

Selecting Mixed Signal Component for Digital Communications Systems

IV. Receiver Architecture Considerations

by Dave Robertson

Part I introduced the concept of channel capacity and its dependence on bandwidth and SNR; part II summarized briefly different types of modulation schemes; and part III discussed approaches to sharing the communications channel, including some of the problems associated with signal-strength variability. This installment considers some of the architectural trade-offs used in digital communications receiver design for dealing with dynamic range management and frequency translation problems.

System Constraints: In a digital communications system, the function of the receiver circuitry is to recover the transmitted signal and process it for introduction to the demodulator, which then recovers the digital bits that constitute the transmitted message. As the last installment illustrates, obstacles to signal recovery show up as the signal travels through the transmission medium. These “impairments” can include signal attenuation, reflections, distortion, and the introduction of “interferers” (other signals sharing the transmission medium). The nature of the transmission impairments is a strong function of the medium (wireless, coaxial cable, or twisted pair wire), the communications scheme being used (TDMA, FDMA, CDMA, etc.) and the particular circumstances of the transmitter/receiver pair (distance, geography, weather, etc.). In any event, the important receiver design considerations are present to some extent in all receivers, simply to differing degrees. For this discussion, two examples will be used to illustrate the various receiver design issues. Figure 1 illustrates the relevant portions of the signal spectrum at the transmitter outputs and receiver inputs for two very different systems: a GSM cellular telephony application (Figure 1a and 1b) and an ADSL twisted-pair modem application (Figure 1c and 1d).




Figure 1. Transmitted and received spectra.

GSM uses a combination of FDMA (frequency division multiple access) and TDMA (time division multiple access) for multiplexing and a variation of quadrature phase shift keying for modulation.

In 1b, the amplitude is significantly reduced—a result of distance from the transmitter. In addition, several strong interfering signals are present—signals from other cellular transmitters in nearby bands that are physically closer to the receiver than the desired transmitter.

The ADSL modem in this example (Figure 1c) uses FDMA to separate upstream and downstream signals, and transmits its signal in a number of separate frequency bins, each having its own QAM (quadrature amplitude modulation) constellation (discrete multi-tone, or DMT modulation). The ADSL signal is attenuated by the twisted pair wire; attenuation is a strong function of frequency. In addition, an “interferer” is present. This might seem anomalous in a dedicated wire system, but in fact the interferer is the duplex (travelling in the opposite direction) signal of the modem leaking back into the receiver. This is generally referred to as *near-end echo*, and for long lines it may be much stronger than the received signal (Figure 1d).

These two examples illustrate critical functions of the receiver processing circuitry:

Sensitivity represents the receiver’s ability to capture a weak signal and amplify it to a level that permits the demodulator to recover the transmitted bits. This involves a gain function. As was discussed in Part 3 of this series, signal strength may vary significantly, so some degree of variable or programmable gain is generally desired. The way gain is implemented in a receiver usually requires a tradeoff between noise, distortion, and cost. Low-noise design dictates that gain be implemented as early in the signal chain as possible; this is a fundamental principle of circuit design. When calculating the noise contribution from various noise sources in a system, the equivalent noise of each component is referred to one point in the system, typically the input—referred-to-input (RTI) noise. The RTI noise contribution of any given component is the component’s noise divided by the total signal gain between the input and the component. Thus, the earlier the gain occurs in the signal path, the fewer stages there are to contribute significant amounts of noise.

Unfortunately, there are obstacles to taking large amounts of gain immediately. The first is distortion. If the signal is in the presence of large interferers (Figures 1b, 1d), the gain can’t be increased beyond the point at which the large signal starts to produce distortion. The onset of distortion is described by a variety of component specifications, including THD (total harmonic distortion), IP3 (third-order intercept point: a virtual measurement of the signal strength at which the power of the 3rd-order distortion energy of the gain stage is as strong as the fundamental signal energy), IM3 (a measure of the power in the 3rd order intermodulation products), and others. For an A/D converter or digital processing, “clipping” at full-scale produces severe distortion. So these strong signals must usually be attenuated before all the desired gain can be realized (discussed below).

