Low Power RF Integrated Circuits: Principles and Practice

A. A. Abidi and H. Darabi
Electrical Engineering Department
University of California
Los Angeles, CA 90095-1594

Abstract

Ultralow power wireless communicators, such as paging receivers operating from a single cell, have prompted the development and discovery of circuit techniques and architectures which lower power consumption of RF IC's ten-fold or more compared to today's norms. This paper illustrates many of these design principles, mainly in a CMOS context. It argues for seeking strategic combinations of high quality off-chip passives with RF integrated circuits, and searching for better architectures in wireless receivers (and transmitters) to lower power. The principles are illustrated with specific examples.

Ultralow Power Wireless Devices

Power consumption is the foremost concern in mobile wireless devices. Cellular telephones, for example, use power conscious RF, baseband and DSP IC's, aided by sophisticated power management algorithms to prolong the time between battery recharge. There is another class of wireless communicators, perhaps not as widely used as telephones which operates at similar radio frequencies, but whose power consumption must be 10 to 50× lower. These communicators often work at a low duty cycle, but they must be dependably available over many months without need to recharge or replace the battery. The radio-paging receiver is one example [1]. A 900 MHz FLEX paging receiver operates for 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 on a daily basis. Wireless communicators built into implantable biomedical devices or in miniature remote sensors [2], 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* discovered in the course of recent research which enable lowering the power consumption of RF, IF, and baseband building blocks in integrated wireless communicators by orders of magnitude. 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 wireless devices such as mobile 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.

Dynamic Range and Power Consumption

The fundamental lower limit to power consumption is tied to circuit performance. The main circuit specification in the signal chain of a wireless receiver 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 (and to normalize for gain, this is referred to the receiver input). As an example, take the simple circuit consisting of two common-source MOSFETs, driven with balanced input signal voltages and producing output currents that are differentially sensed. The dynamic range is limited at the lower end by the voltage noise spectral density and channel bandwidth, and at the upper end by the large signal swing that compresses gain, and thereby defines the intercept point. The equivalent input noise voltage density is $v_n^2=4kT\gamma/g_m$, where kT is a physical constant and γ a noise factor depending on the FET channel length and bias. The gain of the circuit distorts significantly when the signal swing shuts off one of the FETs, i.e. when $v_s=V_{GS}-V_t$. Therefore, the g_m and bias $V_{GS}-V_t$ of the FETs determine dynamic range. If the I-V characteristics of the transistors conform even roughly to the classic square law, then the bias drain current in saturation is:

$$I_D = \tfrac{1}{2}\mu C'_{ox}\tfrac{W}{L}(V_{GS}-V_t)^2 = \tfrac{1}{2}g_m(V_{GS}-V_t)$$

This leads to the important conclusion that in 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 current consumption of a practical circuit can approach this limit. Ultimately the system specifications on dynamic range or channel bandwidth limit receiver power dissipation. Broadly speaking, receivers associated with low data rates will dissipate low power.

This argument linking dynamic range to current drain is also true for bipolar transistor circuits. Just as $V_{GS}-V_t$ extends the large-signal handling capability of a FET, resistor degeneration

ISLPED99, San Diego, CA, USA

linearizes a bipolar transistor. The degeneration voltage, IR, 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. Thus R sets the lower end of dynamic range, IR the upper end, and once again dynamic range sets the minimum current drain.

Lowering Power in Tuned Circuits

Many RF circuits such as low noise amplifiers and oscillators are tuned to a narrow frequency band with an *LC* resonant circuit load, and sometimes at the input port for impedance matching. The passive resonant circuit can profoundly influence the circuit power dissipation. As an example, consider the various requirements for a low-noise amplifier matched to 50Ω at the input port.

