

Rapid Yield Estimation as a Computer Aid for Analog Circuit Design

Tamal Mukherjee, *Student Member, IEEE*, and L. Richard Carley, *Senior Member, IEEE*

Abstract—A rapid yield estimation methodology that aids the analog circuit designer in making design trade-offs that improve yield is presented. This methodology is based on using hierarchical evaluation of analysis equations, rather than simulation, to predict circuit performance. The new analog rapid yield estimation (ARYE) method has been used to predict the yield of two-stage op amps, and has been incorporated into the Carnegie Mellon University (CMU) analog design system (ACACIA). An example of how ARYE allows analog designers to quickly explore the impact of design changes on yield will be presented. The primary goal of ARYE is to make numerous early predictions of parametric yield economical for the analog circuit designer.

I. INTRODUCTION

AN EVER-INCREASING fraction of application-specific integrated circuit (ASIC) designs incorporate analog circuit blocks. A significant fraction of these analog blocks may fail to satisfy their performance specifications due to global variations in the fabrication process (parametric yield loss); in part, this is because the performance of analog circuits is more sensitive to variations in process parameters than the performance of digital circuits. Although analog circuits can also fail due to functional causes [1], this paper will only address parametric yield failures. Since the yield of an ASIC is an important component in determining the manufacturer's profit, and that yield is dominated by the blocks which have the worst yield, estimating and optimizing the yield of analog circuit designs prior to manufacturing is an important task.

The problem of estimating the parametric yield of analog circuits is made difficult by the complex dependence of their performance on process parameters; i.e., it is not sufficient to use circuit simulation to verify the analog circuit's performance at a few process corners (e.g., slow-slow, fast-fast, etc.) as is typically done with digital circuits. Yield modeling and optimization for digital circuits has been greatly debated in recent years, as overviewed in [1]. Currently, analog circuit yield is typically estimated using a Monte-Carlo approach, i.e., by simulating the performance of many instances of the design, each instance representing a possible outcome of the fabrication process. The general process of yield estimation will be discussed in detail in Section II. Unfortunately, this approach typically requires hours, days, or even weeks of computer time, severely limiting its use. Because of its expense, this yield estimation process is usually performed only

once, typically at the very end of the design process, to detect major yield problems before fabrication.

The goal of the proposed analog rapid yield estimation (ARYE) method, which will be described in detail in Section III, is to allow the analog circuit designer to rapidly compare the parametric yield loss of many possible circuit alternatives early in the design process. This allows the designer to consider the impact on yield of major decisions such as circuit topology selection before they are finalized. The primary innovation of the ARYE method is its replacement of circuit simulation to determine analog circuit performance with hierarchical analysis equations. Although these analysis equations may be only approximate, as will be seen, they offer an important alternative to circuit simulation, which is orders of magnitude slower. A discussion of the accuracy limits of analysis equations and methods for constructing analysis equations will be presented in Section III-A.

The ARYE method has been used for yield prediction of op amps, and has been incorporated into the Carnegie Mellon University (CMU) analog design system (ACACIA) to allow analog designers to automatically design op amps and then to quickly explore the impact of design changes on the circuit yield. The specific details of this implementation of the ARYE methodology appear in Section IV. In addition, the use of the ARYE/ACACIA environment to design an operational amplifier, predict its yield, and then to modify the design will be illustrated in Section V.

II. BACKGROUND: YIELD PREDICTION

Two different classes of faults, reasons for a circuit to fail to perform within specified limits, are normally distinguished: functional faults and parametric faults [2]. Functional faults are the result of random local problems in the IC fabrication process in any of the IC layers, i.e., missing or extra pieces of material on any layer (spot defects). Functional faults generally result in structural changes in the connectivity of the IC, either open circuiting a connection or shorting two or more nodes together. For example, in the case of an operational amplifier a functional fault such as an oxide pinhole between the gate and source of one transistor might result in a circuit that cannot amplify a signal at all.

Parametric faults, on the other hand, are caused by global deviations in the process such as variations in substrate doping density or wafer-wide polysilicon line width variation from nominal. Parametric faults cause the performance of the analog circuit to fall outside the specified range; e.g., in an op amp, excessive gate oxide thickness, in all of the transistors, can lead to a decrease in g_m resulting in a failure to meet the unity-gain frequency specification.

Manuscript received August 3, 1990; revised November 1, 1990. This work was supported in part by the Semiconductor Research Corporation, by DARPA under Grant MDA972-88-J-1006, and by the National Science Foundation under Grant MIP-8451496.

The authors are with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213.
IEEE Log Number 9041557.




Fig. 1. Monte-Carlo yield estimation strategy.

For many mixed signal ASIC's the die area is dominated by digital circuitry; a typical ASIC might be 90% digital circuitry and 10% analog circuitry by area. Assuming that the incidence of spot defects is uniform across the die and that analog and digital circuits are roughly equivalent in their sensitivity to failure due to these spot defects,¹ then the likelihood of a functional fault due to a spot defect in the analog circuitry is proportional to the fraction of the die area occupied by the analog circuitry, e.g., one-ninth that of the likelihood of a functional fault in the digital circuitry for the above example. Therefore, if the area of the die is predominantly digital and the functional yield loss for the digital circuitry is acceptably small, then the functional yield loss of the analog circuitry is negligible. Therefore, in the remainder of this paper only yield loss due to parametric faults will be discussed.

