THE BEST OF ICCAD

HIGHLIGHTS IN PHYSICAL SIMULATION AND ANALYSIS AT ICCAD

Kenneth S. Kundert
Cadence Design Systems, Inc.
San Jose, CA

Jacob White
Massachusetts Institute of Technology
Cambridge, MA

Abstract
Six papers were chosen to represent twenty years of research in physical simulation and analysis, three papers addressing the problem of extracting and simulating interconnect effects and three papers describing techniques for simulating steady-state and noise behavior in RF circuits. In this commentary paper we will try to describe the contribution of each paper and place that contribution in some historical context.

1. Introduction

Research in computer-aided design (CAD) is, by its nature, application-driven. Algorithmic and methodological innovation has almost always been a response to the changing challenges faced by designers. It should therefore come as no surprise that many of the best papers presented at ICCAD over its twenty year history are associated with design issues that emerged over the same time frame. For example, most of the physical simulation and analysis papers nominated for this commemorative proceedings addressed either interconnect effects or radio-frequency (RF) design, two topics that have grown enormously in importance over the last twenty years.

2. Inductance Extraction

The first paper in our physical simulation and analysis commemorative collection is *Efficient Techniques for Inductance Extraction of Complex 3-D Geometries* by Matt Kamon, Michael Tsuk, Christopher Smithhisler, and Jacob White. It was published at ICCAD in 1992 [25] and can be found on page **??**. This was the first public description of the *FastHenry* 3-D inductance extractor and its algorithms.

The *FastHenry* project started as an obvious follow-on to the *FastCap* project [42, 43]. *FastCap* combined the fast-multipole algorithm with

iterative Krylov-based linear solvers to provide an algorithm whose complexity increased linearly with the size of the problem. *FastHenry* was expected to be a relatively small project that applied the multipole and iterative methods of *FastCap* to the PEEC-based inductance algorithms of Ruehli [61]. The project was led by Matt Kamon, a student at MIT, and his advisor, Prof. Jacob White. The initial challenge was to determine which approach should be used to formulate the magnetic equations. Ruehli had already demonstrated the use of nodal analysis in his PEEC methods, but Kamon and White instead decided on going with mesh analysis. It proved to be a much better choice for several reasons, but perhaps most important is that it was more compatible with iterative solvers. Next, they applied GMRES, a Krylov-based iterative method, which appeared easy to do in concept, but was actually quite difficult in practice. It is vastly more difficult to construct preconditioners for inductance than for capacitance because of the long range interactions that are inherent to inductance. The initial preconditioner was based on the assumption that the couplings between the mesh currents within an individual conductor are much tighter than those that occur between conductors. It was this, along with the mesh current formulation, that was initially implemented in *FastHenry* and documented at ICCAD in 1992 [25]. In subsequent work they moved on to adapting the multipole methods to the magnetic problem while continuing to improve the preconditioner [26, 27, 28, 29].

The *FastHenry* program would eventually make a large contribution to the research community, particularly in the area of interconnect analysis. While not completely general, *FastHenry* was able to handle a wide variety of 3-D structures. It was not difficult to write an interface to and had sufficient capacity to solve very large problems by the standards of field solvers of the day. For example, if you were willing to wait you could use it to extract the inductance of relatively complex IC package lead frames. *FastHenry* had most of the physical effects that were needed, such as skin and proximity effects and provided an out-of-the-box solution for researchers interested in inductance. It worked, was reliable, was reasonably fast, was freely available, and it had little to nothing in the way of competition. As such, it quickly became a critical tool for those exploring the effect of inductance on delays in interconnect. It was used by researchers to explore and understand inductance in complex geometries, and was a golden reference for anyone developing models of interconnect that include inductance.

The impact of *FastHenry* was different in character from its forerunner, *FastCap*. *FastCap* was the first solver whose computational complexity did not increase super-linearly with the size of the problem, and so was the first truly high-capacity solver. However, capacitance is fairly easy

to understand and it is easy to write solvers using a variety of methods that address the capacitance problem. As such, *FastCap* quickly had many competitors. The contribution of *FastCap* lay mainly in the ideas it embodied and the fact that it started a wave of innovation in solver technology. Inductance was much trickier, especially in 3-D. It was exceptionally difficult to write a solver that was both reliable and efficient, and there were few choices of algorithms. And so while the ideas and algorithms were less ground-breaking, the use of the program itself resulted in a deeper understanding of inductance by the user community and a wave of innovation in interconnect analysis. A remarkably large number of the papers on inductance that followed have used *FastHenry* as the benchmark reference [31, 32, 20, 64, 2].