First, the amplifier must provide adequate gain in the frequency band to which it is tuned. When the stage following the amplifier has input capacitance C, a load inductor (L) is chosen to resonate with it in this frequency band. The resulting peak voltage gain is $g_m Z_0$, where Z_0 is the tuned circuit impedance at resonance. This impedance is usually limited by the inductor quality factor, Q, and is given by $Z_0=Q(\omega_0 L)=Q/(\omega_0 C)$. The voltage gain is thus $Qg_m/(\omega_0 C)$. Now if the amplifier FET is scaled down in width to lower power dissipation, then g_m and the voltage gain will also scale down, unless compensated by inductor Q. Figure 1 illustrates the specific case of a common-source amplifier with a capacitor load, tuned to 930 MHz by a 22 nH inductor. For various inductor Q's, the width of the FET (whose channel length is 0.25μm) is scaled to obtain 20 dB voltage gain. A bias $V_{GS}-V_t$ of 200 mV fixes the large signal handling capability at the FET. Therefore, the FET current scales with its width.

If the 22 nH inductor is implemented on-chip with a Q of say 3, the FET must be biased at 3.3 mA. On the other hand if the Q is raised to 20, the bias current may be lowered to 0.5 mA, a 6.6× reduction in power. Although a spiral inductor with this large a Q cannot be integrated today in a standard IC process, it can be realized as a discrete inductor (Figure 2). This suggests that low power RF circuits must use low-loss discrete inductors fabricated in specialized technologies, a fact that has been noted elsewhere [2].

However, parasitics exact a price in using discrete components. A signal connecting to an off-chip component encounters capacitance and loss at the bond pads, bondwire inductance, the inductance and capacitance of the package leads and of the PC board traces (Figure 3). These parasitics may be large compared to the on-chip reactances, and they do not scale with technology. While an electrostatic shield under the bond pad substantially eliminates substrate loss, the other parasitics, which are mainly reactive and lossless, remain.

A circuit node with large loaded Q is also susceptible to strong detuning when parasitics are either inaccurately estimated, or suffer from large spreads. If variation in gain due to the reactive spreads is to be kept below, say, 1 dB, then the loaded Q at a tuned node must be kept below some upper limit. In turn, this means that the active device attached to that node consumes a minimum current. Some form of active tuning circumvents this limitation. For example, a switched array of small binary-weighted capacitors under control of a digital word re-tunes the resonant circuit to the desired frequency (Figure 4).

A low noise amplifier must be matched at the input port to a characteristic impedance such as 50Ω. At first glance this too is linked to current drain, because FET g_m sets the input resistance of a common-gate amplifier, bias $V_{GS}-V_t$ the large signal handling and together these two factors set the current. However, the LNA may be designed with a higher than wanted input impedance, that is, with a lower bias current, and the impedance may be scaled by a reactive impedance transformer to the desired value (Figure 4). This can be a simple *LC* narrowband transformer, which needs neither coupled magnetic elements nor bias current. This impedance transformation is implicit in a common source tuned LNA. Impedance transformation with modest scale factors is quite feasible in practice, but in seeking too large an impedance scaling the transformer Q becomes very large and is sensitive to parasitics. Transforming up the impedance penalizes intercept point, because the incident input *voltage* also transforms up by the square root of the impedance transformation to conserve energy across the lossless transformer. Thus, referred to the input of the amplifier, the intercept point is proportionally lowered.

Oscillators benefit from high Q inductors in both power dissipation and phase noise. To ensure startup, the g_m of the oscillator FETs must be larger (usually by a factor of 2 to 3 for safety) than the loss conductance of the tuned circuit at resonance. The $V_{GS}-V_t$ determines the oscillation amplitude, and this is usually fixed. Thus, the lower the loss conductance, that is, the higher the Q, the lower the oscillator's power consumption. Again, this argues in favor of a low-loss discrete inductor. The oscillator tuning range must encompass both the frequency band of interest and the spreads in parasitics.

Low Power Mixers

There are various strategies to lower the power dissipation in mixers. Scaling down FET size is the most straightforward. The large signal handling remains unchanged if $V_{GS}-V_t$ is held constant, but the mixer noise scales up. To overcome this, the LNA gain should be larger, which once again argues for the use of high impedance, that is high Q, inductor loads. Although differential LNAs and doubly balanced mixers reject common-mode pickup of stray signals and disturbances, a single-ended LNA consumes half the power of a differential LNA and simplifies connection to the antenna. The LNA output most naturally drives a single-balanced active mixer (Figure 5(a)), whose RF input is single-ended and downconverted output is differential.

