

CMOS RF Integrated Circuits: Past, Present and Future (Invited)

Thomas H. Lee
Stanford University
<http://www-smirc.stanford.edu>

Abstract

CMOS RF ICs continue their transition from academic curiosities to practical devices. Recent milestones at the device- and building block-level include: Noise figures for single-ended LNAs of ~ 1 dB in the low-GHz range; fully integrated oscillators with phase noise compliant with GSM specifications at under 10mW power consumption; 5GHz injection-locked frequency dividers with sub-mW power consumption and large ($\sim 30\%$) tuning range; MOSFET ring mixers with over 10dBm IIP3 and "zero" power consumption; shielded spiral and helical inductors with improved Q ; accumulation-mode MOS varactors with Q -frequency products in excess of 100-200GHz and tuning ratios approaching 2:1; fixed lateral flux capacitors with still higher Q - f products; coplanar transmission lines with < 0.3 dB/mm attenuation at 50GHz; microwave-compatible ESD structures; and constant- g_m biasing for process independence. These advances have combined to enable a single-chip GPS receiver in $0.5\mu\text{m}$ to exhibit 2.8dB overall NF, 57dB SFDR at 115mW, and a 5GHz HIPERLAN receiver to exhibit 5dB NF at < 45 mW using a $0.25\mu\text{m}$ technology. Scaling forces that are already firmly in place will continue to drive still greater improvements in performance.

Introduction

Two important scaling trends are making CMOS increasingly attractive for RF applications. One is the well known dramatic shrinkage of feature size, so that transistors in the $0.18\mu\text{m}$ generation of digital CMOS presently ramping into large-volume production have peak f_T values in excess of 55GHz (and still higher f_{max}), in keeping with the trend of doubling every three years that has prevailed for over two decades. Less widely appreciated, but arguably as important, is the *reverse* scaling of interconnect. More, and thicker, layers of wiring are continually available, enabling the realization of high-quality passive components so critical for RF circuits. These two trends have driven rapid progress, so that building blocks with laughably inferior characteristics just a few years ago have evolved into today's highly integrated receivers with performance competitive with those built in more traditional, "superior" RF technologies. The tremendous investment in keeping CMOS conformant with Moore's law assures continued improvements in performance for at least several more generations.

Supplementing those developments are an appreciation for induced broadband gate noise, which leads to a definite optimum width for LNA devices; a new theory of oscillator phase noise, which has identified the key role played by symmetry in controlling upconversion of the huge $1/f$ noise of MOSFETs; simple, accurate analytical formulas for spiral inductors that greatly facilitate design and optimizations; and a renewed appreciation of injection locking as a low-power, high-frequency circuit technique.

Because of their importance, we begin with a brief survey of some relevant scaling trends and their effects on the performance of both transistors and passive components. We continue with an examination of passive components (inductors and capacitors) and key building blocks (LNA, mixer, oscillator/synthesizer), followed by a discussion of some of the recent theoretical developments that have enabled designers to squeeze additional performance out of CMOS. We conclude with an examination of two system-level implementations (one older, one recent), and a few remarks about the likely evolution of CMOS RF.

Scaling Trends in Brief

We have already noted in the introduction that Moore's law continues to be obeyed. While transistors continue to get faster as a consequence (> 100 GHz f_T MOSFETs will be in production in 2-3 years, and 150-200GHz devices have already been demonstrated in the laboratory), several other important phenomena accompany this scaling. One is the increasing prominence of velocity saturation. This effect, already long enjoyed in the GaAs world, by itself tends to cause device transconductance to approach a constant value. Although the resulting departure from square-law behavior thus reduces device transconductance, the linearization provided by velocity saturation is a welcome compensating trait for amplifiers.

In addition to Moore's law, of course, there is the somewhat more powerful and universal one due to Murphy. It's operative here in the worsening vertical-field mobility degradation that accompanies scaling. As the gate overdrive voltage ($V_{gs} - V_T$) increases, mobility and transconductance diminish. Thus, as the gate overdrive increases from zero volts, a modern short-channel MOSFET starts off exhibiting square-law behavior, makes a transition to a constant- g_m regime, and then displays decreasing transconductance.:




Figure 1: Idealized I_D - V_{DS} curves

This tri-mode behavior challenges the LNA designer who seeks to maximize linearity. Maximum IIP3 results when the DC gate overdrive biases the transistor near the center of the tri-mode regime. The great spread in IIP3 values

conventional vertical MIM structures. Depending on the particular fractal geometry chosen, Q -frequency products in excess of 200GHz are readily achievable.

