Design automation methodology and rf/analog modeling for rf CMOS and SiGe BiCMOS technologies

D. L. Harame
K. M. Newton
R. Singh
S. L. Sweeney
S. E. Strang
J. B. Johnson
S. M. Parker
C. E. Dickey
M. Erturk
G. J. Schulberg
D. L. Jordan
D. C. Sheridan
M. P. Keene
J. Boquet
R. A. Groves
M. Kumar
D. A. Herman, Jr.
B. S. Meyerson

The rapidly expanding telecommunications market has led to a need for advanced rf integrated circuits. Complex rf- and mixed-signal system-on-chip designs require accurate prediction early in the design schedule, and time-to-market pressures dictate that design iterations be kept to a minimum. Signal integrity is seen as a key issue in typical applications, requiring very accurate interconnect transmission-line modeling and RLC extraction of parasitic effects. To enable this, IBM has in place a mature project infrastructure consisting of predictive device models, complete rf characterization, statistical and scalable compact models that are hardware-verified, and a robust design automation environment. Finally, the unit and integration testing of all of these components is performed thoroughly. This paper describes each of these aspects and provides an overview of associated development work.

1. Introduction

The rf CMOS and SiGe BiCMOS process technologies are at the leading edge of today's rapidly expanding telecommunications marketplace. Two key application areas in telecommunications are the wireless and wired areas. Examples in the wireless area include cellular telephone radios with protocols such as GSM (Global System for Mobile Communications [1]) and WCDMA (Wideband–Code Division Multiple Access [2]), location systems such as GPS (Global Positioning Satellite System), and wireless connectivity applications such as Bluetooth (2.4-GHz low power connectivity standard [3]) and 802.11x (IEEE wireless LAN standards [4]). Examples of the wired area are synchronous data transmission over optical networks using various protocols such as SONET (synchronous optical network transmission standard [5]) and SDH (synchronous digital hierarchy [6]). Both of these areas require the use of state-of-the-art rf/mixed-signal process technologies and design automation environments.

There is some differentiation between the wired and wireless requirements. The requirements for monolithic wireless chips emphasize technology with superior high-Q passives (inductors, varactors, and capacitors) in addition to the active devices, compact models with accurate noise-figure and distortion analyses, signal integrity analysis, and *RLC* parasitic extraction. The requirements for monolithic wired chips emphasize process technologies with very-high-speed active devices, rf/analog models of all devices—particularly FET devices, and distributed interconnect models including transmission-line models and field-solver solutions.

As integrated circuit (IC) design becomes more complex and application frequencies continue to rise, the different points in the supply chain must be highly integrated for final product success. The components necessary to successfully enable a silicon chip design are illustrated in **Figure 1**. There has to be a silicon technology base, a set of vendor CAD tools, and both modeling and design automation activities.

©**Copyright** 2003 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/03/$5.00 © 2003 IBM

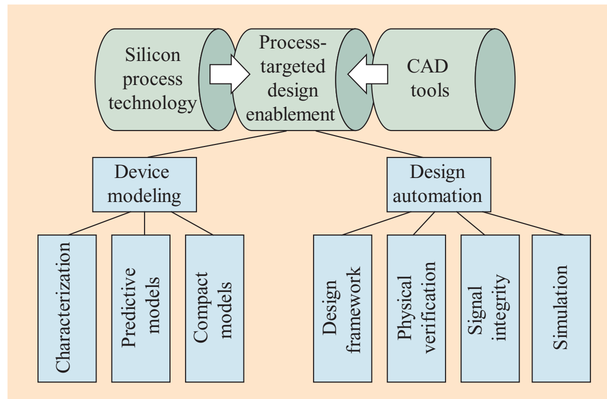

Figure 1

Components required to successfully enable a silicon chip design.

The development of this rf/analog mixed-signal methodology at IBM dates back to the support for the early bipolar technology used in bipolar-based mainframes and, more recently, as an outgrowth of the analog BiCMOS processes initially developed for magnetoresistive (MR) preamplifier applications in the early 1990s. For bipolar mainframes, IBM developed internal modeling and circuit simulation tools[1] [7] with an efficient Monte Carlo statistical simulation package. Today, high-performance analog applications utilize the entire BiCMOS device menu and require accurate analog models for all devices. Digital models used to support the standard digital CMOS technologies do not provide sufficient accuracy in predicting the device characteristics in these regimes. In addition, analog circuit designers requested scalable bipolar junction transistor (BJT) models to remove the limitations imposed by an npn device library. These product design requirements defined the direction for the development of more advanced analog models.

In 1990, IBM began using workstation-based OEM design automation tools. In 1992, IBM released the first BiCMOS Cadence**-based[2] design kits (starting with the design kit for CBiCMOS[3]), which included model libraries, symbol libraries, model/layout call-back routines, SKILL** routines (SKILL is the Cadence application extension language), parameterized cells for layout, layout versus schematic checking, design-rule checking, parasitic extraction, and custom graphical user interface tools. This early kit contained all of the basic elements found in today's design kits. In 1994, work began on the development of complex partial response maximum likelihood (PRML) read channel chips, which have a large digital as well as analog content. This required the development of new methodologies to handle analog mixed-signal designs. Work has continued on these tools over the past ten years, leading to a very flexible and robust design kit.

This paper describes the device-level rf/analog enablement methodology practiced at IBM. This enablement methodology is roughly divided into three major sections. The first section focuses on the co-development of the technology using a robust methodology that starts with technology concepts and evolves to early compact models (predictive modeling). Building and calibrating this methodology is a complex process, involving much greater effort than the simple use of vendor tools. The second section describes a characterization and modeling methodology that focuses on a strong rf/analog infrastructure with a statistical emphasis in order to fully simulate the product in manufacturing and predict performance yields (characterization and modeling). Finally, the third section discusses a detailed integrated design automation methodology/framework that includes physical design and verification (with inclusion of complete rf requirements), accurate signal integrity analysis, rf and mixed-signal simulation, and final test verification.

2. Predictive modeling

A fundamental difference between digital technology development and analog/rf technology development is the sensitivity of analog/rf circuits to many manufacturability/performance tradeoffs that must be made. Thus, while feedback between IC designers and technologists during technology definition is critical for timely product development, it is difficult to realize in practice, because to an IC designer the technology is the design kit—a complete and accurate set of compact models for both active and passive devices as well as interconnects. It is expected to reflect a stable and well-characterized process, while to the technologist the technology is a process recipe that is roughly characterized, often in flux, and, in practice, often not actually realized until the final technology qualification stage. The role of predictive modeling is to utilize detailed process and device simulation, or technology CAD (TCAD), in place of hardware to facilitate the feedback loop between circuit designer and process technologist until definitive hardware data is finally incorporated into the design kit. This feedback path provides timely notification to the technologist of potential shortcomings in the targeted technology design point, thus optimizing the use of the available experimental wafer budget. For the circuit designer, it provides more design turns with the

[1] A Statistical Analysis Program (ASTAP), an IBM internal circuit simulation tool with compiled model interface.
[2] Cadence Design Systems, Inc., San Jose, CA.
[3] S. Strang, IBM Microelectronics Division, Essex Junction, VT, private communication.

technology and thus a greater likelihood of meeting circuit performance targets.

Technology CAD

Semiconductor TCAD originated in the early 1960s [7] with efforts to understand and optimize bipolar transistors. This effort continues today, with the increases in computing power available leveraged to understand and engineer devices with higher operating speeds and fabricated with more complex processes. The TCAD paradigm is applied to all conceivable types of active and passive devices, and intensive TCAD studies are now part of all semiconductor technology development efforts.

The TCAD paradigm can be described as follows: Detailed process simulation creates one-, two-, or three-dimensional device representations, consisting of structural (film thicknesses and shapes) and impurity concentrations used as input for device simulation. Device simulation produces the dc and ac characteristics of interest, which are in turn used to define compact models for use in a prototype design kit for the technology. Thus, process options can receive circuit-design feedback before expending the time and budget to define these options in silicon. While TCAD is typically used to assist technology development, it is leveraged to its fullest extent when combined with compact model development to provide early technology access for circuit design.

Process simulation

Physical process simulation is the critical component in a predictive TCAD capability. Research and development of existing process simulation capabilities are due to the last decade's worth of logic CMOS scaling. Ever more sophisticated process simulation capabilities are being developed as semiconductor processing capabilities, driven by an extremely competitive microelectronics industry, continually progress. However, despite intensive efforts to bring higher-level modeling capabilities such as molecular dynamics and kinetic Monte Carlo codes into practical use, continuum codes based on silicon process physics are still the primary platform for semiconductor process simulation and are focused on here. The critical silicon process operations are impurity implantation and diffusion, oxidation, and material deposition and etching. Silicon and silicon–germanium epitaxy, an increasingly critical silicon process step, is treated via a series of material depositions and diffusions.

Despite the intensive effort to understand, characterize, and model silicon process physics driven by the CMOS logic microelectronics industry, industrial use still requires continual calibration of model parameters, and the predictive range of any process modeling capability can vary significantly between process modules. The most effective approach to TCAD calibration begins with optimizing essential process simulation models, such as impurity diffusion coefficients, implant models, and film thicknesses, to available physical data such as that obtained from structural cross sections and secondary ion mass spectrometry (SIMS). This initial calibration provides a good starting point for multidimensional process and device simulations, the goal of which is to fit electrical data considered critical for the target use of the simulation capability. A feedback loop is then established and exercised, in which a few critical process and device model parameters are optimized to fit the target electrical data. A successful calibration effort provides a self-consistent TCAD capability that links process levers to electrical characteristics through simulation. This has been found to be a very powerful tool to provide the ability to explore design space and optimize device performance outside the lengthy and expensive silicon fabrication process.

Device simulation

Much like process simulation, device simulation has been implemented at several levels of physical sophistication. The approach used in the pioneering continuum-based device simulation efforts [7–9] continues to be the foundation of most silicon technology CAD efforts. This is particularly true in the predictive modeling mode, which for compact model generation demands simulation of the complete device structure, including all parasitic capacitances and resistances, and often represents a considerable expanse of silicon, over a wide range of bias conditions.

Continuum-based device simulation consists of the solution of Poisson's equation along with two or more equations accounting for the transport, generation, and recombination of holes and electrons in the semiconductor. Additional equations can be added to account for carrier energy exchange with the silicon lattice, generating average carrier temperatures, and thus address nonstationary transport effects such as velocity overshoot and device lattice self-heating. A further enhancement for simulation of field-effect devices is the incorporation of some form of Schrödinger's equation to account for carrier quantization effects. An immense amount of theoretical and experimental work has been done to formulate and calibrate the many physical models required in a typical continuum device simulation, such as carrier mobility as a function of electric field, doping concentration, lattice temperature, and surface roughness, and effects of impurity concentration and species on the silicon bandgap [10].

Continuum-based device simulation is able to duplicate experimental ac measurements, such as scattering-parameter (S-parameter) extraction, by application of the ac small-signal approximation to the Poisson and transport

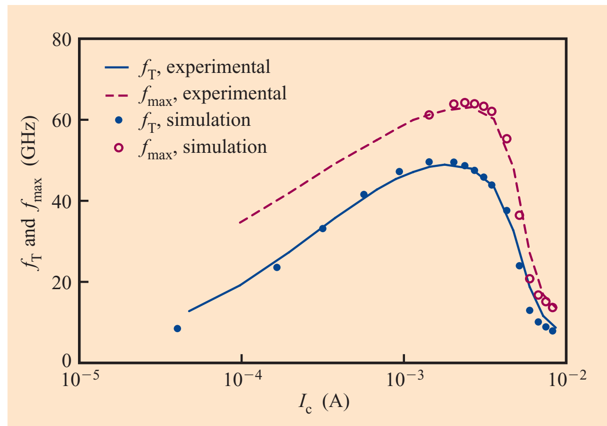

Figure 2

Comparison of experimental and simulated f_T and f_{max} for 50-GHz/60-GHz f_T/f_{max} SiGe HBT. The device parameters depicted here and in Figure 3, critical levers for rf performance optimization, were extracted from device-level ac measurements and small-signal ac simulations of two-dimensional device cross sections, generated by detailed process simulation and appropriately scaled in length.

equations [11]. Implementation of the small-signal approximation completes the demand on the TCAD capability of supplying all necessary electrical data for building a compact model. When device simulation is combined with process simulation and suitably calibrated, the system so formed provides the technology development effort the ability to quickly, and with a useful range of self-consistent physical accuracy, connect process levers to all critical device electrical characteristics.

Semiconductor device scaling has driven future-generation devices into a physical regime in which typical operation is dominated by physical transport effects that are not well addressed by enhanced continuum approaches. A greater amount of physical accuracy and detail is provided by Monte Carlo particle transport codes [12], albeit on (typically) reduced computational regions and requiring significantly longer solution times.

Examples

An example of TCAD application in the technology development mode follows. A SiGe heterojunction bipolar transistor (HBT) with ac performance characterized by f_T = 50 GHz and f_{max} = 60 GHz was assumed to be the technology starting point. The objective of the analysis was to size the ac-performance consequences of simplifying the extrinsic base formation. It began by calibrating the process and device simulation to the technology. The results of the concurrent two-dimensional process and device calibration are shown in **Figure 2**. The figure compares the experimental and simulated summary ac parametrics of cutoff frequency f_T and maximum frequency of oscillation f_{max}. Both parameters were well matched by their respective simulations, and they involved all aspects of the device performance—transport in both low- and high-doped and depletion regions, across signal frequencies and bias conditions—a strong indication that the ac performance of the device was captured with an adequate degree of accuracy. **Figure 3(a)** shows the experimental and simulated base resistance R_b for the target device. The base resistance had a strong three-dimensional component that could not be fully captured by the two-dimensional simulations; however, the simulation showed the same qualitative characteristics and was within 10 to 15% of the experimental values at the current densities at which peak ac performance (f_T) was observed. The collector base capacitance C_{bc} [**Figure 3(b)**] could be modeled by the two-dimensional approach. Both simulation and experimental parameters were extracted from the low-frequency Z parameters. From this

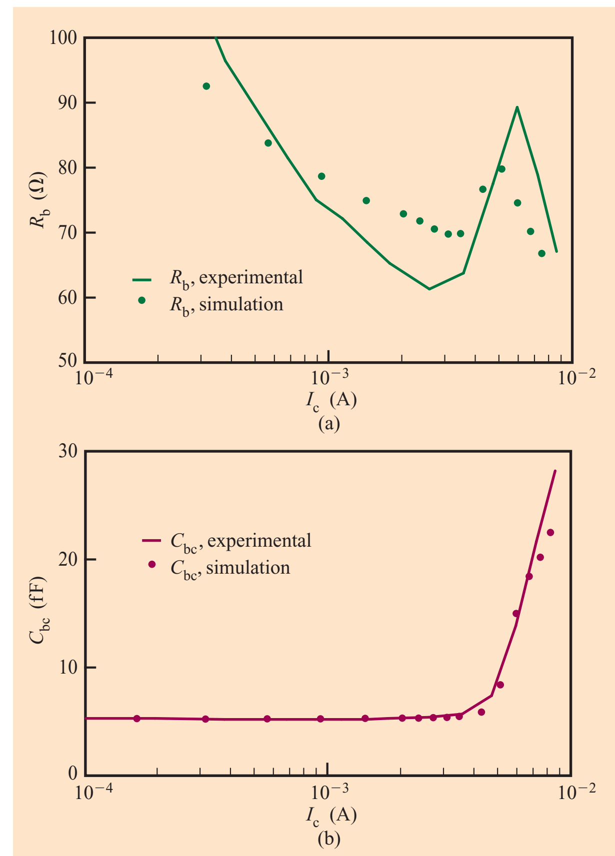

Figure 3

Comparison of simulated and experimental R_b (a) and C_{bc} (b) for 50-GHz/60-GHz f_T/f_{max} SiGe HBT.

calibrated starting point, the consequences for the device ac performance due to a simplified, non-self-aligned (NSA) process for the extrinsic base were investigated.

Figure 4(a) depicts the process used to form the extrinsic base in the device of Figures 2 and 3. In this self-aligned process, a pedestal structure is formed to define the future position of the emitter polysilicon, and it is surrounded by a sidewall spacer. The extrinsic base is then implanted, with the spacer controlling the separation, independently of the emitter width, between the emitter and extrinsic base. An alternative extrinsic-base formation method is shown in **Figure 4(b)**. In this NSA method, the extrinsic base implant occurs immediately after the emitter sidewall spacer is etched. This latter process is simpler and cheaper than the self-aligned process but necessarily results in a higher base resistance, since process tolerances demand that the distance between the emitter out-diffusion and the extrinsic base implant be larger. A more subtle issue is the effect of the extrinsic-base implant on the ac performance of the device. Both experimental results and understanding of silicon process physics suggest that point defects introduced by the extrinsic base implant can induce unwanted additional diffusion in the intrinsic base, thus degrading the ac performance of the transistor [13]. It is expected that the calibrated process and device simulation capability will accurately reflect that phenomenon and provide a method to size the many tradeoffs involved in this option.

Table 1 lists comparative values of several parameters. The simulation accurately predicted the increase in base resistance and simultaneous reduction in base-collector capacitance arising from the alternate extrinsic-base process. Further, it accurately captured the improvement in cutoff frequency f_T resulting from the extrinsic base implant of the NSA process that differed in both dose and energy, as well as being spaced farther from the emitter-base junction. Finally, these modifications in device resistance, capacitance, and cutoff frequency were accurately reflected in the predicted maximum oscillation frequency f_{max}.

Another example of the TCAD predictive approach pertains to the 200-GHz SiGe HBT technology performance path [14]. On the basis of the previous-generation f_T = 120-GHz technology [15], exploration of process and device modifications was performed with TCAD to identify promising ultrahigh-performance SiGe HBT design points [16]. In particular, significant modifications to the extrinsic base process module and epitaxial-base transistor structure were suggested. Because the new process modules required significant process development effort, test structures based on a simple NSA structure were used to prototype the vertical device structure. Initial samples showed promising ac performance results for an f_T = 200-GHz SiGe HBT.

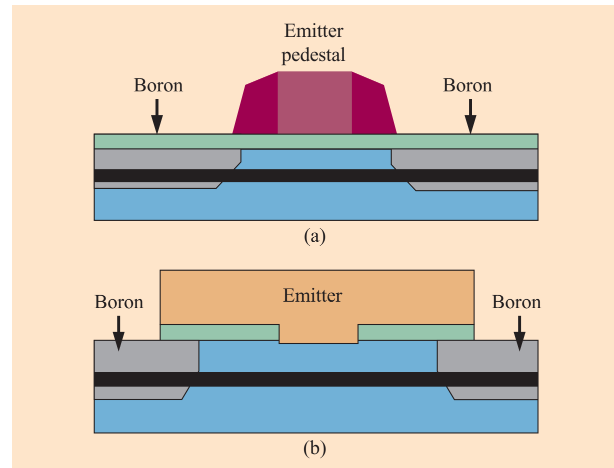

Figure 4

(a) Formation of a self-aligned extrinsic base for a 50-GHz/60-GHz f_T/f_{max} SiGe HBT. The emitter pedestal defines where the polysilicon emitter is to be formed later in the process. The spacers on either side of the emitter pedestal provide control over the tradeoff between reduced extrinsic-base resistance and deleterious effects of the point defects introduced by the extrinsic base implant. (b) Alternate extrinsic-base formation for 50-GHz/60-GHz f_T/f_{max} SiGe HBT. In this process, the extrinsic base implant is performed after the polysilicon emitter has been deposited and etched. This is simpler and cheaper than the self-aligned process depicted in (a), but is expected to result in significant degradation of the extrinsic-base resistance. How this base-resistance degradation would interact with possible changes in device capacitances and reductions in point-defect-enhanced diffusion in the intrinsic base to modify the critical ac performance metrics of f_T and f_{max} was the question addressed by the TCAD simulations. A comparison of simulated and experimental results is shown in Table 1.

Table 1 Comparison of TCAD evaluation of ac-performance consequences of non-self-aligned extrinsic base process.

Parameter	*Experiment*	*Simulation*
C_{bc}	~0	~0
C_{be}	~0	~0
R_b	+125 Ω	+110 Ω
f_{max}	−21 GHz	−18 GHz
f_T	+7 GHz	+8 GHz

While process development continued apace, an early design kit based on calibrated TCAD was formulated. The process to produce this early NSA hardware was used to calibrate the process and device simulation. The complete target process was then simulated, and a scalable compact vertical bipolar inter-company (VBIC) [17] model was

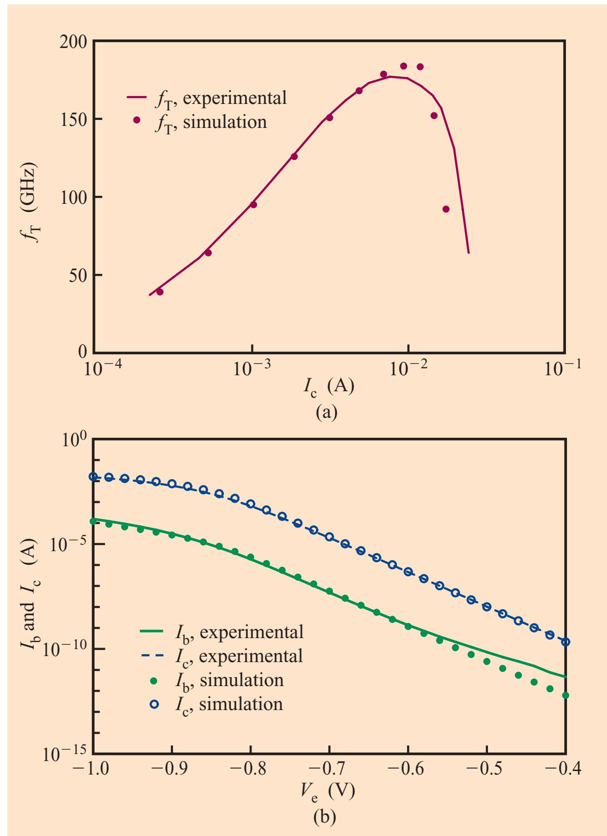

Figure 5

Non-self-aligned TCAD calibration in preparation for TCAD definition and model extraction for f_T = 200-GHz SiGe HBT device development. Process/device simulation and measured characteristics are compared. Parts (a) and (b) pertain respectively to ac and dc characterization.