A single-balanced mixer suffers from LO feedthrough to the output. A large LO feedthrough can overdrive and saturate the

following stage. Here again the mixer output impedance may be tuned with an inductor load to the frequency band of the downconverted signal, thus attenuating the higher frequency LO feedthrough. From this mixer's output onward, the signal path can be fully differential (Figure 5(b)).

Low Power Frequency Dividers

The frequency prescaler, normally consisting of a variable modulus digital divider clocked at the LO frequency, is the most power hungry digital circuit in a receiver front-end. At high frequencies, or when low power dissipation is important, it is often implemented in BiCMOS as a bipolar ECL circuit. For obvious reasons a high frequency, low power CMOS prescaler is of great interest.

Past experience with GaAs FET prescalers shows that the Source Coupled FET Logic (SCFL) current-steering circuit modeled on an ECL divider works best. However, to switch a CMOS differential pair the internal logic swings are much larger than ECL (Figure 6(a)). A 1V single-ended voltage swing is typical for CMOS. If the swing is smaller, the size of the differential pair FETs must be scaled up to commutate a given current, but this increases the input capacitance. On the other hand, a larger swing implies longer risetime and falltime. These quantities must be optimized. The bias current in subsequent stages of the frequency divider, as the frequency diminishes, may be scaled down.

Technology scaling benefits logic circuits in general, and specifically frequency dividers. With linewidth scalling the parasitic capacitance of source and drain junctions and wiring is also smaller, and this usually dominates the load on the divider stages. In one 0.25-μm CMOS implementation, the 1V swing is set by a 100μA current switched into a 10kΩ load (Figure 6(a)). 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 must never be forced into saturation, it is acceptable to operate a CMOS switch in the triode region. A divide-by-2 circuit using two flip-flops (Figure 6(b)) toggles beyond 1.5 GHz.

Low Power Baseband Circuits

Variable gain amplifiers and active filters are important baseband blocks in a direct conversion or low-IF receiver. As large signal handling is more important than noise in these circuits, their operating dynamic range is translated up with respect to the front-end circuits. The design of an amplifying active lowpass filter for channel selection in a zero-IF receiver is considered here.

Past experience shows that at reasonable power dissipations, the active filter is the noisiest block in the receive signal chain. The noise level of a multi-stage filter can only be reduced at the expense of relatively large increase in power. It is usually more efficient to amplify the signal to overcome filter noise, which referred to the receiver input does not substantially degrade noise figure. Now the problem is that the filter must handle amplified large interferers at its input, which although they lie in the stopband and eventually will be removed at the filter output, may produce intermodulation distortion lying in the filter passband. Filter linearity is hence paramount.

Active filters using op amps are usually more linear than filters using open-loop transconductors. The filter is as linear as the passive components, such as resistors and capacitors, in feedback around the op amp. In a well-designed op amp the gain compresses at a maximum signal swing almost equal to the power supply. There are two ways to realize an op amp-based filter: as a switched capacitor circuit, or as an active *RC* circuit. The switched capacitor circuit is ruled out because the filter in a wireless receiver must handle out-of-band signals across a much wider bandwidth than the single channel of interest. Therefore, either the switched capacitor circuit must be clocked at a large oversampling factor to simplify the anti-alias filter, otherwise an anti-alias filter of high order must be used. Either option raises the power consumption.

The continuous-time active *RC* filter [3] (Figure 7) is interesting because with op amps biased in weak inversion it is possible to accurately produce the desired passband, transition band, and close-in stopband characteristics. In most cases, the natural attenuation through the circuit at high frequencies guarantees a large loss in the far away stopband. A two-stage op amp must be used to drive the filter resistors. The output stage swings rail-to-rail in voltage. The fully differential op amp is frequency compensated with a dominant pole and a zero for stable operation in feedback.