In addition to the development of a new fixed capacitor, varactors with excellent characteristics have been developed recently in CMOS. In the past, varactors have been fashioned out of either a p+/n-well diode, or from a MOSFET gate capacitor biased near threshold. Both structures can exhibit capacitance tuning ratios on the order of 2, but also demonstrate poor Q (often well below 10), owing to high series resistance in the well and channel, respectively. Much higher Q values (or, more specifically, Q -frequency products) are possible if one violates some layout rules (without requiring additional process steps, however) to make an accumulation-mode MOS capacitor (see Figure 1). Operation of this MOS structure is near the flatband voltage, rather than near threshold. As a result, the conductivity of the channel is enhanced (on average), relative to the standard MOS gate capacitor, and Q improves greatly. One may obtain Q -frequency values of the order of 100-200GHz using multiple parallel-connected fingers of minimum-length devices. Capacitance tuning ranges approach 2:1 in the limit, and 1.6:1 to 1.8:1 in practice [5].




Figure 3: Accumulation-mode varactor

From the foregoing, it is clear that capacitors in CMOS are generally satisfactory. Sadly, this statement does not apply as readily to inductors. As is widely appreciated, planar spiral inductors in CMOS technology have poor Q (generally below 10) because energy can couple into the lossy substrate in two ways. One mechanism simply involves capacitive coupling, where the inductor windings represent one plate, and the substrate acts as a lossy second plate. The loss due to this mechanism can be similar in magnitude to the conductor loss of the windings.

A second loss mechanism involves eddy currents induced to flow in the substrate by the currents flowing in the inductor windings. The loss associated with this mechanism is especially serious at higher frequencies and with the heavily-doped substrates used in epi processes.

Optimal design has been frustrated by the lack of accurate and simple analytical formulas, so Q values achieved in practice are typically widely distributed, with many falling considerably short of the (already unimpressive) maximum achievable values. By eliminating the lossy capacitance of

the substrate, however, the patterned ground shield (PGS) simplifies the modeling enough to facilitate development of a simple optimization procedure, while reducing the amount of energy coupled into the substrate, thereby boosting Q [6].




Figure 4: Patterned ground shield

Slots are cut in the PGS to prevent eddy currents from being induced in the grounded shield. These slots must be wide enough so that edge-to-edge capacitance does not inadvertently provide a path for circulating currents at high frequencies; this requirement is readily satisfied in practice.

The PGS acts as a low-loss pseudosubstrate as far as electric fields are concerned; the actual silicon substrate becomes largely irrelevant, at the expense of increased parasitic capacitance. When the inductor is used as part of a tank, this additional capacitance is not a serious problem, as it may be readily absorbed into the resonator. In any case, reverse scaling allows us to use interconnect layers that are an increasing distance away from the substrate, reducing this penalty with each successive generation. It is important to note, however, that the PGS is most effective at lower frequencies, and when the substrate is not heavily doped. When these conditions are satisfied, a doubling of resonator Q to near-GaAs values is an entirely realistic expectation.

Unfortunately, epi processes with very heavily-doped substrates represent a sizable fraction of modern technology offerings. A typical substrate resistivity in such a process might be as low as 10m Ω -cm, corresponding to a skin depth of only 160 μ m at 1GHz. Because substrates are certainly thicker, and desired operational frequencies often higher, than these values, the effect of eddy currents generally will be significant, and the PGS will be less (perhaps much less) effective. In such cases, a smaller, but still often substantial, improvement remains possible if the inductor diameter is *reduced*. Although this action ultimately exacerbates conductor loss, the reduction in eddy current loss will usually be larger, leading to a net increase in Q . This idea can be taken further by exploiting both the third dimension and reverse scaling. By building a vertically oriented solenoidal structure (or a stack of planar spirals) out of multiple layers, the overall diameter can be kept small without grossly increasing the conductor loss

[7]. Combining this strategy with the PGS restores much of the lost Q typically encountered with epi processes. Although Q values of the order of 10 can hardly be considered large, they are often large enough to be useful, nonetheless.