3. Including Interconnect in Timing Analysis

Although inductive effects have been important in integrated circuit packaging for some time, it has only recently become a on on-chip concern for designers of high-performance digital circuits. However, interconnect resistance has been a design issue for more than two decades. Interconnect resistance was a problem in the early 1980's primarily because the then available fabrication technology provided only one layer of metal, making it necessary to use less conductive polysilicon for interconnect [40]. Technology improvements soon provided multiple layers of metal, thereby eliminating the need to use polysilicon for interconnect, but these same improvements also allowed for much higher circuit densities. To achieve these high circuit densities, it was necessary to reduce the cross-sectional area of metal interconnect, but since some signals still crossed the entire integrated circuit, maximum interconnect lengths did not scale like cross-section areas. The result was a resistance-preserving transition from fat interconnect in poor conducting polysilicon to skinny interconnect in good conducting aluminum to even skinnier interconnect in better conducting copper.

During the 1980's, one approach dominated for estimating the delay from the output of a driving logical gate, through the resistive interconnect, to the possibly many receiving gate inputs. The approach involved generating a model resistor-capacitor (RC) network, and then approximately analyzing the model [70]. In the model, the driving gate became a voltage source in series with a resistor, the interconnect became a collection of series resistors and grounded capacitors, and the receiving gates became grounded capacitors. The required delays were then estimated by approximating the step response of the model RC network, usually by exploiting the fact that the resistor network was almost always a tree and

then using a recursive algorithm to walk the tree and efficiently compute Elmore delays [60, 38, 17].

In present-day timing analysis programs, the techniques used to compute delays depend on the design methodology. In analyzing full-custom designs, delays are often estimated by numerically integrating the differential equations associated with the network formed from the driver and receiver transistors and the interconnect resistors and capacitors. For designs based on cell libraries, the numerical integration approach may not be the most efficient and can even be infeasible; the designer may not have access to transistor-level descriptions of the cells. Instead, delays in cell-based designs are analyzed using a strategy similar to the twenty-year old RC model strategy described above. The two main differences between the current strategy and that of twenty years ago are associated with determining the voltage source plus series resistor model for the driver gate, and with analyzing the resulting RC network.

The most widely used approach for generating voltage source plus series resistor models for cell library outputs is described in [8]. This approach presumes that the parameters of the simplified model must depend on the true interconnect load impedance and not just on total capacitance. This concept was first presented in the second paper in our physical simulation and analysis commemorative collection, *Modeling the Driving-Point Characteristics of Resistive Interconnect for Accurate Delay Estimation* by Peter R. O'Brien and Thomas L. Savarino, which can be found on page **??**. The main contribution of O'Brien and Savarino's paper was to establish that the behavior of a driving gate is substantially modified when much of the interconnect capacitance is "screened" from the gate output by the interconnect resistance.

O'Brien and Savarino's paper had an additional contribution that is more generic. In the years before their paper appeared, many researchers were using moment-matching methods, but most of those papers were matching only the first order moments to estimate delays [60, 38, 24]. O'Brien and Savarino's paper was one of the earliest to suggest that matching progressively higher order moments, which is equivalent to matching progressively higher terms in the Taylor series expansion of a transfer function, could be used to estimate input impedance. In addition, they suggested the idea that moment-matching could be used more generally to generate a reduced order model of the interconnect. In particular, they used moment-matching to determine values for the resistor and two grounded capacitors in a π circuit, and then suggested that the generated circuit was a reduced model of the interconnect. As will be discussed in the next section, general model order reduction has now become a central

part of strategies for handling interconnect, and O'Brien and Savarino's early observation about the subject was years ahead of its time.

4. Model Order Reduction for Interconnect

In order to examine signal propagation and coupling effects due to on- or off-chip interconnect, it is usually necessary to couple an electromagnetic analysis of the three-dimensional interconnect with a circuit-level analysis of the transistors connected to that interconnect. Although there are techniques and commercial tools that couple time-domain electromagnetic simulation programs directly to circuit simulators [67], these approaches are too computationally expensive to use in any but the simplest of scenarios. Instead, layout extraction tools are combined with model-order reduction techniques to generate low-order models for the interconnect, and then these low-order models are included with the transistors in circuit- or timing-level simulation and analysis.

When including interconnect effects in timing verification of an entire digital integrated circuit, the extraction-plus-reduction strategy must be fast enough to analyze millions of interconnect lines, but the accuracy requirements are modest. The commonly used approach is first to subdivide the interconnect into a large number of small sections, and then to apply a formula- or pattern-based algorithm to convert each of those small sections into resistors, capacitors and inductors. For complicated interconnect geometries, there can be a very large number of small sections, in which case the resulting extracted circuit will have a very large number of elements. Model reduction is used to generate low-order models of those large circuits, while still preserving the large circuit input-output behavior [37].