Low frequency flicker ($1/f$) noise is of particular concern in a CMOS zero IF receiver. The op amp circuit contains several measures to lower the input-referred flicker noise. The input stage is PMOS, because the flicker noise is usually lower than in NMOS, possibly due to buried channel conduction. Furthermore, as input-referred flicker noise voltage is inversely proportional to gate area W×L, the width and the length of all FETs in the op amp are scaled up while keeping W/L constant. The longer L lowers f_T roughly in inverse proportion. However, in a baseband filter this is not of much concern until f_T falls to about ten times the highest frequency applied to the filter.

The final op amp sizes are scaled to deliver 80 dB DC gain when driving the actual resistors in the filter (Figure 7). In turn, the resistors determine the noise spectral density in the filter passband, that is, the filter's noise figure. An important way to improve dynamic range is by inserting gain into the filter stages. Interleaving gain with filtering usually produces the largest dynamic range. When filter noise, not intercept point, is the limitation, it is best to concentrate all the gain into the filter's input stage, which significantly lowers the noise contributions of the subsequent filter stages, particularly of those stages whose noise is boosted by high-Q poles.

The pole and zero frequencies of the active RC filter are tunable to overcome process spreads. The lowest power tuning method is with a binary-weighted array of switched capacitors

at every filter node (Figure 7). The number of elements in the array depends on the desired accuracy of filter poles. For example, a 5b array in parallel with a fixed capacitor equal to the largest array capacitor covers ±50% variation in dielectric thickness yet tunes the filter frequency response to 3% accuracy. This is good enough in most practical cases. A filter implemented with this approach gives a 100 kHz passband, a stopband beyond 5 MHz with a loss of 60 dB, while dissipating only 0.3 mA.

Low Power Radio Architectures

Noise is generally of concern in low power circuits, particularly flicker noise in CMOS receivers with low or zero IF. The sections above discuss noise in building blocks; in addition, the correct choice of receiver architecture can alleviate the cumulative effects of noise. For example, consider low frequency noise contributed by the mixer in a direct conversion receiver.

Although this is not widely recognized, the switches in an active CMOS mixer produce flicker noise at the mixer output. This can substantially degrade the signal-to-noise ratio of the received signal which downconverts to zero IF after amplification only in the LNA. The straightforward solution is to lower the mixer noise by scaling up the transistor size, but this means a disproportionate rise in mixer power consumption. A better solution is to revisit the direct conversion architecture and search for some way around mixer noise.

Analysis shows that the switches contribute noise at the mixer output in inverse proportion to S×T, where S is the slope of the LO voltage waveform when the switches commutate, and T is the period of the LO. As the S×T product rises, flicker noise at the output of the mixer falls. Thus, at the expense of one additional low power mixer and its associated LO buffers, the signal can be *indirectly* converted to zero IF by way of some suitable first IF (Figure 8(a)). When the S×T product is small at the high frequency 1st LO the flicker noise is high at the output of the first mixer, but this does not degrade the signal-to-noise ratio because the downconverted signal lies at the first IF beyond the flicker noise corner. On the other hand because of the lower frequency 2nd LO the S×T product is larger at the second mixer. When the signal is finally translated to zero IF at the 2nd mixer output, it competes with a much lower flicker noise.

The signal level diagram in a 1.5V, 3 mA 900 MHz CMOS receiver [4] (Figure 8(b)) shows how gain is allotted to the various blocks to obtain a cascade noise figure and intercept point which meets FLEX receiver specifications.

Conclusions

This paper has surveyed some of the circuit design techniques and system methodologies that have led to ultralow power wireless receivers implemented as bipolar and CMOS IC's [1,4-6]. 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 was 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.

Baseband and IF building blocks are often very challenging, and may consume as much power as the RF sections to realize the necessary dynamic range. A low or zero IF, if the modulation scheme allows, is most desirable from the point of view of lowering power consumption. The dynamic range of circuits up to, and including, the channel-select filter is usually comparable to the receiver front-end dynamic range, except that it is translated up to handle amplified signals. For narrowband channels, circuits in weak inversion may be suitable. Although their noise level is high, these circuits can swing rail-to-rail at the output. An example is given of a baseband amplifier and active *RC* filter. In CMOS, discrete tuning with switched element is an attractive low power option to continuously tunable circuits.

