An 800-MHz–6-GHz Software-Defined Wireless Receiver in 90-nm CMOS

Rahim Bagheri, *Student Member, IEEE*, Ahmad Mirzaei, *Student Member, IEEE*, Saeed Chehrazi, *Student Member, IEEE*, Mohammad E. Heidari, Minjae Lee, *Student Member, IEEE*, Mohyee Mikhemar, *Student Member, IEEE*, Wai Tang, *Student Member, IEEE*, and Asad A. Abidi, *Fellow, IEEE*

***Abstract*—A software-defined radio receiver is designed from a low-power ADC perspective, exploiting programmability of windowed integration sampler and clock-programmable discrete-time analog filters. To cover the major frequency bands in use today, a wideband RF front-end, including the low-noise amplifier (LNA) and a wide tuning-range synthesizer, spanning over 800 MHz to 6 GHz is designed. The wideband LNA provides 18–20 dB of maximum gain and 3–3.5 dB of noise figure over 800 MHz to 6 GHz. A low 1/f noise and high-linearity mixer is designed which utilizes the passive mixer core properties and provides around +70 dBm IIP2 over the bandwidth of operation. The entire receiver circuits are implemented in 90-nm CMOS technology. Programmability of the receiver is tested for GSM and 802.11g standards.**

***Index Terms*—CMOS, direct conversion, flicker noise, frequency synthesis, GSM, IEEE 802.11a, IEEE 802.11b, IEEE 802.11g, linearity, low-noise amplifier, mixer, RF transceiver, software-defined receiver, wideband matching, wireless LAN, zero IF.**

I. Introduction

A. The Need for Software-Defined Radio

WIRELESS bands and services are proliferating across the world. Every six months it seems a new use for wireless appears, often leading to a new standard. Manufacturers of mobile handsets have a hard time keeping up, because the end user wants to access an increasing number of services from a single handset, and have it adapt to global roaming.

Today these handsets consist of receiver or transceiver boards, often from different suppliers, packed densely together. Some sharing of functions is occasionally possible, for example, in a quad-band GSM/DCS radio which shares a common IF and baseband [1]. The RF portions resolutely defy sharing. A 900-MHz receiver will use a separate surface acoustic wave (SAW) filter, low-noise amplifier (LNA), and mixer from that of an 800-MHz receiver. The handset will only grow in size and cost as it provides more services, which is not what the user or the manufacturer wants. In the face of this proliferation, a universal programmable software-defined radio (SDR) is urgently needed.

Manuscript received June 27, 2006; revised August 26, 2006.

R. Bagheri, A. Mirzaei, S. Chehrazi, and M. E. Heidari are with the Department of Electrical Engineering, University of California, Los Angeles, CA 90095 USA, and also with Wilinx, Inc., Los Angeles, CA 90025 USA. (e-mail: r.bagheri@wilinx.com; a.mirzaei@wilinx.com).

M. Lee, M. Mikhemar, and A. A. Abidi are with the Department of Electrical Engineering, University of California, Los Angeles, CA 90095 USA (e-mail: abidi@icsl.ucla.edu).

W. Tang was with the Department of Electrical Engineering, University of California, Los Angeles, CA 90095 USA. He is now with Silicon Laboratories, Inc., Austin TX 78735 USA.

Digital Object Identifier 10.1109/JSSC.2006.884835

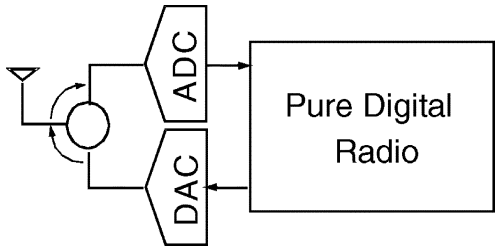

Fig. 1. Classic SDR as defined by Mitola [3].

The term "software-defined radio" has different meanings to different people. What we mean is a radio receiver or transmitter that can be tuned to carrier frequencies over a wide, though not infinite, range, and which can support any reasonable modulation over a range of data rates. In other words, the SDR is a versatile platform that could be the universal module to build tomorrow's handsets. Bits in a register and a clock frequency should configure this platform to any application, whether it is cellular voice, WLAN, TV reception, or a short-range link. A platform tunable from 800 MHz to 6 GHz will cover all major bands in use today [2].

Also of interest is on-demand use of unallocated or presently unoccupied wireless spectrum, a form of spectrum squatting that goes under the label of cognitive radio. Without a low-cost device that embodies all the features of the SDR platform described above, cognitive radio will not become reality.

B. Mitola's Concept

The software-defined radio concept originates from military needs, where a radio must tune to one of multiple bands to communicate, and one or multiple bands simultaneously to eavesdrop. Mitola [3] defined the SDR as one where the only analog components are an RF A/D converter (ADC) at the receiver, and an RF D/A converter at the transmitter (Fig. 1). All radio functions are realized in a programmable digital signal processor (DSP). Given the rapid progress in CMOS scaling and advances in data converters, this looks like a good approach at first sight not just for the military, but for all future radios.

Although with state-of-the-art IC technology, the DSP and DAC are almost within sight, the ADC is far from realizable. Extrapolating from Walden's curves of 1999 [4], the 12-GHz, 12-bit ADC required to digitize all possible signals incident on the receiver will dissipate 500 W. The dynamic range of the

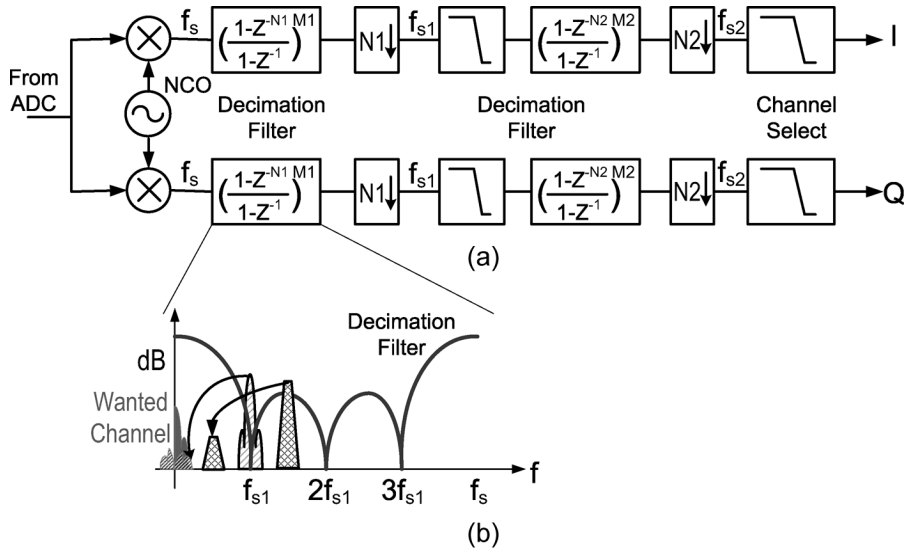

Fig. 2. (a) Classic SDR digital front-end downconverts and lowers sample rate of the signal. (b) Example of decimation filter response.

ADC is set by the quantization noise floor, which determines receiver noise figure. Progress in ADCs takes place at a much slower rate than Moore's Law, which mainly benefits the DSP.

Realizing the huge overhead associated with the ADC, the idea of a radio that needs an ADC at the antenna lost credibility among IC designers. In recent years, it only crops up in panel discussions and proposals as something which is extreme. Meanwhile, a small community of DSP researchers has continued to develop network-on-a-chip architectures for the SDR baseband, and others have explored ideas on the digital front-end that might follow such an ADC if it were realizable.

It is well worth asking why this has proved intractable. When DSP is fast enough that it can execute functions of the digital front-end and baseband in software, Mitola's concept of the SDR can receive every channel incident on the receiver *concurrently*. To receive one channel, a software digital downconverter first centers it to zero IF with quadrature phases of a numerically controlled oscillator. Then, decimating and interpolating filters convert the ADC sample rate (12 GHz) to the symbol rate for equalization and detection. In software, this process can be parallelized for concurrent reception of any number of channels, all of which are available at the ADC output.

This may be what the military wants, but it far exceeds civilian needs. In the next section, we look at scaling down the SDR for mobile terminals.

C. SDR for Mobile Terminals

When we use a personal communications device, we know beforehand what service we want. There may be times when we want more than one service, but this need for concurrent access is limited at most to a few services. It makes sense, then, to narrow the SDR concept to one that receives and transmits a single channel at a time. What distinguishes it from a conventional transceiver is that the SDR operates in any band, on channels with any reasonable bandwidth and modulation scheme. The A/D conversion rate can be lowered by many orders of magnitude from that necessary to digitize *every* incident channel to that required to digitize any *one* channel. Today, the wanted channel's bandwidth may be anywhere from 200 kHz for GSM [5] to, say, 20 MHz for 802.11g WLAN [6]. The savings in ADC power will be enormous.

This assumes, however, that the wanted channel is isolated from the multitude of unwanted channels around it that also pass into the receiver's wideband front end. But an SDR has no RF preselect filter. If adjacent channels are to be suppressed by a baseband analog channel-select filter, its passband must be variable by 100:1. The filter transition band and stopband must also be adjusted from band to band to conform to different blocker profiles. The resulting conventional continuous-time filter would require so many degrees of programmability that it is no longer practical. It is soon obvious that we must seek inspiration elsewhere. As we will see, this comes from the way that digital front-ends following the ADC are designed.

At this early stage of development, we restrict ourselves to a receiver that operates in half-duplex, or time-division duplex. This means that the receiver can service all standards except IS-95 and CDMA2000, which are full duplex.

The most promising path to find the right architecture is to work upstream from the A/D converter towards the antenna. Given the almost unlimited extent of unwanted channels, the question looms large of where and how to sample the incoming multi-channel waveform at a reasonable rate? This is the main question, and it is here that developments in the digital front-end prove a useful guide.

II. Digital Baseband Front-End: Operation

In anticipation of the RF/analog front-end of SDR becoming feasible, researchers have explored in some depth the digital front-end that would follow the hypothetical ADC [7]. The digital front-end takes a large array of digitized channels, and isolates the channel of interest, capturing its modulation pre-envelope as complex samples at zero IF. This sample stream is equalized and demodulated. The process lowers the raw ADC output rate to one channel's symbol rate [Fig. 2(a)].

A. Sample Rate Conversion

Multi-rate signal processing is well established in DSP. It deals with the transition between two discrete-time systems which operate at different clock rates. Decimation is the process of *downsampling*, and interpolation that of *upsampling*. A combination of the two can convert between rates which are the ratio of two integers.

When the Nyquist band shrinks, it folds in channels that would be left outside. In the context of bandpass wireless channels, the main question is how to downsample while protecting the desired channel from aliasing. As aliasing imperils the channel of interest, a filter must suppress all possible aliases *before* downsampling. In conventional wireless receivers, this could be a fixed filter. However, what is needed in SDR is a *tracking filter* that knows the frequencies of all possible aliases and suppresses them for any downsampling factor.

