

Power Conscious Design of Wireless Circuits and Systems

Asad A. Abidi, Gregory J. Pottie, and William J. Kaiser

Electrical Engineering Department
University of California, Los Angeles

ABSTRACT

The great importance of power consciousness is well understood in mobile wireless communications. However, with growing experience the fundamental principles underlying power conscious design of RF circuits, systems, and networks are only now becoming known. Using as example ultralow power wireless devices for messaging such as paging receivers and wireless sensor networks, in the first part this paper presents the relationship between current consumption and dynamic range of low noise amplifiers, mixers, oscillators, and active filters. The second part of the paper covers issues of modulation, protocols, and networking that would be required in dense networks of wireless sensors, which communicate using very little energy. These ideas are expected to find use in most forms of digital wireless communications.

Power Conscious Design of Wireless Circuits and Systems

Asad A. Abidi, Gregory J. Pottie, and William J. Kaiser

Electrical Engineering Department
University of California, Los Angeles

I. Introduction

Power consumption is perhaps the major engineering concern in the design of mobile wireless devices. User expectations of what they will accept as mobile in terms of weight and volume have sharply risen in the last few years, spurred on by advances in consumer devices such as portable tape recorders and CD players, notebook computers and PDA's, and mobile telephones. Weight and physical volume principally determine whether a device may be called "mobile", but it must offer sophisticated features as well.

To prolong the time between battery recharge, mobile telephones must use power-conscious IC's for RF, baseband and DSP, supervised by sophisticated power management algorithms. However, as human speech is highly redundant and time-inefficient, devices such as mobile telephone which support this form of communication fundamentally do not make the best use of the energy source which powers them. At a minimum, the ideal continuously "online" communicator should be capable of receiving brief messages in an efficient binary format, and like that other familiar electronic device, the wristwatch, it should be wearable on the human body, and should worry its user least on the state of charge of its battery. The radio pager, which predates the mobile telephone, is one such device. Another is the wireless transceiver attached to a remote sensor, which must communicate a possibly seldom-occurring event to a command center. These two examples have much in common, although pagers are often receive-only, while sensor communicators are usually two-way.

A new generation of wireless paging receivers is now approaching this ideal shape and size. Entire receivers are built into a wristwatch, where the antenna consists of a filament of wire embedded into the bezel.

The energy storage element, the battery, almost entirely determines the final volume and weight of the wireless communicator. Low instantaneous and average receiver power consumption is therefore the most important way to miniaturization. Transceiver architecture and circuit design determine *instantaneous* power consumption, whereas *average* power consumption depends mainly on good power management and sensible communication protocols.

Paging receivers consume minute amounts of current. For instance, a state of the art receiver operating in the 930 MHz band is powered from a single AAA cell, which lasts up to many months. The receiver can detect signals as low as 1 μV . It comes as something of a surprise, then, that although in the same frequency band as what mobile telephones use today, and almost the same minimum detectable signal, the power consumption of the paging receiver is at least $10\times$ lower than the best receivers used in cellular telephones. The question is: How? What is special about the paging receiver, or the technologies it uses, that result in order-of-magnitude greater energy efficiency compared to a voiceband wireless receiver? That is mainly the subject of this paper.

To fulfill the vision of communication “anytime, anywhere”, the mobile telephone will fulfill a certain function. However, ultralow power wireless devices capable of at least receiving messages may be embedded into a mobile telephone, which may actually be used for short e-mails. This two-tier approach will likely but devices small enough to be worn like a wristwatch will supplement it.

There is another class of wireless communicators, perhaps not as widely used as telephones which operate at similar radio frequencies, but whose power consumption must be 10 to $50\times$ lower. These communicators often work at a low duty cycle, but they must be dependably available over many months without battery replacement or recharge. The radio-paging receiver is one example. A 900 MHz FLEX paging receiver operates for periods as long as six months from a single AAA-size cell. Like any wearable electronic device such as a wristwatch, the user need not be concerned with the state of the battery for long periods of time. Wireless communicators built into implantable biomedical devices, or in miniature remote sensors, must also be very low power. This class of wireless device must also often operate at low voltages such as 0.9V, which is the lower limit of useful life for certain common batteries.

This paper describes circuit techniques and radio architectures that have been discovered in the course of recent research which enable the power consumption of RF, IF, and baseband building blocks to be lowered by orders of magnitude in integrated wireless communicators. It does not cover other important ways to lower power consumption, such as signaling schemes, protocols, and power management. A total low power design must intelligently use all these techniques.

An understanding of the principles and practice of low power techniques is expected to not only benefit the circuits in existing mobile wireless devices such as cellular telephones, it holds the potential to practically realize an altogether new class of wearable communicator which makes wireless ubiquitous in ways that have so far only been imagined.

II. Lowering Power across the Hierarchy

A low power *system* is the result of comprehensive power awareness cutting across the entire hierarchy of a system. In a communication device, this hierarchy descends from the system definition, signaling protocols and choice of modulation, through architecture of the receiver and transmitter, to the transistor-level circuits comprising the RF and IF analog circuits and baseband DSP. It is easy to speak of these concepts in the abstract, but difficult to illustrate them with a specific example of an operational system. Fortunately, there exists a good example in the world of wireless paging: the evolution of the POCSAG paging system into today's FLEX [1].

The POCSAG paging protocols was introduced in the UK in the 1970's, and served for many years as the paging standard there, in the US, and in many other countries of the world. It transmits data at 100 b/s using binary frequency-shift keying of the carrier. Most POCSAG transmissions are broadcast from strong transmitters in the 200 or 400 MHz bands. As paging use grew, there arose the need for bandwidth-efficient modulation and higher data rate to enable a greater number of users to receive messages more frequently. This has led to the development of the FLEX paging protocol.

For higher throughput, FLEX uses 4-FSK modulation enabling the data rate to step up from 100 b/s to 1.1 kb/s. The most significant advances over POCSAG are at the system level [2, 3]. Data frames rely on time-synchronized reception. Internal clocks synchronize

the operation of all receivers within range of a particular transmitter to $5\mu\text{s}$ or better. This means that compared to asynchronous reception in POCSAG, where data frames must contain preambles long enough for any given receiver to wake up randomly and then synchronize, the preamble in FLEX may be much shorter. Synchronization also means that receivers may be powered down for multiple frames at a time. The radio section contains enough logic to decode the address field, and a FLEX receiver only powers up the microcontroller when it has recognized its own address.

Together, these techniques lower a FLEX receiver's power $10\times$ over a POCSAG receiver, while using the *same* RF circuits. This proves that good system planning can extend battery life by an order-of-magnitude without evolution in RF circuits.

The right modulation scheme, too, helps to lower power dissipation. For instance, on-off keying (OOK) or amplitude-shift keying (ASK) need only a threshold detector, which is simple to construct and consumes low power. Low-end applications such as wireless remote keyless entry use these modulations in simple, ultralow power transceivers [4].

However, the transmitted signals are susceptible to additive noise or interference.

Wideband frequency-shift keying (FSK) is less susceptible, but requires a frequency discriminator for demodulation. As the peak energy in its frequency spectrum lies on either side of the carrier frequency away from the center, it is well suited to a zero-IF direct-conversion receiver [5]. Flicker noise and DC offset in baseband circuits may be suppressed by AC coupling without removing valuable signal spectrum.

The direct conversion receiver (Fig.1) is most amenable to full integration on a single chip. It requires no image-reject filter, and only a minimum RF section comprising a low-noise amplifier (LNA) and two mixers. All subsequent amplification and filtering takes place at low frequencies, close to DC, at the price of duplicating circuits in quadrature branches. The power consumption of an active channel-select filter is least when the channel of interest lies at a low frequency [6]. As the active filter is often one of the largest consumers of power [7], zero IF is the clear choice in this respect.

III. Power Conscious RF and Baseband Circuits

This section explores the fundamental limits on power dissipation in many of the building blocks in a wireless receiver. The receiver must usually be more power conscious than the transmitter, which in the messaging-type communications

A. Dynamic Range and Power Consumption

The fundamental lower limit to a circuit's power consumption is tied to its performance. For circuits in the signal chain of a wireless receiver, the main specification is on spurious-free *dynamic range*. In decibels, this is proportional to the difference between the 3rd-order intercept point and the noise floor as measured at the receiver output, then referred to the receiver input to normalize for gain. Take, for example, the simple circuit consisting of two common-source MOSFETs driven with balanced input signal voltages, and producing output currents that are differentially sensed (Fig.2(a)). The dynamic range is limited at the lower end by the voltage noise spectral density integrated across the channel bandwidth, and at the upper end by the large signal swing that distorts gain and defines the intercept point. The equivalent input noise voltage density is $\hat{v}_n^2 = 4kT \gamma / g_m$ where kT is a physical constant and γ is a noise factor associated with the FET channel length and bias [8].

If the FETs obey the pure square law characteristics, this circuit is a perfectly linear large signal differential transconductor. Its input-output characteristic distorts significantly when the signal swing shuts off one of the FETs. If this were the only distortion, however, the circuit would show infinite 3rd-order intercept at small signal well below the shut off condition. In fact, this is not so. Gate electric field dependence of inversion layer mobility introduces a small 3rd-order nonlinearity that defines a finite intercept point even at small test signals [9]. The governing equation is:

$$I_D = \frac{W}{2L} \frac{\mu_0}{(1 + \theta V_{eff})} C'_{ox} V_{eff}^2 \quad (1)$$

where $V_{eff} \triangleq V_{GS} - V_t$. The 3rd-order intercept for a single 0.5- μm NMOSFET simulated using the Philips MOS 9 model fits well with the intercept point this equation predicts with $\theta=0.1\text{V}^{-1}$ (Fig.3). Roughly speaking, a higher intercept point in dBm calls for an almost proportionally larger bias V_{eff} in volts. Therefore, in an *open-loop* circuit as used

in the RF and IF sections, a given dynamic range fixes the FET's g_m and V_{eff} . Now assuming that the I-V characteristics roughly conforms to the classic square law, the bias drain current in saturation is:

$$I_D = \frac{1}{2} \mu C'_{ox} \frac{W}{L} V_{eff}^2 = \frac{1}{2} g_m V_{eff} \quad (2)$$

This leads to the very important conclusion that in open-loop small-signal circuits such as amplifiers, mixers, and active filters, the specification on *dynamic range determines current drain*, independently of the FET channel length and technology scaling! With good design, the actual current consumption of an amplifier can approach this limit. Therefore, on fundamental grounds a system employing very low power receivers must relax requirements on dynamic range. Alternatively, when the noise spectral density is large because of low bias current, the system may improve sensitivity by using narrowband channels, that is, low data rates.

This relationship also holds true for bipolar transistor (BJT) circuits. The intrinsic input-referred IP3 of an ideal BJT is -12.7 dBm at any bias voltage. Just as greater V_{eff} at bias extends a FET's capability to amplify large signals with low distortion, so resistor degeneration linearizes the BJT (Fig. 2(b)) and raises its intrinsic IP3. The degeneration voltage, $I \times R$, determines the maximum input voltage at the onset of gain compression, while the input-referred voltage noise is $4kT\Gamma R$, where Γ is some noise factor. As R sets the lower end of dynamic range and IR the upper end, the current drain is once again fixed.

In fact, a degenerated bipolar differential pair gives the same transconductance per unit bias current as the MOSFET transconductor, when both are designed for equal linear full-scale (Fig. 2(b)).

B. Lowering Power in Tuned Circuits

Strictly speaking, the link between current consumption and dynamic range described above applies to baseband or wideband circuits.