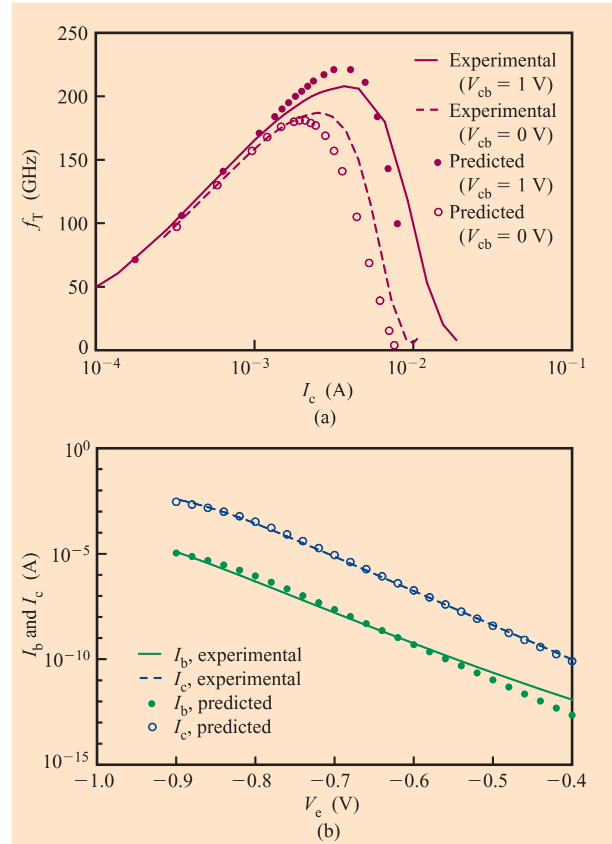

Figure 6

Comparison of (a) ac and (b) dc characteristics of f_T = 200-GHz SiGe HBT hardware and compact models based on TCAD simulations of target process produced nine months prior to hardware completion.

extracted from electrical simulations, following the methodology depicted in Figure 1. Basic ac and dc electrical results used in the calibration are shown in **Figure 5**. A comparison with the compact model characteristics arising from the complete process fabricated a year later is shown in **Figure 6**. Working ring-oscillator circuits, showing record silicon stage delays, were subsequently designed on the basis of the modeling [18].

Summary

Technology CAD is becoming an increasingly critical part of rf technology development. With careful calibration and recognition of the predictive range of the many models and assumptions that constitute the process and device simulation approaches relied upon for industrial technology development, TCAD can significantly improve the process learning and technology performance progress achieved with a given experimental budget. When leveraged to its fullest extent, TCAD can assume an important strategic role in rf product development by providing an efficient and inexpensive link between circuit designers and technologists, resulting in a tighter coupling between the technology and its target products and increasing the likelihood of meeting time-to-market goals. By providing early design kit information prior to the appearance of experimental hardware, more design turns are available to circuit designers, further ensuring the likelihood of timely product development.

3. Characterization

Characterization is an important step in the development of device compact models. Compact model development and parameter extraction involve both dc and ac data; the latter is usually in the form of two-port S-parameter measurements. Test sites are designed and fabricated to provide a full complement of device test structures to

allow for on-wafer measurements in support of model development. Test structures are designed to measure specific process and model parameters over a range of biases and temperatures.

Finally, it is essential to verify model functionality and circuit simulation accuracy. For specific applications, this may include low- and high-frequency noise measurements, large-signal measurements, and two-tone intermodulation distortion measurements.

In-line wafer measurements

A complete set of dc measurements is taken in line during wafer fabrication, and this provides an accurate method for determining the process parameters for each individual wafer. Data is collected for each wafer run through the fabrication line that then provides a large statistical database for many process and modeling parameters. The database parameters include inputs from device attributes, measurement data, and calculation formula. Electrical parameter data analysis benefits enormously from object-oriented techniques that make it easy to create comparison charts. A detailed description of the problems associated with collecting and storing in-line data and the software programs needed to analyze the data has been presented by Freeman [19] and is not reviewed here. This data is later used in the selection of modeling wafers to be slated for characterization. Since it is essential that process variations be fully determined, wafer characterization from the in-line data is an important first step in the data-acquisition and device-characterization process.

DC, C–V, and matching measurements

DC and capacitance measurements are taken on a full set of on-wafer test structures for each device type, and a statistical database is built. Numerous wafers from several lots are tested to generate a large sample size to allow for the employment of statistical modeling techniques. DC measurements are performed on devices to characterize and model device performance, whereas capacitance–voltage (C–V) measurements, in addition to allowing capacitors themselves to be characterized, are also performed to determine specific device model circuit parameters. For example, capacitance measurements on specially designed structures allow for the gate oxide and overlap capacitances to be determined and included in the circuit model for FET devices. Matching measurements are important to fully determine the variation of device parametrics between adjacent devices.

Measurements are taken on each type of device across operating conditions and at numerous bias conditions to fully represent all possible operating modes. The full set of test structures includes a complete range of device sizes. This allows geometry dependencies to be incorporated into the models. Data is also collected at a full range of temperatures (−50°C to +145°C) to allow device temperature dependencies to be fully determined.

At IBM, a program has been developed in object-oriented language, referred to as the Device Measurement And Characterization System (DMACS), which has been described in detail elsewhere [20]. Some key attributes of the software program are storage of large quantities of different kinds of data; storage of automated testing programs; the ability to drive a large variety of instruments under program control; display of the raw data in numerous types of plots and charts as well as manipulation of the data with various types of calculations; and finally, the ability to output data into a wide variety of formats so that it may be ported to other modeling software tools. The program in effect serves as an engineering development tool and also as a manufacturing characterization tool.

AC S-parameter measurements

Small-signal equivalent-circuit models have been developed to enable devices to be represented by lumped elements rather than complicated nonlinear equations. High-frequency two-port parameters such as S-parameters are measured to enable the development of these small-signal equivalent-circuit models and determine their associated model parameters. As with dc measurements, S-parameter measurements are taken on both active and passive devices, across a full range of both geometries and temperatures, and, for active devices, at different operating conditions and across numerous biases. For passives, S-parameters are also converted to Y-parameters (admittance parameters) and are used to determine the input and output characteristics as a function of frequency. In addition, S-parameter measurements are used to calculate two important figures of merit for transistors: the cutoff frequency of the ac-current gain, f_T, and the cutoff frequency of the maximum power gain, also called the maximum oscillation frequency, f_{max}.

Present applications demand measurement of device ac characteristics at frequencies as high as 110 GHz. Special steps must be taken at these frequencies to ensure data integrity and quality. The reference plane must be firmly established. In addition, at these frequencies, padset parasitics become an issue. A probe-tip calibration is carried out first in order to move the measurement reference plane to the probe tips. Then a two-step de-embedding procedure is used to subtract both the series and parallel parasitics of the padsets, thus obtaining the high-frequency characteristics of the device itself. This procedure is described in detail elsewhere [21]. The two-port S-parameters are measured using a vector network analyzer and ground–signal–ground (G–S–G) probes. For active devices, special care must be taken regarding the

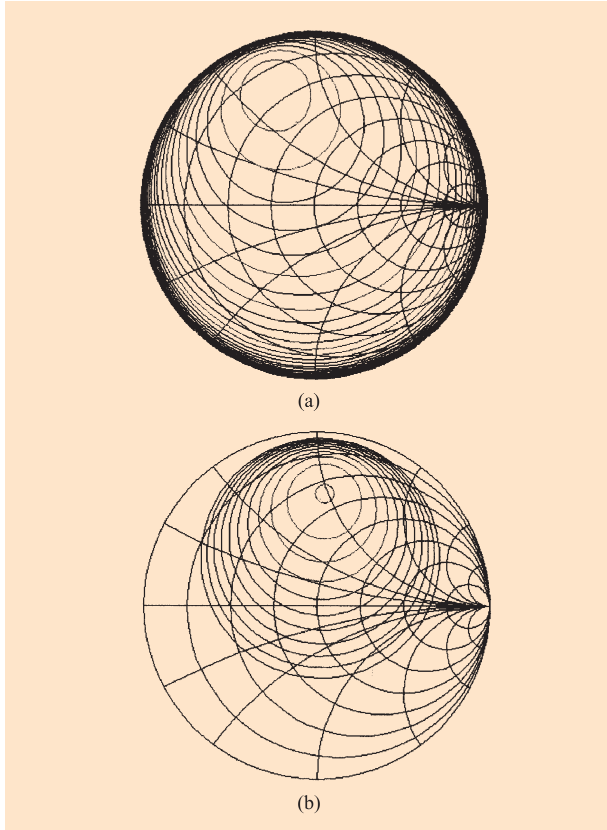

Figure 7

(a) Constant-gain circles vs. source impedance for an SiGe FET; contour start: 7 dB, contour step: 1 dB. (b) Constant-noise circles vs. source impedance for the same device; contour start: 5 dB, contour step: 0.5 dB. The circles illustrate the design tradeoffs that must be made between minimum noise and maximum gain.

power levels to prevent gain compression and to ensure that the device remains in the small-signal regime.

To calculate f_T, the S-parameters are converted to H-parameters and, in graphical terms, the ac current gain, $|H_{21}|$ (in dB), is plotted on a linear scale as a function of frequency on a log scale. The f_T of the transistor is the point at which $|H_{21}|$ crosses the x-axis. To facilitate the determination of f_T, the IBM instrument control and data acquisition software (DMACS) calculates this parameter in the following manner. The $|H_{21}|$ curve is assumed to have perfect single-pole, 20-dB/decade, roll-off characteristics. The f_T is calculated using the base transit time τ_T, where

$$\tau_T = \left| \frac{\sin \angle H_{21}}{|H_{21}|\omega} \right| . \tag{1}$$

Then f_T is given by

$$f_T = \frac{1}{2\pi\tau_T} . \tag{2}$$

It can be shown from circuit theory that these equations are valid at any frequency, though care must be taken to carry out the calculation at a high enough frequency to avoid inaccuracies due to instrumentation phase limitations. The maximum oscillation frequency, f_{max}, is determined as the frequency at which the maximum available power gain (MAG), also a quantity exhibiting single-pole transfer function characteristics, is unity, where

$$MAG = \left| \frac{S_{21}}{S_{12}} \right| (k - \sqrt{k^2 - 1}) . \tag{3}$$

The term k is Kurokawa's stability factor and is given [22] by

$$k = \frac{1 - |S_{11}|^2 - |S_{22}|^2 + |S_{11}S_{22} - S_{12}S_{21}|^2}{2|S_{12}||S_{21}|} . \tag{4}$$

Using the calculation of MAG across the frequency range, DMACS then uses a linear regression routine to determine the x-intercept and thus f_{max}. Again, care must be taken to ensure that an appropriate frequency range is used when performing the linear regression.

Noise characterization

Both low-frequency flicker noise ($1/f$ noise) and high-frequency, broadband noise parameters, including noise figure, associated gain, optimum reflection coefficient, and noise resistance, must be measured to facilitate the modeling efforts and the design of integrated telecommunication circuits in BiCMOS technologies. Flicker noise measurements are made to determine several subcircuit parameters in the modeling of the SiGe BiCMOS technologies. In addition to model verification, the broadband noise performance characterization gives the circuit designer a measure of the signal-to-noise-level degradation that will result when the device is added to the circuit, an important consideration in telecommunications systems, which typically process very low-level signals.

The associated noise figure describes the degradation of the signal-to-noise ratio between the input and output of the device. For a linear two-port device, the dependence of the noise figure on the source reflection coefficient is described by

$$F(\Gamma_s) = F_{min} + 4r_n \frac{|\Gamma_s - \Gamma_{opt}|^2}{(1 - |\Gamma_s|^2)|1 + \Gamma_{opt}|^2} , \tag{5}$$

where Γ_s is the source reflection coefficient, F_{min} is the minimum noise figure, r_n is the normalized noise resistance (the sensitivity of the noise figure to changes in

the source reflection coefficient), and Γ_{opt} is the optimum source reflection coefficient which gives the minimum noise figure. The equation represents a parabola, indicating that a value of source impedance can be determined for which a minimum noise figure is achieved. Thus, high-frequency, broadband noise characterization, together with the gain characterization of the active device, can provide the circuit designer with the information needed to determine what kind of tradeoff must be made when impedance-matching to the device for minimum noise figure vs. maximum gain. This tradeoff is illustrated in **Figure 7**. Figure 7(a) shows the constant-gain circles for an SiGe FET, and Figure 7(b) shows the corresponding constant-noise circles for the same device. The center of the smallest circle indicates the optimum source impedance for each parameter, and the spacing between the circles is an indication of how much the designer must give away when moving away from that optimum point.

Flicker ($1/f$) noise characterization is important because certain types of circuits in telecommunications systems are particularly sensitive to this low-frequency noise. Bipolar transistors and FETs exhibit very different levels of flicker noise. Therefore, it is important to characterize these devices when considering their use in circuits sensitive to such noise.

Large-signal measurements

Determining how active devices behave at different power levels is also an important consideration when designing telecommunications systems. Many rf amplifiers are designed to operate in the weakly nonlinear region, where power-added efficiency (*PAE*) peaks. The power-added efficiency is defined as the ratio of the additional power provided by the amplifier to the dc power [23], viz.,

$$PAE = \frac{P_{\mathrm{rf\,out}} - P_{\mathrm{rf\,in}}}{P_{\mathrm{dc}}}. \tag{6}$$

Large-signal measurements provide output power, gain, and efficiency information at a given input power level. **Figure 8** shows the results of large-signal measurements made on an SiGe HBT. For these measurements, both the input and output impedance were set to 50 Ω. The figure illustrates the power levels at which the device enters compression and the power-added efficiency in that region. The 1-dB gain compression point, usually given in terms of output power, is an important quantity when considering the dynamic range of the transistor.

In addition to power level considerations, impedance matching throughout the system is an essential aspect of rf circuit design. Thus, it is also important to explore how the input and output impedance presented to each device in the system affect the performance of that device in the circuit. Furthermore, a designer may want to determine the necessary impedance-matching conditions to achieve a specific desired performance from the device. Once either the desired input or output impedance is determined, contours can be generated for the other termination to illustrate tradeoffs that must be made between, for example, maximum power-added efficiency and maximum output power. **Figure 9(a)** shows the output power contours versus load termination for an SiGe transistor whose input impedance was conjugately matched for maximum gain. **Figure 9(b)** shows the power-added efficiency contours for the same device. Note that the optimum load termination for maximum output power and that for maximum power-added efficiency are near each other on the Smith chart, suggesting that a designer should be able to use this device in a circuit without a significant tradeoff between output power and power-added efficiency.

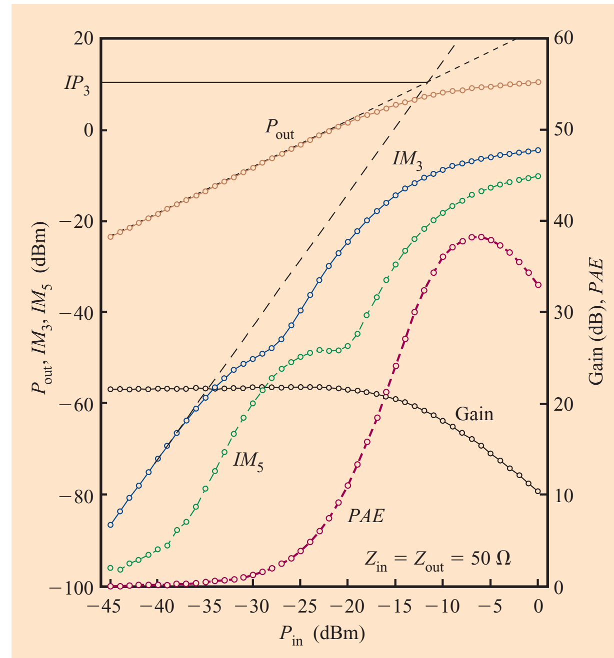

Figure 8

Measured large-signal and intermodulation distortion data for an SiGe HBT. The source and load terminations were set to 50 Ω and the determination of IP_3 is illustrated. P_{out} denotes the output power, IM_3 denotes the third-order intermodulation product, IM_5 denotes the fifth-order intermodulation product, *PAE* denotes the power-added efficiency, and P_{in} denotes the input power.

Distortion

Many components of telecommunication systems receive numerous signals closely spaced in frequency at their inputs. The nonlinearities inherent in all active devices

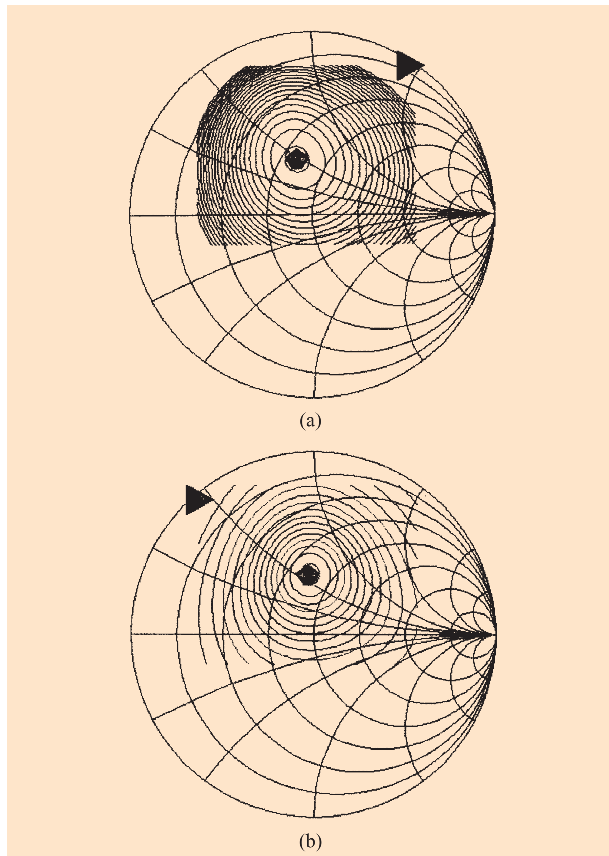

Figure 9

(a) Output power (dBm) and (b) power-added efficiency contours vs. load termination for an SiGe transistor. Frequency = 900 MHz; source impedance set to complex conjugate match. In (a), contour start = 14.5 dBm and contour step = 0.25 dBm. In (b), contour start = 45% and contour step = 2.5%. Note that the optimum load impedance values for both are very close together, and there is a minimum tradeoff when matching for maximum output power and maximum power–efficiency.

lead to certain undesirable effects, such as intermodulation and harmonic distortion, which in turn lead to the transfer of power to other frequencies near the frequency of interest. For a device with two signals at its input, one at a frequency f_1 and the other at a frequency f_2, it is traditionally the third-order (at frequencies $2f_1$–f_2 and $2f_2$–f_1) and fifth-order ($3f_1$–$2f_2$ and $3f_2$–$2f_1$) intermodulation products that are of most concern, because they are near the two frequencies of interest (f_1 and f_2) and therefore will be the most difficult to filter out of the system. Therefore, rf and telecommunications applications, such as power amplifiers, require devices that exhibit highly linear operating characteristics. Two-tone measurements must be performed on the device offerings in the SiGe BiCMOS and rf CMOS technologies to fully analyze the linearity of the devices offered. Use is made of an ATN LP2 load-pull system to make these measurements. The third-order and fifth-order intermodulation products are commonly measured, and the third-order intercept point (IP_3), an important figure of merit for describing linearity, is obtained. The intermodulation products are shown in Figure 8 for an SiGe HBT, and the extrapolation to IP_3 is illustrated. IP_3 is then used to determine the spurious free dynamic range (DR_f), defined as the difference between the output power at the fundamental and the output power at the third-order intermodulation product when the output power at the third-order intermodulation frequency is equal to the minimum detectable output signal, $P_{o,mds}$ [24]. DR_f is given by

$$DR_f = \frac{2}{3}(IP_3 - P_{o,mds}). \tag{7}$$

Thus, IP_3 provides the power-amp designer with a metric for determining a distortion-free operating range for the device.

Test site

The key to generating accurate, scalable, full-featured models, as described in the next section, lies in the availability of test structures from which to make measurements. For example, test-site characterization macros needed to construct an n-FET model would include

1. A length and width array macro for dc extraction/optimization of the Berkeley short-channel IGFET model (BSIM) parameter set.
2. A capacitor array to extract gate-oxide, overlap, and source-to-drain capacitance and leakage.
3. A set length and width array macros to measure threshold voltage and mobility mismatch.
4. A set of rf S-parameter structures of varying length, width, and number of finger configurations.
5. Open/short "de-embedding" structures to go with the above S-parameter macros.
6. DC and S-parameter gate resistance extraction macro.
7. Macros to gauge proximity to n-type wells and other process-specific effects.

