UNLEASHING THE POTENTIAL OF SOFTWARE-DEFINED RADIOS IN MODERN WIRELESS NETWORKS

With the exponential growth in the ways and means by which people need to communicate - data communications, voice communications, video communications, broadcast messaging, command and control communications, emergency response communications, etc. – modifying radio devices easily and cost-effectively has become business critical. Software defined radio (SDR) technology brings the flexibility, cost efficiency and power to drive communications forward, with wide-reaching benefits realized by service providers and product developers through to end users.

Software Defined Radio - Defined:

A number of definitions can be found to describe Software Defined Radio, also known as Software Radio or SDR. The SDR Forum, working in collaboration with the Institute of Electrical and Electronic Engineers (IEEE) P1900.1 group, has worked to establish a definition of SDR that provides consistency and a clear overview of the technology and its associated benefits.

Simply put Software Defined Radio is defined as1 :

"Radio in which some or all of the physical layer functions are software defined"

A radio is any kind of device that wirelessly transmits or receives signals in the radio frequency (RF) part of the electromagnetic spectrum to facilitate the transfer of information. In today's world, radios exist in a multitude of items such as cell phones, computers, car door openers, vehicles, and televisions. Traditional hardware based radio devices limit cross-functionality and can only be modified through physical intervention.

This results in higher production costs and minimal flexibility in supporting multiple waveform standards. By contrast, software defined radio technology provides an efficient and comparatively inexpensive solution to this problem, allowing multimode, multi-band and/or multi-functional wireless devices that can be enhanced using software upgrades.

Characteristics and Benefits of a Software Radio:

Implementation of the ideal software radio would require either the digitization at the antenna, allowing complete flexibility in the digital domain, or the design of a completely flexible radio frequency (RF) front-end for handling a wide range of carrier frequencies and modulation formats. The ideal software radio, however, is not yet fully exploited in commercial systems due to technology limitations and cost considerations.

A model of a practical software radio is shown in Figure 1.

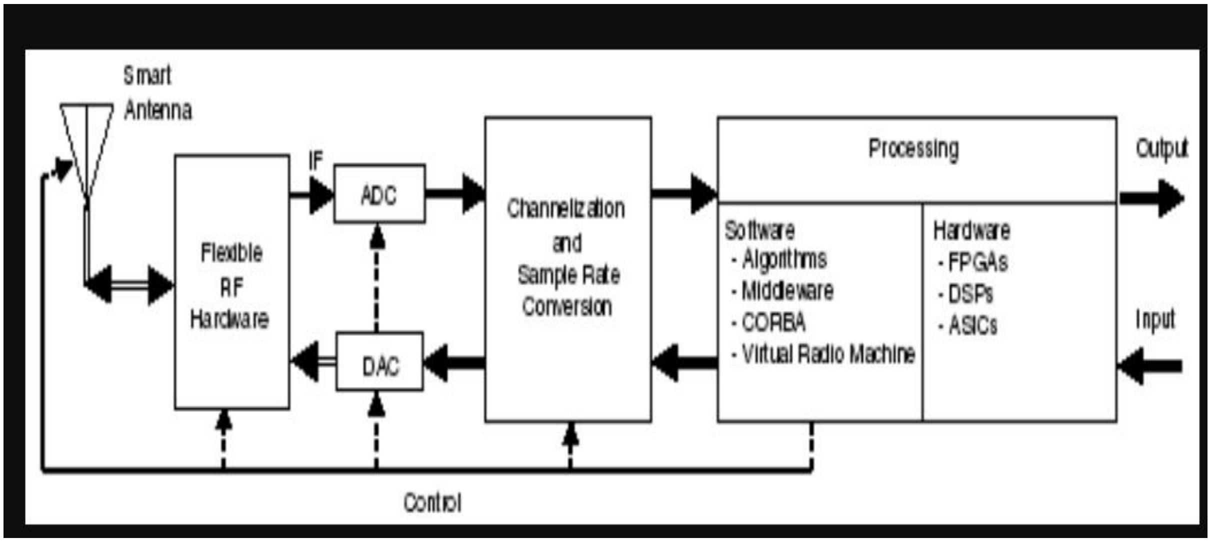

Figure 1. Model of a Software Radio

The receiver begins with a smart antenna that provides a gain versus direction characteristic to minimize interference, multipath, and noise. The smart antenna provides similar benefits for the transmitter. Most practical software radios digitize the signal as early as possible in the receiver chain while keeping the signal in the digital domain and converting to the analog domain as late as possible for the transmitter using a digital to analog converter (DAC). Often the received signal is digitized in the intermediate frequency (IF) band. Conventional radio architectures employ a super heterodyne receiver, in which the RF signal is picked up by the antenna along with other spurious/unwanted signals, filtered, amplified with a low noise amplifier (LNA), and mixed with a local oscillator (LO) to an IF. Depending on the application, the number of stages of this operation may vary.

