Software Defined Radio

Brad Brannon, Analog Devices, Inc.

What is software defined radio?

Over the last decade as semiconductor technology has improved both in terms of performance capability and cost, new radio technologies have emerged from military and research and development labs and become mainstream technologies. One of these technologies is software defined radio. Although much has been discussed in recent years, a good definition of software radio is difficult to generate. This is largely due to the flexibility that software defined radios offer, allowing them to take on many different forms that can be changed to suite the need at hand.

However, software defined radios or SDRs, do have characteristics that make them unique from other types of radios. As the name implies, a SDR is a radio that has the ability to be transformed through the use of software or re-definable logic. Quite often this is done with general purpose DSPs or FPGAs as discussed later in the chapter. In order to take advantage of such digital processing, traditional analog signals must be converted to and from the digital domain. This is accomplished using analog-to-digital (ADC) and digital-to-analog converters (DAC). To take full advantage of digital processing, SDRs keep the signal in the digital domain for as much of the signal chain as possible, digitizing and reconstructing as close to the antenna as possible, which allows digital techniques to perform functions traditionally done by analog components as well as others not possible in the analog domain. There are limits to this however. Despite the fact that an ADC or DAC connected directly to an antenna is a desirable end goal, there are issues with selectivity and sensitivity that an analog front end can remedy. The alternative to digitizing at the antenna is the use of a completely flexible analog front end (AFE) capable of translating a wide range of frequencies and bands to that which the data converters themselves can adequately process [1].

SDRs are ideal candidates to be used for multi-carrier, single-carrier, single-band, multi-band and multi-mode transceivers. Some of these issues will be covered later. The key point is that SDRs have the ability to go beyond simple single channel, single mode transceiver technology with the ability to change modes arbitrarily because the channel bandwidth, rate, and modulation are all flexibly determined through software. These characteristics may be changed by direct input, floppy disk, over the air download or through the use of careful signal analysis to determine analytically how the information is coded through a process termed as Cognitive Radio [2]. Regardless of the means by which the radio is reconfigured, a fully implemented SDR will have the ability to navigate a wide range of frequencies with programmable channel bandwidth and modulation characteristics. The table below lists some of the possible characteristics of a SDR. In addition to RF tuning, a transceiver must include the ability to take advantage of one or more of these characteristics to be considered as an SDR.

Channel Bandwidth
Data Rate
Modulation Type
Conversion Gain

Aspects of software defined radio

As the table above indicates, there are a number of characteristics that an SDR possesses. While it is not required that an SDR have all of these characteristics, having one or more of them is. Additionally, the categories above can be further broken down as detailed below. It should be kept in mind that since software defined implies a high degree of flexibility and variability, the list below is not all encompassing and subject to change over time, but serves as a starting point at understanding the different facets of what SDR can be.

Multi-Band

Most traditional radio architectures operate on a single band or range of frequencies. There are many applications where multiple frequencies of operations are desired. These include cellular communications, government and non-government agencies, and intelligence collection to list a few. Where these situations exist, the norm is to utilize multiple radios; each designed to operate in one specified band. A multi-band radio has the ability to operate on two or more bands either sequentially or simultaneously as in the case of a basestation that may be linking handsets from different bands.

Multi-Carrier

A multi-carrier or multi-channel radio has the ability to simultaneously operate on more than one frequency at a time. This may be within the same band or in the case of a multi-band radio, in two different bands at the same time. Quite often, multi-carrier applies to a basestation that may be servicing many users at once, but can also apply to a user terminal that my be processing both voice and data on different RF carriers.

Multi-Mode

Multi-mode implies the ability to process several different kinds of standards. Examples of standards include AM, FM, GMSK, CDMA but is limited to none of these. An SDR has the ability to work with many different standards and be continuously reprogrammed. Therefore, a better term than multi-mode, which implies a discrete number of modes, may be variable mode, which implying a continuously changeable mode of operation. As with other characteristics, these modes may be sequentially or simultaneously in the case of a multi-carrier radio.

Multi-Rate

Multi-rate is closely related to multi-mode. A multi-rate radio is one that either processes different parts of the signal chain at different samples rates as in a multi-rate filter or one where the radio has the ability to process different modes that require different data rates. An example of a multi-rate radio would be one that can process GSM at 270.833 kSPS or CDMA at 1.2288 MCPS. As with other characteristics, this can be sequentially or at the same time on different carriers.

Variable Bandwidth

Variable bandwidth is also another aspect of multi-mode. A traditional radio determines the channel bandwidth with a fixed analog filter such as a SAW or ceramic filter. An SDR however determines the channel bandwidth using digital filters that can be altered. While a series of switched analog filters could be used to change the channel bandwidth in a traditional receiver, only a small number would be practical. Additionally, digital filters have the potential to implement filters not possible in the analog domain. Lastly, digital filters can be tailored to both adapt around interferers and compensate for transmission path distortion, both features that analog filters are hard pressed to accomplish.

History and Evolution of software defined radio

The history of SDR began in the mid 1980's. One of the first major developments for SDR was the SpeakEasy, a transceiver platform designed by Hazeltine and Motorola, based on SDR technology for Rome AFB. The SpeakEasy was designed to provide tactical military communications from 2 MHz to 2 GHz and to provide interoperability between the different air interface standards of the different branches of the armed forces. To achieve this goal, the SpeakEasy utilized many of the techniques discussed in this chapter to provide multi-band, multi-modes of operations. Although many people contributed to the concept and development of software defined radio, Joe Mitola of Mitre is generally credited with being the 'father of software defined radio' [2].

Figure 1: SpeakEasy Picture from web. Need to secure permission or omit.
http://www.rl.af.mil/div/IFB/techtrans/datasheets/SPEAK-PIX.html

Although the SpeakEasy was a fully developed SDR, it is fair to say that simpler and more rudimentary forms of SDR existed before this program. By taking a look at how systems are being developed in the commercial realm, it is easy to see how they also may have evolved in military and non-military programs.

Although there are many enabling technologies that have come online in the last decade, one of the key technical driving forces was the development of low cost Digital Signal Processors. From a market point of view, the rapid growth of the telecommunications industry, particularly cellular communications, provided a demand for low cost equipment both from a user and infrastructure point of view. Although first generation cellular was based on analog modulation schemes (which did not require significant digital processing), it became clear that due to the limit amount of spectrum and the relative inefficiency of those standards that more efficient means of spectral usage were required. Therefore second generation cellular systems such as GSM and IS-95 were developed that took advantage of the emerging DSP technologies. In these early systems, the DSP became the MODEM function and was responsible for taking the complex baseband data (I and Q) and determining what bit stream was being sent and correcting for errors introduced due to noise, interference and fading.

Conceptually, these modem functions were based on programs running on a DSP and therefore could be changed simply by changing the program. In fact, over time these standards evolved and variations of the standards were introduced that allowed better efficiency and higher data transmission rates. Many of these improvements were offered simply by updating the modem software. While consumers seldom experienced these benefits for a number of economic reasons, the infrastructure side did benefit from these upgrades and is benefiting from many of these software updates in the migration from 2G, to 2.5G and ultimately to 3G, most notably in the evolution of the CDMA2000 and UMTS standards [15, 16].

While the evolution of the modems used for GSM and CDMA is an aspect of SDR, other factors such as incompatibility of these two standards drives the second aspect of SDR. While CDMA is primarily a North American and Asian standard; and GSM is a European and rest of world standard, in reality both of these standards are over-layed in many countries. Ideally, service providers would like to purchase one piece of equipment that would work with both standards. Unfortunately, these (and other) standards are incompatible in terms of bandwidth, modulation format and data rate. Traditional radios, even those with DSP modems, operate with fixed bandwidths and therefore prevent cross functionality. A typical GSM system works with a 200 kHz bandwidth while an IS-95 system operates on 1.25 MHz bandwidth. Both systems typically utilize surface acoustic wave (SAW) filter technology to set the bandwidth. Since these devices are fixed, it is not possible (aside from electronically switching filters) to change the channel bandwidth characteristics. Therefore, aside from the modem function, an SDR needs additional circuitry that allows other properties of the air interface to be adapted. In the example here, the channel bandwidth must be adapted. In practice, other terms must also be adapted as well. In practice, the optimal way to do this is to digitize the signal and use digital techniques for manipulating the channel of interest. This manipulation often

occurs in a general purpose DSP or more frequently in a digital ASIC with the ability to accommodate a near continuous range of channel bandwidth, data rates and other physical characteristics in a fully digital manner.

The following figures show the practical evolution path of an Rx SDR architecture. In figure 2, a traditional super-heterodyne receiver is shown with analog detection. Figure 3 adds the DSP function that operates in place of the analog detector and can also function as the baseband modem. This allows the exact demodulation functions to change as needed while the other channel characteristics are fixed. Figure 4 includes a wideband ADC and a digital pre-processing before the modem that allows the physical channel characteristics to be adapted as necessary. Figure 4 also shows what a full SDR might look like. An architecture such as this may find use in diverse areas such as multi-mode systems capable of simultaneously processing several standards at once or simply as a manufacturing option to simplify production and inventory issues.

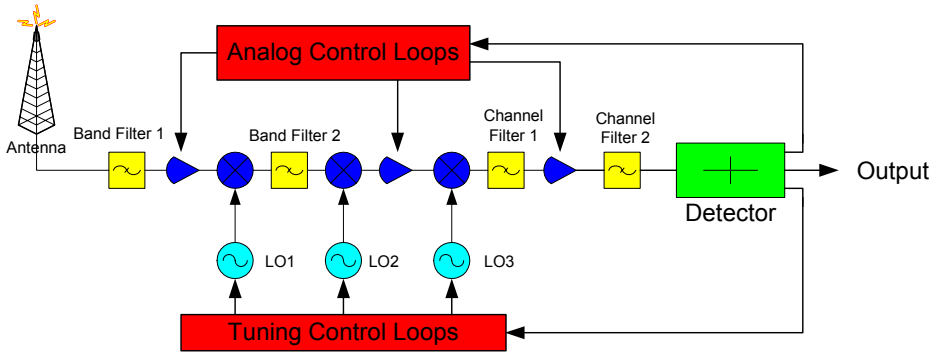

Figure 2: Traditional Super-heterodyne with analog detection

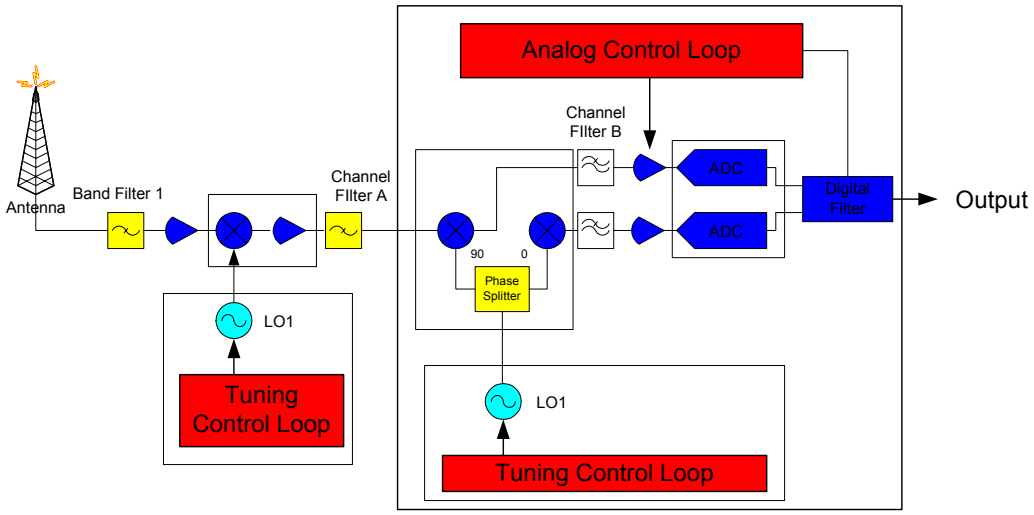

Figure 3: Super-heterodyne with baseband IQ sampling.