Cost is another limiting factor affecting where gain can occur in the signal chain. As a general rule of thumb, high-frequency signal processing is more expensive (in dollars and power) than low frequency or baseband signal processing. Hence, systems that include frequency translation are generally designed to try to implement as much of the required gain as possible at the IF or baseband frequencies (see below). Thus, to optimize the location of gain in the signal path, one must simultaneously trade off the constraints of noise, distortion, power dissipation, and cost.

Specifications used to evaluate gain stages include the gain available (linear ratio or dB) and some description of the noise of the component, either in RTI noise spectral density (in nV/√Hz) or as *noise figure* (basically, the ratio of the noise at the output divided by the noise at the input, for a given impedance level).

Selectivity indicates a receiver's ability to extract or select the desired signal in the presence of unwanted interferers, many of which may be stronger than the desired signals. For FDMA signals, selectivity is achieved through filtering with discrimination filters that block unwanted signals and pass the desired signal. Like gain, filtering is generally easier at lower frequencies. This makes intuitive sense; for example, a 200-kHz bandpass filter implemented at a 1-MHz center frequency would require a much lower Q than the same 200-kHz filter centered on 1 GHz. But filtering is sometimes easier in certain high-frequency ranges, using specialized filter technologies, such as ceramic or surface acoustic wave (SAW) filters.

As noted above, filtering will be required early in the signal path to attenuate the strong interferers. Such filters will need to combine the required frequency response and low noise. Figures of merit for a filter include bandwidth, stop-band rejection, pass-band flatness, and narrowness of the transition band (the region between pass-band and stop-band). Filter response shape will largely be determined by the channel spacing and signal strength variations of the communications channel. Most FDMA cellular standards seek to ease filter requirements by avoiding the use of adjacent frequency channels in the same or adjacent cells, to permit wider transition bands and lower-Q (cheaper) filters.

Part of the selectivity problem is *tuning*—the ability to change the desired channel, since in most applications the signal of interest could be in any one of a number of available frequency bands. Tuning may be accomplished by changing the filter bandpass frequencies, but it is more commonly realized as part of the mixing operation (see below).

Frequency planning (mixing): Radio frequencies are selected based on radio transmission characteristics and availability of bandwidth for use for a given service, such as FM radio or cellular telephony. As was noted earlier, signal processing at high radio frequencies tends to be expensive and difficult. Besides, this added trouble seems unnecessary, since in most cases the actual signal bandwidth is at most a few hundred kHz. So most radio receivers use frequency translation to bring the signal carriers down to lower, more manageable frequencies for most of the signal processing. The most common means of frequency translation is a *mixer* (Figure 2).




Figure 2. Mixing—the image problem.

Mixing means using a nonlinear operation, usually multiplying the input signal and a reference oscillator signal, to produce spectral images at the sum and difference frequencies. For example: if we “mix” an RF signal at 900 MHz with an oscillator at 890 MHz, the output of the mixer will have energy at 1790 MHz (sum of frequencies) and 10 MHz (their difference). The 10-MHz signal becomes the signal of interest at the 10-MHz *intermediate frequency* (IF), while the sum frequency is easily filtered out. If the oscillator

frequency is increased to 891 MHz, it will translate an RF signal at 901 MHz to the IF; hence, channel selection, or tuning, can be realized by varying the oscillator frequency and tuning the output to the IF, using a fixed-frequency bandpass filter.