This problem is well understood in the DSP art. When the channel of interest occupies a narrow band around DC the decimation filter is easily realized (Fig. 2, $M = 1$). The simplest filter comprises an N-stage FIR structure with uniform tap weights, whose output is sampled at $1/N$th the clock rate f_S of the input. The transfer function of the filter is

$$H(z) = \frac{1 - z^{-N}}{1 - z^{-1}} \Rightarrow |H(f)| = \left| \frac{\sin(\pi f N / f_S)}{\sin(\pi f / f_S)} \right|. \quad (1)$$

This filter nulls input frequencies at the downsampled rate f_S/N and every multiple. This guarantees that after downsampling, nothing aliases to DC. This also suggests the right way to use this in a receiver. Let us assume that all channels of interest are of very narrow bandwidth compared to the initial sample frequency. First, the NCO tunes the wanted channel to zero IF. Next, the high initial sample rate is decimated. The decimation filter attenuates unwanted channels as the first step in channel selection and suppresses aliasing channels [Fig. 2(b)]. Partly suppressed unwanted channels beyond f_S/N become adjacent channels after aliasing. After decimation to the final rate, they are deeply suppressed.

So far, we have described a simple decimation filter. Information-bearing channels occupy non-zero bandwidth, which means that a null at one frequency does not suppress across the entire unwanted bandwidth; a residue of the channel remains. It is impractical to strive for a null across a non-zero band; what is more reasonable is to ensure a minimum attenuation across that band. For this purpose, higher order decimation filters can be used, which attenuate more deeply around the nulls.

This digital receiver is very flexible. It can handle any channel by selecting the clock frequency, the number of filter taps, and decimation factor. This is why, as much as possible, the DSP should be responsible for radio signal processing, while the RF/analog front-end should be the minimum necessary to support this with ADCs that dissipate milliwatts instead of watts.

III. Analog Baseband Signal Conditioner

What kind of RF/analog front-end will give a comparable flexibility to the digital front-end just described? Moving up the signal processing chain from the digital front-end lies the ADC. In a mobile terminal, the product of ADC resolution and sample rate must be reasonable. These should be chosen so as to simplify, as much as possible, the analog pre-processing prior to A/D conversion. This means that the channel of interest, and adjacent channels, should be sampled with a minimal filtering and amplification consistent with low-power ADCs. Conventional baseband amplifiers and continuous-time analog filters require substantial design effort to meet noise and linearity, and they are increasingly difficult to design in scaled technologies at low supply voltages. Therefore, the analog circuits used should be as simple as possible.

It soon becomes clear that from the ADC upstream, it is most advantageous to locate the channel of interest at zero IF. This is because the ADC and preceding analog circuits are by nature low-pass. Although in principle these circuits can be extended to bandpass, there is no advantage in doing so—given the advances in zero-IF signal processing—and only disadvantages arising from limited image rejection. Zero IF requires quadrature channels and analog circuits suffer from gain mismatch in two or more channels. Quadrature accuracy cannot be guaranteed beyond a certain limit, but at zero IF this matters least because the effect is to superpose a small fraction of the spectrally inverted wanted channel on to itself. The effect is considerably more benign than if this were a portion of an unwanted channel of some unpredictable, large relative strength.

A. Low Power ADCs at Baseband

Let us budget 10 mW of power from a 1-V supply to the ADC. Going by recent publications, at this power consumption it is possible to realize a 9-bit, 40-MHz Nyquist ADC, or a 14-bit, 100-kHz noise-shaped delta-sigma ADC that samples at 9 MHz. We will assume that both are available, and one or the other is chosen based on the wanted channel's bandwidth and blocker profile.

B. Absorbing Variable Gain Into ADC

As the first example of how DSP can lead to a simpler analog front-end, let us consider the matter of automatic gain control (AGC). Every wireless receiver must adjust its gain to the strength of the wanted channel, which can vary over orders of magnitude. Traditionally, this duty is fully assigned to the analog front-end, which uses amplifiers whose gain is either programmable in discrete steps, or is continuously variable; usually, it is a mixture of both. The gain range usually matches the range of expected signal strengths, and it may be distributed between some variable gain at RF, and the remaining at baseband.

AGC can be implemented in DSP as well. When the dynamic range of the A/D converter exceeds the dynamic range of AM, signals of varying strengths can be digitized and their values normalized. With advances in A/D converters, it is reasonable to expect digital AGC to assume a greater role. Let us illustrate this with the case of GSM.

According to the GSM standard [5], the channel of interest can be anywhere from -102 to -15 dBm in strength at the receiver antenna (Fig. 3). This is an 87 dB dynamic range. In the spirit of "digital as much as possible", let us determine the minimum requirement on analog AGC. Starting from the A/D converter, we assume 14-bit resolution with $+1$ dBm full scale. The

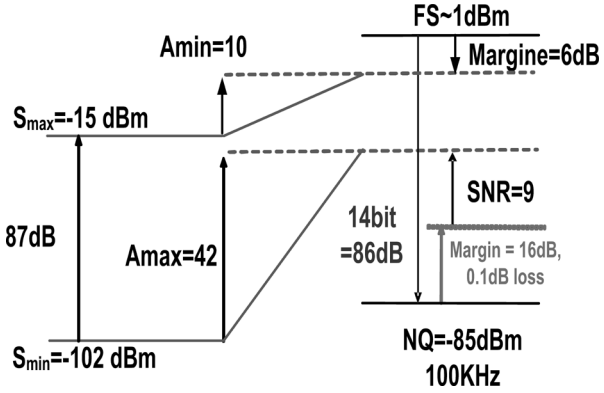

Fig. 3. Programmable gain amplification requirement for GSM example.

quantization noise floor in the 100 kHz bandwidth of the GSM channel is then at −85 dBm. GSM needs a detection SNR of 9 dB, so if quantization noise sets the noise floor at the ADC input, the RF/analog sections prior to the A/D should amplify by 26 dB. However, in reality thermal noise in the receiver front-end sets the detection SNR, and to limit further degradation by quantization noise only to 0.1 dB, the signal must be amplified by another 16 dB, leading to maximum gain of 42 dB (Fig. 3).

If the receiver gain were to remain constant at 42 dB, the largest GSM input signal would overload the ADC. For large inputs, gain must be lowered from 42 to 10 dB, which leaves a 6 dB margin below full-scale, allowing for the envelope variations in EDGE and programmable gain amplifier (PGA) gain setting error. Thus, by exploiting the available A/D dynamic range, a 32 dB variable or programmable gain in the receiver's RF/analog portion is sufficient to capture a single input channel with 87 dB dynamic range; the DSP absorbs the rest (Fig. 3).

Similarly, the smallest range of RF/analog variable gain can be found for 802.11g reception. For this wideband system, the minimum detectable signal at the receiver input is −82 dBm in normal mode, and −65 dBm at high data rate. The largest input signal is −20 dBm [6]. To digitize the 20-MHz-wide channel, a resolution of 9 bits or so is feasible within the power budget. Allowing for the 26 dB SNR for detection at acceptable error rate, and given the smaller dynamic range of the ADC compared to the case of GSM, the RF/analog gain should be variable from 8 to 47 dB. In the wideband case, then, the analog part carries a larger burden. This is a reasonable trade-off; because to shift the burden to DSP would cost a disproportionate rise in A/D power consumption.

C. Sampling a Wireless Channel at DC

We want to sample a narrowband channel at zero IF, which is surrounded by unwanted channels in the same band, unwanted channels in other bands, and unknown channels everywhere else. The impulse sampler is deeply rooted in the electronics tradition because it samples all frequencies with equal fidelity and gain. In the frequency domain, its output is obtained by convolving an infinite sequence of discrete spectral lines at DC and all positive and negative multiples of the sampling frequency with the input spectrum.

An anti-aliasing filter must precede the sampler. This is ideally a brick-wall low-pass filter with a cutoff conservatively at the Nyquist frequency, but more aggressively at the sampling frequency less half the wanted channel bandwidth. Considering the limited dynamic range of the following blocks, especially the ADC, the filter has to suppress the non-aliasing blockers as well. It is not easy to realize this filter, particularly at 1-V supply, if it is to meet demands on low noise and high linearity. In other words, the impulse sampler offers the wrong paradigm for a wireless receiver.

A sampler is needed that intrinsically de-emphasizes adjacent, and particularly far away, unwanted channels. It will be a great bonus if anti-aliasing is built into this sampler. Such a sampler is known, although seldom used. We describe it next.

D. Windowed Integration Samplers

The impulse sampler is based on capturing the instantaneous value of a waveform at the moment the sampling switch is opened. The track-and-hold sampler is its practical realization. The sampling switch closes for a window as large as the sampling period, to enable the sampling capacitor to settle into tracking the dynamic input waveform. Then the switch opens to capture the instantaneous tracked value on the capacitor. The faster the switch shuts off, the more accurately the capacitor acquires any fast dynamics on the waveform. Instantaneous turn off amounts to impulse sampling, subject to bandlimiting caused by the non-zero switch resistance.

Another form of sampling is based on *integrating* the input waveform over a fixed time window [8], [9], and the resulting integral is taken as a sample [Fig. 4(a)]. The window can be as wide as the sample period, and its repetition rate defines the sampling frequency. For a window T_W, the output sample is

$$v_{\text{out}}[n] = \frac{1}{\tau} \int_{nT_S - T_W}^{nT_S} v_{\text{in}}(t)dt \tag{2}$$

where τ is a normalizing time constant. This relation can be transformed into a transfer function $H_{WI}(s)$, and thus to a frequency response:

$$H_{WI}(f) = \frac{T_W}{\tau} \frac{\sin(\pi f T_W)}{(\pi f T_W)}. \tag{3}$$

This means that the transfer function is low-pass, with nulls at frequencies $f = 1/T_W, 2/T_W, \ldots$ to infinity. If the signal of interest clusters around DC this arrangement guarantees built-in anti-aliasing for a sample rate $f_S = 1/T_W$. This is another good reason to downconvert the channel of interest to zero IF [Fig. 4(b)].

Others also have recognized the use of a windowed integrator to sample a narrowband channel at zero IF [10], [11]. However, they assume a preselect filter, and do not fully exploit the power of built-in anti-aliasing and inherent filtering. Whereas in this work, we show how the anti-aliasing sampler and discrete-time filter can *displace* the RF prefilter.

