Using Software Defined Radio for Faster Speeds and Increased Bandwidth

Read About

Demanding SDR applications

Analog receiver architecture

Sub-sampling options

The future of SDR

Introduction

Software Defined Radio (SDR) is a construct for performing radio functions by means of software on a computer or embedded system. SDR isn't a new concept, and many of the techniques are well established. However, there have been challenges when it comes to the practicalities of the more demanding SDR applications driven largely by ADC/DAC limitations trade-offs. With ever faster ADC and DACs at better resolution, many of these compromises are being overcome – or at least being limited to higher parts of the frequency spectrum where the fast, high resolution digitizers with current generation wideband front end cannot reach.

Wideband is an important element for SDR to avoid limiting it to single frequency ranges and to be able to deal with several channels (or channel) at once. This is especially true for Cognitive radio. A cognitive radio introduces intelligence, allowing it to adapt as needed and even learn when appropriate. Having wideband performance is needed to allow room to dynamically manage the spectrum and radio parameters. SDRs offer many elegant advantages, outlined below, but true wideband spectral coverage at good resolution requires current and emerging technology.

Figure 1: VPX3-530 digitizer/FPGA used for SDR

Fortunately, high end SDR is becoming a more accessible solution without needing to resort to scanning tuners or stacked digitizers with fixed parallel tuners. Today's ADCs are able to fully digitize parts of the multi-GHz including L-band and some S-band frequencies at native 12-bit resolution, and into X-band at 8-bit resolution. At these speeds and resolutions, tuner-less radios are becoming a more practical solution.

Info

curtisswrightds.com

Email

ds@curtisswright.com

Built on generic building blocks, the philosophy behind SDR is that common hardware, typically a computer, can change the overall radio function by making changes to the software. An ideal SDR would be able to digitize the frequency spectrum of interest directly from an antenna, present the data to a DSP processor and output to an application – and also the reverse for a transmitter. The main benefits of SDR are:

- **Flexibility**: Little or no hardware changes required for a different function
- **Interoperability**: Can be used with a multitude of old and new systems (with the correct software)
- **Ease of upgrade**: New features can be added through software, such as data encryption
- **Efficiency**: No need to support multiple radios
- **Higher level interfaces**: Access to GUI and network interfaces can be provided
- **Next generation**: The launch pad for next generation radios, such as cognitive radio.

Basic Analog Receiver Architecture

A basic analog receiver employs a mixer or down converter to perform frequency translation from RF to IF and then to audio frequencies in stages by performing an analog multiplication of a local oscillator frequency with the incoming RF signal using a mixer. The output of a linear mixer would produce the sum and difference of the frequency components. By using filters, the unwanted images can be removed. Without filters, then multiple channels on different carriers would be selected resulting in noise and distortion. Further filters are also required to remove aliased images too. Changing the local oscillator frequency changes the channel selection of course. Once filtered, all that remains to do is to detect the signal, demodulate, for example FM demodulation, and to provide suitable output amplification.

Although Figure 2 outlines the ideal analog receiver, from a practical point of view, there can be multiple tuner stages, first collecting the higher frequency RF before a second stage with usually a fixed tuning frequency which provides an IF of a more manageable bandwidth. This is shown in Figure 3.