Now consider an inductor degenerated differential pair, where the resistor in Fig.2(b) is replaced by an inductor. As the feedback inductor is noiseless, the input-referred noise voltage is the same as the original differential pair but the degeneration extends the

linear range. Reactive degeneration thus decouples dynamic range from power consumption. This benefit from the inductor is, however, frequency dependent. Degenerating reactance is lower with frequency, and at high frequencies is limited by the onset of resonance with whatever device and parasitic capacitance appears across it. This technique has been used in some RF circuits, such as transformer-coupled stages [10].

C. Importance of Passives Quality in Resonant Circuits

The quality factor of passive tuned circuits can profoundly influence the circuit power dissipation. As a starting point, consider a low-noise amplifier whose input port impedance is *not* to be impedance matched.

First, the amplifier must provide sufficient gain in the frequency band to which it is tuned. Suppose the input capacitance C of the next stage is known and fixed. Then the load inductor (L) is chosen to resonate with C in the narrow frequency band of interest. The resulting peak voltage gain is $g_m Z_0$, where Z_0 is the tuned circuit's impedance at resonance. This impedance is usually limited by the inductor quality factor, Q , and given by $Z_0 = Q(\omega_0 L) = Q / (\omega_0 C)$. The voltage gain is thus $Q g_m / (\omega_0 C) = 2QI / (\omega_0 C V_{eff})$. To lower power dissipation, if the amplifier FET is scaled down in width, then g_m and the voltage gain will also scale down, unless Q of the inductor is raised to compensate. Fig. 4 illustrates a specific case of a 22 nH inductor tuning a common-source amplifier to 930 MHz while driving a capacitor load. For various inductor Q 's, the width of the FET (length fixed at 0.25 μm) is scaled to obtain 20 dB voltage gain. Bias V_{eff} is fixed at 200 mV to define a constant large-signal V_{GS} handling capability at the FET input port. If the 22 nH inductor is fabricated on-chip with a Q of, say 3, the FET must be biased at 3.3 mA. On the other hand if the Q were 20, the required bias current at that gain falls to 0.5 mA, resulting in a 6.6 \times reduction in power. Therefore, use of high quality inductors lowers the power dissipated in tuned RF circuits, a fact that has also been noted elsewhere [11]. Although an inductor with this large a Q cannot be integrated today as a spiral in a standard IC process, it may be an off-chip discrete component, or it may be fabricated in some specialized technology [12, 13].

Discrete wire wound chip inductors are most commonly used at RF and microwave frequencies [14, 15]. The inductor consists of a solenoid of high conductivity metal wire wound on a ceramic header, or lately with an air core. Inductance as large as 100 nH may be fit into a “chip” about 1 mm on a side (Fig.5). The small size lowers the parasitic capacitance across the inductor, extending its self-resonant frequency and the highest attainable Q . For example, a representative 68 nH chip inductor gives a Q of 60 at 1 GHz.

Every signal line on an integrated circuit connecting to an off-chip component suffers loading by parasitic capacitance and loss due to the bond pads, bondwire inductance, and the inductance and capacitance of package leads and PC board traces. The parasitics are usually large compared to the on-chip reactances, and they do not shrink with transistor scaling. Mounting bare die with solder balls on to ceramic substrates (the flip-chip technique) cuts down on many of these parasitic reactances. It is also possible to fabricate planar spiral inductors using thick, low-loss metal traces on, or within, the ceramic substrates (Fig.6); the latter multi-layer technology is called low-temperature co-fired ceramic (LTCC) [16]. This approach has been successfully used in ultralow power wireless transceivers associated with wireless sensors [11].

On-chip, a grounded shield under the bond pad substantially eliminates substrate loss [17]. For effective shielding, the grounded layer under the pad may have to be connected to off-chip ground with multiple bond wires. Now the pad appears almost purely capacitive, and the remaining parasitics are mainly reactive and lossless.

There are repercussions to using off-chip tuning elements. With inaccurate estimates of the parasitics or because of spreads in their values, a circuit node with a high loaded Q is susceptible to frequency detuning. If the resulting variation in gain due to spreads is to be kept below, say, 1 dB, then the loaded Q at that node must not be too large. Using the classic 2nd-order frequency response of an LCR circuit, this upper limit on loaded Q is:

$$Q < \frac{1}{4 \frac{\Delta\omega}{\omega_0}} = \frac{1}{2\sqrt{\left(\frac{\Delta L}{L}\right)^2 + \left(\frac{\Delta C}{C}\right)^2}} \quad (3)$$

The active device attached to this resonator must consume a minimum current to provide the required gain. One way to circumvent this limitation is to allow a higher Q ,

but compensate for spreads with some form of active tuning. For example, following auto-calibration a digital word can switch an array of small binary-weighted capacitors to re-tune a resonant circuit to the desired frequency.

D. Low Noise Amplifiers

A low noise amplifier must be matched at the input port to some characteristic impedance such as 50Ω . The first candidate for the LNA is the common-gate amplifier circuit degenerated by an inductor, L_s (Fig.7). Its input impedance is matched by resonating the net capacitive reactance of C_{GS} and L_s in series with a gate inductor, L_G [18]. The Q of the matching network amplifies the input voltage, and in the absence of any other losses this produces the LNA noise factor [19]:

$$F = 1 + \gamma g_m Z_0 \left(\frac{\omega_0}{\omega_T} \right)^2 \quad (4)$$

At first sight this suggests that if FET g_m , and therefore its size and bias current are shrunk, and at the same time the matching network is adjusted to maintain impedance match, then this will result in the lowest noise *and* a bias current approaching zero.

Potentially, this is of great value in a low power receiver.

However, by neglecting other forms of loss this limiting case is seriously in error. Series resistor R_G representing loss in the gate inductor L_G , and the inversion layer resistance R_{GS} of the MOSFET, which is a fundamental loss, both scale *up* as the LNA MOSFET shrinks in size. When they are equal to the source impedance Z_0 , the LNA noise figure is 3 dB. As the FET is scaled down further, the noise figure worsens. This implies an optimum FET size for least noise figure. Taking into account these losses, the noise factor is:

$$F = 1 + \frac{R_{GS} + R_G}{Z_0} + \frac{\gamma}{g_m Z_0 Q^2} \quad (5)$$

Q is the quality factor of the input-matching network. The main advantage of the inductively degenerated LNA is that it can achieve a low noise figure limited by the ratio of the unity current gain frequency (ω_T) to the band of operation (ω_0). There is no guarantee, though, that this optimum lies at a low bias current.

Suppose a receiver needs an LNA with 3-dB noise figure, but with lowest possible power consumption. To simplify the analysis, assume L_g is a very low loss external inductor so that the loss at the input port is mainly due to the inversion layer resistance, R_{GS} in series with C_{GS} . It has been shown that $R_{GS} \approx 1/(5g_m)$ [19, 20]. It is a reasonable assumption when the FET size is aggressively scaled down for low power that R_{GS} is the main contributor to noise figure. To justify this, the expression for noise factor is re-cast as follows, where $\omega_0 < \omega_T$ always.

$$F = 1 + \frac{R_{GS}}{Z_0} + 4\gamma g_m Z_0 \left(\frac{\omega_0}{\omega_T} \right)^2 \simeq 1 + \frac{R_{GS}}{Z_0} \quad (6)$$

For LNA noise figure of 3 dB, that is $F=2$, choose $R_{GS} \simeq Z_0$. Then the relation between R_{GS} and g_m specifies the current consumption:

$$I_D = \frac{g_m V_{eff}}{2} \approx \frac{V_{eff}}{10R_{GS}} \simeq \frac{V_{eff}}{10Z_0} \quad (7)$$

This estimate of current consumption is based only on consideration of noise figure. However, a practical receiver also specifies a minimum LNA dynamic range. Of concern here is the Q of the matching network, which amplifies the incident voltage. This lowers the input-referred 3rd-order intercept point of the LNA ($IIP3_{LNA}$) relative to the 3rd-order intercept point of the FET ($IIP3_{FET}$), as follows:

$$IIP3_{LNA} = IIP3_{FET} / (V_{gs} / V_{in})^2 = IIP3_{FET} \times (g_m Z_0)^2 \left(\frac{\omega_0}{\omega_T} \right)^2 \quad (8)$$

Third-order nonlinearity in a FET arises from field-dependent mobility, and $IIP3_{FET}$ in dBm is roughly proportional to V_{eff} . Therefore, LNA dynamic range (DR_{LNA}) is given by the following expression, which shows that at constant impedance match, the dynamic range is proportional to the bias current of this amplifier.

$$DR_{LNA} \propto IIP3_{LNA} / (F - 1) = \frac{g_m Z_0}{\gamma} IIP3_{FET} \propto \frac{g_m Z_0}{\gamma} V_{eff} = \frac{2IZ_0}{\gamma} \quad (9)$$

Thus, stronger reactive feedback merely *slides* the dynamic range by lowering noise figure and intercept point together. The only remaining possibility to lower current at a given dynamic range is by raising the LNA input impedance.

A common-gate amplifier with transformer feedback may be similarly configured for impedance match and low noise [21] (Fig.8). Assuming $g_m Z_0 \gg 1$, the feedback circuit sets the insertion voltage gain (s_{21}) and input impedance:

$$\begin{aligned} s_{21} &= M \\ Z_{in} &= Z_0 \text{ if } N = M^2 - M - 1 \end{aligned} \quad (10)$$

At impedance match, the noise factor of the amplifier is:

$$F = 1 + \frac{\gamma}{g_m Z_0} \quad (11)$$

The noise figure may be lowered towards 0 dB by raising $g_m Z_0$, until, like the inductively degenerated amplifier, it is limited by parasitic losses which scale adversely. However, large g_m is usually not consistent with low current consumption. Also, the termination resistor Z_0 itself must be low noise, for instance implemented at the input terminals of a transimpedance amplifier. Because it is difficult to realize transformers on-chip, this circuit topology is rarely used in integrated LNA's. Transformer coupling remains interesting because the zero DC voltage drop across inductors enables operation at low supply voltage. This is evident in the few integrated LNA's with transformer-coupling reported to date [10, 22].

Next consider a simple common-gate amplifier (Fig.9(a)), whose noise factor is $(1+\gamma)$ at impedance match [17]. FET g_m sets the input resistance, large signal handling determines the gate overdrive V_{eff} , and together these two factors specify the bias current. To circumvent this constraint the LNA may be designed with a higher than wanted input impedance ($Z_1 > Z_0$), which implies a lower bias current, and then it may be scaled to the required value (Z_0) at the amplifier input with a reactive impedance transformer (Fig.9(b)). The transformer may be a simple LC narrowband circuit [23] which needs neither coupled magnetic elements nor bias current. The relations governing this approach are as follows, where I_D , g_m , V_{eff} refer to an LNA matched to Z_0 with no transformer, and the same quantities with the prime symbol refer to an LNA whose input impedance is Z_1 instead of Z_0 . FET V_{eff} is adjusted to V'_{eff} to maintain constant LNA IIP3.

$$\begin{aligned}
I_D &= \frac{1}{2} g_m V_{eff} = \frac{1}{2} \frac{V_{eff}}{Z_0} \\
\text{Choose } V'_{eff} &= V_{eff} \sqrt{\frac{Z_1}{Z_0}} \text{ for constant } \text{IIP3}_{\text{LNA}} \\
I'_D &= \frac{1}{2} g'_m V'_{eff} = \frac{1}{2} \frac{V_{eff}}{Z_1} \sqrt{\frac{Z_1}{Z_0}} = \frac{1}{2} \frac{V_{eff}}{\sqrt{Z_0 Z_1}} \\
\Rightarrow I'_D &= I_D \sqrt{\frac{Z_0}{Z_1}}
\end{aligned} \tag{12}$$

Impedance transformation with modest scale factors is quite feasible in practice. However, when scaling impedance by a large factor, the required quality factor of the matching circuit grows quadratically and is sensitive to parasitics. Including the voltage gain in the matching network, the total LNA insertion gain to the output is:

$$\text{Gain} = \frac{Q_i g_m Q_0}{\omega_0 C} = \frac{Q_0}{50 Q_i \omega_0 C} \tag{13}$$

where Q_i is the quality factor of the LC matching network at the input port, and Q_0 of the resonant load. Compared to the case of no matching network at the input, a load with higher Q_0 is required to obtain a specified voltage gain. As the input port of the common-gate amplifier is equivalent to a two-terminal resistor of value $1/g_m$ terminating the reciprocal matching network, the LNA noise factor at impedance match remains $(1+\gamma)$ independent of the voltage gain in the matching network. γ is 0.67 for long-channel FETs, which implies an LNA NF of 2.2 dB. γ is higher at short channels [8], and in practice, a typical NF for this LNA is 3 dB. For this LNA to bias at the same current as the inductively degenerated common-source amplifier with 3 dB NF, $Z_1:Z_0=25:1$. It may be difficult in practice to realize the required impedance transformer to present Z_0 at the LNA input. Fig.10 shows a practical realization of a low current LNA tuned to 900 MHz.