The windowed integration defined so far is the simplest possible, comprising a uniform window. This idea can be generalized to shaped weighting functions, which then realize more sophisticated filters on the wideband input. That is, suppose the

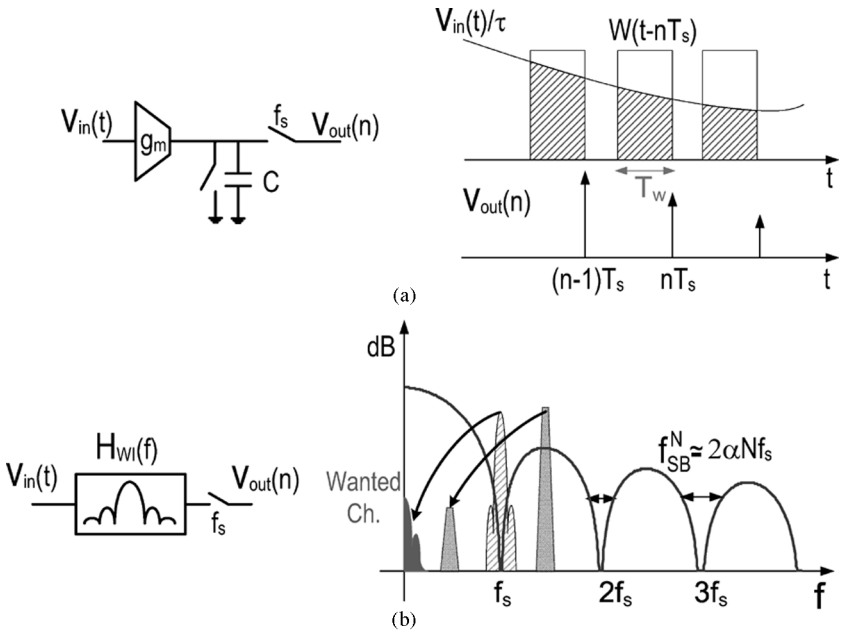

Fig. 4. (a) Simple windowed integration sampler. (b) Windowed integration sampler built-in anti-alias filter utilized for zero-IF receiver.

input is first weighted by a carefully synthesized waveform $w(t)$ which spans over $[-T_W, 0]$, and then integrated; it follows that

$$v_{\text{out}}[n] = \frac{1}{\tau} \int_{nT_S - T_w}^{nT_S} v_{\text{in}}(t) w(t - nT_S) dt. \tag{4}$$

This can be written in another way, in terms of the compact, time-reversed function $h(t) = w(-t)$:

$$v_{\text{out}}[n] = \frac{1}{\tau} \int_{-\infty}^{\infty} v_{\text{in}}(t) h(nT_S - t) dt. \tag{5}$$

When cast into this form, we can recognize that the weighted integration is, in fact, a convolution with a finite impulse response $h(t)$; that is, the weighted integrator is a linear filter with transfer function $H(s)$, the Laplace transform of $h(t)$.

Let us take a specific example. Suppose we wish to realize an anti-aliasing filter whose frequency response is $H(j\omega) = \text{sinc}^2(\omega T)$. Its sidelobes roll off at twice the rate of a sinc filter—which means greater suppression of unwanted channels—and deeper attenuation surrounds the nulls. Then we can take the inverse Laplace transform of the transfer function to obtain the weighting function, or we can derive it using reasoning familiar to electrical engineers. Following the second approach, we recognize that the desired filter transfer function is the cascade of two windowed integrators. The trouble is that two windowed integration samplers cannot be physically cascaded, because the first produces a discrete-time output, while the second expects a continuous-time input. The impulse response of the cascade is the convolution of the impulse response of each block in the cascade. The finite impulse response of a simple weighted integrator is a rectangle of unit height, and width T_w; this is known to physicists as the boxcar. Convolving a boxcar with itself leads to an isoceles triangle with base $2T_w$. This is sufficient to guide the design of a second-order anti-aliasing filter circuit [12].

One may extrapolate from this to yet higher order filters. We know from the frequency response of digital FIR filters that the stopband is non-monotonic; that is, the frequency response has an infinite series of diminishing sidelobes.

When there are very strong unwanted channels, it might seem that a high order anti-alias filter is required whose sidelobes diminish very rapidly. Though, high order anti-alias filter can be implemented [12], its linearity suffers in scaled technologies. However, as we will show, it is possible to obtain a comparable filtering in a cascade of a first-order sampler and discrete-time filters which also has high linearity. For this to work, we must select the right initial sample rate.

E. Choice of Sample Rate

We define the anti-aliasing stopband at the Nth null of the sinc filter as the band f_{SB} across which a certain attenuation α is guaranteed [Fig. 4(b)]. For small α

$$f_{SB}^N \simeq 2\alpha N f_s. \tag{6}$$

A given attenuation is obtained across wider bandwidths at higher nulls. Alternatively, across a given bandwidth surrounding each null, the attenuation is larger at higher nulls. The worst case is the stopband associated with the first null.

The channel of interest at zero IF must be protected from aliasing across its entire bandwidth; this sets the minimum f_{SB}^N. The tolerable co-channel interference and the blocker level determine α. Irrespective of the blocker bandwidth, sufficient anti-aliasing must be guaranteed across the wanted channel's bandwidth.

Clearly, with a first-order integrating sampler, a certain stopband and α dictate the minimum sample frequency. Thus, the sample rate may turn out to be orders of magnitude higher than the wanted channel bandwidth. This should not raise an alarm,

because it is relatively easy to clock and sample at GHz frequencies in scaled CMOS.

In a practical circuit, the filter gain at the nulls is actually non-zero. This is because of leakage in the integration arising from finite output resistance $r_{\rm out}$, and non-zero rise and fall times of the window-defining clock. Leakage causes the zeros of the transfer function to shift by $1/(r_{\rm out}C)$ to the left of the $j\omega$ axis of the s-plane. Non-zero rise and fall times of the clock have opposite effects on the integration, but when the times are equal, their effect cancels. However, as long as the minima are less than the desired α, the sample rate will determine the effective f_{SB}.

Sometimes the blockers are so large that it is impossible to obtain the required alias suppression from a practical windowed integrator. That is when we must seek the assistance of a simple, robust, linear low-pass filter with monotonic stopband, that is, the higher the frequency, the greater the attenuation. A simple, passive RC filter fits the bill. The cutoff frequency of an RC filter can vary by 50%, which means that the nominal pole frequency must be placed well above the bandwidth of the wanted channel. In cases like GSM where out-of-band blockers must be suppressed by more than 100 dB, a cascade of two RC poles may be required. The sample rate must be selected based on the achievable α, the desired suppression of the blocker at an offset equal to the sample rate, and the attenuation at that offset provided by the RC filter.

F. Analog Sample Rate Conversion

Based on the considerations above, the sample rate finally chosen may be many hundreds of MHz. Although ADCs can operate at these rates, their power dissipation will exceed the budget. Given that the channel of interest is anywhere from hundreds of kHz to 20 MHz wide, it should be possible to digitize it at a much lower rate provided there is no aliasing. This needs downsampling, or decimation, of the anti-aliasing sampler's output. With downsampling, aliasing is unavoidable. Just as in the integrating sampler, the wanted channel must be protected from aliasing's adverse effects. On the other hand, there is no harm if aliasing contaminates unwanted channels because they will eventually be filtered. A decimating filter offers this protection. The most well-known decimation filter of Nth order has a sinc^N frequency response, with nulls at the output sample rate and its multiples. This protects aliasing at DC and its vicinity.

This anti-aliasing action is built-in to the weighted integration sampler, already discussed in a previous section. What is different here is that the input to the decimation filter is already sampled. Discrete-time analog decimation filters are better known as charge-domain realizations of FIR filters. To decimate by M and suppress aliases by sinc filtering, M successive input samples are stored on M equal capacitors, which are then connected in parallel. This realizes the weighted sum of charge that corresponds to convolution in discrete-time, and M times downsampling, because the summing takes place once every M samples.

The extent of the impulse response and the decimation factor can be chosen independently. When the impulse response spans many samples but decimation is across a smaller number of samples, two or more weighted summations must be interleaved in time.

In addition to anti-aliasing, the decimation filter gives useful low-pass filtering. The initial sampling frequency is the lowest that will attenuate aliasing blockers across the channel bandwidth. It follows that at a lower rate, a sinc filter *alone* does not sufficiently attenuate these blockers. Yet downsampling lowers the rate, and the aliasing blockers are attenuated by the cascade action of the sampler and decimation filter. This is illustrated with two examples.

G. Examples of Synthesis of Baseband Chain

From Section III-B, we know that a 32 dB variable or programmable gain in the receiver's analog circuits is sufficient to capture a single input channel of 87 dB dynamic range, because the DSP can absorb the rest of the AGC. Suppose that as specified in the GSM standard, a strong unwanted GSM channel of -23 dBm co-exists with a minimum wanted channel of -99 dBm. To reach the required detection SNR, the wanted channel should be amplified by 39 dB before it is digitized. The nearby unwanted channel will also be amplified, but it will overload the ADC, unless an analog filter attenuates it first by 22 dB. This sets one filter specification. There are others. For example, the filter must attenuate out-of-band blockers as large as 0 dBm that might alias on to the wanted channel.

Specifications for variable gain and filtering are developed in the same way for other standards. In this work, we have taken GSM and 802.11g to represent the two extremes: narrowband GSM with well-specified in-band and out-of-band blockers, and wideband 802.11g with no such specifications.

We will show in detail the evolution of the sampler and filter for the 20-MHz-wide 802.11g channel as a series of plots (Fig. 5).

Only four 20-MHz-wide 802.11g channels can populate the ISM band, which spans 2.4–2.48 GHz. The relative levels of these channels are not prescribed. The FCC only limits total radiated power in the ISM band. The nearest well-specified out-of-band blockers are strong cellular channels in the 1.9-GHz CDMA band. To protect the minimum 802.11g channel of -82 dBm from suffering co-channel interference due to aliasing, all cellular channels in a contiguous 20-MHz-wide band at an offset equal to the sampling frequency or its multiples must be suppressed by more than 80 dB. This results in a filter specification with deep nulls at all aliasing frequencies [Fig. 5(a)].

Search by trial and error leads to an initial sample rate of 480 MHz. This high rate widens the sampler's stopband, although not enough for 80 dB attenuation across 20 MHz—that requires a much higher rate [Fig. 5(a)]. A two-pole passive RC filter with a cutoff frequency beyond the channel of interest at zero IF, say, at 20 and 40 MHz, bolsters attenuation around the null at 480 MHz by another 50 dB to meet the specification [Fig. 5(b)]. *RC* poles are chosen high enough so that even with high process–voltage–temperature (PVT) variations in-channel droop remains negligible. However, the filter violates the specification at offsets of 200–300 MHz. This region lies at roughly