It is also important to note that the patterned ground shield significantly reduces coupling through the substrate. Experiments reveal reductions of 20dB or more in the low-GHz frequency range, relative to conventional unshielded structures. The important implications of this improved isolation for mixed-signal implementations hardly require emphasis.

The substrate problem also affects the design of on-chip transmission lines as well [8]. By using somewhat narrower conductor spacings than are common in other technologies, lateral flux increases, once again reducing the coupling of energy into the substrate. Using this approach, coplanar waveguides with attenuation of $\sim 0.3\text{dB/mm}$ at 50GHz have been demonstrated. This level of loss compares favorably with the 0.2dB/mm values achieved in GaAs. These lines are good enough to enable the fabrication of a K -band traveling-wave amplifier (23GHz unity-gain frequency), as well as a 17GHz traveling-wave oscillator in $0.18\mu\text{m}$ technology [9].

It is clear that, despite the lossiness of the typical CMOS silicon substrate, the large (and increasing) number of interconnect layers nonetheless makes it possible to build passive components whose properties are adequate (and, in some cases, excellent) for most RF applications.

Low-Noise Amplifiers in CMOS

It is unfortunate that $1/f$ noise MOS noise has received vastly more attention than has broadband noise. This situation has begun to change only quite recently. Although induced gate noise has been well understood in the GaAs MESFET world for quite some time, it is only in the last couple of years that the broader CMOS design community has begun, grudgingly, to acknowledge its existence in MOSFETs. Even so, there remains a certain level of controversy concerning its importance. Clarifying this issue continues to be a challenge, so perhaps the author will be forgiven for spending a little time here to outline the situation in some detail.

First, if one may neglect the noise generated by the resistive gate material, there are two principal sources of broadband noise in a MOSFET. Both ultimately trace back to thermal agitation of channel charge. These noise generators may be represented by current sources tied between the gate and source, and between the drain and source, respectively. The noise power spectral densities are given by the following expressions:

$$\overline{i_{nd}^2}/\Delta f = 4kT\gamma g_{d0} \quad (1)$$

$$\overline{i_{ng}^2}/\Delta f = 4kT\delta g_g \quad (2)$$

The parameter g_{d0} is the drain-source conductance at zero drain-source voltage. Because it is evaluated at zero bias, it is a more reliable measure of channel charge (the ultimate source of the noise) in the short-channel case than is g_m , although the latter is still frequently found in the literature.

The gate conductance factor, g_g , is given by

$$g_g = \frac{\omega^2 C_{gs}^2}{5g_{d0}} \quad (3)$$

This noisy shunt gate conductance is a directly measurable quantity since both it and the noisy gate current are consequences of distributed feedback from the channel to the gate. When transformed to an equivalent series resistance, one finds a value of $1/5g_{d0}$. Resonating out the gate capacitance with a high- Q inductance permits a measurement of this resistance. This measurement has been made, and the results are in excellent accord with this theory, providing additional support for the existence of gate noise [10].

For long-channel devices, values for γ and δ in saturation are $2/3$ and $4/3$, respectively. Because thermally agitated channel charge is the fundamental noise source, it is not surprising that the two noise currents are partially correlated. Again for long devices, the correlation coefficient is $0.395j$. Measurements have confirmed that long devices do indeed behave according to these expectations.

As mentioned previously, it has been known for a decade that short-channel MOSFETs in saturation exhibit excess broadband RF noise. The excess noise is generally attributed to hot electron effects that are known to become more prominent in the deep submicron regime, so there has been speculation that increasing excess noise might offset improvements in f_T , causing noise figures either to improve weakly, or perhaps even worsen, as devices scale to ever smaller dimensions.