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, scaling down of technology can significantly lower the power consumption of overdriven binary circuits such as frequency dividers.

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 is illustrated with a specific receiver example.

[1] J. F. Wilson, R. Youell, T. H. Richards, G. Luff, and R. Pilaski, "A Single-Chip VHF and UHF Receiver for Radio Paging," *IEEE J. of Solid State Circuits*, vol. 26, no. 12, pp. 1944-1950, 1991.

[2] 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.

[3] A. M. Durham, J. B. Hughes, and W. Redman-White, "Circuit Considerations for High linearity Monolithic Continuous-Time Filtering," *IEEE Trans. on Circuits & Systems-II*, vol. 39, no. 9, pp. 651-657, 1992.

[4] 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.

[5] T. Melly, A. S. Porret, C. C. Enz, and M. Kayal, "A 1.3V Low-Power 430 MHz Front-End using a Standard Digital CMOS Process," in *Custom IC Conf.*, Santa Clara, CA, pp. 503-506, 1998.

[6] S. A. Sanielvici, K. R. Cioffi, B. Ahrari, P. S. Stephenson, D. L. Skoglund, and M. Zargari, "A 900 MHz Transceiver Chipset for Two-Way Paging Applications," *IEEE J. of Solid-State Circuits*, vol. 33, no. 12, pp. 2160-2168, 1998.

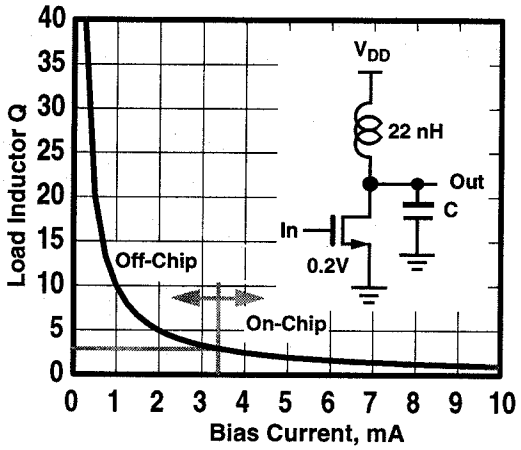

Figure 1. Bias current for a common-source NMOS amplifier with tuned load, as a function of Q of the load inductor.

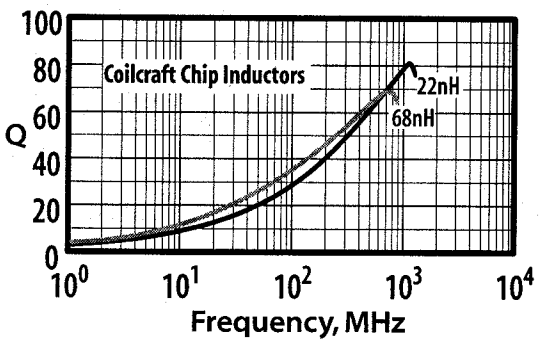

Figure 2. Measured Q vs frequency of discrete surface mount chip inductors (manufactured by Coilcraft).

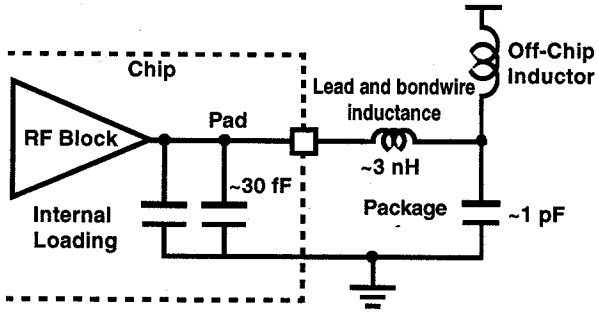

Figure 3. Representative parasitic capacitors and inductors associated with a line carrying a signal to an off-chip inductor load.