In order to simplify the design task, many designers often assume that device parameters, such as the device threshold voltage, are single valued. They then attempt to achieve good circuit parametric yield by anticipating all of the worst-case combinations of device model parameters, leading to conservative designs that can fall far short of the feasible performance in a given technology. State-of-the-art designs require statistically accurate predictions of the parametric yield of a circuit design.

The most common method for predicting the yield of an analog circuit is based on a Monte-Carlo approach (see Fig. 1). First, a number of hypothetical instances of the circuit are generated with a statistical distribution that matches, to the extent possible, that of the actual process line. More specifically, a set of device model parameters is generated for each device in each design instance. Next, the performance of each design instance is predicted. This is usually done by running circuit simulations on each design instance and extracting the circuit's performance from the simulator's outputs. For each design instance, the predicted performance is compared to the performance specifications, and the number of designs which satisfy all of the specifications are counted. The estimated parametric yield loss is simply the fraction of simulated designs that fails to meet all of their specifications. Accurate estimation of parametric yield often requires that thousands of instances of the circuit be generated and simulated [3]. There are two main components to using this approach for yield estimation. The first component is to generate device models for instances that are as close as possible to the actual process line. The second

component is to evaluate the performance of each circuit rapidly and efficiently.

It is important to note that, due to the correlations between the device model parameters, it is not sufficient to obtain the means and variances of each device model parameter in a given process and to generate independent random distributions for them. Instead, it is necessary to use a set of independent random variables that can be mapped in an algorithmic manner into the desired device model parameters. Process simulators do this by choosing the variables related to the fabrication process (e.g., diffusion times and temperatures) as the independent random variables [4], [5]. In order to use a process simulator, the probability distributions of the set of independent fabrication process variables must be determined for the processing line [6]. The process simulator generates independent random variables for the process variables and maps them into sets of properly correlated device model parameters (such as k' and V_{T0}).

Simulating the performance of each of the generated instances of the analog circuit can be a major problem, particularly for complex analog circuits. Circuit simulations of analog circuits frequently experience numerical difficulties, particularly in converging to the initial dc node voltages. Analog designers often coerce the circuit simulator to the right starting point by providing good initial guesses for node voltages. However, as variations in the process occur, these initial guesses may not be close enough to enable the simulator to converge. In these cases, even though the circuit might work correctly, it will appear to be a failure. In fact, because of this difficulty, a frequently employed alternative to the Monte-Carlo approach is to simulate the analog circuit for a small set (e.g., 10) of *very* carefully selected instances whose device model parameters were chosen to be representative of the typical variations measured for the process. Obviously, this method is very ad hoc, and can only be successful for a process line with which designers have acquired significant experience.

III. ANALOG RAPID YIELD PREDICTION METHOD

The ARYE method for yield estimation follows the Monte-Carlo approach presented in the previous section. However, in order to dramatically increase the speed of yield estimation, two important modifications have been made. First, the very time-consuming process of simulation and performance extraction has been replaced with a set of analysis equations that is solved to determine performance. Second, the process of generating design instances has been removed from the yield prediction loop. Instead, libraries of instances are generated once for the desired process, and instances are used as needed during yield prediction. The details of the implementation, and the implications for accuracy and speed, of these two modifications will be presented in detail in this section.

A block diagram for the ARYE method is shown in Fig. 2. A design instance, i.e., a set of device model parameters, is selected from the library. The analytical prediction of the performance, using the selected device model parameters, is compared with the performance specifications to determine which instances met the specifications (passed) and which did not (failed). The number of chips that passed (i.e., the number of device parameter sets for which the nominal design meets the performance specifications) divided by the

¹This is a reasonable assumption given that both types of circuitry are typically patterned with the same physical design rules concerning line widths and separation between lines on various layers.




Fig. 2. Device model parameter libraries.

total number of chips (i.e., the total number of device parameter sets) is the yield prediction.

A. Analysis Equations for Performance Prediction

As mentioned above, parametric yield estimation for complex analog circuits is computationally expensive because of the large number of circuit simulations required. The emerging importance of sets of analysis equations for estimating circuit performance in terms of device parameters for circuit synthesis, e.g., [7]–[11], suggested an alternate methodology for rapid yield estimation. Although approximate, analysis equations that predict performance require comparatively little computational time, making rapid parametric yield estimation attractive. Because this method for rapid yield prediction is based on approximate analysis equations, a final yield estimation using full circuit-level simulations would still be desirable for verification of the yield prediction prior to actual fabrication. The primary goal of this research is to make numerous early predictions of parametric yield economical for the designer. In this section the accuracy of analysis equations for performance prediction and the difficulties encountered when developing accurate analysis equations for new circuit topologies will be discussed.

The analog circuit whose yield is being predicted is usually, though not necessarily, decomposed into a number of common subblocks. For example, a two-stage operational amplifier can be decomposed into the several current mirrors, a differential pair, a level shifter, and a transconductance amplifier. Performance prediction begins by traversing upward through the hierarchical decomposition, starting with the lowest level subblocks and working upward until the performance prediction of the circuit is completed. For instance, starting with the specific topology of current mirror selected (e.g., simple or cascoded) and the device sizes, the analysis equations predict the current mirror's output resistance, current gain, minimum output voltage, bandwidth, etc. Once the performance of each subblock has been computed, that information is used to compute the performance of the components that are higher up in the hierarchy, eventually leading to the predicted performance for the desired circuit. The information flow for the performance prediction of a two-stage operational amplifier is shown in Fig. 3. Note that this information flow need not be purely upward. If two subblocks interact, for example in determining a dc voltage, then the analysis equations at the next higher level must iterate on these two subblocks until they converge.

