-subspace method will require to achieve convergence for the preconditioned problem. Although Krylov-subspace methods have guaranteed convergence properties, they can be so slow as to be ineffective. Both practical results given below, as well as some limited theoretical results [11], indicate that the above preconditioned

Krylov-subspace method converges rapidly for periodic steady-state problems. However, the subject deserves a more significant investigation.

B. Recycling the Krylov Subspace

A designer is typically interested in sweeping the frequency of $u_s(t)$ and then examine the steady-state responses from the periodic time-varying linear system. Frequency sweeping is also used to determine the noise spectral density. One could then ask whether the Krylov-subspace method used to solve (15) at one frequency can be of any help in solving it at the next frequency. In general, in a problem where the matrix and the right-hand are functions of the swept parameter, the answer is no. Exactly how the matrix changes in this problem is constrained, however, and so previous iterative solutions can be exploited. To see this, we will make use of the following theorem:

Theorem 1 *The space spanned by the vectors*

$$\{p^0, (I+\alpha(f_s)H)p^0, (I+\alpha(f_s)H)^2p^0, ..., (I+\alpha(f_s)H)^{k-1}p^0\} \tag{16}$$

is identical to the space spanned by the vectors

$$\{p^0, Hp^0, (H)^2p^0, ..., (H)^{k-1}p^0\} \tag{17}$$

independent of α.

To prove Theorem (1), one need only note that adding the identity to H is inconsequential because p^0 is already in the space, and multiplying H by a scalar does not change the resulting vector directions. ■

In order to describe how to make use of the result in Theorem (1), we give a modified generalized conjugate-residual algorithm for using the recycled Krylov-subspace vectors. First note that

$$\beta(I+\alpha(f_s)H)p^0 + \gamma p^0 = (I+\alpha(\tilde{f}_s)H)p^0 \tag{18}$$

where $\beta = \alpha(\tilde{f}_s)/\alpha(f_s)$ and $\gamma = 1 - \beta$. This implies that a matrix-vector product computed using the matrix associated with frequency f_s can be converted into a matrix-vector product using the matrix associated with frequency $\tilde{f}_s$ by a simple scalar multiplication.

The GCR algorithm with recycling seems like it could generate an enormous number of vectors as the frequency is swept, but the invariance of the span of the Krylov subspace with respect to α insures that the recycling index never gets very large. In fact one has a bound on the number of vectors, or the recycling index, which is independent of the number of frequency points computed.

Theorem 2 *In exact arithmetic, the number of vectors generated by the GCR algorithm with recycling is bounded by N, where N is the dimension of H.*

Algorithm II
(Recycling Generalized Conjugate Residual Algorithm for Solving $(I + \alpha_l H)x = b_l$.)

```
RecycleIndex = 0.
Repeat for l = 1 to L /* This sweeps the frequency. */
  r^0 = b_l. /* Initialize the residual */
  k = 1. /* Initialize the iteration index. */
  Repeat
    If k > RecycleIndex
      p^k = r^k. /* Initialize the search direction. */
      ap^k = (I + α_l H)p^k. /* Matrix-vector Product */
      RecycleIndex = k. /* Adds to the subspace. */
    else
      ap^k = β ap^k + γ p^k. /* Use the recycled vector. */
    end
    for j = 1 to k − 1
      ap^k = ap^k − (ap^j)^H ap^k ap^j. /* Orthogonalize. */
      p^k = p^k − (ap^j)^H ap^k p^j.
    end
    r^{k+1} = r^k − (ap^k)^H r^k ap^k. /* Update r. */
    x_l^{k+1} = x_l^k + (ap^k)^H r^k p^k. /* Update x. */
    k = k + 1.
  Until ||r^k|| < tolerance.
end
```

The theorem follows directly from Theorem (1) and the fact that there are no more than N orthogonal N-length vectors. ■

It should be noted that if the GCR algorithm with recycling requires N vectors, then its cost is similar to the direct factorization approach given in the previous section. However, in practice, many fewer than N are needed, unless N is very small.