The vast number of structures required to cover the characterization needs of all of the devices in a given chip technology can be enormous. For example, an SiGe or rf CMOS modeling test site might be as large as 20 mm × 20 mm. **Figure 10** shows a top-level view of a test site (designated as the "Granite" test site) which is the primary modeling test site for the IBM rf CMOS technology.

4. Compact model development

A mathematical model that predicts the electrical characteristics of a semiconductor device as a function of the conditions and constraints applied to it is designated as a *compact model*. In the case of the MOSFET device, a compact model predicts the output current (I_{ds}) and its derivatives (g_m, g_{ds}, g_{mb}) as a function of temperature, voltage bias, channel length, and device width. A compact model may be composed of a single element, such as an ideal resistor, or a complex network of interdependent sources, resistors, capacitors, and diodes used to model a bipolar junction transistor (BJT). To analyze a model or circuit containing more than one element, a matrix solver such as HSPICE**[4] or Spectre**[5] must be used. Implementing the compact model may require the extraction of less than ten parameters, as is the case for a junction diode, or it may require the elaborate extraction and optimization needed to extract the dozens of parameters of a BSIM-based MOSFET model.

The primary goal of a compact model development effort is to provide physics-based, scalable models that are fully integrated into the overall design kit environment. The models must be capable of predicting the full complement of device characteristics and behavior as a function of bias and temperature and must represent the statistical process window of the technology being modeled. Another important consideration is to make use of industry-standard or common elements in building the model to allow for more efficient translation of the models to multiple simulators while maintaining consistency in simulation results.

The emphasis on models which are physics-based dictates that the development effort must employ direct parameter extraction, rather than empirical or numerical optimization, whenever possible. It is also important to make use of process information obtained from technology development, such as vertical profile dimensions or doping concentrations. The use of a more physical model is more effective for predicting results for conditions beyond those used during the initial model parameter extraction, such as different device geometries. The demands of the rf/analog design environment have also led to the use of more complex subcircuit topologies for both active and passive devices to better predict high-frequency behavior. However, it is important to minimize the number of elements required in order to maximize simulation efficiency.

All of the key building blocks for the development of the scalable, statistical, and physics-based models used in IBM are described throughout this section. Note that this information, from the physical device layout and design rules to the development of scaling equations and the incorporation of device characterization and in-line electrical data, is required regardless of the type of device that is being modeled.

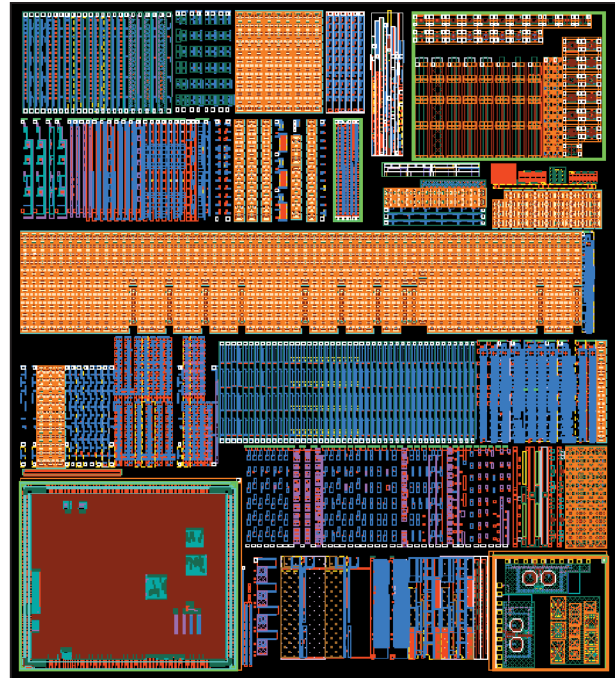

Figure 10

"Granite" test-site layout for 0.25-μm rf CMOS chip technology; contains structures for obtaining information on process parameters, interconnect parameters, in-line test macros, modeling macros, etc.

Statistical modeling

The compact device models being developed for the IBM SiGe BiCMOS and rf CMOS technologies have the ability to support standard Monte Carlo (statistical), process corner, and wafer-specific simulations. The basic structure of the model library makes use of a "skew file," which defines all of the statistical distributions, process corner parameters, and other model parameters that are shared across multiple devices. The definition of these distributions is dependent on the cooperation between technology development and compact model teams to determine the most dominant process parameter variations and the proper correlation of process and device model parameters. These correlations account for effects across multiple devices that may share a common process step, as well as multiple parameters within a single device that exhibit a strong physical correlation.

The primary input for specifying nominal and tolerance specifications of process parameters comes from in-line

[4] Synopsys, Inc., Mountain View, CA.
[5] Cadence Design Systems, San Jose, CA.

wafer (kerf) parametric data. This data provides the necessary statistical sample and establishes a direct connection between skew file parameters and measurements that are used as acceptance criteria for product wafers during manufacturing. Additional wafer characterization is used to supplement the in-line data and provide the basis for correlation of the process parameters and key device metrics, such as npn f_T and f_{max}.

The conventional method for process corner simulation has involved using a pair of model parameter sets to represent the process extremes, often referred to as "fast"/"slow" or "high"/"low" corners. The assumption in this method requires that, for a "fast" process corner, the device parameters are skewed so as to maximize active device currents and minimize other capacitances and resistances. While this may be valid for analyzing the process variation of the characteristic time constant of an analog circuit, the drawback of this method for corner-file generation is that this definition of the model parameter combinations will not always yield an extreme in the circuit performance for all types of analog applications.

To enable process corner simulation, the IBM skew file approach (patent pending) includes multiple corner parameters corresponding to each of the device types, such as resistors, capacitors, and BJTs. This structure supports simulation of different combinations beyond a single "fast" and "slow" corner pair and enhances a designer's ability to assess the sensitivity of the circuit performance. This sensitivity analysis can be done by repeating a simulation with each of the individual corner parameters set to +1 and −1 (corresponding to a $\pm 1\sigma$ variation) and comparing the results against the nominal simulation. With only one corner parameter set to be nonzero at a time, the total number of simulations will be twice the number of corner parameters. These single-parameter simulations are carried out to determine the appropriate sign, positive or negative, for each corner parameter necessary to maximize (or minimize) overall circuit performance. In this methodology, equal weights are given to the variations of all device types, so all of the corner parameters are set to the same magnitude. This magnitude is determined by first finding the 3σ variation limits of the circuit performance using a statistical simulation and then setting the magnitude of the corner parameters to match these limits. By using an initial Monte Carlo simulation to calibrate the results from the corner analyses, the designer acquires an efficient means to account for the effects of the process variation and include the necessary design margin. As a result, this approach provides the benefits of both conventional corner and Monte Carlo simulations and requires only a few additional simulation iterations.

In addition to the process statistics in the skew file, the individual model files also include distributions to represent device mismatch effects. The skew file statistics represent the global process variation across all wafers, while the mismatch represents the local variation observed on a typical wafer. Specific test-site structures are used to measure the mismatch and are designed using good layout practices such as use of same orientation of near-adjacent devices with symmetric wiring. As with all other aspects of the modeling, every attempt is made to define these mismatch effects and account for geometric and bias dependence in a manner that is consistent with the physical nature of the devices.

One key aspect of the extraction methodology necessary to maintain the statistical integrity of the final models is the assessment of the hardware used to extract the model parameters. It is important to establish any offset that may exist between the defined nominal process values and the measurements of test-site wafers. Following the device characterization and completion of the model extraction, skew file parameters are then recentered to represent nominal process and device specifications. This concept of "recentering" also enables the models to support simulation analyses using skew file parameter adjustments that are based on a set of process parameters and single bias point measurements that are taken from the in-line wafer parametric data.

The overall flow of the steps of the compact model development process is illustrated in **Figure 11**. Note that model parameter recentering and the inclusion of the process distributions to enable full statistical simulation follow directly after parameter extraction, as shown in the lower left corner of the figure. Although the figure pertains to the use of the BSIM model in support of MOSFET devices, the overall development flow is applicable to all IBM SiGe and rf CMOS technology devices.

HBT

Until the mid-1990s, the semiconductor industry relied almost exclusively on the simulation program with integrated circuit emphasis (SPICE) Gummel–Poon (SGP) model for BJT circuit design. The model included effects important in analog design not found in the earlier Ebers–Moll-type models such as low-current nonidealities and high injection effects, and replaced the underlying physical model with equations based on the more complete integral charge-control relation (ICCR) [25]. But recently, with the revival of BJT and HBT technology for high-speed communications and rf applications, the SGP model was found to be increasingly inadequate and had to be revised to include more accurate modeling of the physical effects found in high-speed devices operating at high current densities. These effects include better Early-effect modeling (output conductance), quasi-saturation,

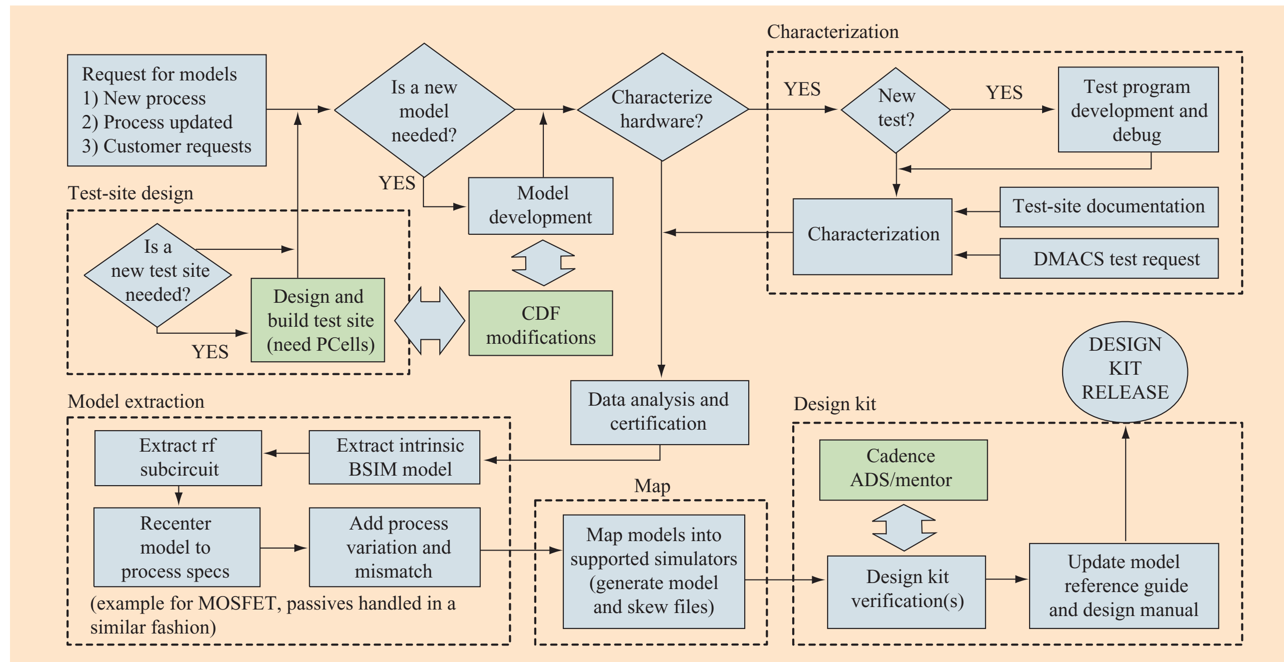

Figure 11

Overall flow of the steps of the compact model development process.

avalanche multiplication, thermal self-heating, and accurate transit-time modeling.

This needed revision of the SGP model for modern bipolar transistors was addressed at the Bipolar/BiCMOS Circuits and Technology Meeting, which started a committee to establish a public-domain improved bipolar compact model. This resulted in the vertical bipolar inter-company (VBIC) model formally presented in 1995 [17]. The VBIC model was physically based on the same ICCR that underlies the SGP model, but also included several additional model elements built around the core model. Additionally modeled effects include a parasitic pnp, self-heating, bias-dependent Early voltages, temperature scaling, a Kull-based model for quasi-saturation, and additional parasitic capacitances found in aggressively scaled modern devices. Also, in contrast to the SGP model, which used separate equations to model the transistor in each operating regime, the VBIC model was constructed with continuous smooth functions over all bias ranges for enhanced numerical stability. However, in an effort to keep a partial backward compatibility with the SGP model, the extra physical modeling structure increased the internal model node count from 3 to 7 and approximately doubled the required number of parameters to 70. The primary recognized inadequacy of the VBIC model revolves around the poor implementation of the Kull model for device operation in strong quasi-saturation [26].

IBM SiGe technology design kits currently integrate both SGP and VBIC models for the SiGe HBT, but the rapidly growing suite of SiGe HBT technologies, with an extremely wide range of device performance targets, has placed additional questions on the validity of the physical assumptions used to derive the standard VBIC model. For example, in IBM SiGe technologies, the model must correctly predict the strong quasi-saturation and avalanche breakdown of the IBM SiGe 5PA high-voltage (6.4-V) HBT, as well as model non-quasi-static transport and ac current crowding of the recently announced 200-GHz eight-generation SiGe HBT. This is a difficult task for even the most complex models.

For these reasons, two additional HBT models, HiCUM and MEXTRAM, are currently under evaluation by IBM and the Compact Modeling Council as potential successors to the current VBIC standard. The HiCUM (*Hi*gh *CU*rrent transistor *M*odel), developed at Ruhr University in Bochum, Germany, and first implemented in 1981, was developed initially for design of high-speed ECL circuits that operate at high current densities [27]. Based on the ICCR, the model was extended to include SiGe HBT structures with the General Integral Charge-Control Relation (GICCR) that now provides the physical basis for the model [28]. A most important impact of the GICCR is

Table 2 Comparison of modeled physical effects and requirements for current HBT compact models. Adapted from [30], with permission; © 2001 IEEE.

	SGP	*VBIC*	*HiCUM*	*MEXTRAM*
HBT/SiGe modeling	—	√	√√	√√
Quasi-saturation	—	√	√√	√√
f_T modeling	—	√	√√√	√√
Self-heating	—	√	√	√
Substrate modeling	—	√	√	—
β-E breakdown	—	√	√	—
Parasitic pnp	—	√	—	−√
No. of internal nodes	3	7	4	5
No. of parameters	35	80+	90+	67

the implementation of weighting factors that account for the change in mobility and intrinsic carrier concentrations affecting the charge storage in the neutral regions affected by the high Ge content. Other important effects present in the HiCUM model include accurate modeling of the quasi-saturation region and extensive physical description of bias-dependent transit times and non-quasi-static behavior.

The MEXTRAM (*M*ost *EX*quisite *TR*ansistor *M*odel) model, developed and implemented at Phillips [29], is also based on the ICCR of the SGP model. The modeling of the collector epilayer differentiates the MEXTRAM model from VBIC. MEXTRAM extends the modified Kull model by adding the effects of velocity saturation in the collector at high current densities, correctly predicting quasi-saturation and the onset of Kirk effect. This collector model is also implemented in a smoother mathematical description that is beneficial to the calculation of higher-order derivatives, important for harmonic distortion analysis. MEXTRAM, like HiCUM, has implemented additional parameters to take into account bandgap grading in SiGe devices, as well as an extra parameter to model the changes in I_b due to neutral-base recombination.

Table 2 contains a brief summary of the physical effects included in each of the BJT models and a comparison of the number of internal nodes and model parameters required. The existing IBM SiGe design kits have implemented the VBIC model as the primary element within the npn subcircuit for technologies that typically include both high-f_T and high-breakdown types of devices. The weak avalanche effect in the VBIC model is based on the assumption that the peak E-field occurs at the base–collector interface, which is valid for the highly doped collector of the high-f_T device. For the lower-doped collector of the high-breakdown npn, the peak E-field occurs at the collector–substrate interface. To account for this difference and to overcome this limitation in the VBIC model, the IBM model topology was modified to move the physical location of the weak avalanche current generator inherent in the VBIC model for high-breakdown devices. The standard VBIC avalanche current is still used for high-f_T devices. While there are distinct differences in the performance of these devices, a single model file has been used with an input parameter that is passed in from the circuit netlist to specify which device type is being modeled. This is another way in which the overall modeling methodology tries to reflect the physical realities of the devices, since it allows for commonality in many of the key calculations and model parameters that are derived from the basic process flow.

Typical model-to-hardware correlation plots help to illustrate the success achieved in using the VBIC model to represent the various device characteristics of the SiGe HBT npn devices. The dc forward Gummel and output curves for the high-f_T npn device in the IBM 0.25-μm SiGe process are shown in **Figure 12**, and for the high-breakdown npn device in the same process in **Figure 13**. The figures show the effects of self-heating, weak avalanche current, and the difference in the quasi-saturation behavior of the two types of devices. Note that the model parameter extraction process begins with fitting of these dc characteristics. Once initial parameter values associated with the basic dc characteristics have been determined, extraction and model optimization continues using S-parameter, noise-figure, and large-signal measurements. **Figure 14** shows the V_{cb} dependence of f_T vs. I_c, for both high-f_T and high-breakdown devices, as extrapolated from S-parameter data at a fixed frequency using a simple gain-bandwidth product. Examples of minimum noise figure and rf power gain for the high-f_T npn are included in **Figure 15**. Finally, the initial results of the model vs. large-signal data correlation can be seen in **Figure 16** for both device types.

Another key aspect in the development of the SiGe HBT models is the inclusion of extensive geometric scaling equations. While the more advanced BSIM-based MOSFET models have multiple parameters to represent short-channel or narrow-channel effects, the built-in scaling of VBIC is not adequate to support the range of layout geometries offered by the IBM SiGe design kit parameterized cells (PCells). The use of a single "area" factor to scale all of the current density, capacitance, and resistance parameters does not provide the flexibility necessary to accurately differentiate the dimensional changes in the various regions within the device. Separate calculations are included as a part of the subcircuit model to determine the proper effective dimensions, such as intrinsic and extrinsic junction areas or perimeters, to generate the final set of VBIC parameters, given a specific emitter size and device layout configuration.

Modeling V_{be} and β mismatch of the SiGe npn devices is important for providing designers with information

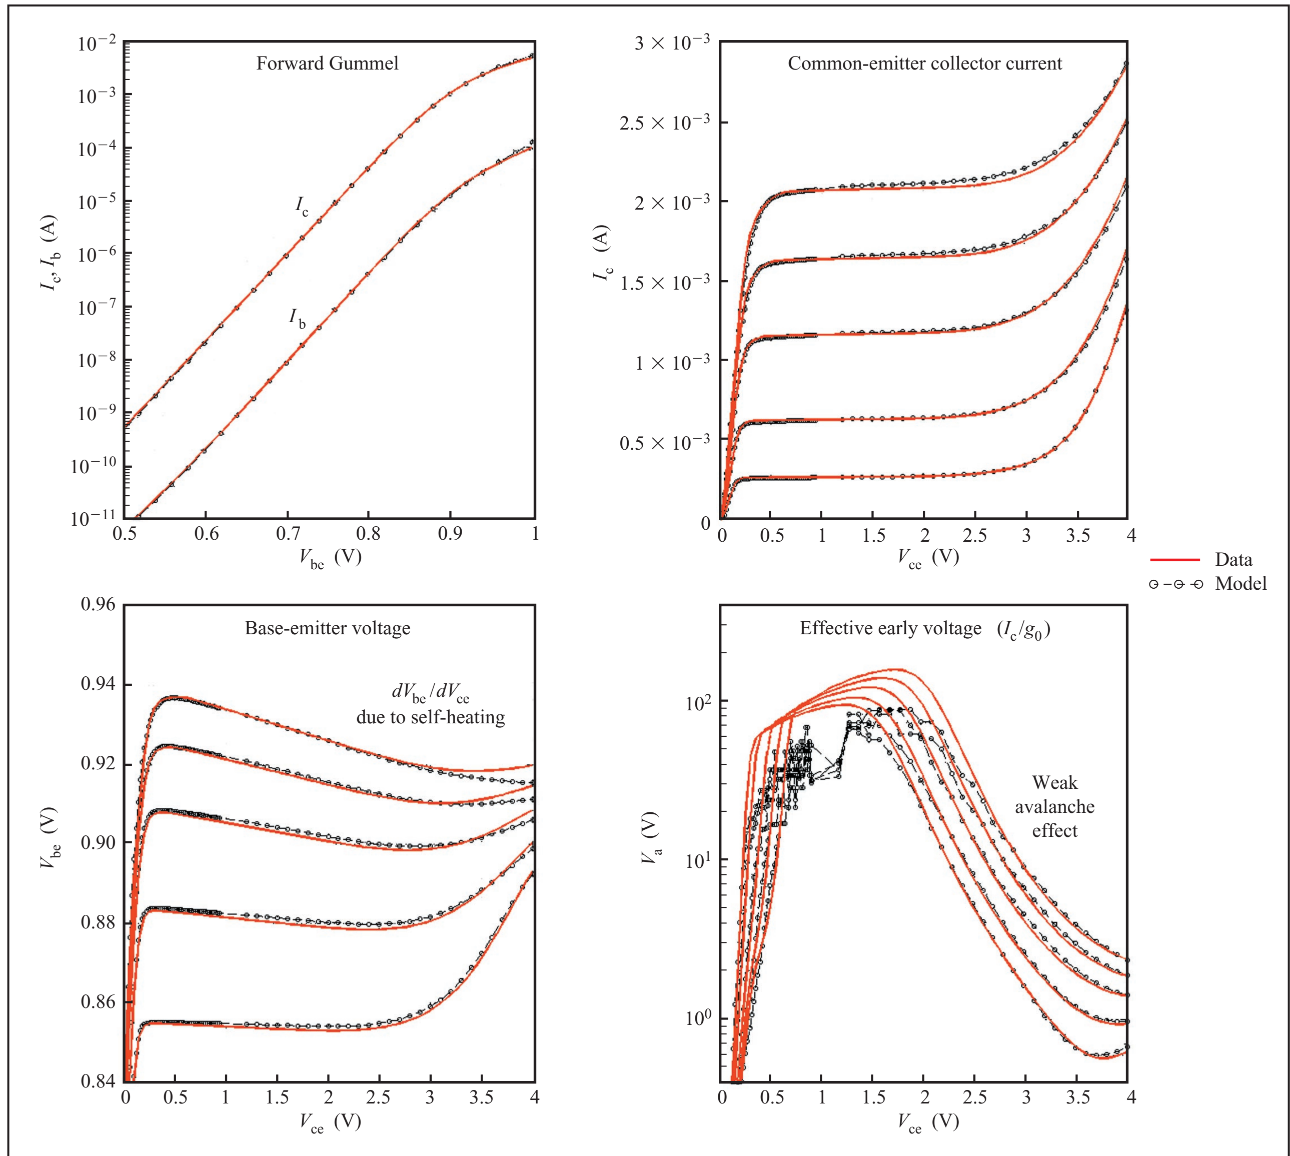

Figure 12

DC characteristics of high-f_T npn SiGe HBT fabricated using the IBM 0.25-μm SiGe process.

necessary to assess performance in many typical small-signal analog applications such as high-speed A/D converters, bandgap voltage references, and differential circuits. The scaling equations as implemented also include statistical distributions to represent low-current V_{be} mismatch as a decreasing function of increasing emitter area. Another factor is defined to account for the increase in V_{be} mismatch that is observed as the current bias increases. A third distribution is used to represent the β mismatch, which is also modeled as a function of emitter area.