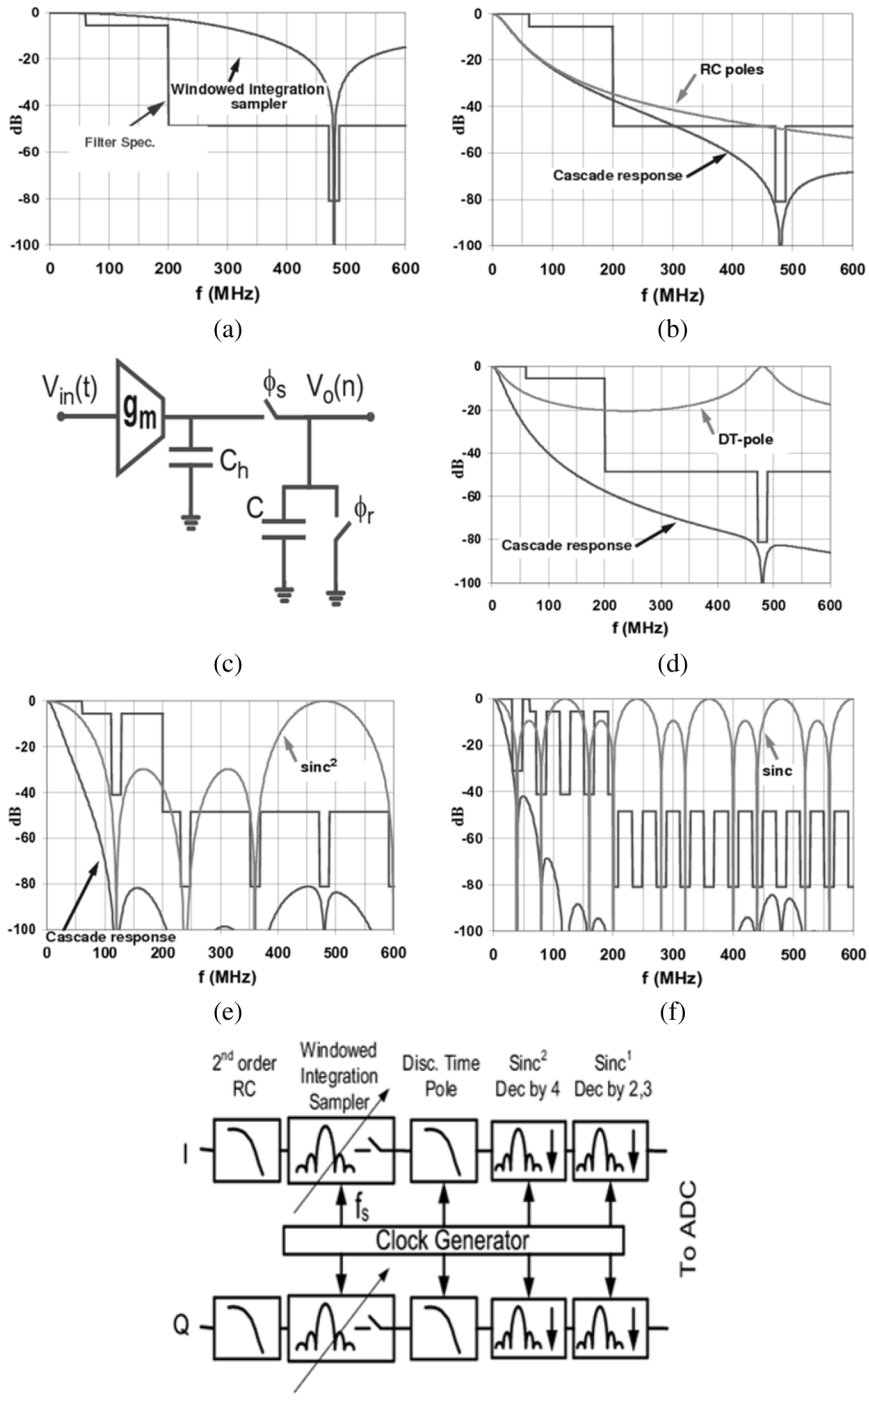

Fig. 5. Evolution of the sampler and filter for 802.11g mode.

half the sample rate, and coincides with the Nyquist frequency at the sampler output.

We note that the continuously connected parasitic capacitance at the output of the transconductor, in parallel with the resistance of the periodically switched sampling capacitor, forms a discrete-time pole [Fig. 5(c)]. Another capacitor can be added to the output of the transconductor to adjust the pole frequency. As this is a discrete-time pole, its magnitude response is uniquely defined up to input frequencies equal to the Nyquist rate, followed by images around the clock frequency and its multiples. Thus, it has minima in its magnitude characteristic at the Nyquist rate and odd multiples. With a pole frequency of 13 MHz, the filter minimum at 250 MHz is at -20 dB. This pole is placed precisely, independent of PVT variations. In cascade with the sampler frequency response, this meets the filter specification everywhere [Fig. 5(d)].

But our task is not yet complete, for 480 MHz is too high a rate for digitization by a low-power ADC. The rate must be lowered by a decimation filter. With each downsampling by M, the number of anti-aliasing notches in the specification also multi-

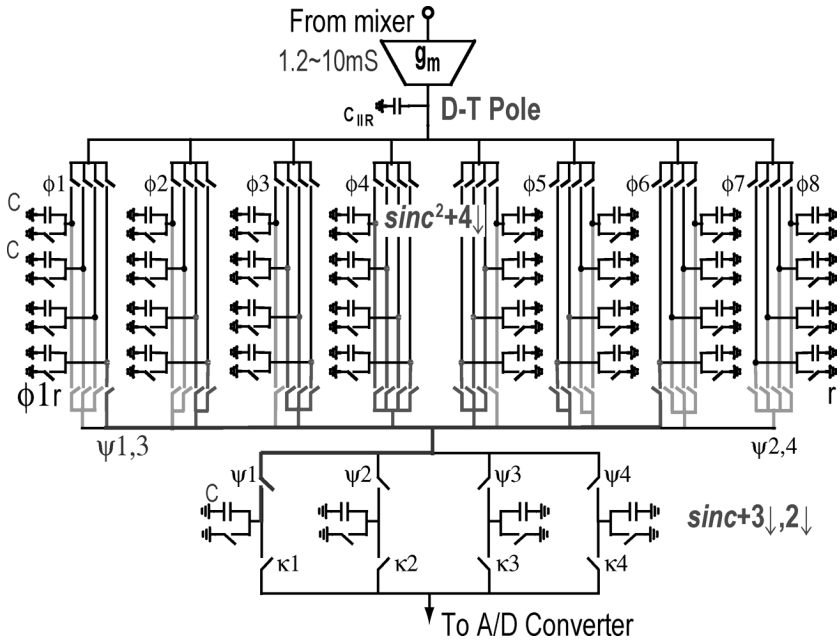

Fig. 6. Realization of anti-aliasing sampler and subsequent decimation stages.

plies by M. A decimation filter of the right order to fulfill the specification is again found by trial and error. In general, the larger the downsampling factor M, the higher the order of the required filter. It is cumbersome in practice to realize a filter of order higher than 2. For 802.11g, then, with a sinc^2 decimation filter the downsampling is limited to $4\times$ [Fig. 5(e)]. Now the sample rate is 120 MHz, but this is still too high. The filter's output must be downsampled again, and this time it can be downsampled by $3\times$ using a first-order sinc decimation filter [Fig. 5(f)]. The final sample rate is 40 MHz, which can be easily digitized at 9-bit resolution by a Nyquist ADC consuming less than 10 mW [Fig. 5(g)].

Using the same design process, GSM can be received in the presence of 0 dBm blockers (as high as -53 dBm/Hz blocker power spectral density) with the same hardware described above, but with different clock rates and a slight reconfiguration, as follows. An initial sampling by windowed integration at 72 MHz is followed by downsampling by $4\times$ with sinc^2 filtering, and downsampling by $2\times$ with sinc filtering. The output sample rate is 9 MHz, which is very reasonable for a milliwatt delta-sigma ADC that resolves 14 bits in a 100-kHz bandwidth. The two-pole passive RC filter's cutoff frequencies are about 550 kHz and 1.1 MHz.

The hardware arrangement is versatile enough to handle most wireless standards, except when the adjacent channel is very strong. Programming is with clock frequency and decimation factor, not capacitance or resistance. We have reached our goal of modeling an analog front-end on a digital receiver.

Ultimately, what is most appealing about the SDR baseband filter is its very simple circuit. Sampling, decimation, and filtering all together require a single transconductor, followed by switches and capacitors only.

The actual realization is differential, but a single-ended half-circuit is shown for clarity in Fig. 6. As sampling is to be followed by decimation, the transconductor integrates on four equal subcapacitors in parallel, repeated in eight time-interleaved channels. The impulse response of the sinc^2 filter is an isosceles triangle whose base spans two *decimated* sample periods. The input is accumulated over seven input clock periods across seven channels, and in the eighth period it is read out to the next stage. During the single clock period it takes for readout, the next input sample is integrated on the spare eighth channel.

It turns out that eight channels, with four unit capacitors per channel, is the minimum necessary for the desired filter and decimation. The seven weights corresponding to samples of the isosceles triangle impulse response are 1:2:3:4:3:2:1. Thus, convolution consists of combining one capacitor from channel 1, two capacitors from channel 2, and so on. Furthermore, as the two impulse responses overlap, continuous throughput requires that the next output sample must start being weighted by 1:2:3, while the previous sample is being weighted 3:2:1. This is readily done with the unused of the four capacitors in each channel: one unit capacitor carrying a certain input sample is available to contribute to the next output, while its three companion capacitors supply the previous output, and so on. Before the next windowed integration, the capacitors in that channel are reset. Thus, the process repeats.

The second decimation is easier. The output of the sinc^2 filter is taken successively on four channels each with an equal capacitor. Then two or three of these capacitors are shorted in parallel to obtain the required decimation of two or three, respectively, at the same time realizing a sinc filter.

This arrangement needs a multi-rate, and furthermore a multiphase, clock. The clock is generated with a ring counter driving combinational logic shown in Fig. 7. Clocks for decimation by 4 and then 3 need a 96-stage ring counter, whereas decimation by 4 and then 2 needs only 32 stages. These two counters and their associated combinational logic are implemented on-chip, and one clock generator is active at a time.

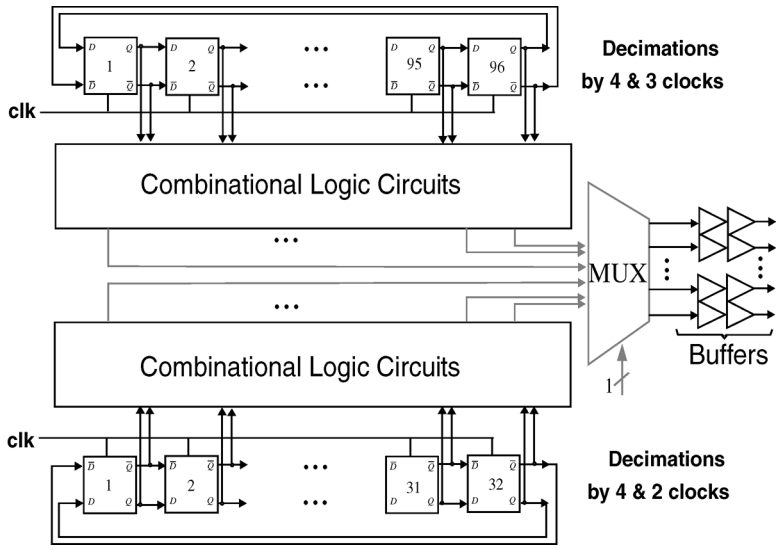

Fig. 7. Clock generator of the anti-alias and decimation filters.

The fixed capacitor $C_{\rm IIR}$ at the output of the transconductor in combination with the switched capacitors realizes the discrete-time pole (Fig. 6). With this capacitor continuously present, the integration window of the initial sampling is determined by the time difference between when one channel is disconnected to when the next channel is disconnected. It does not matter when the channel capacitors are turned on, as long as there is enough time to fully share charge with $C_{\rm IIR}$ before they turn off.

IV. Wideband LNA and Mixer

The SDR receiver needs a wideband LNA that gives relatively uniform gain and input impedance close to 50 Ω from 800 MHz to 6 GHz. The LNA and mixer must handle the full dynamic range of the wideband spectrum incident on the antenna, without circuit distortion corrupting the wanted signal. The LNA must be wideband at input and output ports, and the mixer should be wideband at its input and local oscillator (LO) ports. The output of the mixer is narrowband low-pass, because the wanted channel is at zero IF.

A. Amplifier

We briefly describe the evolution of this wideband low-noise amplifier. The methods employed are quite different than the standard narrowband circuit techniques in RFIC design. The starting point is an exploration of circuits that readily give a wideband impedance match. In this respect, the common gate (C-G) MOS amplifier excels, because biasing it so that its $g_m = 1/50\ \Omega$ gives a well-controlled resistance at the input port. The input capacitance C_{GS}, as well as parasitic capacitances at the input port, disturb impedance match at high frequencies. We have shown in related work [13] that when a wideband termination resistor is present, the capacitances can be embedded into a low-pass LC-ladder network by inserting inductors. In this way, the typically parabolic s_{11} frequency response of tuned amplifiers can be flattened dramatically, giving very wideband impedance match.