So far, the LNA scaling has been considered with the intent of maintaining a constant dynamic range. However, often two different circuit blocks in a complete receiver chain may determine the upper and lower limits to dynamic range. For example, the LNA may principally determine the cascade noise figure, and the following mixer or some other downstream circuit the intercept point [7, 24]. Now it is appropriate to lower LNA

noise figure alone, without the consequently lower intercept point becoming a major concern. The common-source LNA with inductor degeneration offers a lower noise figure than the simple common-gate LNA, however accompanied by some practical limitations. A high Q matching network must precede the LNA FET, whose width is small to lower the bias current. The capacitance to ground due to the bond pad and package lead (Fig.7), which is much larger than the small FET C_{GS} , significantly alters the real part of the LNA input impedance and disturbs the intended impedance match. This may be corrected by re-tuning the matching network, but in practice the additional branch to ground due to the relatively large parasitic tends to degrade the lowest achievable noise figure.

E. Oscillators

It is well known that for lower phase noise oscillators benefit from high Q inductors. In fact, this forms an important part of conventional methodology of oscillator design [25]. However, what is not as well understood is the impact of inductor Q on current consumption, when the oscillator delivers a certain phase noise at a given frequency. A complete analysis of oscillator phase noise in some recent work [26] sheds light on this question.

Consider the differential switched current oscillator (Fig.11), which is popular in RF-ICs. There are three sources of noise in this circuit: loss in the LC resonator, noise in the switching differential pair, and noise in the tail current source. The ratio of oscillation amplitude to noise determines the oscillator phase noise sidebands. Once oscillation starts, the differential pair switches the tail current, I , into the resonator. The resultant differential amplitude is $(4/\pi)IR$, where R is the parallel loss resistance of the differential inductor. The power supply limits the largest achievable amplitude to $2V_{DD}$ differential, when the tail current source FET shuts off at each negative peak of the oscillation. Least phase noise is obtained when the current is just large enough to reach this amplitude [26, 27]. This optimum current is:

$$I_{opt} = \frac{\pi V_{DD}}{2R} = \frac{\pi V_{DD} \omega_0 C}{Q} \quad (14)$$

Q is the loaded quality factor of the resonator, usually determined by the inductor. C is the net capacitance in parallel with each inductor. On an integrated oscillator the

loading of the following stage often determines C . Clearly, therefore, the larger the inductor Q , the lower the optimum tail current.

However, the question remains how an oscillator with a lower optimum current compares with another oscillator with a higher optimum current, as measured by some meaningful figure-of-merit? One such figure-of-merit is the phase noise $\mathcal{L}(f_m)$ at offset frequency f_m normalized to oscillation frequency f_0 and to power consumption [28].

$$FOM = \left(\frac{f_0}{f_m} \right)^2 \frac{1}{\mathcal{L}(f_m) V_{DD} I} \quad (15)$$

Analysis of the differential oscillator [26] shows that

$$\mathcal{L}(f_m) = \frac{2FkT}{\left(\frac{1}{2}V_0^2 / R\right) \left(2Qf_m\right)^2} \quad (16)$$

V_0 is the amplitude and F is the noise factor representing, respectively, the contributions of resonator loss, differential pair switches and current source to phase noise, as follows:

$$F = 2 + \frac{8\gamma RI}{\pi V_0} + \gamma \frac{8}{9} Rg_{mbias} \quad (17)$$

At the optimum bias current, I_{opt} , the figure of merit is then:

$$FOM = \frac{4}{\pi} \frac{Q^2}{kT} \frac{1}{2 + 4\gamma + \frac{32}{9} \gamma \pi \frac{V_{DD}}{V_{eff,bias}}} \quad (18)$$

The FOM improves as Q^2 , but depends on little else that is circuit-specific other than V_{eff} in the 3rd term in the denominator. This means that for high- Q resonators, both the optimum current and the normalized phase noise go down together. This is a doubly favorable outcome. In practice, the resonator Q must include varactor loss, and the tuning range should encompass both the frequency band of interest and spreads in parasitics.

If thermal noise in the differential pair and in the current source dominates F , and the tail current drives the differential oscillation amplitude to its largest possible value, $2V_{DD}$, limited by the power supply, then:

$$\mathcal{L}(f_m) = kT \frac{\gamma I}{\pi V_{DD}^2} (\omega_0 L)^2 \left(\frac{1}{V_{DD}} + \frac{2}{V_{eff}} \right) \left(\frac{f_0}{f_m} \right)^2 \quad (19)$$

V_{eff} is the effective gate voltage biasing the current source FET. As $V_{eff} \ll 2V_{DD}$, noise in the current source dominates phase noise. Raising V_{eff} lowers current source compliance, causing the highest oscillation amplitude to limit at some value lower than $2V_{DD}$. This worsens phase noise. The equation also shows that at a fixed tail current, I , forces the oscillation amplitude into the V_{DD} -limited regime, the phase noise is independent of inductor Q . In fact in this regime as the tail current is made larger the phase noise worsens (Fig.12). The lowest current to produce the V_{DD} -limited amplitude, $I_{opt} = \pi V_{DD} / 2R = \pi V_{DD} / (2Q\omega_0 L)$, produces the least phase noise \mathcal{L}_{min} :

$$\mathcal{L}_{min}(f_m) = kT \frac{\gamma}{V_{DD}} \frac{\omega_0 L}{2Q} \left(\frac{1}{V_{DD}} + \frac{2}{V_{eff}} \right) \left(\frac{f_0}{f_m} \right)^2 \quad (20)$$

Furthermore, at a given oscillation frequency, the *lower* the (L/Q) ratio of the resonator's inductor, the lower the \mathcal{L}_{min} . On the other hand, to lower power with smaller I_{opt} the $(L \times Q)$ product should be *greatest*. In discrete inductors Q is almost independent of L at frequencies in the range of 1 GHz (Fig. 5). Thus, L should be low for better phase noise \mathcal{L}_{min} but high for low current I_{crit} . For a given Q , phase noise scales down with resonator inductance, that is, inversely with resonator capacitance, but at the optimum bias current is higher.

Two practical realizations of low current oscillators at 800 MHz and 1 GHz are shown, respectively, in Figs. 13(a) and 13(b). The switched capacitor tuning method [29] in Fig. 13(b) is a means to cover process variations and resonator spreads with digital control, without using a high swing varactor, which degrades phase noise.

F. Mixers

The passive commutating switch mixer, consisting of four FETs used as analog switches (Fig. 14), needs no bias current, which makes it the ultimate low power circuit.

However, this mixer is non-unilateral, that is, signals flow bi-directionally from input to output. This poses several practical problems. For example, if this mixer is directly connected to the inductor output node of an LNA, the subsequent amplifier stage may substantially load the inductor. Similarly, two mixers downconverting a common input

by quadrature LO phases would tend to load each other's outputs during the overlap period of LO-I and LO-Q when switches turn ON in both mixers. For these reasons unilateral active mixers are preferred in a receiver.

The most straightforward way to lower power dissipation of an active mixer (Fig.15) is to scale down FET size. The large signal handling remains unchanged if V_{eff} is held constant for the input transconductor stage, but the mixer noise goes up inversely with bias current. The following expressions gives the total input-referred DSB white noise spectral density of a double-balanced MOS active mixer [30], and compares it with the input IP3 of the mixer:

$$\begin{aligned} \hat{v}_{on}^2 &= 8kTR \left(1 + \frac{2\gamma IR}{\pi A_{LO}} + \gamma g_m R \right) \\ \text{SSB } \hat{v}_{in}^2 &= \frac{\hat{v}_{on}^2}{\left(\frac{2}{\pi} g_m R\right)^2} \simeq 2\pi^2 kT \frac{\gamma}{g_m} = \pi^2 kT \gamma \frac{V_{eff}}{I} \\ \text{IIP3} &\sim V_{eff} \end{aligned} \quad (21)$$

γ is the FET noise factor and the amplitude of a sinewave LO, A_{LO} , is assumed large enough to completely commutate the mixer switches. The three terms in the expressions for \hat{v}_{on}^2 give the output noise voltage due to the load resistors, the mixer differential switches, and the input transconductor stage. As the noise due to the input transconductor stage dominates in most cases, the expression for input-referred noise may be simplified. When compared to the expression for IIP3, this shows that increasing V_{eff} merely slides the mixer dynamic range, that is noise and IP3 go up together. A larger dynamic range requires more bias current. In low-power receivers, the mixer is often the bottleneck to overall IP3.

The LNA insertion gain must be raised to overcome large mixer noise, which once again argues for the use of high impedance, that is high Q, inductors in the LNA load. A differential LNA followed by a doubly balanced mixer rejects common-mode pickup of stray signals and disturbances. On the other hand, a single-ended LNA takes half the power of a differential LNA and it is simple to connect to the antenna. A single-ended LNA output readily drives a single-balanced active mixer, which then produces a differential downconverted output. The subsequent baseband stages may be fully differential.

The main problem in using a single-balanced mixer is that the large LO feedthrough at its output may saturate the following stage. Here, too, an inductor load tuned to the intermediate frequency at the mixer output can attenuate the LO frequency. However, the higher the IF, the sharper the required transition band of the LC filter load to attenuate the LO feedthrough.

It has been said before that the direct conversion architecture is of great interest in low power implementations. However, flicker noise in active CMOS mixers poses a threat to the achievable sensitivity in this receiver, particularly for narrowband channels. This is because the switch FETs (Fig.16) contribute flicker noise to the mixer output without frequency translation [30]:

$$v_{n,o} = \frac{4IR}{ST_{LO}} v_{n,sw}, \text{ where } \hat{v}_{n,sw}^2 \propto 1/(WL) \quad (22)$$

$v_{n,sw}$ is the input-referred flicker noise voltage of the switch FET's, IR is the maximum voltage across the mixer's load resistor, S is the slope of the LO voltage at differential zero crossing, and T_{LO} is the LO period. So as not to load the LO buffers and yet to switch quickly, the switch FETs must be of small width and of minimum channel length. As a result, $v_{n,sw}$ is large, and this produces high flicker noise $v_{n,o}$ at the mixer output. If the LNA amplifies the RF signal of interest by only a modest gain and the mixer downconverts it to zero IF, flicker noise may easily overwhelm the signal at the mixer output. Clearly, the narrower the signal bandwidth, the closer the valuable signal energy lies to DC, and the worse the SNR.