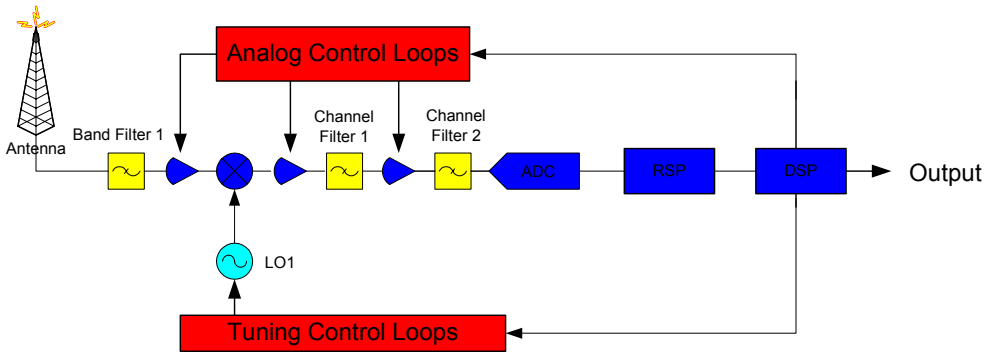

Figure 4: Super-heterodyne software defined radio

Applications and need for SDR

Interoperability

The military is not the only agency in need of interoperability. As numerous agencies, both domestic and international, have responded to various natural and manmade disasters around the world, communications between the different responding groups has often been hindered by the fact that different communications systems rarely work with one another because the frequency and air interfaces are different. SDR provides an ideal solution to these dilemmas. A centrally deployed basestation could be used to receive the transmissions of one agency and reformat and rebroadcast them on the frequencies of the other responding agencies. Since the system would be reconfigurable, as new agencies arrive or depart, the SDR can be rapidly changed to support the required services. When the disaster is over, the system can easily be stowed and re-deployed at a later time when required.

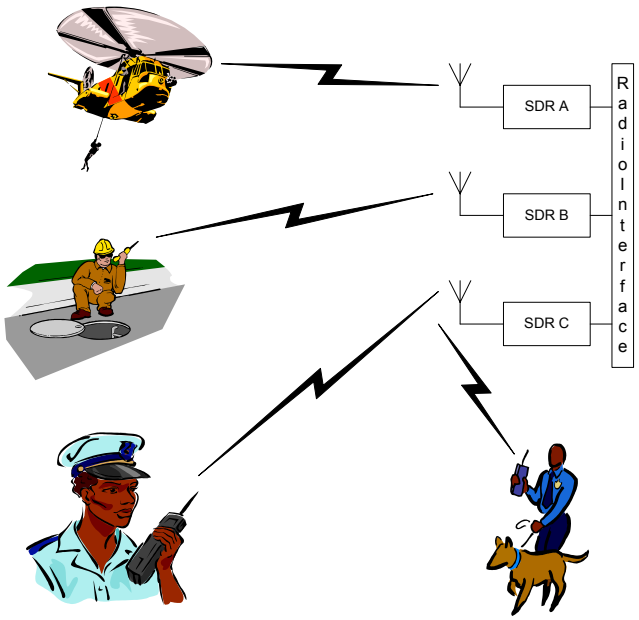

Figure 5: Interoperability

Interception

As outlined above, early applications were for military interoperability. Another military application for SDR is the interception of communications. Since the frequency and modulation format of these transmissions are often unknown, a flexible receiver platform capable of rapid self-adjustment is a benefit. Since an SDR can rapidly be reconfigured they are ideal for the interception of wireless communications [3]. Additionally, since they already employ high speed DSP, the DSP can also be utilized for advanced interception functions such as voice recognition and code decryption. Additionally, if a network of SDRs is used, then triangulation can be used to aid in the location of the rogue transmission.

Manufacturing platform

Although there are many applications where the dynamic configuration of an SDR is required, perhaps one of the most practical applications is that of a standardized communications platform. For example, most manufacturers of cellular infrastructure equipment sell platforms for a variety of air standards such as GSM, CMDA, IS-136 and AMPS to name but a few. Typically each of these is a different piece of hardware that must be manufactured and inventoried. If a single design could be fabricated that could have identical hardware, the cost of manufacturing could be significantly reduced because only one system would need to be inventoried. The hardware could then be configured prior to shipment or in the field for the air interface required. While it may not be practical to look to standards of the past in and of themselves, systems of the future are prime candidates. With the competing 3G standards a manufacturer with limited resources could build a single system capable of supporting either CDMA2000 or UMTS while continuing to support the legacy standards from which these have evolved. From a user point of view, such a system is also valuable because if the user should want to change standards, all that is required is that the system be reprogrammed for the new standard, preserving all of the investment made in the original equipment.

Of course other areas could benefit from the economies of scale offered here. Other such areas include devices for the reception of competing satellite broadcast of audio and video content, and two-way communications systems to name a few [4].

Architectures

Ideally the designer of an SDR would like to put the data converters directly on the antenna. However as stated previously, this is not a practical solution. In reality, some analog 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 many decades old, new semiconductor technology and high levels of integration have kept this architecture vitalized 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 basestation side. It is possible that future developments will enable direct conversion on the receive side as well. Until then, the super-heterodyne architecture will continue to be used in one form or another.

Receiver

High performance SDR receivers are typically constructed from some variant of the super-heterodyne 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 analog techniques and multiple front ends would allow operation across different RF bands. In the case of multicarrier applications, this could be done simultaneously if necessary.

Multicarrier

Depending on the applications, one or more receive channels may be desired. Traditional applications may only require 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 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 digitized. With this much bandwidth, it is also possible to process multiple channels. The figure below shows a typical multi-carrier receiver example. 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 channel 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.

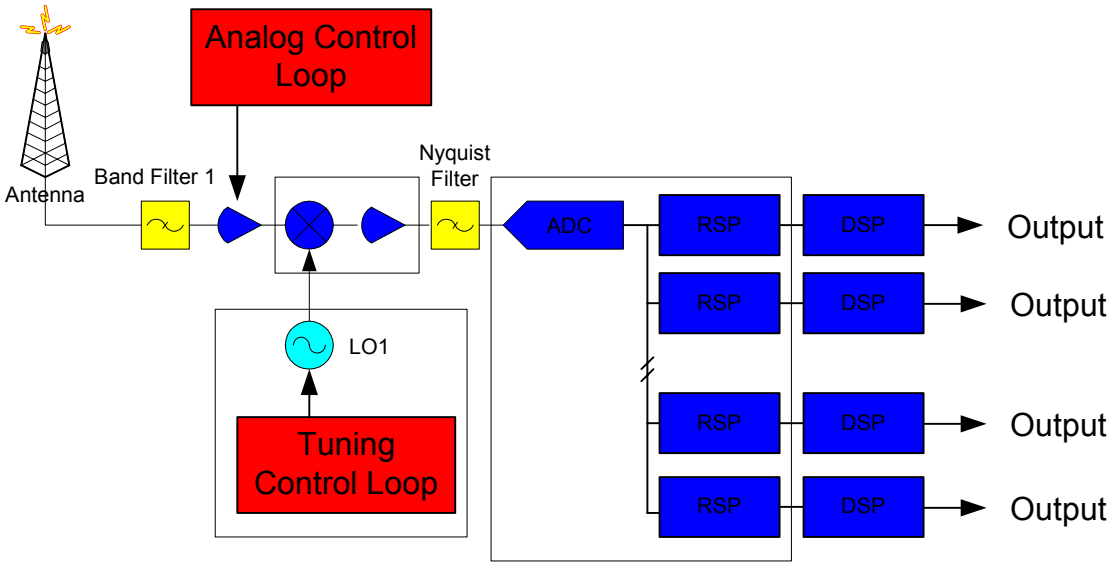

Figure 6: Multi-carrier CDMA example

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 pre-processing and the general purpose DSP are both 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 200 kHz 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 under-utilized, 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!

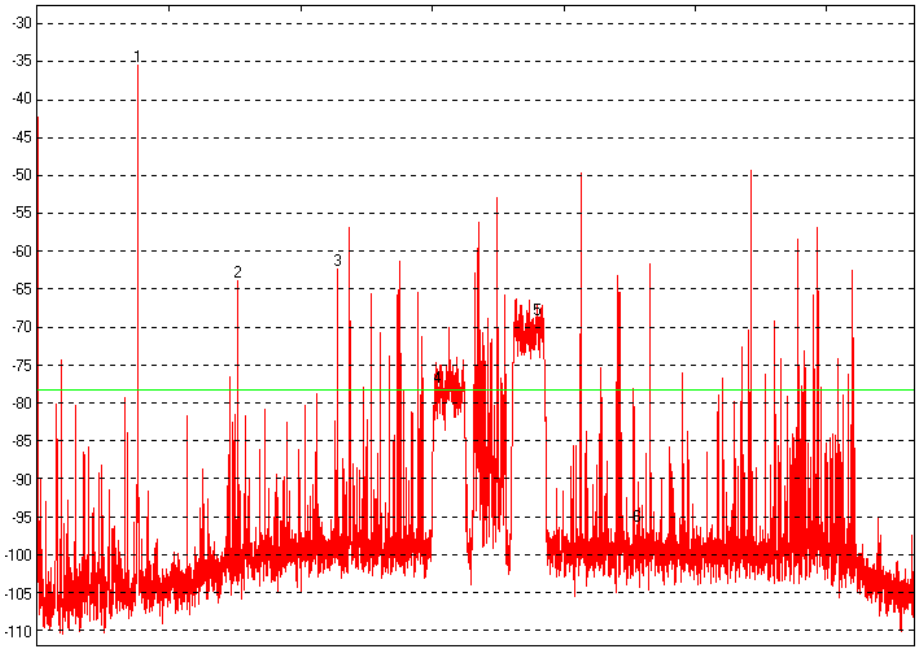

Figure 7: multi-mode spectrum with IS-95 and narrowband carriers

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 few kHz or as wide and 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 3 times the bandwidth. In this example, a sample rate of 30.72 MSPS or higher would allow signal bandwidths from a few kHz up to 10 MHz to be processed. Aside 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.