However, as has been discussed in the literature at length, the noise figure of a common-gate amplifier is usually higher than 3 dB. For a wideband LNA, the load impedance has to be relatively low, and this results in a noise figure of higher than 5 dB, which is too high for most receivers. The common-gate amplifier must evolve further.

1) Noise Cancellation: Before inventing feedback as a way to lower amplifier distortion, Black was investigating methods of distortion cancellation, or compensation, using feedforward. These ideas were further developed by McMillan [22]. The idea was to sense the distortion as an internally generated signal, amplify it through a parallel path, and cancel it at the output. This concept never caught on because of the large overhead of an auxiliary amplifier [14].

However, the concept finds new life in the form of noise cancellation to lower the noise of an LNA [15]. When a single noise source stimulates two amplifiers whose outputs are then added or subtracted, and if it experiences the same gain through both, then it can be cancelled by addition (or subtraction). What is more, if the *signal* couples into the two amplifiers with opposite polarities to the noise, it reinforces at the two outputs. This arrangement cancels a single source of noise, but reinforces the signal.

Our wideband LNA (Fig. 8) consists of a non-inverting C-G stage, comprising a FET with transconductance g_{M1} and load resistance R_{L1}, in parallel with an inverting common-source (C-S) amplifier with g_{M2} and R_{L2}. We want to cancel the noise of the C-G FET, which sets the input resistance. This is because noise associated with g_{M1} that matches impedance is too large. The circuit, of course, does not become noiseless, because now it is noise of the auxiliary amplifier that dominates. The benefit is that noise and impedance matching arise from two different devices, and are decoupled. The noise can, in principle, be lowered arbitrarily without affecting impedance match.

Part of the C-G FET's noise current flows through the FET itself, and the remainder through the resistors at the source R_S and drain R_{L1}. The voltage it induces on R_S drives the C-S

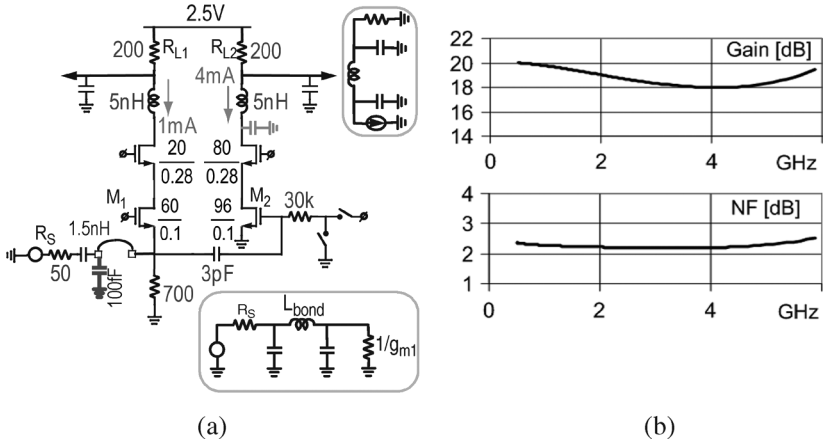

Fig. 8. Wideband LNA. (a) Circuit and *LC* ladder bandwidth extension. (b) Gain and noise figure.

stage. The resulting voltage swings on R_{L1} and R_{L2} are in the same polarity, and in fact are equal in magnitude when

$$R_{L1} = R_S \times g_{M2} R_{L2}. \tag{7}$$

Now if the output is sensed differentially, this noise will appear common-mode and is canceled.

Meanwhile, after attenuation by the amplifier input resistance, the signal source voltage v_S is amplified in-phase at R_{L1} by a gain $g_{M1}R_{L1}$ and through the C-S stage by a gain $-g_{M2}R_{L2}$. Therefore, when noise is cancelled exactly, *and* when the input resistance exactly matches R_S, the amplified signal output is balanced. Compared to the single-ended output of the C-G amplifier alone, the gain is doubled.

It is not for the higher gain alone that this effort is warranted; it is for the fact that if the noise is now determined by the C-S amplifier, g_{M2} can be enlarged to bring down its noise to arbitrarily low levels, as long as R_{L2} is adjusted for balance.

As the C-G and C-S input nodes bias at different voltages, they are capacitively coupled. Cascode FETs are inserted for better isolation, but also to enable switched dumping of signal currents to lower the gain.

This LNA topology serves as a very useful single-ended to differential converter, that takes the single-ended input from the antenna and drives the differential input of the next mixer stage. However, because the circuit topology is not itself balanced, it does not suppress second-order nonlinearity, as fully balanced circuits do. Nevertheless, there can be cancellations of distortion arising in the C-G and C-S amplifiers, which was the reason that cancelling topologies were invented in the first place.

2) Embedding in LC Ladder Filters: With a desired upper cutoff frequency of 6 GHz, FET and parasitic capacitances, as well as parasitic inductances due to bondwires, can be a major limitation.

Instead of trying to minimize these parasitics which soon reaches a point of diminishing returns for the effort involved, it is more productive to embed them into a well-understood network with desirable properties. In this case, as the signal enters the amplifier, it sees, first, the lead parasitic capacitance to ground, then a series bondwire inductance, then the pad capacitance to ground in parallel with C_{GS} of the C-G and C-S FETs. Terminating this is the resistance $1/g_{M1}$. In this configuration, the elements resemble a doubly-terminated all-pole third-order LC ladder filter. By adjusting the component values in simulation, one can approximate a maximally-flat transfer function. This embedding can improve the bandwidth of acceptable impedance match, that is when $s_{11} < -10$ dB by as much as two octaves. It also removes droop in the transfer function from v_S to v_{gs1} and v_{gs2}. The desired effect is obtained in this amplifier by adjusting the FET sizes, thus their C_{GS}, and effective gate bias voltages.

The same concept can be applied at the load resistor, which drives the input capacitance of the mixer. A broadbanding effect is obtained by inserting 5-nH on-chip spiral inductors strategically at the drains of the cascode FETs to form a singly-terminated third-order maximally flat *LC* ladder low-pass filter. Again, this extends the useful bandwidth significantly.

B. Mixer

No known FET mixer could meet the SDR receiver's needs, in terms of low 1/f noise and high IIP2. Much like the LNA, several candidate circuit topologies were compared before arriving at the mixer finally used.

1) Virtues of the Passive FET Mixer: Recently, we have investigated the noise mechanisms in the passive FET mixer [16]. The study was based on a differential input voltage source commutated into resistor loads by a double-balanced arrangement of FET switches. What we learned was that the white noise in the mixer was simply that due to the channel resistance of the FETs that are on at a given time, while flicker noise at the mixer output arose from random duty cycle modulation of the commutating waveform. This is readily seen when FET flicker noise is modelled as a fluctuating voltage in series with the gate. It is flicker noise at the mixer output that is particularly troublesome at zero IF, and in the passive mixer its magnitude is proportional to the RMS input voltage.

Given that the RF input voltage swing in a receive mixer is very small, the passive mixer will show much lower 1/f noise than an active mixer which commutates a large bias current. So our attention focused on ways to improve the passive mixer further.

2) The Right Topology: The input voltage-dependent 1/f noise at the mixer output can be suppressed if the input voltage

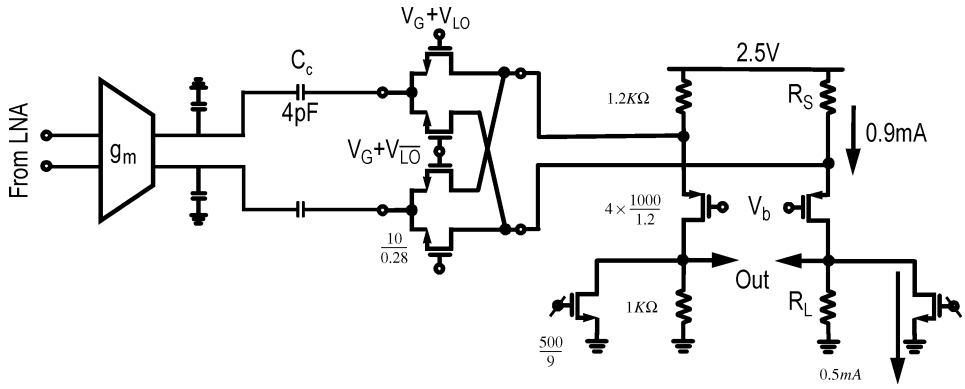

Fig. 9. High linearity and low 1/f noise current driven passive mixer.

swing is lowered to negligible levels. A residual 1/f noise will remain because of random duty cycle variation of the commutating waveform when the FET gate is driven by an LO waveform with finite rise and fall time. As the edges of the LO waveform become sharper, the output 1/f noise will be lower.

It is possible in principle to operate a mixer without the FET switches seeing *any* voltage swing. This is because the analog switch can equally well commutate voltage *or* current. In the study of passive mixer noise, voltage commutation is assumed. That is, the periodically switched FETs commutate a signal voltage source into a buffer with high input impedance. We have since realized that the passive MOSFET mixer is at its best when commutating currents—and then only, signal currents. This requires embedding the FET switches into the dual circuit: that is, a current source input, and a buffer at the output with zero input resistance (Fig. 9).

In practice, a transconductor gives a current source-like drive, and a transresistance feedback amplifier, or more straightforwardly a common-gate FET amplifier with large g_{m0}, can terminate the mixer in a low resistance. As the switches in a double-balanced mixer connect one input to one output at all times, the bias voltage forces itself on the transconductor output. If the input transconductance is g_{mi}, then the resulting signal voltage swing at the source and drain of the commutating FETs is $(g_{mi}/g_{m0}) \times v_{RF}$, where the v_{RF} is the voltage applied by the LNA output to the mixer. By choosing this ratio to be $\ll 1$ in magnitude, the already small RF voltage is attenuated. However, the mixer's conversion gain G_{mix} can still be healthy, depending on the choice of the output buffer's load R_L:

$$G_{\text{mix}} = \frac{2}{\pi} g_{mi} R_L. \tag{8}$$

It is assumed that in the pure current commutating mixer, only signal-bearing RF current and no bias current flows through the switches. This means that the input transconductor should provide an internal path for bias current to flow from supply to ground.

Nor is the pure current-commutating mixer free of flicker noise. 1/f noise originating in the commutating switches still appears at the output in an amount inversely proportional to the slope of the gate voltage waveform at LO transitions. This can be seen by modelling the flicker noise as a slowly varying offset voltage in series with the gate. When subject to a perfectly periodic LO waveform with slope of magnitude S at the instant the FET switches turn on, flicker noise v_n modulates the commutation window by

$$\Delta T = v_n / S. \tag{9}$$

Thus, the signal current is commutated by a low-frequency randomly varying pulse-width modulation. As our previous work shows, at the mixer output this noise appears superimposed on the wanted channel at baseband, and its magnitude is proportional to the signal. In other words, if no signal current is being commutated, the 1/f noise in the FET switches does not appear at the mixer output.