Finally, the IF is then mixed exactly to baseband. Digitizing the signal with an analog to digital converter (ADC) in the IF range eliminates the last stage in the conventional model in which problems like carrier offset and imaging are encountered. When sampled, digital IF signals give spectral replicas that can be placed accurately near the baseband frequency, allowing frequency translation and digitization to be carried out simultaneously. Digital filtering (channelization) and sample rate conversion are often needed to interface the output of the ADC to the processing hardware to implement the receiver. Likewise, digital filtering and sample rate conversion are often necessary to interface the digital hardware that creates the modulated waveforms to the digital to analog converter. Processing is performed in software using DSPs, field programmable gate arrays (FPGAs), or application specific integrated circuits (ASICs). The algorithm used to modulate and demodulate the signal may use a variety of software methodologies, such as middleware, e.g., common object request broker architecture (CORBA), or virtual radio machines, which are similar in function to JAVA virtual machines. This forms a typical model of a software radio.

1. **Multi-functionality—**With the development of short-range networks like Bluetooth and IEEE 802.11, it is now possible to enhance the services of a radio by leveraging other devices that provide complementary services. For instance, a Bluetooth-enabled fax machine may be able to send a fax to a nearby laptop computer equipped with a software radio that supports

the Bluetooth interface. Software radio's recon-figuration capability can support an almost infinite variety of service capabilities in a system.

2. **Global mobility—**A number of communication standards exist today. In the 2G alone, there are IS-136, GSM, IS-95/CDMA1, and many other, less well known standards. The 3G technology tried to harmonize all the standards. However, there are many standards under the 3G umbrella. The need for transparency, i.e., the ability of radios to operate with some, preferably all, of these standards in different geographical regions of the world has fostered the growth of the software radio concept. Military services also face a similar issue with incompatible radio standards existing between as well as within branches of the military.
3. **Compactness and power efficiency—**Multifunction, multimode radios designed using the "Velcro" approach of including separate silicon for each system can become bulky and inefficient as the number of systems increases. The software radio approach, however, results in a compact and, in some cases, a power-efficient design, especially as the number of systems increases, since the same piece of hardware is reused to implement multiple systems and interfaces.
4. **Ease of manufacture—RF** components are notoriously hard to standardize and may have varying performance characteristics. Optimization of the components in terms of performance may take a few years and thereby delay product introduction. In general, digitization of the signal early in the receiver chain can result in a design that

incorporates significantly fewer parts, meaning a reduced inventory for the manufacturer.

5. **Ease of upgrades—**In the course of deployment, current services may need to be updated or new services may have to be introduced. Such enhancements have to be made without disrupting the operation of the current infrastructure. A flexible architecture allows for improvements and additional functionality without the expense of recalling all the units or replacing the user terminals. Vocoder technology, for example, is constantly improving to offer higher quality voice at lower bit rates. As new vocoders are developed, they can be quickly fielded in software radio systems. Furthermore, as new devices are integrated into existing infrastructures, software radio allows the new devices to interface seamlessly, from the air-interface all the way to the application, with the legacy network.

RF Implementation issues:

A. Receiver Architectures

The primary distinction between receivers is the number of stages taken to downconvert a signal to baseband. Direct conversion takes one down-conversion, superheterodyne receivers employ two or more. Complexity increases with the number of down-conversions. The simplicity of direct conversion brings with it several technical problems which would appear to make direct conversion architecture inappropriate for an SDR receiver.

i) Direct Conversion Architecture

A basic direct conversion receiver architecture is shown in Figure 2. The receiver consists of a low noise amplifier (LNA) which provides modest RF gain at a low noise figure. The output signal from the LNA is filtered in a preselect filter, and down converted in a complex (I,Q) mixer. The majority of the gain and automatic gain control (AGC) is provided in a high gain baseband amplifier.

