

Embedding Mixed-Signal Design in Systems-on-Chip

Innovative approaches and new design methodologies are needed to integrate digital, analog and RF components in CMOS systems-on-a-chip smaller than 100 nm.

By JAN M. RABAEY, *Fellow IEEE*, FERNANDO DE BERNARDINIS, ALI M. NIKNEJAD, BORIVOJE NIKOLIĆ, *Senior Member IEEE*, AND ALBERTO SANGIOVANNI-VINCENTELLI, *Fellow IEEE*

ABSTRACT | With semiconductor technology feature size scaling below 100 nm, mixed-signal design faces some important challenges, caused among others by reduced supply voltages, process variation, and declining intrinsic device gains. Addressing these challenges requires innovative solutions, at the technology, circuit, architecture, and design-methodology level. We present some of these solutions, including a structured platform-based design methodology to enable a meaningful exploration of the broad design space and to classify potential solutions in terms of the relevant metrics.

KEYWORDS | Mixed analog-digital integrated circuits

I. INTRODUCTION

The “systems-on-chip” (SoC) concept has been quite successful, particularly in the digital arena. Hundreds of designs combining numerous processors, configurable and dedicated accelerators, I/O interfaces, and on-chip networks have been designed, fabricated, and brought to market. Reduction of overall system cost, chip size, and power together with performance increase have been the main driving force behind this continuing integration of functionality onto a single die. Yet, true SoC integration requires not only the inclusion of the digital computing, but also a seamless integration of the mixed-signal (MS) processing that embodies the periphery to the outside world and the environment. Some very complex and

sophisticated mixed-signal SoCs have been designed and manufactured over the past few years. This is true in the domain of wireless and/or wireline communications, where cost considerations have led to the combination of high-frequency RF components, high-performance analog, and complex digital functionality [1]–[3]. Besides cost reductions, SoC integration has led to some additional rewards—for example, lower jitter and phase noise is attainable by not having to go off chip, and large power savings can be achieved by eliminating two low-voltage differential signaling (LVDS—an interconnect standard) interfaces for high-performance analog–digital converters (ADCs), to name some. Yet, each of these mixed-signal SoCs is typically the result of a painstakingly executed and time-consuming custom design process. Most MS modules are still custom designed using tool flows that are only modestly more advanced than the ones we used over the last decades. Reuse of modules, design exploration, and the adoption of higher abstraction levels, going beyond logic synthesis—techniques that have been proven to be so effective in the digital world—have made little inroad. Analog building blocks are still “dropped” into the SoCs as hard macros, and integration of analog modules is limited to medium-performance hard-wired blocks with frozen layouts.

If the situation is difficult today, all indications are that in the future the integration of analog and RF components using an SoC design style will be next to impossible both economically and technically, if we do not change radically the way we think about mixed-signal design. Technology scaling, which is primarily driven by the digital components, is making the life of the mixed-signal designer increasingly harder. Power and reliability concerns require a scaling of the supply voltage proportional to the scaling in feature sizes (even though this is bound to saturate in the near future). This dramatically reduces the design margins

Manuscript received December 1, 2005; revised January 6, 2006.

J. M. Rabaey, **A. M. Niknejad**, **B. Nikolić**, and **A. Sangiovanni-Vincentelli** are with the University of California at Berkeley, Berkeley, CA 94720 USA (e-mail: jan@eecs.berkeley.edu; bora@eecs.berkeley.edu; niknejad@eecs.berkeley.edu; asv@eecs.berkeley.edu).

F. de Bernardinis is with the University of Pisa, Pisa 56126, Italy (e-mail: f.debernardinis@ing.unipi.it).

Digital Object Identifier: 10.1109/JPROC.2006.873609

available to the analog designer. Design margins are further compromised by the increase in process variability, which is becoming pronounced in the sub-100-nm space. These technical difficulties make the adoption of advanced design methodologies borrowed from digital design even more challenging than it is today. Lacking an effective design methodology when complexity grows yields inevitably increased nonrecurring engineering costs that will make the design of mixed-signal SoC impractical. Some suggest that integration of analog and RF components may be best accomplished by abandoning the SoC paradigm, and by turning to advanced packaging techniques to reap some of the benefits of SoC integration [4]. We, on the other hand, believe that exploring ways to make mixed-signal walk hand-in-hand with scaled digital is of paramount importance. Doing so is bound to lead to novel ways of realizing mixed-signal functionality, and will ensure that the mixed-signal scaling will continue well into the next decade.

The goal of this paper is explore a number of plausible paths that lead to truly scaleable mixed-signal SoCs. First, the technology challenges and constraints of mixed-signal design in a sub-100-nm setting are discussed. Next, we analyze the impact of these trends on a number of mixed-signal and RF building blocks. We proceed with an overview of when and where scaled mixed-signal makes sense, and we illustrate it with some concrete examples. Finally, the paper explores the options on how to raise the abstraction levels in mixed-signal design, which we consider to be of essence for scaleable flexible analog to come to realization. We conclude with some reflections and conclusions.

II. MIXED-SIGNAL DESIGN IN A SUB-100-NM TECHNOLOGY SETTING: CHALLENGES AND APPROACHES

In this section, we present the challenges that the sub-100-nm technology brings to mixed-signal SoCs. We will then examine a set of possible solutions.

A. Challenges

1) *Technology Divergence*: Cost reduction has always been and still is the fundamental motivation behind technology scaling [88]. This is accomplished by integrating more functions on the same die, or by moving an existing product to a finer resolution technology (obviously assuming that the manufacturing cost per wafer does not increase significantly). As a side benefit, technology scaling of digital functionality also led to increased performance and reduced power per function for each consecutive technology node. Cost concerns were again the main motivation behind the move to the systems-on-chip paradigm. By integrating complete solutions on a single die leads in general to a reduced system cost. In

addition, SoC solutions often lead to improved performance or reduced power. For example, the faster and wider interconnects on a chip enable tighter integration of components, which often leads to more efficient and effective system architectures.

This proposition may not hold in the current scaling regime. This point is best understood by introducing the concept of the *optimal technology*. The simplest set of parameters that characterize a technology are the minimum channel length, transistor threshold, the nominal supply voltage and the reliability-imposed voltage constraints (one may also include the f_T in this set, although it is to a first degree proportional to some of the above parameters). An optimal technology can be determined as the one that best fits a function or module (such as digital, SRAM, DRAM, analog, etc). In the case of an SoC, the optimal technology must fit a complete chip, that is the ensemble of its building blocks. It has been a common practice to choose the technology option most appropriate for the majority of the chip. As long as technologies offer plenty of headroom, or when most components on a die have similar characteristics, the impact of doing so bears little penalty or overhead. For instance, when the supply voltages are high enough (as was the case in the past), almost all analog functions can be implemented with little or no penalty on the same die as the digital, or even the memory functions.

In the last few years, we have entered the so-called *power-limited scaling regime*, where power (energy) considerations are as important as the overall performance of the product. Subsequently, power has become a dominant factor in defining the process parameters [5]–[8]. One unfortunate side effect of this regime is that scaling may not offer the improvements in performance (as measured in device speed) and/or power we were accustomed to, and cost considerations are the predominant (or maybe the only) reason behind further scaling. Another crucial property of the power-limited scaling regime is that the optimal technologies for the different functions or modules are gradually diverging.

To illustrate this point, consider, for instance, the “optimal” supply voltage. One of the most prominent properties of technology scaling is the accompanying reduction in supply voltages. According to the constant-field scaling theory [9], the supply voltage is supposed to scale with the same factor as the minimum device feature. This has been true in general, although, for a while, supply voltages were kept at a higher value to provide a performance boost (Fig. 1). For the sub-100-nm regime, the International Technology Roadmap for Semiconductors (ITRS) [10] predicts a slowdown in the scaling of the maximum supply voltage, which can be mostly attributed to the inability to further reduce the threshold voltage due to leakage considerations and process variations. The ITRS predicts that the finest resolution devices will be able to sustain 1 V or higher supplies down to the 45-nm

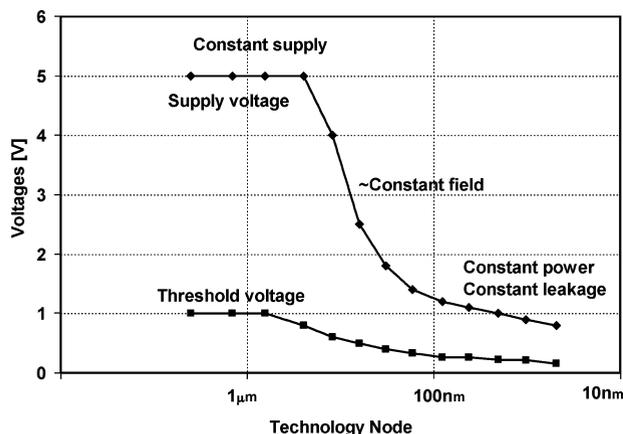


Fig. 1. Voltage scaling regimes.

technology node. While 1 V is high enough to enable analog designers enough margin to provide most analog functionality, scaling sharply below 1-V supplies (as is the preferred low-power digital scenario) will have a dramatic impact on what can be accomplished. Even at 1 V, high-resolution, high dynamic range analog is already becoming exceedingly hard. Hence, while supply voltage reduction still is very attractive in the digital space, it becomes less and less so for analog (as well as for memory) modules.

This divergence of the technology roadmaps is strongly reflected in the latest version of ITRS, which prescribes different scaling scenarios for virtually every application (Fig. 2). Even within the digital logic realm three different technology roadmaps are identified (high performance, low operating power, low standby power). In the mixed-signal arena, we can identify options ranging from high-bandwidth analog, high-resolution analog, RF, and MEMS.

Technology divergence has some important repercussions on the SoC paradigm. The widely diverging requirements of the different components of the SoC can be resolved by adding extra process steps—that is, by creating a more complex technology that supports the needs of the subcomponents. For instance, the availability of transistors with thicker gate oxides allows for higher operating voltages for some of the analog components (fortunately, digital I/O may need similar devices) These extra steps (the number of which can be quite substantial) adversely affects the cost of the SoC, and may undermine the cost advantage offered by system-level integration.

This may make other solutions such as system-in-package (SiP) more attractive. To understand the trade-offs, it is worthwhile examining the impact of deep-submicrometer scaling on digital, analog, and RF subsystems first. This can help us to identify what components can be integrated in the mainstream, “digital” technology at a reasonable cost, and what components would be better off being integrated at the package or board level.

2) *Digital and Memory Scaling*: Since digital logic represents the bulk of the integration complexity, it is worth briefly recapitulating the digital scaling trends first. Digital logic applications benefit greatly from scaling of the transistors, and will continue to do so for a considerable time. Most digital transistors use the minimum channel lengths available in a technology node, as this produces the smallest gates and the highest logic density. In addition, shorter transistors are faster and/or can save power by operating at a lower supply voltage.

As stated in the previous subsection, there exists an optimal supply and threshold voltage for each digital application, which depends upon the required frequency of operation and the switching activity ([8], [11]–[13]). If the active energy dominates, the total energy dissipation can be minimized, while maintaining the speed, by simultaneously lowering the supply voltage and thresholds. The optimum is achieved when the active and the leakage energy are balanced. For typical CMOS devices, leakage should contribute to approximately a third of overall energy dissipation [12]. High-activity, high-frequency blocks prefer to operate with low thresholds and low supplies, while low-activity or low-performance circuits, such as some mobile applications and memory, suppress the leakage by using high-threshold transistors. Architectural techniques such as the use of concurrency, and dynamic voltage and threshold scaling, have been developed to better utilize the dependence of the technology sweet-spot on application parameters.