Recent work with a hydrodynamic carrier transport model has verified that carrier heating by the large electric fields present in short devices does indeed cause both γ and δ to increase [11]. However, the correlation coefficient also increases in magnitude at the same time. As is known from classical two-port noise theory, increases in correlation reduce NF [12], offsetting the increases in γ and δ with increasing drain-source voltage, and with device scaling. Measurements on a number of LNAs built in a variety of technologies verify that the attainable noise figures are not very much worse than expected from long-channel theory [3], [13]. For example less than a doubling of γ explains the observed LNA NFs of [13]. Hence, one may indeed expect device noise figure to continue improving with scaling according to [12]:

$$F_{min} = 1 + \frac{2}{\sqrt{5}} \frac{\omega}{\omega_T} \sqrt{\gamma\delta(1-|c|^2)} \quad (4)$$

This equation yields an F_{min} of 0.2dB and 0.4dB for ω_T/ω equal to 20 and 10, respectively, assuming a doubling of γ and δ from their long-channel values, and an increase in $|c|$ to 0.8. Given the high values of f_T , it is clear that MOSFETs can already provide excellent noise figures in the low-GHz range.

Now it is well understood that device and circuit noise figures are generally two very different things. A particular challenge is to make a primarily capacitive device present a real-valued input impedance. Of the many possible ways to do so, the best by far is the use of inductive degeneration:




Figure 5: Inductively-degenerated LNA (simplified)

This circuit traces its ancestry to a vacuum tube progenitor first analyzed by van der Ziel in the late 1930s. This circuit was largely forgotten for several decades, although it is discussed in volume 18 of the *MIT Radiation Laboratory* series. It was rediscovered by engineers working in GaAs technology around 1980, and subsequently adapted for CMOS implementations. The inductor L_S provides series feedback, delaying the flow of source-drain current relative to the application of a gate voltage. The delay adds to the conventional quadrature phase relationship between voltage and current that normally prevails in a capacitor, giving rise to a real term in the impedance. Because this shift is accomplished with a reactance, the real impedance is not accompanied by the thermal noise of an equivalent resistor. A proper choice of degenerating inductance produces the desired resistance, while inductance L_g resonates out the remaining capacitive reactance, leaving a pure resistance. Part of this resistance ($1/5g_{d0}$) is due to nonquasistatic phenomena as discussed earlier, and should be properly taken into account in designing for an input match.

Using the correct noise model with this circuit, one may develop a surprisingly simple set of design equations. Perhaps the most important is that there is a unique device width that minimizes noise figure for a fixed power budget and for a given input match. The existence of this optimum is easily explained by considering what happens at extremes of device width. Very narrow devices will have high f_T because of the higher current density, but will require very large inductances to produce resonance. The impedance of the network as seen by the gate noise current will therefore be high, leading to a dominance by gate noise.

Very wide devices, on the other hand, suffer from low f_T . Power gain is consequently degraded, leading to lower output signal power. This condition allows the drain current noise to dominate. In between these two extremes is the optimum condition, which balances gate and drain noise properly. Derivation of the optimum condition is somewhat complicated, but its expression is not: The product of device width and operating frequency should be approximately $500\mu\text{m}\cdot\text{GHz}$ in a 50Ω system. This choice fortunately corresponds to an input Q in the range of approximately 2-3, so the percentage bandwidth is fairly large, relaxing requirements on component tolerances.

As summarized earlier, CMOS LNAs have improved greatly in performance as technology has scaled. Early GHz-level LNAs exhibited $>10\text{dB}$ noise figures, in contrast with sub-1dB values that are now achievable for $\sim 10\text{--}15\text{mW}$ of consumption. Simultaneously, the linearization that accompanies velocity saturation can be exploited to provide excellent IIP3 values as well. Values of approximately 0dBm are readily achieved. If some degradation of NF is permissible, the device width may be increased (and the inductances decreased) to lower the input Q and thereby increase IIP3 further. Continued scaling will permit the attainment of ever-lower noise figures at a given frequency and power consumption, providing more freedom to effect this tradeoff if desired.

CMOS Mixers

Perhaps the most common mixer topology in any IC technology is based on the Gilbert multiplier. This mixer performs multiplication in the current domain, and so must provide a voltage-to-current transformation of the incoming RF signal. The requirement for current switching is the direct result of the limitations of bipolar transistors: they can't switch voltages well. The high-quality switches available in CMOS offer an alternative implementation of mixers, thereby evading the noise and linearity degradation inherent in any $V\text{-}I$ conversion scheme.