MOSFET

The technical development of the MOSFET compact model has closely followed the increased demands placed on it by circuit designers. As MOSFET-based designs have evolved from purely digital to analog and to analog rf, corresponding MOSFET models have become increasingly more complex.

Traditionally, the first MOSFET model was the so-called digital model. In the development of the digital model, the primary focus was on accurately predicting the on-current (I_{ds} at $V_{gs} = V_{ds} = V_{dd}$) and the switching

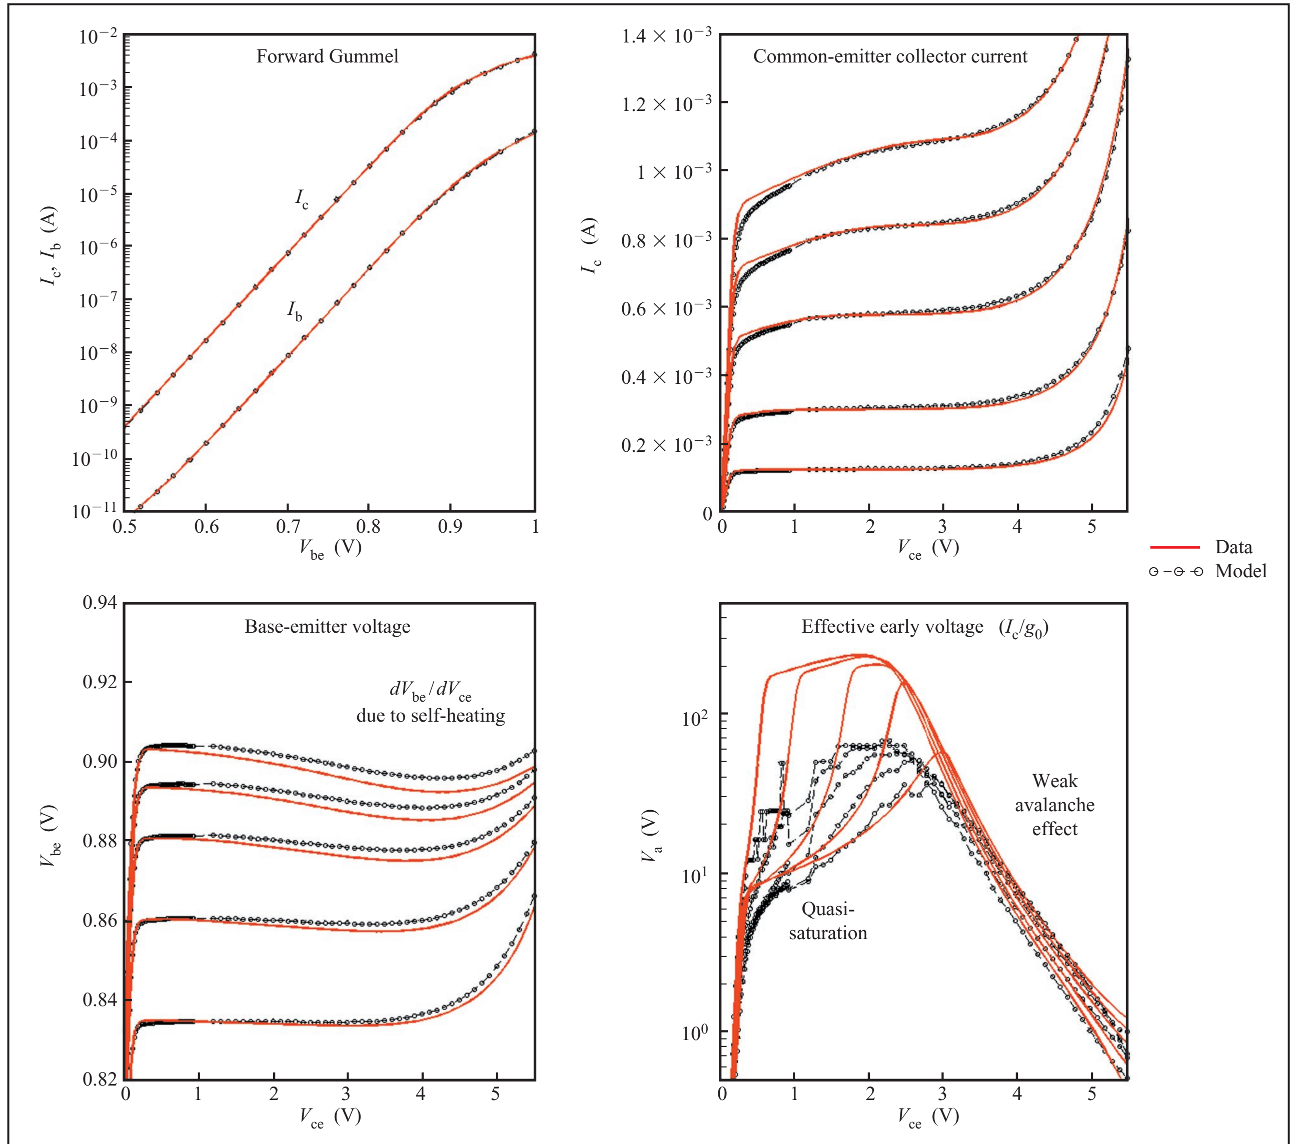

Figure 13

DC characteristics of high-breakdown npn SiGe HBT fabricated using the same IBM process.

speed of the device. Issues such as scalability and I_{ds} accuracy near threshold were not addressed, since the focus was on minimum-channel-length (i.e., fast) devices that were either on or off. This model served as a good introduction point because the equations comprising the model had their basis in semiconductor device physics. This is in contrast to an empirical model where the equations may be composed of splined polynomials whose coefficients have been optimized to provide the best fit between the measured data and the simulation results. As CMOS technologies have progressed into the submicron region, the limitations of the existing "digital" models have become much more apparent. The following list highlights many of the physical effects that the existing models are unable to predict [31]:

- Short- and narrow-channel effects on threshold voltage.
- Effects of nonuniform doping.
- Mobility reduction due to vertical field.
- Bulk charge effects.
- Velocity saturation effect.
- Drain-induced barrier lowering (DIBL).
- Channel-length modulation (CLM).
- Substrate-current-induced body effect (SCBE).

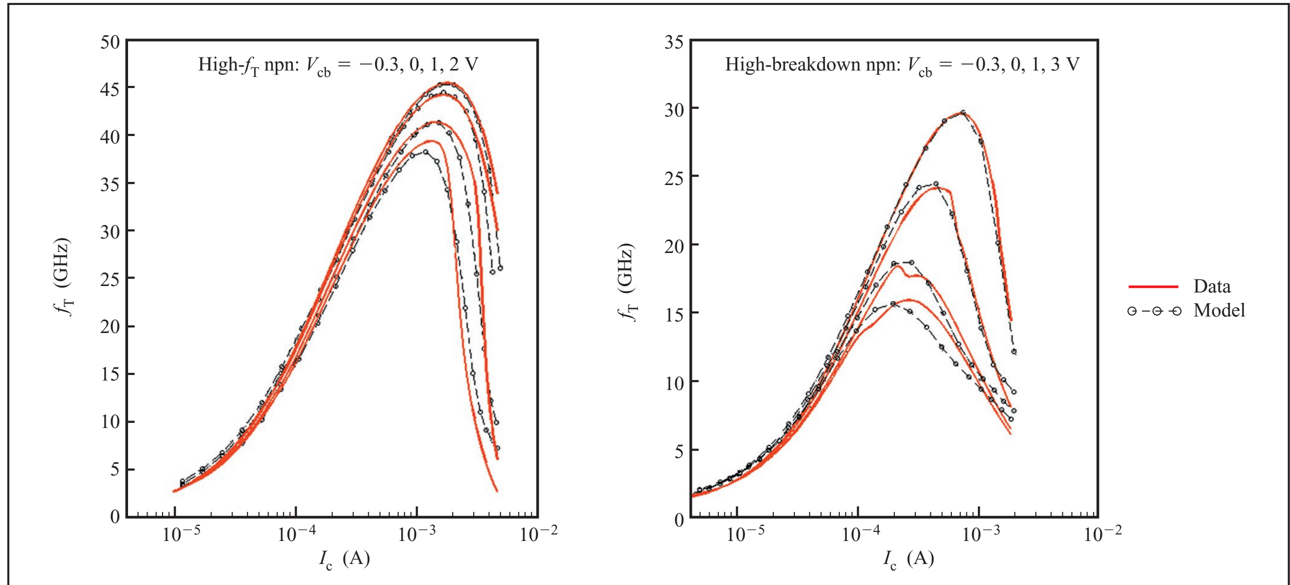

Figure 14

Associated f_T characteristics vs. V_{cb} for high-f_T and high-breakdown npn devices.

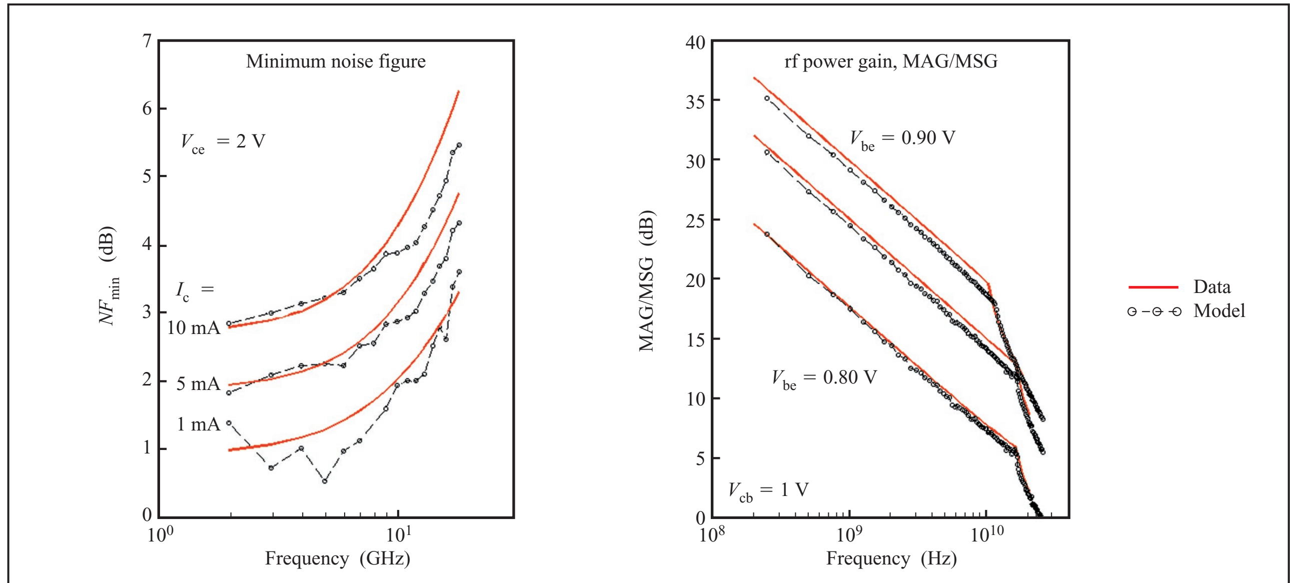

Figure 15

Associated minimum noise figure and rf power gain for high-f_T npn device.

- Subthreshold conduction.
- Source/drain parasitic resistance.

The inability of the digital model to deal with these effects results in a model that does not scale with channel length and/or width. In fact, because of the operating region and voltage bias dependencies of these effects, the drain current is not accurately modeled even for a fixed-dimension device. Even worse, the derivatives of the drain current (g_m, g_{ds}, and g_{mbs}) can become discontinuous

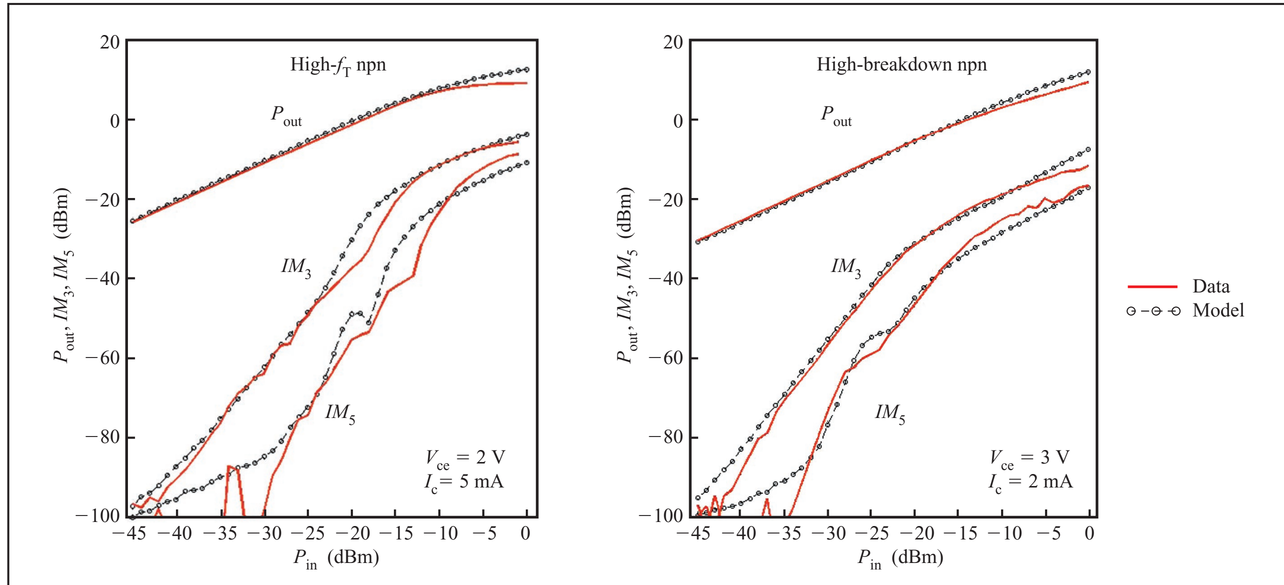

Figure 16

Associated large-signal characteristics for high-f_T and high-breakdown npn devices; $Z_{in} = Z_{out} = 50\ \Omega$.

owing to the "splining" approach that has been applied to the digital model equations. This means that, in addition to the inherent inaccuracies, many analog circuit simulations do not converge to a result when these digital models are used. This results in a severe handicap for the analog designer whose designs depend upon an accurate representation of MOSFET drain current, gain (g_m) and output impedance (I/g_{ds}).

Consequently, several attempts were made through the late 1980s and early 1990s to create an analog model. Of these, the Berkeley short-channel IGFET model (BSIM) and its successors are the most widely accepted and used. In developing the BSIM, a "start from scratch" approach was used that placed an emphasis on three areas of importance:

1. *Device physics* The robustness of a compact device model can usually be traced to how much of its fundamental basis is tied to the physics of the device.
2. *Scalability* Width and length scalability is incorporated into nearly all of the equations that compose the structure of the BSIM model. This enables a single model, and thus a single model parameter set, to be used to predict the performance of a device over a wide range of geometries.
3. *Robustness* The equations used to represent the various effects were combined in such a way as to create one continuous expression for I_{ds}. This eliminated the discontinuous derivative problems found in many of the earlier MOSFET models mentioned above. From a circuit simulation standpoint, the robustness of the BSIM model can also be tied to its formulation as a charge-conserving model. In a charge-conserving model, the nodal equations are expressed in terms of charge instead of capacitance. These equations are considered balanced when all of the charges sum to 0. Capacitance-based models are not charge-conserving because capacitance is an incremental quantity that accurately predicts the change in charge versus voltage for only infinitesimally small changes in voltage [32]. The development of the BSIM model to the BSIM3v3 model that is in widespread use today is well documented [33].

The MOSFET model used by IBM in its SiGe BiCMOS- and rf CMOS-based design kits is the third-generation BSIM3 model from U.C. Berkeley (BSIM3v3.2). Although its name implies that it is a continuation of the earlier BSIM and BSIM2 models, it is really a total rewrite from scratch. The goal here was to introduce more accurate physics-based equations into the model in a way that was still mathematically robust from a simulator standpoint. In general, the best compact models are those whose representative equations have their basis in the physics behind the devices they are representing. These models tend to scale better and to be more accurate at biases and geometries outside the bounds being measured. This was the approach used in developing the BSIM3 model. Where possible, the equations used to represent the effects

listed above were based on the solution of Poisson's 2D equation for the distribution of charge across the channel of a MOSFET. As a result, the fundamental equations for threshold voltage, mobility, and velocity saturation tended to be very physically based.

The semi-empirical nature of the BSIM model dictates that a combined extraction and optimization approach should be used to extract the BSIM parameter values. First, the physical "process" parameters such as oxide thickness, base threshold voltage, channel length and width, and series resistance are extracted from either single- or multi-device measurements. Next, an optimization approach is used to produce a set of BSIM parameters that minimize the model-to-hardware error across the device size, bias range, and temperature space being fit. To date, two different optimization approaches have been used. The more "traditional" approach is the so-called local optimization method. In local optimization, parameters are fitted, a few at a time, by varying their value to minimize the model-to-hardware difference over a very specific range of device sizes and operating regions in which the effects of these parameters are dominant. This "local" optimization approach is preferred over the standard "global" optimization approach because it tends to avoid the local minima phenomena that can occur when very large number of parameters are optimized over a very large space. The second optimization method used more recently at IBM is a genetic algorithm (GA)-based approach [34]. The genetic algorithm avoids using the "stair-climbing" optimization approach that causes other global optimization approaches to become trapped in local minima. The GA is able to optimize all of the BSIM parameters concurrently by means of a "fitness function" that weights its target criteria accordingly. It is very CPU-intensive, but minimizes the human effort needed to extract the models. The key to its success lies in the definition of the fitness function.

Length and width scalability have been built into the threshold voltage equation, the effective mobility equation, and most of the other fundamental equations behind the BSIM model. However, if the physical mechanisms underlying the behavior of the device do not match those of the scaling equations, the models may not scale across the entire geometry range being offered in a given technology. To circumvent this problem, a technique known as "binning" has been incorporated into the BSIM model. As its name implies, binning allows the channel width/length geometrical space to be broken up into several regions, or bins. For example, if a model is broken down into three length regions and three width regions, it is said to be composed of nine bins (3×3). This approach enables the modeler to extract nine separate BSIM parameter sets to cover the entire width/length space. Since each of the nine BSIM parameter sets only has to cover its width/length region, the scalability requirements placed on the intrinsic physics-based scaling are relaxed. The downside of using this approach is the increased development overhead associated with creating and maintaining binned models. Although such models are not currently in use in the effort described here, they may be useful in the future, as decreasing channel lengths continue to challenge the intrinsic scalability of the BSIM model.