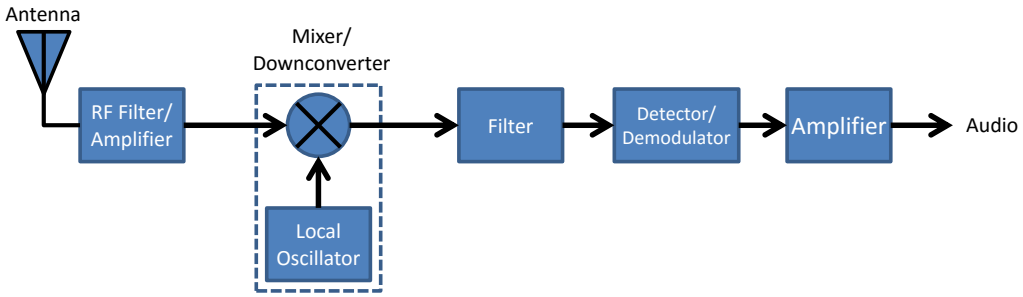

Figure 2: Basic analog receiver block diagram

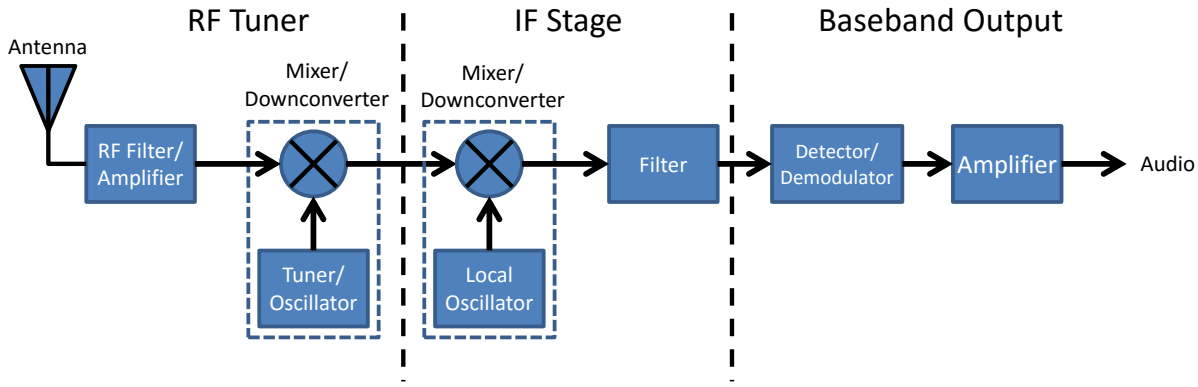

Figure 3: Practical analog receiver block diagram

Digital Receivers Including Subsampling Options

For a digital solution, an ADC is used to sample the input signal at a speed of at least twice the bandwidth of the frequency range which is the basic criteria of the defined Nyquist-Shannon. In practice, the recommended sampling frequency could be much higher because of imperfect anti-aliasing filter characteristics and other design parameters. A common misunderstanding is the ADC sample rate must always be at least twice the highest frequency of interest, but this is only true for baseband applications working in the first Nyquist zone only and in an ideal world. Provided the bandwidth of interest fits completely within a single Nyquist zone, and the ADC has sufficient analog bandwidth, then sampling frequencies below input frequency are possible with appropriate filtering to remove components from other Nyquist zones which could alias the signal. This is known as sub-sampling and useful because slower lower cost ADCs can be used and corresponding less data hits the DSP.

Each Nyquist zone is defined as being half the sample frequency (Fs) with Nyquist zone 1 frequency range is 0 to 0.5*Fs, zone 2 is 0.5*Fs to 1*Fs, zone 3 is 1*Fs to 1.5*Fs and so on (see Figure 4). If the digitizer is sampling input signals higher than 2*Fs, then you cannot unambiguously determine which Nyquist zone the signal is from. However, if all other Nyquist zone signals are filtered out before digitization, then sub-sampling is an option – provided the Fs>2*signal bandwidth. Sub-sampling can be used to translate the signal down to Nyquist zone 1 for baseband processing. If an even number Nyquist zone is used, then frequency reversal will take place, but this is easily handled in the software domain when using FFTs (see Figure 5). If the signal bandwidth does not fully fit in a Nyquist zone, then frequency folding or aliasing will occur.

For sub-sampling applications to work:

- The ADC front end must have sufficient bandwidth to “see” the signal in the first place
- The sample rate must be >2*signal bandwidth (ideally more)
- The signal of interest should fit entirely within a single Nyquist zone (and in practice have guard bands away from the Nyquist zone edges because filters are not perfect and have roll-offs)
- Bandpass filters are used to remove signal outside the all other Nyquist zones

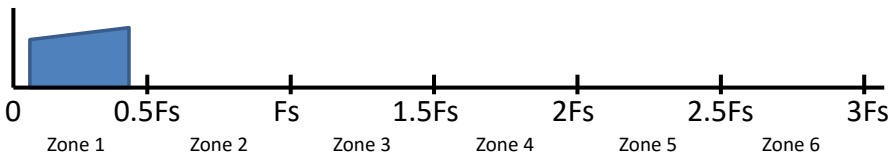

Figure 4: Baseband signal complying to the Nyquist sampling criteria

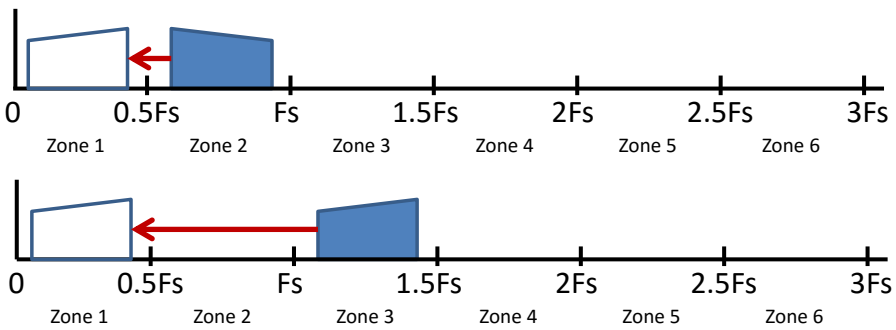

Figure 5: Sub-sampling in odd or even Nyquist zones

Depending on the speed and resolution of ADCs (and DACs in the reverse direction), either the baseband or better still the IF stage can be digitized. Digitizing means all the lower level tuning and processing can be done in software – this is an SDR. See Figure 6.