A straightforward way to lower noise is by quadratic scaling up of the switch FET area. This, however, translates into a quadratic increase in power consumption in the LO buffer driving the switches. A more power-efficient way is to replace "direct" by "indirect" conversion to zero IF (Fig.17). The RF input is first converted by a high frequency LO (small T_{LO}) to some IF, and then by a lower frequency 2nd LO (large T_{LO}) to zero IF. Suppose CMOS inverters buffer the LO buffers to drive the mixers. The inverter outputs will slew with a maximum slope S , which is the same in the expression above for both 1st and 2nd mixers. With all else the same, flicker noise at the output of the 2nd mixer is then lower in proportion to its LO frequency. After the first downconversion, the signal lies at some IF which is easily chosen far away from the

frequency band where the flicker noise is large at the first mixer output. Compared to direct conversion, this scheme requires one additional mixer and an LO buffer. Both LO frequencies can be derived from a single VCO by tapping off the prescaler in the PLL. The choice of frequencies is flexible, as no fixed filters are involved. The first IF must be high enough so that the RF preselect filter and the tuned LNA load satisfactorily suppress the image signal. However, it must not be too high, otherwise the LO frequency in the 2nd mixer must be higher, which would lead to more flicker noise at its output.

G. Frequency Dividers

The frequency divider can easily be the most power-hungry digital circuit in the radio portion of the receiver. It tunes the LO frequency to a particular channel. In the worst case, it comprises a variable modulus digital divider clocked at the LO frequency. However, usually the LO frequency is first divided by a fixed small integer (often $\div 2$ or $\div 4$) and then by the variable modulus. At high frequencies, or when low power dissipation is important, the prescaler is implemented as a bipolar ECL circuit. A high frequency, low power CMOS prescaler is of great interest.

High frequency, low power GaAs FET prescalers are designed with Source Coupled FET Logic (SCFL), a current-steering circuit modeled on an ECL divider (Fig.18). The SCFL is extendable to CMOS. The internal logic voltage swings must be large enough to switch a CMOS differential pair. A 1V single-ended voltage swing is very typical. If the swing is smaller, the size of the differential pair FETs must be scaled up to rapidly switch the tail current, but this increases the input capacitance. On the other hand, too large a swing only lengthens the rise and falltime.

Technology scaling directly improves the performance of logic circuits in general, specifically of frequency dividers. As linewidth scales down, so does parasitic capacitance of source and drain junctions and wiring which usually dominate the load on the divider stages.

In a 0.25- μm CMOS implementation, the 1V swing can be set by, say, a 100 μA current switching into a 10k Ω load. The load capacitance limits the lowest current. With only a 1.5-V supply, this swing drives the conducting device in the differential pair deep into

the triode region. Unlike a bipolar transistor that should not be driven into saturation, it is acceptable to force a CMOS switch into the triode region. The bias current may be scaled down in subsequent stages of the frequency divider, which toggle at lower frequencies.

H. Baseband Circuits

Variable gain amplifiers and active filters are the important baseband blocks in a direct conversion or low-IF receiver. As the input signal and accompanying interferers in the circuits up to and including channel-select filter have been previously amplified, large signal handling in these circuits is of greater concern than noise. The dynamic range should slide up relative to the front-end. To illustrate this point, an active lowpass filter for channel selection in a zero-IF receiver is considered.

Past experience [7] shows that in well-designed integrated receivers, noise in the active filter dominates the receive signal chain. The input noise level of a multi-stage filter v_{in} can only be lowered at the expense of a quadratic increase in current consumption and capacitor area. This is because the noise spectral density in the passband of, say, a g_m - C filter depends on the g_m , and to implement a fixed pole frequency, the filter capacitance is proportional to g_m . As a result, the input-referred noise in the passband and the filter capacitors are related as follows [6]:

$$\begin{aligned} \hat{v}_{in}^2 &= 4kT\Gamma Q^2 / g_m \propto Q^2 V_{eff} / 2I \\ C &= g_m / \omega_p \propto Q^2 / (\omega_p \hat{v}_{in}^2) \end{aligned} \quad (23)$$

Γ is some noise factor associated with the transconductor circuit, and Q is the pole quality factor. A power-efficient receiver will use an active channel-select filter with the lowest Q poles, that is, whose passband is centered at the lowest frequency. The filter should therefore select the desired channel at zero or low IF. Rather than to lower noise by scaling the filter, it is usually more power-efficient to amplify the input signal before it arrives at the filter. However, now the filter must handle amplified large interferers at its input, which although they lie in the stopband and are eventually suppressed at the filter output, may induce intermodulation distortion that falls in the filter passband. Filter linearity is therefore paramount.

Op amp-based active filters are usually more linear than open-loop transconductor-based active filters. An op amp-based filter is as linear as the passive feedback components such as resistors and capacitors, and in a well-designed op amp the gain compresses at a maximum output voltage swing almost equal to the power supply. There are two well-known types of op amp-based filter: switched capacitor, or active RC . The switched capacitor circuit is ruled out because the filter in a wireless receiver must handle out-of-band signals lying far away from the narrow passband needed for a single channel. Therefore to simplify the anti-alias filter, the switched capacitor circuit must be clocked at a large oversampling factor, otherwise sampling at the filter input may alias the out-of-band signals into the filter passband. The high clock rate proportionally raises the current consumption.

The continuous-time active RC filter (Fig.19) is interesting because with op amps biased at low currents in weak inversion, it is possible to accurately produce the desired passband, transition band, and close-in stopband characteristics at low or zero IF. In most cases, the circuit's inability to respond at high frequencies ensures large loss in the high frequency stopband, without requiring high loop gain in the op amp circuits. A two-stage op amp drives the filter resistors (Fig.20). The output stage must be capable of rail-to-rail voltage swing. A dominant pole and a zero compensate the fully differential op amp for stable operation in feedback. The pole and zero frequencies of the active RC filter must be tunable to overcome process spreads. The lowest power method to do so is with a binary-weighted array of switched capacitors at every filter node [31]. The number of array elements depends on the desired accuracy of filter poles. For example, a 5b switch-selectable capacitor array in parallel with a fixed capacitor equal to the largest array capacitor encompasses variation of up to $\pm 50\%$ in dielectric thickness, while tuning the filter frequency response to 3% accuracy. This is good enough in most practical cases.

Low frequency flicker ($1/f$) noise is of particular concern in a CMOS zero-IF receiver. The op amp circuit must be designed to lower the input-referred flicker noise. The flicker noise in a PMOS input stage is usually lower than in an NMOS stage, an effect ascribed to buried-channel conduction in PMOS. Furthermore, as input-referred flicker noise voltage is inversely proportional to gate area $W \times L$, the width and length of all

FETs in the op amp is scaled up while keeping W/L constant. The increase in L lowers f_T roughly in inverse proportion. The noise corner frequency, where flicker noise density intercepts white noise, is proportional to f_T (Fig.21). In baseband circuits such as the filter, lower f_T is not of much concern until it approaches roughly ten times the highest signal frequency applied to the circuit.

The final op amp sizes must be scaled to deliver some large DC gain when driving the filter resistors. These resistors determine the noise spectral density in the filter passband, that is, the filter's noise figure. A powerful way to improve dynamic range is by embedding gain into filter stages. The largest dynamic range is usually obtained by uniformly interleaving gain with filtering. When the bottleneck is filter noise but not intercept point, it is best to concentrate gain into the filter's input stage, which then significantly lowers the noise contributions of the subsequent filter stages, particularly of stages with high-Q poles where noise is boosted near the pole frequency.

Another important baseband circuit is the A/D converter and demodulator. Using sigma-delta techniques, it is relatively easy to implement a high-resolution A/D converter in the baseband section of a zero IF receiver. However, for wideband FSK demodulation it is simpler still, because a low-power limiting amplifier [32], which is effectively a 1b A/D converter, may be followed by an efficient but very low power detector [33]. This approach can be extended to constant envelope modulations with lower index, such as GFSK.

I. On-Chip Inductors

The foregoing sections point out the vital role of large value and high quality inductors in lowering power dissipation. It has also been said that these inductors must usually be off-chip. However, there is increasing need to integrate all the elements in a receiver, among other reasons to obtain a physically compact form-factor. On-chip inductors that are substantially better than the ones customary today will come about from advances on two fronts: CMOS processes which address the needs for better inductors; and efficient software to design spiral inductors, which accurately captures the various losses. Recently, there has been progress on both fronts.

Some BiCMOS and also pure CMOS processes now offer lightly doped substrates. Traditionally CMOS substrates were heavily doped ($\sim 10 \text{ m}\Omega\text{-cm}$) to avoid latchup, and the active devices were fabricated in a lightly doped epitaxial layer (usually p-epi on p+ bulk). However, today the trend is to eliminate the epitaxial layer and to use a lightly doped bulk ($\sim 15 \text{ }\Omega\text{-cm}$) accompanied by dual or triple wells. Also, it is relatively straightforward to add a thicker than usual film of interconnect metal at the uppermost layer of a system of multi-level interconnect, where planarization is less of a concern. The thicker metal cuts down on the usually dominant series ohmic loss in an on-chip spiral inductor.

Spiral inductors fabricated on lossy substrates are poorly modeled. In addition to ohmic losses, they induce displacement currents in the substrate, and eddy currents too as the inductor flux substantially penetrates the substrate in planes parallel to the inductor (Fig. 22). The latter are distributed effects, and difficult to model analytically. General-purpose software to solve Maxwell's equations is usually too slow to be useful. To this end, a simulator customized to spirals has been developed to model self-inductance, capacitance, and all losses including skin effect [34].

With the right features in the technology and the appropriate CAD tools, it becomes possible to obtain better inductors than the usual. One example is shown here, an 80 nH spiral inductor with a Q of 4.3, with a self-resonance frequency of about 1 GHz (Fig. 23). The inductor consists of four identical spirals stacked in series to obtain a solenoid-like coupling, whose total self-inductance grows quadratically with the number of layers. The high self-resonance is due to the much smaller footprint of the structure compared to a single-layer spiral, and the consequently lower capacitance to substrate. This type of high value integrated inductor is frequently required in low power circuits.

J. Examples of Low Power Radio Implementations

Various CMOS building blocks and a complete receiver embodying the principles described above have been reported [24]. For instance, an indirect conversion receiver tuned to 900 MHz takes a total of 2.1 mA from 1.5 V, at least a factor of ten lower current than what is reported in other well-designed CMOS receivers. Eventually the receiver drained 3 mA because the on-chip coupling capacitors were much smaller than

expected. The details and performance summary is available in the original publication [24].

What is interesting for the purposes of this paper is that this receiver gives experimental verification of the principles outlined above. The RF front-end of this receiver was built in different CMOS technologies spanning two generations of feature size scaling: $0.6\mu\text{m}$, and $0.25\mu\text{m}$. The front-end consists of the 900 MHz LNA, a first mixer, and an 830 MHz VCO. Both implementations use the same external inductors. As Table 1 shows, the performance of both circuits is almost equal, and they dissipate essentially the same power. This proves our earlier assertion that once dynamic range is fixed, shrinking technology does not lower the power dissipated by the analog circuits in the RF front-end. What matters more is the quality of the passive elements.

K. Conclusions: Circuits

In the first part, this paper has surveyed some of the circuit design techniques and system methodologies underlying ultralow power wireless receivers implemented in bipolar and CMOS IC's. High quality passive components are shown to be very important in lowering power dissipation of RF IC's, a fact that was either not realized or set aside in the rush to fully integrate wireless transceivers on a single chip. Dissipation-less passive components can also transform high impedance levels associated with low power circuits to lower off-chip characteristic impedance.