Analysis equations must predict both large-signal and small-signal performance characteristics of the circuit. In order to make this possible, simplified device models are used to facilitate the process of writing analysis equations. In general, errors on the order of 5 to 10% normally occur in

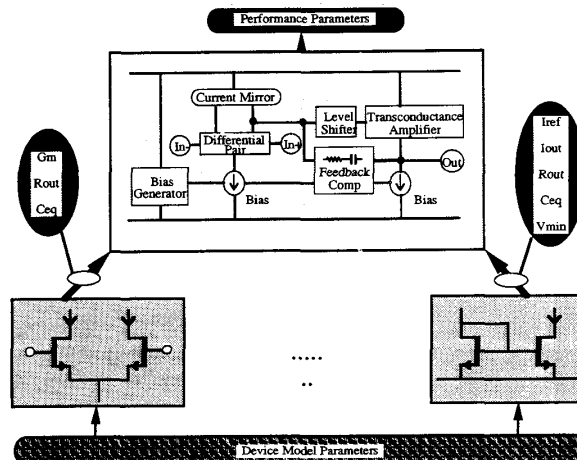


Fig. 3. Information transfer in ARYE performance prediction.

the process of generating a mathematically tractable set of analysis equations, although this can be made smaller by increasing the complexity of the equations employed. For example, in this research two levels of complexity have been employed in writing analysis equations for two-stage op amps: analysis equations have been created that employ SPICE Level 1 models for MOS devices and more complex but more accurate BSIM models [12] for MOS devices. Tables I and II give an example of the accuracy of these two models for two typical two-stage op-amp designs with different device-level circuit topologies. Comparing these analyses to circuit simulation using the BSIM model in the circuit simulator HSPICE [13], only very small errors occur for op-amp circuit 1, except for the dc gain which is a very sensitive function of both the device model and the exact bias point. However, only the minimum value of the dc gain of an op amp is normally important in op-amp design so the underestimation of the gain is acceptable. For op-amp circuit 2, which employs a much more complex topology with cascoded current mirrors and level shifters, the worst-case error occurs when predicting the phase margin: a 12.5% overestimation by the Level 1 model and a 7.75% overestimation by the BSIM model. In both cases the error is caused by a failure to accurately model the extra phase lag caused by the level shifter. By further adding analysis equations that predict this extra phase shift, these errors could be reduced still further. Note that the price of using the more complex models is apparent in the computer time required to predict performance (see Tables I and II), which increased by roughly a factor of 5 when going from the simple model to the complex one; however, even the complex analysis equations are still more than an order of magnitude faster than using a circuit simulator.

Because of the large amount of computer time required for final yield verification, rapid yield estimation should err on the conservative side. After the analysis equations are characterized in terms of their worst-case error over the desired range of topologies and performance specifications, the threshold of acceptability for each performance specification should be biased by the worst-case error on that specification in order to guarantee that the desired yield will

TABLE I
ACCURACY OF LEVEL-1- AND BSIM-BASED ANALYTICAL EQUATIONS—CIRCUIT 1

Performance Parameter	Level-1-based Analytical Model	BSIM-based Analytical Model	HSPICE™
DC Gain (dB)	56.564	56.0889	60.817
Unity Gain Frequency (MHz)	11.44	12.63	12.25
Phase Margin (degrees)	59.28	59.87	59.27
Slew Rate (V/μs)	15.85	17.10	17.02
Execution Time (s)	0.1	0.6	11.4

TABLE II
ACCURACY OF LEVEL-1- AND BSIM-BASED ANALYTICAL EQUATIONS—CIRCUIT 2

Performance Parameter	Level-1-based Analytical Model	BSIM-based Analytical Model	HSPICE™
DC Gain (dB)	68.229	64.301	69.402
Unity Gain Frequency (MHz)	1.193	1.252	1.266
Phase Margin (degrees)	66.976	63.53	58.62
Slew Rate (V/μs)	1.873	1.871	1.86
Execution Time (s)	0.2	0.7	14.4




Fig. 4. Block diagram of hierarchical two-stage op-amp model.

be achieved, despite small inaccuracies in the analysis equations.

In order to make use of the proposed rapid yield estimation technique for a new circuit, it is necessary to create analysis equations that can accurately predict the new circuit's performance. In some cases, formation of accurate analysis equations can be quite difficult, which is a potential obstacle to using this method. This difficulty can be compounded by the diversity of topologies found, even for a single circuit block. For example, there are literally hundreds of op-amp topologies. Instead of creating analysis equations that predict the performance of every possible op-amp topology, hierarchy can be used to greatly reduce the effort needed to create these analysis equations.

For example, even within the narrow range of two-stage transconductance op amps, the current ARYE system can predict the performance of 128 different device-level topologies using a hierarchy of blocks. The subblock topology of the two-stage transconductance op amp is shown in Fig. 4. Analysis equations exist to predict the performance of simple and cascoded current mirrors, simple and cascoded differential pairs, simple and cascoded transconductance amplifiers, and simple and source-follower level shifters from MOS device sizes. Also, analysis equations exist to predict the performance of the overall two-stage op-amp block given the performance of each of the six subblocks. Each subblock has two possible topologies² and the overall op amp can have

²Note that the choice for the current mirror acting as the load for the differential pair and the choice for level shifter are typically coupled in order to maintain a low systematic offset in the face of variations in the

either NMOS inputs or PMOS inputs resulting in 128 different device-level two-stage op amps whose performance can be predicted from just nine sets of analysis equations. Even more important, the set of analysis equations needed for each of these blocks is much simpler than the set of analysis equations needed to express a single op amp without hierarchy. Essentially, because of the hierarchical decomposition, fewer variables are dealt with in creating any one set of analysis equations. This represents a tremendous savings in the time required to create the analysis equations for a new circuit.