A very different extraction-plus-reduction strategy is used when examining interconnect coupling effects in packaging or for analog circuits. The objective of the coupling analysis is to determine if too much noise will be injected into a victim signal due to the proximity of the victim signal's interconnect to the interconnect of simultaneously active aggressor signals. In order to capture the "ganging-up" effect of many simultaneously active aggressors, it is necessary to extract interconnect coupling terms that would be too small to be of concern in timing analysis; but there is no need to extract more than a few hundred interconnect lines. The modest speed and high accuracy needed to investigate packaging and analog circuit coupling problems has led to extraction-plus-reduction strategies in which model reduction is embedded in a three-dimensional field solver, and the combination directly produces reduced models [63, 5, 50].

Although the approach to extraction is very different in digital circuit timing analysis versus packaging and analog circuit coupling examination, the approach to model reduction is very similar. In both cases the interconnect is viewed as a linear multi terminal device. The device terminals are usually transistor-interconnect interface locations, but may also simply be convenient separation points. Regardless of application, the goal of model reduction is to construct a representation of the interconnect that is inexpensive to evaluate, yet still accurately represents terminal behavior. Finally, the form of the reduced model should be appropriate for circuit or timing simulation. The reduced model could be another circuit, as in [46, 44], or a state-space model, but not tables of frequency-domain data.

The classic approach to model reduction is to select a fixed circuit topology, like the π model in [46], and then use some kind of fitting procedure to determine the element parameters. Such fitting techniques have uncertain accuracy, and there is no way to increase the that accuracy. In the early 1980's, researchers in system theory developed methods for reducing the order of state-space models that were, in a carefully chosen metric, provably optimal in preserving input-output behavior [22]. Although these optimal methods produced excellent reduced models, the computational cost of the numerical algorithms used in the reduction grew cubically with the number of states in the *unreduced* model, making the method much too slow to use on large interconnect problems.

In [52], an algorithm was presented for generating low-order Padé approximates [1] to circuit transfer functions. The algorithm was reasonably efficient even for large circuits, and it was expected that increased accuracy could be achieved by increasing the order of the Padé approximate. Also, since a q^{th} order Padé approximate matches $2q-1$ moments of the original transfer function, the method could be viewed as a generalization of the moment techniques used in interconnect delay estimation. Early implementations of these Padé-based methods were sometimes unreliable when applied to interconnect problems, and occasionally generated inaccurate and/or unstable representations of the interconnect. Since a q^{th} order Padé approximate matches $2q-1$ terms in the *zero-frequency centered* Taylor series expansion of the original frequency response, accuracy problems were mitigated by generalizing the approximate to match Taylor expansions at multiple center frequencies [7].

Using multiple center frequencies improved the robustness of the methods in [52, 3], but the fundamental difficulty was not made clear until the seminal paper by Feldmann and Freund [18] (though less well known, [21] appeared nearly simultaneously). In [18], it was shown that many of the accuracy issues associated with the Padé approximates stemmed from the

numerically unstable way in which those approximates were being computed. In addition, it was shown that the Padé approximate for a transfer function could be computed in a numerically stable manner by starting with a state-space description, and then constructing bi-orthogonalized Krylov subspaces associated with the system matrix, its transpose, the input vector and the output vector.

Once the connection was made between moment-matching and the Krylov subspaces, results using Krylov subspaces for model reduction appeared rapidly. Block full orthogonalization and bi-orthogonalization methods were developed to directly generate state-space representations of multiple-input multiple-output reduced models [63, 19], methods were generated with guaranteed stability properties [62], techniques were developed that allowed multiple expansion centers, and Krylov-subspace reduction was coupled to fast electromagnetic solvers [63, 50]. These algorithmic variants were then organized into a unified theory of Krylov subspace reduction methods, the projection framework, which only appeared in a unique Ph. D. thesis that is required reading in the field [23].

In the coupled circuit-interconnect simulation scenarios described above, it is certain that many reduced-order models will be combined to describe a larger system, and therefore the reduced-order models must be passive. Stability is insufficient because the combination of stable systems is not necessarily stable, but the combination of passive systems is passive. The issue of passivity in reduced-order models first appeared in a widely circulated but still unpublished manuscript [4]. In that manuscript, an algorithm was given for generating guaranteed passive reduced models of RC circuits, and that algorithm was connected to the notion of congruence transforms in [30]. The race was then on to find a guaranteed passive reduction strategy for RLC circuits, and the third paper in our physical simulation and analysis commemorative collection is the winner of that race *PRIMA: A Passive Reduced-order Interconnect Macromodeling Algorithm*, by Altan Odabasioglu, Mustafa Celik, and Lawrence Pileggi (see page **??**.

The PRIMA algorithm was more than just first, it had an elegant simplicity that stemmed from combining two key steps. First, the reduction was applied to a circuit equation formulation that generated two positive semi-definite matrices, one operating on the state vector and one operating on the state vector's derivative. Then, each of the positive semi-definite matrices were reduced with the same congruence transform. This two-step approach was influential because it became a strategy for generating passivity-preserving methods with other desirable properties, such as matching multiple moments [16] or having simple update formulas [37]. The two-step approach was also used to adapt the PRIMA algo-

rithm to other applications, such as extraction from 3-D electromagnetic analysis [39].