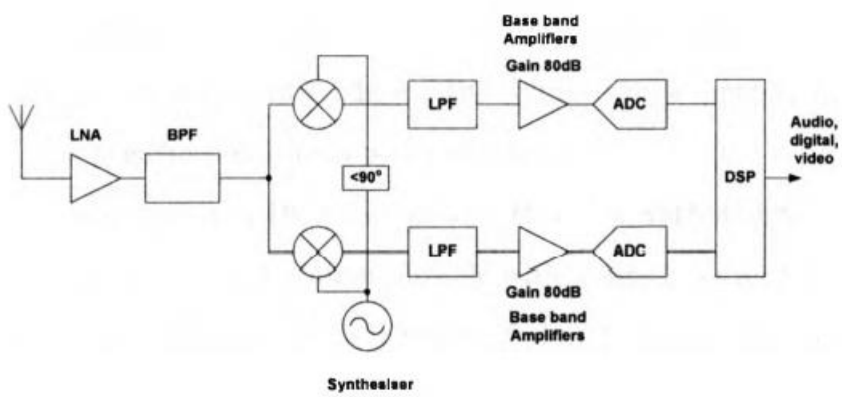

Figure 2. Direct conversion Receiver Architecture

The advantages of having a direct conversion architecture is that it has low complex architecture and does not require an IF stage. Image signals in a direct conversion architecture turns out to be a frequency-inverted image of the wanted signal itself and, although some image rejection advantages can be achieved, the image cannot be ignored. If the I and Q components of the local oscillator are in precise phase quadrature and precise amplitude balance, then the image signal will be eliminated. But if this is not achieved, then a small sideband will be superimposed on the wanted sideband, resulting in magnitude and phase errors.

A significant problem with direct conversion architecture is the introduction of a DC offset. This DC offset can arise from a number of sources. It can be seen that the leakage of the local oscillator signal to the input of the LNA, occurs through an imbalance in the mixer. Once these signals are in the system, they are then mixed with themselves to produce a DC voltage at the input to the baseband amplifier. This phenomenon is known as self-mixing.

ii) Multiple Conversion Architecture

A multiple conversion receiver is shown in Figure 3. The receiver consists of a LNA, a variable local oscillator, two band pass filters and an IF amplifier that amplifies the intermediate frequency signal.

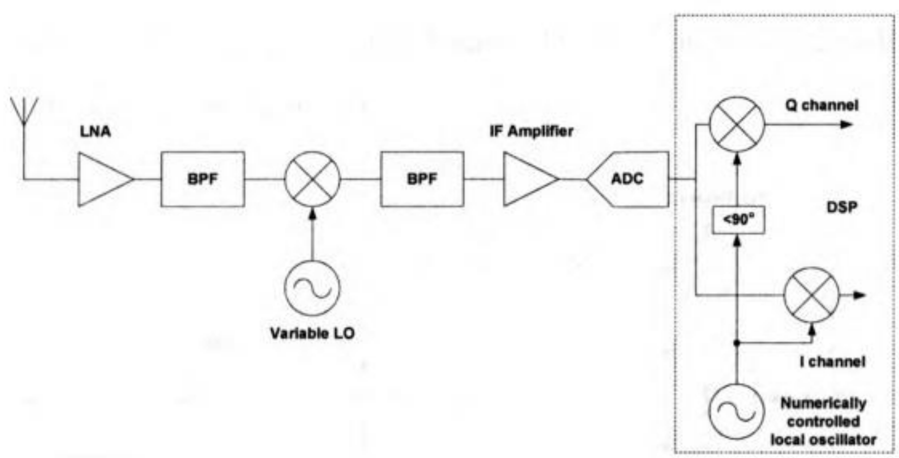

Figure 3. Multiple conversion Receiver Architecture

Although the multiple conversion stage only shows two explicit down-conversions (one in the RF hardware and one in digital signal processing(DSP), further conversions can be done via the processes of decimation and/or subsampling. In this architecture, the first conversion may be done in RF hardware, and all of the others are done in FPGA/DSP.