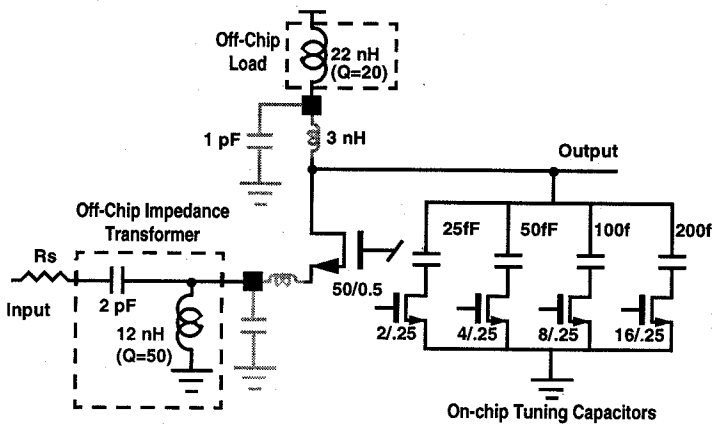

Figure 4. Common-gate LNA tuned to 900 MHz and biased at 1 mA. Narrowband reactive network matches input impedance to 50Ω. Off-chip inductor load raises LNA gain. On-chip trim capacitor array centers resonant frequency of high-Q output node.

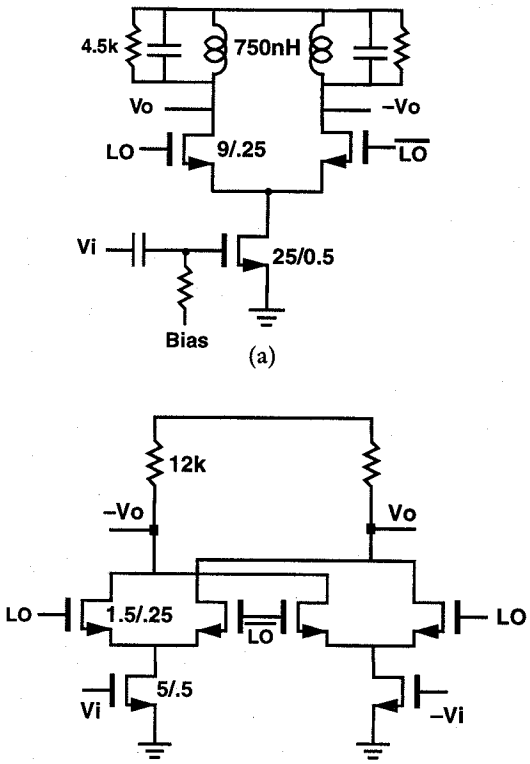

Figure 5. Low power mixers. (a) Single-balanced mixer capacitively coupled to 900 MHz LNA output, biased at 0.5 mA. Load tuned with large off-chip inductor to reject LO feedthrough. (b) Double-balanced mixer with 100 MHz input, biased at 0.3 mA. Note the very small FET sizes.

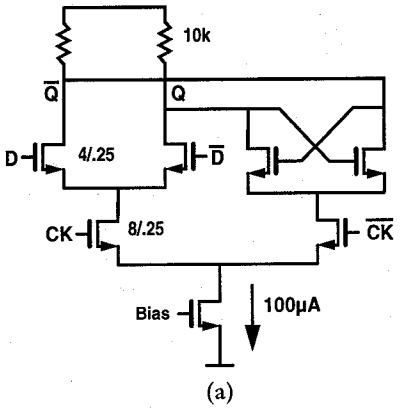

Figure 6. (a) Low power, high frequency SCFL D flip-flop clocks at 1.5 GHz. Logic swing is 1V, nominal supply voltage 1.5V. (b) Divide-by-2 circuit.