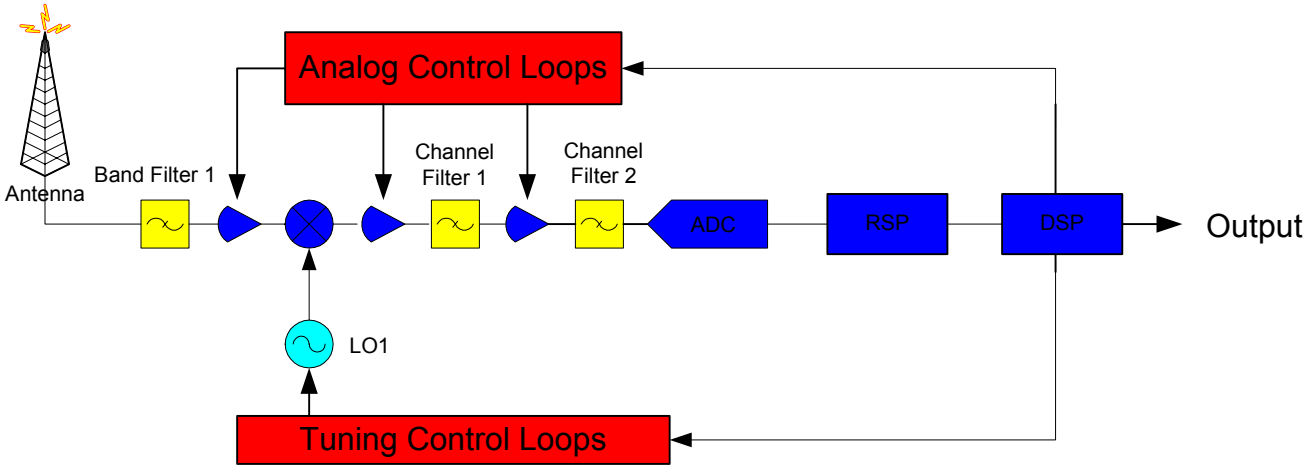

Figure 8: Single Carrier Rx Example

SDR Receiver Elements

Referring to the single carrier block diagram above (but keep 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, the antenna is one of the weakest elements in an SDR [1]. Since most antenna structures have a bandwidth that is a small percentage of it center 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 such as the SpeakEasy discussed earlier, the antenna must be tuned by some means to track the operating frequency to maintain operating efficiency. 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 potentially resulting 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 minimize 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. A 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 this 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. While only 1 mixer is shown in this diagram, many receivers may use two or three mixer stages, each successively generating a lower frequency. [Note: Receiver IF's are not always lower than the RF signal. A common example is found in HF receivers where the desired RF signal may only be 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. Analog channel filtering is not possible in the case of a multi-carrier receiver where the channel bandwidths are not known in advance. Therefore, 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 utilize 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 in not a problem, however, in the analog domain, 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 utilizing 'real' sampling (as opposed to complex sampling) as shown in the single carrier example and use a digital quadrature demodulator in the digital pre-processor 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 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 optimize 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 analog control loop. However, a digital control loop can also be used to implement difficult control loops not possible using analog 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 analog 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]. Often times, the designer is forced to select the best available ADC, realizing that under many conditions the ADC may be over specified. Still other times, air interface standards may not be directed towards multi-carrier receivers and require a 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 pre-processor 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 pre-processor capable of tuning the entire range of the ADC Nyquist band and filtering a signal with bandwidths from a few kHz to several MHz. When multiple channels are required, the design can be repeated to fill the FPGA. If a lower cost option is required, a variety of ASICs are available that perform these functions. They are often referred to as channelizers, 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 equalization, 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 as well as controlling all of the features in the other elements of the block diagram. As DSP clock rates increase, DSPs may well take over many of the functions within the digital pre-processors.

Transmit

Transmit functions like the receive are typically based on some form of super-heterodyne or direct conversion. The figures below illustrate these two options. The multi-carrier option is best suited to single and multi-carrier applications while the direct conversion offers an excellent, low cost solution for single carrier applications. As integration technology improves, multi-carrier direct conversion may become a possibility, however, such a transmit configuration requires sideband suppression about 15 dB better than the spurious requirements to prevent images on one side of the center frequency from overpowering a potentially weak carrier on the other.

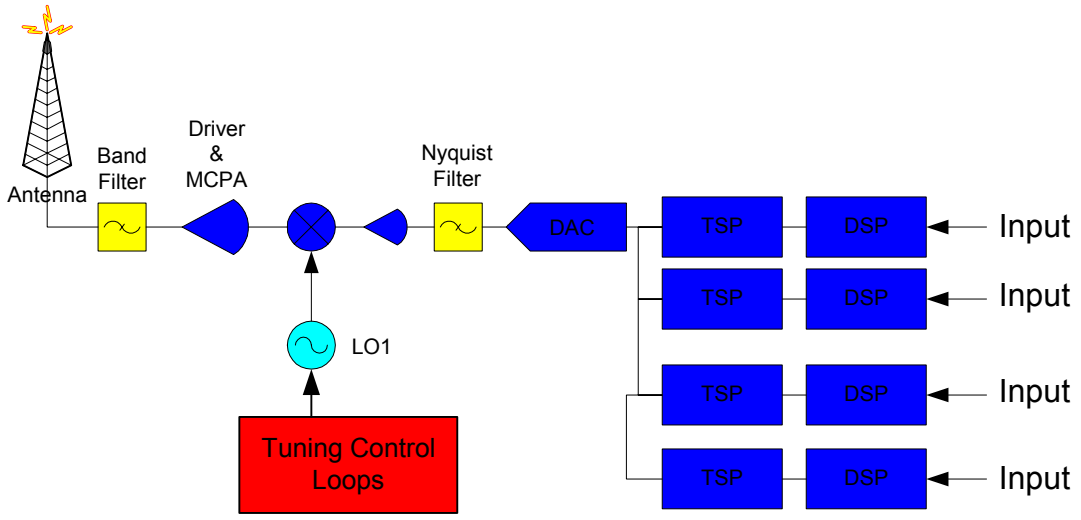

Figure 9: Multi-channel transmit with single up-convert super-heterodyne

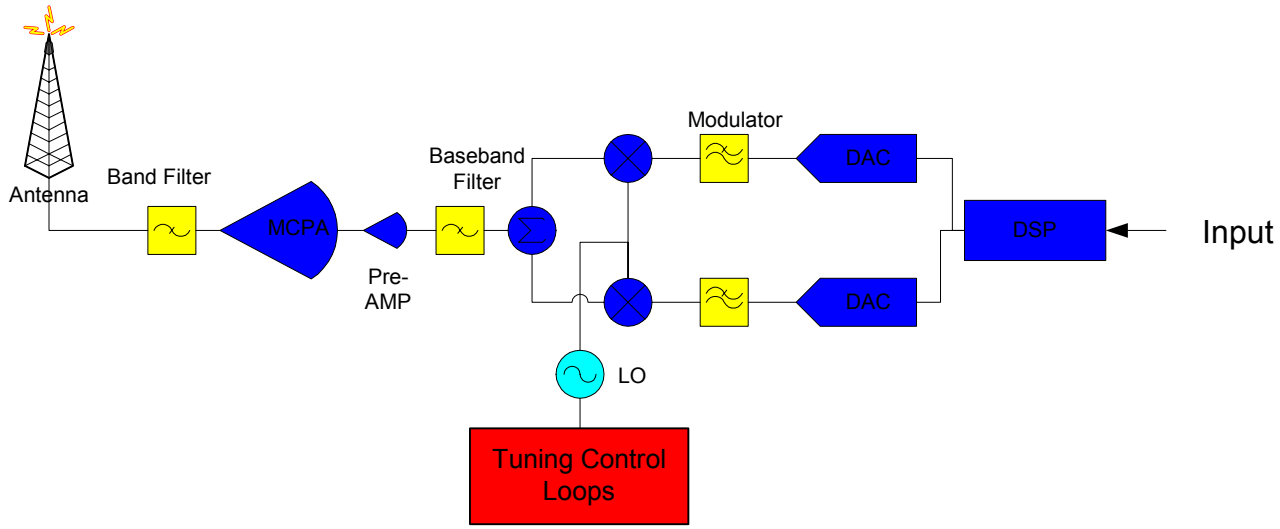

Figure 10: Single carrier direct conversion transmit

In either application, a DSP or baseband ASIC is used to generate the modulated baseband data. This data is fed either directly to a pair of baseband DACs (I and Q) for direct RF modulation or to a digital processor responsible for digitally translating them to a suitable digital IF. Depending on the application, the DSP alone or in conjunction with digital processor can be used to digitally pre-distort the baseband data in such a manner that distortion products generated later in the signal chain will be cancelled.

If an IF stage is employed, the baseband data generated by the DSP must be up-converted either digitally with an FPGA or ASIC (also know as TSPs or DUCs) or alternately with a traditional mixer or modulator to the desired IF. This traditional technique is being replaced by digital means because of the added flexibility offered through digital logic and the availability of good cost effective digital to analog converters. As with the related receive function, the purpose of this device is to shape the bandwidth of the desired channel and then up-convert by digital means to the desired IF frequency. If

multiple channels are required, they can be synthesized on one chip. After translation, each of the channels can be summed together and interpolated to the desired data rate and then sent to a DAC. If desired, digital pre-distortion can be added in conjunction with the DSP to correct for distortion later in the signal chain.

Either a mixer or a modulator is used for frequency translation to the final RF frequency. If direct RF modulation employed, an RF modulator will be used. If an IF is used (either directly from a DAC or a traditional IF up-conversion), a mixer will be used to translate to the final RF frequency. As with the receive mixer/demodulator, it may be desirable to change the bias levels or the drive level of the data or LO levels to optimize distortion.