IV. Results

A. High-Performance Receiver

A high-performance image rejection receiver was simulated with the periodic time-varying linear analysis. The receiver consists of a low-noise amplifier, a splitting network, two double-balanced mixers, and two broad-band Hilbert transform output filters combined with a summing network that is used to suppress the undesired sideband. A limiter in the LO path is used for controlling the amplitude of the LO. It is a rather large RF circuit that contains 167 bipolar transistors and uses 378 nodes. It is simulated at the circuit level using the Gummel-Poon model for each transistor. It generated 986 equations (378 nodes, 102 inductors, and 495 additional nodes internal to BJTs to support their parasitic resistors).

A 780MHz LO was applied and a periodic steady-state analysis [12] was performed in order to build the linear pe-

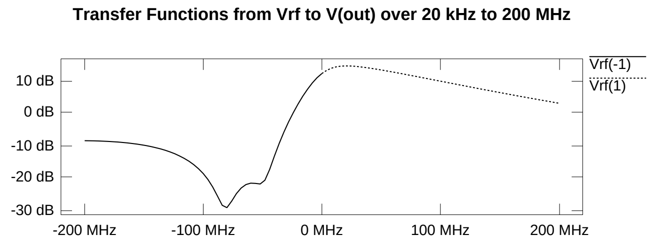

Fig. 1. Conversion gain of sideband suppression mixer. Output frequencies from -200MHz to $+200$MHz correspond to input frequencies of 580MHz to 980MHz. The results both with and without search direction reuse are plotted. They fall right on top of each other indicating that the results are virtually identical.

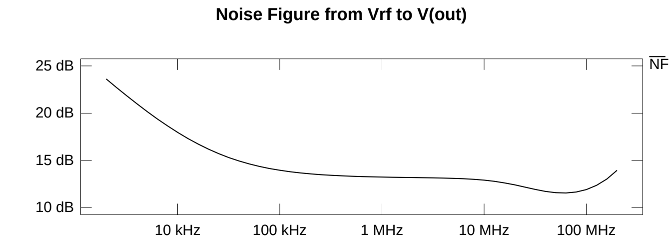

Fig. 2. The noise figure from Vrf, assuming 50Ω source resistance. The flicker or $1/f$ noise of the devices is clearly visible.

riodically time-varying representation of the circuit. Periodic transfer function analysis was applied to compute the conversion gain of the receiver. The input frequency was set to $780\text{MHz} + f_{\text{off}}$ where f_{off} is swept from -200MHz to $+200$MHz using 51 equally spaced points. The output is translated down in frequency by the 780MHz LO, meaning the output sweeps from -200MHz to $+200$MHz. The conversion gain is shown in Figure 1 while the timing results are given in Table I. Notice that the gain for negative output frequencies (input frequencies below the 780MHz LO) is considerably less that for positive output frequencies (input frequencies greater than the LO) because the receiver is designed to suppress the negative sideband. The null near -80MHz corresponds to frequencies where the two signals coming out of the Hilbert transform filters cancel each other.

The noise figure, as computed by periodic noise analysis, is shown in Figure 2. This analysis is similar to the traditional noise analysis, except it includes the effect of noise moving from one frequency to another as it mixes with the LO and its harmonics. In Figure 2, the noise figure at each frequency point f_{noise} includes contributions from noise at frequencies $f_{noise} \pm kf_{LO}$, where $|k| < 7$. Since this noise analysis is performed on the detailed periodically varying linearized circuit, it includes a great deal of subtle effects in the calculation. For example, at the points in time where the LO signal is near 0V, the differential pairs that make up the mixer are balanced and so exhibit a great deal of gain from the LO to the output. Any noise in the LO is greatly magnified at this time. Once LO moves away from 0V, the differential pairs saturate and exhibit a great deal of attenuation from the LO to the output. Thus, the faster the LO moves through zero, the less noise is transferred from LO to the output. This is why switching mixers are generally preferred over multiplying mixers. This effect is one of many that are accurately accounted for by the periodic noise analysis.