When designers place a MOSFET device of a given width/length ratio in the schematic, a level of confidence is obtained that the I_{ds}, g_m, or g_{ds} values predicted by the model are valid for that device. If the device they are using has an rf PCell (described below), they have confidence that the device parasitics are also represented accurately within the model. If designers are using a standard "digital" MOSFET PCell, they must rely on an estimation of device parasitics that have been incorporated into the design kit. To predict the source/drain (S/D) parasitic capacitance, an estimate of the area and perimeter of both the source and drain is passed to the S/D diode model contained within the core BSIM model subcircuit. As with any estimate, some assumptions must be made. To estimate the size of the S/D diffusions, we make the following assumptions: 1) Contact sizes and spacing are at a minimum, thus enabling the smallest possible diffusion. 2) In the event of an even number of gate "fingers," there will be one less drain diffusion than source diffusion. 3) Applied photoresist bias is nominal [35]. Within the SiGe BiCMOS 6HP design kit, the designer has the choice of two different parasitic estimation approaches. In the first approach, the S/D area and perimeter estimates are calculated in the Cadence component description format (CDF) associated with each schematic level device instance. The values of drain area (*AD*), drain perimeter (*PD*), source area (*AS*), and source perimeter (*PS*) are then passed into the device model. In the second approach, the values of *AD*, *PD*, *AS*, and *PS* are calculated in the device model itself. Each approach has its own advantages and disadvantages. Calculating the area and perimeter values in the CDF simplifies the model code and reduces overall simulation time, since the values are calculated once prior to netlisting. Calculating the area and perimeter in the model enables statistical process variation to be applied to the area and perimeter during Monte Carlo simulation and allows the user to sweep the channel length (or width) as a design variable during simulation and still have the estimated parasitics included.

In addition to increased focus on analog model accuracy, the BSIM generation of compact device models has also placed an increased focus on modeling noise. Prior to the use of BSIM, the noise contributions of a MOSFET were modeled as the standard kT/q thermal noise contribution. In BSIM, flicker noise ($1/f$ noise), channel thermal noise,

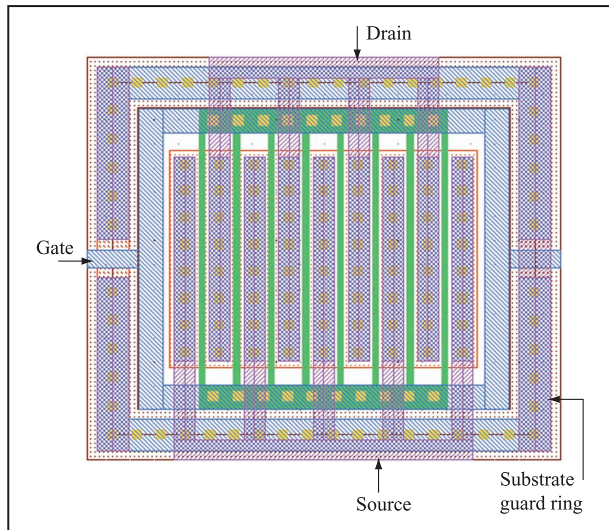

Figure 17

Example of an rf MOSFET PCell-based device layout.

induced gate noise, and the thermal noise associated with parasitic gate and diffusion resistances are all modeled.

At its conception, the IBM rf CMOS technology was intended to be a variation of the base IBM CMOS technology, which offered improved rf-quality models and extra front end of the line (FEOL) passives to give the rf CMOS designer an adequate "toolbox" of devices. RF CMOS is targeted at the cost-sensitive customers whose analog design points enable them to implement their designs in CMOS only. Typically, these design points include rf circuits with relatively little analog content or rf circuits that are well understood and have been implemented using bipolar devices in the SiGe BiCMOS technology.

From a MOSFET model standpoint, the needs of the rf CMOS analog designer are even more critical than the needs of the BiCMOS analog designer. In a BiCMOS technology, a designer would likely use a bipolar device in situations where high speed, high gain, and good matching are important. For an rf CMOS technology, this is not an option. This forces the rf CMOS designer to use more innovation in design—e.g., to use elaborate on-chip timing approaches to minimize signal latency and DSP approaches to "clean up" signal issues related to distortion.

In concept, the transition from an analog model to an rf-quality model is not really as big as that from a digital model to an analog model. From a designer's perspective, the rf model really involves extending the analog model accuracy from the low-frequency region into the rf region. At low frequencies, the capacitive load of the gate on the preceding circuit element is important. At high frequencies, the gate capacitance, channel resistance, and diffusion parasitics all combine to form a complex load on the prior stage. From an output impedance standpoint, a similar analogy can be made. At low frequencies, the output resistance is dominated by the channel resistance. At high frequencies, the S/D parasitics and bulk resistance contribute largely to the complex output impedance. This presents a unique problem to MOSFET designers. At low frequencies, the performance of the MOSFET can be adequately represented by modeling the intrinsic device, the overlap capacitance, and the S/D parasitic diode. This can be done because these effects are independent of device layout (as long as the number of fingers and multiplicity are known). At high frequencies, the layout cannot be ignored. The approaches used to wire to the gate and contact the substrate greatly affect the high-frequency performance of the device. This creates a problem for MOSFET designers: Without knowledge of the exact layout of the MOSFET, how can its high-frequency performance be modeled? Our answer is to provide designers with two MOSFET PCell layouts. The first is the standard "digital" layout, in which the designers are free to wire to the gate and to the substrate. The model provided with this device can be verified at lower frequencies. The second PCell is similar to the one shown in **Figure 17**. In this PCell, the gate wiring and substrate contact scheme are defined and controlled. The model provided with this device scales as a function of channel width, length, and number of "fingers" at both low frequency and high frequency. By using this layout and its associated rf model, designers can be certain that the rf model they are using represents the device they have designed. The purpose of this rf PCell is not to provide the highest-performance layout, but rather to provide a controlled device configuration that has been well characterized and modeled.

The substrate circuit defined in the model is critical to accuracy at high frequency. A comparison of S-parameter measurements for identically sized devices with varying substrate contact schemes has demonstrated two key points. First, the output impedance at the drain of the device can vary significantly depending upon the proximity of the substrate contact. Second, this variation can be represented accurately by using a model structure similar to the one shown in **Figure 18**. Note that the model parameters needed to model the substrate resistance effects are layout-dependent, as previously described. Therefore, a unique set of substrate network parameters must be used for the layout shown in Figure 17.

Inductors

Inductors are essential devices in many rf circuits (voltage-controlled oscillators, or VCOs, impedance-matching networks, etc.). Their on-chip implementation decreases

packaging parasitics, reduces the number of pins required, makes possible more compact circuit boards, and paves the way for system-on-a-chip (SoC) solutions. Any design flow for rf applications would be incomplete without adequately optimized and modeled on-chip inductors [36].

The performance of an inductor is best judged by the inductance and Q (quality factor) [37] across a frequency spectrum of interest. Q is inversely proportional to the power dissipated by the inductor; therefore, power losses associated with the inductor should be minimized for increased performance. The two main causes of power loss in the inductor are its resistance and loss within the substrate beneath the inductor due to capacitive and/or magnetic coupling into the substrate. Losses associated with the inductor's series-resistive elements increase with frequency because of skin-effect loss and magnetic-field-induced proximity-effect loss.

There are a variety of options with which inductors can be implemented. While inductors can be fabricated using regular metal interconnects for low-Q, low-frequency applications, very-low-resistance or one or more thick Cu or Al layers must be used for most rf applications to minimize losses due to series resistance. Furthermore, the inductor metallization should be separated from the substrate with as thick a dielectric stack as possible, preferably with a low dielectric constant (to minimize capacitive coupling, and hence losses within the substrate underneath the inductor). A Faraday shield [38] between the inductor and the substrate should also be an option, since it may increase the peak Q of an inductor at the expense of a reduced self-resonant frequency and added layout complexity.

Designers seek accurate compact models for all device types, and the inductor is no exception. The inductor model should be inclusive of all of the various configurations possible for the device. The model should support various metallization options—fabrication of the inductor with wiring interconnects or with dedicated thick Cu metallization, various dielectric stack heights, possible Faraday shielding (ground-plane) options, and a range of values for the parameters defining the inductor planar geometry. For an octagonal spiral, which can be used for on-chip inductors, the planar geometry parameters are outer dimension (d), turn width (w), turn spacing (s), and number of turns (n).

The low-frequency inductance of an on-chip spiral can be calculated using various methods, as outlined in [39]. One common method is known as current sheet approximation [39], in which the inductance, L, is calculated via

$$L = \frac{\mu_0 n^2 d_{\text{avg}} c_1}{2} \left[\ln\left(\frac{c_2}{\rho}\right) + c_3 \rho + c_4 \rho^2 \right], \qquad (8)$$

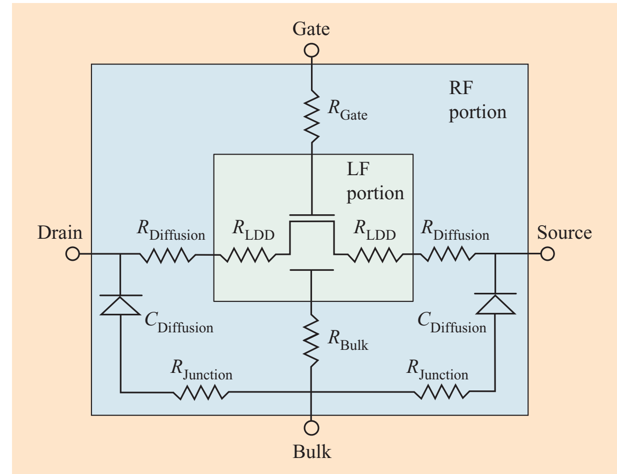

Figure 18

Schematic overview of an rf MOSFET model.

where

permeability of free space $\mu_0 = 4\pi \times 10^{-7}$;

inner diameter of spiral $d_{\text{in}} = [x - 2nw - 2*(n-1)*s]$;

outer diameter $d_{\text{out}} = x$;

averaged diameter $d_{\text{avg}} = (d_{\text{in}} + d_{\text{out}})/2$;

fill ratio $\rho = (d_{\text{out}} - d_{\text{in}})/(d_{\text{out}} + d_{\text{in}})$;

c_1, c_2, c_3, and c_4 are layout-dependent coefficients (for an octagonal spiral, they are chosen as $c_1 = 1.27$, $c_2 = 2.29$, $c_3 = 0$, and $c_4 = 0.19$).

The inductance calculated this way does not include the vertical thickness of the spiral metallization (the self-inductance of a metal diminishes with increased thickness). In an accurate inductor model, the metal thickness and its effect on the inductance should be addressed by the use of scaling constants.

The fact that the resistance and inductance of an on-chip inductor are frequency-dependent should also be addressed in the model. Skin-effect losses are caused by the current in the spiral turns flowing increasingly on the surface of the metal as frequency increases. Since the cross-sectional area through which the current flows decreases with increasing frequency (it flows in an increasingly narrow annular ring), the effective resistance and inductance of the conductor undergo a change with frequency. The fact that the current flows toward the surface of the metal causes the resistance of the conductor to increase as a function of the square root of the

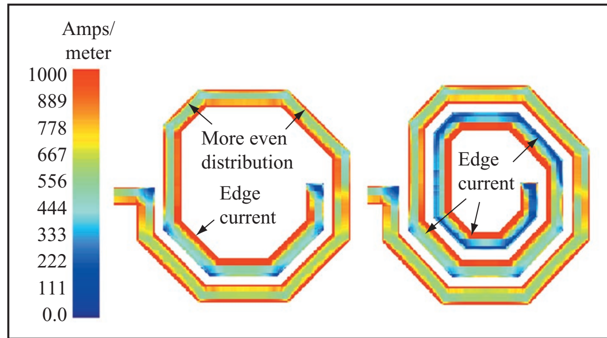

Figure 19

Electromagnetic simulation of two inductors operating at 2 GHz, showing nonuniform current distributions. Adapted from [40], with permission; © 2002 IEEE.

frequency. This surface current flow also excludes the magnetic field from the interior of the conductor, causing a slight reduction in the self-inductance with increasing frequency, until it is completely determined by the magnetic field external to the conductor.

The enhanced magnetic field that exists in the central and outer portions of the spiral tends to cause nonuniform current flow in the turns. For the inner turns, current flows only on the innermost edges of the turns, while for the outer turns, current flows on their outermost edges. **Figure 19** shows electromagnetic simulations of two inductors ($n = 1.5$ and $n = 2.5$). The current flow is shown at a frequency of 2 GHz. The simulations are done using Sonnet**,[6] a planar electromagnetic (EM) simulator that accurately handles vertical current flow. Such a simulator is known as a 2.5D simulator (compared to a 3D simulator, for which current flow is allowed in all directions).

This effect is frequency-dependent because the induced voltage in the turns (eddy current) increases with increasing frequency. This nonuniform current flow, typically called the "proximity effect," tends to cause the effective spiral resistance to rise faster with increasing frequency than can be attributed to skin effect alone. In addition, the net inductance decreases because of an effective reduction in the radius of the spiral caused by the current crowding to the innermost edge of the inner turns. The proximity effect can be the dominant loss mechanism at frequencies of interest for multi-turn spirals.

The modeling of the frequency dependence of the inductance and resistance of an inductor can be quite challenging. The standard approach is to calculate the resistance and inductance at a predetermined frequency point—the frequency at which the inductor is expected to dissipate the most power. While this approach produces accurate results for single-frequency simulations, for broadband simulations it is rather inaccurate. A better approach is implementing the frequency dependence via a network of fixed resistors and inductors. Such a network then will have the desired frequency dependence in its effective resistance and inductance. Circuit implementation of skin and frequency effects (CISP) produces accurate results for all frequencies of interest and is superior to the standard approach, which is accurate at only a single frequency point.

A scalable inductor model provides a mapping from all of the input model parameters (i.e., type of metallization, dielectric stack properties, type of Faraday shield, if any, planar geometry parameters) to a subcircuit for which the circuit element values are functions of the input model parameters. The subcircuit should employ CISP to capture the frequency dependence of the resistance and inductance, should include all of the various capacitances associated with the device, and should account for substrate-related losses. Such a subcircuit would then be an accurate representation of the device for the frequency spectrum of interest. **Figure 20** is a schematic diagram of a general inductor model subcircuit. In the current IBM model, the block for series inductance and resistance elements is replaced by a CISP network that greatly improves broadband accuracy.

Our inductor models meet the requirements outlined above. Along with our design automation tools, the designer receives a scalable inductor model that spans a very wide range of inductances (from hundreds of picohenries to hundreds of nanohenries) with quality factors over 30 (available with standard IBM dual-metal BiCMOS and rf CMOS offerings). Designers can see "on-the-fly" the low-frequency inductance and peak Q frequency values as they alter the PCell parameters (the input model parameters mentioned above). The designer receives full documentation on the model showing how the device is predicted to behave with statistical process variations and changes in temperature, and how the model correlates with measurements for a large variety of device configurations. The designer is also given guidance on how to achieve the desired inductance, peak Q, peak Q frequency, and self-resonant frequency, which is critical for efficient use of the model. The inductor is widely scalable. It employs CISP to capture skin and proximity effects, and therefore is accurate for all frequencies of interest. It also supports various ground-plane, metallization, and dielectric options. In its development and verification, designer/user feedback, electromagnetic simulations (e.g., Sonnet), and on-wafer S-parameter measurements are utilized. For geometries that are not yet covered by our scalable broadband model, such as

[6] Sonnet Software, Inc., Liverpool, NY.

symmetric cross-coupled interleaved structures [41], we provide S-parameter blocks and/or geometry-optimized subcircuits that can be used to model a specific structure.

Capacitors

The SiGe BiCMOS and rf CMOS technologies support different types of capacitors, including MOS, poly–poly and metal–insulator–metal (MIM) devices. The variety of the device offering allows designers to optimize their circuit performance and layout by selecting a capacitor on the basis of requirements for density, maximum voltage use, voltage bias sensitivity, or higher Q. Each of these capacitor types requires a unique subcircuit topology to represent the physical nature of the structures. In addition to a main capacitor element, each model includes resistive, inductive, and capacitive elements, as necessary, to account for the fringe effects, parasitic capacitance under the bottom plate of the capacitor, series resistance of the top and bottom plates, and, in the case of the MIM, the inductance associated with the metal plates. Examples of the basic model topology for MOS and MIM capacitors are shown in **Figures 21(a)** and **21(b)**, respectively.

Each element of the subcircuit in the model is coded to support the scalable dimensions of the capacitor, based on the defined width and length of the structure. Capacitance data is used to extract parameters for the voltage and temperature dependence of the main capacitor element. The final R–C or R–L–C configuration of the subcircuit is verified with S-parameter measurements, which show that the MIM capacitor has a higher quality factor (Q), while a MOS capacitor of similar geometry yields a higher capacitance. These results illustrate the tradeoff in density and performance that different capacitor types often provide. Additional model features include capacitance mismatch as a function of geometry and the flexibility to generate the netlist of the capacitors by specifying length and width or length and capacitance.

Varactors

While competitive technology comparisons often focus on the performance of the active devices, the features of the passive devices, such as varactors, MIMs, and inductors, are essential for offering a complete technology solution. VCO designs are dependent on key varactor characteristics, including the linearity, tuning range, and capacitance density. Current technology offerings include both collector–base (C–B) junction and MOS accumulation varactors.

The model topology, process parameters, and scaling equations for the C–B junction varactor are taken directly from the structure of the npn model. An example of this commonality is that the calculations for the extrinsic collector resistance in the npn are also used for the C–B varactor cathode resistance. The parasitic capacitance of the cathode to substrate is also modeled in the same manner as the collector–substrate capacitance. Modifications are made to reflect the PCell layout options for multiple anode and cathode devices, but the physical correlation between the npn and C–B varactor is preserved.

The n-MOS varactor is a tunable capacitor that uses a thin-oxide n-FET in an n-well with the n+ source and drain shorted together. The variation of the capacitance is controlled by the gate-to-diffusion potential that takes the silicon surface under the gate from depletion into accumulation. The typical device model that has been

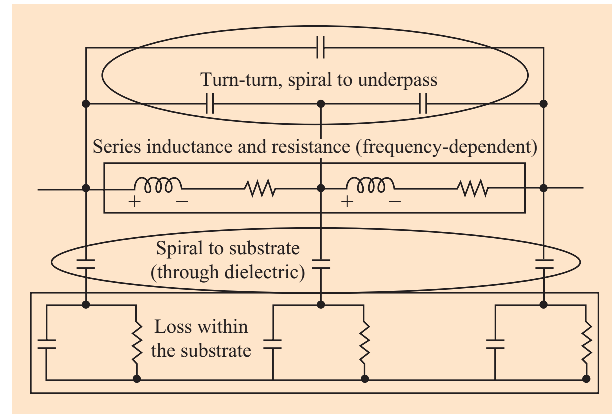

Figure 20

Inductor model subcircuit.

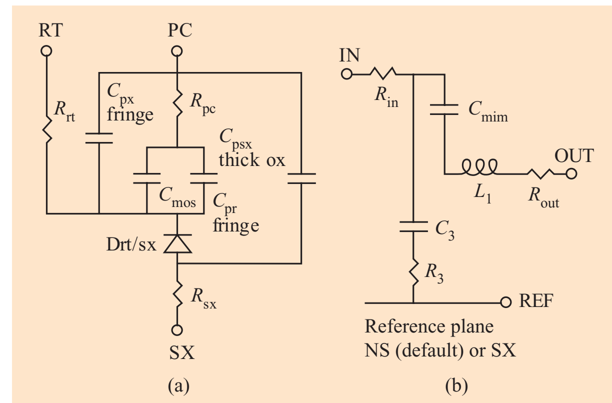

Figure 21

Typical (a) MOS and (b) MIM capacitor model subcircuits. The MIM model supports different reference-plane options, including buried subcollector (NS) or substrate (SX) options, as defined for the technology. The terms RT, PC, and Drt/sx respectively denote reachthrough (bottom plate), polysilicon (top plate), and reach-through to substrate parasitic diode.

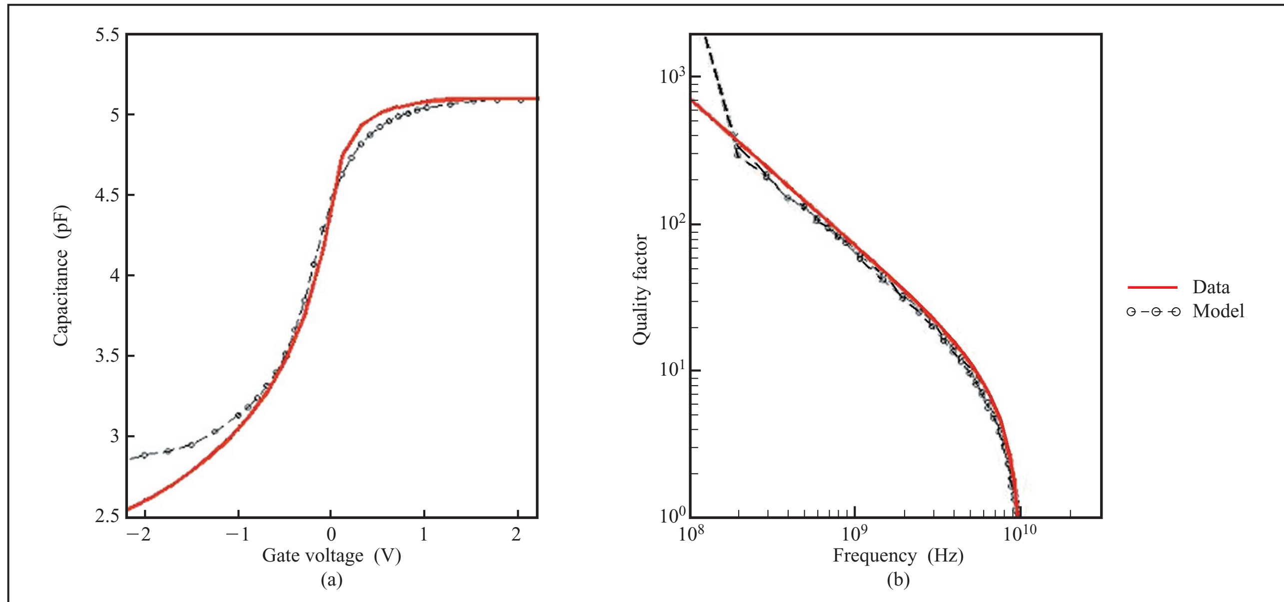

Figure 22

Typical capacitance vs. voltage (a) and frequency response (b) of an n-MOS varactor; 600 fingers, each having a length of 0.5 μm and a width of 5 μm.

supported for standard CMOS technologies makes use of a high-order polynomial to represent the variation in the capacitance as a function of applied voltage. Though it is possible to achieve a reasonable fit to the measured C–V curve, this implementation is highly nonphysical in nature and can lead to convergence problems during circuit analyses because the polynomial is numerically unstable near the boundaries of the voltage range over which the varactor is biased.