As with the receive LO, the transmit LO is variable in frequency and easily programmable via software control using PLL or DDS techniques. Here too, it may be desirable to change the LO drive level to optimize spurious performance under a variety of signal conditions. As with the single band operation of the receiver, there may be cases where a fix LO is required. Such an example would be for operation within a single band where tuning is accomplished within the TSP, DUC or FPGA.

As with the receive path the data converter or DAC is often the bottleneck. However since dynamic range requirements for the transmit signal path are much lower (typically 25 to 45 dB) than the receive path, component selection is not quite is difficult. Many DACs are available that facilitate a wide range of adjustments include gain and offset correction so that I/Q imbalances in the transmit signal chain can be minimized. Other desired features include data rate interpolation and I/Q phase correction.

Finally, power gain is achieved through a pre-amp and PA. Aside from the fact that these devices must operate across a wide range of frequencies, it is desirable to adjust the RF output power. There could be regulatory issues that require some frequencies to be transmitted at lower power than others. While the PA gain is usually fixed, the pre-amp may be in the form of a VGA.

Implementation issues

The reality is that without improvements in semiconductor technology through the late 1990's, SDR as outlined above would still be a concept and not a reality. Although the evolution of DSP technology has certainly been key to SDR, it is not the only technology that has had to 'grow up' in order to support this development. Because the level planning is different in these systems, most of the analog components are stressed to a higher degree and better performance is required than that found in traditional architectures.

Analog Front End

Quick conversion

In order to take advantage of digital processing, a software-defined radio seeks to convert from the RF domain to digital domain as soon as possible. This is true both for receive and transmit. By doing this, as much of the processing can be done digitally as possible.

Advantages

If most of the processing is done digitally, then reconfiguration can be quite simple. Filter coefficients can be changed, different software run or even in the case of FPGAs, they can be completely reconfigured for the required format. In the analog domain, space and resources limit the reconfiguration options available.

In the analog domain, only a small number of modulation/demodulation schemes are possible. However, in the digital domain, the possibilities are limitless if the functions are configurable through software. Even where complex functions are implemented in the analog domain, various errors such as quadrature errors can cause performance issues.

In the digital domain, quadrature (and other functions) can be exactly generated. Once in the digital domain, the accuracy of the function is limited only by the bit precision of the math used to implement the function. For example, it is always much easier to add a few more bits to a multiply than it is to improve the linearity of an analog mixer.

Finally, because most of the signal chain is digital, performance would be more consistent for each system manufactured, eliminating much of the product variation and yield. Since the performance is more consistent, many of the factory trim and alignment issues would be eliminated, potentially reducing a large part of the manufacturing and test cost.

Higher gain at RF frequencies

These advantages are good for providing a more consistent transceiver at a lower production cost. However, as already mentioned, the level planning for an SDR is sufficiently different than a standard super-heterodyne that different strategies are required.

Level planning

In a traditional super-heterodyne transceiver, the conversion gain is distributed throughout the signal chain. Typically, gain in the front end is balanced between high enough for a low NF but not so high as to overdrive the remainder of the signal chain and cause excessive intermodulation distortion. Similarly, as much gain is run after the channel filtering as possible so that interfering signals have already been eliminated. Throughout the signal chain, only enough gain is used to offset the accumulated loses and

to prevent those elements from significantly contributing to the overall NF, thereby allowing for IP3 to be carefully balanced against NF.

By contrast, in an SDR many of the intermediate stages have been eliminated because sampling occurs as close to the antenna as possible. Since most SDR applications are also wideband, there are no channel filters thereby allowing many of the neighboring signals to also pass the signal chain. Because all of the conversion gain must occur in the presence of many signals, intermodulation performance is inherently more important.

In a transmitter with multi-stage up conversion, the issues are very similar to the super-heterodyne receive above. As with receivers, both noise and intermodulation are very important specifications. However, the active dynamic range of most transmit requirements are only on the order of 60-70 dB whereas most receivers require 100 dB or more of dynamic range. The real difficulty of the transmit signal path is maintaining the noise and linearity performance in the RF power amplifiers when the power level reaches several hundred or even several thousand watts. The discussion of the PA is beyond the scope of this discussion.

Fixed versus variable gain

In a traditional receiver, the total conversion gain is quite often variable. This reduces the required dynamic range of the circuitry following the AGC, reducing the required linearity requirement of those components. Also, the AGC action allows for optimum signal levels over a wider range of input signal condition. While an AGC is still quite useful for SDR, there are certain restrictions on their use.

Tradeoffs vs. MC and SC

In a single carrier receiver, there are two main issues with setting the gain. The first issue is that it is desirable not to overdrive the front end when a strong signal of any frequency is present. After this issue is accounted for, the conversion gain of can be increased or decreased as necessary to achieve the sensitivity required. While it is possible that a nearby signal will pass the front-end filters of the receiver and cause a reduction of the gain and subsequent loss of the desired signal, this is typically managed through the use of ‘band select’ or RF tracking filters that filter all but the desired signals. However SDR and multi-carrier receivers typically have a ‘wider’ front-end bandwidth and therefore allow many more signals to pass full analog signal chain. As a result it is much easier for a strong signal at one frequency to desensitize the desired signal at another frequency. Since the receiver has a limited noise floor (thermal and other noise sources) the gain can only be reduced to the point that the weakest signal retains the minimum SNR required for detections. Since the design has already been configured for multiple carriers, it is likely that the gain has been reduced to a minimum so that the largest expected signal will not overdrive the signal chain. Because the gain is limited, the noise floor of the receiver becomes limited by that of the data converter.

Tradeoff vs. converter resolution

In a traditional receiver, if the signal level was not large enough to be adequately detected by the ADC, then additional gain is used to boost the level above the ADC noise floor using an AGC topology. However, it is just as valid to lower the noise floor of the converter. There are several ways to do this as discussed below. The easiest is to just specify a converter with more bits. Unfortunately, the more bits a converter has, the more expensive it is and the more power it dissipates. Therefore, balancing the conversion gain and converter resolution is a very important task; too much gain and the ADC is overdriven, too little gain and the ADC directly sets the noise floor which is an undesirable situation [11]. Ideally, the conversion gain of the receiver should place non-ADC noise 10 dB above the ADC noise. Therefore, given an ADC converter noise floor, an ideal minimum gain can be determined that prevents the ADC from dominating overall noise performance.

$$G_{\max} = P_{ADC\_Fullscale} - P_{\max-signal}$$

EQ 1: Maximum Conversion gain in dB

$$G_{\min} = NSD_{ADC} + 10 - 10\log\left(\frac{k*T*1Hz}{0.001}\right) - NF_{ana\log\_front\_end}$$

EQ 2: Minimum Conversion gain in dB

These equations outline the desired maximum and minimum conversion gain. To achieve gains beyond these bounds, an AGC can be carefully used. Even with an AGC, the data converter will determine what the instantaneous dynamic range of the transceiver will be based on the difference between the noise floor and the fullscale of the converter. If G_{min} is greater than G_{max}, either the fullscale of the ADC must be increased or the NSD (noise spectral density) of the ADC must be lowered indicating than a better converter may be required. While there are means of increasing the maximum input of the converter or reducing the noise floor of the ADC, it is often easier to specify an ADC with better performance or more bits of resolution.

Higher IP3 requirements

While neither option comes easily, increasing the fullscale of the converter may have other undesirable consequences. If the input range is increased, then larger signal swings are required to take advantage of this increase range. This implies that high signal powers are required. Therefore, in order to keep the intermodulation products at the same level, the IP3 specification of the drive circuitry must also be increased to take full advantage of the increase signal range, otherwise, what signal dynamic range is gained will be quickly lost to increasing intermodulation distortion, most notably the 3rd order products. For IF sampling, even order intermodulation products are generally not a problem because they most often fall away from the signal of interest and are easily filtered from the spectrum of the input to the ADC input.

Signal dynamic range

Signal dynamic range is the difference between the largest and the smallest signal that can be detected. If the receiver is properly designed, the ADC will largely dominate this. The fullscale range of the ADC will determine the largest signal as already established by reworking equation 1. Likewise, the smallest detectable signal will be set directly or indirectly by the noise floor of the converter. Ideally, the noise from the analog front end will dominate the total noise because it has been placed as much as 10 dB above the converter noise floor. If the front-end noise floor is much less than 10 dB, then the contribution to total noise from the ADC will increase and must be included in the overall noise calculation of the receiver.

The largest signal is determined by the fullscale of the ADC and the applied conversion gain. Similarly, the smallest signal to be detected can be calculated by the noise in the channel of interest

$$XDS = P_{ADC\_Fullscale} - G$$

EQ 3: Maximum Detectable signal before clipping

$$MDS = 10\log\left(\frac{k*T*BW_{signal}}{0.001}\right) + NF_{cascaded\_total} + SNR_{required}$$

EQ 4: Minimum Detectable signal

For example, if the ADC has a fullscale of +4 dBm and the conversion gain is 35 dB, the maximum input power is –31 dBm. Similarly if the channel of interest is 200 kHz wide, the total NF is 3 dB and the required SNR is 5 dB, then the MDS is –112.8 dBm. Taking the difference between equation 3 and 4 will estimate the dynamic range of the receiver. In this example, the dynamic range is found to be 82 dB.

$$DR = P_{ADC\_Fullscale} - G - 10\log\left(\frac{k*T*BW_{signal}}{0.001}\right) - NF_{cascaded\_total} - SNR_{required}$$

EQ 5: Signal Dynamic Range

There are many factors that will reduce both the MDS and the DR. A key point to remember is that as shown here, XDS, MDS and DR are ‘static’ tests and in reality that more than one carrier may share the dynamic range of a multi-carrier receiver. Because of constructive interference, the fullscale power of the converter must be shared between each of the signals, thereby effectively reducing the largest possible input signal. As a guideline, if all signals are at the same level, each time the number of carriers is doubled, the largest any of them can be is reduced by 3 dB of power. For example, if two signals are present, each signal must be 3 dB below the clipping point. If 4 are present, 6 dB and 8 must have 9 dB and so on. Therefore, for applications where many signals are present, the effective dynamic range is limited.

