Evolution of a Software-Defined Radio Receiver's RF Front-End

Asad A. Abidi
Electrical Engineering Department
University of California, Los Angeles

Introduction

The software-defined radio has remained elusive so far. In the form envisaged by Mitola [1], all RF and baseband signal processing is digital, enabled by an A/D converter (ADC) with extraordinary specifications. For example, coverage of the frequency band from 800 MHz to 5.5 GHz, where all of today's cellular and WLAN channels lie, requires a 12b, 10 GS/s ADC. Not only is this ADC impossible today, it will remain so in the foreseeable future. This bottleneck can only be overcome by re-examining the underlying paradigm.

What motivates Mitola is the flexibility and programmability of digital signal processing (DSP) in both radio front-end and baseband to adapt to any modulation, channel bandwidth, or carrier frequency. But what he proposes is capable of much more. With a number of digital downconverters, a single ADC of this type enables *simultaneous* reception of *every* channel in the band. This is much more than what is needed in a mobile handset. The most likely scenario is that the user will select one channel at a time, or perhaps two or even three—for example, while an incoming cellular call is relayed to his headset by Bluetooth, the user watches television on the handset or downloads a large file via WLAN in the background.

With this more modest goal, it is sufficient to realize a software-defined receiver that enables reception of *any one* channel at an arbitrary carrier frequency with any modulation. This receiver should be less power hungry, and with the appropriate architecture may even be realizable in current IC technology. A reasonable number of receivers in parallel can enable that many channels to be received simultaneously.

The subject of this paper is how to construct this flexible receiver. We survey the steps that have been taken over the years in this direction—some very tenuously—leading up to a recent realization of a software-defined RF front-end. We do not cover the work that has gone into the digital front-end which follows the ADC [2], nor on the software-defined baseband—both are essential for a full radio. Furthermore, an efficient software-defined transmitter without which the software-defined radio project is incomplete is still in the research phase.

HF SDR Receiver

The military must often receive a large number of channels simultaneously across a band. The UK DERA has developed a prototype for the HF band (3-30 MHz) [3], where many military communications take place. This is the only receiver (Fig.1(a)) we have found of Mitola's type, and in the following description its strengths and its limitations should be clear.

The HF antenna is followed by a 30 MHz lowpass filter, which suppresses signals that are out of the HF band. The purpose is to limit the total dynamic range, but we note its role as an *anti-aliasing filter*. This is followed by a 12b, 75 MHz A/D converter, state-of-the-art for the time, developed for IF conversion with high spur-free dynamic range (SFDR) in wireless basestations[1]. After digitally controlled attenuation to align the input dynamic range with the ADC's full scale, the entire HF band is digitized. Any number of channels can now be scanned simultaneously. It is instructive to follow the signal processing flow through the digital circuits after the ADC (Fig.1(b)), what has been called the digital front-end.

DSP channelizers characterize modulation by a complex number stream at DC. In each channelizer, a complex (quadrature) number-controlled digital oscillator downconverts the channel of interest to zero IF. There is no offset or noise in the digital domain to corrupt the signal samples.

The downconverted sample stream is at the ADC rate. It represents the desired (narrow) channel at zero surrounded by all other channels in the HF band which are not wanted. The *symbol* rate of the desired narrow channel is some small fraction, in general incommensurate, of the *ADC sample* rate. Downstream equalization and demodulation takes place at this symbol rate. Therefore the task of the digital front-end is to decimate the sample rate by this large non-integer factor, suppressing the aliasing inherent to decimation, and filtering down the relative strength of the unwanted channels [2].

It is well known in the DSP art that the cascaded integrator comb is an efficient decimation filter. The filter has a lowpass characteristic described by the *sinc* function up to the *input* Nyquist rate, with nulls at the *output* (decimated) Nyquist rate and multiples. This ensures that after decimation nothing aliases to DC, where the spectral density of the desired signal is at a peak; although small unwanted energy will alias on to the wanted spectrum away from DC. This dictates the filter specifications, which are fulfilled by an appropriate cascade of sections.

This HF software-defined receiver prototype performed almost as well as a conventional analog-tuned triple-conversion superheterodyne. Its main advantage is programmability, and the ability to add any number of digital channelizers for simultaneous reception. Its main limitation stems from the maximum sample rate of the A/D converter. If the highest carrier frequency is limited to half this sample rate, then with the state-of-the-art ADC today, an SDR receiver cannot operate beyond the VHF broadcast FM band.