B. Transmitter Architectures

Basically the same choice applies to transmitter architectures as applies to receiver architectures. The advantages and disadvantages associated with receiver architectures more or less translate to transmitters.

i) Direct Conversion Transmitter

The block diagram of a direct conversion transceiver is shown in Figure 4. The transmitter consists of Digital-to-Analog converters and an IQ modulator that converts the baseband frequency directly to RF. The output signal is then filtered in a band pass filter and then amplified using a High Power Amplifier (HPA).

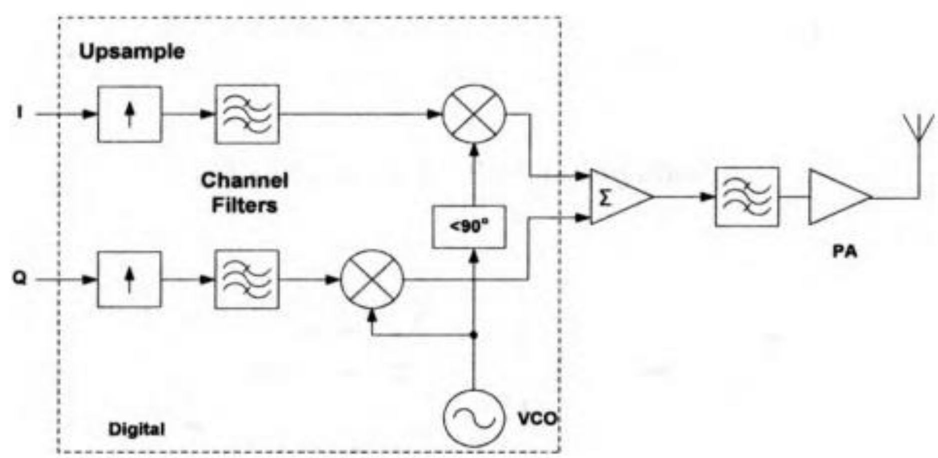

Figure 4. Direct up-conversion transmitter

Low complexity is one of the main advantage of a direct conversion transmitter. It also has simple filter requirements and it's suitable for integrated circuit realization. Due to lower number of parts used in this architecture, it helps to reduce the current consumption. Moreover, unwanted signals or image signals and sideband problems are dealt with easily.

The direct-conversion architecture nonetheless suffers from an important drawback: disturbance of the local oscillator by the power amplifier output. This issue arises because the PA output is a modulated waveform having a high power and a spectrum centered around the LO frequency. Despite various shielding techniques employed to isolate the VCO, the noisy output of the PA still corrupts the oscillator spectrum. This corruption occurs through injection pulling or injection locking, whereby the frequency of an oscillator tends to shift towards the frequency of an external stimulus. The local oscillator leakage through the mixer will be radiated through the antenna. The power amplifier linearization and the final mixers have to operate over a wide frequency band.

ii) Multiple Conversion Transmitter

Another approach to circumventing the problem of LO pulling in transmitters is to upconvert the baseband signal in multiple steps so that the PA output spectrum is far from the frequency of the VCOs. A multiple conversion architecture is shown in Figure 5.

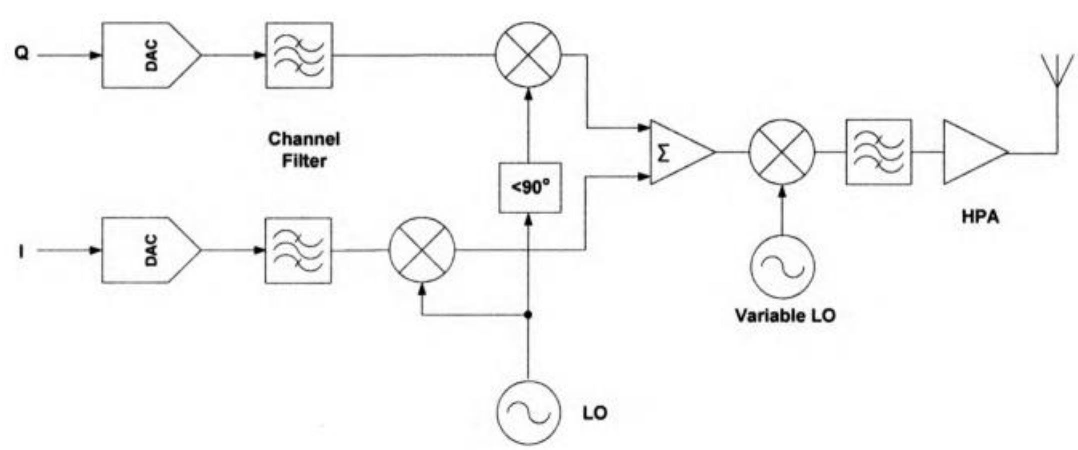

Figure 5. Multiple conversion transmitter

Here the baseband I and Q channels undergo quadrature modulation at a lower intermediate frequency, and the resulting output is upconverted by mixing and bandpass filtering.