To achieve a certain target for cascade noise figure, a low power receive signal chain must carefully choose between inserting gain prior to the most noisy building blocks or scaling up power consumption to lower their noise level. This tradeoff is illustrated with specific examples.

Baseband and IF building blocks are often very challenging, because to realize the necessary dynamic range they may consume as much power as the RF sections. If the modulation scheme allows, a low or zero IF is most desirable from the point of view of lowering power consumption. The dynamic range of blocks up to, and including, the channel-select filter is usually comparable to the receiver front-end dynamic range, except that as the signal is amplified, the dynamic range must slide up. For narrowband channels, circuits in weak inversion are well suited to implement these blocks. If the

circuits can swing rail-to-rail at the output without significant gain compression, the input signal can tolerate a higher noise level, which implies that circuit power may be scaled down. An example is given of a baseband amplifier and active RC filter.

One important finding is that technology scaling brings limited benefits in lowering the power consumption of the receiver. To the first order the required dynamic range sets the current consumption of RF and IF circuit blocks, and improving the f_T of the transistors brings no major benefit. However, technology scaling can significantly lower the power that overdriven binary circuits such as frequency dividers consume at a certain clock frequency.

IV. Power Conscious System Design

The variability of wireless channels in terms of both the propagation conditions between intended partners and the interference due to other users of a shared band must be considered in the design of reliable low-power radios. Flexibility in terms of transmitter power levels, antenna beam patterns, equalizer capabilities, symbol rate, and constellation size can lead to large improvements in the ability of many users to share the available time/frequency radio resource, and vastly improved robustness to channel variations. Network topology also has a considerable impact on capacity and reliability. Given the large performance improvements that are available with the channel and interference mitigation techniques, a trade in favor of increased signal processing at baseband can allow for a dramatic relaxation of design specifications for the RF sections, for example, a combination of increased noise figure or reduced transmitter power. After a brief description of the fundamental constraints on network capacity, and principles that can be used to achieve a large fraction of that capacity, we present low-power wireless sensor networks as an application.

A. Dealing with Variability in Radio Channels

The throughput in radio communications is fundamentally limited by constraints on power, bandwidth, and latency, and by channel losses, noise, and the inability to estimate channel and interference variations. Maxwell's laws and Shannon's capacity theorem together dictate a maximum bit rate, given a certain bandwidth, noise, power,

and range. Limits such as the Cramer-Rao bound [35] similarly apply to how well channel conditions can be estimated. Variability in the instantaneous SNR due to either the channel (e.g. multipath, shadowing, distance loss) or interference (e.g. bursty and/or fading traffic) can limit a simple-minded system to the worst-case conditions. This can represent many orders of magnitude of loss in power efficiency compared to a Gaussian channel. Achieving performance similar to the Gaussian channel requires an ability to spread the information, with the gap closing as better estimates of the channel are utilized.

The most important parameter for quality of service of a radio communications link is the signal to interference ratio (SIR), where we include background noise as interference. Variations in both the signal and interference strength affect the SIR, and consequently methods designed to deal with variable signal strength such as time, frequency, and space diversity [36] also apply to some extent to variations in interference levels. Indeed, were the interference independent of the actions taken to mitigate its variability, exactly the same techniques would suffice. However, the interference level may be dependent, as for example may occur if transmitter power is adjusted to try to maintain a fixed SIR. One user raises its power; the other responds by raising its own, leading to a further response [37-40]. Quite different estimation techniques are needed for independent versus dependent interference; probing is one effective means [41, 42].

Two basic principles can be invoked for dealing with variable interference in a multiple access setting: *interference averaging*, and *interference avoidance*. When due either to the high mobility or rapid change in the interference (e.g. short packets) we cannot form good estimates of the SIR, then rather than accept the worst case SIR conditions, we spread the information over time, frequency or space so that the aggregate channel looks more like the average SIR condition. This is the principle at work in spread spectrum communications, channel coding with interleaving, ARQ schemes, equal gain antenna combining, and a large number of other techniques. When better estimates can be formed at the receiver, coherent combining techniques may be used to increase the diversity benefits. This further permits explicit interference cancellation techniques, such as multi-user detection [43-45], and improved channel coding methods such as space-time codes [46]. When the change is slow enough for the transmitters and

receivers to cooperate in a feedback fashion, then interference avoidance is possible-- information is sent on the channels or in the antenna beam directions that both leads to low interference to other users and at the receiver of the communicating pair. Methods include dynamic channel and power allocation, variable bit allocation [47, 48], and adaptive transmit and receive antenna arrays [42, 49].

The severe effects of multipath and interference must be effectively dealt with to ensure robust, reliable, low-power design. A combination of diversity techniques and feedback control of variable bit and power allocation can lead to vast improvements. In practice, constraints on the delay, dynamic power range, and complexity limit the improvements that can actually be obtained. In addition, the rate of change of the channel imposes fundamental limits on the quality of the channel state estimator. For example, in mobile communications it is difficult to form accurate estimates of the instantaneous SIR. Yet compensation for multipath in the receiver allows the transmitter to perform reliable allocation using an average channel state, i.e., the residual channel after diversity combining, which has smaller variations than the original channel. A mix of feedback and diversity techniques can lead to a lower overall complexity or consistently higher quality of service than reliance on only one approach. The payoffs for adaptive techniques are very high in radio systems, and as these techniques are implemented at baseband rather than RF, result in a favorable chip area tradeoff compared to simply using higher power RF front ends.

B. Networking and Application Issues

Mitigation of multipath and interference can both be dealt with up to the link layer. Shadowing and distance losses by contrast are dealt with in the provisioning of the network infrastructure, and in the selection of the network topology, at least for low-power radios in which high transmit power is not available to overcome distance losses. Thus for example the base stations may be sited in a cellular system to deal with gross shadowing conditions (e.g., by raising their elevation), and the power constraints of the mobile units (e.g. by placing them closer together). A wired backbone then connects the base stations and avoids the need for high power RF units. A commercial example where a great deal of attention has been placed on low-energy design is radio paging. Protocols are used to enable power-hungry circuits to be active for as short a time as

possible [1]. Thus, time division is employed so that radio receivers can be off for a large fraction of the time. Asymmetry between the resources in the base station and mobile units is also exploited. In the FLEX and POCSAG protocols, there are beacons that are broadcast at much higher duty cycles than those of the receivers, to enable power savings in the pagers. The protocol and hardware designs are harmonized so that in so far as possible no bits are wasted, and processing takes place in an efficient fashion. Knowledge of the type of traffic being carried is essential to this undertaking. Thus, the design is optimized from physical layer through to application. A number of these principles carry over to a more far-reaching application, which we now detail.

It has been observed that processors are embedded virtually everywhere, but seldom can communicate with each other. However, with the recent advances in the technology of low-power radios, potentially the physical world can be networked, leading to new applications in medicine, security, factory automation, machine maintenance, environmental monitoring, and many other areas. This general topic goes under a number of names, including pervasive computing, invisible computing, smart spaces, and wireless integrated network sensors. The same advances in semiconductor technology that enable the fabrication of micro power RF components in silicon and high speed CMOS circuits also permit the construction of complete micro electromechanical systems (MEMS) on a chip. This further allows the bundling of communications, signal processing, sensing, and actuation capability in very small form factors, with mass production techniques. However, low cost is not simply a matter of making the silicon cheap; the installation and energy system costs must also be low. This implies a need for both self-organization, and attention to energy efficiency at all levels of the design. To focus the discussion, in the remainder of this article we consider the example of wireless sensor networks, in which the sensor nodes must be compact and parsimonious in their use of energy, while alerting outside agencies when priority events are observed.

Clearly long-range transmission will not be accomplished in a single hop by such devices. Rather, by multi-hopping the messages through a sequence of nodes, only short-range transmission is required. This approach both allows dense spatial re-use of frequencies and accommodates transmitter power constraints. The power advantage can

be very large. For example, assuming fourth power distance propagation losses, using N equal hops requires a factor of N^3 less transmission energy in total than sending the same message in one long hop. Further, with a dense network, it is possible to route around obstacles that may produce shadowing. Thus, such networks may ameliorate both shadow fading and distance losses. Additionally, the short transmission range implies that equalization will not be required, further simplifying the radios. There is of course a limit in how dense the network should be, since the up and down conversion costs in the transceiver will eventually dominate the total communications energy cost.

Sensor networks with thousands of nodes can generate far more data than can be handled by any end user; it must be reduced at some point, or information overload results. Given the large energy and bandwidth cost of actually transporting the raw data, it is far better to do most of this data reduction at source, with further trimming as it passes through the network. Long-term trends in feature size indicate that the energy cost of doing this processing in the remote nodes will decline over time, while we have earlier observed that there are fundamental lower limits to the energy cost of transport of bits over a given distance. It is of course not true that processing at source is free, and thus there are a number of interesting architectural choices for accomplishing the desired task. For example, in a security application, the goal might be to detect and identify possible threats in the field of view. Very simple processing based upon energy thresholds could suffice for much of the time, but will lead to unacceptable false alarm rates. Further processing is performed when the energy threshold is passed, possibly through multiple layers. At some point, a cooperative decision among a group of sensors that are close to the target may be required, necessitating the high energy cost step of transmitting raw data locally. Finally, a decision can be multi-hopped back to either a gateway or the end user. Along the way, the decision may be aggregated with the decisions of other sensors, so that the end information is a summary report of activity in some area. In very rare circumstances, the end user may request the original raw data, e.g., to update the detection and classification algorithms. This hierarchy with feedback enables a lowering of the average energy and bandwidth usage, while preserving the ability to apply high-level processing when absolutely needed. Some of this processing may be performed on

special purpose devices (e.g., for frequently invoked operations, or those requiring real-time response), while others may be run on general-purpose machines.

In both the communications and signal processing tasks, a hierarchy that is in part physical and in part constructed by means of protocols enables a better matching of resource usage against the probabilities of events. Most of the time, there is nothing to report, but yet constant vigilance is required. However, to meet reliability targets, sophisticated processing together with large data transfers may be needed on occasion, albeit with allowances for latency. A hardware separation of real-time and simple functions from non-real time but sophisticated functions offers the advantage that the real-time circuits can be made very efficient, and the sophisticated functions can be performed in software for maximal flexibility. The split of functions intimately depends upon the application, and further while layering is performed in both hardware and software, for efficiency there must be access to the physical layer all the way through to the application.

Low bit rate communication together with short transmission ranges and compact form factors limits the range of diversity techniques available. If the nodes are stationary, time diversity is not a reliable option, while the small size rules out antenna diversity. Thus, coding in combination with frequency spreading is the only plausible means for realizing diversity for peer-to-peer communications.

For very short hops the transceiver power consumption for reception is nearly equal to that of transmission, and thus there are limits in how short the hops should be made. Further, this implies that the protocol should be designed so that radios are off as much of the time as possible, that is, the MAC should include some variant of TDMA. This requires that the radios periodically exchange short messages to maintain local synchronism. It is not necessary for all nodes to have the same global clock, but the local variations from link to link should be small to minimize the guard times between slots, and enable cooperative signal processing functions such as fusion and beam forming. The messages can combine health-keeping information, maintenance of synchronization, and reservation requests for bandwidth for longer packets. The abundant bandwidth that results from the spatial re-use of frequencies and local

processing ensures that relatively few conflicts will result in these requests, and so simple mechanisms can be used.