Similarly there are reasons that the noise floor will increase above that calculated in the static equation above. One example of this is seen through reciprocal mixing between the phase noise of the local oscillator and a nearby blocking signal resulting in an increase in the noise floor of the mixer. A similar example is seen in the increase in the converter noise floor associated with the 'reciprocal mixing' between the same blocker and the aperture jitter of the ADC. Fortunately, if the converter noise floor is adequately below that of the analog front end, variations of several dB in the noise of the ADC will have only limited effects in overall performance. If however, the ADC was designed to dominate overall noise or guarding of much less than 10 dB was used between the two, the effects on the overall receiver performance should be revisited to determine the effects of ADC noise versus a variety of signal conditions [11].

IF selection/tradeoffs

Selecting an IF frequency for a traditional single carrier receiver can be challenging enough. However, in a multi-carrier receiver, traditional issues such as determining what the level of the high order intermodulation products of all signals is complicated by the fact that now entire bands are being translated from one frequency to another. This problem is further complicated by aliasing within the ADC. In a typical single carrier receiver, the IF signal into the ADC is chosen such that any aliased harmonics of the input signal fall away from the input signal. This is important at low signal levels because when an ADC is stimulated by very low-level inputs, it is possible that spurious generated within the ADC can be larger than the desired input. Thus if the harmonics are designed to fall away from the desired input, this problem can be averted. However, with a multi-carrier receiver, the harmonics can cover a very wide band of frequencies. Generally, even order harmonics are not much of a problem. At very low signal levels, the key harmonics are the 3^{rd} and 5^{th} harmonics of the input spectrum. Since the input may be a wide band, the third and fifth harmonics are 3 and 5 times as wide respectively. Given this, it becomes difficult to try and place these signals in a part of the spectrum where they will cause no problems. In cases where careful signal planning is not possible, dithering techniques provide relief to many of these problems [12]. The graph below shows how it is possible to place some harmonics out of band if sufficient over sampling is possible. In this example, the second and third harmonics are placed in such a manner that the aliased components fall away from the desired fundamentals.

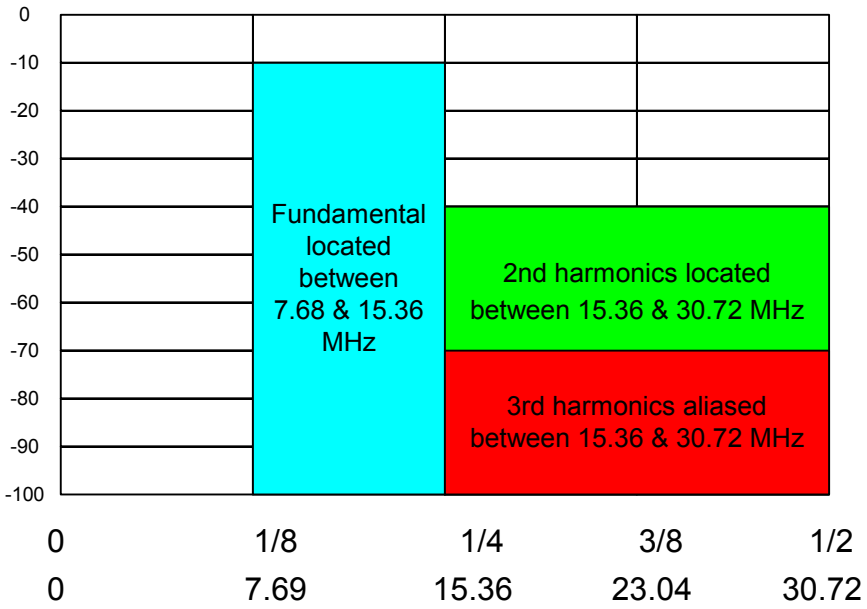

Figure 11: Aliased wideband signal and harmonics

Transition band requirements

Since all of the channel selection in an SDR is done in the digital domain, the analog filter requirements are different. Their purpose is primarily to prevent the overlap of images either in the mixing process or the sampling process. In mixer stages, care must be taken to suppress the undesired mixer images. In the ADC, signals both above and below the desired band may be sampled and aliased into the usable spectrum of the ADC. For either mixer images or aliased signals within the ADC, these signals must be filtered so that they are below the minimum detectable signal. If they are not, then it is possible that they will overpower the desired signal. In the case of the ADC, this required rejection must be achieved before the aliasing becomes critical. Therefore as shown in the diagram below, the full rejection of the undesired signals must be achieved before the spectrum is folded upon itself.

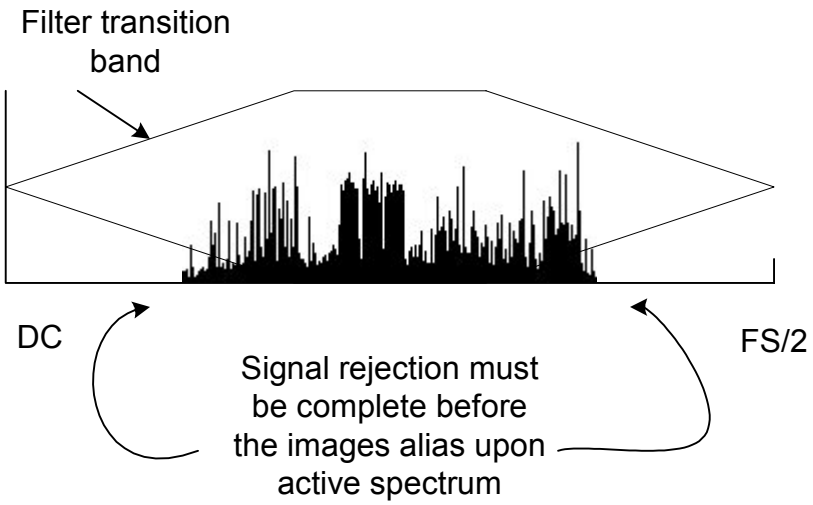

Figure 12: Aliased transition bands

Data Conversion

For both transmit and receive, the data converters are usually key components in the signal plan. The key elements for both are the dynamic range, which is bound on one end by the noise floor of the converter and the maximum input, or output range on the other. Other equally important issues include distortion, both harmonic and intermodulation. Although related, have somewhat different effects in the limitation of performance.

There are many different converter topologies and each has their benefits and limitations [13, 14]. While there is no set architecture that provides better performance for SDR, the selection is best made based on the performance requirements for the application and then a study of the available data sheets of the potential converters. Since data sheets cannot fully represent the actual performance of a data converter, it is always best to take a look at them on the bench in an environment similar to that of the end product.

Because data converters are somewhere between the digital and analog domains, they are often poorly understood by both the analog designer and the digital designer. For this reason, their effects on transceiver design are often over estimate, under estimate or both in different areas. The next few topics will sort through many of the issues to help determine exactly how converter performance determines performance.

General converter requirements

Bits & Noise and a little signal

The number of converter bits is the most visible specification. From a mathematical point of view the number of bits the converter contains limits performance physically. An ideal converter will exhibit an SNR that is determined by the equation below.

$$SNR = 6.02N + 1.8$$

EQ 6: SNR = 6.02N + 1.8

In reality, there are many other issues that determine the performance of the converter including clock jitter and thermal noise. This equation provides the noise due to ideal quantization and does not take into account any of the other sources of noise in a converter. A modification of this equation provides a more insightful measure of converter performance. This equation takes into account clock jitter, thermal noise and the effects of non-ideal quantization, which are the major limitations in converter performance.

$$SNR = -20\log\left[\left(2\pi F_{ana\log} t_{j_{rms}}\right)^2 + \left(\frac{1+\varepsilon}{2^N}\right)^2 + \left(\frac{v_{noise_{rms}}}{2^N}\right)^2\right]^{1/2}$$

Figure 7: Modified converter SNR

Although this equation is a little more complicated, it does take into account many of the important factors in converter performance. In the equation, $F_{ana\log}$ is the analog IF frequency, $t_{j_{rms}}$ is the aperture uncertainty, ε is the average DNL of the converter, $v_{noise_{rms}}$ is the thermal noise of the converter and N is the number of bits.

Many data sheets will include information on the effective number of bits that a converter possesses. This number is usually determined by solving the SNR equation above for N, the number of bits. While an effective bits measurement is a convenient comparison tool, it has limited usefulness in radio design. SNR is a better tool because of its direct link to noise. A better still measurement is that of noise spectral density or NSD for the converter. NSD provides the amount of noise energy in a 1 Hz bandwidth for a converter. This number is not usually specified in the data sheet because it is dependent on the actual sample rate used and the input termination condition.

$$NSD = P_{ADC\_Fullscale} - SNR_{ADC\_Fullscale} - 10\log\left(\frac{Sample\_Rate}{2}\right)$$

EQ 8: NSD of a data converter

Once the noise spectral density has been determined, it can be used to either validate if the converter meets the noise floor requirements or to determine the minimum gain required from the front end of the transceiver design.

As an example, an ADC is selected with an SNR of 70 dB at the selected input frequency. With the required input termination, a fullscale of +4 dBm is achieved. The configuration requires a sample rate of 61.44 MSPS. Using the equation above, the noise spectral density is –140.9 dBm/Hz.

As outlined above, it is desirable that the converter noise not limit the performance of the transceiver, then the front-end noise needs to be approximately 10 dB higher than that of

the converter. Therefore, to safely use this converter, noise generated by the front end must be about –131 dBm/Hz. Modifying equation 2 above because the NF of the analog front end is not yet known, this gives the equation below. With this equation, the combined front-end gain and noise figure must be 43 dB to ensure that the ADC does not dominate performance.

$$G_{\min} + NF_{ana\log\_front\_end} = NF_{ADC} + 10 + 10\log\left(\frac{k*T*1Hz}{0.001}\right)$$

EQ 9: Gain and Noise requirements

Noise Figure for an ADC

If it is not possible to design the system such that the converter noise is significantly below the remainder of the system, then the noise must be included in the calculation. This can be accomplished by using equation 9 above to determine the noise from the ADC or the NF of the converter can be calculated and included in the cascaded NF of the signal chain analysis with the other linear devices. While an ADC is not a power device, the NF can be estimated and should only be considered valid for the set of operating conditions specified. Therefore, if the conditions are changed, then the NF will change appropriately.

$$NF_{ADC} = P_{ADC\_Fullscale} - SNR_{ADC\_Fullscale} - 10\log\left(\frac{Sample\_Rate}{2}\right) - 10\log\left(\frac{k*T*1Hz}{0.001}\right)$$

EQ 10: Equivalent ADC Noise Figure

If it is determined that the conversion gain required to offset the converter noise is large enough that the converter is overdriven, this is an indication that a better converter is required.