5. Harmonic Balance

The fourth paper in our commemorative collection is *Nonlinear Circuit Simulation in the Frequency Domain* by Kundert and Sangiovanni-Vincentelli [33] and can be found on page **??**. This paper and the follow-on journal paper [34] were not the first papers on harmonic balance. However, they were the first to attempt to apply harmonic balance to large scale circuits, and as such directly led to the introduction of the commercial RF simulators that are so heavily used today. Previous attempts focused on microwave circuits that contained a very small number of transistors, one or at most two, and a large number of passive components. This made sense for discrete microwave circuits, but was not appropriate for the monolithic microwave integrated circuits (MMICs) of the day, nor would it be appropriate for the coming radio frequency integrated circuits (RFICs). Engineers had been successfully designing discrete microwave circuits for years by iterating prototypes, and there was not a strong need for nonlinear circuit simulators from this community. However, for MMIC designers, iterating prototypes was both very expensive and time consuming, so they tended to use SPICE to verify their circuits before fabrication.

There were several problems with using SPICE for microwave circuits [35]. It is a time-domain simulator and so has difficulty including models of distributed components such as transmission line structures, particularly if they include loss or dispersion. It is a transient-based simulator, and so is inefficient when high frequency carrier signals are combined with low-frequency modulation signals. And, it is incapable of predicting the noise of circuits such as mixers and oscillators. *Nonlinear Circuit Simulation in the Frequency Domain* did not really address any of these issues directly, but rather it showed how harmonic balance, which was known to be an answer to the first problem, could be extended so that it had the capacity to be applied to integrated circuits. It was the ability to simulate integrated circuits that gave harmonic balance its commercial appeal. It was left to follow-on work to address the remaining issues.

The increase in capacity was achieved by reducing the computational cost of factoring the frequency-domain harmonic balance Jacobian by using a more easily factored near block diagonal matrix, one which ignored some of the coupling between frequencies. This represents the second generation harmonic balance. Many years later there was another substantial increase in the capacity of harmonic balance when it was combined with

fast Krylov subspace based methods [41]. These methods, which are now pervasive, also require an approximate Jacobian to act as a preconditioner. The approach pioneered by Kundert and Sangiovanni-Vincentelli continues to live on as one of a few commonly used preconditioners for this third generation of harmonic balance simulators.

Nonlinear Circuit Simulation in the Frequency Domain also reported on the development of *Harmonica*, a harmonic balance simulator that was to be very influential; spawning several of today's leading simulators. The name was later changed to *Spectre* when the *Harmonica* name was co-opted by Compact Software for the name of its independently developed harmonic balance simulator. *Spectre* was successful, at least in part, because it followed a formula pioneered by SPICE. *Spectre* had a SPICE-like netlist and use-model, it placed no artificial limits on the number of nodes or components in the circuit, and most importantly, its source code was made freely available to anyone willing to pay a nominal fee. The idea that Berkeley should simply give away the source code to its simulators was first championed by Don Pederson, and is credited with much of the broad success of SPICE [49]. This allowed Berkeley *Spectre* to become the foundation of what are currently the two leading RF simulators, Agilent's *ADS* and Cadence's *SpectreRF*, and a leading integrated circuit simulator, Cadence's *Spectre*. It is interesting to note that even though Cadence's simulators are direct descendents of Berkeley's *Spectre*, neither of them currently uses harmonic balance. Rather, *Spectre* is a SPICE-class simulator and *SpectreRF* uses shooting algorithms that were developed as part of follow-on research to the original harmonic balance work.

In another interesting aside, this paper started a spirited competition between the CAD community and the microwave community [56, 55]. Each produced papers at a relatively rapid pace on variations of, and extensions to, harmonic balance in an attempt to lay first claim to the resulting innovations. Even competing simulators were produced. While both communities had their successes, in the end it was the CAD community's deep knowledge of numerical algorithms and large-scale programming techniques that produced the biggest advances and ultimate success. In particular, it was the CAD community that substantially advanced the capabilities of RF simulators by enhancing shooting methods [65], introducing Krylov methods, and developing efficient methods for time-varying noise analysis [66, 58, 14].

6. Noise Analysis

In 1971, Rohrer et al [57] established the noise analysis method that eventually made its way into SPICE [45] and became the de facto stan-

dard for three decades. This method was accurate, robust, and efficient, but was limited in the type of circuits it could handle. With their approach the circuit was first linearized about the DC operating point. The noise analysis is performed on the resulting linear time-invariant (LTI) representation by computing the frequency dependent transfer functions from noise source to the output of interest. Their key insight, still important nearly thirty years later, is that for typical circuits there are many noise sources but noise is measured at only a few outputs. This many-input few-output case is much more efficiently analyzed using an adjoint formulation.