The Role of Downconversion

This limitation can be overcome if the band of interest is first downconverted, so that it lies within the ADC's Nyquist band. Usually we think of this as the (lowpass) first Nyquist band, but equally well it could be a higher Nyquist band that the ADC undersamples. In either case, the information bandwidth should not exceed the Nyquist band. The antenna RF prefilter is the customary way to bandlimit the receiver input. This is exactly how Toshiba realizes a mainly digital flexible receiver for PDC and DCS (Fig.2) [4]. Switch selectable RF filters 1.5 and 1.9 GHz pass 27 MHz-wide bands, respectively, requiring, if these filters are brickwall, an ADC with a minimum sample rate of 54 MHz.

The preselected bands cover all PCS or DCS channels. Quadrature (complex) analog mixers with fixed LO frequency block downconvert the *band* to DC. A complex ADC (one 12b, 64 MHz ADC in each of the two quadrature branches) digitizes the band, and a number-controlled complex oscillator translates the channel of interest to DC. Then there follow decimation and channelization in the digital domain, as described above.

The limited image rejection of the first complex analog downconverter is overcome by translating the 27 MHz-wide band to one side of DC, so the images on the other side of DC lie in the RF filter's transition or stopband and are attenuated. Gain mismatch in the two ADCs also worsens image rejection, but with all said and done the image is suppressed by 70 dB.

This receiver is one step in the right direction, but as we will see, there is farther to go.

Subsampling or Undersampling

The bandpass sampling theorem says that a signal with an information band situated entirely within any integer translation of the 1st Nyquist band on the frequency axis can, in principle, be digitized by a conventional ADC; provided the front-end sample-and-hold can track

[1] Wireless basestations have always been software-defined radios that receive all channels in the band simultaneously. Radio power dissipation is however not of primary concern.

the signal—that is, the sampling bandwidth should encompass the signal band.

Without limiting the input spectrum to the ADC's Nyquist bandwidth, digitization will suffer from aliasing, even if it is wideband noise that aliases. Consider the case of *two* bandlimited spectra, each lying in possibly different translations of the Nyquist band, sampled at a rate greater than the sum of the two information bandwidths. If the sampling bandwidth encompasses the highest frequency in the two bands, their aliases after sampling will appear in the 1st Nyquist band. In fact, with the right sample rate, they will alias side-by-side without overlap. Now both bands can be digitized at baseband, and channelized for simultaneous reception.

This is the principle that underlies the front-end of a positioning receiver that receives signals from GPS (1575 MHz) and GLONASS (1602 MHz) satellites (Fig.3) [5]. The prefilter is placed after the LNA to attenuate wideband LNA noise that would otherwise accumulate after sampling. The prefilter, in fact, is two custom prefilters that pass these bands. A sampling rate of 24.2 MHz positions the two bands side-by-side in the effective IF range from DC to 12 MHz.

This example is interesting because it illustrates how a single sampling action effects two different frequency translations. We note, though, that the need for RF prefilters limits the receiver's flexibility.

Subsampling with Analog Decimation

It may be that the passband of available RF filters far exceeds the relatively narrow band containing the channels of interest. Then to avoid aliasing, the input must be sampled at a higher than minimum rate. However, it wastes precious power to digitize at a surplus rate, especially when a later digital channelizer will discard much of the spectrum because it is unwanted. In wireless receivers, the ADC power dissipation often falls sharply with the sample rate, because the clock frequency and required dynamic range both become smaller. This suggests that the analog samples should be decimated to a lower rate before they are digitized.

In the same way as DSP-based decimation, analog sample rate conversion narrows in on the desired band while suppressing aliasing that arises from downsampling. Analog decimation is realized by choosing an initial (high) sample rate, some multiple of which is centered on the band of interest. Aliasing translates the band to DC. Then, downsampling takes place through a discrete-time (D-T) analog decimating filter.

How to realize this filter? The most direct way is to sample the input successively into N capacitors, and at the end short all capacitors to redistribute the N samples of charge (Fig.4) [6]. This forms an N-tap FIR filter with unit tap weights that decimates by N, and its transfer function is a *sinc* with nulls at all integer multiples of the output (lower) sample rate.