Channel Noise in a receiver

Once the total receiver noise level has been determined, sensitivity of the receiver can be found. If conversion gain is known then sensitivity with respect to the antenna can be found, otherwise, it will be with respect to the ADC input. In the typical SDR signal chain, a digital tuner or channelizer will follow the ADC. In this block, the desired signal is tuned and all other signals in the Nyquist band are filtered from the spectrum. Typically the data rate is also reduced to a speed that is suitable to the data rate of the modulation being carried. If a quiet channel is selected all that should be on the output of the channelizer is the noise from the analog front end plus the ADC. Since the noise spectral density has already been established in a prior section, the total noise in the channel can be determined by integrating this over the channel bandwidth. In log math, this equation is very simple.

$$N_{channel} = 10\log\left(\frac{k * T * 1Hz}{0.001}\right) + G_{\min} + NF_{total} + 10\log\left(\frac{BW_{channel}}{1Hz}\right)$$

EQ 11: Receiver channel noise

With the total integrated channel noise and the required SNR for the modulation standard, the reference sensitivity can be determined. Keep in mind that the required SNR may be positive or negative depending on the amount of digital gain provided by detection algorithm. As an example, GSM requires about 5 dB SNR while IS-95 requires an SNR of about –16 dB.

Continuing the example from earlier, if the analog front end generate a noise density of –131 dBm/Hz (the ADC is 10 dB below this and not a contributing factor) and the channel bandwidth is 200 kHz, then the total channel noise is –78 dBm/200 kHz. If the required signal level is 5 dB above this, the smallest signal that can be detected as presented to the ADC will be –73 dBm. If the conversion gain of our signal chain is known, then the sensitivity at the antenna can be calculated. In order to achieve the noise of –131 dBm/Hz, a gain plus NF of 43 is required. At this point the NF is not known, but may be estimated based on available technology. A good typical NF would be about 3-4 dB. This would place the conversion gain at 40 dB. Therefore if the –73 dBm signal is referred back to the antenna, it will be 40 dB smaller or –113 dBm, a very good sensitivity for a channel 200 kHz wide.

Digital Processing

Parts

There are two categories of digital parts. Both could be called digital signal processors. The traditional DSP is a computational unit that consists of program and data memory. A program is executed from the program space operating on data from I/O ports and data stored in the data memory. This type of DSP is the most common, however, this type of DSP is limited in the data throughput. While great advances in parallel computing and core speeds have increased the rate at which real time data can be processed, general purpose DSPs can only process limited amounts of data.

Fixed Function DSP

To augment the processing power of a general purpose DSP, fixed function DSPs are designed to process very large amounts of data very fast and efficiently. While a general purpose DSP can be infinitely reprogrammed, the signal flow within a fixed function DSP must be restricted to a single architecture. Programming is also limited to configuration registers and memory coefficients. However, since most radios are based on some form of super-heterodyne architecture, this is not such a limitation due to the high degree of similarity between different designs. Therefore, the fixed function DSPs can be designed to represent a very large class of receiver or transmitter designs. These fixed function DSPs are often implemented either in FPGAs or ASICs. Processors

designed for the receive function are called Receive Signal Processors (RSP) and transmit functions are called Transmit Signal Processors. In general, both TSP and RSP contain exactly the same elements, only the order is reversed. For either device, there are three key sub-functions found in these devices.

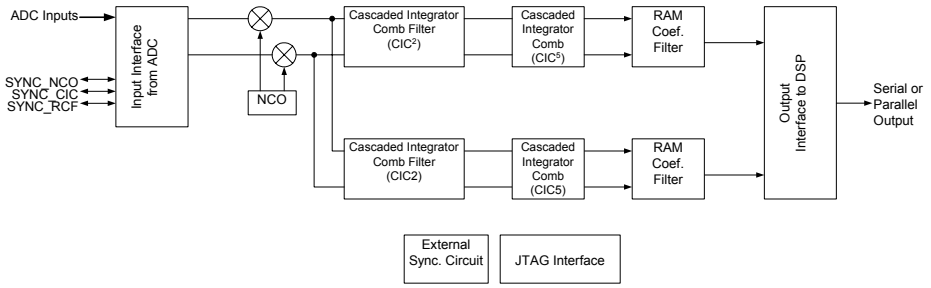

Figure 13: Fixed Function Receive Signal Processor (RSP)

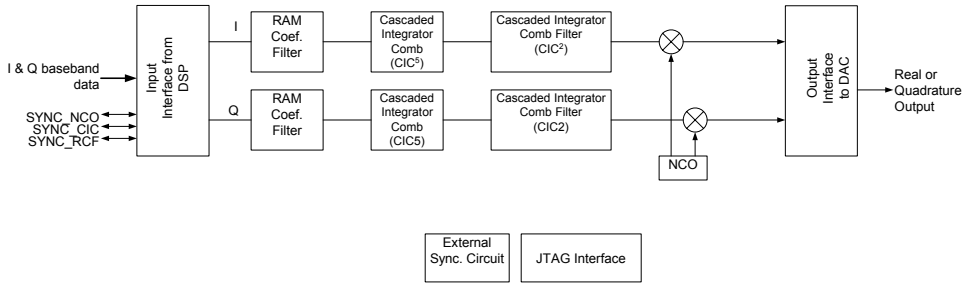

Figure 14: Fixed Function Transmit Signal Processor (TSP)

The first function is the frequency translation. In the analog domain, the frequency is translated with a mixer or modulator/demodulator. This function is used to mix two inputs together in such a way that the sum and difference frequencies are generated on the output. In the digital domain, this is represented by a multiplication. If the function is a mix, then a single 'real' multiply is performed, but most often, the multiply is a complex multiply used to generate quadrature data and thus separate positive and negative frequencies. In an IF sampling receiver, a real digital IF is applied to one of the complex inputs of the multiply. The other input is the output from a complex Numerically Controlled Oscillator (NCO). The NCO is tuned to the desired frequency such that the result is a complex signal at DC and at the sum frequency.

Following the NCO and complex mixer (demodulator) is a low pass filter. This filter serves two purposes. The first purpose is to remove the undesired noise, signals and spurious. By doing so, the all of the wideband noise on the output of the ADC is removed except that which lies within the passband of the filter, giving rise to what is often referred to as processing gain. Second, the filter shapes the passband characteristics. Quite often, the passband must be a matched filter or otherwise shape the characteristics of the incoming spectrum. This is easily accomplished with the digital filter, a task that is often difficult with analog channel filters. In fact, since these filters are digital, they can implement any filter that can be realized using FIR or IIR techniques.

$$G_{Noise\_Processing} = 10\log\left(\frac{sample\_rate/2}{BW_{channel}}\right)$$

EQ 12: ADC processing gain

Following channel filtering, the bandwidth will be relatively small compared to the data rate because of the high oversampling rate in the ADC. Therefore, it is advantageous to reduce the data rates. This has several benefits. First, with the reduced data rate, the computational burden on the general purpose DSP is reduced. Second, in CMOS technology, lower data rates results in lower power. Therefore, following the channel filters, data decimation is performed. The decimation must be consistent with the Nyquist, but significantly reduces the computation by the general purpose DSP that follows the RSP.

In the transmit direction, the data flow is reversed. First the data is filtered and then interpolated to the desired data rate. Then, the data is translated to the proper frequency using a modulator and complex NCO.

For a typical RSP/TSP channel, the computational load may be as high as 1.5 Giga-operations per second. If multiple channels are required, then the process scales linearly. At the present time, this load exceeds the capabilities of a general purpose DSP, however, as DSP technology improves, it may be possible to take on some or all of the processing in the future.

General purpose DSP

General purposes DSPs like microprocessors are designed to execute a software program. The software for a DSP is developed in the same manner as that for a microprocessor using program languages such as 'C' and assembly. However, DSPs are designed specifically to execute code as fast as possible. In fact, DSPs are usually designed to execute programs or tasks in real time. Since the DSP may be processing real-time data such as voice or video, program execution is required to keep up with incoming data, otherwise, throughput will be sluggish, intermittent or simply come to a halt as the DSP struggles to keep up with incoming data.

To prevent this from happening, DSPs are especially designed to improve data throughput taking advantage of a number of techniques. Often, one vendor will focus on one technique and refine that while a different vendor will focus on a different optimization technique. Both result in faster throughput, but with slightly different advantages. Some of these techniques are listed in the table below.

Program and Data look-ahead caching
Multiple Address Generation
Separate Program and Data Memory
Multiple Arithmetic Logic Units
Multiply Accelerators
Shifters
Hardware Accelerators
On-chip Memory
Floating/Fixed Point Optimization

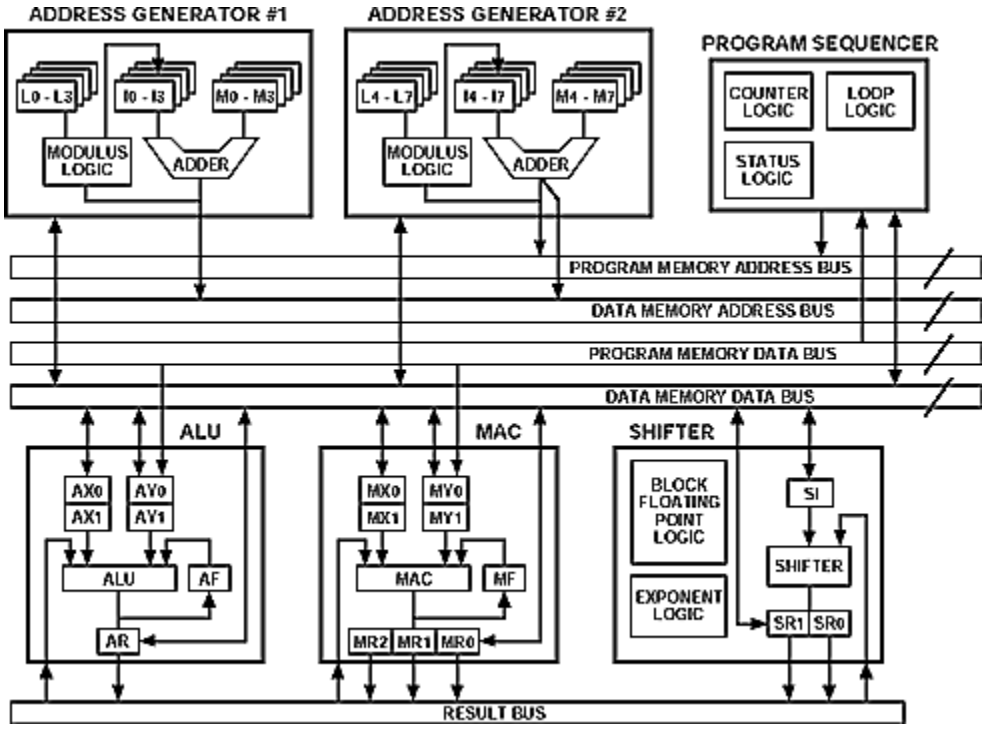

Figure 15: General Purpose Digital Signal Processor (DSP)

In an SDR, the general purpose DSP is generally tasked to perform the Nyquist rate processing. That is the signal processing required at data rates that supports the Nyquist rate of the raw data. In our SDR application that may support a channel as wide as 10 MHz, the actual data rate may be as high as 20 MHz. While not all applications may require this much processing, some applications may. The actual processing requirements will depend on the application and functions instantiated. As with other components, if a wide range of processing is expected, the design has to consider the maximum requirement, even in the case where excess processing capability may exist in some operating modes.