Memory presents an entirely different set of challenges. Density and cost/bit have long been the dominant metric. Especially DRAM represents an extreme example of technology divergence. Typical standalone DRAM technology has branched off from the main technology stream quite a long time ago, so that it has diverged into a process that is fundamentally different than standard logic CMOS. A DRAM process uses multiple polysilicon layers or deep

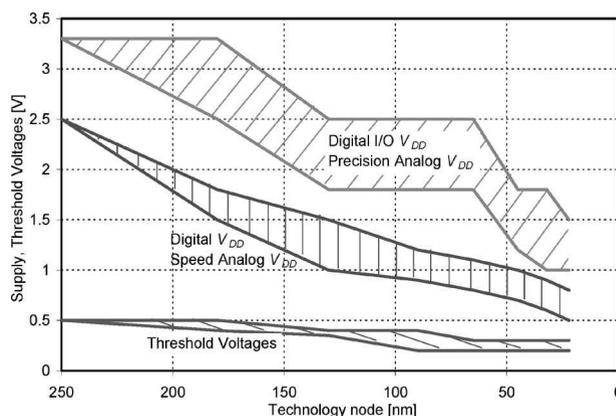


Fig. 2. Diverging supply voltage scaling trends (as projected in the ITRS [10]).

trenches to implement large storage capacitances. Interconnect is few and regular, so that a low count of metal wiring layers suffices. To embed DRAM into a logic process, some additional process steps can be included to add storage capacitance between the wiring layers (e.g., [14]). Yet, embedding DRAM into standard processes has rarely been considered cost effective, and SiP or board-level integration strategies have been the preferred option.

In contrast, SRAM arrays occupy a large fraction of the chip area in many of today's SoCs, and will continue to do so for a large fraction of future designs. In fact, memory integration is quite an essential component of the SoC concept, predominantly because of the higher bandwidth and lower latency of on-chip memory. It is hence of paramount importance that memory density track the scaling trends of logic. Already, high-density embedded SRAM frequently utilizes additional mask layers to increase the memory density [15]. Leakage and process variations present major challenges for the further scaling of the conventional six-transistor (6-T) SRAM cells. The increasing spread in transistor thresholds effects both the static power dissipation (which is the dominant source of power consumption in large memories), as well as the memory robustness. While a lower supply voltage can be used to address the power consumption, especially in standby mode [16], [17], reducing the supply voltage compromises the cell stability, measured as the static noise margin [18]. SRAM stability can be maintained through increasing the transistor widths, invoking the fundamental tradeoff between density and robustness. Another option is to keep the supply and threshold voltages steady, or to even increase them. Hence, as the technology scales, SRAM exhibits a reverse voltage-scaling trend [19]. This obviously challenges future integration of complex logic and large memories.

3) *Analog Scaling*: We have argued the existence of "optimal technologies" for digital building blocks. While considerably less documented, there exist optimal technologies for analog circuit blocks as well. Technology scaling increases the unity-gain frequency f_T of the transistors (Fig. 3), but adversely affects virtually everything else relevant to analog designs, as discussed below. High-speed analog circuits can exploit the higher f_T , if they can cope with the obstacles that come along. On the other hand, high-dynamic-range data converters (such as Nyquist converters) suffer greatly from the lowered supply voltages. As a result, they will most likely not be implemented in a low-voltage, minimum channel-length digital technology.

Reduced SNR: Reduction in supply voltages lowers the available voltage excursions in analog circuits, fundamentally affecting the signal-to-noise ratio. In addition, the transistor noise gradually increases with scaling. This is illustrated by (1), where the consecutive terms represent the thermal noise component of the drain current, and the

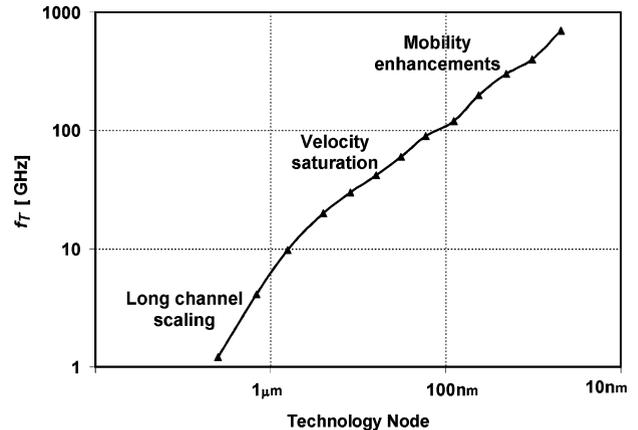


Fig. 3. Trends in f_T [10].

flicker ($1/f$) noise, respectively. Parameter γ , which equals to $2/3$ in long-channel devices, increases to modestly higher values in short-channel devices [20]. Even though early measurement results indicated dramatically higher values for γ , recent measurements and modeling shows a more gradual increase

$$\overline{i_d^2} = 4kT\gamma g_{ds0}\Delta f + K\frac{I_D^\alpha}{f}\Delta f. \quad (1)$$

Unfortunately, in scaled technologies, pMOS transistors are no longer buried channel devices, and their flicker noise is no further an order of magnitude lower than nMOS. Hence, flicker noise increases at the same pace both nMOS and pMOS devices.

In sampled-data circuits, a reduction of kT/C noise component requires an increase in the capacitor sizes, thus paying a penalty in power consumption. For example, an increase in 6 dB in SNR requires doubling the capacitor size and, consecutively, the power dissipation. For example, the results of [21]–[23] clearly show that technology scaling beyond an optimal point negatively affects the power of sampled-data pipeline A/D converters.

Lower intrinsic gain: The intrinsic voltage gain of the transistor is proportional to $g_m r_o$, which has been decreasing with scaling. The dominant reason is the lowered device output resistance resulting from drain-induced barrier lowering (DIBL) and hot-carrier impact ionization.¹ The gain limitation presents a serious challenge to precise charge manipulation as needed in sampled-voltage circuits, such as those needed for high-precision data converters.

Device leakage: One of the most attractive features of the MOS technology is the natural availability of a

¹Observe that the ITRS predicts $g_m r_o$ to stay constant over the next decade.

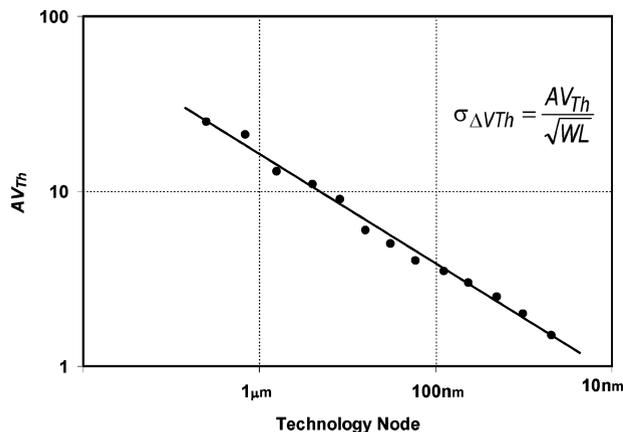


Fig. 4. Transistor matching properties as expressed in AV_{Th} (or indirectly standard deviation) as a function of the technology node.

high-quality switch, which lends itself to precise manipulation of charge [24], [25]. With continued technology scaling both the drain leakage and the gate leakage currents are projected to increase. Increased drain-to-source leakage of an off-switch can decrease the low-frequency gain of an amplifier that drives it by lowering its equivalent output resistance. The fidelity of charge stored on a node connected to leaky transistors becomes hard to maintain. As a particular challenge, elevated gate leakage violates the high-impedance “summing-node” assumption that underlines the operation of switched-capacitor circuits, especially at low speeds.

Reduced matching: Transistor matching properties are proportional to the device area, and improve with a reduction in oxide thickness [26], [27], which allows for size reduction of low-precision analog circuits (Fig. 4).

However, devices with small geometries also experience larger mismatch due to higher order terms with either short W or L [28]. When the oxide thickness reduces to a few atomic layers, quantum effects and random dopant fluctuations become significant and matching parameters degrade.

Passives: Scaled CMOS processes, mostly targeting digital applications, frequently lack high-quality capacitors. Sampled-data systems, for example, rely on linear high-density, high-Q capacitors. With process scaling, gate and diffusion capacitances can no longer be used in precision applications (because of linearity and parasitic resistance concerns). Double poly or MiM capacitors are frequently offered as an option in the process, albeit at an increase in cost. Fringe capacitors utilizing the dense array of metal layers in a modern process offer an alternative [29]. These capacitors benefit from the increased number of metal layers and decreased metal spacing offered by advanced processes.

Models: The immaturity of the device models in scaled processes offers a particular challenge to the mixed-signal

designer. Deep-submicrometer effects combined with evolving process parameters (even while the process is already in production) make it hard to create stable accurate models to the designer. The digital design methodology allows for the design of components, while the process node is being refined. This is not the case in mixed-signal where mature device models are needed. As a result, there has been a noticeable trend to skip a technology node, especially during the years of acceleration (for instance, move directly from 180 to 90 nm, skipping the 130-nm node).

4) *RF Scaling:* It is interesting to observe how technology scaling, while presenting more and more of a challenge to analog design, is the engine behind increasingly higher performance integrated CMOS RF. In a sense, the situation is somewhat similar to what happened to mixed-signal design in the late 1970s and 1980s, where newer technology nodes allowed MOS to take over functions that were traditionally implemented in bipolar. It is undeniable that devices have become faster pushing the f_{max} over the 100 GHz barrier with the introduction of 130-nm CMOS (see Fig. 3). Raw speed is a boon for RF integrated circuits, allowing sloppy analog style amplifier design rather than optimal microwave design approaches.

While a microwave designer examines important metrics such as G_{max} and stability, power gain, and noise circles in the complex impedance space in order to arrive at the optimal transistor bias, size, and impedance matching network, an analog “RF” designer simply maximizes voltage gain without matching considerations. Since the analog circuit is operating at a small fraction of f_T there is much margin for error enabling simpler design techniques. In addition to higher frequency of operation, scaling also allows RF circuits to operate at lower power and with lower noise at a given frequency. Measurements show that continued scaling will lead to improvements in RF device metrics such as NF_{min} and f_{max} .

Passive devices play a key role in RF and microwave integrated circuits. A schematic of a discrete solid-state radio is peppered with inductors, capacitors, transformers, and transmission lines. In these designs, discrete devices such as vacuum tubes, bipolar, or FET transistors, were relatively expensive and unreliable. Due to insufficient gain or excessive phase delay in RF transistors, feedback could not be used to correct for analog impairments. Hence, most signal processing is performed with passives rather than actives. A diametrically opposed philosophy evolved in the analog integrated circuit community, where transistors were effectively free and passive devices consumed inordinate amounts of die area. Any means to reduce the size of or eliminate capacitors and resistors were quickly adopted, such as Miller pole-splitting compensation. When resistors or capacitors were needed, good designs could tolerate large process variations in the absolute value of any component by relying on precision

matching of relative quantities, such as the ratio of the sizes of two components.

As Si technology progressively improved and the upper operating frequency approached the microwave “GaAs” (III-V) region, a battle of philosophy was fought between the Si analog IC design ($g_m R_o$ centric), and the GaAs microwave (s -parameter community). The Si community won this battle, because in the end, their products were cheaper and they were more adept at producing working Si for mass consumer markets. Hence, the GaAs community has retreated into niche markets such as power amplifiers and radars, although even this territory is now coming under dispute [30].

Today billions of RFICs are fabricated in CMOS and BiCMOS/SiGe BiCMOS technologies, and such chips are often dominated in area by inductors, capacitors, and transformers. Integrated inductors on Si were advocated first by Meyer and Nguyen [31], and have since played an integral role in RFICs. Inductively degenerated cascode LC tuned amplifiers are the workhorse low-noise amplifiers. LC tank VCOs are the building blocks for low phase-noise frequency synthesizers. Interstage matching networks are often composed of LC circuits in power amplifiers. LC filters are also used for wideband, concurrent, and multi-mode components [32], [33]. In general, due to the charge-voltage nature of MOS and bipolar devices, the intrinsic capacitance of such structures is tuned out by inductors. Inductors are also playing an increasingly important role in wireline and optical high-speed communication circuits, as key elements in clock and data recovery PLLs and shunt-peaked amplifiers.

As Si technology continues to scale, there is a real question is whether it is wise to devote large areas of the chip to inductors and capacitors. This is doubly true in a deeply scaled technology where the area of inductors can be replaced with hundreds of thousands of logic gates. A ring oscillator, for instance, can replace an LC tank oscillator at a small fraction of the area. If one can devise a way to overcome the intrinsically higher phase noise of the ring oscillator using digital logic circuitry, then there is a compelling economic reason to do so. It is therefore likely that digitally assisted RF circuits will emerge.

5) *Power*: Power conditioning is a major component of electronic systems. Power devices have to be able to sustain voltages that correspond to battery supplies (3.6–5 V), and are fundamentally incompatible with deeply scaled mainstream CMOS. In addition, large inductors or capacitors are often needed. All of these seems to indicate that an SiP approach is more attractive for the realization of advanced power distribution and conversion networks.

