Meeting the Design Challenges of 4G

Implementation of an OFDM Transceiver using an SDR Platform

By J. Cyrus Sy, Spectrum Signal Processing
Even though the roll-out of third generation (3G) wireless networks and services is just beginning, equipment vendors, service providers, OEM manufacturers and communications research bodies are working to define the next generation of wireless networks. The move towards what is being coined fourth generation (4G) wireless is complicated by the fact that a single 3G standard upon which to build does not exist. However, most industry experts agree that the future of wireless is one in which voice, video, multimedia and broadband data services traveling across multiple wireless air interfaces are meshed into one seamless network. 4G Wireless networks will be characterized by the following:
1. Seamless network of multiple air interfaces and protocols
2. Improved spectral efficiency
3. IP Based
4. Higher data bandwidths (&lt100Mbps)

Several technologies based on multi-carrier modulation (MCM) have come to the forefront in order to achieve the above characteristics. We shall examine one such technology known as Orthogonal Frequency Division Multiplexing (OFDM).
But the fact that OFDM is just one of many competing technologies highlights the need for reconfigurable and flexible software defined radio systems as a development platform. Here we will present an architecture for a software development platform by first looking at the challenges of designing an OFDM-based system. The same architecture, however, could also be used to implement other interfaces such as wideband code division multi-access (W-CDMA), and multi-carrier code division multiple access (MC-CDMA).

Orthogonal Frequency Division Multiplexing

OFDM is a communications technique that divides a communications channel into a number of equally spaced frequency tones (bands). OFDM is a form of multi-carrier modulation (MCM) where a sub-carrier within each frequency band is modulated to carry a portion of the user information. A communications data stream is effectively split into N parallel low bandwidth modulated data streams (Figure 1). Each sub-carrier overlaps, but they are all orthogonal to each other, such that they do not interfere with one another.

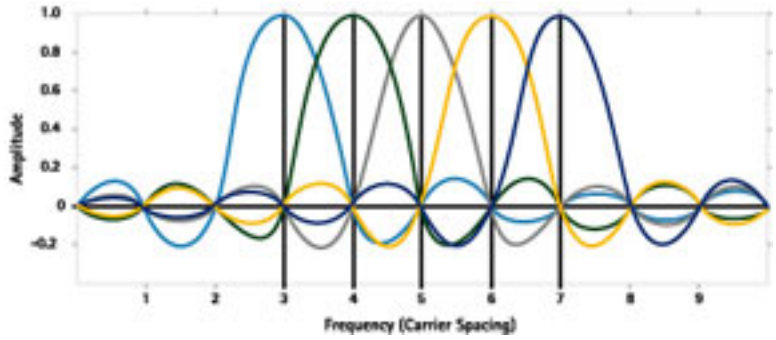

Figure 1. OFDM Modulation Scheme

Each of the sub-carriers has a low symbol rate. But the combination of sub-carriers carrying information in parallel allows for high data rates. The other advantage of a low symbol rate is that inter-symbol interference (ISI) can be reduced dramatically since the symbol time represents a very small proportion of the typical multi-path delay.

The transmitter stage of an OFDM transceiver takes data from an IP network, converts, and encodes it into a serial stream before modulation (Figure 2). The OFDM signal is generated using an Inverse Fast Fourier Transform (IFFT) into an IF analog signal which is then sent to the RF transceiver. The receiver stage of the transceiver simply reverses the process.