Sampling with built-in anti-aliasing

All the schemes described so far need an anti-aliasing prefilter to start with which is often an RF bandpass filter. This filter restricts the software-defined receiver from tuning to arbitrary carrier frequencies. Flexibility requires that we find a way to eliminate the prefilter.

Classic impulse sampling, which is implied in all the instances above when "sampling" is mentioned, is inherently wideband. Its weakness is that it samples, with equal fidelity, the narrow wanted spectrum and all unwanted spectrum that surrounds it. Ultimately it is this which imposes the need for an anti-aliasing prefilter.

Is there a way, then, to sample a continuous-time input of the type found in wireless channels with built-in anti-aliasing? This question has been asked before [7], and Yuan gives the simplest answer [8] (although perhaps not in the simplest terms). First, the channel of interest is translated to DC, accompanied by unwanted channels on either side that may take up unlimited bandwidth. This is a mixing process, although when realized with passive FET switches the mixer can be interpreted as (or even confused as) an RF sampler. Second, this shifted spectrum, represented after quadrature mixing by a complex analog waveform, is integrated over a window T (= $1/f_s$) of some suitable width (Fig.5). The result of this integration comprises an output sample. The integrator is reset (or dumped), and after time T is ready with the next sample. This process can continue without interruption by, for instance, time-interleaving two identical circuits alternating between integration and readout/reset. The windowed integration sampler takes continuous-time input and produces D-T analog samples at rate f_s. Its frequency transfer function is

$$|H(f)| = \frac{g_m T}{C}\left|\frac{\sin(\pi Tf)}{\pi Tf}\right| \tag{1}$$

This transfer function peaks at DC, and is zero at f_s, $2f_s$,... This protects the input at DC from aliasing. Inputs lying in the sidelobes will alias into the main lobe, but away from DC. The stopband $H_\alpha(kf_s)$ around the k^{th} null is defined by the bandwidth across which a specified attenuation α is obtained relative to DC. Thus,

$$H_\alpha(kf_s) = 2\alpha kf_s \tag{2}$$

This relation specifies the minimum sample rate f_s. For if the desired channel at DC occupies a bandwidth $\pm B/2$ around DC, and an anti-aliasing attenuation of α is desired *across this bandwidth*, then the sample rate must satisfy

$$f_s \geq B/(2\alpha) \tag{3}$$

An IF receiver based on this principle has been built (Fig.6) [9], although it uses a preselect bandpass filter and therefore does not allow wideband tuning. Furthermore, the frequency $1/4T$ of mixer switching is linked by a simple counter to the output sample rate $1/NT$, where N is programmable. It is straightforward to generate these clocks at 100 MHz IF, but may be difficult at RF, say at 5.5 GHz.

The windowed integration circuit merges quadrature mixing. It consists of a transconductance amplifier with IF input, tri-level switches that commutate its output current, and an op amp-based current integrator that is read and reset every N cycles of commutation.

These techniques are used in the Bluetooth [10] and GSM receivers [11] from Texas Instruments. Both use preselect filters for the respective bands, and therefore are not, in our sense, software-defined receivers; but a large amount of programmable filtering and decimation give them more flexibility than traditional receivers. Furthermore, the very simple realizations of the RF/analog blocks enable the radio to be integrated on the same chip as the DSP in deep submicron CMOS.

In the first stage following the LNA, the Bluetooth receiver (Fig.7) merges downconversion, integration and decimation filtering. The initial f_s~2.4 GHz (actually equal to the wanted channel frequency) positions the first alias deep in the RF prefilter's stopband. Decimation by 8× through a *sinc* filter similar to [6] follows, and an unavoidable continuously connected capacitor realizes a D-T pole. A first-order D-T *sinc* filter decimates by 4, and a third-order D-T *sinc*^{3} filter decimates by 2.

At the final rate of 37.5 MHz, it is possible that some other channel in the 80 MHz-wide ISM band might alias on to the wanted channel—but by this point the cumulative filtering will reduce it to a negligible level. As the first LO tunes to different channels, the conversion to symbol rate requires different non-integer interpolations in DSP.