Dramatically faster computers now allow designers to routinely simulate much more complicated analog circuits, with commensurately complicated noise spectra, possibly with multiple sharp peaks and nulls. In the approach in [57], the frequency-dependent output-to-noise conjugate transfer functions are computed by solving the transpose of the linearized circuit equation for each frequency of interest. In order to capture sharp features in the noise spectra, it would be necessary solve at a large number of frequencies. The fifth paper in our commemorative collection, *Circuit noise evaluation by Padé approximation based model-reduction techniques* by Peter Feldmann and Roland Freund (see page **??**), cleverly applies their well-known Padé-via-Lanczos algorithm for model reduction [18] to the problem of avoiding the multiple frequency solves in noise analysis. In this approach, a rational function description for the noise spectra is computed directly, and therefore sharp spectral features are easily captured. Extensions of this method can be used for computing noise in more complicated time-varying cases described below [59].

Analyzing the circuit about the DC operating point is adequate for simple circuits such as amplifiers and passive circuits such as filters, but cannot be used on circuits for which the noise performance is strongly affected by large signals, such as with mixers, oscillators, samplers, switched-capacitor filters, and the like. The large, generally periodic, signals present in these circuits act to modulate both the noise sources and the transfer characteristics of the circuit from the noise sources to the output. The result is cyclostationary noise, or noise with time-varying statistics, at the output of the circuit. With the recent rise of importance of RF circuits, a more general type of noise analysis became essential. What was needed was the ability to perform a noise analysis not about the DC operating point, but about a time-varying operating point.

Okumura et al [48] extended Rohrer's method to support prediction of noise of circuits linearized about a periodic operating point. These ideas were later commercialized by Telichevesky [66] and others. In this way the noise of circuits such as mixers and switched-capacitor filters could

be predicted accurately. However, there were still a great number of fundamental questions left unanswered, particularly regarding oscillator phase noise.

Time-Domain non-Monte Carlo Noise Simulations for Nonlinear Dynamic Circuits with Arbitrary Excitations was published at ICCAD in 1994 by Alper Demir, Edward Liu, and Alberto Sangiovanni-Vincentelli. It is the last physical simulation and analysis paper highlighted in this commemorative collection and can be found on page **??**. In a break from EDA tradition, Demir et al proposed to perform noise analysis by formulating and solving the stochastic differential equations for the circuit [10, 11]. The approach offered a unique advantage in that it did not require a periodic operating point, or even that the circuit be in steady state. However, his method was also perceived as being complex and was computationally expensive for large circuits. The added generality of the method was not seen as a compelling advantage in light of the disadvantages, so the method he proposed has not seen much use. Nevertheless, in these papers, Demir et al were the first in the EDA field to advocate a more rigorous approach to noise analysis. They were also the first to introduce to the design community the theoretical foundation that would be needed to address the difficult questions of which they were just becoming aware.

These papers were just the first of several by Demir that used stochastic differential equations to deeply explore questions of noise, particularly oscillator phase noise [12, 13, 14, 15]. This flurry of results led, either directly or indirectly, to papers by many others that together advanced the fundamental understanding of noise and improved both the analysis [58, 9, 68, 51, 69] and design of low-noise circuits [36, 53].

7. Conclusions and Acknowledgments

The area of physical simulation and analysis has been enlivened in recent years by the growing importance of problems associated with RF design and interconnect effects. Since many of the key papers in these fields did not initially appear at ICCAD, and therefore were not considered for this commemorative collection, we hope that this commentary both celebrates the selected papers *and* recognizes the importance of contributions that have appeared elsewhere.

It is not clear what emerging problems will provide the stimulation to open up new directions in this area, but given how often the CAD community has incorrectly predicted that physical simulation research has plateaued, the authors would like to wait and see.

We would like to thank Joel Phillips for his insightful technical comments while evolving this commentary, and for providing some first-hand historical information. In addition, we would like to thank the many contributors to physical simulation and analysis both at ICCAD and elsewhere.

References

[1] George A. Baker, Jr. and Peter Graves-Morris, *Padé Approximants, 2nd Edition* Cambridge University Press, 199

[2] M. Beattie, B. Krauter, L. Alatan, and L. Pileggi. Equipotential shells for efficient inductance extraction. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*. vol. 20, no. 1, January, 2001, pp. 70-79.

[3] J. E. Bracken, V. Raghavan, and R. A. Rohrer. Interconnect simulation with asymptotic waveform evaluation. *IEEE Trans. Circuits Syst.*, 39(11):869-878, November 1992.

[4] J. E. Bracken. Passive modeling of linear interconnect networks. *Widely circulated notes*, 1995.