The use of a hierarchical formulation also allows reuse of analysis equations: there are three current mirrors in the two-stage op amp, but only one set of analysis equations is needed for each possible current mirror topology. Analysis equations for a subblock can even be reused in the creation of analysis equations for new analog circuit blocks. For example, analysis equations for a comparator could make use of the differential pair and current-mirror analysis equations already developed for the op amp, substantially decreasing the time required to create complete comparator analysis equations. The hierarchical formulation also has the advantage that all of the analysis equations relating to a particular subblock are located in one place, making maintaining, updating, and debugging much easier.

Another technique that can greatly facilitate the creation of analysis equations for the linear part of an analog circuit block or subblock is the use of a symbolic simulator such as ISAAC [14] which generates and simplifies equations for arbitrary linear circuits. For op amps, a symbolic simulator can be used to automatically generate almost half of the necessary analysis equations. However, the analog designer must still provide analysis equations that predict the dc operating point and any large-signal performance specifications.

process delta-L; i.e., when the current mirror is simple then the level shifter is a simple wire, and when the current mirror is cascoded then a source follower is used as a level shifter. However, for designs in which the op-amp offset is unimportant, this choice is not mandated; the choice of level-shifter topology can be made independent of the choice of current-mirror topology.

B. Generating Design Instances

The global random disturbances that affect the fabrication process can be divided into die-to-die disturbances and device-to-device disturbances within the die. Although, die-to-die disturbances can be further divided into die-to-die (on the same wafer) disturbances, wafer-to-wafer disturbances, and lot-to-lot disturbances, for the purposes of this paper these will all be lumped as die-level disturbances. In order to improve the speed of yield estimation, the process of generating design instances for die-level disturbances—a set of device model parameters for each instance—has been removed from the yield prediction loop. Because the die-level process disturbances do not depend on the specific circuit being fabricated,³ a set of instances for them can be generated once for a given process, and need not be generated each time the yield of a specific circuit is to be predicted. This is consistent with the goal of this approach, which is to make frequent predictions of the yield of a single circuit as its topology and device sizes are adjusted so as to improve its performance and manufacturability.

There are several possible methods that can be employed to generate this library of design instances for die-level disturbances. The most obvious method to generate these design instances, the correlated sets of device model parameters, is to simply compile a large library of device models extracted from actual fabrication runs. In order that the distribution match that of the actual fabrication process, the number and distribution of test cells being characterized on each wafer should correspond to the number and distribution of die on each wafer. Although this method is guaranteed to give a realistic prediction for that fabrication line, extraction of the data and fitting of model parameters for the required (typically thousands) number of design instances is a formidable task.

As mentioned in Section II, a second method for generating the correlated sets of device model parameters is to use a statistical process simulator, e.g., FABRICS II [5]. However, this requires that the means and variances of the independent variables used by the process simulator, as well as various parameters of the process simulator, be properly adjusted for the actual fabrication line [6]. Note that although it would be possible to place the process simulator within the Monte-Carlo loop (see Fig. 2), this would not be desirable because process simulators, in general, require significant computer time to generate each new design instance.

A third possible method is based on measuring a small (relative to the number of instances that will be used during Monte-Carlo yield estimation) number of design instances from the actual fabrication line, extracting the underlying statistical relationships, and then generating a large number of design instances that are drawn from the estimated statistical distribution of the actual fabrication line. If the probability distributions of the measured device model parameters can be approximated by Gaussian distributions, then it is possible to use standard statistical techniques such as principle components analysis [15]–[18] to extract the eigenvalues

and eigenvectors of the covariance matrix between the instance model parameters. Straightforward null hypothesis tests can be used to determine how many principle components (eigenvalues) must be used in order that the generated covariance matrix is indistinguishable from the measured covariance matrix with any desired confidence level [15]. Just as for the method of direct measurement, in order that the generated distribution of design instances match that of the actual fabrication process, the number and distribution of test cells used as the basis for the principle components analysis on each wafer should correspond to the number and distribution of die on each wafer. Because generating random design instances from principle components requires little computational time⁴ it can be used to generate a library or it can be placed in the Monte-Carlo loop, unlike either of the first two methods of generating design instances.

In general, far fewer orthogonal components are needed to generate the device-model instances when using either a process simulator or principle components because the variations in device parameters are strongly correlated [19], [20]. For example, as will be seen in Section V, when using this method to generate BSIM device models with approximately 60 variable parameters, only 11 principle components were needed.

This leaves unaddressed an extremely important part of design instance generation, from the perspective of the analog circuit designer: predicting the mismatch of model parameters between devices on the same die. The performance of many analog circuits depends critically on circuit element matching and it has been the topic of much research (e.g., [21]–[24]). Device mismatch can be modeled by creating a unique set of device model parameters for every device in the circuit rather than a single set of device model parameters for each die as proposed for modeling die-level disturbances. Essentially this can be viewed as a two-stage process. First, a single die-level set of device model parameters is generated and then a model for the distribution of random intra-die device mismatch can be used to generate the deviations of the individual device model parameters from the die-level ones.