Stand-alone SDR Transceiver:

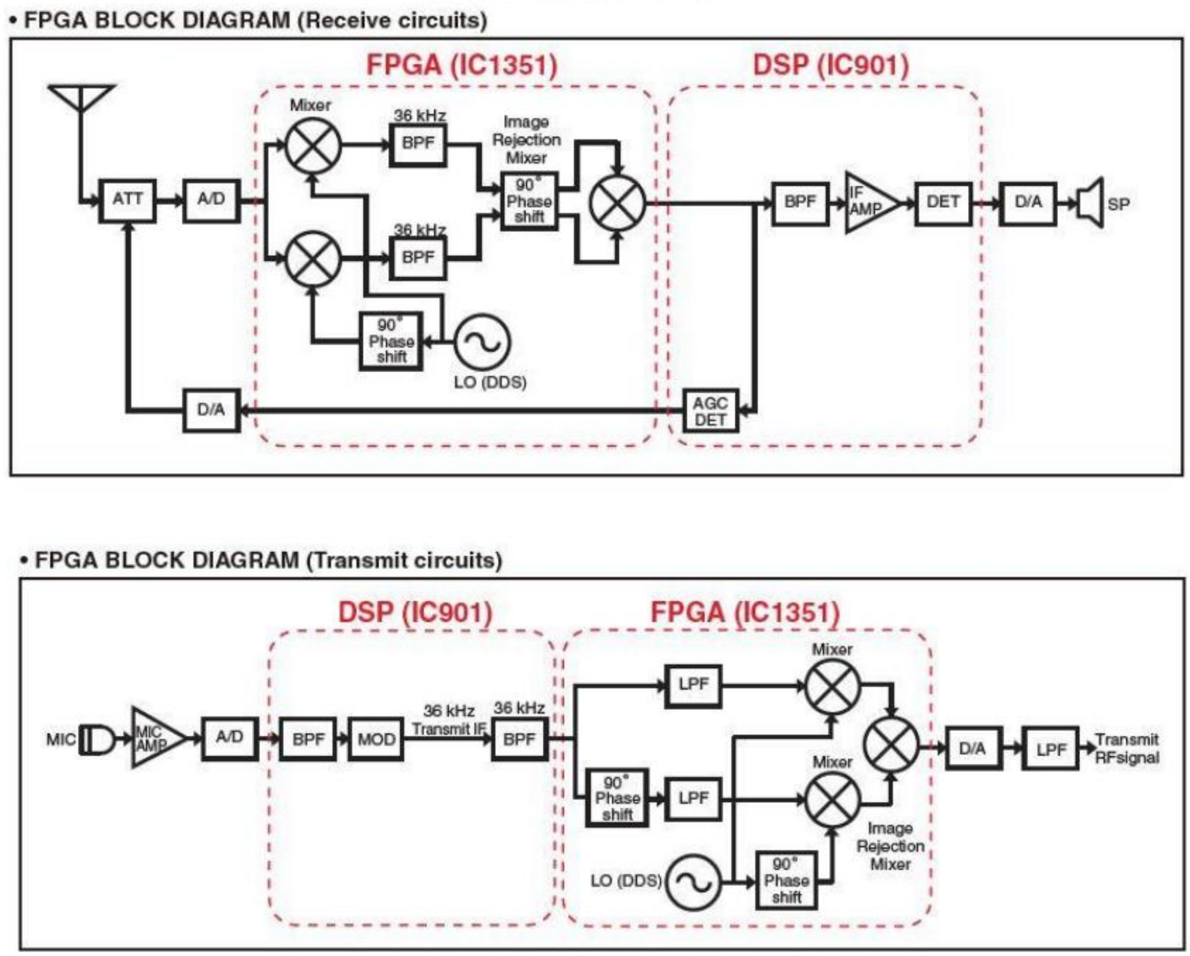

Figure 6. SDR Transceiver Block Diagram

A basic SDR system may consist of a personal computer equipped with a sound card, or other analog-to-digital converter, preceded by some form of RF front end. Significant amounts of signal processing are handed over to the general-purpose processor, rather than being done in special-purpose hardware (electronic circuits). Such a design produces a radio which can receive and transmit widely different radio protocols (sometimes referred to as waveforms) based solely on the software used.