If batteries power the nodes, the network will have a life cycle that begins in a boot-up, proceeds through a phase of maximum functionality, decline, and finally failure. Every bit that is exchanged hastens the end of the network. Particular nodes may be more heavily stressed by traffic than others (e.g., those in the vicinity of a gateway or other long-range link). Thus, routing protocols must to some extent be energy-aware, to sustain useful operation as long as possible. The minimum energy path is not necessarily the most desirable; rather we must ordinarily route to extend operation, although high priority messages may be routed for low latency, even if this exhausts precious network resources. The predictability of flow to and from a relatively small number of gateways enables infrequent construction of sets of paths to these data consumers, minimizing overhead.

These considerations are behind a protocol suite [50] developed in a joint effort of UCLA and the Rockwell Science Center. A TDMA-like schedule is gradually built up, with nodes initially having no awareness of their neighbors, and with no synchronism. To conserve energy, synchronism and channel assignments take place in the same procedure. Within a superframe, some fraction of the frames is devoted to invitation frames, and others to synchronized communication. At the beginning, most of the frames are devoted to the invitation and response cycle. As a node forms more connections, a decreasing fraction of the frames are devoted to adding new connections. Node to node, node to network, and network-to-network connections take place, the latter requiring adjustments of timing and channel assignments in the neighborhood of the connection. Since the algorithm is distributed, worst-case local conditions, rather than the absolute number of nodes in the network govern the time until all nodes are connected. Simulations reveal a very short time to achieving a conflict-free connected network, which is independent of the number of nodes for network sizes above 200. The speed is aided by the availability of many more channels (frequency/time/code slots) than are needed, which reduces conflicts. The super frames have a sparse allocation of the time/bandwidth resource, with most available for reservations to accommodate local exchanges of raw data for purposes such as beam forming.

A distributed energy-conserving routing protocol has also been developed, for the purpose of extending network life [50]. Overlapping spanning trees are constructed outward from gateways, by messages that note how many hops outward they have moved. This establishes which tier the nodes belong to; the tree is then constructed by allowing only connections within a tier or down to a lower level. By this means, nodes have multiple routes to get packets to the gateway, which are chosen based on a metric that reflects a combination of the packet priority, and the energy reserves of nodes along the path. A significant issue is how to update the routes, given that nodes can catastrophically fail, and in any case gradually lose energy. The updating consumes energy, but failing to update in a timely fashion can also drain the reserves of critical nodes; simulations showed that the split traffic algorithm using spanning has a fairly broad maximum in network duration in terms of update rates. Centralized updating is extremely costly in energy, and so it was found to be better to rebuild around failures in the local neighborhood.

These protocols indicate the feasibility of achieving distributed low-power operation in a flat multi-hop network. It is clear however that for a wide range of applications means must be found to conveniently link sensor networks to the Internet, so that standard tools such as browsers can be used for access and control, and resources such as networked databases can be made available. Inevitably then some layering of the protocols (and devices) will be needed to make use of these standard interfaces, provided these software and hardware interfaces make available to upper layers control over the functions down to the physical level (e.g., radio transmit power, energy availability, etc.). An example of such a device is reported in [51].

Hierarchy in communications resources can also be exploited to reduce latency or lessen the power burden for the more numerous small devices. For example, provision of higher power long-range links reduces the number of hops required to transmit data over a given distance. It can also improve scalability, since higher levels can aggregate large volumes of traffic without the need to process it down to ever more compact representations. A cluster architecture where the clusterhead has multiple antennas and increased signal processing capability can lead to reduced transmission power requirements by the remote nodes, using adaptive array techniques. Thus, the

proliferation of low power radios in the environment may actually imply the need for yet higher power and sophisticated radios as well.

A further issue that arises for low power radios is how they can coexist with their higher power cousins who may be sharing the same transmission bands (e.g., ISM bands). Many protocols assume that all other radios in interference range use the same protocol. This is clearly a very poor assumption, given the heterogeneity in the devices now on the market. Sensor networks fortunately have lower quality of service requirements in terms of latency and throughput than the higher power applications, and transmit over a short range. Frequency diversity together with coding provides also some degree of interference diversity, which can be extended with retransmissions. In effect, the methods that are effective against intentional jammers can be applied to improve the robustness of these systems. Yet at some point, a dense collection of high power users may crowd out low power users. Fixed or dynamic channel planning based on application classes may be needed.

C. Conclusions: Systems

The variability of radio channels poses many challenges to the design of a system that includes low-power radios. However, a wide variety of techniques exist which can provide trades among digital signal processing complexity, latency, transmitter power, noise figure, and bandwidth. Low-cost low-power RF design is enabling for a potentially vast range of new applications involving densely distributed networks, which can in turn be designed in such a way as to simplify the demands placed upon the radio. This technological synergy will have profound consequences for years to come.

V. References

- [1] W. Mangione-Smith, "Low power communications protocols: paging and beyond," in *Symp. on Low Power Electronics*, San Jose, CA, pp. 8-11, 1995.
- [2] L. I. Williams, "System integration of the Flex paging protocol. I. System design constraints," *Mobile Radio Technology*, vol. 14, no. 6, pp. 10, 12, 14, 16, 18, 20, 22, 24, 26, 1996.
- [3] L. I. Williams, "System integration of the Flex paging protocol. Part 2: System design recommendations," *Mobile Radio Technology*, vol. 14, no. 7, pp. 12, 14, 16, 18, 1996.
- [4] D. L. Ash, "SAW-Based Hybrid Transceivers in SLAM Packaging with Frequency Range from 200 to 1000 MHz," in *IEEE Ultrasonics Symp.*, Sendai, Japan, pp. 389-398, 1998.

- [5] A. A. Abidi, "Direct-Conversion Radio Transceivers for Digital Communication," *IEEE J. of Solid-State Circuits*, vol. 30, no. 12, pp. 1399-1410, 1995.
- [6] A. A. Abidi, "Noise in Active Resonators and the Available Dynamic Range," *IEEE Trans. on Circuits and Systems*, vol. 39, no. 4, pp. 296-299, 1992.
- [7] A. Rofougaran, G. Chang, J. J. Rael, J. Y.-C. Chang, M. Rofougaran, P. J. Chang, M. Djafari, J. Min, E. Roth, A. A. Abidi, and H. Samueli, "A Single-Chip 900 MHz Spread-Spectrum Wireless Transceiver in 1- μ m CMOS (Part II: Receiver Design)," *IEEE J. of Solid-State Circuits*, vol. 33, no. 4, pp. 535-547, 1998.
- [8] A. A. Abidi, "High Frequency Noise Measurements on FETs with Small Dimensions," *IEEE Trans. on Electron Devices*, vol. ED-33, pp. 1801-1805, 1986.
- [9] Q. Huang, F. Piazza, P. Orsatti, and T. Ohguro, "The Impact of Scaling Down to Deep Submicron on CMOS RF Circuits," *IEEE J. of Solid-State Circuits*, vol. 33, no. 7, pp. 1023-1036, 1998.
- [10] J. R. Long, R. A. Hadaway, and D. L. Harame, "A 5.1-5.8GHz Low-Power Image-Reject Downconverter in SiGe Technology," in *Bipolar Circuits & Technology Mtg.*, Minneapolis, pp. 67-70, 1999.
- [11] T.-H. Lin, H. Sanchez, R. Rofougaran, and W. J. Kaiser, "CMOS Front End Components for Micropower RF Wireless Systems," in *Int'l Symp on Low Power Electronics and Design*, Monterey, CA, pp. 11-15, 1998.
- [12] A. K. Agrawal, R. D. Clark, J. J. Komiak, and R. Browne, "Microwave module interconnection and packaging using multilayer thin film/thick film technology," in *Intl. Microwave Symp.*, Albuquerque, NM, pp. 1509-11, 1992.
- [13] L. Zu, L. Yicheng, R. C. Frye, M. Y. Lau, S. C. S. Chen, D. P. Kossives, L. Jenshan, and K. L. Tai, "High Q-factor inductors integrated on MCM Si substrates," *IEEE Transactions on Components, Packaging and Manufacturing Technology, Part B: Advanced Packaging*, vol. 19, no. 3, pp. 635-43, 1996.
- [14] B. Breen, "Multi-layer inductor for high frequency applications," in *Electronic Components and Technology Conf.*, Atlanta, GA, pp. 551-554, 1991.
- [15] M. Sakakura and S. Skiest, "Ultra-Miniature Chip Inductors Serve at High Frequency," in *J. of Electronic Engineering*, pp. 48-51, December 1993.
- [16] C. Q. Scrantom, J. C. Lawson, and L. Liu, "LTCC technology: where we are and where we're going. II," in *MTT-S Intl. Topical Symp. on Technologies for Wireless Applications*, Vancouver BC, Canada, pp. 193-200, 1999.
- [17] A. Rofougaran, J. Y.-C. Chang, M. Rofougaran, and A. A. Abidi, "A 1 GHz CMOS RF Front-End IC for a Direct-Conversion Wireless Receiver," *IEEE J. of Solid-State Circuits*, vol. 31, no. 7, pp. 880-889, 1996.
- [18] R. E. Lehmann and D. D. Heston, "X-Band Monolithic Series Feedback LNA," *IEEE Trans. on Microwave Theory & Techniques*, vol. MTT-33, no. 12, pp. 1560-1566, 1985.
- [19] D. K. Shaeffer and T. H. Lee, "A 1.5-V, 1.5-GHz CMOS Low Noise Amplifier," *IEEE J. of Solid-State Circuits*, vol. 32, no. 5, pp. 745-759, 1997.
- [20] C. Enz and Y. Cheng, "MOS Transistor Modeling Issues for RF Circuit Design," in *Analog Circuit Design — (X)DSL and other Communication Systems; RF MOST Models; Integrated Filters and Oscillators*, W. Sansen, J. Huijsing, and R. van de Plassche, Eds. Boston: Kluwer, 1999.