However, when mixing the 900-MHz RF with an 890-MHz local oscillator (LO), any 880-MHz interference present on the RF signal will also be translated to a difference frequency of 10 MHz. Clearly, any RF signal at the “image” frequency of 880 MHz must be suppressed well below the level of the desired signal before it enters the mixer. This suggests the need for a filter that passes 900 MHz and stops 880 MHz, with a transition band of twice the intermediate frequency. This illustrates one of the trade-offs for IF selection: lower IFs are easier to process, but the RF image-reject filter design becomes more difficult. Figures of merit for mixers include gain, noise, and distortion specifications like those used for gain stages, as well as the requirements on the oscillator signal input.

Other mechanisms of dealing with the image rejection problem are beyond the scope of this short treatment. One worth mentioning, though, because of its widespread use is *quadrature downconversion*. In-phase and quadrature representations of the input signal are mixed separately and combined in a way to produce constructive interference on the signal of interest and destructive interference on the unwanted image frequency. Quadrature mixing requires two (or more) signal processing channels well-matched in both amplitude and frequency response, because mismatches allow the unwanted image signal to leak into the output.

Equalization: Real-world transmission channels often have a more severe impact on signals than simple attenuation. Other channel artifacts include frequency-dependent amplitude and phase distortion, multi-path signal interference (prevalent in mobile/cellular applications), and bandlimiting/intersymbol interference from the receiver processing circuits. Many receiver systems feature “equalization” circuits, which provide signal processing that attempts to reverse channel impairments to make the signal more like the ideal transmitted signal. They can be as simple as a high frequency boost filter in a PAM system or as complicated as adaptive time- and-frequency-domain equalizers used in DMT ADSL systems. As capacity constraints push system architectures towards more complicated modulation schemes, equalization techniques, both in the analog and digital domains, are increasing in sophistication.

Diversity: In mobile applications, the interference patterns from a mobile transmitter can vary the strength of the signal at the basestation receiver, making the signal difficult or impossible to recover under certain conditions. To help reduce the odds of this occurring, many basestations are implemented with two or more receiving antennas separated by a fraction of the RF wavelength, such that destructive interference at one antenna should represent constructive interference at the other. This diversity improves reception at the cost of duplicating circuitry. Diversity channels need not be closely matched (matching is required for quadrature channels), but the system must have signal processing circuitry to determine which of the diversity paths to select. *Phased-array* receivers take the diversity concept to the ultimate, combining the signal from an array of receivers with the proper phase delays to intentionally create constructive interference between the multiple signal paths, thereby improving the receiver's sensitivity.

Conventional Receiver Design: Figure 3a illustrates a possible architecture for a GSM receiver path, and Figure 3b illustrates that of an ADSL modem. As noted earlier, the task of the receive

circuitry is to provide signal conditioning to prepare the input signal for introduction to the demodulator. Various aspects of this signal conditioning can be accomplished with either digital or analog processing. These two examples illustrate fairly traditional approaches, where the bulk of signal processing is done in the analog domain to reduce the performance requirements on the A/D converter. In both examples, the demodulation itself is done digitally. This is not always necessary; many of the simpler modulation standards can be demodulated with analog blocks. However, digital demodulation architectures are becoming more common, and are all but required for complicated modulation schemes (like ADSL).

The GSM receiver signal path shown in Figure 3a illustrates the use of alternating gain and filter stages to provide the required selectivity and sensitivity. Channel selection, or tuning, is accomplished by varying the frequency of the first local oscillator, LO1. Variable gain and more filtering is applied at the IF frequency. This is a narrowband IF system, designed to have only a single carrier present in the IF processing. The IF signal is mixed down to baseband, where it is filtered once more and fed to a sigma-delta A/D converter. More filtering is applied in the digital domain, and the GMSK signal is digitally demodulated to recover the transmitted bit stream.

The ADSL receiver has different requirements. Frequency translation is not required, since the signal uses relatively low frequencies (dc to 1.1 MHz). The first block is the “hybrid”, a special topology designed to extract the weak received signal from the strong transmitted signal (which becomes an interferer—see Figure 1d). After a gain stage, a filter attempts to attenuate the echo (which is in a different frequency band than the desired signal.) After the filter, a variable-gain stage is used to boost the signal to as large a level as possible before it is applied to the A/D converter for digitization. In this system, equalization is done in both the time and frequency domains before the signal is demodulated. This example shows the equalization taking place digitally (after the A/D converter), where it is easier to implement the required adaptive filters.