B. Transmitter

A transmitter chip that contains an up-convertion mixer along with associated circuitry was simulated to compute transfer and noise characteristics. The circuit contained 800 bipolar transistors and thousands of resistors and capacitors that were generated during parasitic RC extraction. The circuit generated 6200 equations, which is so large that it could not be analyzed using direct methods. The transmitter modulated a 1.5GHz carrier with voice-band audio. Periodic small-signal analysis was used to compute the conversion gain, LO feed-through, and noise performance. The noise performance was of particular interest. In transmitters, flicker noise in the transistors is translated up to the carrier frequency by the mixer. Flicker noise near the carrier dominates over other noise sources and can only be predicted if frequency translation effects are included in the noise analysis. A periodic small-signal version of noise analysis does account for frequency translation due to the circuit as well as the effect of periodic modulation of the bias currents in the devices. Computed result agreed to within 30% of measured results, which is considered quite good for this type of analysis.

C. Switched-Capacitor Synchronous Detector

One half of a switched-capacitor synchronous detector was simulated using periodic time-varying linear analysis to determine its frequency response, including parasitic passbands that result from undersampling effects [4]. The synchronous detector is used to determine the magnitude and phase of a 100kHz signal. The circuit is fully differential and consists of an interleaved mixer running at 800kHz and a 5-pole leapfrog 4kHz low-pass Bessel filter that has a 400kHz clock on the first stage, 200kHz clock on the second stage, and a 100kHz clock on the remaining three stages.

The 100kHz clock rate on the last three stages of the filter implies that there will be parasitic passbands at 100kHz intervals, however filtering in the earlier stages reduces the amplitude somewhat. Finally, use of bilinear integrators causes nulls to be placed at 100kHz intervals.

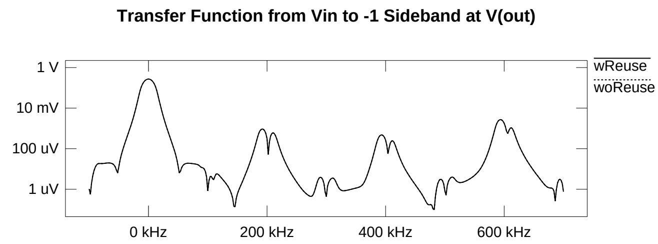

Fig. 3. Frequency response of switched-capacitor synchronous detector. Notice the parasitic passbands caused by undersampling and the nulls generated by the bilinear integrators. The input was swept from 0 to 800kHz. The clock acts to translate the input signals down in frequency by 100kHz, so the output sweeps from -100kHz to 700kHz. The results both with and without search direction reuse is plotted and are virtually identical.

Circuit		receiver	trans	SC Synch
Statistics	Eqns	986	6242	375
	TP	200	200	1039
	FP	51	51	401
Direct	TT	9962	—	6164
no Reuse	MVP	766	6164	2807
	TT	1015	208199	7036
with Reuse	MVP	41	116	18
	TT	231	7450	2212
Direct/Reuse		43.1	—	2.8
no Reuse/Reuse		4.4	27.9	3.2

TABLE I

Results with and without Krylov subspace reuse. *Eqns* gives the number of circuit equations, *TP* gives the number of time-points used in the periodic steady-state analysis, and *FP* gives the number of frequency points in the periodic linear time-varying analysis. Direct gives the estimated CPU time for the optimized direct factorization approach, *MVP* gives the cumulative number of matrix-vector products needed to solve the linear systems of equations. *TT* gives the total time needed for the analysis.