B. Addressing the Mixed-Signal SoC Challenges

Following our line of reasoning, it seems clear that the concept of integrated mixed-signal system-on-chip faces

some crucial challenges in the coming decade. One can envision three different strategies to address these concerns.

- For the lowest cost, it would be ideal to use only the *standard digital process*, which offers a single oxide thickness with two threshold voltages. By necessity, this approach requires the introduction of novel mixed-signal architectures, which exploit the abundance of digital gates to compensate for the lack of high-quality analog components.
- SoC in the most *advanced process with additional features*. While this increases the manufacturing cost, adopting a more complex process may help to accomplish better performance, power, or system size. The cost increase of additional process features depends upon the type of modification. Some changes can be made quite easily. For instance, implementing a device with a different threshold requires another masking step, to control the threshold-adjust implant. The design rules for this layer are typically of noncritical-dimension, increasing the mask cost by approximately 2% in a 90-nm technology. The addition of a second oxide thickness allows for higher voltage operation and larger signal swing. Fortunately, many 130- and 90-nm processes already include that option, and provide 2.5-V I/O transistors with 0.25- μm minimum channel lengths. An extra mask step is needed to add high-density linear metal-insulator-metal (MiM) capacitors, which are quite attractive in sampled-data and RF systems. This option trades off increased process cost for reduced die size. Other process options are a lot more expensive, however. Implementing embedded DRAM requires six extra mask steps, and most of them are at critical dimensions. This increases both the nonrecurring mask costs and fabrication costs by about 20%–25%. SRAM density is improved by adding a capacitor, requiring two extra masks. More aggressive process steps can be considered as well, such as the postprocess addition of high-Q MEMS resonators (for instance, using SiGe). This requires a substantial number of (typically noncritical) extra process steps.
- The SiP approach is attractive when the process needs become too divergent. This allows each system component to be optimized individually, leading to better and often lower cost solutions. This was not the case traditionally. SiP technologies tended to be expensive, and major hurdles prevented a widespread adoption (such as the availability of bare die, for instance). This changed substantially with advances in mobile telephone technology with its tight size and aspect ratio constraints. The huge volumes of mobile handsets helped to bring the cost down significantly, and SiP

has become the technology of choice for low-cost integration of analog functions and memory with the digital baseband.

In the RF space, the SiP approach allows for the use of higher quality passive devices external to the RF chip. Today this can be done, either by using a multichip module with a less expensive low-loss substrate housing passives, or by simply absorbing the passives into the package substrate. Such an approach is actively pursued by the industry to lower the cost and to add flexibility to the design of RF ICs. High-Q package-based inductors, for instance, can be realized in relatively inexpensive LTCC packages [34], or housed in high-resistivity Si, quartz, or SOS [35]. The SiP approach is also attractive for the integration of high-Q RF-MEMS components such as the FBAR [36], which require nonsilicon materials such as AlNi and need encapsulation for reliable operation [37].

One limitation, though, are the parasitics of the package and the undesirable coupling that occurs in the package. The parasitic inductance of bond wires is prohibitive and requires very careful EM simulation and modeling. For instance, if we wish to integrate an inductor with value $L = 5$ nH, then placing this inductor in the package is possible if the parasitics (which vary) are a small fraction, say, below 1 nH. This is viable with bondwire technology and easily done with solder bumps where the parasitic inductance is of the order of hundreds of pH. Furthermore, switched capacitor tuning techniques can be used to overcome process variations. LC-based oscillators using CMOS switches for coarse tuning are now common [38]. Placing an inductor smaller than 1 nH off chip is much more difficult, though. In addition, as the frequency of operation increases from low GHz to tens of GHz, the SiP solution must be examined very carefully. The moment a signal is sent “off-chip,” there are EM coupling consequences that complicate the design. For instance, if the off-chip component is the tank for a VCO, unwanted signals couple into the off-chip VCO and generate spurs in the frequency spectrum of the frequency synthesizer. Again, careful modeling and simulation is necessary to minimize such effects. In contrast, integrated inductors exhibit substrate coupling, which is typically an order of magnitude smaller. Efficient simulation of substrate-induced noise, especially supply and ground noise, remains a challenging task, although research efforts have made important headway toward solving this problem, e.g., [39]. Another concern with respect to the SiP approach is that today and in the near future the number of dies that can be integrated in a package is limited by the bonding technology and the relative sizes of the dies.

While both advanced processes and SiP approaches are attractive, the standard digital process promises the lowest cost and the highest integration, yet faces the largest design challenges. The rest of our discussion hence focuses on the latter.

III. MIXED-SIGNAL SoC DESIGN IN A DIGITAL WORLD

The growing interest in implementing analog functions using transistors with minimum channel length and minimum oxide thickness, as available in a standard digital technology, has led to the conception of many innovative circuit and building blocks, some of which are presented in this section. Even more importantly, this has led to a new analog design philosophy, often called the “digitally assisted analog” design approach, which advocates a departure from the traditional precision analog design [23]. The concept is to avoid precision analog to start with and use the abundance of digital gates to correct and compensate for the inaccuracies or distortion. This approach has become popular in a number of modules, ranging from A/D converters (e.g., [45]) and PLLs to RF transceivers (e.g., [42], [64]). Some examples illustrating this novel design philosophy are also included in this section.

Switches: Charge-based analog circuits rely on the availability of precise sample-and-hold (S/H) circuits. The preferred implementation strategy of such an S/H circuit in CMOS is the switched capacitor. The fidelity of sampled charge on a capacitor relies on the performance of an MOS switch, which is measured by its on-state small-signal bandwidth, and the ratio of the gate capacitance of the switch to that of the sampling capacitor. Technology scaling reduces the associated capacitance while keeping the on-resistance of the switch nearly constant [28]. This improves the tracking bandwidth of the S/H circuit, and simultaneously alleviates the charge injection problem during the turn-off of the sampling switch. Hence, the increase of f_T through technology scaling improves the linearity of the sampling switch. A rule of thumb for

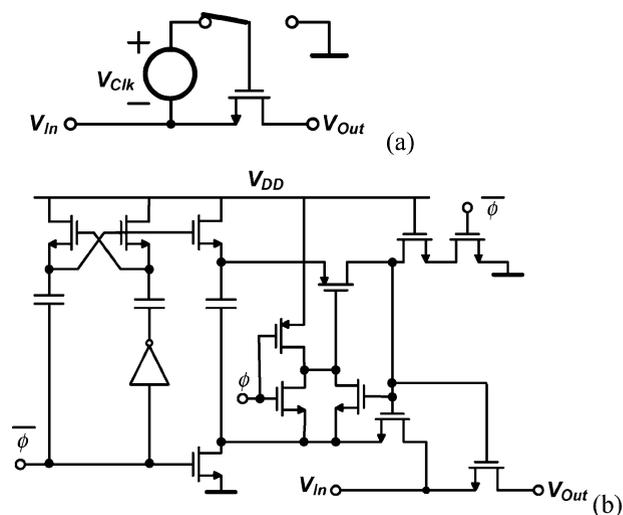


Fig. 5. Clock bootstrapping. (a) Principle. (b) Circuit implementation.

analog designers is to use minimum channel length for switches when matching and leakage requirements are not critical.

A major drawback of the simple S/H circuit is the variation in the on-resistance of the input switch, which introduces distortion. In recent technology nodes, threshold voltage scaling lags that of the supply (see Fig. 2), resulting in a larger on-resistance variation in a switch. This causes the bandwidth of the switch to become increasingly input-signal dependent, causing signal-dependent distortion. One approach to combat this problem is to adopt clock bootstrapping [40], [41]. This technique keeps the gate-source voltage of the switch constant and equal to the supply voltage, as shown in Fig. 5(a). One challenge in this approach is ensuring that the reliability of the circuit is not compromised. An example circuit implementation that limits the gate-to-source and drain-to-source voltage excursions to V_{DD} is shown in Fig. 5(b).

Bandgap references: Designing precision voltage references with sub-1.2-V supplies presents a significant challenge. Many designs therefore resort to using precise off-chip references or avoid using them altogether [42].

Preamps and comparators: The offset in preamps and comparators constitutes the major source of error in parallel ADCs. The simple differential structure with thin oxide devices continues to be the predominant preamplifier architecture in newer technologies [Fig. 6(a)]. Its input transistors are made large to minimize the offset. The comparator, on the other hand, is sized small to maximize the bandwidth, which is limited by the preamplifier-comparator interface. Small comparators, such as the one shown in Fig. 6(b) can be used in low-voltage pipeline ADCs, since those have a built-in correction to offset errors.

Amplifiers: Precision op-amps are one of the workhorses of analog design. For instance, in multistage ADCs they are almost invariably employed to relay the input signal (or the residue signal) to the following conversion circuits. The accuracy and speed of the analog subsystems are often dictated by the performance of these amplifiers. Operating on the edge of the performance envelope, these op-amps are the subject of an intense tradeoff between dynamic range, linearity, settling speed, stability, and power consumption.

The most prominent challenge in amplifier design is the reduced supply voltage, which limits the number of transistors that can be placed in a cascode stack. The available signal swings are further reduced by nonscaling of the transistor thresholds, and transistor overdrive voltages. For example, a differential cascode amplifier with tail current source (Fig. 7) has five transistors in the stack, and with 150-mV overdrives can achieve a larger than 2 V_{p-p} output voltage (for a 2.5-V supply). When the same design is scaled to a supply of 1 V with the same overdrive, the voltage swing is reduced to 0.5 V_{p-p}. To maintain the same SNR, such a design has to employ

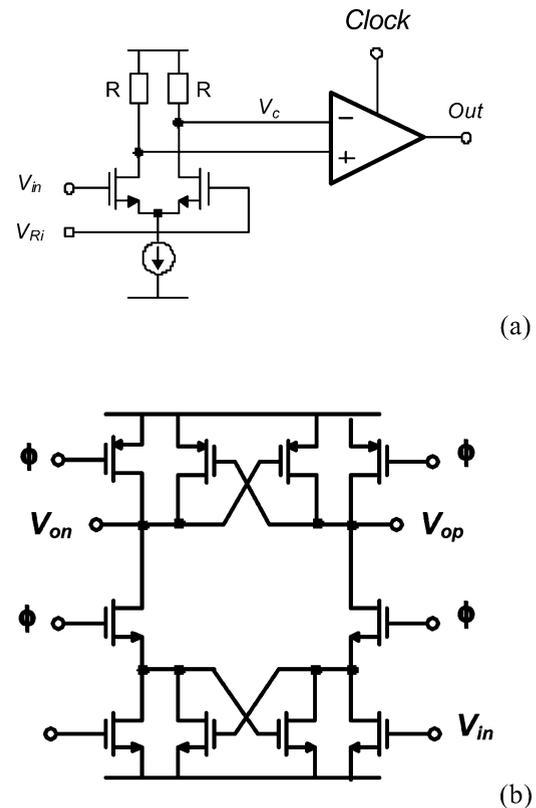


Fig. 6. (a) Differential resistively loaded preamplifier. (b) dynamic comparator.

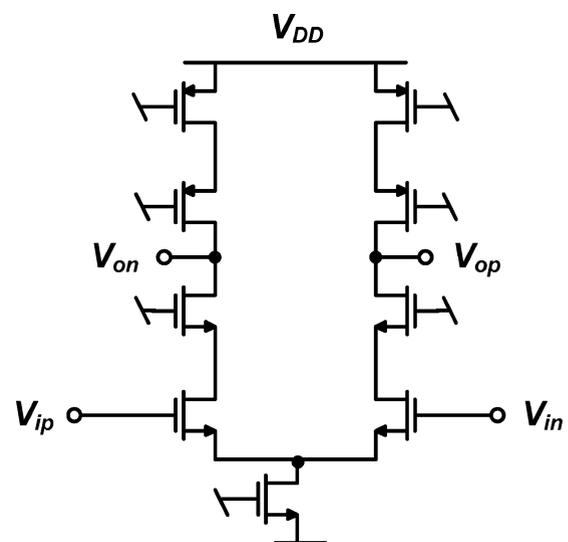


Fig. 7. Differential cascode amplifier with tail current source.

capacitors that are 16 times larger, which severely affects the power consumption and the bandwidth of the system.