1/f noise in the common-gate FET amplifier would not appear at the output, if it is driven by an ideal signal current source. However, a finite bias resistor R_S is present at the source. Now the input-referred noise voltage of this FET scaled by $1/R_S$ adds to the wanted signal current. The bias current, and R_S, cannot be changed much. But 1/f noise in the FET can be lowered by scaling up both W and L, because the flicker noise depends inversely on the FET's gate area. One consequence of enlarging L is that the f_T of the FET drops, but this does not matter because the wanted signal is now at zero IF. The polysilicon resistors R_L and R_S contribute no flicker noise.

3) Second-Order Nonlinearity: Second-order distortion in the baseband sections makes a zero-IF receiver vulnerable to AM detection of every unwanted channel. In a later section we will specify the magnitude of tolerable IIP2. For now, we consider the physical mechanisms.

The mixer load is a narrowband passive *RC* low-pass filter, whose cutoff frequency roughly tracks twice the channel bandwidth. So, to the first order, we will assume that this filter will attenuate the troublesome unwanted AM channels. FET switches, and the following common-gate FET, process the wanted signal at baseband, and here the vulnerability to AM detection remains.

Both these circuits are differential, so ideally second-order distortion will produce common-mode outcomes that are rejected by subsequent differential stages. However, offsets and mismatch between the two halves of any differential circuit cause common-mode to convert into differential mode.

We reason qualitatively why the IIP2 should be very high in this mixer; in fact, it should be the highest possible for a

wideband circuit where no tuned circuits are used [17]. In the mixer, threshold voltage offset and mismatch between the four FETs contribute imbalance. These effects can be modelled by an equivalent input voltage source at the gate of one in each pair of otherwise identical FETs. This is, in fact, the location of the equivalent input flicker noise voltage source; and we have shown that as the signal current becomes small, the flicker noise contributed by the mixer FETs disappears. Therefore, by extension, this also applies to static offset between FETs.

Similarly, when gate area of the two common gate FETs is enlarged to the limit to lower flicker noise, their matching also improves. This reduces any differential second-order distortion they may cause.

C. Linearity in SDR Front-End

Upon entering the receiver, far away channels are filtered for the first time at the mixer load. Second- and third-order nonlinearity in the LNA and mixer can cause unwanted channels to corrupt the desired channel through AM detection and cross-modulation, respectively. There is the additional problem of downconversion by harmonics of the LO. In this section, we define these problems and explain how they specify the upper limits on acceptable nonlinearity.

We will assume that the input-output characteristic of a circuit block is defined by

$$v_{\rm out} = \alpha_1 v_{\rm in} + \alpha_2 v_{\rm in}^2 + \alpha_3 v_{\rm in}^3 \tag{10}$$

and in terms of these variables, the input-referred intercept points in power are given by

$$\text{IIP2} = \frac{1}{2}\left(\frac{\alpha_1}{\alpha_2}\right)^2, \quad \text{IIP3} = \frac{2}{3}\left|\frac{\alpha_1}{\alpha_3}\right|. \tag{11}$$

1) Limits to AM Detection: The problem of AM detection at zero IF by second-order nonlinearity in baseband sections is well known. More precisely, the modulated signal

$$x(t) = a(t)cos\left(\omega t + \phi(t)\right) \tag{12}$$

when subject to second-order distortion produces a baseband input-referred voltage of

$$X_{bb}(t) = \frac{\alpha_2}{2\alpha_1} a^2(t) \tag{13}$$

whose power is given by

$$P_{Xbb} = \left(\frac{\alpha_2}{2\alpha_1}\right)^2 E(a^4) = \frac{1}{8\text{IIP2}} E(a^4) \tag{14}$$

where $E()$ is the expected value of a random variable.

In the case of SDR, the unwanted signal may occupy a different bandwidth (BW_{UW}) than the bandwidth of the wanted signal (BW). Then, using (15) we can specify the needed IIP2 so that detection of the AM content of unwanted signal with power of P_{UW} leads to a certain SNR due to co-channel interference. It is assumed that $\text{BW} \leq \text{BW}_{\text{UW}}$.

$$\text{IIP2} = 2P_{\text{UW}} - (P_W - \text{SNR} - M - 10\log(\text{BW}/\text{BW}_{\text{UW}}) + 4 + R_2) \tag{15}$$

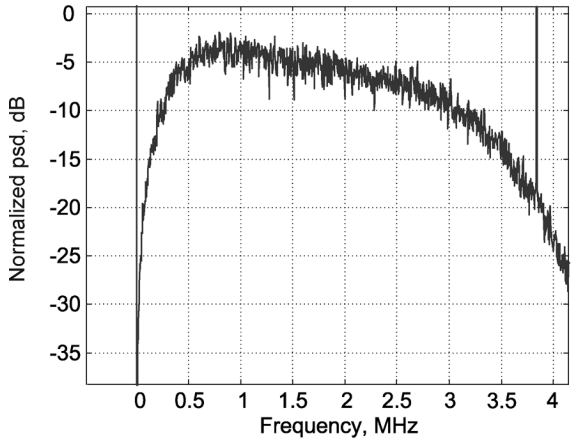

(a)

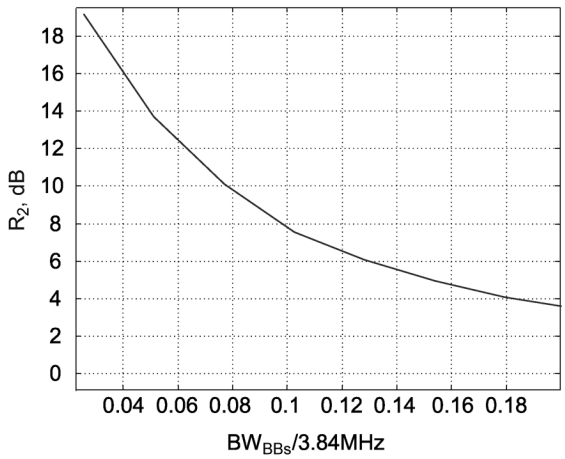

(b)

Fig. 10. (a) Simulated normalized PSD of AM distortion of WCDMA channel. (b) R_2 for WCDM 6DPDCH interference versus normalized bandwidth of a wanted channel with a flat PSD.

where R_2 is a relaxation factor. This factor depends on the statistical properties of the wanted and unwanted signals, and it also captures the relationship between the AM power in the unwanted modulated signal to the AM power in bandlimited white Gaussian noise of an equal power. For example, if the unwanted signal itself resembles bandlimited Gaussian noise (BGN), then R_2 is about -1 for a wanted signal with a flat power spectral density (PSD) across its bandwidth. M is a safety margin.

A -15 dBm unwanted 802.11g channel entering the receiver is, because of a large number of OFDM subcarriers, well-approximated by BGN. To receive a -99 dBm GSM signal in its presence with $M = 6$ dB, the receiver IIP2 must be at least $+61$ dBm. Use of matched filter in demodulator can improve the SNR, but is not considered here.

On the other hand, a WCDMA modulated signal does not resemble BGN. Simulations show that AM detection of a single WCDMA channel with 6 equal gain dedicated physical data channel (DPDCH) has lower energy around DC [Fig. 10(a)] and this leads to an $R_2 = 19$ dB for a wanted GSM signal occupying ± 100 kHz around DC [Fig. 10(b)]. If the receiver IIP2 is better than $+48$ dBm, a wanted GSM channel at -99 dBm can be demodulated with a margin of 6 dB in the presence of a -15 dBm WCDMA unwanted signal.

If the same unwanted signal power is concentrated into a narrower bandwidth, it will create worse co-channel interference.

2) Limits to Cross Modulation: Cross-modulation is when an unwanted AM signal, through third-order nonlinearity, causes spurious AM on a wanted signal. This effect is described by the following analysis, using the same notation as in the previous section.

Consider two input signals to a nonlinear receiver at different frequencies, specified in the form of (10). The output at the wanted signal's frequency, ω_1, is given by

$$y(t) = \alpha_1 a_1(t)\left(1 + \frac{3}{4}\frac{\alpha_3}{\alpha_1}\left(a_1^2(t) + 2a_2^2(t)\right)\right) \times \cos\left(\omega_1 t + \phi_1(t)\right) \simeq \alpha_1 a_1(t)\left(1 + \frac{3}{2}\frac{\alpha_3}{\alpha_1}a_2^2(t)\right) \times \cos\left(\omega_1 t + \phi_1(t)\right) \quad (16)$$

where the unwanted signal amplitude $a_2(t)$ is assumed to be much greater than the wanted signal $a_1(t)$.

The unwanted signal thus modulates the envelope of the wanted signal. We can say that the unwanted signal causes a change in average gain, that is, it desensitizes the wanted signal and imposes envelope fluctuation. The average gain and the fluctuation component are

$$\alpha_1' = \alpha_1 + 3\alpha_3 P_2$$
$$d(t) = 3\alpha_3\left(\frac{a_2^2(t)}{2} - P_2\right)a_1(t)\cos\left(\omega_1 t + \phi_1(t)\right). \quad (17)$$

P_2 is the power of the unwanted signal, and $d(t)$ is the distortion caused at the wanted signal's frequency by cross modulation. From (17) it can be noted that for a constant envelope unwanted signal, $d(t)$ has zero power.

Equation (17) can be translated into a specification on IIP3 of the receiver.

$$\mathrm{IIP3} = P_2 - 5\log\left(\frac{\max(\mathrm{BW}_1, \mathrm{BW}_2)}{\mathrm{BW}_1}\right) - R_3 + \frac{\mathrm{SNR} + M}{2} + 3.5. \quad (18)$$

R_3 is the relaxation factor in this case, similar to R_2 in the previous section. $R_3 \simeq 0$ for BGN interferences. It should be pointed out that the IIP3 specification is independent of the strength of the wanted channel; that is, the corruption caused by cross-modulation scales with the wanted signal, and in relative terms stays constant.

Using this expression, we can see that for a -99 dBm GSM channel to be received in the presence of an unwanted 802.11g channel of -15 dBm, the receiver IIP3 must be -14 dBm. This assumes a 6 dB margin. Similarly, the same GSM signal can be detected in the presence of a -15 dBm WCDMA unwanted signal in a receiver with IIP3 of -18 dBm. This assumes $R_3 = 16$ dB, obtained from simulation.

In addition, we must consider the impact on noise figure of gain desensitization due to the presence of the unwanted signal. Desensitized gain in the LNA causes the noise of subsequent circuits to raise overall receiver noise figure. On this basis, we find that to receive a -99 dBm GSM channel in the presence of a -15 dBm unwanted channel, a receiver IIP3 of -5 dBm is required. This requirement on IIP3 is higher than the one based on cross-modulation; that is, in this case gain desensitization is the primary vulnerability.

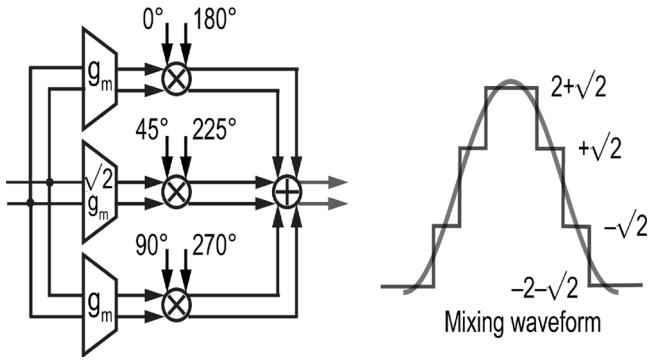

Fig. 11. Harmonic-suppressing mixer block diagram and its sinusoidal approximated mixing waveform.