FABRICS II [5], a statistical process simulator, is able to generate the required sets of device model parameters for each device in each design instance. However, it assumes that all matching devices are drawn from a single distribution, regardless of whether the devices were in intimate proximity or on opposite sides of the die, which is in contradiction with experimental observations [21]–[24]. Accurate predictions of device mismatch cannot be performed given only the sized transistor schematic; in addition, information about the IC mask geometry is required as well. Because requiring complete mask geometry before yield prediction violates the spirit of using rapid yield estimation to guide early design decisions, a simplified conceptual model that allows the salient features of the mask geometry to be derived from the device sizes and some user-supplied choices about the style of device geometry and position has been developed.

The simple model for intra-die variations of device parameters proposed by Pelgrom *et al.* [21] is employed in this

³This assumes that the device models do not depend on device parameters like W and L for an MOS device, or if they do that a complete set of models that spans all possible device parameters is generated for each design instance.

⁴All it involves is a weighted sum of random numbers that are generated in the principle component space, followed by a transformation into the device model parameter space.

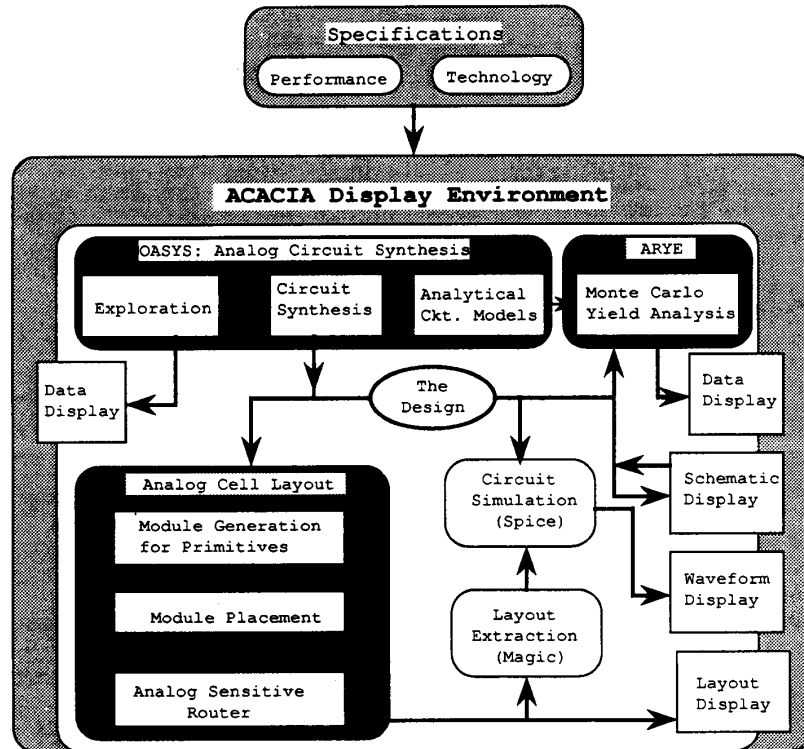


Fig. 5. Block diagram of ACACIA analog design system.

paper. This model includes one disturbance component that models parameter variations which are uncorrelated at lengths equal to transistor dimensions, e.g., the distribution of ion-implanted substrate charges. A second component models the variations with low spatial frequency that are typically found in experimental measurements, e.g., the gradual variation of oxide thickness with position on the die. These two sources of parameter fluctuations must then be convolved with the spatial distributions of the devices themselves. For example, the variance of a device parameter P between two simple rectangular devices is given by

$$\sigma^2(\Delta P) = \frac{A_p^2}{WL} + S_p^2 D_x^2$$

where W and L are the length and width of each device, respectively, and D_p is the center-to-center distance between the devices. A_p and S_p are the power of the variations for the broad-band spatial frequency disturbance source and the low spatial frequency disturbance source, respectively.

In order to compute the distribution of device parameter variations given only device sizes, it is necessary to make assumptions about how the matching devices will be placed on the die. The first assumption is that matching devices will be always be placed as close together as allowed by design rules. Second, only a limited set of device placement styles will be allowed (e.g., simple, one fold, two folds, etc.). Finally, only a limited set of device combinations will be allowed (e.g., side-by-side and cross-coupled quad). The designer specifies which device form and device combination

will be used for each set of matching devices. Using this information and the design rules for the fabrication process, it is possible to generate the actual device-to-device variance for each device model parameter. For example, for a side-by-side pair of simple devices which have their sources tied together,

$$D_x = L + 2D_{\text{diff}} + D_{\text{cont}}$$

where W and L are the given device sizes and D_{diff} is spacing between diffusion contacts and the gate and D_{cont} is the width of a minimum metal-diffusion contact.