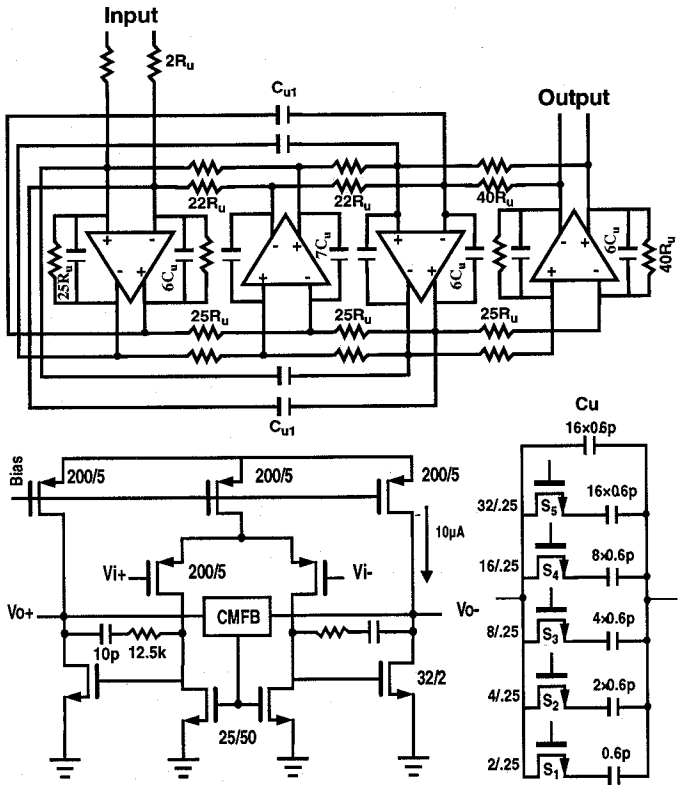

Figure 7. Example of a fourth-order elliptic active *RC* filter. Resistors are polysilicon, detail of unit capacitor shown. Micropower two-stage op amp drives resistor load.

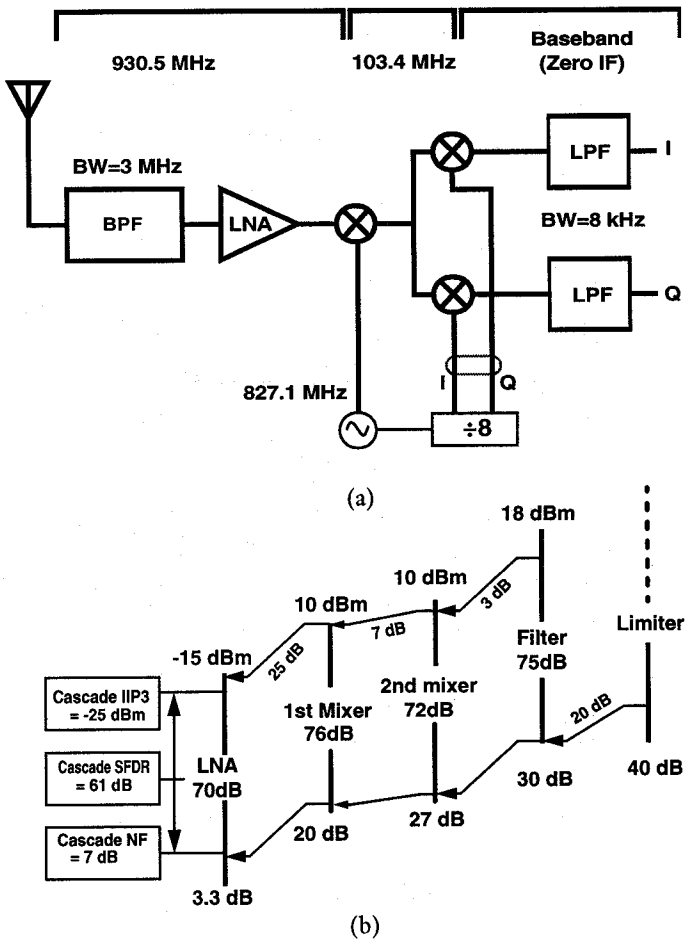

Figure 8. (a) Receiver to "indirectly" convert RF signal to zero IF. By downconverting the RF signal to zero 2nd IF from a 1st IF at 100 MHz, this architecture circumvents SNR degradation due to elevated flicker noise at the output of the first mixer. (b) Signal level plan in receiver, showing noise figure and IIP3 of each block wrt 50Ω, and spurious free dynamic range (SFDR).