The most recent efforts to model this n-MOS varactor for the IBM SiGe BiCMOS and rf CMOS technologies have focused on the definition of a more physical calculation of the capacitance. In the improved model, the value of the primary capacitance comes from the total series capacitance of a fixed-oxide capacitor and a voltage-dependent capacitor that is modeled using the reverse bias C–V equation found in the standard SPICE diode element. While a varactor layout using a large polysilicon area would yield a high capacitance with better tunability, multiple finger devices of smaller area wired together are often used instead of a single large device to obtain a higher Q. This layout configuration dictates that the model must properly account for the fringe capacitance, the presence of which decreases the tunability. The subcircuit includes resistors to model the resistances of the polysilicon gates, the n-well, the metal lines connecting the gates, and the metal lines connecting the source/drains. Finally, the subcircuit includes a diode element to model the parasitic capacitance of the n-well to substrate junction.

The plots in **Figure 22** show a typical C–V curve and the frequency response for a large multi-finger n-MOS varactor. Though the model fit of the capacitance degrades steadily in the depletion region (gate voltage below -1 V), the use of this device is restricted to a maximum reverse bias voltage of -0.75 V because of the instability of the inversion layer. As a result, the inability of the model to properly predict the voltage dependence in this region is not significant.

Resistors

Key issues for modeling the different types of diffused and polysilicon resistors available in the IBM SiGe BiCMOS and rf CMOS technologies include accounting for the resistance of the silicided end regions, calculation of the parasitic capacitance, and determining the proper partitioning of the capacitance and resistance across the elements of the subcircuit. Standard dc measurements of resistance across voltage bias and temperature conditions are made on a set of resistors with varying width and length dimensions to extract parameters necessary to describe the total resistance as a function of geometry, bias, and temperature.

As with other passive devices, the frequency response of the model is verified using S-parameter measurements. The actual number of resistive and capacitive element

segments and the division of the total resistance and capacitance across those elements determine the frequency dependence predicted by the model. The standard model subcircuit contains four resistance elements to represent the body and end regions of the resistor. Parasitic capacitance is split across three capacitors connected from the nodes between the resistors to the body of the resistor. This model topology has been successful in predicting the resistance roll-off across frequency for both polysilicon and TaN resistors. *S*-parameter measurements to date have shown that, for resistors of similar geometry, the polysilicon resistor begins to roll off in a frequency range near 1–2 GHz, while the TaN resistor maintains its low-frequency resistance value up to nearly 10 GHz.

Additional model features include resistance mismatch as a function of geometry, support for parallel or series bars, and the flexibility to netlist resistors by specifying width and length or width and resistance.

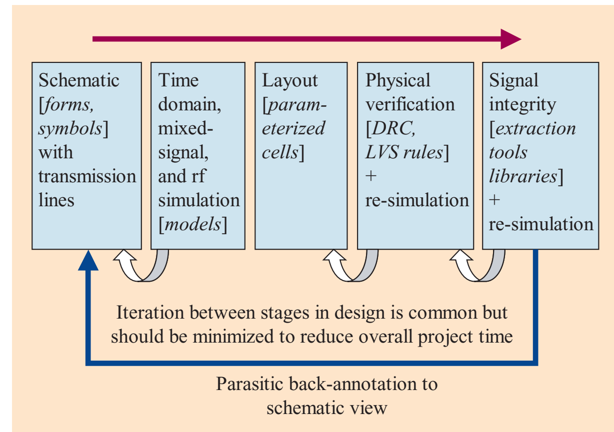

Figure 23

High-level view of rf design flow and its enablement.

5. Design automation

Circuit designers interface with the technology through a collection of design tools linked through a common design framework. The emphasis here is on a device-level design system with all devices indicated by basic layouts and base compact models, as opposed to higher-order design systems concerned with timing closure, methodology checks, etc. Macro-level designs are used for characterized circuit structures such as electrostatic discharge (ESD) designs. The key components of this device-level toolset are schematic capture, layout, physical verification tools including design rule checking (DRC), and layout versus schematic (LVS) tools, simulation tools for both time domain and frequency domain analysis, and interconnect parasitic extraction/modeling tools.

A typical design flow (as shown in **Figure 23**) has the circuit designer first using schematic entry and simulation, followed by layout, verification, parasitic extraction, and resimulation with the parasitics back-annotated into the schematic. The circuit schematic simulation is expected to differ from the parasitic extracted view, and some design iterations are usually required to meet the circuit performance and specifications.

An important aspect of the design framework is the ability to implement different vendor tools.

Design framework

System- and circuit-level design methodologies require many CAD tools for the various design stages, viz., system design and architecture, circuit-level design and simulation, physical layout, device and circuit design rule checking, physical-to-logical design comparison, and signal integrity. A design framework uses a design database to link these CAD tools. The design framework provides a high-level extension language which can manipulate the design database and provide a means for data translation between different tools, as well as a mechanism for developing custom software utilities and enhancements. The following pertain to examples of associated custom programs developed by IBM:

1. Parameters define the characteristics of a device and have an interdependence on other parameters. These parameters may become outdated in migration to new versions of the process kit. A procedure has been developed to propagate through the design hierarchy, recalculate the parameters, compare the parameters to the previous value, update the database, and report mismatches.
2. Programs have been developed to migrate designs from one technology to another, aiding in the development of a base design library of proven characterized circuits. Such programs map the devices and update the manufacturing design rules to the migrating technology. Parameters are calculated for the technology.
3. All design data for simulations and layouts are stored in a framework database. To enable ready access to the database, query programs to display and manipulate the database have been developed.
4. The design flow is controlled by library properties and automatic trigger functions. The library properties determine the technology and level of design. The trigger functions are procedures that are evaluated when the database is accessed via design editors.
   IBM has implemented an extensive verification flow incorporating several verification tools and multiple job runs per tool. A flow of required verification checks is established to ensure that the designs run each check.

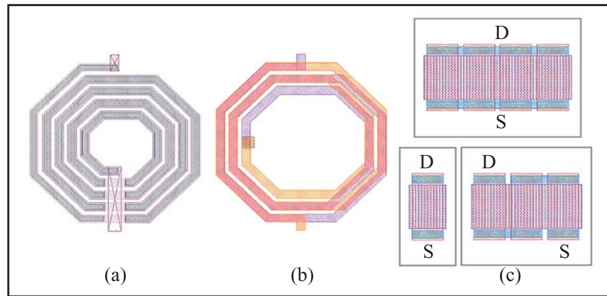

Figure 24

Inductors, implemented as PCells, stacked or in series and having various metal widths and numbers of turns (a) and (b). The resistors in (c) may be defined with series or parallel bars for matching and form-factor considerations. Devices may have optional guard rings and ground planes.

Multiple checks include design-rule checking, insulator integrity checks, pattern density checks, orthogonality checks, device extraction and comparison, and signal integrity checks.

The design flow discussed in this paper pertains to custom and semicustom circuit design. (The key design tool elements for an rf custom design were shown in Figure 23.) The overall goal is to reduce or eliminate hardware fabrication and design iterations through the use of optimal tools and design methodologies.

Design entry and simulation

For entry of a design into the framework, a library of hardware-characterized and modeled device primitives for, e.g., BJTs, MOSFETs, resistors, capacitors, inductors, diodes, and transmission lines are developed. Each component contains information stored in the database for schematic representation, parameter definitions, and netlisting.

Industry-standard symbolic representations of these active and passive components are used together with ideal components to input the electrical representation of the design (schematic). The artwork used to define a component may be configured by the designer to meet design requirements. For example, if a MOSFET device is to be used for a high-voltage application, a parameter is set, and the artwork changes characteristics such as color or shape to indicate usage.

Each device is defined by a model name, ports, and a set of parameters determining device dimensions and electrical properties. Parameters include geometric dimensions, device current rating, resistance/capacitance/inductance, device model parameters and subcircuit construction, reliability characteristics, and device options such as ground planes, trenches, and guard rings.

Each device is defined with a procedure to create a netlist entry of the component for a given circuit simulator. Device parameters may be entered as variables to enable a simulation sweep of the parameter to optimize designs. Parameters such as resistance values may be defined as variables, and a series of simulations automatically scale the resistance. The resistance parameter may be varied over a wide range and the results overlaid. This information can be used to optimize the circuit. Each component has a defined netlist procedure that writes a netlist with the nodes and parameters for a given circuit. These procedures are customized to incorporate specific enhancements such as device multiplicity.

A post-circuit-simulation analysis is performed to determine critical operating range and indicate to the designer the need to adjust device dimensions to tolerate circuit conditions based on current, voltage, or temperature. A custom program is used to define the design within the safe operating region of the circuit. Other simulation enhancement tools are defined to input manufacturing process bias and variability to target device performance after fabrication.

Digital and analog circuitry developed within a library may be written into parameterized cells. The benefit from schematic PCells is the ability to alter symbol graphics based on a property and to reuse underlying schematics on an instance-by-instance basis. For example, this methodology is employed in developing a library of ESD structures [42]. The base circuits may be expanded through device properties to increase the voltage protection by selecting the circuit type, number of stages, and number of devices. Parameters defining layout construction are defined during the schematic entry cycle and may be overridden in the physical environment. Once the circuit is designed in schematic form and optimized in simulation, the design is ready for physical layout.

Physical layout integration and interconnect is essential to rf design in order to control parasitic effects and match circuits for optimal performance [43]. All physical device geometric dimensions input for device placement are checked for physical design rules. This prevents the use of devices that do not conform to design rules and eliminates design rule checking (DRC) and layout versus schematic (LVS) errors. Early detection of these errors can significantly reduce design cycle time.

Physical design and verification

Parameterized cells

Schematic-driven layout helps streamline the physical design process. A schematic-driven layout tool places the

components in a predefined aspect ratio. Net connectivity information and device parameters are copied along with the device. The devices are thus placed within the layout exactly as they appear in the schematic.

PCells, as discussed earlier, are programmable component layouts that may be stretched through parameter inputs. Device primitives are defined as such cells. **Figure 24** shows various inductor PCells that have been implemented. When the schematic is transferred to the layout, these parameters generate the PCell component defined by the circuit designer. The PCell is designed in accordance with process design rules, and, when placed, the component is DRC-correct (discussed below) by construction. The design rules may be input into the database for access by all tools within the framework. These rules are input into the PCell as variables that enable easy migration to technologies with a database update. Layout options such as guard rings (Figure 17), well connections, ground planes (**Figure 25**) and multi-stripe connections are passed to the PCell as optional parameters.

Net information is input to a wiring utility that highlights connections between nodes. Auto-wiring tools are available, but at high frequencies custom wiring is recommended to control current flow, noise coupling, and parasitic effects. LVS ensures that all nets are correctly wired and that all devices appearing in the schematic are included in the layout.

Guard rings are designed for noise isolation and latchup protection [44]. The manual construction of these elements can be tedious, particularly from the ground-rule perspective. Additionally, to manually enclose all circuitry requires an intricate polygon layout. Using the interpretive language of the framework, design-rule-correct guard-ring paths may be rapidly constructed around sensitive circuitry using a PCell, as shown in **Figure 26**. Furthermore, the conductor may be cut in a designated region to allow wiring between design stages.

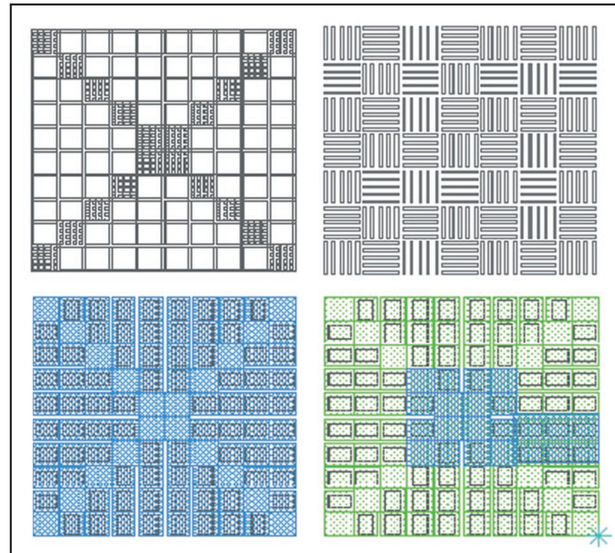

Figure 25

Variations in ground planes. The patterned ground planes shown are options for existing cells, and have been modeled for decreased coupling parasitics.

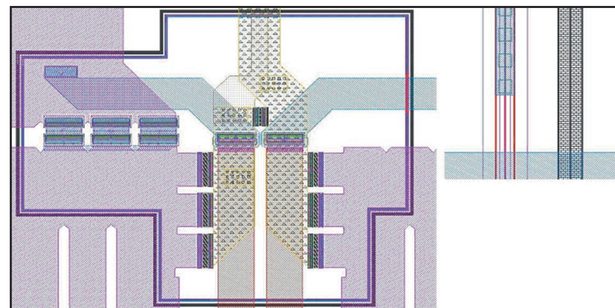

Figure 26

Circuits sensitive to noise or coupling effects may be isolated with a guard ring. A utility to draw a design-rule-correct guard ring encloses the circuitry shown. The conductor may be cut from the guard ring to allow a wiring channel (left).

Design-rule checking (DRC)

Design-rule verification involves checking mask-layer interactions to meet manufacturing processing and tooling requirements. Process layers are checked to width, length, area, and separation limits as well as overlay tolerances with other process steps such as those pertaining to the formation of contacts within diffusions, polysilicon regions, and interconnections. Robust manufacturability also pertains to voltage and current considerations, as well as density and uniformity. Some manufacturing process steps use high static charges that may require design rules to protect the wafer during fabrication through final packaging. Proper device construction is verified to ensure that all required design levels are present. Layers that may alter the device characteristics are checked for omission from the device, marking improper abutment or device interaction. Design-rule violations are flagged by the DRC tool as errors or warnings.

Copper-wiring-based chip interconnections necessitate special requirements for local pattern density [45], including minimum density percentages of certain layers and the ability to increase or reduce density. Density is increased by adding metal patterns to sparse regions or decreased by cutting holes in wide metal structures. This requires specialized verification procedures to ensure that

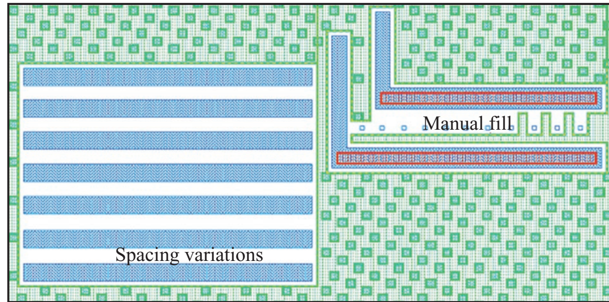

Figure 27

Typical layout obtained from the use of an automated pattern-filling process for achieving desired metallization density in nonsensitive circuits.

density requirements are met both regionally and globally. An analysis of the foundry pattern density routines must be made available to designers prior to data submission in order to customize fill patterns to ensure that critical regions are protected and circuit performance has not been affected [46] (**Figure 27** shows an example).

Layout versus schematic checking

LVS checking verifies that the physical design is consistent with the schematic. A netlist is generated from the schematic, and a physical netlist that identifies devices, parameters, and connectivity for each section is extracted from the layout. These netlists are compared or matched, and differences in parameters or connectivity are identified. It is important to have a common framework and netlisting environment to ensure that the design sections are accurately checked.

The layout must be processed through device-recognition and parameter-extraction routines that tolerate variable design styles in order to successfully identify all shape interactions that physically represent devices and connectivity. Standard comparison procedures may be used for simple devices, but many devices employ complex or unusual parameterization (especially in SiGe BiCMOS technologies). Therefore, custom comparison procedures may be required to process the results and report any differences. Custom programs to analyze the results enable the designer to diagnose design errors more efficiently.

Multiplicity is the repetition of like-sized devices with identical node connections. Device symbols define multiplicity as a parameter to indicate the number of devices in parallel. By this convention, m devices are simulated, while a single device instance is shown in the schematic. Schematic-driven layout creates m devices and tags, each device having a recognition shape. Device extraction recognizes devices in multiplicity separately from non-multiplicity devices. LVS checking compares the multiplicity parameter from the schematic and checks the quantity of devices and the exact device connection to that in the layout.

As technology and design complexity has increased, it has become necessary to adopt higher-performance "hierarchical" checking tools. "Flat" checking tools typically are limited to designs containing less than tens of thousands of devices. Hierarchical tools enable designers to check circuits containing millions of devices and are particularly useful for rf CMOS and BiCMOS configurations with high levels of integration.

Signal integrity

High-frequency IC design has traditionally required accurate modeling of active and passive devices, as described in previous sections. With the development of submicron CMOS and BiCMOS technologies and the move toward rf and microwave frequency applications, accurate estimation of parasitic effects for interconnect wiring, substrate interaction, and the package have become increasingly critical to meeting targeted design specification. These areas are part of an important field called signal integrity (SI). In this section, many of the significant SI issues are discussed, and best-in-class solutions are presented, including specifics on how they are implemented in the IBM SiGe and rf CMOS design kits. An in-depth background of this field can be found in [47].

For rf IC design, some key areas of concern are the following:

- *Interconnect delay* Solving the interconnect modeling problem is the first priority, for the full band of signal frequencies. This section discusses this area in detail.
- *Substrate noise coupling* Substrate noise coupling is an especially difficult problem to model in telecommunication ICs, and must be handled in a practical manner that is useful to the designer. A summary of the issues and guidelines is presented here.
- *Package delay and coupling (PCB)* IC parasitic wiring issues are highly integrated with package (and even PCB) effects and must be accounted for in simulations throughout the design phase. The topic is outside the scope of this paper.

Interconnect modeling

The need for interconnect awareness

As shown in **Figure 28**, the metal stack for IC processes is rapidly becoming complex for rf CMOS and BiCMOS designs. Traditionally, the modeling of the effects of these

metal interconnects has been treated as an afterthought at the last step of the design phase, with little consideration early in the design flow of the possible impacts of the interconnect impedance. Four key changes have led to the critical need for designers to consider interconnect effects early in the design process:

1. A more complex and varied metal/dielectric stack with denser pitches and thick dielectric/metal add-on modules (analog metal) [48].
2. Higher levels of integration, leading to larger ICs and longer chip-wide interconnects.
3. Higher signal frequencies, driven by shrinking technologies.
4. Lower V_{dd} levels, driven by shrinking technologies.

Add to this list the higher complexities of standards such as third-generation cell phone, SONET (optical network infrastructure), and the intense time-to-market pressures associated with these markets, and it becomes apparent that the effective modeling of the interconnect effects is critical to design success [47, 49–51].

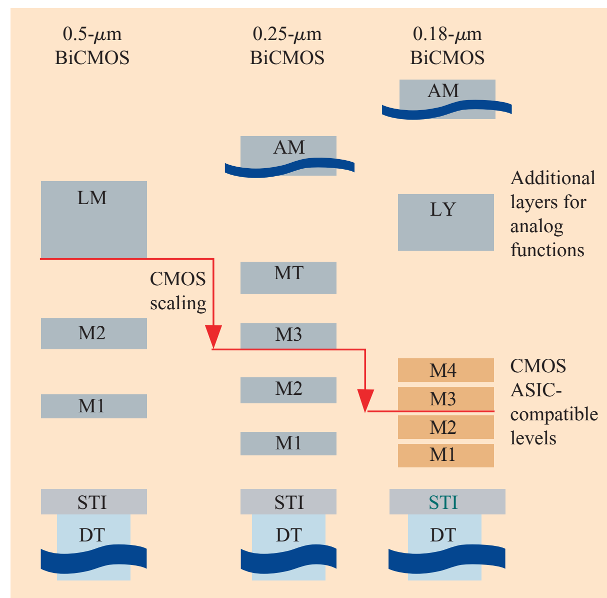

Figure 28

Evolution of interconnect metal stacks with device technology advances. LM denotes the last metal layer, AM denotes an analog metal layer, and LY denotes another metal layer. STI denotes a shallow-trench isolation, and DT denotes a deep trench. Aluminum layers are depicted in gray and copper layers in orange. Thicknesses of the LM, AM, and LY layers are nominally 1.4, 4, and 2.1 μm, respectively.