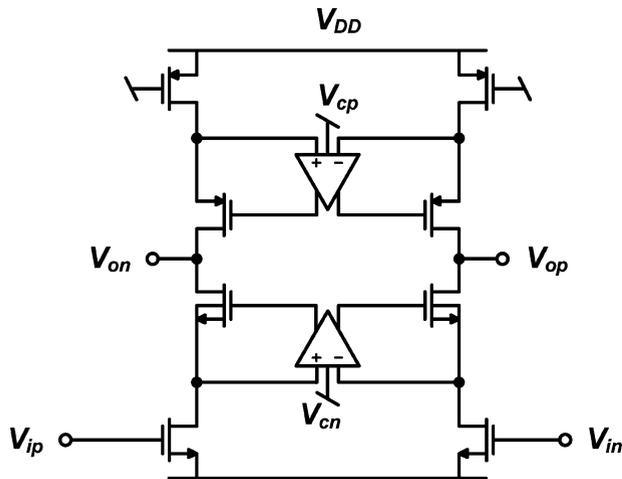


Fig. 8. Gain-booster op-amp (pseudodifferential).

In order to maintain the dynamic range and improve the noise immunity, all deep-submicrometer designs currently exploit differential topologies. As the supply voltages are approaching 1 V, folded cascodes and pseudodifferential structures are gaining popularity for their potential to improve the voltage headroom. Pseudodifferential amplifiers are particularly attractive. However, with this architecture, one must be careful in regulating the common-mode biasing.

The switched op-amp presents another approach to achieve high-output swing and to save power [65]. While the slow turn-on following a complete op-amp turn-off led to low operating speeds in early implementations, recently an 8-bit, 200-MS/s pipeline ADC with partially switched op-amps has been demonstrated [43].

As discussed earlier, the aggressive scaling of channel length leads to enhanced short-channel effects and high device output conductance, causing the intrinsic device gain g_{m,r_o} to drop to lower and lower values. To maintain the necessary open-loop gain, designers have resorted to alternative or novel configurations, such as gain-booster cascode devices and multigain stage amplifiers. The gain-booster amplifier (Fig. 8) uses an active cascode configuration to boost gain without incurring excessive voltage drop [44]. This approach exploits the fact that, when the load is purely capacitive, the dc gain of an op-amp can be increased by using negative feedback to enhance the cascoding effect. Today we even see recursive gain boosting, where the amplifier used for gain boosting is itself gain boosted [45]. The ultimate limit of this approach would be the direct current path to ground at the drains of the cascode devices due to hot-carrier-induced substrate currents.

It is interesting to note that using three or more stages of amplification is another simple scheme to realize higher gain blocks. However, compensation of feedback ampli-

fiers with more than two stages is tedious, requiring careful nested Miller compensation (rendering the approach somewhat less power-efficient). This approach may become the only viable op-amp architecture at supply voltages of less than 1 V.

A/D converters: The impact of technology scaling on the performance of various ADC architectures depends on their resolution. Despite the improvement in f_T , the reduction of supply voltages and the increased channel noise are shrinking the dynamic range achievable. For high-resolution converters, this inevitably leads to an increase of power consumption to maintain SNR (Table 1). On the other hand, the accuracy of lower resolution ADCs is limited by component mismatch. For fixed conversion speeds and nonscaled voltage references, this leads to a power and area scaling trend similar to that of digital circuits, as is illustrated in Table 1 [46]. In 180-90-nm technologies, the breakpoint between matching and SNR limited is around the 8–10-bit resolution. Table 1 also indicates the scaling trend of a commonly used metric for ADCs, called the *figure of merit* (FOM), which is defined as

$$FOM = \frac{P}{2^{ENOB} \cdot f}, \quad (2)$$

where P is the power, ENOB the effective number of bits, and f the sample rate (for a Nyquist ADC) or twice the effective-resolution bandwidth (for an oversampled ADC).

As scaling continues, there is a noticeable migration trend of the boundaries between the various A/D architectures (Fig. 9). Oversampled converters are migrating into the zone that used to be dominated by pipeline ADCs [47]. Although the front-end of an oversampled converter is negatively affected by a lowered supply, the decimation filter greatly benefits from scaling. Hence, oversampled converters become increasingly more power efficient than pipelined converters at high resolutions and low conversion speeds. The latter are now reporting resolutions as low as 5–8 bits, which were considered the preferred space for flash-type architectures [48].

Table 1 Scaling Model of AD Converters (and Mixed-Signal Circuits in General)

Scaling Parameter	Digital Circuits	Matching-Limited	SNR-Limited
Dynamic	Wordlength	$\propto \sqrt{WL}$	$\propto \sqrt{C/kT}$
V_{DD}	$1/S$	$1/S$	$1/S$
Speed	S	1	1
LSB	n/a	$1/S$	1
Area	$1/S^2$	1	$1/S^2$
Power	$1/S^2$	1	$1/S^2$
FOM	n/a	1	$1/S^2$

Calibration of an ADC is one of the best opportunities for the aforementioned digitally assisted analog design philosophy. Consider, for instance, the pipeline ADC converter. In a classical analog approach, where only a gain of around two is required in each amplifier stage, large open-loop gain is used with feedback to overcome the distortion of the operational amplifier stages. Thus, to simultaneously meet the settling time requirements and accuracy, a complex analog design with high power consumption is required. In the digitally assisted approach [23], though, simple low-precision open-loop amplifiers are used instead, and the error due to the nonlinearity of the amplifiers is corrected in the digital domain. While the idea of self-calibration of an ADC is well established [49], [50], the underlying philosophy differs in that earlier approaches begin with a precision ADC converter and calibrate the ADC to improve the last few bits in converter accuracy. The key feature of this new philosophy is strong reliance on adaptive signal processing in the digital domain to correct for large errors in the analog domain. In other words, very little attempt is made to produce a high-quality ADC up front.

Power amplifiers: The linearization of power transmitters has been an active area of research for many years. Early work started with the linearization of traveling wave amplifiers [51]. Techniques such as predistortion, feed-forward [52], envelope elimination and restoration (EER), and Cartesian feedback emerged as viable solutions. Most of these techniques, though, were applied to immobile base stations rather than mobile transmitters due to the high cost and complexity of the digital signal processing. This trend is now being reversed, as mobile transmitters are capable of orders of magnitude more computation than before. In contrast, little work has been done on the linearization of receivers such as low-noise amplifiers and mixers due to the asymmetry of the problem. In a transmitter, one can examine the output of the PA to determine

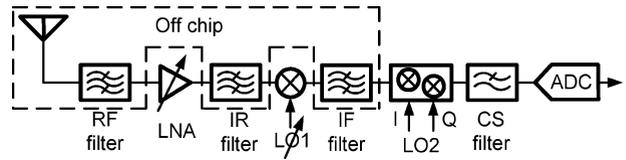


Fig. 10. Super-heterodyne receiver architecture.

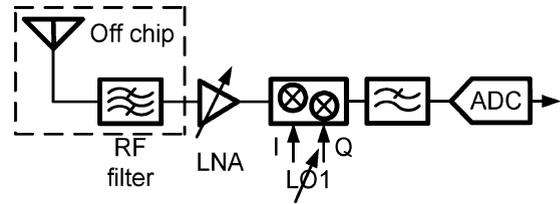


Fig. 11. Direct conversion/low-IF receiver.

the error, since the transmitted signal is generated by the same device. In a receiver, though, the received bits are unknown, with the exception perhaps for the header of a packet. We expect that advanced “blind” digital processing techniques will be applied to every block in a typical mixed-signal IC, including the RF front end and the analog baseband.

RF transceivers: We have already indicated that the need for a larger scale of integration has instilled notable changes in the architecture of mixed-signal systems. Radio transceiver architectures are a great example of this trend. Most of the receivers in the past adopted the conventional super-heterodyne approach. The super-heterodyne transceiver decouples the requirements for selectivity and sensitivity, inherent to heterodyne receivers, by performing dual-IF mixing. When migrated to integrated circuit technologies, several filtering operation have to be performed off chip, as is shown in Fig. 10. By using high-quality off-chip passive components, the super-heterodyne architecture addresses issues such as out-of-band/channel energy and image rejection, and lowers the requirements for the on-chip channel-selection filters. In addition, the architecture often requires off-chip components for tunable RF synthesizers.

Several architectures have been explored to reduce the need for off-chip components in integrated transceivers. Just to mention the predominant ones: direct conversion [66], low-IF [67], wideband-IF [68], and Weaver [69]. The direct-conversion architecture (Fig. 11) (also known as homodyne or zero-IF) eliminates most off-chip components in the receive signal path by translating the desired signal directly to baseband. It naturally avoids the problem of the image frequency, since its IF is at zero. However, direct-conversion receivers have several major drawbacks,

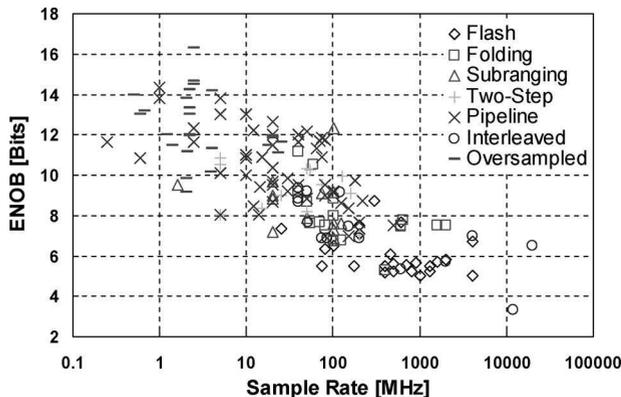


Fig. 9. Effective number of BITS (ENOB) versus sample rate of ADC architectures, published between 1987 and 2005.

including large sensitivity to dc offsets (induced through the self-mixing process), and sensitivity to $1/f$ noise. Most of these concerns are effectively addressed by adopting an active offset cancellation approach, as well other circuit and system design techniques.

The low-IF integrated receiver architecture (Fig. 11) resolves many of the dc offset problems of the direct conversion receiver. A single mixing stage translates all the desired channels to a low IF, with a bandwidth totaling several channels. A wider bandwidth, high dynamic range ADC (or a bandpass sigma-delta modulator), is used to extract the desired channel. The key feature of the low-IF architecture is that the desired channel is moved away from dc, hence alleviating the dc offset issue. Since the IF frequency is not at zero, some image rejection is still needed, typically through an image-rejection mixer.

It is also in the RF transceiver domain that we find one of the more extreme examples of the digitally assisted analog design methodology. Proliferation of wireless communication is creating a demand for multistandard, multimode operation, presenting a great opportunity for high levels of integration. In [42], [64], a highly integrated direct-conversion architecture is presented for both the Bluetooth and GSM standards. It exploits the inherent mixing performed by the sampling function, and follows it by discrete-time, passive IIR and FIR anti-aliasing and decimating filters. To address the many challenges offered by the scaled digital CMOS process, discrete time signal processing is leveraged throughout the transceiver chain.

IV. MIXED-SIGNAL DESIGN FLOWS WITH PLATFORMS

We argued in the introduction that there is a widening chasm between the design methodologies and tools used in the digital domain and the ones used for analog and RF

design. One can easily enumerate a number of reasons why this is the case. Foremost, the rich parameter space of analog circuit design makes simple abstractions hard to come by. In fact, whereas Boolean logic and RTL have served as fruitful abstractions for digital circuits, the abstraction for analog circuit design, albeit system-level models have been proposed in the literature, is typically a SPICE BSIM model, containing about 10 000 lines of device model code. Hence, hierarchical design is non-trivial, as too many parameters/constraints have to be carried between abstraction layers. In addition, building a complex system through simple composition of library models has failed in general in the mixed signal design, due to cross-coupling effects between modules—this means that the “separation of concerns” concept, again very successful in the digital realm, is hard to accomplish.

Much has been done in the area of analog design methodologies and tools in the past 20 years [70]. Yet, the impact of these developments on the design community has been small. We review here the most relevant contributions, and argue as to why these approaches had difficulties to be adopted by circuit designers.

