SDR Basics: Receivers

By Bruce A. Fette

Architectures

Ideally the designer of an SDR would like to put the data converters directly on the antenna. However, this is not a practical solution. In reality, some analogue front end must be used before the ADC in the receive path and after the DAC in the transmit path that does the appropriate frequency translation. The most common of these architectures is the super-heterodyne architecture. Although it's many decades old, new semiconductor technology and high levels of integration have kept this architecture vitalised and in popular use both in the transmit and receive signal paths [5, 6].

Other architectures such as direct conversion both for transmit and receive are seeing some popularity in applications that are not as demanding. Currently, direct conversion (Tx and Rx) is found in user terminals for cellular communications as well as for Tx on the base station side. It is possible that future developments will enable direct conversion on the receive side as well. Until then, the superheterodyne architecture will continue to be used in one form or another.

Receiver

High-performance SDR receivers are typically constructed from some variant of the superheterodyne architecture. A super-heterodyne receiver offers consistent performance across a large range of frequencies while maintaining good sensitivity and selectivity [7, 8]. Although not trivial to design, the possibility of combining wideband analogue techniques and multiple front ends would allow operation across different RF bands. In the case of multi-carrier applications, this could be done simultaneously if necessary.

Multicarrier

Depending on the applications, one or more receive channels may be desired. Traditional applications may require only a single RF channel. However, applications that require high capacity or interoperability may require a multi-carrier design. SDRs are well suited for multi-carrier applications, since they employ a highly oversampled ADC with ample available bandwidth.

An oversampled ADC is one in which the sample rate is operating beyond that which is required to meet the Nyquist criterion [18], which states that the converter sample rate must be twice that of the information bandwidth. Since an SDR may not have advance knowledge of the bandwidth of the signal it will be used to receive, the sample rate must be appropriately high enough to sample all anticipated bandwidths.

Current ADC technology allows high dynamic range bandwidths of up to 100 MHz to be digitised. With this much bandwidth, it is also possible to process multiple channels. Figure 18.5 shows a typical multi-carrier receiver example, and Figure 18.6 shows a spectral display.

In this example, the sample rate of the ADC is set to 61.44 mega-samples-per-second (MSPS), which gives a Nyquist bandwidth of 30.72 MHz. If each RF channel is 1.25 MHz wide, then Nyquist indicates that the number of potential channels is about 24.5. In practice, by allowing for reasonable transition bands on the anti-aliasing filters, the typical available bandwidth is one-third the sample rate instead of the Nyquist one-half. Thus, the available bandwidth for our example is 20.48 MHz, which is just over 16 channels at 1.25 MHz.

Since the channel characteristics can be changed, it is easy enough to change the CDMA example to a GSM example. In this case, both the digital preprocessing and the general-purpose DSP are reconfigured, respectively, by changing the digital channel filter from GSM to CDMA and by loading the new processing code into the DSP. Since GSM channels are 200kHz wide, this example could easily be reconfigured as a 102-channel GSM receiver.

While both such examples would provide a lot of utility, perhaps a more interesting example would be to configure the receiver such that part of the channels could be CDMA while the other would be configured as GSM! Furthermore, if one of the configurations is at capacity and the other is underutilised, CDMA channels could be converted into several GSM channels or vice versa, providing the flexibility to dynamically reallocate system resources on an as-needed basis (a key goal of software-defined radio).

Single Carrier

Not all SDR applications require more than one channel. Low-capacity systems may require only one carrier. In these applications, a high oversampling is still desired. If the channel is reprogrammable, it is possible that it may be as narrow as a fewkHz or as wide as 5 to 10 MHz. In order to accommodate this range of bandwidths, the sample rate should be suitable for the highest potential bandwidth, in this case 10 MHz. From the multi-carrier example, we would typically sample at least three times the bandwidth. In this example, a sample rate of 30.72 MSPS or higher would allow signal bandwidths from a fewkHz up to 10 MHz to be processed. Aside from the fact that only one channel is processed, the single-carrier receiver has all of the capacities of that of a multi-carrier receiver; it can be reconfigured as necessary.