IV. YIELD PREDICTION WITHIN ACACIA

ACACIA is an experimental design framework developed at Carnegie Mellon University with the aim of providing a set of computer-aided design (CAD) tools that facilitate the analog circuit design process [25]. ACACIA can generate IC mask geometry for frequently used analog modules (e.g., op amps), starting from performance specifications at nominal device model parameter values. As can be seen in the ACACIA block diagram (see Fig. 5), there are four major components: 1) OASYS, including the analytical circuit model (ACM) which generates sized transistor schematics; 2) KOAN/ANAGRAM II, which transforms sized schematics to mask geometries; 3) ARYE, a tool that implements the rapid yield estimation method presented in this paper; and 4) a graphical interface that enables the analog circuit designer to explore various design trade-offs. Graphic displays showing the yield statistics give the designer the necessary information to adjust the design to improve yield. A screen

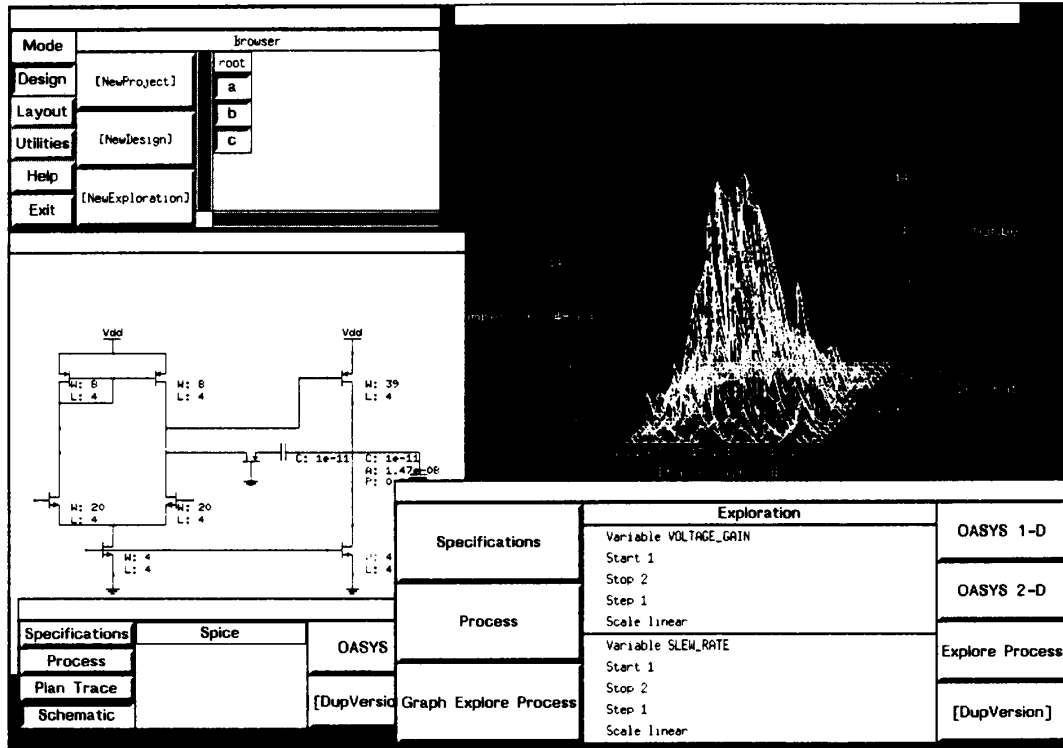


Fig. 6. ACACIA designer interface for yield exploration.

showing the interactive interface between ACACIA and the analog designer performing yield enhancement is shown in Fig. 6

Automated synthesis in ACACIA is performed by OASYS [8], [9] which takes specifications (e.g., for an op amp these would be gain, unity-gain frequency, etc.), for the selected analog circuit blocks (currently, OASYS implements op amps and comparators) using nominal device model parameters for the process in which the circuit will be implemented. In OASYS, as in several other recent analog synthesis methodologies (e.g., [7], [11]), synthesis is based on approximate analysis equations that predict the performance of a fixed topology in terms of the design parameters (e.g., transistor sizes and operating points) and device model parameters. However, one unique feature of OASYS is its use of fine-grained hierarchy, which allows it to have an equation model for current mirrors and differential pairs that can be reused and can itself contain multiple styles (e.g., simple or cascoded mirror).

The ACM part of OASYS includes a hierarchical set of analysis equations that predicts the performance of a sized circuit topology. Integrating the ARYE methodology into the ACACIA framework makes it possible to use the ACM's to provide a dual use of these analysis equations. As new synthesis targets are added, predicting their yield with ARYE will be straightforward. It is also possible to add analysis equations that predict performance for circuits for which OASYS does not have complete synthesis equations. This is desirable because the analysis equations that predict performance are only a subset of the equations used to perform

design. ARYE currently operates on all 128 possible two-stage op-amp topologies. The yield of any op-amp circuit, provided it matches one of the 128 possible topologies, can be predicted by ARYE. ARYE automatically determines the correct topology from the SPICE description of the circuit which can be provided by the user or automatically generated by OASYS.

This implementation of ARYE uses two alternate methods for generating the library of design instances. The FABRICS II process simulator [5] was used to generate a library of device instances for SPICE Level 2 device model parameters, assuming the process flow for a hypothetical 2- μm CMOS process. Intra-die variation between devices was modeled by assuming that all matching devices are side-by-side pairs of simple devices. The second method for generating the library of design instances used a principle components analysis to extract principle components for the BSIM short-channel MOS model for the Mosis/Orbit 2- μm p-well CMOS process. These principle components were used to generate a library of design instances.