Some applications need to process several channels or bigger channel carrying more than audio data rates. For these wideband applications, very high speed converters have become main stream with typical parts operating in the Gsps range and therefore needing the processing and interfaces to deal with the multi-GBps of data. A typical converter would be a 12b 4 Gsps device and able to fully digitize the L-band spectrum used for many telecommunications applications. In reality, even though the application may look to digitize a large piece of the spectrum, not all of it is useful. As a result, digital filters and DDCs are used to extract the meaningful data to keep the volume of data being moved and processed to a minimum.

ADCs and DACs are continually getting faster and faster with better and better resolution. This allows bigger chunks of the frequency spectrum to be seen and processed. From a hardware point of view, this is useful because there is less need for analog tuner stages, or even no tuners at all.

Having no tuner and ADC sampling at RF is often referred to as “direct RF”. Tuner-less direct RF is the Holy Grail for SDR because it saves cost, reduces the size of the system and simplifies the overall architecture (see Figure 7). For example, a by-product of using a tuner is that the complex phase outputs (I&Q) double up on the number of ADC converters needed to digitize the data, which also need to be phase and gain matched to avoid compromising the system performance. Replacing the RF tuner to provide a direct RF can also be absorbed into the SDR philosophy as shown in Figure 5. Given the higher raw data rate, a corresponding power DSP, other processors or FPGAs would be recommended for digitizing as RF would equate to multi-GBps of data. However, since the tuner stage has been bypassed, the footprint would be much smaller and a wider part of the spectrum would be directly available without any analog re-tuning.

The combined tasks of interfacing to high speed ADCs (and DACs) and DDC implementation is ideally suited to FPGAs. If the FPGA is large enough, multiple parallel DDC or channelizers can be implemented and each data stream associated with its own modulation or encryption scheme.

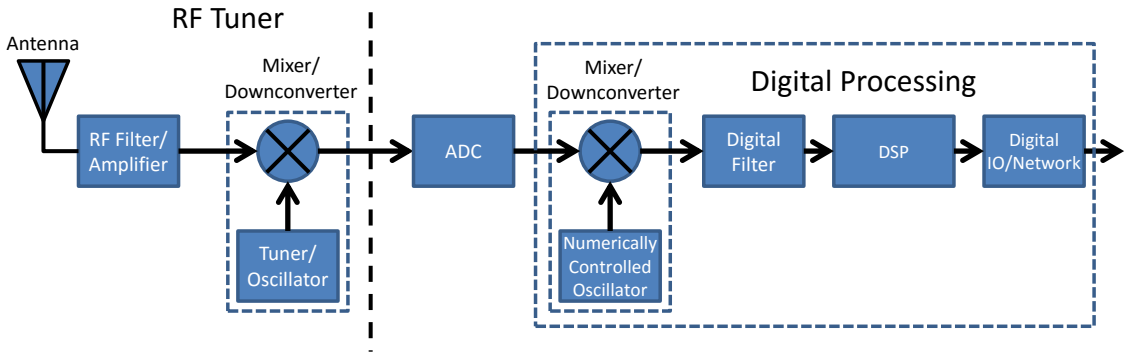

Figure 6: An SDR receiver

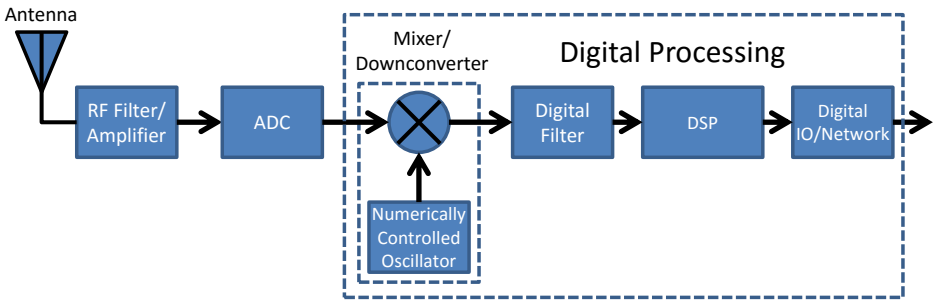

Figure 7: Direct RF SDR receiver

Some applications require very wide bandwidth operation. As bandwidth, and therefore sample rate, increases, the choice of digitizers reduces along with resolution though faster and faster parts are always being developed. For these applications, tuner front ends are used in a stacking arrangement with multiple tuners set up to deal with pieces of the spectrum with a corresponding digitizer for each tuner. The wider the sample rate and bandwidth of each digitizer, the fewer the number of digitizers required to achieve full spectral coverage and less likely to have a channel which overlaps the digitizer to digitizer overlap. The latter is because there are fewer spectral overlaps.