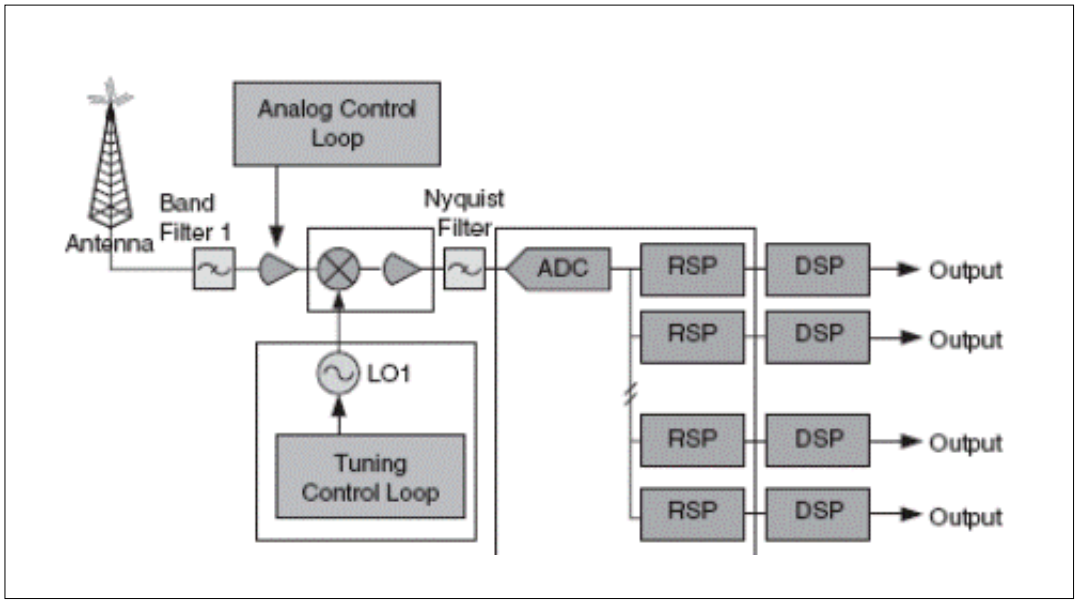

Figure 18.5: Multicarrier CDMA example.

SDR Receiver Elements

Referring to the single-carrier block diagram in Figure 18.7, while keeping in mind that this applies to the multi-carrier example as well, a fully developed SDR will have all signal elements that are programmable.

The antenna is no exception, and unfortunately it is one of the weakest elements in an SDR [1]. Since most antenna structures have a bandwidth that is a small percentage of its centre frequency, multi-band operation can become difficult. In the many applications where single bands of operation are used, this is not a problem. However, for systems that must operate across several orders of frequencies, the antenna must be tuned by some means to track the operating frequency to maintain operating efficiency.