New twists—receivers “go digital”: Advances in VLSI technology are making more-sophisticated receiver architectures practical; they enable greater traffic density and more flexibility—even receivers that are capable of handling multiple modulation standards. An important trend in this development is to do more and more of the signal processing in the digital domain. This means that the A/D “moves forward” in the signal chain, closer to the

antenna. Since less gain, filtering and frequency translation is done prior to the A/D, its requirements for resolution, sampling frequency, bandwidth, and distortion increase significantly.

An example of this sophistication in modems is the use of *echo cancellation*. The spectrum of Figure 1d shows the strong interferer that dominates the dynamic range of the received signal. In the case of a modem, this interference is not a random signal, but the duplex signal that the modem is transmitting back upstream. Since this signal is known, signal processing could be used to synthesize the expected echo on the receive line, and subtract it from the received signal, thereby cancelling its interference. Unfortunately, the echo has a strong dependence on the line impedance, which varies from user to user—and even varies with the weather. To get reasonable cancellation of the echo, some sort of adaptive loop must be implemented. This adaptivity is easier to do in the digital domain, but it requires an ADC with sufficient dynamic range to simultaneously digitize the weak received signal and the echo; in the case of ADSL, this suggests a 16 bit A/D converter with 1.1 MHz of bandwidth. (e.g., the AD9260). As a significant reward for this higher level of performance with a sufficiently accurate echo canceller, upstream and downstream data can simultaneously occupy the same frequencies, dramatically increasing the modem’s capacity, particularly on long lines.

In the case of GSM, there are various approaches to advanced receivers. As the ADC moves forward in the signal chain, instead of capturing a baseband signal around dc, it has to digitize the IF signal, which would typically be in the range of 70 MHz to 250 MHz. Since the bandwidth of interest is only a few hundred kHz, it is unnecessary (and undesirable) to run the ADC at 500 MHz; instead, undersampling is used. If the ADC is clocked at 20 MHz with the signal of interest at 75 MHz, the signal will alias down to 5 MHz ($= 4 \times 20 - 75$) MHz; essentially, the undersampling operation of the ADC acts like a mixer. As with a mixer, there is an image problem, so signal content at 65 MHz ($= 3 \times 20 + 5$ MHz) and 85 MHz ($= 4 \times 20 + 5$ MHz) would need to be filtered out ahead of the ADC. (An AD6600 dual-channel gain-ranging ADC—available by winter—would be useful here).

An even greater advancement on cellular receivers is to implement a wideband receiver. In the example shown in Figure 3b, the single carrier of interest is selected by varying the LO frequency and using very selective filters in the IF signal processing. A wideband radio (available soon) seeks to digitize *all* the carriers, allowing the tuning and signal-extraction functions to be implemented digitally. This imposes severe requirements on the ADC’s performance. If a 15-MHz-wide cellular band is to be digitized, an ADC sampling rate of 30-40 MSPS is required. Furthermore, to deal with the near/far problem, the converter dynamic range must be large enough to simultaneously digitize both strong and weak signals without either clipping the strong signals or losing the weak signals in the converter quantization noise. The converter requirements for a wideband radio vary with the cellular standard—anywhere from 12 bits, 40 MSPS for the U.S. AMPS standard (AD9042) to 18 bits, 70 MHz for GSM. The great advantages to this kind of implementation make the tradeoff worthwhile; one receiver can be used to simultaneously capture multiple transmissions, and—since the selection filtering is done digitally—programmable filters and demodulators can be used to support a multi-standard receiver. In radio industry jargon, this is a move towards the “software radio”, where most of the radio processing is digital. ▶




Figure 3. Typical receiver architectures.