The CMOS ring mixer is hardly novel; its forbears extend back to copper oxide diode ring modulators developed for carrier telephony in the 1920s and 1930s. Successful GaAs MESFET implementations abound. In its simplest form, such a mixer appears as follows:




Figure 6: CMOS ring mixer

A key advantage of full integration is that the load at the IF port tends to be primarily capacitive in CMOS circuits. This may be exploited to provide considerably more freedom in sizing the switches than if a 50Ω load must be driven.

Achievable IP₃ values are a function of device width. Larger widths result in lower values of switch resistance and thus improved linearity, but also increase the necessary LO power. For narrowband applications, the gate capacitors can be resonated out to reduce the required LO power. With only a few hundred μW of LO drive, IIP₃ values in excess of 10dBm are routinely achieved. Another attractive attribute of this topology is the absence of any standing DC current, reducing headroom requirements to the bare minimum. It thus scales more gracefully than do Gilbert-based mixers as supply voltages diminish. Furthermore, one would expect that the absence of a DC bias current implies a similar absence (or at least a diminution) of $1/f$ noise as well.

The small conversion loss of the commutating ring mixer is often tolerable, and can be readily compensated for in any case by the addition of a post-amplifier (whose design at IF is presumably easier than at RF). The noise figure, to a first approximation, is roughly equal to the power conversion loss.

Phase Noise of CMOS Oscillators

There are two significant challenges in designing integrated CMOS oscillators. One is the relatively poor Q factor of spiral inductors, and the other is the large transistor $1/f$ noise. The latter is upconverted into close-in noise near the carrier. MOSFETs, like MESFETs, are surface-controlled devices and therefore are exquisitely sensitive to the trapping phenomena that give rise to $1/f$ noise.

Until quite recently, the precise mechanism whereby this upconversion takes place was poorly understood. It was easy in such a circumstance to assume that the close-in phase noise of CMOS oscillators would always be hopelessly inferior. This assumption has been frequently used to dismiss "CMOS RF" as oxymoronic (or simply moronic). Fortunately, clearing up how this upconversion occurs has also pointed out the way to suppress it. To appreciate this new insight, it is necessary to revisit older theories of phase noise, identify hidden assumptions, then correct subtle (but important) errors.

Let us begin by considering Leeson's famous phase noise equation [14]:

$$L\{\Delta\omega\} = 10\log\left[\frac{2FkT}{P_{sig}}\left\{1 + \left(\frac{\omega_o}{2Q\Delta\omega}\right)^2\right\}\left(1 + \frac{\Delta\omega}{|\Delta\omega|^{1/f^3}}\right)\right] \quad (5)$$

His equation contains a factor F to account for increased noise in the $1/(\Delta\omega)^2$ region, an additive factor of unity (inside the braces) to account for a noise floor at large offsets, and a multiplicative factor (the term in the second set of parentheses) to provide a $1/|\Delta\omega|^3$ behavior at sufficiently

small offset frequencies. With these modifications, the phase noise spectrum appears approximately as follows:




Figure 7: Leeson's spectrum

Unfortunately, F is empirical. Furthermore, the Leeson model asserts that $\Delta\omega_{1/f^3}$ is precisely equal to the $1/f$ corner of device noise. However, measurements frequently show no such equality, so $\Delta\omega_{1/f^3}$ is empirical as well. Finally, the frequency at which the noise flattens out is not always equal to half the resonator bandwidth, $\omega_o/2Q$.

The inability to make quantitative predictions is a hint that there are unsupported assumptions that we need to revisit and correct. Specifically both linearity and time invariance have been assumed.

Amplitude nonlinearity is clearly a fundamental property of all real oscillators, as it is necessary for amplitude limiting. However, we will assume that if a certain amount of injected noise produces some amount of disturbance, doubling the injected noise would produce double the disturbance. Thus linearity would appear to be a reasonable assumption *as far as the noise-to-phase transfer function is concerned*.