- [21] D. E. Norton, "High Dynamic Range Transistor Amplifiers using Lossless Feedback," in *Int'l Symp. on Circuits and Systems*, Newton, MA, pp. 438-440, 1975.
- [22] J. R. Long and M. A. Copeland, "A 1.9 GHz low-voltage silicon bipolar receiver front-end for wireless personal communications systems," *IEEE J. of Solid-State Circuits*, vol. 30, no. 12, pp. 1438-1448, 1995.
- [23] J. Smith, *Modern Communication Circuits*. New York: McGraw-Hill, 1986.
- [24] H. Darabi and A. A. Abidi, "An Ultralow Power Single-Chip CMOS 900 MHz Receiver for Wireless Paging," in *Custom IC Conf.*, San Diego, CA, pp. 213-216, 1999.
- [25] D. B. Leeson, "A Simple Model of Feedback Oscillator Noise Spectrum," *Proceedings of the IEEE*, vol. 54, pp. 329-330, 1966.
- [26] J. J. Rael and A. A. Abidi, "Physical Processes of Phase Noise in Differential LC Oscillators," in *Custom IC Conf.*, Orlando, FL2000.
- [27] A. Hajimiri and T. H. Lee, "Phase noise in CMOS differential LC oscillators," in *Symposium on VLSI Circuits*, Honolulu, HI, pp. 48-51, 1998.
- [28] P. Kinget, "Integrated GHz Voltage Controlled Oscillators," in *Analog Circuit Design: (X)DSL and other Communication Systems; RF MOST Models; Integrated Filters and Oscillators*, W. Sansen, J. Huijsing, and R. van de Plassche, Eds. Boston: Kluwer, 1999, pp. 353-381.
- [29] A. Kral, F. Behbahani, and A. A. Abidi, "RF-CMOS Oscillators with Switched Tuning," in *Custom IC Conf.*, Santa Clara, CA, pp. 555-558, 1998.
- [30] H. Darabi and A. A. Abidi, "Noise in RF-CMOS Mixers: A Simple Physical Model," *IEEE J. of Solid-State Circuits*, vol. 35, no. 1, pp. 15-25, 2000.
- [31] A. M. Durham, J. B. Hughes, and W. Redman-White, "Circuit architectures for high linearity monolithic continuous-time filtering," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 39, no. 9, pp. 651-7, 1992.
- [32] S. Khorram, A. Rofougaran, and A. A. Abidi, "A CMOS Limiting Amplifier and Signal-Strength Indicator," in *Symp. on VLSI Circuits*, Kyoto, pp. 95-96, 1995.
- [33] J. Min, H.-C. Liu, A. Rofougaran, S. Khorram, H. Samueli, and A. A. Abidi, "Low Power Correlation Detector for Binary FSK Direct-Conversion Receivers," *Electronics Letters*, vol. 31, no. 13, pp. 1030-1032, 1995.
- [34] J. Lee, A. A. Abidi, and N. G. Alexopoulos, "Design of Spiral Inductors on Silicon Substrates with a Fast Simulator," in *European Solid-State Circuits Conf.*, The Hague, The Netherlands, pp. 328-331, 1998.
- [35] H. L. van Trees, *Detection, Estimation, and Modulation Theory*. New York: Wiley, 1968.
- [36] J. G. Proakis, *Digital Communications*. New York: McGraw-Hill, 1989.
- [37] J. Zander, "Distributed cochannel interference control in cellular radio systems," *IEEE Transactions on Vehicular Technology*, vol. 41, no. 3, pp. 305-11, 1992.
- [38] S. A. Grandhi, R. Vijayan, and D. J. Goodman, "Distributed power control in cellular radio systems," *IEEE Transactions on Communications*, vol. 42, no. 2-4, pt.1, pp. 226-8, 1994.
- [39] G. J. Foschini and Z. Miljanic, "A simple distributed autonomous power control algorithm and its convergence," *IEEE Transactions on Vehicular Technology*, vol. 42, no. 4, pp. 641-6, 1993.
- [40] N. Bambos and G. J. Pottie, "Power control based admission policies in cellular radio networks," in *IEEE Global Telecommunications Conference*, Orlando, FL, pp. 863-7 vol.2, 1992.

- [41] C. J. Hansen, *Probing Techniques for Multiuser Channels with Power Control*, PhD Thesis, Electrical Engineering, University of California, Los Angeles, 1997
- [42] G. J. Pottie, "Wireless Multiple Access Adaptive Communication Techniques," in *Encyclopedia of Telecommunications*, Froelich and Kent, Eds. Amsterdam: Dekker, 1999.
- [43] S. Verdu, "Minimum probability of error for asynchronous Gaussian multiple-access channels," *IEEE Transactions on Information Theory*, vol. IT-32, no. 1, pp. 85-96, 1986.
- [44] A. Duel-Hallen, "Decorrelating decision-feedback multiuser detector for synchronous code-division multiple-access channel," *IEEE Transactions on Communications*, vol. 41, no. 2, pp. 285-90, 1993.
- [45] U. Madhow and M. L. Honig, "MMSE interference suppression for direct-sequence spread-spectrum CDMA," *IEEE Transactions on Communications*, vol. 42, no. 12, pp. 3178-88, 1994.
- [46] V. Tarokh, N. Seshadri, and A. R. Calderbank, "Space-time codes for high data rate wireless communication: performance criterion and code construction," *IEEE Transactions on Information Theory*, vol. 44, no. 2, pp. 744-65, 1998.
- [47] C. C. Wang and G. J. Pottie, "Interference avoidance and power control strategies for coded frequency hopped cellular systems," in *Intl. Conf. on Communications*, Seattle, WA, pp. 1737-41 vol.3, 1995.
- [48] C. C. Wang and G. J. Pottie, "Bit allocation algorithms for FH-CDMA wireless communication systems," in *Allerton Conf. on Communication, Control, and Computing*, Monticello, IL, pp. 652-61, 1996.
- [49] E. Perahia, *Diversity Combining, Adaptive Antennas, and Equalization for Digital Radio*, PhD Thesis, Electrical Engineering, University of California, Los Angeles, 1995
- [50] K. Sohrabi, J. Gao, V. Ailawadhi, and G. Pottie, "A Self-Organizing Sensor Network," in *Allerton Conf. On Comm., Control, and Computing*, Monticello, IL 1999.
- [51] G. J. Pottie and W. J. Kaiser, "Wireless Integrated Network Sensors: Principles and Practice," *Communications of ACM*, vol. 43, no. 5, 2000.








Fig.1. Direct Conversion Receiver



↓

$$g_{m(\text{eff})} = 2/R$$

↓

$$g_{m(\text{eff})} = 2/R$$

Fig.2. Comparison of degenerated bipolar and MOS large-signal transconductors, designed for the same gain compression point.

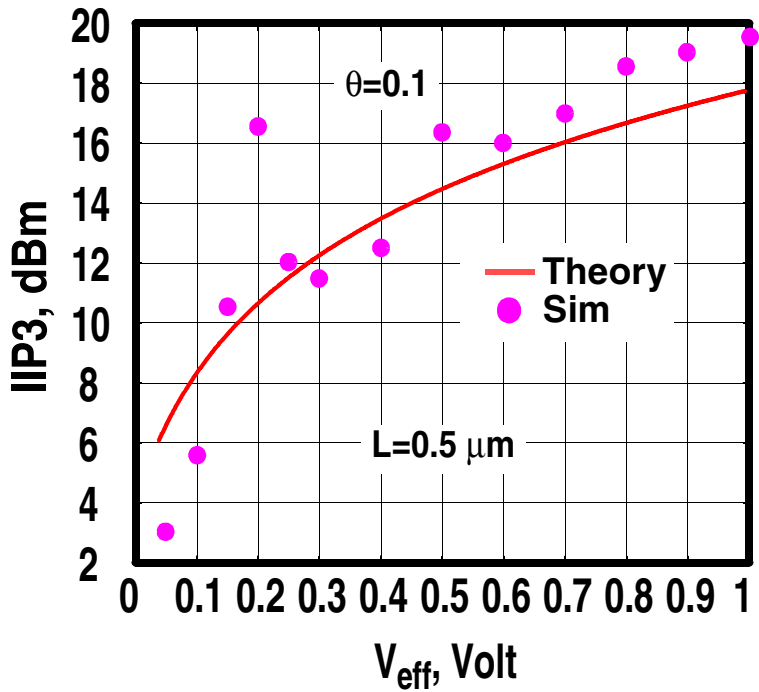


Fig.3. Simulated IIP3 of a MOSFET compared with analytical expression, as a function of effective gate voltage, $V_{eff} = V_{GS} - V_t$.

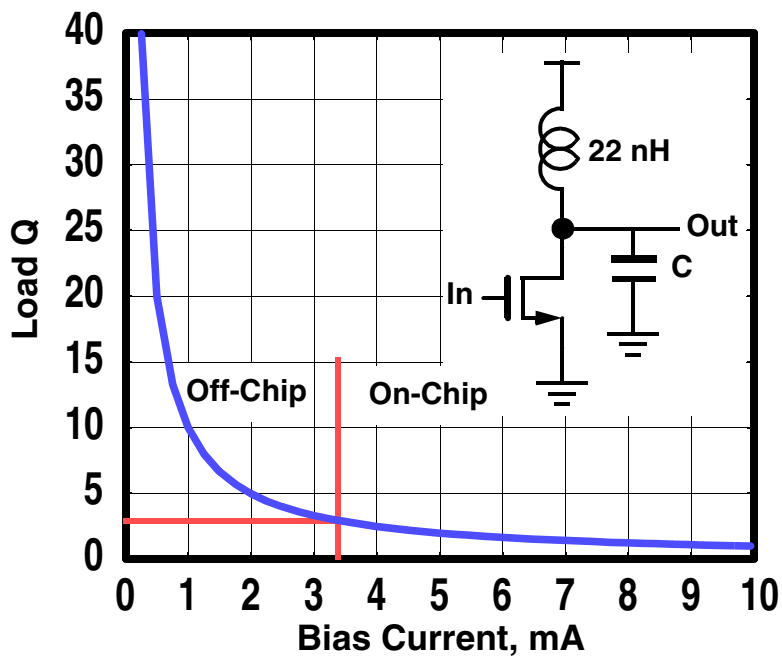


Fig. 4. Bias current versus inductor Q for a simple amplifier at constant gain.

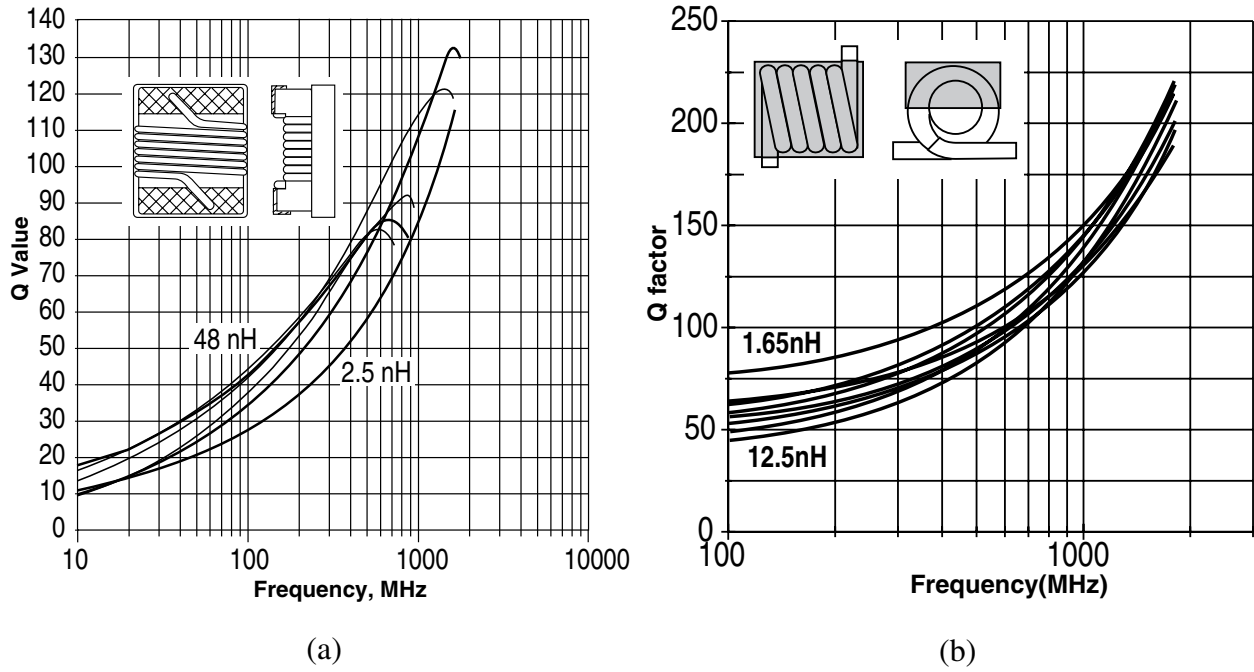


Fig.5. Discrete surface-mount chip inductors, (a) Ceramic core, (b) Air core. Physical structure, and typical Q versus frequency. (Courtesy Coilcraft Corporation).




Fig.6. Ceramic substrate with embedded printed spiral inductors, and flip-chip die attach.




Fig.7. Narrowband common-source amplifier with inductor degeneration. Various parasitics are also shown.






Fig.8. Wideband common-gate amplifier with transformer feedback.



(a)



(b)

Fig.9. (a) Simple common-gate LNA; (b) Low current common-gate LNA with impedance transformer at input port.