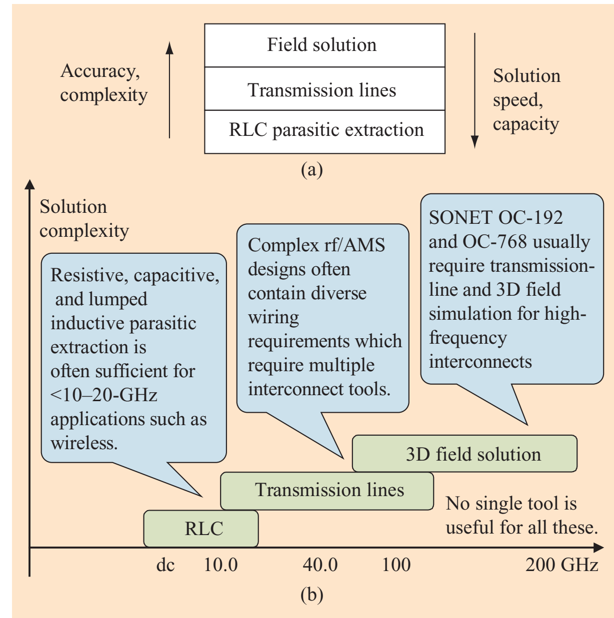

Figure 29

Tradeoffs in interconnect modeling methodologies: (a) Tradeoff between accuracy and speed when performing interconnect modeling; (b) interconnect solution alternatives.

Interconnect tool tradeoffs and alternatives

Since telecommunication IC designs typically contain many hundreds to thousands of devices with a broad frequency range, designers need to employ the appropriate tool on relevant nets. Generally, there is an effective accuracy versus speed tradeoff, as shown in **Figure 29(a)**:

- 3D field simulation is very accurate but very slow.
- 2D field solution is faster than 3D field solution.
- Resistive, capacitive, and inductive (*RLC*) extraction is fast and fairly accurate until nets become large.
- Transmission lines can offer good accuracy and performance.

The parasitic effects of interconnects vary depending on the size of the lines and the frequencies of the signals traveling through the lines. Frequency-dependent effects are threefold for rf interconnect design:

1. Skin resistance causes significant changes in impedance with changes in frequency.
2. Inductance becomes significant in the gigahertz frequency range.
3. In general, process variations cause unwanted performance degradation at higher frequencies.

Table 3 Interconnect modeling options.

Tool	*Functionality*	*Strength*	*Weakness*	*Application*
Parasitic *RLC* extraction	*RC* and *L* wiring estimation, in extraction phase	Useful for large numbers of nets, integrated in design	Inaccurate for higher frequencies; tools sometimes inaccurate	<10–20 GHz Typically *L* needed above *RC* for accuracy
Transmission lines	Distributed frequency and time domain interconnect effects	Arbitrary accuracy and efficiency	Typically poor physical design integration	Field solver engines more robust and accurate
2D field solution	Modeling of 2D and simple 3D structures	Accurate and efficient for 2D effects	Inaccurate for 3D and fringe effects	Packages, PCB, and off-chip effects; straight interconnects
3D field solution	Modeling of complex 3D structures	Very accurate	Very slow; often difficult to set up	>10 GHz Used for highest-frequency on-chip structures and complex structures such as vias

Figure 29(b) reflects how the different approaches to interconnect modeling can be mapped to the frequency range of the signal. Dimensions of the structures, noise-isolation requirements, impedance matching, and operating frequency all play a part in determining the best approach. For example, a quasi-static solution is accurate up to a quarter wavelength of the signal, above which a full-wave solution is preferable.

Table 3 summarizes the strengths, weaknesses, and applications of the various modeling techniques discussed here. The discussion covers IC, package, and PCB modeling and is focused on modeling approach and functionality as opposed to any specific tool offering.

Further parasitic tool requirements

There is a further requirement for estimation of parasitics at three points in the design flow:

1. Pre-layout estimation of interconnect parasitics. Schematic designers frequently need tools to estimate interconnect parasitics early in the design process—prior to layout. This can be done using transmission-line models or through parasitic *RLC* extraction tools.
2. Check of accuracy of integrated interconnect extraction tools. Questions naturally arise about whether such tools are accurate for a given structure.
3. Investigation of parasitics within devices and at device-to-interconnect boundaries. Accurate 3D *RLC* extraction tools are important for designers when investigating the parasitic effects in key circuit structures, such as the output stage of a power amplifier. This is important for both accurate modeling of parasitics and the compensation and optimization of parasitic effects through layout adjustment.

Verification of model and tool accuracy

Emphasis must be placed on accuracy and performance verification, whether it is directed at an *RLC* extraction tool or a transmission-line model, etc. Thorough testing across the range of interconnect layer combinations is required, and accuracy across frequency must also be characterized. Attention must be paid to outliers—the root cause of inaccuracies above the acceptable range, typically 10% for capacitance calculations. For example, five 25% errors in wide-metal-line testing may appear statistically insignificant in a test suite of 1500 structures. However, if those five structures are attributed to a certain size and formation of interconnects, there may be a bug in the model/tool that must be corrected. In contrast, outliers might also be attributed to unusual structures that are seldom encountered in practical design layouts.

Transmission-line development

There is a critical need for accurate integrated transmission-line models in telecommunication IC design. In high-end designs, such as SONET communications multiplexer/demultiplexer (MUX/DEMUX) design, there is a need for concurrent transmission-line design and circuit design—a demonstration of how parasitics are critically important to the performance of the design.

Transmission lines are by their nature constrained in structure, allowing for higher modeling accuracy [47]; typical examples are microstrip lines over a ground plane and coplanar waveguide structures. For rf CMOS designs, these structures may have to be different to allow for higher density. **Figure 30** shows (transmission-line) chip microstrip structures currently supported by IBM. These models are precharacterized parameterized frequency-dependent *RLC* lumped equivalent models that have

been correlated (using hardware test sites and field-solver simulations) up to 40 GHz for the single-line and 20 GHz for the coupled-line structures. They are usable in both the schematic and layout views of the design, through the use of schematic symbols and PCells. The models have also been implemented for use with periodic steady-state (PSS) and harmonic balance simulation algorithms.

The development of the transmission-line models at IBM has been introduced in the following order, providing for efficient deployment of a stable and accurate flow:

1. Black boxes, LVS clean single-ended and coplanar-coupled. At the schematic level only, providing a symbol, no models or model callbacks. DRC rules apply.
2. Simple but effective approach—straight lines, shielded below, optional lateral shielding. Single-ended and coplanar-coupled. Parameterized layout cells, LVS, model callbacks.
3. Frequency-dependent model development, including model–hardware correlation.
4. Discontinuities, such as vias, elbows, and junctions—through a phased-in approach.

Further enhancements and structures were planned at the time this paper was written. Examples include support of waveguide structures for very-high-frequency signals, with signals in the 100-GHz+ range, as well as optimized structures for leading-edge rf CMOS designs.

Substrate coupling issues

The principle of substrate coupling is illustrated in **Figure 31** [47, 52–54]. Substrate coupling occurs frequently in rf and mixed-signal designs but until recently has not had much of an effect on sensitive signals. With very complex new designs being built, such as wideband–code division multiple access (WCDMA), there is a need to design proactively, with substrate isolation in mind. In some cases, there is a need for up to 100 dB isolation between transmit and receive chains in the design, requiring an accurate understanding of the effects of the substrate.

Substrate, ground, power supply, and package noise injection are interlinked effects. Since both novice and experienced designers are facing difficult substrate noise issues today, studies have been performed on substrate coupling issues to ascertain whether general guidance data can be provided. Currently it appears that good design practices, as well as careful package choices, are the best way to avoid substrate coupling issues. This conclusion arises because initial investigations of commercially available substrate modeling tools (as of 2001) have not been encouraging.

In rf CMOS designs, where today analog/rf blocks are separately designed and then integrated into large digital designs, prediction and modeling of substrate effects in the block design phase is critical for integration to succeed. Practical specifications are needed for the block designer and integrator in order to ensure success [55].

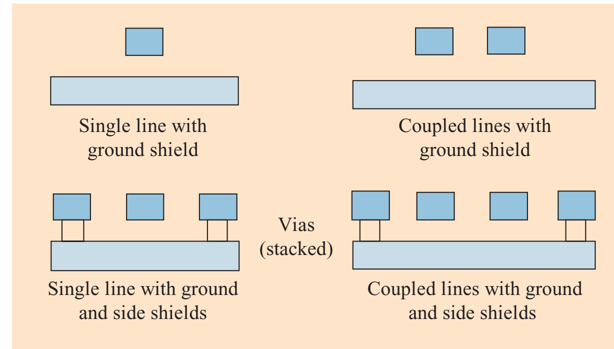

Figure 30

Transmission-line structures supported by IBM rf/analog design kit.

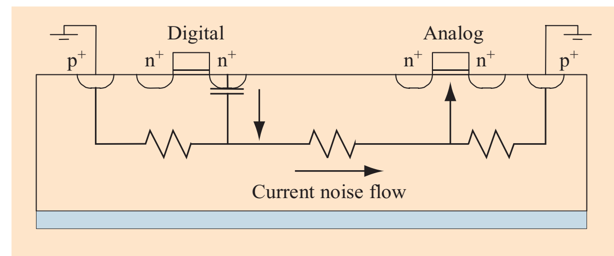

Figure 31

Illustration of the principle of substrate coupling. From [53], with permission; © 1999 IEEE.

For analysis of the isolation effectiveness of the substrate, with and without isolation structures such as deep trenches (DTs), a combination of TCAD simulations and test-site structures are continuously being developed and implemented. From these activities, guidelines for substrate impedance calculations and isolation effectiveness can be determined.

Substrate impedance calculation

For impedance extraction, with no isolation structures, some initial test-site data leads to Equations (9) and (10). This information pertains to ongoing work and is currently targeted at the IBM SiGe process [56]:

$$R_{sub} = [1152 + 330*\log(d) - 860*\log(w)] \text{ ohms;} \qquad (9)$$

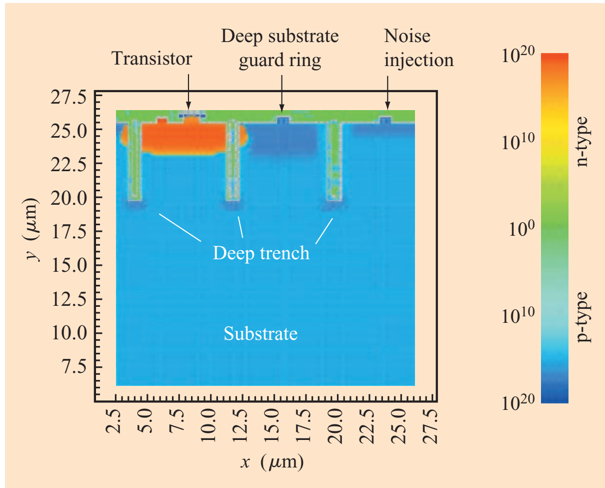

Figure 32

Structure and doping profile of typical two-dimensional large-signal simulation structure. A BiCMOS 6HP HBT is on the left, with a substrate injection port on the right. Separating them is a substrate guard ring, the heavily doped region of which has been extended into the substrate via additional implants. The temperature scale on the right depicts the next doping concentration.

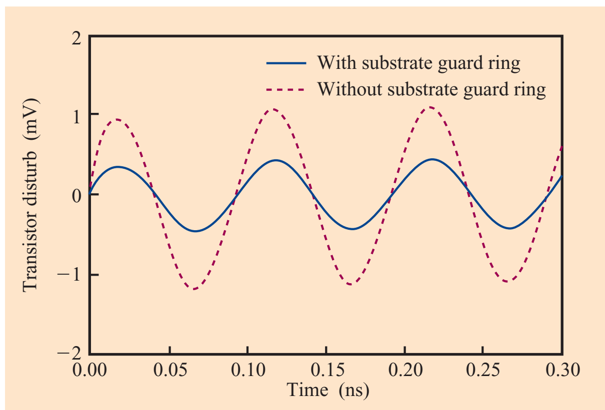

Figure 33

Electrical simulation of structure of Figure 32. The input was assumed to be a 100-mV/10-GHz sinusoidal noise signal injected as indicated in that figure.

$$C_{sub} = \epsilon_{Si} \times \rho_{sub}/R_{sub} \text{ farads.} \quad (10)$$

The substrate resistance guideline was extracted on the basis of measurements on a square substrate contact width between 8 μm and 75 μm. The distances between the substrate contacts (of identical sizes) measured were between 300 μm and 3000 μm. The substrate was 3000 μm thick.

Some interesting observations affecting accuracy that are seen in the measurement data include the following:

- The substrate resistance calculation using the relation $R = R_s\, l/w$ can be off by an order of magnitude.
- The substrate resistance reaches a plateau for device distances larger than about ~2000 μm.
- For distance less than ~2000 μm, substrate resistance varies logarithmically with the distance between devices.
- The substrate resistance varies inverse-logarithmically with the width of the devices.

Isolation effectiveness

Figure 32 shows an example of a TCAD simulation structure used to model the isolation characteristics of the substrate [56]. **Figure 33** shows corresponding results in the time domain. Many such simulations are required to understand the substrate isolation effect and develop process-specific guidelines for designers.

TCAD and test-site experience [56] leads to conclusions that strongly biased ac guard rings are very effective (up to 30 GHz+), with bipolar block layers and n-wells less effective but still useful. Deep-trench structures are useful for isolating <1-GHz digitally injected noise, but are transparent at higher frequencies. It is important to note that distance is also very effective for isolation of embedded cores. In addition, initial results show that substrate current runs at deeper levels in the substrate, i.e., is not limited to the surface of the substrate for distances equivalent to or greater than the substrate thickness. In these cases, the deep trench may be totally ineffective because of its 6-μm depth.

Design practices

Some good design practices that are commonly used include the following:

- Use of differential signal lines.
- Careful noise-aware floorplanning of chip layouts.
- Use of low-noise digital I/O buffers, since they are a key source of noise in the circuits.
- Use of dedicated ac-grounded circuit guard rings around key sensitive circuitry with careful routing and grounding; attention must be paid to noise coupling at bond pads and antenna effects.
- A clear understanding of the ac impedances to ground and supply structures, and package impedances.
- A clear understanding of the frequencies (and relative magnitudes) of signals being injected into the substrate, and how they can affect other parts of the design. For example, a low-frequency digital signal block may inject enough noise to affect an analog dc bias point on the other end of the IC, whereas a high-frequency noise component may not significantly affect a digital circuit.

- The IBM SiGe and rf CMOS design kits include a substrate methodology that allows designers to create multiple substrate regions efficiently. Such a methodology is needed for any design in which full or partial substrate effects must be modeled.

6. Concluding remarks

The rapid growth in the telecommunications market is driving increased integration levels and complex rf/analog mixed-signal integrated circuit products. Unlike digital design, rf/analog has typically required several fabrication iterations, resulting in longer design schedules and slower time to market. It is therefore critical that circuit designers and process technology personnel engage at a very early point in the product design cycle. A predictive modeling TCAD methodology has been described that produces supporting process and device simulation for the development of the IBM SiGe BiCMOS and rf CMOS process technologies, and in addition develops compact models and an early predictive design kit for evaluation before initial hardware is run. When hardware is run, it must be accurately characterized and modeled. Characterization combines data from many sources, including in-line data for a larger statistical base, bench-test low-frequency measurements, and rf measurements. The fundamental measurements and important aspects of the measurements to ensure accuracy and the methodology for creating statistical compact models have been described. The models must be embedded in a design environment to allow the circuit designer to make maximum use of the technology. A complete set of rf and analog mixed-signal simulation domain tools are required. A methodology has been described for creating a flexible design system supporting early test-site development as well as schematic entry, layout, DRC, and LVS tools. Library development begins at the inception of the design kit, with callbacks to the scalable statistical compact models that give the designer flexibility. For rf/analog design, it is imperative to pay close attention to signal integrity aspects, especially for wired applications at 40 Gb/s or higher. Transmission-line models and accurate *RLC* interconnect extraction are critical to this. Finally, a robust final test verification methodology has been described.

Acknowledgments

The authors wish to acknowledge the many individuals in IBM who contributed to the SiGe program and are responsible for the current position of SiGe technology in the marketplace. Specifically, we acknowledge the IBM Semiconductor Research and Development Corporation (SRDC) and Communication Research and Development Corporation for support. We acknowledge the many early access SiGe BiCMOS customers for helping shape design kit requirements. We acknowledge the SiGe BiCMOS and rf CMOS development organizations at the IBM East Fishkill, New York, and Essex Junction, Vermont, facilities for the technology development. We acknowledge the application engineering teams in Encinitas, California, and Essex Junction, Vermont, for supporting the technology and the design kits with customers. We acknowledge the RFIC design centers in Encinitas, California, and Lowell, Massachusetts, for technology direction and early evaluation of the technology and design kits. Finally, we acknowledge the IBM Business Line Management groups for marketing the technology and the design kits.

References

1. See the GSM Association Web site: *http://www.gsmworld.com/*.
2. See the National Institute of Standards and Technology Web site: *http://www.nist.gov/*.
3. See the Bluetooth Special Interest Group Web site: *http://www.bluetooth.com/*.
4. See *http://grouper.ieee.org/groups/802/11/*.
5. See the American National Standards Institute Web site: *http://www.ansi.org/*.
6. See the International Telecommunications Union Web site: *http://www.itu.int/*.
7. H. K. Gummel, "A Self-Consistent Iterative Scheme for One-Dimensional Steady State Transistor Calculations," *IEEE Trans. Electron Devices* **ED-11,** 455–465 (1964).
8. D. L. Scharfetter and H. K. Gummel, "Large-Signal Analysis of a Silicon Read Diode Oscillator," *IEEE Trans. Electron Devices* **16,** 64–77 (1969).
9. E. M. Buturla and P. E. Cottrell, "Two-Dimensional Finite Element Analysis of Semiconductor Steady-State Transport Equations," presented at the International Conference on Computational Methods in Nonlinear Mechanics, Austin, TX, 1974.
10. Siegfried Selberherr, *Analysis and Simulation of Semiconductor Devices*, Springer-Verlag, New York, 1984.
11. S. E. Laux, "Techniques for Small-Signal Analysis of Semiconductor Devices," *IEEE Trans. Electron Devices* **38,** 2028–2037 (1985).
12. M. V. Fischetti and S. E. Laux, "Monte Carlo Simulation of Electron Transport in Si: The First 20 Years," *Proceedings of the European Solid State Device Research Conference (ESSDERC'96)*, G. Baccarani and M. Rudan, Eds. (Editions Frontières, Gif-sur-Yvette Cedex, 1996), pp. 813–820.
13. J. B. Johnson, A. Stricker, A. J. Joseph, and J. A. Slinkmanet, "A Technology CAD Methodology for AC Performance Optimization of SiGe HBTs," *IEDM Tech. Digest*, pp. 234–238 (2001).
14. B. Jagannathan, M. Khater, F. Pagette, J.-S. Rieh, D. Angell, H. Chen, J. Florkey, F. Golan, D. R. Greenberg, R. Groves, S. J. Jeng, J. Johnson, E. Mengistu, K. T. Schonenberg, C. M. Schnabel, P. Smith, A. Stricker, D. Ahlgren, G. Freeman, K. Stein, and S. Subbanna, "Self-Aligned SiGe NPN Transistor with 285GHz f_{max} and 207GHz f_t in a Manufacturable Technology," *IEEE Electron Device Lett.* **23,** 258–260 (2002).
15. A. Joseph, D. Coolbaugh, M. Zierak, R. Wuthrich, P. Geiss, Z. He, X. Liu, B. Orner, J. Johnson, G. Freeman, D. Ahlgren, B. Jagannathan, L. Lanzerotti, V. Ramachandran, J. Malinowski, H. Chen, J. Chu, P. Gray, R. Johnson, J. Dunn, S. Subbanna, K. Schonenberg, D. Harame, R. Groves, K. Watson, D. Jadus, M. Meghelli, and A. Rylyakov, "A 0.18 mm BiCMOS Technology Featuring 120/100 GHz (f_T/f_{max}) HBT and ASIC-