To analyze this circuit, a periodic steady state analysis was run first with a $10\mu s$ period. Then a 'small' sinusoidal input was applied and periodic small-signal analysis was run while sweeping the input signal from 0 to 800kHz. 200 frequency points were needed in order to cleanly resolve the passbands and the nulls. The results of the periodic small-signal analysis are given in Figure 3.

As shown in Table I, the recycled preconditioned Krylov method is much faster than the optimized direct factorization, particularly as the problem size increases. The 986-node receiver circuit runs nearly forty times faster with the new method.

V. Conclusion

Many important two-tone analyses for communication circuits, such as intermodulation distortion or mixer frequency translation, can be computed using the periodic time-varying linear approach described by Okumura et al [7]. In this paper, we examined using Krylov-subspace methods to accelerate the solution of the large linear system generated by the periodic time-varying linear approach when one tone can be assumed to be small. We presented a recycled preconditioned Krylov-subspace method, and showed that the new method is as much as forty times faster than direct factorization on large circuits.

Possible future enhancements to this method would be to replace the GCR algorithm with more robust Krylov-subspace methods such as GMRES [9], and developing more efficient approaches to computing the right-hand sides. Extensions to multiple tone excitations are also under investigation[1]

References

[1] H. C. Elman. *Iterative methods for large sparse nonsymmetric systems of linear equations.* Ph. D. thesis, Computer Science Dept., Yale Univ., New Haven, CT, 1982.

[2] Pauli Heikkilä. *Object-Oriented Approach to Numerical Circuit Analysis.* Ph. D. dissertation, Helsinki University of Technology, January 1992.

[3] Chung-Wen Ho, Albert E. Ruehli, Pierce A. Brennan. "The modified nodal approach to network analysis." *IEEE Transactions on Circuits and Systems*, vol. CAS-22, no. 6, pp. 504-509, June 1975.

[4] Kenneth S. Kundert, *A Switched-Capacitor Synchronous Detector,* Master's thesis, University of California at Berkeley, June 1983. Available through Electronic Research Laboratory Publications, U. C. B., 94720, Memorandum No. UCB/ERL M83/6.

[5] Kenneth S. Kundert, Jacob K. White and Alberto Sangiovanni-Vincentelli. *Steady-State Methods for Simulating Analog And Microwave Circuits.* Kluwer Academic Publishers, Boston 1990.

[6] R. Melville, P. Feldmann, and J. Roychowdhury. "Efficient multi-tone distortion analysis of analog integrated circuits." *Proceedings of the 1995 IEEE Custom Integrated Circuits Conference*, May 1995.

[7] M. Okumura, T. Sugawara, and H. Tanimoto. "An efficient small signal frequency analysis method for nonlinear circuits with two frequency excitations." *IEEE Transactions of Computer-Aided Design of Integrated Circuits and Systems*, vol. 9, no. 3, pp. 225–235, March 1990.

[8] M. Okumura, H. Tanimoto, T. Itakura, and T. Sugawara. "Numerical Noise Analysis for Nonlinear Circuits with a Periodic Large Signal Excitation Including Cyclostationary Noise Sources." *IEEE Transactions On Circuits and Systems - I Fundamental Theory and Applications.*, vol. 40, no. 9, pp. 581–590, September 1993.

[9] Y. Saad and M. H. Schultz. "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems." *SIAM Journal on Scientific and Statistical Computing*, vol. 7, pp. 856–869, July 1986.

[10] Yousef Saad. *Iterative Methods for Sparse Linear Systems,* PWS Publishing Company, 1995.

[11] Ricardo Telichevesky, Kenneth S. Kundert, Jacob K. White. Manuscript in progress.

[12] Ricardo Telichevesky, Kenneth S. Kundert, Jacob K. White. "Efficient Steady-State Analysis based on Matrix-Free Krylov-Subspace Methods." *Proceedings of the 1995 Design Automation Conference,* June 1995.

[1] Ingve Thodesen, private communication.