1) *Review of Relevant Analog Design Tools and Methodologies:* Fig. 12 shows the timeline of the best known university-developed tools. To put these tools in a proper perspective, consider the simplified analog design flow, as shown in the diagram of Fig. 13. The first tools devoted specifically for analog circuits were for circuit-sizing optimization. AOP [71] and, later, APLSTAP [72] and Delight.Spice [73] developed in the late 1970s and early 1980s were in fact addressing this problem. These methods did not scale to large circuits, because the complexity of the algorithms was in disproportion to the computational power available at that time. Hence, the achieved “optimal circuits” were subject to numerical noise and/or local

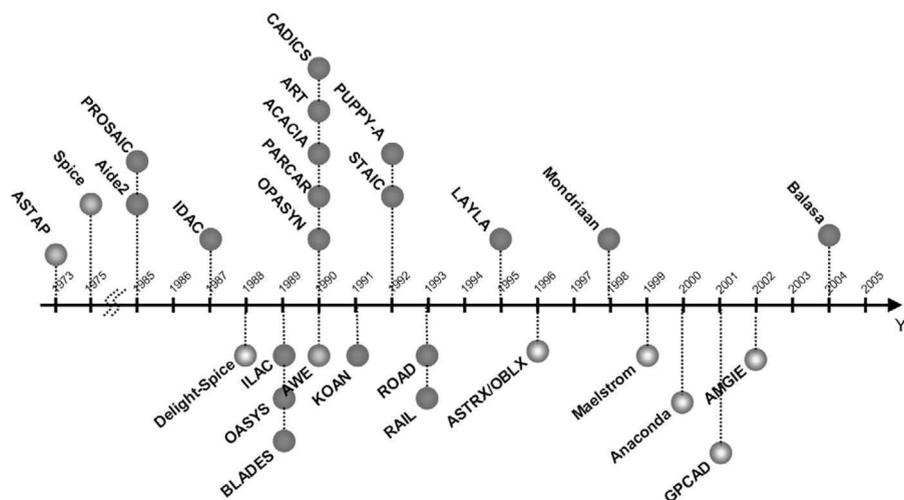


Fig. 12. Historical development of CAD tools for analog circuits.

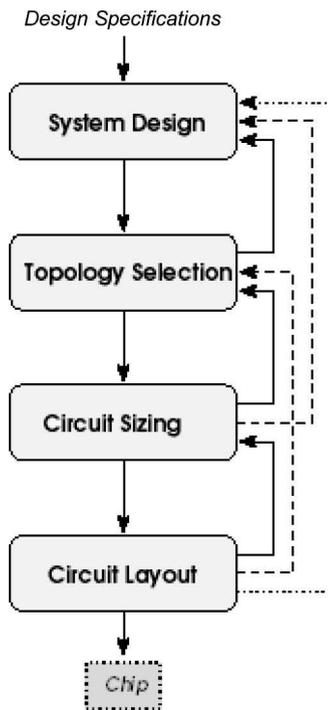


Fig. 13. Standard analog design flow.

minima. Furthermore, the resulting circuits were not robust with respect to parasitics and process variations.

Initial attempts at providing automated “synthesis” solutions to analog designs dealt with *ad hoc* tools to synthesize specific components with hard-coded topologies and layout styles. (see AIDE2 [74], PROSAIC [75], and BLADES [76]). The first examples of complete solutions to circuit sizing and layout generation can be traced back to IDAC/ILAC [77], OPASYN [78], and OASYS/ANAGRAM [79]. IDAL/ILAC and OASYS/ANAGRAM were generic synthesis tools for analog integrated circuits, while OPASYN was a module generator capable of producing complete OpAmp layouts through circuit selection, parametric optimization and layout generation phases. However, it was limited to analytical models and digital-like routing mechanisms. As often the case in CAD tools, all the previous tools relied on hard-coded libraries. However, in contrast to digital circuits, the rigidity of hard-wired libraries tends to lead to poor performance in analog circuits.

Astrx/Oblx [80] was a synthesis environment using simulation to evaluate the performance of the proposed circuits. Astrx/Oblx was used to synthesize several circuits with complexities equivalent to industrial applications. The achieved results were usually comparable if not better than manual designs. However, the class of reported circuit was mostly confined to OpAmps with linear specifications. Moreover, the original tool was not capable of handling process variability and mismatch efficiently,

although some extensions for robustness were introduced in [86]. This is a typical problem of simulation-based synthesis approaches. Later simulation-based approaches focused on improving the simulation accuracy and the optimization techniques (see, for instance, [81]). Eventually, these concepts made it from academia (Carnegie-Mellon University, to be more precise) to the commercial domain with NeoCad, which was eventually acquired by Cadence.

In the past decade, a specific class of nonlinear optimization, known as convex optimization [82], has received its share of attention. Due to large improvements in computational efficiency achieved in [83], large-scale problems (thousands of variable and tens of thousands constraints) can be solved in minutes on a standard workstation providing a global optimum solution. In the 1980s, researchers from Bell Labs developed accurate models for circuits using ratios of polynomials (posynomials). Convex optimization can indeed be applied to posynomials. This was the basis for a set of tools developed in the late 1990s [87], which formed the core offering of the Barcelona Design EDA company. While more general than linear problems, convex problems do, however, limit designers in stating their problems. Expressing a design problem in posynomial form may be a very complex task, if not unfeasible. While some specifications are naturally expressed in posynomial form, others cannot be immediately captured. This (amongst other reasons) ultimately led to the unfortunate demise of Barcelona Design.

A. The Platform Paradigm

In this section, we are exploring the potential of a mixed-signal platform paradigm, which presents a radical departure from the way mixed-signal design is typically done.

1) *The Precursor to the Analog Platform-Based Design Concept*: The origins of what is now called analog platform-based design (A-PBD) can be traced to the top-down, constraint-driven design methodology for analog integrated circuits described in [84]. In this methodology, optimization is performed on system-level behavioral models that are architecture independent. To partition system-level requirements into circuit specifications, an approach based on the so-called *flexibility functions* was adopted. The goal was to maximize the likelihood that a feasible design be achieved at the end of the design process. A flexibility function captured, heuristically, information of the “difficulty” of achieving a set of performances. Solving an appropriate optimization problem, system specifications are decomposed into a set of lower level requirements. Flexibility functions allow raising the level of abstraction of analog design to capture analog systems rather than coping with individual circuits or layouts. In this sense, the methodology provides the first rigorous approach to system-level analog design with

possible extensions to mixed signal designs. However, the quality of the result was critically dependent on the flexibility functions. Failing to recognize the difficulty of a behavioral model parameter (e.g., overestimating the flexibility function for some parameters or ignoring correlations between other parameters) may lead to unnecessary problems in circuit design or even failures and costly iterations. Conversely, overconstraining a parameter with an excessively strict flexibility function may lead to designs far from what is attainable. Therefore, the aid of an expert analog designer is essential in the development of the flexibility functions. In addition, it is intrinsically difficult to model interrelationships among flexibility functions for different parameters as required to perform topology selection at the system level. Yet several circuits of moderate complexity were developed with this methodology.

2) *Platform-Based Design for Analog Mixed-Signal Circuits:* The lesson learned while developing the top-down, constraint driven methodology was that defining the proper level of abstraction, finding accurate and efficient behavioral models, capturing the design space and providing means for efficient exploration were all essential to solve the design methodology issue for mixed-signal designs. The classic dichotomy top-down design/bottom-up verification followed in [68], fell short in these systems, where the top-down phase is severely limited by the difficulty of introducing analog architectural constraints at the system level, and the bottom-up verification has problems at bridging the abstraction gaps left behind during the top-down phase. The problem cannot be decomposed cleanly, since the top-down and the bottom-up phases are so tightly coupled that it is difficult even to talk about the one without referring to the other. A paradigm shift was then required to cope with mixed signal designs.

Platform-based design (PBD) as presented in [53] and [54] has emerged as a novel paradigm in the digital domain to allow designing at higher level of abstraction while considering lower level physical properties. The PBD paradigm is a *meet-in-the-middle* approach consisting of a bottom-up characterization phase and a top-down mapping phase. A platform is a *library* of components and interconnects along with composition rules, determining legal compositions of components (platform instances). The bottom-up characterization phase abstracts architectures as library components providing a set of models for the services that can be implemented on it and their cost and performance. The top-down phase consists of selecting the optimal platform instance (according to some cost function) that can support the requested functionality while satisfying all system and architecture constraints.

B. Analog Platforms

An analog platform (AP) is a set of components each decorated with a set of behavioral models (\mathbf{F}), configura-

tion and performance models (\mathbf{C} , \mathbf{P}), and validity laws (\mathbf{L}) [55]. This rich set of models helps to address the concerns raised earlier, leveraging behavioral models to perform design optimization while considering architectural constraints and costs. Essential also is that the resulting abstractions ensure the following properties.

- *Flexibility*—any analog component can be encapsulated as an AP: newly specified circuits, analog IPs (possibly third party's), module generators, circuit synthesizers and optimizers.
- *Accuracy*—the AP abstraction requires a set of models that introduce architectural effects at the system level while guaranteeing composability.
- *Hierarchy*—AP components allow building high-level hierarchical models while preserving information on the actual architecture space. Therefore, correct abstraction levels can be selected for MS designs and enable efficient design space explorations.
- *Implementability*—a notion of feasible performance is propagated bottom-up into the design hierarchy, thus restricting (and characterizing) the actual design space.

The *behavioral model* \mathbf{F} allows for the abstract computation of the system response without being directly constrained to a specific architecture. Very general techniques can be adopted to implement behavioral models, ranging from hand written block models (requiring deep insight on the designer/developer) to *model order reduction* approaches, which are based on a mathematical formulations and can be fully automated [56].

Even more essential to the AP abstraction is the introduction of *configuration* (\mathbf{C}) and performance (\mathbf{P}) models. For a given module or component, the configuration model outlines the space of the feasible realizations (in terms of design parameters such as transistor sizes, bias currents, supply voltages, etc.). The corresponding performance model maps these configuration constraints into a set of feasible performance vectors. For example, for OTAs, the performance model is specified in terms of the {gain, noise, bandwidth, power} n -tuples, which accurately identify the feasible performance range of a given OTA architectures in a given technology. Having quantified bounds available is more attractive and reliable than the recursive estimation and optimization based approaches (such as advocated in [56]), especially in light of the many secondary parameters emerging in today's deep submicrometer processes. In hierarchical designs, where several components are connected together (defining a platform instance), the performance model of level l is the configuration space of level $l + 1$, hence a direct relation exists between the performance model \mathbf{P}^l and the configuration model \mathbf{C}^{l+1} .

The efficient and accurate mapping of configuration models into performance models is one of the main challenges of the A-PBD approach. Traditional performance

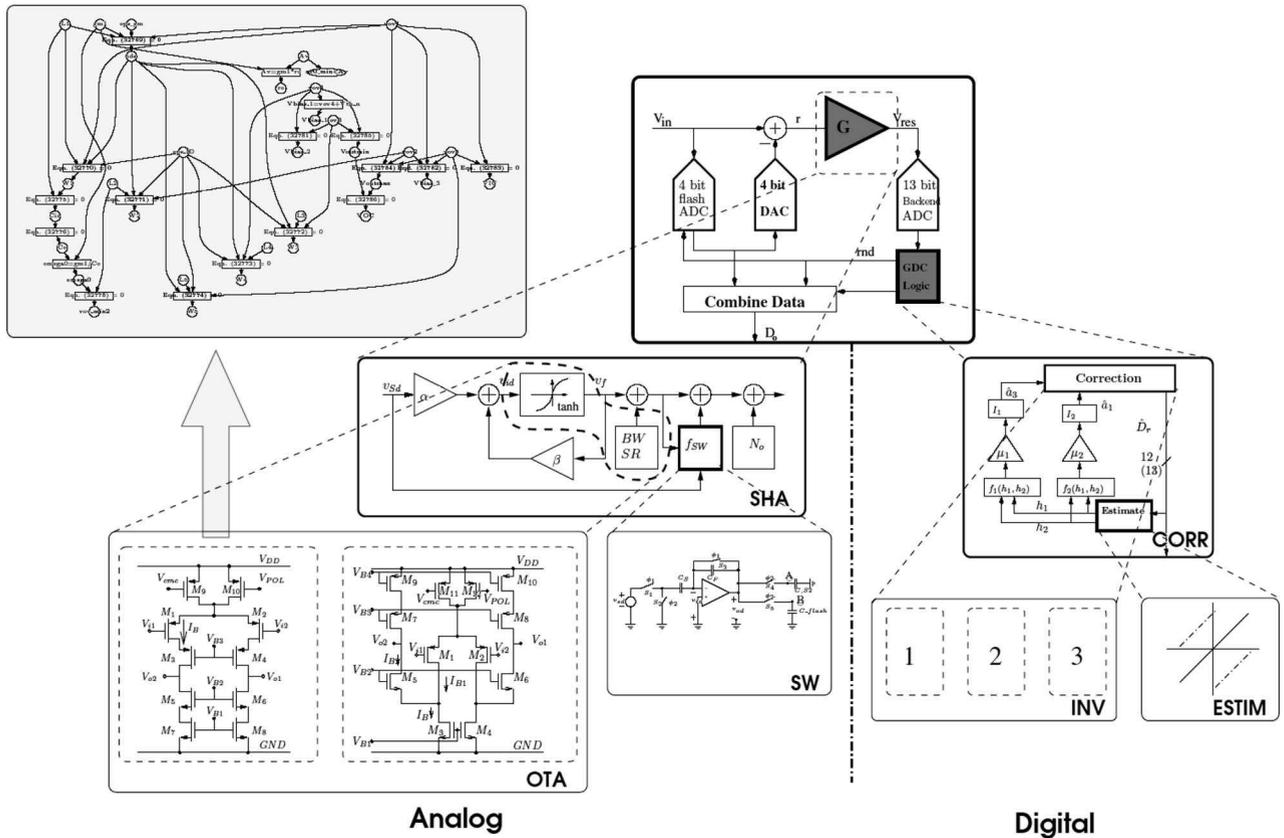


Fig. 14. Platform stack for the optimization of the first stage residue amplifier. Starting from the bottom left corner, an analog platform stack is built from circuit level components generating instances and new components at higher levels of abstraction. The top left graph shows the analog constraint graph (ACG) used to sample performance of the telescopic operational transconductance amplifier (OTA). The digital part of the mixed signal platform is generate in a similar way as shown on the right.