Envelop Detection (AM)
Phase/Frequency Detection (PM/FM)
Phase/Frequency Correction
Equalization of a TDM burst
Spread/De-spread a CDMA signal
Voice Coding/Decoding

The table above lists just a few of the functions that are typically performed. Since the DSP is programmable, any function that can be coded can be executed. Additionally, since the code is software, it can be upgraded or changed at any time to further support the SDR.

Case Study – A close look at a CDMA2000 & UMTS SDR receiver

Now that many of the facets of SDR have been discussed, the final section will cover an example of a multi-carrier SDR receiver. While this is not a full analysis, it will cover many of the issues that surround the design and development of an SDR that are not covered in a typical receiver design. As with any design, the first place to start is with the specifications. The table below summarizes a few of the critical specifications for both CDMA2000 and UMTS (WCDMA).

	CDMA2000	UMTS
Reference Sensitivity	-117 dBm	-121 dBm
Bandwidth	1.25 MHz	5 MHz
Chip Rate	1.2288 MCPS	3.84 MCPS
Signal Spectral Density	-177.9 dBm/Hz	-186.8 dBm/Hz
Sample Rate	61.44 MSPS 50x oversample	61.44 MSPS 16x oversample
De-spread and Coding gain	21 dB	25 dB
Narrowband Blocker	-30 dBm	Na
CDMA Blocker	-52 dBm	-40 dBm
Two-tone blocking	-45 dBm, 2 tone CW	-48 dBm, 1 CW and 1 CDMA

The goal of this exercise will be to design a multi-carrier, multi-mode, single band receiver RF through baseband that is capable of processing both of these standards either independently or at the same time. Such a design would be useful for manufacturers of 3G basestation equipment where it is desirable to have a single piece of hardware that is capable of processing both standards, thereby eliminating duplicated design efforts.

One of the most direct ways of accomplishing this is to compare the two specifications and determine which will limit performance. One of the first issues will be to determine the largest signal that requires processing. The CDMA2000 standard calls for an –30 dBm narrowband signal where as UMTS does not address narrowband blockers. However, it does require that a –40 dBm CDMA signal be correctly processed. While narrowband signals can often be considered to have little envelop, a CDMA signal has between 10 and 12 dB of peak to rms on the envelope. Therefore, a CDMA signal of –40 dBm actually peaks very close to –30 dBm. Therefore, both standards require about the same peak signal capacity.

Since we know that we will need to digitize the signals, initial ADC characteristics may be established. Later in the analysis, the speciation can be validated to determine if the assumptions were correct. Since high performance data converters are expensive, it is desirable to use the lowest performance possible that allows the specifications to be met. From the table of typical converter specifications, the fullscale input range is 2 volts peak to peak differential. If this input is terminated with 200 ohms differentially, the rms power to drive the converter to fullscale will be +4 dBm. Similarly, the converter SNR is 75 dB and the SFDR both single and two tone is –95 dBFS. This performance is maintained out to analog frequencies of 100 MHz providing flat performance.

From this information an initial estimate of the conversion gain required can determined using equation 1.

$$G_{\max} = P_{ADC\_Fullscale} - P_{\max-signal} = +4 - -30 = 34dB$$

In order that the receiver not be overdriven, the conversion gain will be limited to 30 dB. For the moment assuming that the noise figure of the front end, less the ADC, will be 3 dB, the thermal noise delivered to the ADC can be determined. At room temperature, the thermal noise can be calculated to be

$$NSD_{Ana\log\_front\_end} = 10\log\left(\frac{k*T*1Hz}{0.001}\right) + G + NF_{Ana\log\_front\_end} =$$
$$-174dBm/Hz + 30 + 3 = -141dBm/Hz$$

Based on the information above, the NSD of the ADC can be determined using equation 8.

$$NSD = P_{ADC\_Fullscale} - SNR_{ADC\_Fullscale} - 10\log\left(\frac{Sample\_Rate}{2}\right) =$$

$$4dBm - 75dB - 10\log\left(\frac{61.44MHz}{2}\right) = -145.9dBm/Hz$$

Since the NSD of the ADC is less than 10 dB better than the NSD of the analog front end, the noise contributed from the ADC must be included in the overall noise analysis. Therefore, using equation 10 will provide the equivalent NF of the ADC for the configuration used here.

$$NF_{ADC} = P_{ADC\_Fullscale} - SNR_{ADC\_Fullscale} - 10\log\left(\frac{Sample\_Rate}{2}\right) - 10\log\left(\frac{k*T*1Hz}{0.001}\right) =$$

$$4dBm - 75dB - 10\log\left(\frac{61.44MSPS}{2}\right) - 10\log\left(\frac{1.38\times10^{-23}*300*1Hz}{0.001}\right) = 28dB$$

Based on this information and commercially available components, the level planning in the figure below can now be generated. This design features double conversion in the analog domain to allow for more efficient processing of images and out of band blockers. Additionally, dual down conversion offers the possibilities of producing an IF frequency in the range that the ADC can faithfully digitize.

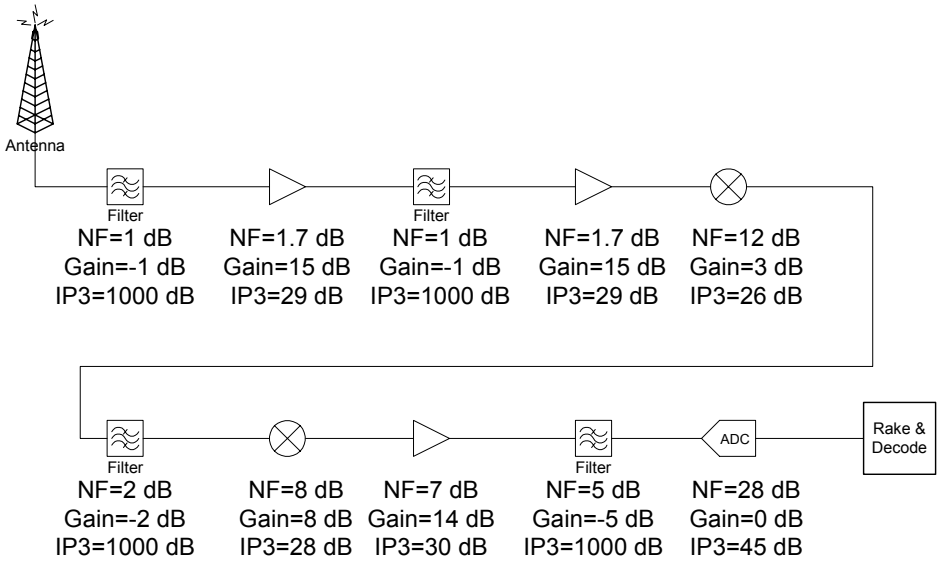

Figure 16: SDR receive signal chain

A traditional numerical analysis of this signal chain provides the following results.

Total NF	4.13 dB
Gain	30 dB
Input IP3	-8.7 dBm
Output IP3	+21.3 dBm

Given this signal chain, the SNR can now be determined for the reference sensitivity posted earlier. Using these updated terms in the equation above for the overall NSD, the total noise can now be determined.

$$NSD_{Ana\log\_front\_end} = 10\log\left(\frac{k*T*1Hz}{0.001}\right) + G + NF_{Ana\log\_front\_end} =$$
$$-174dBm/Hz + 30 + 4.13 = -139.87dBm/Hz$$

If this energy is integrated over the chip rate for each of the standards, the total noise in the channel for each standard is shown below.

	CDMA2000	UMTS
Noise at antenna	-113.1 dBm	-108.2 dBm
Noise after ADC	-83.1 dBm	-78.2 dBm
Signal energy after ADC	-87 dBm	-91 dBm
SNR after ADC	-3.9 dB	-12.8 dB
Required SNR (approx.)	-16 dB	-19

In both cases, adequate SNR is maintained, resulting in adequate sensitivity. Additionally, the excess SNR can be used to increase the sensitivity of the receiver beyond that which the specification calls for.

In addition to sensitivity, the spurious performance of the signal chain must be analyzed. The analysis of spurious performance is a little bit more difficult but can nonetheless be analyzed. In analyzing the CDMA2000 specification, there are two specifications to review. The key specifications are the two tone blocking and single tone blocking tests.