V. EXAMPLE OF RAPID YIELD PREDICTION

In this section, a brief example of how ARYE can aid the designer of a two-stage op amp will be presented. In this example OASYS was first used to automatically select the specific two-stage op-amp topology and device sizes to meet a set of desired performance specifications. ARYE was then invoked to rapidly estimate the yield of the op-amp circuit. In this case, the estimated parametric yield was only 48.6%, based on 1000 design instances. The computation time re-

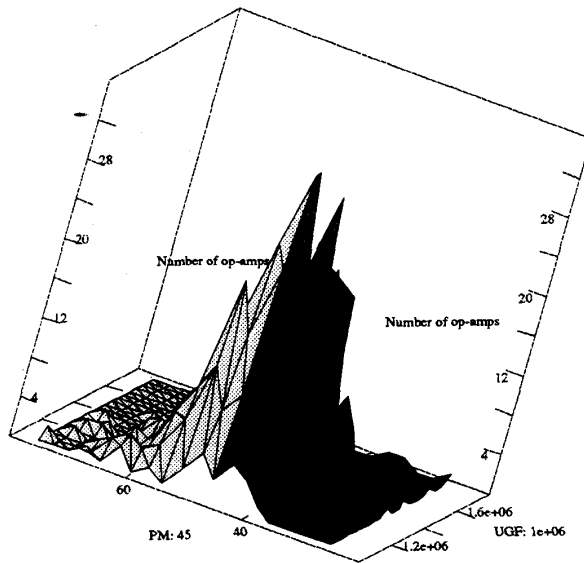


Fig. 7. Yield surface for an op-amp design A.




Fig. 8. Yield surface for an op-amp design B.

quired was 41.5 s on an approximately 13 MIP workstation. The specification that was violated most often was the desired phase margin of 45°. Fig. 7 shows the yield surface for the phase margin and the unity-gain frequency (UGF) of the initial design. The z axis of the graph indicates the number of instances that fell into the bin with the UGF and phase margin indicated by the x and y axes. The dark regions on the graph indicate designs that failed to meet the required performance specifications.

Yield enhancement can proceed in two possible ways. Either the analog designer can directly suggest changes in individual device sizes, or he can make changes in the requested performance specifications to OASYS and rerun the synthesis. In this case, in looking at Fig. 7 it is clear that while the phase margin is being frequently violated, the UGF values are far from their lower limit specification, 1 MHz. Since decreasing the UGF (which, for example, can be achieved by increasing the compensation capacitor) also increases the phase margin, a balance between yield loss due to phase margin and due to UGF can be struck. Fig. 8 illustrates the effect of a 5% increase in the compensation capacitor. The yield in this case increased to an estimated 89%.

The effect of device mismatch on op-amp performance can also be determined given the assumptions stated in Section IV about how matching devices will be positioned on the IC. Fig. 9 shows a histogram of the input offset voltage of a two-stage op amp with input devices that have a W/L of $20 \mu\text{m}/4 \mu\text{m}$.

VI. CONCLUSIONS

In this paper a new rapid yield estimation strategy that uses analysis equations to predict performance was presented. Although approximate, analysis equations require little computational time compared to circuit simulation, making rapid parametric yield estimation attractive. This rapid yield estimation strategy is particularly well suited for

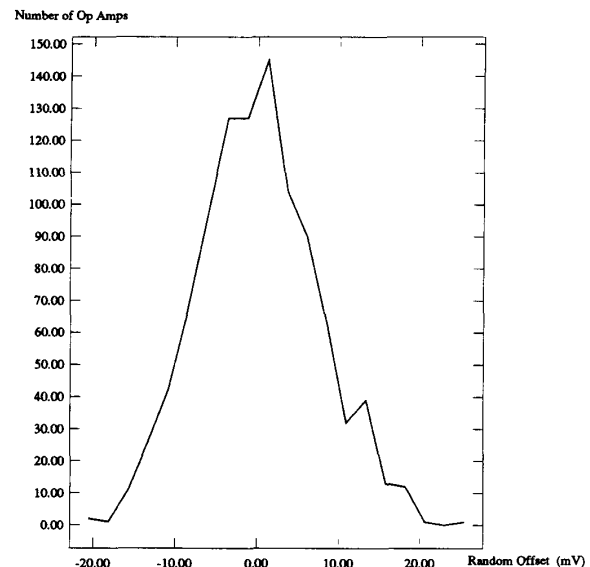



Fig. 9. Histogram of the op-amp input offset voltage.

used with synthesis systems that use analysis equations (e.g., OASYS). The availability of a rapid yield estimation CAD tool, by making numerous early predictions of parametric yield economical, has the potential to significantly improve analog circuit designer productivity.

REFERENCES

- [1] W. Maly, A. J. Strojwas, and S. W. Director, "VLSI yield prediction and estimation: A unified framework," *IEEE Trans. Computer-Aided Design*, vol. CAD-5, no. 1, pp. 114–130, Jan. 1986.
- [2] W. Maly, "Computer-aided design for VLSI circuit manufacturability," *Proc. IEEE*, vol. 78, no. 2, pp. 356–392, Feb. 1990.