models are based on regression schemes, for which a rich literature exists (ranging from simple quadratic models for optimization [57] to advanced data mining techniques [58] and template independent schemes [59]). A more effective strategy is based on *classifiers*, which have the distinct advantage of making it possible to encapsulate architectural alternatives for the same functionality (that obviously share the same performance space) in a very straightforward manner. Only a negligible setup time is required to define a performance model for arbitrary performance figures and circuit topologies can be used. On the other hand, since the approach is based on a sampling scheme requiring accurate simulation of performance, the characterization itself may be expensive. In [60], an approach aimed at pruning the number of samples (simulations) required to characterize a circuit is presented exploiting structural and functional properties of the configuration space. Even if the exponential nature of the problem is not affected, the approach has shown to be practical for real case studies. In addition, the process is easily parallelizable, so the entire characterization time can be reduced to a few machine hours. A classification approach for analog

performance based on support vector machines [62] is presented in [61].

Validity laws **L** form the final element of the AP abstraction. When assembling a platform instance (composing platform library elements), the accuracy of the instance model has to be guaranteed. In fact, the plain composition of behavioral models may not correspond to the behavioral model of the composition. Validity laws limit the scope of behavioral models to enforce correct compositions and accurate modeling of interface effects (e.g., circuit loading due to other circuits). An example of validity laws and interface modeling in the AP context can be found in [63], where an RF receiver platform is built and used to optimize a UMTS system.

C. Mixed-Signal Design Flow With Platforms

The essence of platform-based design is building a set of abstractions that facilitate the design of complex systems by a successive refinement/abstraction process (Fig. 14).

Bottom-up Phase—Without loss of generality, we can assume that the library characterization starts from circuit

level components. In this case, a set of candidate topologies for the required functionality is selected, and the required models (behavioral, configuration, and performance models and validity laws) are generated. Notice that several topologies can be merged into a single higher level component, whose configuration space is the union of the individual topology performance spaces. The process can be iterated at higher levels of abstraction, using configuration models to limit the characterization space and to generate new performance and behavioral models. In this way, several layers of platform libraries can be obtained until the system can be easily mapped on the top-level platform. The “distance” between adjacent abstractions is a critical tradeoff points between the single component and the instance characterization efforts. Important is also that successive abstractions should preserve the conservativeness of the model. A model is conservative if its predicted performance is guaranteed to be achievable by some (lower level) configurations. Conservativeness may be broken because of: 1) inaccurate instance models and approximate validity laws and 2) because of optimistic performance/configuration models. The first issue has to be carefully addressed when abstracting the platform instance behavior and usually requires careful validation (see [63]). The second issue is addressed at the classifier generation level. In fact, the SVM scheme adopted in [61] provides smooth, continuous extrapolation among observed performances that, even mathematically speaking cannot guarantee conservativeness, in practice turn out to be very close to conservative. Moreover, conservativeness of a SVM model can be easily controlled with the heuristics adopted in the same paper.

Top-Down Phase—The top-down phase progresses through successive optimization and refinement. Design goals are captured as constraints and cost functions. At the highest level of abstraction, the constraints are intersected with the feasible performance set to identify the set of achievable performance that satisfy design constraints. This constrained optimization problem yields a point in the feasible configuration space for the platform instances at the highest level of abstraction. These points can be mapped back at a lower level of abstraction where the process is repeated to yield a new point in the achievable configuration set until we reach a level where the circuit diagrams and even a layout is available. If configuration models are conservative, it holds that for each design refinement a consistent solution can always be found. Hence, the design process can be shortened considerably. The question is obviously how many feasible points may not have been not considered because of the conservative approximation. Thus, the usual design speed versus design quality tradeoff has to be explored. Otherwise, at each level of the hierarchy, we have to verify using the performance models, the behavioral models and the validity laws.

A-PBD greatly simplifies the final verification step, since, in the end, models and performances used in the top-down phase were obtained with a bottom-up scheme. Therefore, a consistency check of models, performances, and composition effects is all that is required at a hierarchical level, followed by some more costly (but localized and isolated) low-level simulations that check for possible important effects that were neglected when characterizing the platform.

D. Example

To illustrate the A-PBD approach, we optimize the first stage of a pipeline analog-to-digital converter (ADC) across the analog/digital interface. The case study targets an 80-MS/s, 14-bit ADC designed in 130 nm CMOS that exploits a digital calibration technique scheme similar to the ones mentioned in the previous paragraphs. The first stage residue amplifier in the pipeline converter poses the most challenging constraints in terms of linearity, noise, and performance. Purely analog solutions typically come at a large expense in power consumption. Murman *et al.*, [23] proposed to use low-power open-loop amplifiers instead, and to identify and invert the nonlinearities arising from the lack of negative feedback in the digital domain. Exploiting A-PBD, we rely on a similar approach to optimize a *closed-loop* solution, so that the intrinsic advantages of feedback can be retained but linearity and gain requirements are relaxed to enable a reduction in power consumption [55].

The first step in the design process consists of building a set of platform libraries. Fig. 14 shows the platform stack developed for this system. Both analog and digital platform components are characterized. Focusing on the

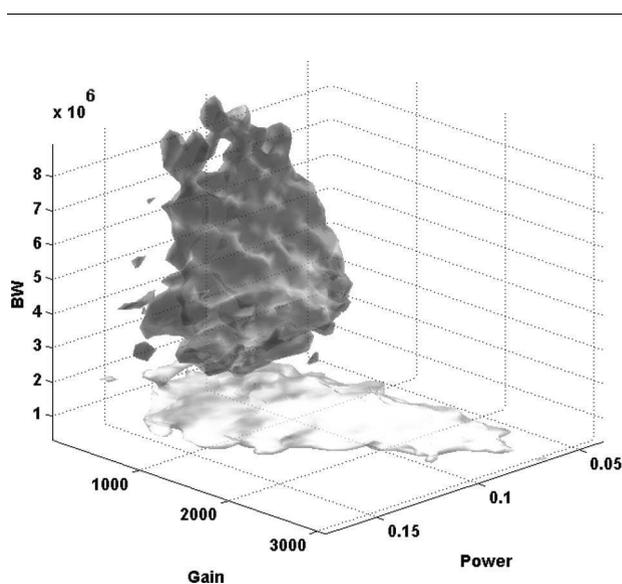


Fig. 15. 3-D feasible performance plots for the telescopic (dark region) and folded cascode (light region) topologies.

analog components first, a level-1 platform library is built first, abstracting circuit functionalities and feasible performance. Two amplifier topologies are included in the library, being a folded-cascade and a telescopic amplifier. A set of constraints including basic biasing requirements, minimum gain and phase margin have been exploited to generate the configuration constraints (represented in an analog constraint graph or ACG). The top-left corner of Fig. 14 shows the ACG used for telescopic topology, as described in [53]. Fig. 15 shows a 3-D projection of the 6D feasible performance model for both amplifiers. A sample-and-hold amplifier (SHA) platform instance is then built at level-1, and characterized as a component for the level-2 *mixed signal* platform library. In this particular case, the validity laws \mathbf{L}^1 of the SHA platform instance are trivially satisfied as the interface (output load of the amplifier) is fixed.

On the digital side, a similar process is taking place. The actual accuracy of the algorithm proposed in [23], which models the overall nonlinear behavior of the converter as the coefficients of a third order polynomial, is characterized through extensive simulations. Three implementations of polynomial corrector modules are proposed and characterized in terms of accuracy and power consumption (through synthesis and mapping on a standard cell library). Finally, a digital corrector platform instance is assembled and characterized as a digital component (GDEC) in the level-2 mixed signal platform. The bottom-up phase therefore builds a platform stack where multiple analog implementation architectures are presented at a common level of abstraction together with digital enhancement platforms, each component being annotated with feasible performance spaces.

The top-down phase consists of solving an optimization problem addressing power minimization of the overall system constrained on linearity and noise requirements from the system level and feasible configuration models on the architectural side. Solving the optimization problem using techniques such as simulated annealing, the results shown in Table 2 are produced. Topology selection of the optimum amplifier and polynomial inversion algorithm is automatically performed, resulting in the telescopic topology and a moderately complex polynomial inversion being selected. A-PBD hence allowed performing concurrent analog/digital optimization and design space explorations through a set of accurate models that export architectural effects at high levels of abstraction. An overall power consumption reduction of 64% was achieved

Table 2 Optimization Results for the Pipeline ADC First Stage Optimization. The mapped column report simulation results achieved after propagating top-level optimal parameters to the circuit level. The reference column reports performance of the initial OTA used in the ADC. Optimization performed on the folded cascade alone (fourth column) shows significantly reduced gain in power consumption.

Performance	Optimal (Telescopic)	Mapped	Folded Cascade	Reference
DNL (LSB)	0.4	0.6	0.44	0.68
INL (LSB)	0.1	0.21	0.15	0.26
SNR (dB)	86.3	86.4	84.3	84.3
Power (mW)	52.5	52.6	102	146
Power _{SHA} (mW)	47.7	47.9	97	146
A_{V0}	220	214	1,186	1,492
Bandwidth (MHz)	3.3	3.3	0.79	1.35
V_{noise} (mV rms)	2.5	2.61	7.9	8.86
G	7.3	7.27	7.8	7.94
Power _{GDEC}	4.8	4.8	4.2	-

with a closed-loop SHA, a result that is comparable to the ones reported for open-loop solutions.

V. SUMMARY AND CONCLUSIONS

This paper presented a summary of the challenges and opportunities awaiting mixed-signal SoC in deep-submicrometer CMOS technologies. The underlying message is quite simple. If further integration of complete systems onto a single die is to continue, designers must embrace circuit, architecture, and system techniques and methodologies that depart from the business as usual. Exploiting the abundance of digital transistors to make up for the shortcomings of scaled analog is one option. New architectures (for instance, using redundancy) is another. Novel devices or components offer further opportunities. Yet, handling the complexities that come with these choices is only possible if a rigorous design flow with clear abstractions is adopted. Only then will mixed-signal emerge from the custom design philosophy that has been its earmark forever. ■

Acknowledgment

The authors wish to acknowledge the contributions of the students and sponsors of the Berkeley Wireless Research center, as well as the SRC and the SIA MARCO centers (GSRC and C2S2) to the topics discussed in this paper.