UCLA SDR Receive Architecture

We define an SDR receiver as one (a) able to tune any channel from 800 MHz to 6 GHz, (b) able to tolerate known blockers as specified by the spectral templates of all bands in its tuning range with no prefilter at the antenna, (c) requiring an ADC that operates at an oversampling factor of the wanted channel's bandwidth without excessive power consumption, (d) exploits the conversion dynamic range of the ADC to relax filtering and programmable gain in the RF and analog front-end. A direct conversion receiver with a wide tuning LO best meets these requirements.

Let us illustrate (d) above with the case of GSM (Fig.8). The 200 kHz-wide wanted channel can be received at any strength from –102 to –15 dBm, a dynamic range of 87 dB. It is relatively easy in today's tech-

nology to envisage a delta-sigma converter that yields 14b resolution (86 dB SNDR) across 100 kHz. Allowing enough pre-amplification for the minimum SNR of 9 dB; then ensuring that amplification raises the input to no more than 4 dB below full scale to allow for EDGE; and finally leaving 16 dB margin, the RF/analog front-end needs a variable gain from 12 to 43 dB. In other words, to receive an input signal of 87 dB dynamic range only 31 dB of variable gain is sufficient and the ADC absorbs the rest [12]. Of course, this underestimates the complete picture, because depending on the filtering, blocking signals will also be incident on the A/D converter, and an expanded range of variable gain may be needed to prevent saturation.

Let us turn to filtering. The desired channel at zero IF is surrounded by adjacent channels in the same band, and active channels in all other bands. The windowed integration sampler is a filter with stopbands surrounding the nulls that when clocked at a sufficiently high frequency, attenuates *every possible* aliasing channel which is potentially a co-channel interferer for the wanted channel. Away from the aliasing frequencies, sidelobes which roll off at 20 dB/decade attenuate adjacent channels.

The initial sample rate is almost always too high, and must be decimated for A/D conversion. The decimation filter's nulls will lie in the sampler's sidelobes, and due to the lower rate, the stopband may not be wide enough to cover the wanted bandwidth. There are two ways to widen the stopband: either with a higher-order decimation filter, or with additional D-T filtering before decimation. We have found that a combination works well for most cases ranging from GSM reception with large out-of-band unmodulated blockers specified for type approval tests, to 802.11g WLAN reception in the presence of strong WCDMA cellular channels. As is usual in filter synthesis, at first solutions are found by trial-and-error, and then with experience synthesis is automated. In our SDR receiver (Fig.9), there are two *RC* poles at the mixer load, cascaded by windowed integration and an embedded D-T pole. A second-order $sinc^2$ decimation filter by 4× follows, and then a first-order decimation filter with selectable rate of 2× or 3×.

In effect, the SDR's baseband filter assumes the role of a passive prefilter at RF (Fig.10). Realized with simple switched-capacitor circuits the baseband filter is clock programmable over decades in cutoff frequency, whereas an RF prefilter is not programmable at all[2].

When large blocking signals are filtered *after* the LNA and mixer, even mild nonlinearity in these circuits can irreparably corrupt the wanted channel. The two phenomena to watch out for are AM detection by 2nd-order nonlinearity in the mixer, and cross-modulation by 3rd-order nonlinearity, usually also in the mixer. The only recourse is to use the best known mixer circuit topology, and to design it well.

To complete the receiver, what remains is a frequency synthesizer with wide tuning range. The tuning range of any *LC* oscillator is limited, but can be expanded with switchable capacitor banks. For continuous coverage of the band from 800 MHz to 6 GHz, two or three oscillators centered at widely different frequencies may be sufficient, accompanied by a bank of dividers and a multiplexer. Only one oscillator would be active at a time. Frequency addition and subtraction by SSB mixers is usually not acceptable in an SDR because of the inevitable spurs at unwanted sidebands.

These components, preceded by a wideband low-noise amplifier, comprise a recently developed single-chip CMOS software-defined receiver [13]. With no prefilter, it receives GSM using a 9 MHz (decimated from an initial f_s=72 MHz) sigma-delta ADC that resolves 14b in 100 kHz; and 802.11g using an 8b, 40 MHz (decimated from initial f_s=480 MHz) Nyquist ADC. In today's state-of-the-art, each ADC consumes about 10 mW.