- [3] T. Mukherjee, "Yield prediction of operational amplifiers within the ACACIA framework," Master's thesis, Carnegie Mellon Univ., Pittsburgh, PA, Mar. 1990.
- [4] W. Maly and A. J. Strojwas, "Statistical simulation of the IC manufacturing process," *IEEE Trans. Computer-Aided Design*, vol. CAD-1, no. 3, pp. 120-131, July 1982.
- [5] S. R. Nassif, A. J. Strojwas, and S. W. Director, "FABRICS II: A statistically based IC fabrication process simulator," *IEEE Trans. Computer-Aided Design*, vol. CAD-3, no. 1, pp. 20-46, Jan. 1984.
- [6] C. J. B. Spanos and S. W. Director, "Parameter extraction for statistical IC process characterization," *IEEE Trans. Computer-Aided Design*, vol. CAD-5, no. 1, pp. 66-78, Jan. 1986.
- [7] M. G. R. DeGrauwe *et al.*, "IDAC: An interactive design tool for analog CMOS circuits," *IEEE J. Solid-State Circuits*, vol. SC-22, no. 6, pp. 1106-1116, Dec. 1987.
- [8] R. Harjani, R. A. Rutenbar, and L. R. Carley, "OASYS: A framework for analog circuit synthesis," *IEEE Trans. Computer-Aided Design*, vol. 8, no. 12, pp. 1247-1266, Dec. 1989.
- [9] R. Harjani, R. A. Rutenbar, and L. R. Carley, "A prototype framework for knowledge-based analog circuit synthesis," in *Proc. 24th ACM/IEEE Design Auto. Conf.*, June 1987, pp. 42-49.
- [10] M. Ismail El Nagar, Ed., *Automated Design of Operational Amplifiers: A Case Study*. Boston MA: Kluwer, 1990.
- [11] H. Y. Koh, C. H. Sequin, and P. R. Gray, "Automatic synthesis of operational amplifiers based on analytic circuit models," in *Proc. 1987 IEEE Int. Conf. CAD*, Nov. 1987, pp. 502-505.
- [12] B. J. Sheu, J. H. Shieh, and M. Patil, "BSIM: Berkeley short-channel IGFET model for MOS transistors," *IEEE J. Solid-State Circuits*, vol. SC-22, no. 4, pp. 558-566, Aug. 1987.
- [13] *HSPICE User's Manual*, Meta-Software Inc., Campbell, CA, 1987.
- [14] G. G. E. Gielen, J. C. C. Walscharts, and W. M. C. Sansen, "ISAAC: A symbolic simulator for analog integrated circuits," *IEEE J. Solid-State Circuits*, vol. 24, no. 6, pp. 1587-1597, Dec. 1989.
- [15] W. Maly and S. W. Director, "Dimension reduction procedure for simplicial approximation approach to design centering," *Proc. Inst. Elec. Eng.*, vol. 127, no. 6, pp. 255-259, Dec. 1981.
- [16] *SAS User's Guide: Statistics*, SAS Tech. Inst., SAS Institute Inc., Cary, NC, 1985.
- [17] D. N. Lawley and A. E. Maxwell, *Factor Analysis as a Statistical Method*. New York: Elsevier, 1971.
- [18] J. D. Finn, *A General Model for Multivariate Analysis*. New York: Holt, Rinehart and Winston, 1974.
- [19] W. Maly and S. W. Director, "Dimension reduction procedure for the simplicial approximation approach to design centering," *Proc. Inst. Elec. Eng.*, vol. 127, no. 6, pp. 255-259, Dec. 1980.
- [20] R. W. Dutton, D. A. Divekar, A. G. Gonzalez, S. E. Hansen, and D. A. Antoniadis, "Correlation of fabrication process and electrical device parameter variations," *IEEE Trans. Computer-Aided Design*, vol. CAD-5, no. 1, pp. 114-130, Jan. 1986.
- [21] M. J. M. Pelgrom, A. C. J. Duinmaijer, and A. P. G. Welbers, "Matching properties of MOS transistors," *IEEE J. Solid-State Circuits*, vol. 24, no. 5, pp. 1433-1439, Oct. 1989.
- [22] J. B. Shyu, G. C. Temes, and K. Yao, "Random errors in MOS capacitors," *IEEE J. Solid-State Circuits*, vol. SC-17, pp. 1070-1075, 1982.
- [23] J. B. Shyu, G. C. Temes, and F. Krummenacher, "Random error effects in matched MOS capacitors and current sources," *IEEE J. Solid-State Circuits*, vol. SC-19, pp. 948-955, 1984.
- [24] K. R. Lakshmikummar, R. A. Hadaway, and M. A. Copeland, "Characterization and modeling of mismatch in MOS transistors for precision analog design," *IEEE J. Solid-State Circuits*, vol. SC-21, no. 6, pp. 1057-1066, Dec. 1986.
- [25] L. R. Carley *et al.*, "ACACIA: The CMU analog design system," in *Proc. IEEE Custom IC Conf.*, May 1989, pp. 4.3/1-5.



Tamal Mukherjee (S'86) was born in New Delhi, India, on February 1, 1967. He received the B.S. and M.S. degrees in electrical engineering from Carnegie Mellon University, Pittsburgh, PA, in 1987 and 1990, respectively. He is currently a Ph.D. candidate in the Department of Electrical Engineering at Carnegie Mellon University.

He has worked as a summer intern engineer for three summers at the IBM T. J. Watson Research Center, Yorktown Heights, NY. His research interests are in the area of automated analog circuit synthesis, and yield prediction and optimization of analog circuits.



L. Richard Carley (S'77-M'84-SM'90) received the S.B. degree from the Massachusetts Institute of Technology (MIT), Cambridge, in 1976 and was awarded the Guillemin Prize for the best EE undergraduate thesis. He remained at MIT where he received the M.S. degree in 1978 and the Ph.D. degree in 1984.

He has worked for MIT's Lincoln Laboratories and has acted as a consultant in the area of analog circuit design and design automation for Analog Devices and Hughes Aircraft among others. In 1984 he joined Carnegie Mellon University, Pittsburgh PA, where he is currently an Associate Professor of Electrical and Computer Engineering. His research is in the area of analysis, design, automatic synthesis, and simulation of mixed analog/digital systems.

Dr. Carley received a National Science Foundation Presidential Young Investigator Award in 1985, and a Best Paper Award at the 1987 Design Automation Conference.