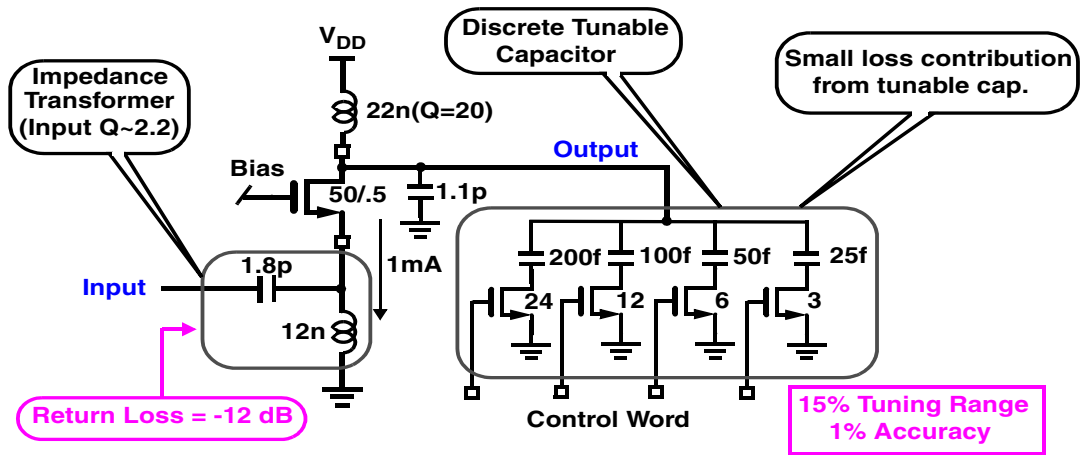


Fig.10. Actual implementation of common-gate LNA, including tuning of high-Q load.




Fig. 11. (a) Differential LC Oscillator; (b) Equivalent circuit




Fig. 12. Amplitude and phase noise dependence on bias current.

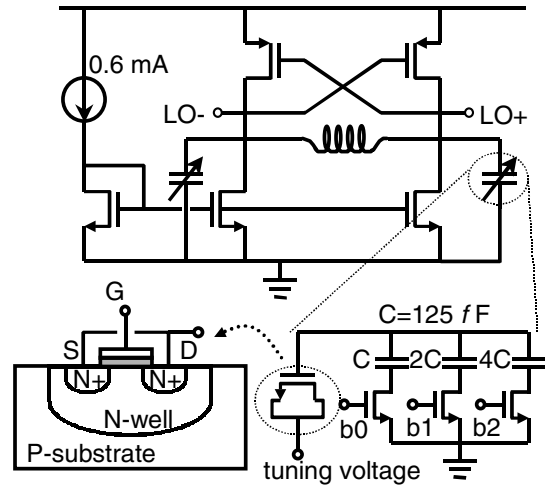



Fig.13. Actual low power oscillator circuits. (a) NMOS differential pair switches with MOS varactor for continuous tuning; (b) PMOS differential pair switches with continuous and discrete tuning.




Fig.14. Passive FET mixer.




Fig.15. Active CMOS mixer example: single balanced.

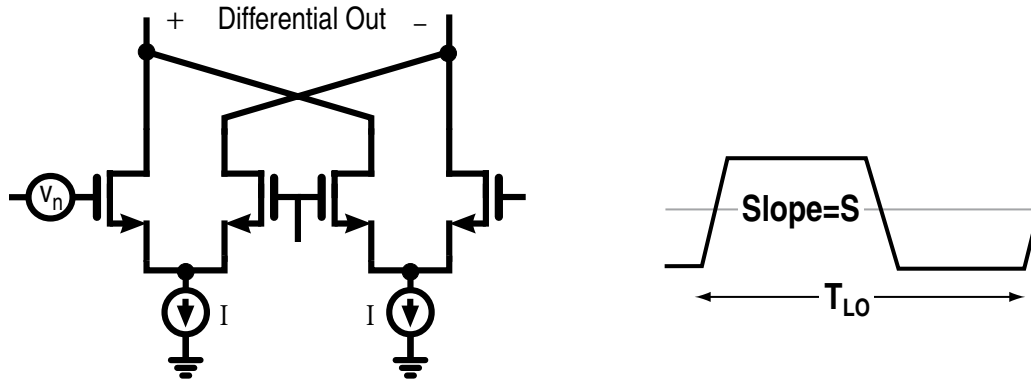


Fig.16. Flicker noise equivalent source at mixer switches, and assumed LO voltage waveform.

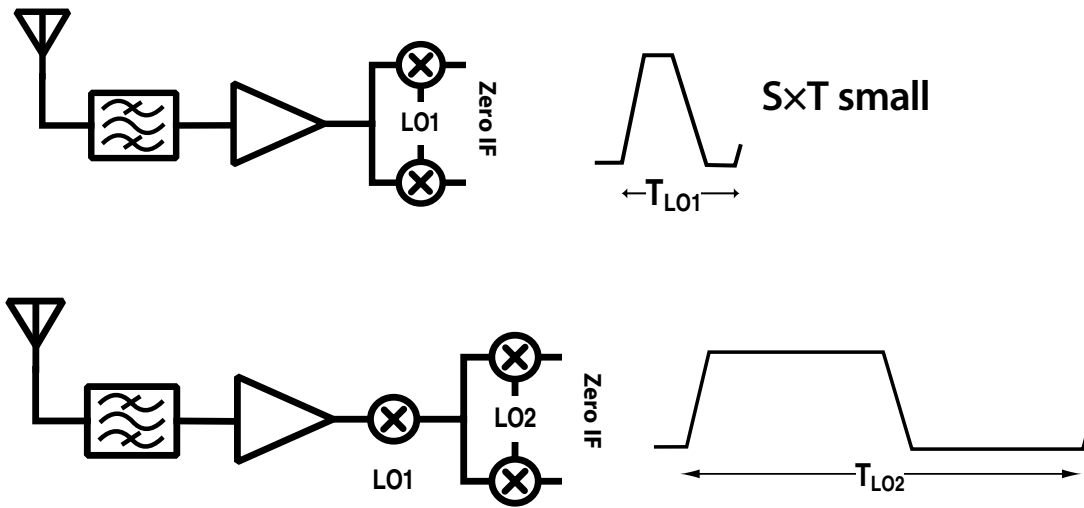


Fig.17. (a) Signal at zero IF in direct conversion receiver competes with large flicker noise at mixer output; (b) Indirect conversion overcomes this.




Fig.18. Low power CMOS divide-by-2, which toggles at 1.5 GHz.

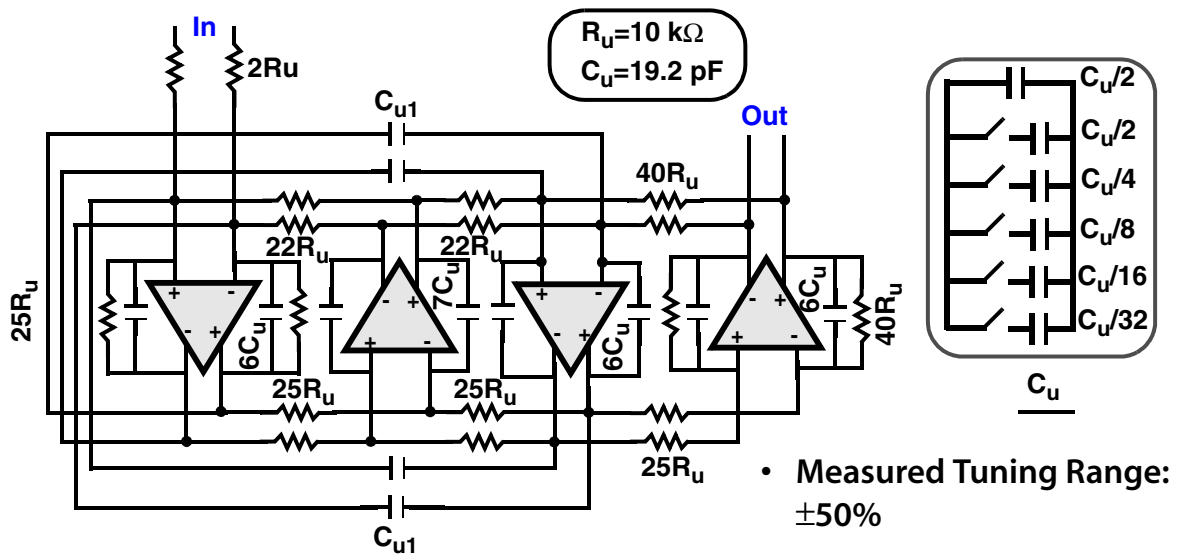


Fig.19. Active RC filter offers best linearity of all time-continuous active filters. Discrete tuning method shown.

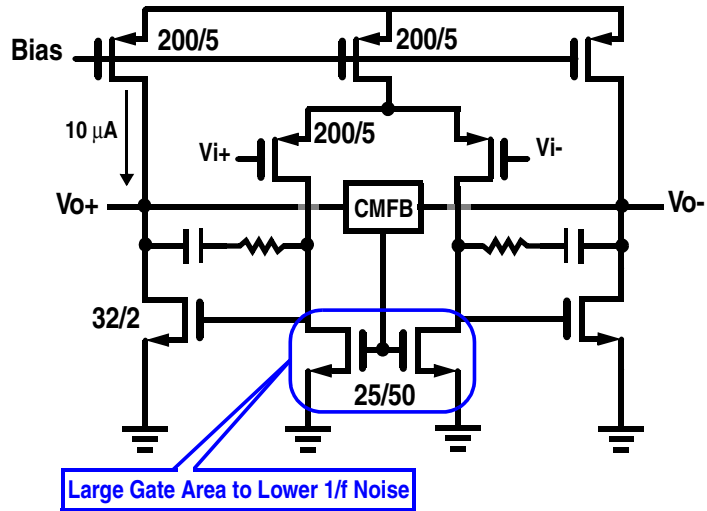


Fig.20. Low power two-stage op amp used in active RC filter.




Fig.21. Input-referred noise of a MOSFET, and its various dependencies on device size.




Fig.22. Three sources of loss in a spiral inductor on silicon: (a) ohmic, (b) displacement current, and (c) eddy current.

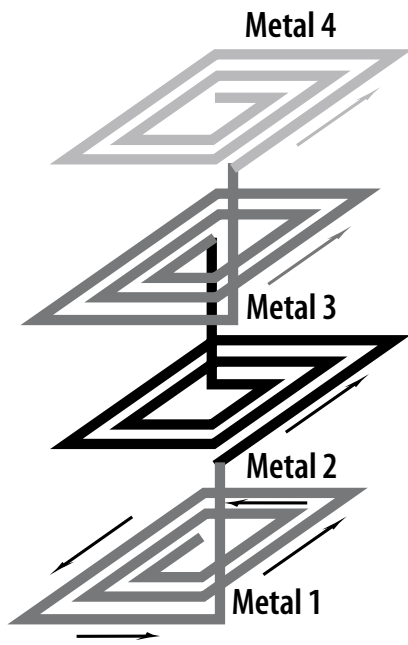


Fig. 23. Multiple spirals in series with coupled magnetic fields realize a vertical solenoid. The self-inductance of the stack increases quadratically with the number of layers.

Comparison across three CMOS generations

	0.6- μm CMOS	0.25- μm CMOS
LNA/Mxr1 Gain	30 dB	32 dB
LNA/Mxr1 NF	5 dB	5 dB
IIP3	-25 dBm	-25 dBm
VCO Tun Range	18%	18%
Power	2 mW/1.25V	2.5 mW/1.25V