D. Harmonic Downconversion

Another problem associated with a wideband receiver is that when the LO tunes a channel at the lower end of the receiver passband, say, around 900 MHz, the square wave commutation mixes unwanted channels at the third and fifth harmonic of the LO which also lie in the passband. These channels also appear at zero IF. This effect can only be suppressed by somehow "linearizing" the mixing action. A hard switching mixer is always preferred because it gives the best conversion gain. The harmonic content associated with hard switching can be lowered by shaping the commutation square wave into a step-wise waveform resembling a sinewave; more precisely, a waveform that corresponds to samples of a sinewave. This is obtained by constructing a mixer from weighted transconducters in the ratio $1{:}\sqrt{2}{:}1$, whose output currents are switched by LO waveforms delayed by 1/8th of the period, and added together (Fig. 11). Theoretically, this commutation waveform has no third or fifth harmonic [18].

However, in the presence of gain mismatch and phase error in three different paths, the harmonic suppression ratios are limited to

$$\sigma_{H_3}/H_1 = \sqrt{(\sigma_\phi/2)^2 + (\sigma_{g_m}/6g_m)^2} \quad (19)$$

$$\sigma_{H_5}/H_1 = \sqrt{(\sigma_\phi/2)^2 + (\sigma_{g_m}/10g_m)^2} \quad (20)$$

where σ_ϕ is variance of LO phases and σ_{g_m} is variance in g_m values (Fig. 11).

V. Frequency Synthesizer

The SDR receiver needs a wide-tuning-range frequency synthesizer. The large tuning range should not be at the expense of phase noise, which must meet specifications for all the standards that it spans. It is difficult to imagine how, without a great deal of switchable adaptation, a single oscillator can span 800 MHz to 6 GHz. The most straightforward solution is choosing one from a collection of voltage-controlled oscillators (VCOs) tuned to separate bands, each optimized for phase noise.

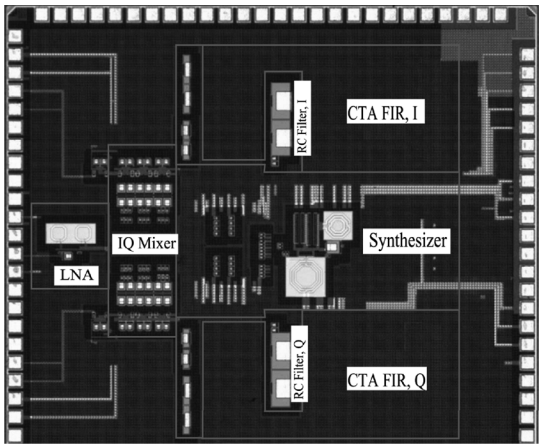

Fig. 12. Die photo.

It is seldom a good idea to turn on two VCOs operating at different frequencies, because even small amounts of parasitic on-chip coupling can cause the two to pull each other in frequency. Nor, in general, is it a good idea to create sum or difference frequencies in a single-sideband (SSB) mixer by combining the VCO output frequency with itself after division. This is because imperfections in the SSB mixer such as mismatch in gain or inaccuracy in quadrature phases create spurious tones well above phase noise levels.

The on-chip local oscillator comprises two VCOs, one centered at 3640 MHz and the other at 5200 MHz. After frequency division by 2, the first VCO's range covers 1.65 to 2.1 GHz with quadrature phases, and after division by 4, it covers 820-1050 MHz. It is in this range of LO frequencies that the harmonic-suppressed mixer must be used. The divide-by-4 naturally gives the phases required to suppress harmonics, -45°, 0°, $+45^\circ$ for the I channel, and $+45^\circ$, $+90^\circ$, $+135^\circ$ for the Q channel.

The second VCO is a quadrature oscillator. It tunes from 4.7 to 5.4 GHz, and after division by 2, it covers 2.35 to 2.7 GHz.

This two VCO arrangement covers all important bands. If continuous coverage from 800 MHz to 6 GHz is required, for instance for a cognitive radio receiver that can tune to all used and vacant frequency bands, then it can be shown that three VCOs are sufficient.

VI. Measurements

The prototype receiver consists of the wideband LNA, the harmonic-suppressing mixer, the tunable anti-aliasing filter with two stages of decimation, and the local oscillator circuits. It is fabricated in the STMicroelectronics 90-nm CMOS process. Active area is 3.8 mm^2, dominated by baseband filter capacitance (Fig. 12). The current consumption of different blocks is summarized in Table I.

The only wideband characterization possible is to measure s_{11} of the LNA across the full band. Although impedance matching is important in itself, this is an indirect way to verify the amplifier's wideband input–output characteristic. An earlier version of this circuit realized in 130-nm CMOS was used to evaluate the standalone amplifier's frequency response, and proved the concept [19].

TABLE I
Power Consumption

Power Consumption			
Circuit		Supply(V)	Current(mA)
LNA		2.5	5
Mixer input gm		2.5	4.6
Mixer output stage		2.5	1.8
Filter	high gain GSM	1	23
	low gain GSM	1	8
	high gain 802.11g	1	28
	low gain 802.11g	1	13
3.8 GHz VCO		1	12.5
Div-by-4		1	6
Div-by-2		1	3

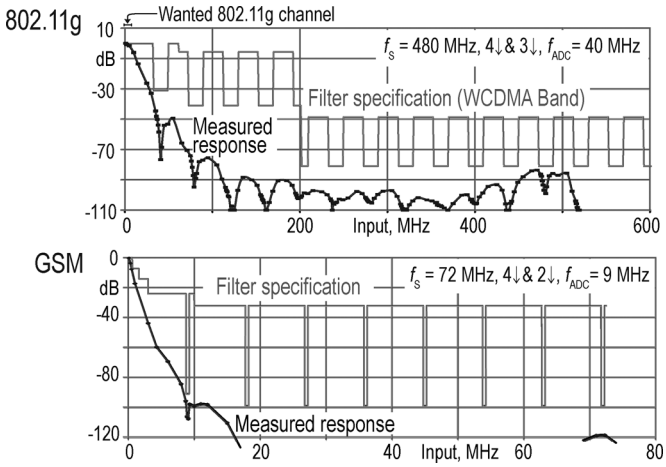

Fig. 13. Measured selectivity of the receiver for GSM and 802.11g modes.

Measurement proceeds by applying single-tone RF inputs at frequencies from 800 MHz to 5 GHz, tuning the LO close to the input frequency, and measuring the low-frequency discrete-time analog filter output. Measured filter response is plotted for 802.11g and GSM against specifications, which show the deep anti-aliasing notches (Fig. 13). The passive RC mixer load was switched between poles at 20 and 40 MHz for 802.11g and 550 and 1100 kHz for GSM reception. The specifications are met with margin, clearly showing that the receiver can operate without an RF prefilter.

Any wireless receiver is vulnerable to spurious response, usually in the mixer. In the SDR receiver, the clocked baseband filter can also contribute spurious responses. Spurious responses were characterized for all anticipated input frequencies. For example, the response to inputs lying in the front-end passband at the third and fifth harmonic of the LO frequencies is measured at baseband. This is compared to the response to an input at the fundamental. The measured harmonic suppression is -38 and -40 dB, respectively, indicating good LO phase accuracy and gain matching in the three paths comprising each mixer.

TABLE II
COMPARISON OF THIS SDR-RX IN GSM AND 802.11G MODE WITH PRIOR PUBLISHED RESULTS

	GSM		**802.11g**	
	This work	ref[20]	This work	ref[21]
Chip NF/System NF[a] [dB]	5/5	3/5.5	5.5/5.5	4.5/6.5
IIP3 [dBm]	-3.5[b]	-13[c]	-4[b]	5
IIP2 [dBm]	$+60$[b]	+37$\sim$+51[c]	$+60$[b]	-

[a]RF preselect filter noise contribution is added to NF. This work doesn't use such filter.
[b]Mid-gain setting for RX, to meet 3dB desensitized RX test in presence of blockers.
[c]Gain setting is not specified.

Similarly, the spurious response of the baseband filter to down-converted inputs at half the final decimated clock frequency is measured, and found to be -60 and -74 dB, respectively, for 802.11g and GSM configurations.

Table II compares this receiver when programmed for GSM and 802.11g versus recent radios which are designed specifically for each standard. The cascade noise figure of the receiver was measured from output SNR. In high gain mode, the 50-Ω noise figure is about 5 dB at 900 MHz, and rises by only 0.5 dB at 2.5 GHz. This is another indirect measure of the wideband nature of the front-end. The choice of mixer topology and the use of FETs with large gate area after the mixer prove themselves because it is very hard to detect any flicker noise above 10 kHz at baseband; that is, the direct conversion receiver's 1/f noise corner is below 10 kHz. This is good enough to receive narrowband channels such as GSM without significant impact on SNR.

Nonlinearity in the front-end matters most when a weak desired channel is being received in the presence of a strong unwanted channel. In a well-engineered receiver, this condition is sensed and the front-end gain is lowered. Therefore, linearity is usually specified at medium or low gains. The receiver IIP3 at medium gain is -3.5 dBm, and its IIP2 is almost $+60$ dBm. The integrating sampler and decimator contributes almost nothing to cascade nonlinearity, because the passive *RC* low-pass load of the mixer attenuates unwanted channels.

In practical terms, these intercept points mean that the 3 dB desensitized minimum detectable 802.11g signal (-62 dBm) can be received in the presence of co-channel interference due to AM detection of a 6DPDCH WCDMA signal as strong as -17.5 dBm, or cross-modulation of a -22.5 dBm signal. As both types of distortion take place simultaneously, the smaller of the two limits applies: in this case, it is due to the third-order nonlinearity.

Similarly, GSM can be received with unwanted 802.11g signals as large as -15 dBm, again limited by AM detection, and -13.5 dBm limited by 1-dB compression. In this case, the effect of second-order nonlinearity dominates.

VII. CONCLUSION

A low-power software-defined radio receiver capable of tuning over 800-MHz–6-GHz frequency band is described. A windowed integration sampler samples at high rate and subsequent discrete-time decimation filters reduce the sample rate low enough to be digitized by a low-power ADC. The baseband anti-aliasing filters are programmable by the clock to adapt for the desired signal bandwidth and its surrounding interferences. Wideband and low-noise RF front-end characteristic is obtained using bandwidth extension and noise cancellation methods. It is shown that current-driven passive mixer provides low 1/f noise and high IIP2 needed for SDR-RX. The concepts are validated through the measurement results of a 90-nm CMOS prototype that achieves full programmability with low power consumption.