[5] A. C. Cangellaris and L. Zhao. Passive reduced-order modeling of electromagnetic systems. *Computer Methods in Applied Mechanics and Engineering* vol. 169, no. 3-4, February 1999, pp. 345-358.

[6] P. K. Chan. An extension of Elmore's delay. *IEEE Trans. on Circuits and Systems*, CAS-33(11):1147-1149, 1984.

[7] Eli Chiprout and Michel S. Nakhla. Analysis of interconnect networks using complex frequency hopping (CFH). *IEEE Trans. CAD*, 14:186-200, February 1995.

[8] F. Dartu, N. Menezes, and L. Pilegi. Performance computation for precharacterized CMOS gates with RC loads, *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 15, No. 5, May 1996.

[9] A. Dec, L. Toth, and K. Suyama. Noise analysis of a class of oscillators. *IEEE Transactions on Circuits and Systems II*, vol. 45, no. 6, pp. 757-760, June 1998.

[10] Alper Demir, Edward W. Y. Liu, and Alberto L. Sangiovanni-Vincentelli. Time-domain non-Monte Carlo noise simulation for nonlinear dynamic circuits with arbitrary excitations. *IEEE/ACM International Conference on Computer-Aided Design*, 1994, pp. 598-603.

[11] A. Demir, E. Liu, and A. Sangiovanni-Vincentelli. Time-domain non Monte-Carlo noise simulation for nonlinear dynamic circuits with arbitrary excitations. *IEEE Transactions on Computer-Aided Design*, vol. 15, pp. 493-505, May 1996.

[12] A. Demir, A. Mehrotra, and J. Roychowdhury. Phase noise and timing jitter in oscillators. *Proceedings of the IEEE 1998 Custom Integrated Circuits Conference* (CICC-1998), pp. 45-48.

[13] A. Demir. Phase noise in oscillators: DAEs and colored noise sources. *IEEE/ACM International Conference on Computer-Aided Design, Digest of Technical Papers*, ICCAD-1998, pp. 170-177.

[14] A. Demir, A. Mehrotra, and J. Roychowdhury. Phase noise in oscillators: a unifying theory and numerical methods for characterization. *IEEE Transactions on*

Circuits and Systems I: Fundamental Theory and Applications, vol. 47, no. 5, May 2000, pp. 655-674.

[15] A. Demir, D. Long, and J. Roychowdhury. Computing phase noise eigenfunctions directly from steady-state Jacobian matrices. *IEEE/ACM International Conference on Computer Aided Design, 2000* (ICCAD-2000). pp. 283-288.

[16] Ibrahim M. Elfadel and D. D. Ling, A block rational Arnoldi algorithm for multipoint passive model-order reduction of multiport RLC networks, *International Conference on Computer Aided-Design*, San Jose, California, November 1997.

[17] W. C. Elmore. The transient response of damped linear networks with particular regard to wide-band amplifiers. *Journal of Applied Physics*, 19(1):55-6:3, January 1948.

[18] P. Feldmann and R. W. Freund. Efficient linear circuit analysis by Padé approximation via the Lanczos process. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 14:639-649, 1995.

[19] Peter Feldmann and Roland W. Freund. Reduced-order modeling of large linear subcircuits via a block Lanczos algorithm. In 32$^{\text{nd}}$ *ACM/IEEE Design Automation Conference*, pp. 474-479, San Francisco, California, June 1995.

[20] K. Gala, V. Zolotov, R. Panda, B. Young, J. Wang, and D. Blaauw. On-chip inductance modeling and analysis. *2000 IEEE/ACM Design Automation Conference*, pp. 63-68.

[21] K. Gallivan, E. Grimme, and P. Van Dooren. Asymptotic waveform evaluation via a Lanczos method. *Applied Mathematics Letters*, 7(5):75-80, 1994.

[22] K. Glover. All optimal Hankel-norm approximations of linear multivariable systems and their L^{∞}-error bounds. *Int. J. Control*, vol.39, No.6, pp.1115-1193, 1984.

[23] Eric Grimme. *Krylov Projection Methods for Model Reduction.* PhD thesis, Coordinated-Science Laboratory, University of Illinois at Urbana-Champaign, Urbana-Champaign, IL, 1997.

[24] M.A. Horowitz. Timing Models for MOS Circuits. PhD thesis, Stanford University, January 1984.

[25] M. Kamon, M. J. Tsuk, C. Smithhisler, and J. White. Efficient techniques for inductance extraction of complex 3-D geometries. *IEEE/ACM International Conference on Computer-Aided Design Digest of Technical Papers 1992* (ICCAD-92), pp. 438-442.

[26] M. Kamon and J. White. Preconditioning for multipole-accelerated 3-D inductance extraction. *Electrical Performance of Electronic Packaging*, 1993, pp. 189-192.