To show that it is time-invariance that fails to hold, consider explicitly how an impulse of current affects the waveform of the simplest resonant system, a lossless LC tank:




Figure 8: LC oscillator excited by current pulse

Assume that the system is oscillating with some constant amplitude until the impulse occurs, then consider how the

system responds to an impulse injected at two different times:




FIGURE 9. Impulse responses of LC tank

If the impulse happens to coincide with a voltage maximum (as in the upper plot), the amplitude increases abruptly by an amount $\Delta V = \Delta Q/C$, but the timing of the zero crossings does not change. An impulse injected at some other time generally affects both the amplitude and the timing of the zero crossings, as in the lower plot. We therefore see that the amount of phase disturbance for a given injected impulse depends on when the injection occurs; time-invariance thus fails to hold. An oscillator is therefore a linear, but (periodically) time varying (LTV) system.

Because of linearity, the impulse response still completely characterizes the system. Noting that an impulsive input produces a step change in phase, the impulse response may be written as:

$$h_{\phi}(t, \tau) = \frac{\Gamma(\omega_0 \tau)}{q_{max}} u(t - \tau) \quad (6)$$

where $u(t)$ is the unit step. Dividing by q_{max} , the maximum charge displacement across the capacitor, makes the function $\Gamma(x)$ independent of signal amplitude. $\Gamma(x)$ is called the *impulse sensitivity function* (ISF), and is a dimensionless, frequency- and amplitude-independent function periodic in 2π . As its name suggests, it encodes information about the sensitivity of the system to an impulse injected at phase $\omega_0 \tau$. In our example of the LC oscillator, $\Gamma(x)$ has its maximum value near the zero crossings of the oscillation, and a zero value at maxima of the oscillation waveform. In general, it is most practical and accurate to determine $\Gamma(x)$ through simulation, but there are also approximate analytical methods that apply in special cases [16]. Once the ISF has been determined (by whatever means), we may compute the phase displacement due to any disturbance through use of the superposition integral. This extension is valid here since superposition depends only on linearity, not time invariance:

$$\phi(t) = \frac{1}{q_{max}} \left[\frac{c_0}{2} \int_{-\infty}^t i(\tau) d\tau + \sum_{n=1}^{\infty} c_n \int_{-\infty}^t i(\tau) \cos(n\omega_0 \tau) d\tau \right] \quad (7)$$

This equation allows us to compute the excess phase caused by an arbitrary noise current injected into the system, once the Fourier coefficients c_n of the ISF have been determined.

Lack of time invariance means that signals (noise) injected at some frequency may produce spectral components at a different frequency. Ultimately one may show explicitly that the following frequency translations of device noise into phase noise occurs:




Figure 10: Evolution of circuit noise into phase noise

Components of noise near *integer multiples* of the carrier frequency all fold into noise near the carrier itself.

Noise near DC gets *upconverted*, weighted by coefficient c_0 , so $1/f$ device noise becomes $1/f^3$ noise near the carrier; noise near the carrier stays there, weighted by c_1 ; and white noise near higher integer multiples of the carrier undergoes *downconversion*, turning into noise in the $1/f^2$ region.

From the previous paragraph we see that, since c_0 is the DC value of the ISF, there is a possibility of reducing by large factors the $1/f^3$ phase noise corner by satisfying the implied symmetry criterion. The ISF is a function of the waveform, and hence potentially under the control of the designer. This result is perhaps the most dramatic insight provided by this LTV model [15].

A configuration that can satisfy the symmetry criterion is the following negative resistance oscillator (6):




Figure 11: Simple symmetrical negative-R oscillator

By selecting the relative widths of the PMOS and NMOS devices appropriately to minimize the DC value of the ISF for each half-circuit, one may minimize the upconversion of $1/f$ noise. Phase noise is also improved through a near doubling of oscillation amplitude that is possible over some range of bias currents. Through exploitation of symmetry, the $1/f^2$ corner can be dropped by orders of magnitude. As a result, a phase noise of -121dBc/Hz at an offset of 600kHz has been obtained with on-chip spiral inductors at 1.8GHz , on 6mW of power in a $0.25\mu\text{m}$ technology [16]. This result rivals what one may achieve with bipolar technologies.

Recent Milestones

A single-chip GPS receiver in $0.5\mu\text{m}$ CMOS which integrates all RF and IF functions except for an LO reference crystal exhibits 2.8dB overall NF, 57dB SFDR at 115mW . Both the integration and performance levels compare quite favorably to GPS implementations in other technologies. More recently, a 5GHz HIPERLAN receiver with $\sim 5\text{dB}$ NF at $<45\text{mW}$ in a $0.25\mu\text{m}$ technology has also been demonstrated. Part of the reason for the good performance at low power consumption is the use of injection-locked oscillators as sub-mW frequency dividers [17].