REFERENCES

[1] E. Götz *et al.*, "A quad-band low power single chip direct conversion CMOS transceiver with $\Sigma\Delta$-modulation loop for GSM," in *Proc. ESSCIRC'03*, Sep. 2003, pp. 217–220.
[2] R. Bagheri, A. Mirzaei, S. Chehrazi, M. Heidari, M. Lee, M. Mikhemar, W. Tang, and A. A. Abidi, "An 800 MHz to 5 GHz software-defined radio receiver in 90 nm CMOS," in *IEEE ISSCC 2006 Dig. Tech. Papers*, San Francisco, CA, 2006, pp. 1932–1941, p. 26.6.
[3] J. Mitola, "The software radio architecture," *IEEE Commun. Mag.*, vol. 33, no. 5, pp. 26–38, May 1995.
[4] R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 5, pp. 539–550, Apr. 1999.
[5] *Digital Cellular Telecommunications System (Phase 2+): Radio Transmission and Reception*, 3GPP TS 05.05, version 8.15.0, release 1999, European Telecommunication Standards Inst. (ETSI).
[6] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, Further Higher-Speed Physical Layer Extension in the 2.4 GHz Band*, IEEE Standard 802.11g, 2002.
[7] T. Hentschel and G. Fettweis, "The digital front-end: Bridge between RF and baseband processing," in *Software Radio: Enabling Technologies*, W. Tuttlebee, Ed. New York: Wiley, 2002.
[8] J. Yuan, "A charge sampling mixer with embedded filter function for wireless applications," in *Proc. 2nd Int. Conf. Microwave and Millimeter Wave Technology*, Beijing, China, 2000, pp. 315–318.
[9] Y. Poberezhskiy and G. Poberezhskiy, "Sampling and signal reconstruction circuits performing internal anti-aliasing filtering and their influence on the design of digital receivers and transmitters," *IEEE Trans. Circuits Syst. I*, vol. 51, no. 1, pp. 118–129, Jan. 2004.
[10] R. B. Staszewski *et al.*, "All-digital TX frequency synthesizer and discrete-time receiver for Bluetooth radio in 130-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 39, no. 12, pp. 2278–2291, Dec. 2004.
[11] K. Muhammad *et al.*, "A discrete time quad-band GSM/GPRS receiver in a 90 nm digital CMOS process," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC'05)*, San Jose, CA, Sep. 2005, pp. 804–807.
[12] A. Mirzaei, R. Bagheri, S. Chehrazi, and A. Abidi, "A second-order anti-aliasing prefilter for an SDR receiver," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC'05)*, Sep. 2005, pp. 629–632.
[13] A. Ismail and A. Abidi, "A 3-10-GHz low-noise amplifier with wideband LC-ladder matching network," *IEEE J. Solid-State Circuits*, vol. 39, no. 12, pp. 2269–2277, Dec. 2004.
[14] H. W. Bode, "Feedback—The history of an idea," in *Selected Papers on Mathematical Trends in Control Theory*, R. Bellman and R. Kalaba, Eds. New York: Dover, 1960.

[15] F. Bruccoleri, E. M. Klumperink, and B. Nauta, "Wide-band CMOS low-noise amplifier exploiting thermal noise canceling," *IEEE J. Solid-State Circuits*, vol. 39, no. 2, pp. 275–282, Feb. 2004.
[16] S. Chehrazi, R. Bagheri, and A. Abidi, "Noise in passive FET Mixers: a simple physical model," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC'04)*, Sep. 2004, pp. 375–378.
[17] M. Brandolini, P. Rossi, D. Sanzogni, and F. Svelto, "A +78 dBm IIP2 CMOS direct downconversion mixer for fully integrated UMTS receivers," *IEEE J. Solid-State Circuits*, vol. 41, no. 3, pp. 552–559, Mar. 2006.
[18] J. A. Weldon *et al.*, "A 1.75 GHz highly-integrated narrowband CMOS transmitter with harmonic-rejection mixers," in *IEEE ISSCC 2001 Dig. Tech. Papers*, Feb. 2001, pp. 160–161.
[19] S. Chehrazi, A. Mirzaei, R. Bagheri, and A. Abidi, "A 6.5 GHz wide-band CMOS low noise amplifier for multi-band use," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC'05)*, Sep. 2005, pp. 796–799.
[20] E. Song, Y. Koo, Y. J. Jung, D. H. Lee, S. Chu, and S. I. Chae, "A 0.25 μm CMOS quad-band GSM RF transceiver using an efficient LO frequency plan," *IEEE J. Solid-State Circuits*, vol. 40, no. 5, pp. 1094–1106, May 2005.
[21] M. Zargari *et al.*, "A single-chip dual-band tri-mode CMOS transceiver for IEEE 802.11a/b/g wireless LAN," *IEEE J. Solid-State Circuits*, vol. 39, no. 12, pp. 2239–2249, Dec. 2004.
[22] B. McMillan, "Multiple-feedback systems," U.S. Patent 2,748,201, May 29, 1956.

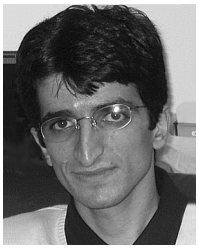

Rahim Bagheri (S'02) received the B.Sc. and M.Sc. degrees in electrical engineering, both with honors, from Sharif University of Technology, Tehran, Iran, in 1997 and 1999, respectively. He is pursuing the Ph.D. degree in electrical engineering at the University of California, Los Angeles (UCLA), working on CMOS wireless transceiver circuits and architectures for reconfigurable radios.

From 1999 to 2000, he was with UCLA MOSFET Research Lab, where he worked on sub-100-nm MOSFET design. He was with Valence Semiconductor Inc., Los Angeles, from 2000 to 2001 as a Design Engineer working on 802.11a CMOS radio. He was with JAALAA Inc., Los Angeles, in 2003 as RFIC team leader developing high speed wireless CMOS chips with proprietary communication schemes. He co-founded WiLinx Inc., Los Angeles, developing CMOS wireless-communication chips for universal-UWB and Bluetooth over UWB systems in 2003. Since then, he has been with WiLinx Inc., as CTO and President. He holds 12 patent applications in the field of RF-CMOS design.

Mr. Bagheri received the Gold Medal in National Physics Olympiad and Honorable Mention Diploma in the XXIV International Physics Olympiad (1993). He is a recipient of the UCLA ICSRR Best Student Presentation Award (2006), Analog Devices Outstanding Student Designer Award (2003), UCLA Graduate Division Fellowship (2000), and UCLA Regent Stipend Fellowship (1999).

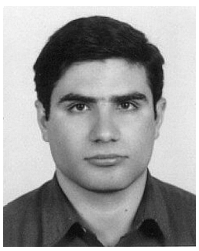

Ahmad Mirzaei (S'02) received the B.Sc. and M.Sc. degrees (with honors) in electrical engineering from Sharif University of Technology, Tehran, Iran, in 2000 and 2002, respectively. He received the Ph.D. degree in electrical engineering from the University of California, Los Angeles, where he was working on reconfigurable CMOS receivers.

Currently, he is with WiLinx Inc., Los Angeles, CA, as the Technical Director, where he is leading the development of the CMOS RFIC for universal UWB systems. His research interests include wideband RF CMOS circuits, programmable filters, accurate quadrature CMOS oscillators, and low-power injection-locked dividers. He holds five patent applications in the field of RF-CMOS design.

Mr. Mirzaei is the recipient of the Silver Medal in National Mathematics Olympiad in 1996 and the Analog Devices Outstanding Student Designer Award in 2003.

Saeed Chehrazi (S'97) received the B.Sc. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 2001, and the M.Sc. degree from the University of California, Los Angeles (UCLA), in 2004, where he is currently working toward the Ph.D. degree.

Since joining UCLA, he has been involved in designing and modeling RF front-end circuits for software-defined radios. During the summer of 2006, he was with WiLinx, Los Angeles, where he worked as an RF/analog circuit designer. His research interests are RF, analog, and mixed-signal circuit design and modeling.

Mohammad E. Heidari received the B.Sc. and M.S. degrees in electrical engineering from Sharif University of Technology, Tehran, Iran, in 1995 and 1997, respectively.

He was with Valence Semiconductor during 2000–2001 as an analog CMOS IC Designer. He is currently working on reconfigurable transmitters for his Ph.D. at the Integrated Circuits and Systems Lab (ICSL), University of California, Los Angeles, and he is also with WiLinx Inc., Los Angeles. His interests are in the area of SDR transceivers, low phase-noise digital PLL/DLL, and over-sampling data converters. He holds five patent applications in the field of RF-CMOS design.

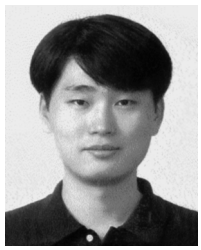

Minjae Lee (S'03) received the B.Sc. and M.S. degrees, both in electrical engineering, from Seoul National University, Seoul, Korea, in 1998 and 2000, respectively.

He was a consultant with GCT Semiconductor, Inc., and Silicon Image Inc., designing analog circuits for wireless communication and digital signal processing blocks for Gigabit Ethernet. He joined Silicon Image Inc. in 2001, developing Serial ATA product. Since 2003, he has been working toward the Ph.D. degree in the field of analog and mixed-mode circuit design at the Integrated Circuits and Systems Lab (ICSL), University of California, Los Angeles.

Mohyee Mikhemar (S'99) received the B.S. and M.S. degrees in electrical engineering from Ain Shams University, Cairo, Egypt, in 2000 and 2004, respectively. Since 2004, he has been with the University of California, Los Angeles, where he is working towards the Ph.D. degree.

His research interests include software-defined radio architectures, wideband RF front-end design, and effects of nonlinearities in wideband receivers.

Wai K. Tang (S'02) received the B.S. degree in electrical engineering from the University of Texas at Austin in 2004, and the M.S. degree from the University of California, Los Angeles, in 2006.

In 2003, he worked as an intern and implemented a jitter measurement system for 10-GHz signals at Silicon Laboratories, Inc., Austin, TX. Currently, he is with Silicon Laboratories, Inc., as a design engineer, where he deals with isolator and gate driver design.

Asad A. Abidi (S'75–M'80–SM'95–F'96) received the B.Sc. (with honors) degree from the Imperial College, London, U.K., in 1976, and the M.S. and Ph.D. degrees in electrical engineering from the University of California, Berkeley, CA, in 1978 and 1981, respectively.

From 1981 to 1984, he was with Bell Laboratories, Murray Hill, NJ, as a Member of Technical Staff at the Advanced LSI Development Laboratory. Since 1985, he has been with the Electrical Engineering Department, University of California, Los Angeles (UCLA), where he is a Professor. He was a Visiting Faculty Researcher at Hewlett Packard Laboratories in 1989. His research interests are in CMOS RF design, high speed analog integrated circuit design, data conversion, and other techniques of analog signal processing.

Dr. Abidi was the Program Secretary for the IEEE International Solid-State Circuits Conference (ISSCC) from 1984 to 1990, and the General Chairman of the Symposium on VLSI Circuits in 1992. He was the Secretary of the IEEE Solid-State Circuits Council from 1990 to 1991. From 1992 to 1995, he was the Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS. He received an IEEE Millennium Medal, the 1988 TRW Award for Innovative Teaching, and the 1997 IEEE Donald G. Fink Award, and is co-recipient of the Best Paper Award at the 1995 European Solid-State Circuits Conference, the Jack Kilby Best Student Paper Award at the 1996 ISSCC, the Jack Raper Award for Outstanding Technology Directions Paper at the 1997 ISSCC, the Design Contest Award at the 1998 Design Automation Conference, an Honorable Mention at the 2000 Design Automation Conference, and the 2001 ISLPED Low Power Design Contest Award. He is a Fellow of the IEEE, and he was named one of the top ten contributors to the ISSCC.