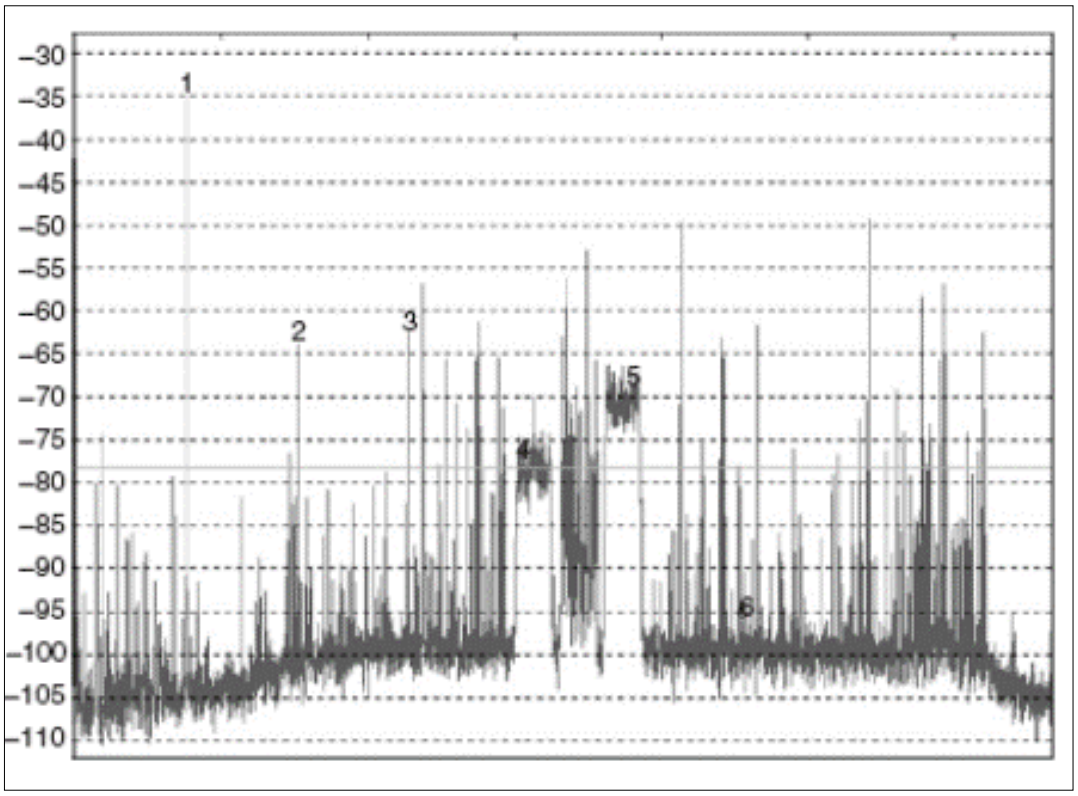

Figure 18.6: Multimode spectrum with IS-95 and narrowband carriers.

While it is true that just about any antenna can be impedance matched to the active electronics, there is usually a sacrifice in the link gain that potentially results in an antenna loss, whereas most antenna designs should actually provide a modest signal gain. Therefore, tuning the electrical length of the antenna is desired over simply changing the matching of the antenna.

Next in the signal chain is the band-select filter electronics. This element is provided to limit the range of input frequencies presented to the high-gain stage to minimise the effects of intermodulation distortion. Even in the case where intermodulation is not a problem, it is possible that strong out-of-band signals could limit the amount of potential gain in the following stages, resulting in limited sensitivity. This is especially true for receivers tuned near television and audio broadcast services where transmit power levels can exceed 100 kW.

This can be especially problematic for multi-carrier receivers where many orders of signal magnitude must be dealt with. If all of the signals are of interest, then it will not be possible to filter the stronger signals, and the resulting receiver must have a relatively large signal dynamic range [8].

Most receivers require a low-noise amplifier or LNA. An SDR should ideally incorporate an LNA that is capable of operating over the desired range of frequencies. In addition to the typical low NF and high IP3, it may be desirable to have the ability to adjust the gain and potentially scale the power down (often NF and IP3 track bias current) when possible, which will allow for a variety of signal conditions that exist across the bands of operation.

Mixers are used to translate the RF spectrum to a suitable IF frequency. Although only one mixer is shown in Figure 18.7, many receivers may use two or three mixer stages, each successively generating a lower frequency. (Note: Receiver IFs are not always lower than the RF signal. A common example is found in HF receivers where the desired RF signal may be only a few MHz. In these cases, they are frequently mixed up to IF frequencies of 10.7 MHz, 21.4 MHz, 45 MHz, or higher IF frequencies because of the availability or performance of the required component.) Each successive stage also takes advantage of filtering that is distributed throughout the chain to eliminate undesired images as well as other undesired signals that may have survived the mix-down process. The filtering should also be appropriate for the application. A traditional single-carrier receiver would generally apply channel filtering through the mixer stages to help control the IP3 requirements of each stage. Analogue channel filtering is not possible in the case of a multi-carrier receiver where the channel bandwidths are not known in advance.

Therefore, the mixing process must preserve the entire spectrum of interest. Likewise, our single-carrier SDR application must also preserve the maximum possible spectrum in case the SDR requirements need the full spectrum. In this case, it is probable that our single-carrier example may be processing many carriers, even if only one is of interest. As with the LNA, it would be desirable for the mixer in an SDR to have an adjustable bias. As with the LNA, this bias could be used to properly set the conversion gain and IP3 of the device to correspond to the desired signal conditions.

Some receiver architectures utilise a quadrature demodulator in addition to, or instead of, a mixer. The purpose of the demodulator is to separate the I and Q components. Once they have been separated, the I and Q paths must maintain separate signal conditioning. In the digital domain, this is not a problem; in the analogue domain, however, the signal paths must be perfectly matched, or I/Q imbalances will be introduced, potentially limiting the suitability of the system.