Thus, the software-defined receiver has become practical reality.

[1] J. Mitola, "The Software Radio Architecture," in *IEEE Communications Mag.*, vol. 33, no. 5, pp. 26-38, May 1995.

[2] T. Hentschel and G. Fettweis, "The Digital Front-End: Bridge between RF and Baseband Processing," in *Software Defined Radio: Enabling Technologies*, W. Tuttlebee, Ed. Chichester, UK: Wiley, 2002, pp. 151-198.

[3] N. C. Davies, "A high performance HF software radio," in *Eighth Intl. Conf. on HF Radio Systems and Techniques*, pp. 249-256, 2000.

[4] H. Yoshida, S. Otaka, T. Kato, and H. Tsurumi, "A software defined radio receiver using the direct conversion principle: implementation and evaluation," in *IEEE Intl. Symp. Personal, Indoor and Mobile Radio Communications*, pp. 1044-1048 vol.2, 2000.

[5] D. M. Akos, M. Stockmaster, J. Tsui, and J. Caschera, "Direct bandpass sampling of multiple distinct RF signals," *IEEE Trans. on Communications*, vol. 47, no. 7, pp. 983-988, 1999.

[6] S. Lindfors, A. Parssinen, and K. Halonen, "A 3-V 230-MHz CMOS decimation subsampler," *IEEE Trans. on Circuits and Systems II*, vol. 50, no. 3, pp. 105-117, 2003.

[7] Y. Poberezhskiy and G. Poberezhskiy, "Sampling technique allowing exclusion of antialiasing filter," *Electronics Lett.*, vol. 36, no. 4, pp. 297-298, 2000.

[8] J. Yuan, "A charge sampling mixer with embedded filter function for wireless applications," in *2nd Intl. Conf. on Microwave and Millimeter Wave Technology*, Beijing, China, pp. 315-318, 2000.

[9] S. Karvonen, T. A. D. Riley, and J. Kostamovaara, "A CMOS Quadrature Charge-Domain Sampling Circuit with 66-dB SFDR up to 100 MHz," *IEEE Trans. on Circuits and Systems I*, vol. 52, no. 2, pp. 292-304, 2005.

[10] R. B. Staszewski, *et al.*, "All-digital TX frequency synthesizer and discrete-time receiver for Bluetooth radio in 130-nm CMOS," *IEEE J. of Solid-State Circuits*, vol. 39, no. 12, pp. 2278-2291, 2004.

[11] K. Muhammad, *et al.*, "A discrete time quad-band GSM/GPRS receiver in a 90nm digital CMOS process," in *Custom IC Conf.*, San Jose, CA, pp. 804-807, 2005.

[12] M. Beach, P. Warr, and J. MacLeod, "Radio Frequency Translation for Software Defined Radios," in *Software Defined Radio: Enabling Technologies*, W. Tuttlebee, Ed. Chichester, UK: Wiley, 2002, pp. 3-22.

[13] 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 90nm CMOS," in *Intl. Solid-State Circuits Conf.*, San Francisco, pp. 480-481,667, 2006.

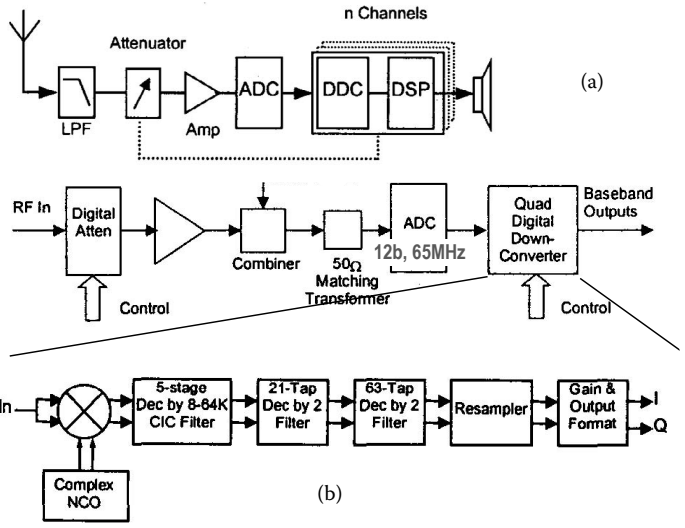

Fig.1. (a) Block diagram of Mitola-style HF software-defined receiver. (b) Details, particularly of the digital downconverter or channelizer.