Compatible CMOS Using Copper Interconnect," *Proceedings of the IEEE Bipolar/BiCMOS Circuits and Technology Meeting (BCTM)*, 2001, pp. 143–146.
16. S. Subbanna, J. Johnson, G. Freeman, R. Volant, R. Groves, D. Herman, and B. Meyerson, "Prospects of Silicon–Germanium Technology for Very High-Speed Circuits," *IEEE MTT-S Microwave Symposium Digest* **3,** 361–364 (2000).
17. C. C. McAndrew, J. A. Seitchik, D. F. Bowers, M. Dunn, M. Foisy, I. Getreu, M. McSwain, S. Moinian, J. Parker, D. J. Roulston, M. Schröter, P. van Wijnen, and L. F. Wagner, "VBIC95, the Vertical Bipolar Inter-Company Model," *IEEE J. Solid-State Circuits* **31,** No. 10, 1476–1483 (1996).
18. J. S. Dunn, D. C. Ahlgren, D. D. Coolbaugh, N. B. Feilchenfeld, G. Freeman, D. R. Greenberg, R. A. Groves, F. J. Guarín, Y. Hammad, A. J. Joseph, L. D. Lanzerotti, S. A. St.Onge, B. A. Orner, J.-S. Rieh, K. J. Stein, S. H. Voldman, P.-C. Wang, M. J. Zierak, S. Subbanna, D. L. Harame, D. A. Herman, Jr., and B. S. Meyerson, "Foundation of rf CMOS and SiGe BiCMOS Technologies," *IBM J. Res. & Dev.* **47,** No. 2/3, 101–138 (this issue, 2003).
19. G. Freeman, J. Kierstead, and W. Schweiger, "Electrical Parameter Data Analysis and Object-Oriented Techniques in Semiconductor Process Development," *Proceedings of the IEEE Bipolar/BiCMOS Circuits and Technology Meeting (BCTM)*, 1996, pp. 81–88.
20. M. Peterson, "Device Measurement and Characterization System: Creating Unity in a Diverse Characterization World," *IBM MicroNews* **7,** No. 2, 29–34 (May 2001).
21. P. J. Van Wijnen, "On the Characterization and Optimization of High-Speed Silicon Bipolar Transistors," Cascade Microtech, Inc., Beaverton, OR, 1995.
22. M. S. Gupta, "Power Gain in Feedback Amplifiers, A Classic Revisited," *IEEE Trans. Microwave Theor. Tech.* **40,** No. 5, 864–879 (May 1992).
23. S. A. Maas, *Nonlinear Microwave Circuits*, IEEE Press, New York, 1997, p. 370.
24. G. Gonzalez, *Microwave Transistor Amplifiers, Analysis and Design*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1984, pp. 174–180.
25. H. K. Gummel, "A Charge Control Relation for Bipolar Transistors," *Bell Syst. Tech. J.* **49,** 115–120 (1970).
26. G. M. Kull, L. W. Nagel, S. Lee, P. Lloyd, E. J. Pendergast, and H. Dirks, "A Unified Circuit Model for Bipolar Transistors Including Quasi-Saturation Effects," *IEEE Trans. Electron Devices* **32,** No. 6, 1103–1113 (1985).
27. See *http://www.iee.et.tu-dresden.de/iee/eb/comp_mod.html*.
28. M. Schroter and H.-M. Rein, "A Generalized Integral Charge-Control Relation and Its Application to Compact Models for Silicon-Based HBT's," *IEEE Trans. Electron Device Lett.* **40,** 2036–2046 (1993).
29. See *http://www.semiconductors.philips.com/models/*.
30. J. Berkner, "Bipolar Parameter Extraction," *IEEE Bipolar/BiCMOS Circuits and Technology Meeting Short Course Notes*, 2001.
31. W. Liu, X. Jin, J. Chen, M. C. Jeng, Z. Liu, Y. Cheng, K. Chen, M. Chan, K. Hui, J. Huang, R. Tu, P. K. Ko, and C. Hu, *BSIM3v3.2.2 MOSFET Model User's Manual*, University of California, Berkeley, 1999, p. 2-1.
32. K. Kundert, *The Designer's Guide to SPICE & SPECTRE*, Kluwer Academic Publishers, New York, 1995, pp. 167–176.
33. D. Foty, *MOSFET Modeling with SPICE*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1997, pp. 213–215.
34. J. Watts, C. Bittner, D. Heaberlin, and J. Hoffmann, "Extraction of Compact Model Parameters for ULSI MOSFETs Using a Genetic Algorithm," *Technical Proceedings of the Second International Conference on Modeling and Simulation of Microsystems*, 1999, pp. 176–179.
35. S. Parker, M. Kumar, A. Palaniswamy, S. Strang, and A. Bertolet, "Parasitic Estimation in Submicron Multifinger MOSFETs," *IBM MicroNews* **7,** No. 3, 32–36 (2001).
36. J. R. Long and M. A. Copeland, "The Modeling, Characterization, and Design of Monolithic Inductors for Silicon RF IC's," *IEEE J. Solid-State Circuits* **32,** 357–369 (March 1997).
37. K. O, "Estimation Methods for Quality Factors of Inductors Fabricated in Silicon Integrated Circuit Process Technologies," *IEEE J. Solid-State Circuits* **33,** 1249–1252 (August 1998).
38. C. P. Yue and S. S. Wong, "On-Chip Spiral Inductors with Patterned Ground Shields for Si-Based RF IC's," *IEEE J. Solid-State Circuits* **33,** 743–752 (May 1998).
39. S. S. Mohan, M. M. Hershenson, S. P. Boyd, and T. H. Lee, "Simple Accurate Expressions for Planar Spiral Inductances," *IEEE J. Solid-State Circuits* **34,** 1419–1424 (October 1999).
40. D. Harame, A. Joseph, D. Coolbaugh, G. Freeman, K. Newton, S. M. Parker, R. Groves, M. Erturk, K. Stein, R. Volant, C. Dickey, J. Dunn, S. Subbanna, H. Zamat, V. S. Marangos, M. M. Doherty, O. Schreiber, T. Tangib, D. A. Herman, and B. S. Meyerson, "The Emerging Role of SiGe BiCMOS Technology in Wired and Wireless Communications," *Proceedings of the Fourth IEEE International Caracas Conference on Devices, Circuits, and Systems*, Aruba, April 17–19, 2002.
41. J. R. Long, "Monolithic Transformers for Silicon RF IC Design," *IEEE J. Solid-State Circuits* **35,** 1368–1382 (September 2000).
42. S. Voldman, S. Strang, and D. Jordan, "An Automated Electrostatic Discharge Computer Aided Design System with the Incorporation of Hierarchical Parameterized Cells in BiCMOS Analog and RF Technology for Mixed Signal Applications," presented at the Electrical Overstress/Electrostatic Discharge (EOS/ESD) Symposium, Electrostatic Discharge Association, Rome, NY, January 2002; see *www.esda.gov/*.
43. D. Ahlgren and J. Dunn, "SiGe Comes of Age," *IBM MicroNews* **6,** No. 1 (2000).
44. IBM Technology Development, *BiCMOS7HP Design Manual*, Section 3.28.3, "Guardring Layout," February 2002.
45. IBM Technology Development, *CMOS7SF Design Manual*, Section 2.10.2, "Copper Pattern Density and Layout Requirements," 2000.
46. S. Strang and D. Jordan, "Custom Verification Using DIVA and DFII," presented at the International Cadence Users Group Conference, San Jose, CA, December 2001.
47. *Signal Integrity Effects in Custom IC and ASIC Designs*, R. Singh, Ed., IEEE Press and Wiley & Sons, New York, 2001.
48. S. A. St. Onge, D. L. Harame, J. S. Dunn, S. Subbanna, D. C. Ahlgren, G. Freeman, B. Jagannathan, J. Jeng, K. Schonenberg, K. Stein, R. Groves, D. Coolbaugh, N. Feilchenfeld, P. Geiss, M. Gordon, P. Gray, D. Hershberger, S. Kilpatrick, R. Johnson, A. Joseph, L. Lanzerotti, J. Malinowski, B. Orner, and M. Zierak, "A 0.25 μm SiGe BiCMOS Mixed-Signal RF Production Technology Featuring a 47 GHz f_T HBT and 0.18 μm L_{eff} CMOS," *Proceedings of the IEEE Bipolar/BiCMOS Circuits and Technology Conference (BCTM)*, 1999, pp. 117–120.
49. A. Deutsch, P. Coteus, G. Kopcsay, H. Smith, C. Surovic, B. Krauter, D. Edelstein, and P. Restle, "On-Chip Wiring Design Challenges for Gigahertz Operation," *Signal Integrity Effects in Custom IC and ASIC Designs*, R. Singh, Ed., IEEE Press and Wiley & Sons, New York, 2001, Part 2, pp. 161–187.
50. B. Kleveland, T. Lee, and S. Wong, "50-GHz Interconnect Design in Standard Silicon Technology," *IEEE MTT-5 Microwave Symposium Digest* **1,** 1913–1916 (1998).

51. W. Kao, C.-Y. Lo, M. Basel, and R. Singh, “Parasitic Extraction: Current State of the Art and Future Trends,” *Proc. IEEE* **89,** No. 5, 729–739 (2001).
52. D. Su, M. Loinaz, S. Masui, and B. Wooley, “Experimental Results and Modeling Techniques for Substrate Noise in Mixed-Signal Integrated Circuits,” *Signal Integrity Effects in Custom IC and ASIC Designs*, R. Singh, Ed., IEEE Press and Wiley & Sons, New York, 2001, Part 4, pp. 339–349.
53. R. Singh, “A Review of Substrate Coupling Issues and Modeling Strategies,” *Proceedings of the IEEE Custom Integrated Circuits Conference*, May 1999, pp. 491–499.
54. M. Pfost and H. M. Rhein, “Modeling and Measurement of Substrate Coupling in Si-Bipolar IC’s up to 40GHz,” *IEEE J. Solid-State Circuits* **33,** No. 4, 582–591 (1998).
55. *Signal Integrity Specifications Document Version 1.0*, Virtual Socket Interface Alliance; see *http://www.vsi.org/*.
56. Raminderpal Singh, Youri Tretiakov, Jeffrey B. Johnson, Susan Sweeney, Robert Barry, Mukesh Kumar, Mete Erturk, John Katzenstein, Carl Dickey, and David Harame, “Parasitic Modeling and Noise Mitigation in Advanced RF/Mixed-Signal Silicon Germanium Processes,” *IEEE Trans. Electron Devices* (May 2003), in press.

Received April 5, 2002; accepted for publication November 11, 2002

David L. Harame *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (dharame@us.ibm.com).* Dr. Harame received a B.A. degree in zoology from the University of California at Berkeley in 1971 and an M.S. degree in zoology from Duke University in 1973. He also received an M.S. degree in electrical engineering from San Jose State University in 1976 and an M.S. degree in materials science and a Ph.D. in electrical engineering, both from Stanford University, in 1984. He then joined the IBM Bipolar Technology group at the IBM Thomas J. Watson Research Center, where he worked on the fabrication and modeling of Si-based integrated circuits. His research interests at Yorktown Heights included Si transistors, SiGe-channel FET transistors, npn and pnp SiGe-base bipolar transistors, complementary bipolar technology, and BiCMOS technology for digital, analog, and mixed-signal applications. In 1993 he joined the IBM Semiconductor Research and Development Center in the Advanced Semiconductor Technology Center (ASTC) in Hopewell Junction, New York, where he was responsible for the development of SiGe technology for mixed-signal applications. He managed SiGe BiCMOS technology development at the ASTC through 1997. In 1998 Dr. Harame joined the IBM manufacturing organization in Essex Junction, Vermont, where he managed a SiGe technology group and the installation of the 0.5-mm SiGe BiCMOS process into manufacturing. In 1999 he co-managed the qualification of a 0.25-μm SiGe BiCMOS as well as 0.18-μm SiGe BiCMOS and two derivative SiGe BiCMOS technologies. In 2000, he became the Senior Manager of the RF Analog Modeling and Design Kit Department. Dr. Harame is a Distinguished Engineer of the IBM Corporation, a Senior Member of the IEEE, an Executive Committee member of the Bipolar/BiCMOS Circuits and Technology Meeting (BCTM), and a member of the Compact Model Council.

Kim M. Newton *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (kimnewt@us.ibm.com).* After receiving a B.S. degree in electrical engineering from Northwestern University in 1984, Ms. Newton joined IBM at the Essex Junction facility to work on model development. She has been involved in the modeling and characterization of IBM SiGe BiCMOS technologies, with an emphasis on bipolar devices and passives, and now manages the compact modeling team.

Raminderpal Singh *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (raminder@us.ibm.com).* Dr. Singh received a Ph.D. degree in efficient substrate modeling techniques in 1997. He then worked with Cadence Design Systems, initially as the lead engineer on the Substrate Coupling Analysis tool and then later as an ASIC Signal Integrity Methodology Expert and a Mixed-Signal CAD Flow Expert. Since March 2001, he has been a Senior Engineer at the IBM RF/AMS Design Kit group, where he leads various projects related to transmission-line model development, substrate modeling, and rf/mixed-signal design methodologies. Dr. Singh has authored numerous technical publications in the area of signal integrity, including the recent book *Signal Integrity Effects in Custom IC and ASIC Designs* (IEEE Press, 2001). He is a co-chairman of the AMS Signal Integrity Working Group in the VSIA SOC Standards body and a Senior Member of the Institute of Electrical and Electronics Engineers.

Susan L. Sweeney *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (slsweene@us.ibm.com).* Ms. Sweeney received a B.A. degree in physics from Saint Olaf College, Northfield, Minnesota, in 1984 and an M.S.E.E. degree in electrical engineering from New Mexico State University, Las Cruces, in 1995. She joined IBM in 1999 as a member of the Communications Research and Development Center (CRDC) and is currently an RF Characterization Engineer in the SiGe Device Development Laboratory, specializing in large-signal, distortion, and noise-parameter measurements at high frequencies.

Sue E. Strang *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (sstrang@us.ibm.com).* After receiving a B.S. degree in electrical engineering from the University of Pittsburgh in 1982, Ms. Strang joined IBM at the Burlington facility to work on DRAM design. She subsequently worked on satellite control and telemetry systems in the IBM Federal Systems Division facility in Gaithersburg, Maryland. She is currently a Senior Engineer, working on design-kit CAD development for the SiGe and rf technologies, designing device libraries and simulation. In 2000, she received an IBM Outstanding Innovation Award for her work on the first commercial SiGe heterojunction bipolar transistor (HBT) technology (SiGe5E). Ms. Strang is a coauthor of a patent for practical distributed transmission-line analysis and of several papers presented at International Cadence Users Group Conferences on CAD development. She is IC chair at the International Cadence Users Group Conference.

Jeffrey B. Johnson *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (jbj@us.ibm.com).* After receiving a Ph.D. degree in electrical engineering from Carnegie Mellon University in 1987, Dr. Johnson joined IBM at the Burlington facility to work on technology CAD. He is currently a Senior Technical Staff Member.

Scott M. Parker *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (smparker@us.ibm.com).* Mr. Parker manages the SiGe TCAD modeling group at the Burlington facility. He had previously worked on ADC design and modeling support for the IBM Read Channel products group and on device modeling in the BiCMOS Device Modeling group. His areas of expertise include device parasitics, BSIM parameter extraction, and device mismatch.

Carl E. Dickey *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (cdickey@us.ibm.com).* Mr. Dickey is a Development Engineering Manager in the IBM Communications Research and Development Center at the Burlington facility, currently managing the back-end design automation development team for IBM SiGe and rf CMOS technologies. He is also the strategy lead for the design kits group.

Mete Erturk *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452, and Department of Electrical Engineering, University of Vermont, Burlington, Vermont 05405 (merturk@zoo.uvm.edu).* After receiving an M.S. degree in electrical engineering from the University of Maryland in 2000, Mr. Erturk joined IBM at the Burlington facility to work on compact modeling of on-chip passive devices. He is currently a Ph.D. candidate at the University of Vermont. His technical interests include devices in rf and analog environments, on-chip inductors, and nontraditional methods in computation using semiconductors.

Greg J. Schulberg *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (gschulb@us.ibm.com).* Mr. Schulberg received a B.S. degree in 1998 and an M.Eng. degree in 1999, both in electrical engineering, from Cornell University. He is currently working on SiGe design kit testing at the Burlington facility.

Donald L. Jordan *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (djordan@us.ibm.com).* After receiving an A.S. degree in electrical engineering from Vermont Technical College in 1966 and a B.S. degree in electrical engineering from the University of Vermont in 1970, Mr. Jordan joined IBM at the Burlington facility to work on manufacturing diffusion engineering. He continued his work in in-line characterization of the manufacturing process and later in test data development support compilers for testers. Mr. Jordan is currently a Senior Engineer, working on design kit CAD development for the SiGe technology. In 1984 he received an IBM Outstanding Innovation Award for his work on the specials test data supply system, and in 2000 he received an IBM Outstanding Technical Achievement Award for his work on the first commercial SiGe HBT technology (SiGe5E). He is a coauthor of a patent for practical distributed transmission-line analysis and coauthor of several papers presented at International Cadence Users Group Conferences on CAD development.

David C. Sheridan *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (davsheri@us.ibm.com).* Dr. Sheridan received B.S. (1995), M.S. (1997), and Ph.D. (2001) degrees in electrical engineering from Auburn University, where his research involved simulation and fabrication of ultrafast SiC rectifiers. He subsequently joined the SiGe predictive modeling team at the IBM Microelectronics Division in Essex Junction, Vermont. Currently he is performing SiGe heterojunction bipolar transistor (HBT) process and device design for next-generation wireless technologies, as well as developing tools for the efficient extraction of predictive device models through the use of genetic algorithms. Dr. Sheridan has authored or coauthored more than 20 technical papers; he is a member of the Institute of Electrical and Electronics Engineers.

Michael P. Keene *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (keenemp@us.ibm.com).* Mr. Keene is a Development Engineering Manager in the IBM Communications Research and Development Center at the Burlington facility. He currently manages the front-end design automation development team for IBM SiGe and rf CMOS technologies.

John Boquet *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (boquet@us.ibm.com).* After receiving a B.S. degree in electrical engineering from Tulane University in 1980, Mr. Boquet joined the IBM Federal Systems Division in Manassas, Virginia, to work on VLSI and radiation-hardened ASICs. In 1994, he joined the IBM ASICs group of the IBM Microelectronics Division Burlington facility to manage the ASIC Verification Department. In 2000, he joined the SiGe Design Kit Development area to manage design kit verification, release, project management, and business operations.

Robert A. Groves *IBM Microelectronics Division, East Fishkill facility, Route 52, Hopewell Junction, New York 12533 (grovesr@us.ibm.com).* Mr. Groves received a B.S.E.E. degree from the State University of New York in 1996. He joined the IBM Corporation at the East Fishkill facility in 1989 as a Development Laboratory Technician. Since 1994, he has worked on SiGe technology development, with an emphasis on high-frequency modeling and characterization. His current focus is on interconnections, capacitors, and inductors (particularly integrated spiral inductor optimization and modeling) for microwave integrated circuits.

Mukesh Kumar *IBM Microelectronics Division, Burlington facility, 1000 River Road, Essex Junction, Vermont 05452 (mk@us.ibm.com).* Dr. Kumar is an Application Engineer in the RF and Analog Product Support group. He received a Ph.D. degree in electrical engineering in 1996 from the University of Cincinnati, where he developed integrated optical channel waveguides in compound semiconductors using focused ion-beam-induced mixing techniques. He subsequently joined the SiGe device modeling group at the IBM Burlington facility and led device model development in the initial IBM SiGe technologies, focusing on rf CMOS model development. His contributions have included the design of integrated CMOS VCOs for high-speed serial link communications.

Dean A. Herman, Jr. *IBM Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (dean@us.ibm.com).* Dr. Herman is Manager of Analog and Mixed Signal Operations, Office of the Chief Technologist, IBM Technology Group. Until recently he managed the groups doing research and development on the IBM SiGe technologies and on the design kits for all IBM analog and mixed-signal technologies; this effort included the IBM high-frequency measurement and modeling efforts. Dr. Herman received a B.S. degree in engineering from the U.S. Military Academy at West Point, New York, in 1960, and M.S. and Ph.D. degrees in electrical engineering from Stanford University in 1967 and 1975, respectively. He taught in the Department of Electrical Engineering at West Point for 15 years and was Professor and Deputy Head of the Department prior to joining IBM in 1986. He is a Senior Member of the Institute of Electrical and Electronics Engineers.

Bernard S. Meyerson *IBM Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (meyerson@us.ibm.com).* Dr. Meyerson is Chief Technologist, IBM Technology Group, also serving as Vice President, IBM Communication Research and Development Center. He received a Ph.D. degree in solid-state physics from the City College of the City University of New York in 1980 and subsequently joined IBM at the Thomas J. Watson Research Center to work on the chemical fundamentals of electronic materials thin-film deposition. His work led to the development of ultrahigh-vacuum/chemical vapor deposition (UHV/CVD), a technique which proved successful for the epitaxial deposition of a wide variety of Si-based alloys under conditions and having chemical compositions which lay well outside accepted bounds on such processes and materials. Most significant was the deposition of production quantities of SiGe epitaxial layers at temperatures of 500°C and below, resulting in the development of extremely high-performance SiGe heterjunction bipolar transistors (HBTs). Under his leadership, a complete SiGe HBT technology was developed, demonstrating high levels of integration at performance characteristics representing a discontinuity in historical trends for bipolar technology. Although the technology was initially intended for computing applications, Dr. Meyerson refocused it into a communication-centric technology. By 1992, he had launched a series of now well-publicized IBM alliances to leverage this technology trend, over the years incorporating organizations as diverse as Alcatel, Analog Devices, Hughes, Harris Semiconductor (now Intersil), Leica, National Semiconductor, Nortel, Tektronix, and many others. Early adoption of the technology has led to numerous commercial and consumer products, 40-Gb/s SONET standard data systems and 802.11b 11-Mb/s wireless LAN cards being limited examples. IBM's lead into this marketplace has resulted in general technology acceptance after ten years of controversy, with all major communications technology providers announcing upcoming, competing market entries. In 1998, the IBM analog and mixed-signal technology development was consolidated under Dr. Meyerson. In 2002, he became Vice President of the Communications Research and Development Center, a multinational organization encompassing IBM's worldwide communications technology and circuit design efforts and supporting a rapidly growing worldwide customer base. In 1992, Dr. Meyerson was designated as an IBM Fellow. He is also a Fellow of the American Physical Society and the IEEE. He holds more than 40 patents and has published several hundred technical papers. He has received numerous awards for his work, including the Materials Research Society Medal (1991) and the Electrochemical Society Electronics Division Award (1993); he was cited as "Inventor of the Year 1997" by the New York State Legislature and was honored as the 1999 "United States Distinguished Inventor of the Year" by the U.S. Patent and Trademark Office. He recently received the 1999 IEEE Ernst Weber Award for the body of work culminating in the commercialization of SiGe-based communications technology. In December 1999, he received the IEEE Electron Devices Society J. J. Ebers Award, and in 2002 he was inducted into the National Academy of Engineering. His team's work has been featured in a long-term ongoing study by the Lally School of Business Management, exploring the management of discontinuous innovation within large organizations.