Many SDR receivers avoid this problem by utilising "real" sampling (as opposed to complex sampling), as shown in the single-carrier example, and using a digital quadrature demodulator in the digital preprocessor that will provide perfect quadrature. The local oscillator is used to generate the proper IF when mixed with the incoming RF signal. Generally, a local oscillator (LO) is variable in frequency and easily programmable via software control using PLL or DDS techniques. There are cases

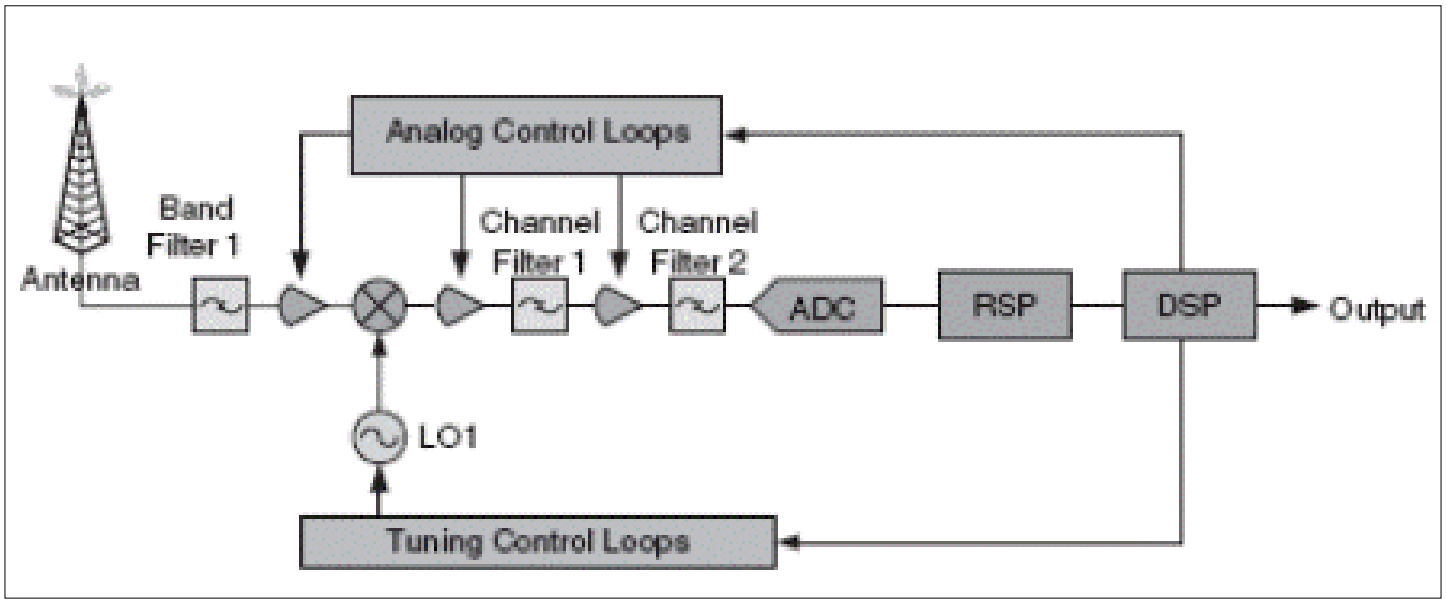

Figure 18.7: Single-carrier Rx example.

where the LO may not require frequency hopping. One such example is for receiving multiple carriers within a fixed band. In this case, the LO is fixed, and the entire band is block-converted to the desired IF. It often may be desirable to change the LO drive level to optimise spurious performance under a variety of signal conditions.

Quite often the IF amplifier is in the form of an AGC. The goal of the AGC is to use the maximum gain possible without overdriving the remainder of the signal chain. Sometimes the AGC is controlled from an analogue control loop. However, a digital control loop can also be used to implement difficult control loops not possible using analogue feedback. In multi-carrier applications, use of an AGC may at best be difficult. If insufficient dynamic range is available in the receiver (determined largely by the ADC), reduction in gain from a strong signal may cause weaker signals to be lost in the noise floor of the receiver. In applications such as this, a digital control loop for the gain is ideal. The control loop can be used as normal as long as no signals are at risk to being lost.

However, if a weak signal is detected in the presence of a very strong signal, the decision could be made to allow a limited amount of clipping rather than reduce the gain and risk total loss of the weak signal. Conditional situations like this are much easier to control with a digital control loop than with an analogue loop, allowing much greater control of total conversion gain of the receiver.