[27] M. Kamon and J. R. Phillips. Preconditioning techniques for constrained vector potential integral equations, with application to 3-D magnetoquasistatic analysis of electronic packages. *Proceedings on the Colorado Conference on Iterative Methods*, Breckenridge, CO, April 1994.

[28] M. Kamon, M. J. Tsuk, and J. K. White. FastHenry: a multipole-accelerated 3-D inductance extraction program. *IEEE Transactions on Microwave Theory and Techniques*, vol. 42, no. 9, part 1 of 2, Sept. 1994, pp. 1750-1758.

[29] M. Kamon, B. Krauter, J. Phillips, L. Pileggi, J. White. Two optimizations to accelerate method-of-moments algorithms for signal integrity analysis of complicated 3-D packages. Electrical Performance of Electronic Packaging, 1995, pp. 213 -216.

[30] K. J. Kerns, I. L. Wemple, and A. T. Yang. Stable and efficient reduction of substrate model networks using congruence transforms. In *IEEE/ACM International Conference on Computer Aided Design*, pp. 207-214, San Jose, CA, November 1995.

[31] B. Krauter, L. T. Pileggi. Generating sparse partial inductance matrices with guaranteed stability. *1995 IEEE/ACM International Conference on Computer-Aided Design Digest of Technical Papers*, pp. 45-52.

[32] B. Krauter and S. Mehrotra. Layout based frequency dependent inductance and resistance extraction for on-chip interconnect timing analysis. *Proceedings of the 1998 Design Automation Conference*, pp. 303-308.

[33] Kenneth S. Kundert and Alberto Sangiovanni-Vincentelli. Nonlinear circuit simulation in the frequency domain. *IEEE International Conference on Computer-Aided Design Digest of Technical Papers*, pp 240-242, November 1985.

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

[35] Ken Kundert. Introduction to RF Simulation and its Application. *IEEE Journal of Solid-State Circuits*, vol. 34, no. 9, September 1999.

[36] T. H. Lee and A. Hajimiri. Oscillator phase noise: a tutorial. *IEEE Journal of Solid-State Circuits*, vol. 35. No. 3. pp. 655-674, March 2000.

[37] H. Levy, D. MacMillen, and J. White, A rank-one update method for efficient processing of interconnect parasitics in timing analysis, *Proceedings of the Design Automation Conference*, Los Angeles, June, 2000, pp 75-7

[38] T.-M. Lin and C.A. Mead. Signal delay in general RC networks. *IEEE Trans. on Computer Aided Design*, CAD- 3:331-349, 1984.

[39] N. Marques, M. Kamon, J. White, L. M. Silveira, A mixed nodal-mesh formulation for efficient extraction and passive reduced-order modeling of 3D interconnects, *Proceedings of the 35th Design Automation Conference*, San Francisco, June, 1998, pp. 297-302.

[40] C. A. Mead and L. Conway. *An Introduction to VLSI Systems.* Addison Wesley, 1980.

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

[42] K. Nabors and J. White. A fast multipole algorithm for capacitance extraction of complex 3-D geometries. *Proceedings of the 1989 IEEE Custom Integrated Circuits Conference*, pp. 21.7/1-21.7/4.

[43] K. Nabors and J. White. FastCap: a multipole accelerated 3-D capacitance extraction program. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 10, no. 11, Nov. 1991, pp. 1447-1459.

[44] K. Nabors, T.-T. Fang, H.-W. Chang, K. Kundert, and J. White. A Gaussian-quadrature based algorithm for RLC-line to RLC-line reduction, *Proceedings of the 34th Design Automation Conference*, Anaheim, CA, June, 1997

[45] L. W. Nagel. *SPICE2: a computer program to simulate semiconductor circuits.* Electronic Research Lab. Report No. ERL-M520, University of California, Berkeley, May 1975.

[46] Peter R. O'Brien and Thomas L. Savarino. *Modeling the driving-point characteristics of resistive interconnect for accurate delay estimation IEEE/ACM International Conference on Computer-Aided Design Digest of Technical Papers 1989* (ICCAD-89), pp. 512-515.

[47] Altan Odabasioglu, Mustafa Celik, and Lawrence Pileggi. PRIMA: A Passive reduced-order interconnect macromodeling algorithm. In *International Conference on Computer Aided-Design*, pp. 58-65, San Jose, California, November 1997.

[48] Makiko Okumura, Hiroshi Tanimoto, Tetsuro Itakura, and Tsutomu Sugawara. Numerical noise analysis for nonlinear circuits with a periodic large signal excitation including cyclostationary noise sources. *IEEE Transactions on Circuits and Systems – 1: Fundamental Theory and Applications*, vol. 40, no. 9, September 1993, pp. 581-919.

[49] Tekla S. Perry. Profile: Donald O. Pederson. *IEEE Spectrum*, vol. 35, no. 6, June 1998.