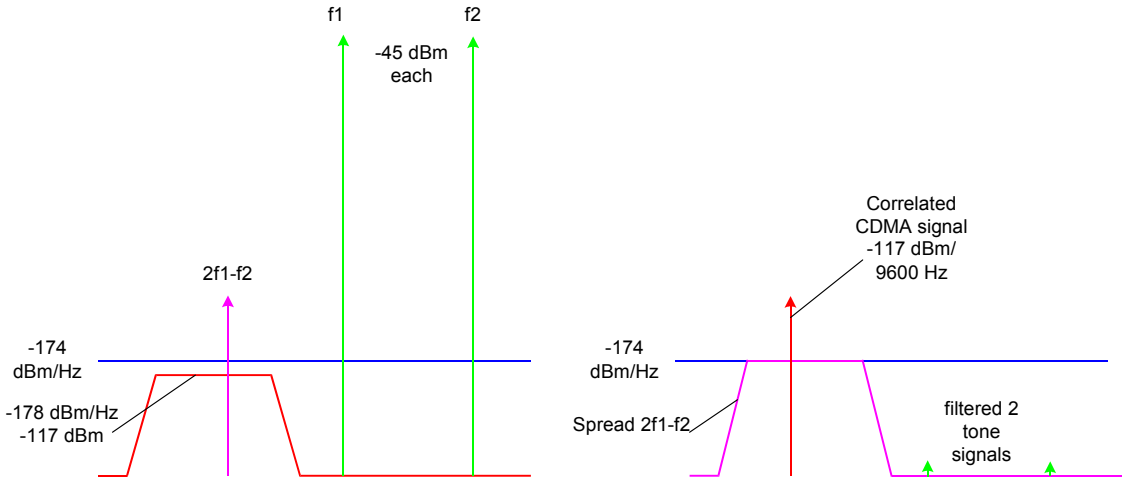

Figure 17: CDMA2000 two-tone blocking requirements

Two tone blocking requires that the receiver tolerate two CW carriers at –45 dBm. Since this is an IF sampling application, the even order term (difference) fall near DC and is filtered. The odd order products are the most critical, especially the third order products that fall in band. In the diagram above, a third order term (2f1-f2) is shown to fall near the channel center of the CDMA channel. Since it falls near the channel center, it cannot be allowed to disrupt the desired CDMA carrier. The goal is to determine how large the intermodulation product can be such that disruption does not occur. Fortunately, because the desired carrier is a CDMA signal, it will pass through a despreading circuit, which will correlate the desired CDMA carrier and de-correlate the undesired spurious term. After de-correlating, the CW signal will resemble white noise as shown in the right half of the drawing while the desired CDMA signal will be rendered as a narrowband and easily filtered and processed. Since the spurious signal becomes pseudorandom noise, it adds to the effective thermal noise at a density of –174 dBm/Hz (kT noise). Furthermore, the mobile power is allowed to increase by 3 dB during this test indicating that the noise generated by the spurious is allowed to equal that of the thermal noise. If it is assumed that the spurious products are generated in the ADC, then the noise figure may be added to the thermal noise before determining how large the spurious signal can be. Reflecting all the spurious to the antenna, the effective thermal noise including the NF of the entire signal chain produces an NSD of –169.87 dBm/Hz. Integrating this over 1.25 MHz will provide the total energy that may be contained in the spurious without adversely effecting performance of the receiver. The total power in 1.25 MHz is –108.9 dBm/Hz. This is the spurious level reflected to the antenna that will not cause blocking in the receiver. Since the receiver is blind as to how the spurious is generated, this number is valid for single or two tone blocking. Therefore, since the two CW tones were at –45 dBm, the input referred IP3 is found to be –13 dBm or –63.9 dBc. Likewise the single tone performance can be calculated with reference to –30 dBm giving –78.9 dBc. Since the ADC performance is listed at –95 dBc for either single or two tones, no performance limitations should be anticipated.

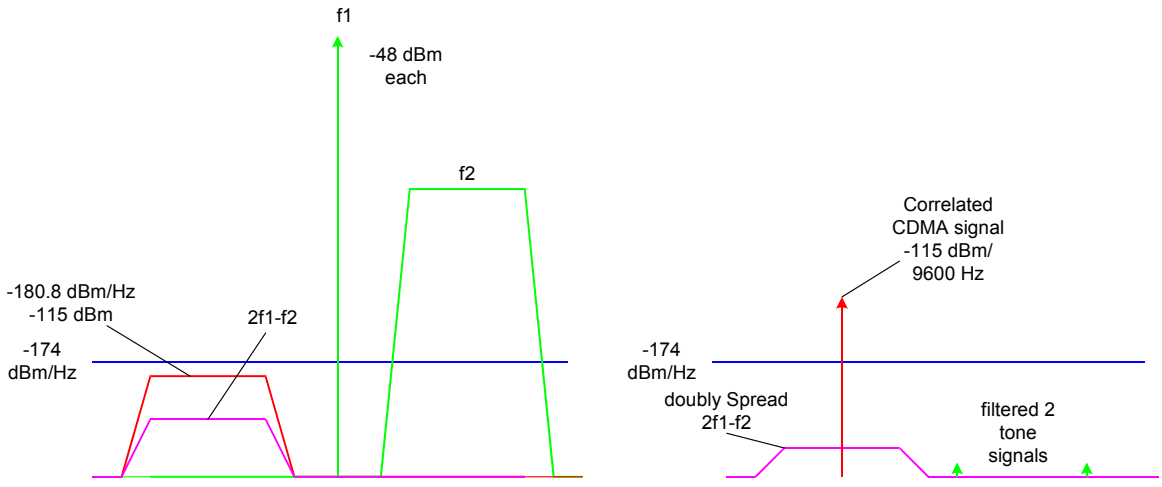

Figure 18: UMTS Intermodulation Performance Requirements

UMTS is a little different. First there is no single tone desensitization specification. Thus the primary specification is the intermodulation testing. This too is different because only one of the tones is CW; the other is a modulated CDMA carrier. Therefore, when 2f1-f2 is generated, it is an image of the CDMA signal shifted by the difference between itself and the CW tone. If this image falls directly on top of a desired CDMA carrier, it is possible that if the chipping sequence is not orthogonal to that of the desired carrier, then the blocking signal could be received instead of the desired signal. More likely however is that the undesired signal will simply increase the effective thermal noise. As shown in the right half of the figure, as with the CDMA2000 example, the desired signal is correlated and the undesired signal is spread. If the data rate is high enough into the despreading device, the original wideband intermodulation product will becomes doubly spread as it is convolved with the orthogonal despreading code. However, most often the oversampling ratio into the despreading function is only 2 or 4 potentially causing much of the doubly spread energy to alias back into the band of interest causing insignificant decrease in the spectral density of the noise. As with CDMA2000 the mobile is allowed to increase its power by 3 dB indicating that the noise due to the intermodulation product can equal the thermal noise.

Integrating the noise of 3.84 MHz gives a total intermodulation noise of –108.1 dBm. Again allowing for the noise figure of the Rx chain of 4.13 dB allows this noise to increase to about –104 dBm. Then comparison to the CW tone, this gives an input referred IP3 of –20 dBm. Reflecting this to the ADC gives an IP3 of +10 dBm or intermodulation performance of –74 dBc.

Since this receiver may also be used for reception of a narrow band signal, a quick check of narrow band performance is a good idea. Total noise referenced to the antenna can be calculated in a 30 kHz band to be –125 dBm. If 5 dB of SNR is required, this is reasonably good performance. In terms of intermodulation rejection, to achieve unrestricted performance at –125 dBm, the intermodulation products from two narrow band terms must be below this level. If the products must be below –130 dBm and they are generated by a –45 dBm to, then an input referred IP3 of –2.5 dBm is required or in terms of single tone performance with a –30 dBm in band blocker, -100 dBFS performance. Clearly from the single tone requirements, narrowband performance will be limited by the harmonics of the blockers, more so than sensitivity.

Clearly, this design maintains good performance for both CDMA2000 and UMTS while retaining the ability to perform reasonably well at narrowband standards. While this review is not exhaustive, it does indicate a methodology for doing looking at multimode radio performance. Since this is a wideband receiver, it is possible that the simultaneous reception in all three modes is possible providing that the digital processing is available (i.e. multi-carrier). Likewise, this receiver is suitable for field configuration between these modes, even if not operated in a multimode manner, providing many deployment options.

Conclusions

As new and more complex communication standards are developed around the globe, the demand for new transceivers architectures will also grow. However, more and more often the available capital, both cash and human, limit the designs that can be tackled. Fortunately, software radio technology is available for a select and growing group of these architectures that allow a single platform to leverage into many diverse designs. As seen here, this has many distinct advantages and is not limited to interoperability, investment retention and great flexibility. As with any software project, quite often the potential is only limited by the imagination of the designer. The great part is that like any software project, if there is a design error, it is just as simple as backspace, type and enter and the problem is fixed.

Fortunately, the last decade has seen significant advances in semiconductor technology that has caused impressive gains [17] not only in performance but also in cost. SDR is one area that has greatly benefited from these varied technologies and will continue to do so as the meaning of SDR is developed just as the history of programming languages has done.

While SDR is not the solution to all communication problems, it will offer robust solutions to challenging design issues in the coming years. These issues include phased array technology, location services, interoperability and complex concepts yet to be defined. However, there are still some challenges preventing full acceptance of this technology. The two main issues are cost and power. Interestingly, these two have a first order positive relationship; solve one problem and the other will only get better. Without low power, user devices will not be able to take full advantage of SDR technology. Clearly, the power issue comes from the need for high performance components. High performance means ultra linear devices. High linearity devices means low efficiency through high standing currents. Therefore, if the issue of how to design high linearity devices with lower power can be solved, and it will, then costs too will also fall, opening the door for many other applications. So the key to continued SDR development and evolution is continued device improvement down the Moore's law curve and continued interest in flexible radio architectures. Despite these challenges, the current state of performance is more than sufficient for engineers and manufacturers to seriously begin to investigate the possibilities of SDR as covered in this text.

References

1. "Software Radio, A Modern Approach to Radio Engineering", Jeffrey H. Reed, Prentice Hall PTR, 2002
2. Software Radio – Cognitive Radio, Dr. Joseph Mitola, III, http://ourworld.compuserve.com/homepages/jmitola/.
3. "A Look At Software Radios: Are They Fact Or Fiction", Brad Brannon, Dimitrios Efstathiou, and Tom Gratzek, Electronic Design, December 1, 1998, pg 117-122.
4. "Software Radio Concepts", Bob Clarke and Kevin Kreitzer, unpublished.

5. “Digital-radio-receiver design requires re-evaluation of parameters”, Brad Brannon, EDN, November 5, 1998, pg 163-170.
6. “New A/D Converter Benefits Digital IFs”, Brad Brannon, RF Design, May 1995, pg 50-65.
7. “Introduction to Radio Frequency Design”, W. H. Hayward, The American Radio Relay League, 1994-1996.
8. “Secrets of RF Circuit Design”, Joseph J. Carr, McGraw-Hill, 2001.
9. “Fast and Hot – data converters for tomorrow’s software-defined radios”, Brad Brannon, RF Design, July 2002, pg. 60-66.
10. “Redefining the Role of ADCs in Wireless”, Brad Brannon and Chris Cloninger, Applied Microwave and Wireless, March 2001, pg. 94-105.
11. “DNL and Some of its Effects on Converter Performance”, Brad Brannon, Wireless Design and Development, June 2001, pg. 10.
12. Analog Devices Applications Note AN-410, “Overcoming Converter Nonlinearies with Dither”, Brad Brannon, www.analog.com.
13. “High Speed Sampling and High Speed ADCs”, Walt Kester, Section 4, High Speed Design Techniques, www.analog.com.
14. “High Speed DACs and DDS Systems”, Walt Kester, Section 6, High Speed Design Techniques, www.analog.com.
15. http://www.qualcomm.com.
16. http://www.3gpp2.org
17. “Analog-to-Digital Converter Survey and Analysis”, Robert H. Walden, IEEE Communications Magazine, Vol. 17, No. 4, April 1999, pg. 539-550.