REFERENCES

- [1] H. Samuelli, "Broadband communications ICs: Enabling high-bandwidth connectivity in the home and office," in *Dig. Tech. Papers IEEE Int. Solid-State Circuits Conf. (ISSCC 1999)*, pp. 26–30.
- [2] O. Ergodan et al., "A single-chip quad-band GSM/GPRS transceiver in 0.18 μm Standard CMOS," in *Dig. Tech. Papers IEEE Int. Solid-State Circuits Conf. 2005 (ISSCC 2005)*, pp. 318–319.
- [3] H. Darabi et al., "A fully integrated SoC for 802.11b in 0.18 μm CMOS," in *Dig. Tech. Papers IEEE Int. Solid-State Circuits Conf. 2005 (ISSCC 2005)*, pp. 96–97.
- [4] W. Krenik, D. Buss, and P. Rickert, "Cellular handset integration—SIP versus SOC," in *Proc. IEEE 2004 Custom Integrated Circuits Conf.*, pp. 63–70.
- [5] S. Borkar, "Design challenges of technology scaling," *IEEE Micro*, vol. 19, no. 4, pp. 23–29, Jul.–Aug. 1999.
- [6] P. P. Gelsinger, "Microprocessors for the new millennium: Challenges, opportunities and new frontiers," in *Dig. Tech. Papers*

- IEEE Int. Solid-State Circuits Conf. (ISSCC'01)*, pp. 22–25.
- [7] H. P. Hofstee, “Power constrained micro-processor design,” in *Proc. IEEE Conf. Computer Design*, 2002, pp. 14–16.
- [8] D. Markovic, V. Stojanovic, B. Nikolic, M. A. Horowitz, and R. W. Brodersen, “Methods for true energy-performance optimization,” *IEEE J. Solid-State Circuits*, vol. 39, no. 8, pp. 1282–1293, Aug. 2004.
- [9] R. H. Dennard, F. H. Gaensslen, H. N. Yu, V. L. Rideout, E. Bassous, and A. R. LeBlanc, “Design of ion-implanted MOSFET’s with very small physical dimensions,” *IEEE J. Solid-State Circuits*, vol. SC-9, no. 5, pp. 256–268, Oct. 1974.
- [10] The 2004 International Technology Roadmap for Semiconductors. [Online]. Available: <http://www.itrs.net/Common/2004Update/2004Update.htm>
- [11] D. Liu and C. Svensson, “Trading speed for low power by choice of supply and threshold voltage,” *IEEE J. Solid-State Circuits*, vol. 28, no. 1, pp. 10–17, Jan. 1993.
- [12] R. Gonzalez, B. Gordon, and M. A. Horowitz, “Supply and threshold voltage scaling for low power CMOS,” *IEEE J. Solid-State Circuits*, vol. 32, no. 8, pp. 1210–1216, Aug. 1997.
- [13] K. Nose and T. Sakurai, “Optimization of V_{DD} and V_{TH} for low-power and high-speed applications,” in *Proc. Asia South Pacific Design Automation Conf.*, 2000, pp. 469–474.
- [14] H. Ishiuchi et al., “Embedded DRAM technologies,” in *Tech. Dig. 1997 Int. Electron Devices Meeting (IEDM'97)*, pp. 33–36.
- [15] Y. Uemoto, E. Fujii, A. Nakamura, K. Senda, and H. Takagi, “A stacked-CMOS cell technology for high-density SRAM’s,” *IEEE Trans. Electron Devices*, vol. 39, no. 10, pp. 2359–2363, Oct. 1992.
- [16] J. Wu, D. Weiss, C. Morganti, and M. Dreesen, “The asynchronous 24 MB on-chip level-3 cache for a dual-core titanium-family processor,” in *Proc. Int. Solid State Circuits Conf.*, 2005, pp. 488–489.
- [17] H. Qin, Y. Cao, D. Markovic, A. Vladimirescu, and J. Rabaey, “SRAM leakage suppression by minimizing standby supply voltage,” in *Proc. 5th Int. Symp. Quality Electronic Design (ISQED'04)*, pp. 55–60.
- [18] E. Seevinck, F. J. List, and J. Lohstoh, “Static-noise margin analysis of MOS SRAM cells,” *IEEE J. Solid-State Circuits*, vol. SC-22, no. 5, pp. 748–754, Oct. 1987.
- [19] E. Morifuji et al., “New guideline of Vdd and Vth scaling for 65 nm technology and beyond,” in *Dig. Tech. Papers 2004 Symp. VLSI Circuits*, Jun. 2004, pp. 164–165.
- [20] A. A. Abidi, “High-frequency noise measurements on FET’s with small dimensions,” *IEEE Trans. Electron Devices*, pp. 1801–1805, Nov. 1986.
- [21] D. W. Cline and P. R. Gray, “A power optimized 13-b 5 Msamples/s pipelined analog-to-digital converter in 1.2 μm CMOS,” *IEEE J. Solid-State Circuits*, vol. 31, pp. 294–303, Mar. 1996.
- [22] Y. Chiu et al., “A 14-b 12-MS/s CMOS pipeline ADC with over 100-dB SFDR,” *IEEE J. Solid-State Circuits*, vol. 39, no. 12, pp. 2139–2151, Dec. 2004.
- [23] B. Murmann and B. E. Boser, “A 12-bit 75-MS/s, pipelined ADC using open-loop residue amplification,” *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2040–2050, Dec. 2003.
- [24] J. L. McCreary and P. R. Gray, “All-MOS charge redistribution analog-to-digital conversion techniques—Part I,” *IEEE J. Solid-State Circuits*, vol. SC-10, no. 6, pp. 371–379, Dec. 1975.
- [25] R. E. Suarez, P. R. Gray, and D. A. Hodges, “All-MOS charge-redistribution analog-to-digital conversion techniques—Part II,” *IEEE J. Solid-State Circuits*, vol. SC-10, no. 6, pp. 379–385, Dec. 1975.
- [26] M. J. M. Pelgrom et al., “Transistor matching in analog CMOS applications,” *Tech. Dig. Int. Electron Devices Meeting*, pp. 915–918, 1998.
- [27] ———, “Matching properties of MOS transistors,” *IEEE J. Solid-State Circuits*, vol. 24, no. 5, pp. 1433–1439, Oct. 1989.
- [28] M. Steyaert et al., “Threshold voltage mismatch in short-channel MOS transistors,” *Electron. Lett.*, vol. 30, pp. 1546–1548, 1994.
- [29] R. Aparicio, “Capacity limits and matching properties of integrated capacitors,” *IEEE J. Solid-State Circuits*, vol. 37, no. 3, pp. 384–393, Mar. 2002.
- [30] I. Aoki, S. D. Kee, D. B. Rutledge, and A. Hajimiri, “Fully integrated CMOS power amplifier design using the distributed active-transformer architecture,” *IEEE J. Solid-State Circuits*, vol. 37, no. 3, pp. 371–383, Mar. 2002.
- [31] N. M. Nguyen and R. G. Meyer, “Si IC-compatible inductors and LC passive filters,” *IEEE J. Solid-State Circuits*, vol. 25, no. 4, pp. 1028–1033, Aug. 1990.
- [32] H. Hashemi and A. Hajimiri, “Concurrent multiband low-noise amplifiers—theory, design, and applications,” *IEEE Trans. Microw. Theory Tech.*, vol. 50, no. 1, pp. 288–301, Jan. 2002.
- [33] A. Bevilacqua and A. M. Niknejad, “An ultrawideband CMOS low-noise amplifier for 3.1–10.6-GHz wireless receivers,” *IEEE J. Solid-State Circuits*, vol. 39, no. 12, pp. 2259–2268, Dec. 2004.
- [34] A. Sutono, D. Heo, Y.-J. Emery Chen, and J. Laskar, “High-Q LTCC-based passive library for wireless system-on-package (SOP) module development,” *IEEE Trans. Microw. Theory Tech.*, vol. 49, no. 10, pp. 1715–1724, Oct. 2001.
- [35] R. A. Johnson, C. E. Chang, P. M. Asbeck, M. E. Wood, G. A. Garcia, and I. Lagnado, “Comparison of microwave inductors fabricated on silicon-on-sapphire and bulk silicon,” *IEEE Microw. Guided Wave Lett.*, vol. 6, no. 9, pp. 323–325, Sep. 1996.
- [36] R. Ruby et al., “Ultra-miniature high-Q filters and duplexers using FBAR technology,” in *Dig. Tech. Papers IEEE ISSCC*, 2001, pp. 120–121.
- [37] R. Ruby, P. Bradley, and J. L. III, “FBAR microcap technology,” in *IEEE ISSCC Dig. Tech. Papers*, 2002.
- [38] N. Pletcher and J. Rabaey, “A 100 μW , 1.9 GHz oscillator with fully digital frequency tuning,” presented at the 2005 ESSCIRC Conf., Grenoble, France, 2005.
- [39] X. Chenggang, R. Gharpurey, T. S. Fiez, and K. Mayaram, “A Green function-based parasitic extraction method for inhomogeneous substrate layers,” in *Proc. DAC*, 2005, pp. 141–146.
- [40] T. B. Cho et al., “A 10 b, 20 MS/s, 35 mW pipeline A/D converter,” *IEEE J. Solid-State Circuits*, vol. 30, no. 3, pp. 166–172, Mar. 1995.
- [41] A. M. Abo et al., “A 1.5-V, 10-bit, 14.3-MS/s CMOS pipeline analog-to-digital converter,” *IEEE J. Solid-State Circuits*, vol. 34, no. 5, pp. 599–606, May 1999.
- [42] R. Staszewski, K. Muhammad, and D. Leipold, “Digital RF processor for cellular phones,” presented at the *Int. Conf. Computer-Aided Design*, San Jose, CA, 2005.
- [43] H. Kim, “A 30 mW 8b 200 MS/s pipelined CMOS ADC using a switched-opamp technique,” in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2005, pp. 284–285.
- [44] K. Bult et al., “A fast-settling CMOS op amp for SC circuits with 90-dB DC gain,” *IEEE J. Solid-State Circuits*, vol. 25, no. 12, pp. 1379–1384, Dec. 1990.
- [45] Y. Chiu et al., “A 14-b 12-MS/s CMOS pipeline ADC with over 100-dB SFDR,” *JSSC*, vol. 39, pp. 2139–2151, Dec. 2004.
- [46] Y. Chiu, B. Nikolic, and P. R. Gray, “Scaling of analog-to-digital converters into ultra-deep submicron CMOS,” presented at the *IEEE Custom Integrated Circuits Conf.*, San Jose, CA, 2005.
- [47] A. Bosi et al., “An 80 MHz $4\times$ oversampled cascaded delta-sigma-pipelined ADC with 75 dB DR and 87 dB SFDR,” in *Dig. Tech. Papers IEEE Int. Solid-State Circuits Conf. 2005 (ISSCC 2005)*, pp. 174–175.
- [48] A. Varzaghani and C.-K. K. Yang, “A 600 MS/s, 5-bit pipelined analog-to-digital converter for serial-link applications,” in *Dig. Tech. Papers 2004 Symp. VLSI Circuits*, pp. 276–279.
- [49] O. E. Erdogan, P. J. Hurst, and S. H. Lewis, “A 12-b digital-background-calibrated algorithmic ADC with -90 -dB THD,” *IEEE J. Solid-State Circuits*, vol. 33, no. 12, pp. 1812–1820, Dec. 1999.
- [50] A. Karanicolas et al., “A 15-b 1-MSample/s digitally self-calibrated pipeline ADC,” *IEEE J. Solid-State Circuits*, vol. 28, no. 12, pp. 1207–1215, Dec. 1993.
- [51] A. A. M. Saleh and J. Salz, “Adaptive linearization of power amplifiers in digital radio systems,” *Bell Syst. Tech. J.*, vol. 62, pp. 1019–1033, Apr. 1983.
- [52] M. Faulkner and M. A. Briffa, “Amplifier linearization using RF feedback and feedforward techniques,” in *IEEE Vehicular Technology Conf.*, 1995, pp. 525–529.
- [53] K. Keutzer, S. Malik, R. Newton, J. Rabaey, and A. Sangiovanni-Vincentelli, “System level design: Orthogonalization of concerns and platform-based design,” *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 19, no. 12, pp. 1523–1543, Dec. 2000.
- [54] F. Balarin, Y. Watanabe et al., “Metropolis: An integrated electronic system design environment,” *Computer*, vol. 36, no. 4, pp. 45–52, Apr. 2003.
- [55] F. De Bernardinis, P. Nuzzo, and A. Sangiovanni-Vincentelli, “Mixed signal design space exploration through analog platforms,” in *Proc. DAC*, 2005, pp. 875–880.
- [56] J. Roychowdury, “An overview of automated macromodeling techniques for mixed-signal systems,” in *Proc. IEEE CICC*, 2004, pp. 109–116.
- [57] R. Brayton, D. Hachtel, and A. Sangiovanni-Vincentelli, “A survey of optimization techniques for integrated circuits,” *Proc. IEEE*, vol. 69, no. 10, pp. 1334–1362, Oct. 1981.
- [58] H. Liu, A. Singhee, R. Rutenbar, and L. Carley, “Remembrance of circuit past: Macromodeling by data mining in large analog design spaces,” in *Proc. DAC*, 2002, pp. 437–442.
- [59] T. McConaghy, T. Eecklaert, and G. Gielen, “CAFFEINE: Template-free symbolic model generation of analog circuits via canonical