Finally, as CMOS continues to scale, it is not altogether ridiculous to speculate about its possible use in millimeter-wave applications. Using low-loss coplanar transmission lines, a distributed amplifier with 23GHz unity-gain bandwidth and a 17GHz distributed oscillator have been constructed. All of these developments have been augmented by microwave-compatible ESD structures and constant- g_m biasing for process independence, finally removing the technical barriers to acceptance of CMOS as a viable RF technology.

Summary

The dual trends of shrinking transistor feature sizes, and reverse scaling of interconnect have solved the most serious problems, allowing digital CMOS to perform well in the RF regime. With the single exception of power amplifiers (due to the ever-diminishing breakdown voltages), all of the criti-

cal components and building blocks of transceivers in the low-GHz range of frequencies can be built out of ordinary CMOS layers. And, as these scaling trends (both forward and reverse) continue along well-established trajectories, performance will only improve, assuring a steady transition to industry from universities.

References

- [1] A. Abidi, "High-Frequency Noise Measurements on FET's with Small Dimensions," *IEEE Transactions on Electron Devices*, Nov. 1986.
- [2] H. Samavati, H. Rategh and T. H. Lee, "A 12.4mW CMOS front-end for a 5GHz Wireless LAN Receiver," *VLSI Symposium*, June 1999.
- [3] G. Brenna, "LNA Research," *Diploma Thesis*, École Polytechnique Fédérale de Lausanne, Feb. 1998.
- [4] H. Samavati et al., "Fractal Capacitors," *IEEE J. of Solid-State Circuits*, Dec. 1998.
- [5] T. Soorapanth et al., "Analysis and Optimization of Accumulation-Mode Varactor for RF IC's," *VLSI Symposium*, June 1998.
- [6] C. Yue, and S. Wong, "On-Chip Spiral Inductors with Patterned Ground Shields for Si-Based RF IC's," *1997 VLSI Circuits Conference Digest of Technical Papers*, June, 1997.
- [7] A. Abidi, "CMOS wireless transceivers: The next wave," *Proc. Norchip*, November 1998.
- [8] B. Kleveland, T. H. Lee and S. S. Wong, "50-GHz Interconnect Design in Standard Silicon Technology," *1998 MTT-S Digest*, June 1998.
- [9] B. Kleveland et al., "Monolithic CMOS Distributed Amplifier and Oscillator," *ISSCC Digest of Technical Papers*, Feb. 1999.
- [10] J. Janssens and Michiel Steyaert, "Design of Broadband Low-Noise Amplifiers in Deep-Submicron CMOS Technologies," *Proc. Advances in Analog Circuit Design*, April 1998.
- [11] J.-S. Goo et al., "RF Noise Simulation for Submicron MOSFETs Based on Hydrodynamic Model," *1999 VLSI Technology Conference Digest of Technical Papers*, June 1999.
- [12] T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*, Cambridge University Press, 1998.
- [13] D. Shaeffer et al., "A 115mW , $0.5\mu\text{m}$ CMOS GPS Receiver with Wide Dynamic-Range Active Filters," *IEEE J. of Solid-State Circuits*, Dec. 1998.
- [14] D. B. Leeson, "A Simple Model of Feedback Oscillator Noise Spectrum," *Proceedings of the IEEE*, vol. 54, pp. 329-330, Feb. 1966.
- [15] A. Hajimiri and T. Lee, "A General Theory of Phase Noise in Electrical Oscillators," *IEEE Journal of Solid-State Circuits*, vol. 33, no. 2, pp. 179-194, Feb. 1998.
- [16] A. Hajimiri and T. Lee, "Design Issues in CMOS Differential LC Oscillators," *IEEE Journal of Solid-State Circuits*, May, 1999.
- [17] H. Rategh et al., "A 5GHz , 32mW Frequency Synthesizer with an Injection Locked Frequency Divider," *1999 VLSI Circuits Conference Digest of Technical Papers*, June, 1999.