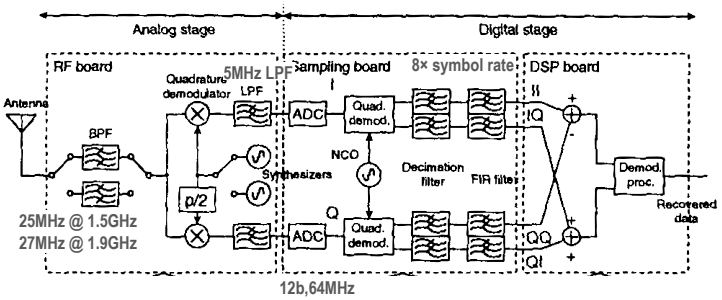

Fig.2. Two band software-definable receiver from Toshiba. With RF block downconversion, it uses the same ADC as the HF receiver to tune to arbitrary carrier frequencies.

[2] This is the case today. Tunable MEMs-based filters are in the research phase.

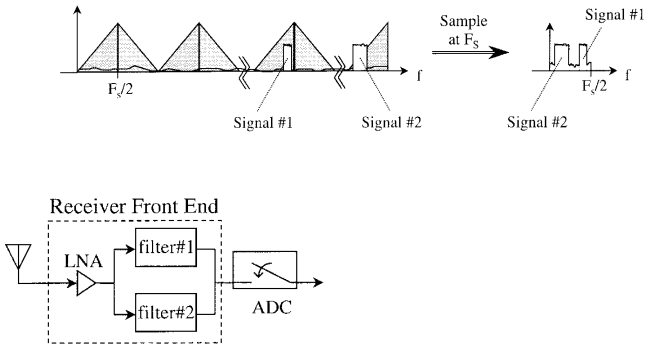

Fig.3. Use of undersampling (with prefilter) to accomplish two different frequency translations with a single LO, enabling simultaneous reception of GPS/GLONASS.

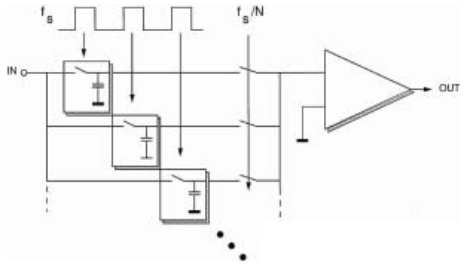

Fig.4. Analog decimation filter with equal tap weights, realizing anti-aliasing by D-T *sinc* transfer function. Images of *sinc* repeat at input Nyquist rate.

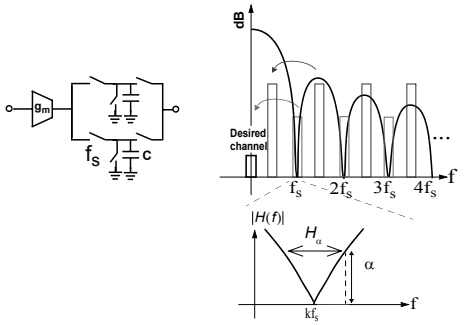

Fig.5. Windowed integration sampler offers lowpass transfer function with infinitely wide built-in anti-aliasing. Suitable to sample narrowband signal at DC (zero IF) surrounded by wide bandwidth of unwanted signals.

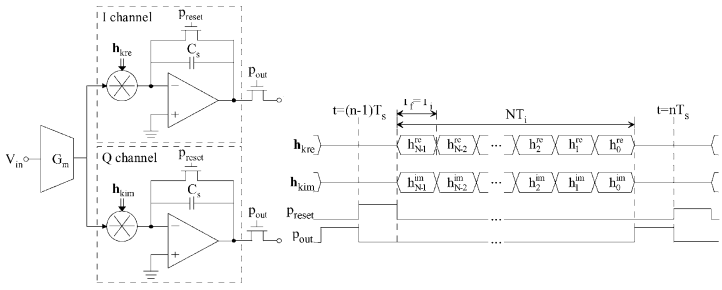

Fig.6. One realization of windowed integration in quadrature downconversion from IF to baseband.