- form functions and genetic programming," in *Proc. DATE*, 2005, pp. 1082–1087.
- [60] F. De Bernardinis and A. Sangiovanni-Vincentelli, "Analog constraint graphs," presented at the *ICCAD 2005*, San Jose, CA.
- [61] F. De Bernardinis, M. Jordan, and A. Sangiovanni-Vincentelli, "Support vector machines for analog circuit performance representation," in *Proc. DAC 2003*, pp. 964–969.
- [62] B. E. Boser, I. M. Guyon, and V. N. Vapnik, "A training algorithm for optimal margin classifiers," *Proc. 5th Annu. ACM Workshop COLT*, D. Haussler, Ed., 1992, pp. 144–152.
- [63] F. De Bernardinis, S. Gambini, F. Vincis, F. Svelto, R. Castello, and A. Sangiovanni-Vincentelli, "Design space exploration for a UMTS frontend exploiting analog platforms," in *Proc. ICCAD*, 2004, pp. 923–930.
- [64] R. Staszewski et al., "All-digital TX frequency synthesizer and discrete-time receiver for bluetooth radio in 130 nm CMOS," *IEEE J. Solid State Circuits*, vol. 39, no. 12, pp. 2278–2291, Dec. 2004.
- [65] J. Crols and J. M. Steyaert, "Switched-opamp: an approach to realize full CMOS switched-capacitor circuits at very low power supply voltages," *IEEE J. Solid-State Circuits*, vol. 29, no. 8, pp. 936–942, Aug. 1994.
- [66] A. Abidi et al., "The future of CMOS wireless transceivers," in *Dig. Tech. Papers 1997 IEEE Int. Solid-State Circuits Conf.*, Feb. 1997, pp. 118–119.
- [67] J. Crols and M. Steyaert, "A single-chip 900 MHz CMOS receiver front-end with a high performance low-IF topology," *IEEE J. Solid-State Circuits*, vol. 30, no. 12, pp. 1483–1492, Dec. 1995.
- [68] J. Rudell et al., "A 1.9 GHz wideband IF double conversion CMOS integrated receiver for cordless telephone applications," *IEEE J. Solid-State Circuits*, vol. 32, no. 12, pp. 2071–2088, Dec. 1997.
- [69] D. Weaver, "A third method of generation and detection of single-sideband signals," *Proc. IRE*, pp. 1703–1705, Dec. 1956.
- [70] G. Gielen and R. Rutenbar, "Computer-aided design for analog and mixed-signal integrated circuits," *Proc. IEEE*, vol. 88, no. 12, pp. 1825–1852, Dec. 2000.
- [71] G. Hachtel, M. Lightner, and H. Kelly, "Application of the optimization program AOP to the design of memory circuits," *IEEE Trans. Circuits Syst.*, vol. CAS-22, no. 6, pp. 196–203, Jun. 1975.
- [72] G. Hachtel, T. Scott, and R. Zug, "An interactive linear programming approach to model parameter fitting and worst case design," *IEEE Trans. Circuits Syst.*, vol. CAS-27, no. 10, pp. 871–881, Oct. 1980.
- [73] W. Nye, D. C. Riley, A. Sangiovanni-Vincentelli, and A. L. Tits, "Delight.Spice: an optimization-based system for the design of integrated circuits," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 7, no. 4, pp. 501–518, Apr. 1988.
- [74] P. Allen and E. Macaluso, "Aide2: An automated analog IC design system," in *Proc. IEEE CICC*, 1985, pp. 198–201.
- [75] R. Bowman and D. Lane, "A knowledge-based system for analog integrated circuit design," in *Proc. IEEE/ACM Int. Conf. CAD*, 1985, pp. 210–212.
- [76] F. El-Turky and E. Perry, "BLADES: An A.I. approach to analog circuit design," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 8, no. 6, pp. 680–692, Jun. 1989.
- [77] M. Degrauwe et al., "IDAC: An interactive design tool for analog CMOS circuits," *IEEE J. Solid State Circuits*, vol. 22, no. 6, pp. 1106–1116, Dec. 1987.
- [78] H. Koh, C. Sequin, and P. Gray, "Opasyn: A compiler for CMOS operational amplifiers," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 9, no. 2, pp. 113–125, Feb. 1990.
- [79] R. Harjani, R. Rutenbar, and L. Carley, "Oasys: a framework for analog circuit synthesis," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 8, no. 12, pp. 1247–1266, Dec. 1989.
- [80] E. S. Ochotta, R. A. Rutenbar, and L. R. Carley, "Synthesis of high-performance analog circuits in astrx/oblx," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 15, no. 3, pp. 273–293, Mar. 1996.
- [81] M. Krasnicki, R. Phelps, R. A. Rutenbar, and L. R. Carley, "Maelstrom: Efficient simulation-based synthesis for custom analog cells," in *Proc. Design Automation Conf.*, 1999, pp. 945–950.
- [82] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge University Press, 2004.
- [83] Y. Nesterov and A. Nemirovsky, "Interior point polynomial methods in convex programming," *Studies Appl. Math.*, vol. 13, 1994.
- [84] H. Chang et al., *A Top-Down Constraint-Driven Design Methodology for Analog Integrated Circuits*. Norwell, MA: Kluwer, 1997.
- [85] H. Chang et al., "Top-down, constraint-driven methodology based generation of n-bit interpolative current source D/A," in *Proc. IEEE CICC Conf.*, 1994, pp. 369–372.
- [86] E. Ochotta, T. Mukherjee, R. A. Rutenbar, and L. R. Carley, *Practical Synthesis of High-Performance Analog Circuits*. Norwell, MA: Kluwer, 1998.
- [87] M. Hershenson, S. Boyd, and T. Lee, "GPCAD: A tool for CMOS op-amp synthesis," in *Proc. ACM/IEEE ICCAD*, 1998, pp. 296–303.
- [88] G. E. Moore, "Cramming more components onto integrated circuits," *Electronics*, vol. 38, no. 8, pp. 114–117, Apr. 19, 1965.

ABOUT THE AUTHORS

Jan M. Rabaey (Fellow, IEEE) received the electrical engineering and Ph.D. degrees in applied sciences from the Katholieke Universiteit Leuven, Belgium, in 1978 and 1983, respectively.

From 1983 to 1985, he was connected to the University of California, Berkeley, as a Visiting Research Engineer. From 1985 to 1987, he was a Research Manager at IMEC, Belgium, and in 1987, he joined the faculty of the Electrical Engineering and Computer Science department of the University of California, Berkeley, where he is now holds the Donald O. Pederson Distinguished Professorship. He was the associate chair of the EECS Dept. at Berkeley from 1999 to 2002, and is currently the Scientific codirector of the Berkeley Wireless Research Center (BWRC), as well as the director of the GigaScale Systems Research Center (GSRC). He has been a Visiting Professor at the University of Pavia (Italy), Waseda University (Japan), Technical University Delft (Netherlands), Victoria Technical University, and the University of New South Wales (Australia). His main research interests include the conception and implementation of next-generation integrated wireless systems.



This includes the analysis and optimization of communication algorithms and networking protocols, the study of ultralow-energy implementation architectures and circuits, and the supporting design automation environments.

Fernando De Bernardinis received the Laurea degree in electrical engineering from the University of Pisa, Pisa, Italy, in 1996 and the M.S. and Ph.D. degrees from the University of California, Berkeley, in 2001 and 2005, respectively.

He is currently Assistant Professor in the Department of Information Engineering at the University of Pisa. His research interests include mixed-signal design, ADC design, analog CAD, and system level design methodologies.



Ali M. Niknejad received the B.S.E.E. degree from the University of California, Los Angeles, in 1994 and the M.S. and Ph.D. degrees in electrical engineering from the University of California, Berkeley, in 1997 and 2000, respectively.

From 2000 to 2002 he worked at Silicon Laboratories, Austin, TX, where he was involved with the design and research of CMOS RF integrated circuits and devices for wireless communication applications. Currently he is an Assistant Professor in the EECS department at the University of California, Berkeley. His current research interests lie within the area of analog integrated circuits, device modeling, and electromagnetics, particularly as applied to wireless and broadband communication circuits.

Dr. Niknejad is an active member at the Berkeley Wireless Research Center (BWRC), and he is the codirector of the BSIM Research Group. He has served as an Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS and on the Technical Program Committee of the CICC and ISSCC conferences.



Borivoje Nikolić (Senior Member, IEEE) received the Dipl.Ing. and M.Sc. degrees in electrical engineering from the University of Belgrade, Yugoslavia, in 1992 and 1994, respectively, and the Ph.D. degree from the University of California, Davis, in 1999.

He was on the faculty of the University of Belgrade from 1992 to 1996. He spent two years with Silicon Systems, Inc., Texas Instruments Storage Products Group, San Jose, CA, working on disk-drive signal processing electronics. In 1999, he joined the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, where he is now an Associate Professor. He is coauthor of the book *Digital Integrated Circuits: A Design Perspective*, 2nd ed. (Prentice-Hall, 2003). His research activities include high-speed and low-power digital integrated circuits and VLSI implementation of communications and signal processing algorithms.

Dr. Nikolić received the IBM Faculty Partnership Award in 2005, the NSF CAREER award in 2003, the College of Engineering Best Doctoral Dissertation Prize, and the Anil K. Jain Prize for the Best Doctoral Dissertation in Electrical and Computer Engineering at the University of California, Davis, in 1999, as well as the City of Belgrade Award for the Best Diploma Thesis in 1992. For work with his students and colleagues he received the Best Paper Award at the ACM/IEEE International Symposium of Low-Power Electronics in 2005, and the 2004 Jack Kilby Award for the Outstanding Student Paper at the IEEE International Solid-State Circuits Conference.



Alberto Sangiovanni-Vincentelli (Fellow, IEEE) received the Dr. Eng. degree (*summa cum laude*) in electrical engineering and computer science from the Politecnico di Milano, Italy, in 1971.

He holds the Edgar L. and Harold H. Buttner Chair of Electrical Engineering and Computer Sciences at the University of California, Berkeley. He was a cofounder of Cadence and Synopsys, the two leading companies in the area of electronic design automation (EDA). He is the Chief Technology Adviser of Cadence. He is a member of the Board of Directors of Cadence and the Chair of the Technology Committee, UPEK, a company he helped spinning off from ST Microelectronics, where he is the Chair of its Nominating and Governance Committee and a member of the Audit Committee, Sonics, where he serves as the Chair of the Nominating and Governance Committee, Gradient, where he is a member of the Compensation committee and Accent, an ST Microelectronics-Cadence joint venture he helped to found. He is the Technology Advisor to the President of the Abruzzo Region. He is a member of the HP Strategic Technology Advisory Board, of the Science and Technology Advisory Board of General Motors and of the Scientific Council of the Tronchetti Provera foundation. He has consulted for many companies, including Bell Labs, IBM, Intel, United Technology, COMAU, Magneti Marelli, Pirelli, BMW, Daimler-Chrysler, Fujitsu, Kawasaki Steel, Sony, ST, and Hitachi. He is the founder and Scientific Director of PARADES, a European Group of Economic Interest supported by Cadence and ST Microelectronics. He is a member of the High-Level Group and of the Steering Committee of the EU Artemis Technology Platform. He is an author of over 700 papers and 15 books in the area of design tools and methodologies, large-scale systems, embedded controllers, hybrid systems, and innovation.

Dr. Sangiovanni-Vincentelli has been a Member of the National Academy of Engineering, the highest honor bestowed upon a U.S. engineer, since 1998. He received the Distinguished Teaching Award of the University of California in 1981. He received the worldwide 1995 Graduate Teaching Award of the IEEE for "inspirational teaching of graduate students." In 2002, he was the recipient of the Aristotle Award of the Semiconductor Research Corporation. He has received numerous best paper awards, including the Darlington and Guillemín-Cauer Awards. In 2001, he was given the prestigious Kaufman Award of the Electronic Design Automation Council for pioneering contributions to EDA.