Operating Principles:

Superheterodyne receivers use a VFO (variable-frequency oscillator), mixer, and filter to tune the desired signal to a common IF (intermediate frequency) or baseband. Typically in SDR, this signal is then sampled by the analog-to-digital converter. However, in some applications it is not necessary to tune the signal to an intermediate frequency and the radio frequency signal is directly sampled by the analog-to-digital converter (after amplification).

Noise and Distortion in the RF Chain, ADC and DAC Distortion:

Real analog-to-digital converters lack the dynamic range to pick up sub-microvolt, nanowatt-power radio signals produced by an antenna. Therefore, a low-noise amplifier must precede the conversion step and this device introduces its own problems. For example, if spurious signals are present (which is typical), these compete with the desired signals within the amplifier's dynamic range.

They may introduce distortion in the desired signals, or may block them completely. The standard solution is to put band-pass filters between the antenna and the amplifier, but these reduce the radio's flexibility. Real software radios often have two or three analog channel filters with different bandwidths that are switched in and out. The flexibility of SDR allows for dynamic spectrum usage, alleviating the need to statically assign the scarce spectral resources to a single fixed service.

Importance of LNA in an SDR

A low noise amplifier is an integral part of an RF system, and it is important to choose an LNA with the right specifications LNAs serve a vital role in software defined radio.

As RF signals travel from the transmitter to the receiver, the channel introduces noise into the signals. RF signals also experience attenuation, distortion, reflection, and other forms of interference. All these degrade the signal quality at the receiver, hence the reason for amplification. The RF engineer must optimize the sensitivity of the LNA to improve the sensitivity and performance of an RF circuit (including the SDR).

The noise figure of an amplifier specifies the noise performance of RF components, including an LNA. The lower the noise figure, the better the amplifier's performance. One of the major goals in RF circuit design is to minimize the noise of other components before this device to ensure that minimal noise is introduced into the circuit.

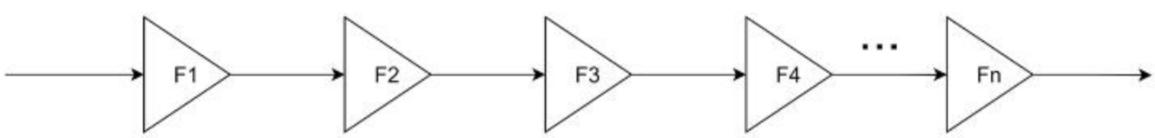

$$F_{total} = F_1 + \frac{F_2 - 1}{G_1} + \frac{F_3 - 1}{G_1 G_2} + \frac{F_4 - 1}{G_1 G_2 G_3} + \cdots,$$

ADC and DAC Distortion:

ADC and DAC components often make or break an SDR system. In general, as speed goes up, resolution goes down. Higher conversion rates allow more bandwidth to be digitized, but higher-resolution converters provide more dynamic range (the difference in amplitude between the strongest and weakest signals that can be simultaneously digitized). A communications system designer may want to carry as much bandwidth as possible, but high usage of the spectrum of interest is likely to include strong interfering signals as well as weaker signals of interest. The dynamic range requirement imposed by the different signal amplitudes likely defines the number of bits required. For example, suppose a small signal (carrier) needs to be recovered in the presence of a much larger signal (interference) and all signals of interest lie in a 40-MHz bandwidth. The dynamic range required to do this is the difference in signal amplitudes (60 dB) plus an adequate operating margin to demodulate the signal, perhaps 10 dB. Sampling theory dictates that the sampling frequency must be at least twice the signal bandwidth. These requirements can be satisfied by a good quality, 14-bit ADC which, when sampled at 100 MHz, provides a dynamic range of at least 80 dB.