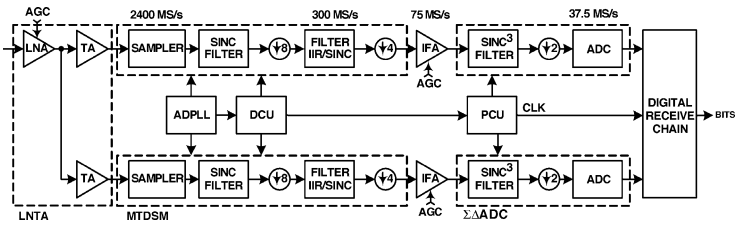

Fig.7. Bluetooth receiver from Texas Instruments embodies many of the techniques illustrated above.

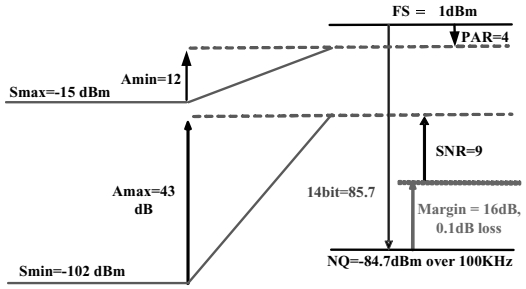

Fig.8. Using the surplus dynamic range of a low-power sigma-delta ADC to relax the requirements on analog variable gain. To receive GSM, for example, 30 dB variable gain is sufficient.

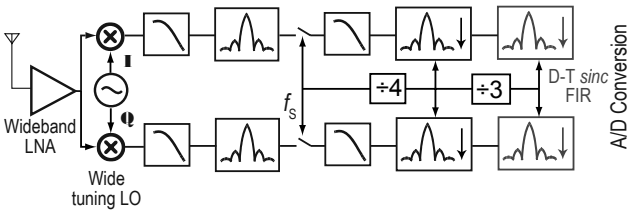

Fig.9. UCLA software-defined receiver tunes from 800 MHz to 6 GHz. On-chip programmable filter replaces role of antenna prefilter.

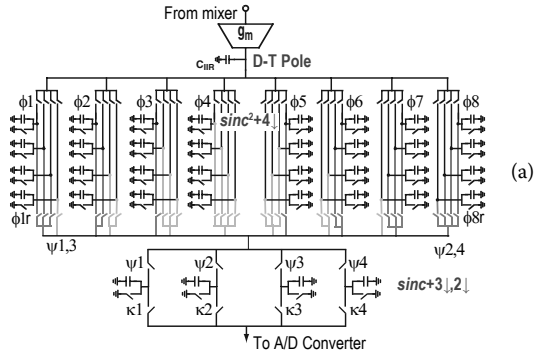

(a)

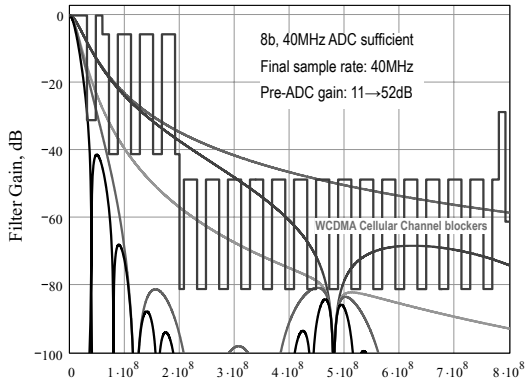

(b)

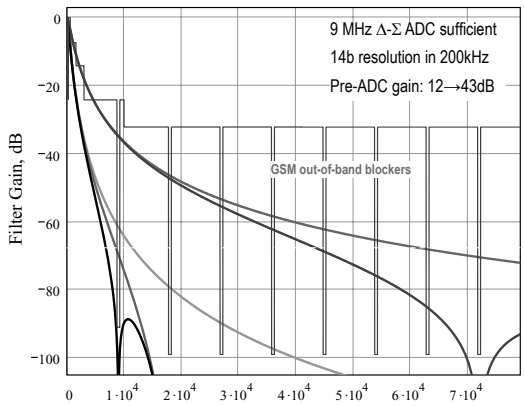

(c)

Fig.10. (a) Complete sampler and filter circuit. Filter settings for (b) 802.11g, and (c) GSM.