The ADC is used to convert the IF signal or signals into digital format for processing. Quite often the ADC is the bottleneck, and selection of the ADC is often a driving factor that determines the architecture of the SDR [1, 9, 10]. Oftentimes, the designer is forced to select the best available ADC, realising that under many conditions the ADC may be overspecified.

Still other times, air interface standards may not be directed towards multi-carrier receivers and require much better ADCs than are required when deployed in the field, simply because of the test methodology specified by the standard. For the ADC it may be desirable to change the sample rate, input range, and potentially the active bandwidth. The digital preprocessor can take many forms. For very high sample and data rates, this is usually implemented as either an FPGA or ASIC. These circuits by nature are quite flexible in their functions and range of parameters. An FPGA can, of course, be programmed for any function desired. Typically, an FPGA would be programmed to perform the quadrature demodulation and tuning, channel filtering, and data rate reduction.

Other functions such as RF power measurement and channel linearization are possible. All of these elements are easily generated using a variety of digital techniques and are readily programmed by loading a variety of coefficients to the FPGA. By doing this, a single-chip configuration can be used to generate a digital preprocessor capable of tuning the entire range of the ADC Nyquist band and filtering a signal with bandwidths from a fewkHz to several MHz. When multiple channels are required, the design can be repeated to fill the FPGA. If a lower cost option is required, various ASICs are available that perform these functions. They are often referred to as channelisers, RSPs, or DDCs. The final element in the SDR is the DSP. Since this is general-purpose DSP, it can be programmed for any required processing task. Typical tasks include equalisation, detection, rake receiver functions, and even network interfacing, to name a few.

Because they are fully programmable, they can be used for just about any signal processing task and control all of the features in the other elements of the block diagram. As DSP processing capabilities increase, DSPs may well take over many of the functions within the digital preprocessors.

References

1. J. H. Reed, Software Radio: A Modern Approach to Radio Engineering, Prentice Hall, Upper Saddle River, NJ, 2002.
2. J. Mitola, III, "Software Radio"Cognitive Radio," http://ourworld.compuserve.com/homepages/jmitola/.
3. B. Brannon, D. Efstathiou, and T. Gratzek, "A Look at Software Radios: Are They Fact or Fiction?" Electronic Design, (December 1998): pp. 117"122.
4. B. Clarke and K. Kreitzer, "Software Radio Concepts," unpublished paper.
5. B. Brannon, "Digital-Radio-Receiver Design Requires Reevaluation of Parameters," EDN, 43 (November 1998): pp. 163"170.
6. B. Brannon, "New A/D Converter Benefits Digital IFs," RF Design, 18 (May 1995): pp. 50"65.
7. W. H. Hayward, "Introduction to Radio Frequency Design," The American Radio Relay League, 1994"1996.
8. J. J. Carr, Secrets of RF Circuit Design, McGraw-Hill, New York, 2001.

9. B. Brannon, "Fast and Hot: Data Converters for Tomorrow's Software-Defined Radios," RF Design, 25 (July 2002): pp. 60"66.
10. B. Brannon and C. Cloninger, "Redefining the Role of ADCs in Wireless," Applied Microwave and Wireless, 13 (March 2001): pp. 94"105.
11. B. Brannon, "DNL and Some of Its Effects on Converter Performance," Wireless Design and Development, 9 (June 2001): p. 10. w.newnespress.com.
12. B. Brannon, "Overcoming Converter Non-linearities with Dither," Analog Devices Applications Note AN-410, www.analog.com.
13. W. Kester, "High-Speed Sampling and High-Speed ADCs," Section 4, High-Speed Design Techniques, www.analog.com.
14. W. Kester, "High-Speed DACs and DDS Systems," Section 6, High-Speed Design Techniques, www.analog.com.
15. About CDMA and CDMA University. Available at http://www.qualcomm.com.
16. Specifications. Available at http://www.3gpp2.org.
17. R. H. Walden, "Analogue-to-Digital Converter Survey and Analysis," IEEE Communications Magazine, 17 (April 1999): pp. 539"550.
18. H. Nyquist, "Certain Topics in Telegraph Transmission Theory," AIEE Transactions, 47 (April 1928): pp. 617"644.
19. AD6645 Datasheet. Available at http://www.analog.com.