[50] J. R. Phillips, E. Chiprout, and D. D. Ling, Efficient full-wave electromagnetic analysis via model-order reduction of fast integral transforms. In *Proceedings of the 33nd Design Automation Conference*, Las Vegas, June, 1996.

[51] J. Phillips, and K. Kundert. Noise in mixers, oscillators, samplers, and logic an introduction to cyclostationary noise. *Proceedings of the IEEE Custom Integrated Circuits Conference, 2000* (CICC-2000), pp. 431-438.

[52] Lawrence T. Pillage and Ronald A. Rohrer. Asymptotic waveform evaluation for timing analysis. *IEEE Trans. CAD*, 9(4):352-366, April 1990.

[53] J. J. Rael and A. A. Abidi. Physical processes of phase noise in differential LC oscillators. *Proceedings of the 2000 IEEE Custom Integrated Circuits Conference* (CICC 2000), pp. 569 -572.

[54] Curtis L. Ratzlaff and Lawrence T. Pillage. RICE: Rapid interconnect circuit evaluation using AWE. *IEEE Trans. CAD*, 13(6):763-776, June 1994.

[55] G. Rhyne, M. Steer and B. Bates. Frequency-domain nonlinear circuit analysis using generalized power series. *IEEE Transactions on Microwave Theory and Techniques*, vol. 36, no. 2, pp. 717-720, February 1988.

[56] V. Rizzoli and A. Neri. State of the art and present trends in nonlinear microwave CAD techniques. *IEEE Transactions on Microwave Theory and Techniques*, vol. 36, no. 2, pp. 343-365, February 1988.

[57] R. Rohrer, L. Nagel, R. G. Meyer, and L Weber. Computationally efficient electronic-circuit noise calculations. *IEEE Journal of Solid-State Circuits*, vol. SC-6, no. 4, pp. 204, August 1971.

[58] J. Roychowdhury, D. Long, and P. Feldmann. Cyclostationary noise analysis of large RF circuits with multitone excitations. *IEEE Journal of Solid-State Circuits*, vol. 33, no. 3, March 1998, pp. 324-336.

[59] J. Roychowdhury, Reduced-order modeling of time-varying systems *IEEE Transactions on Circuits and Systems-II: Analog and Digital Signal Processing*, vol. 46, no. 10, October 1999, pp. 1273-1288.

[60] J. Rubinstein, P. Penfield, Jr., and M. A. Horowitz. Signal delay in RC tree networks. *IEEE Transaction on Computer-Aided Design*, CAD-2(3):202-211, July 1983.

[61] A. E. Ruehli. Survey of computer-aided electrical analysis of integrated circuit interconnections. *IBM Journal of Research and Development*, vol. 23, November 1979, pp. 626-639.

[62] L. M. Silveira, I. M. Elfadel, J. White. A guaranteed stable model-order reduction algorithm for packaging and interconnect simulation. *Proceedings of the IEEE 2nd Topical Meeting on Electrical Performance of Electronic Packaging*, Monterey, CA, October, 1993, pp. 166-168.

[63] L. Miguel Silveira, M. Kamon and J. White, Efficient reduced-order modeling of frequency-dependent coupling inductances associated with 3-D interconnect structures, *Proceedings of the 32nd Design Automation Conference*, pp. 376-380, San Francisco, CA, June, 1995.

[64] K. L. Shepard and Tian Zhong. Return-limited inductances: a practical approach to on-chip inductance extraction. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 19, no. 4, April 2000, pp. 425-436.

[65] R. Telichevesky, K. Kundert and J. White. Efficient steady-state analysis based on matrix-free Krylov-subspace methods. *Proceedings of the 32nd Design Automation Conference*, June 1995.

[66] Ricardo Telichevesky, Kenneth S. Kundert, and Jacob K. White. Efficient AC and noise analysis of two-tone RF circuits. *Proceedings of the 33rd Design Automation Conference*, June 1996.

[67] V. Thomas, M. Jones, M. J. Piket-May, A. Taflove, and E. Harrigan, The use of SPICE lumped circuits as sub-grid models for FD-TD analysis. *IEEE Micro. Guided Wave Letters*, July 1994.

[68] L. Toth, I. Yusim, K. Suyama. Noise analysis of ideal switched-capacitor networks. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 46, no. 3, March 1999, pp. 349-363.

[69] L. Toth. Analytical approach for the exact phase noise analysis of oscillators. *The 2001 IEEE International Symposium on Circuits and Systems* (ISCAS 2001), vol. 2, pp. 771-774.

[70] J. L. Wyatt, Jr. Signal propagation delay in RC models for interconnect, chapter 11.2, pages 254-290. *Advances in CAD for VLSI*. Elsevier Science Publishers B. V., North Holland, 1987.