SDR schemes could be extended further, limited only by processor power or resource to handle multiple channels in parallel from a common data source. An example of a parallel SDR receiver is shown in Figure 8.

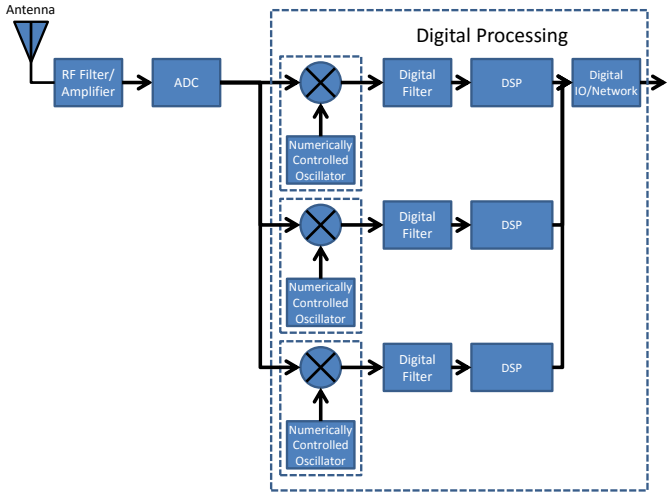

Figure 8: Parallel SDR processing

Example Wideband SDR Hardware

Examples of rugged SDR capable solutions with high-end wideband performance include products like the VPX3-530 which supports up to 4 Gsps analog inputs for either direct L-band frequency digitization. Since product like the VPX3-530 can support dual 4 Gsps inputs, it is possible to use high frequency tuner front ends by supporting I & Q inputs post analog down conversion from the tuner. VPX3-530 also supports two 2.8 Gsps update rate 14-bit DAC with interpolation up convertors modes to 5.6 Gsps.

A local Xilinx® Virtex®-7 FPGA can provide resource for DDCs, filtering and other processing to provide much of the hardware requirements for wideband SDR.

Figure 9: Example SDR hardware solution

SDR Software

Software fundamentally defines the operation of the radio. This demands standardization in its architecture, because the core of SDR is the ability to update or change parts of the software efficiently and allow interoperability. In the SDR domain, “waveform” is the term that defines a range of parameters such as the modulation, carrier frequency and coding scheme. In short, a waveform is a radio function. To add a new radio function, a new waveform is added. Consequently, as new waveforms become available, it is important they are easily incorporated, so standard frameworks are needed along with the ability to do these changes on-the-fly.

The Software Communications Architecture (SCA) is an open standard framework used for SDR largely used by the defense community and defined by the Joint Tactical Radio System (JTRS) program, but there are other architectures. Underpinning SCA is CORBA (Common Object Request Broker Architecture) middleware as a platform for waveforms applications which are platform independent, and therefore interoperable.

In addition to SCA, there are other frameworks. GNU radio has a large following and its code base is supported by a community of developers. For this reason, GNU radio is not favoured by the defense community, but does allow quick prototyping of systems as its free of charge. GNU radio is distributed under the terms of the GNU General Public Licence.

Another framework is provided by REDHAWK as a more complete environment and uses CORBA as an underlying middleware. It can also fit within the SCA framework.

Authors

Jeremy Banks, BSc (Hons) in Electronic and Electrical Engineering
Product Marketing Manager
Curtiss-Wright Defense Solutions

Summary

Technology is close to providing the ideal hardware required for direct RF, baseband processing as wideband, multi-Gsps technology with matched processing resource which is making higher performance ideal SDR solutions a reality. The foundation for this is the ever faster sampling ADC and DAC devices with wide bandwidth analog front ends with the backend processing to match. This means more of the ideal SDR architecture can be realized in the digital domain rather than being reliant on the somewhat fixed function of front end tuners, which at best can be considered software controlled, rather than software defined. The more that can be defined as "software defined" moves us closer to true wideband SDR.

What does the future hold next? Breaking the 10 Gsps barrier. Once ADCs were just that, an ADC, but now the higher speed devices have built in DDCs which are also becoming more sophisticated too. It could be argued such DDCs in ADCs are software controlled rather than defined, but this just gives the SDR additional options as it's such a common requirement for all radios. An ideal SDR moves the ADC and DAC closer and closer to the antenna; this is what we're seeing in the multi-GHz spectrum used in communications, modems and beyond. In a few years from now, we'll likely have all of this at even better resolution for even better performance.

Learn More

Case Study: 3U VPX SDR Solution

White Paper: Wideband DDC as Part of an Integrated Approach

Product: VPX3-530 3U VPX Virtex-7 FPGA ADC/DAC

Technology:

- Modular Open Radio Frequency Architecture (MORA)
- Size, Weight and Power (SWaP)