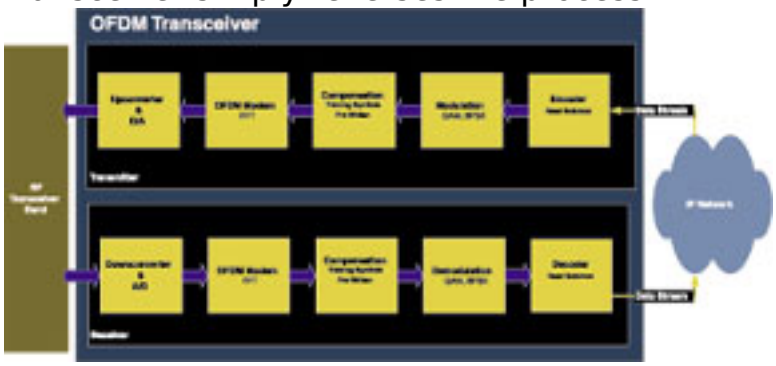

Figure 2. OFDM Transceiver

OFDM provides a particularly robust air interface that is resistant to the effects of multi-path delays while maintaining spectral efficiency. However, designers of OFDM and other MCM-based systems have been forced to deal with a host of challenges. Fortunately, there are a wide range of technologies and techniques available today that address each one of these challenges.

Fast-Fourier Transform

The typical OFDM transmitter converts an information stream to serialized PSK or QAM symbols. The serial stream is then converted to a parallel stream that is modulated using an inverse FFT operation. The outputs of the IFFT are streams that are then serialized and modulated by a single carrier. The receiver side implements a reverse process where the receive data is split back into parallel streams that are processed by an FFT. The output of the FFT is then serialized into a single stream for decoding.

Such a scheme requires an FFT/IFFT processing engine capable of doing continuous calculations in real-time. In the past, only specially designed ASICs optimized for FFT performance gave the real-time performance required. Today, high speed Field Programmable Gate Arrays (FPGA) from companies such as Xilinx have features such as built-in hardware multipliers and expanded memory blocks that provide the

required performance. General-purpose processors and digital signal processors (DSP) such as Motorola's G4 Altivec processor and Texas Instrument's C64xx processor line take advantage of SIMD and VLIW architectures combined with high clock speeds, give the necessary horsepower for real-time FFT performance.

Frequency Offset

Frequency offset occurs when the voltage-controller oscillator (VCO) at the receiver is not oscillating at the same carrier frequency as the transmitter. Any difference in the frequency at the receiver side will lead to increased error rate. While frequency offset is inherent for all modulation schemes, OFDM is particularly sensitive because energy from adjacent sub-carriers that have been offset in frequency will destroy the orthogonality of the sub-carriers. Obviously the result is severe inter-carrier interference (ICI) thereby increasing the error rate. One method of dealing with such a problem in a packet-based network is by adding a training sequence at the beginning of the packet to help the receiver estimate the amount of offset, thereby allowing the frequency of the VCO to be adjusted accordingly. Such a method can easily be implemented in programmable logic devices such as FPGAs and PLDs.

Phase Noise

The use of VCOs at the transmit and receive side not only leads to frequency offset, but also phase noise due to inherent jitter present in practical VCOs. By combining the use of a PLL and the insertion of training symbols the effects of phase noise can be reduced. Some OFDM systems make use of contain training sub-carriers known as pilot tones. These pilot tones are modulated by a BPSK sequence that is known by the receiver. Such modulation schemes can easily be implemented in general purpose processors such as PowerPC and DSPs like the C64xx from Texas Instruments.

Peak to Average Power Ratio

A problem that is inherent to all (multi-carrier modulated) systems such as OFDM is that there is a significant difference between the average and peak power of the signal. The difference, which is also known as the peak-to-average power ratio (PAR) is due to the fact that the multiple carriers can add together to constructively create a very large signal, or destructively to create a very small signal. The wide variation makes for a challenging power amplifier design, as distortion must be minimized while keeping the average power low enough to accommodate the large peaks.

Besides limiting peak signals by either clipping or replacing the peaks with a lower amplitude pulse, there are several techniques that have been developed to reduce PAR. One method restricts the modulation schemes of the sub carriers to some defined phase and amplitude relationship so that PAR is reduced. Another technique involves using block codes to modulate the sub carriers rather than using the data directly.

All of these methods require a fairly complex manipulation of data and modulation schemes. OFDM systems that consist of FPGAs and general purpose or digital signal processors allow developers to implement various schemes to reduce PAR.

Bringing it all together: Software Definable Radio Platform

The challenge of designing a 4G wireless system is compounded by the fact that a single standard for 4G wireless networks does not exist. The drive to be the first out to market has to be balanced against the risk of committing to fixed-function and inflexible hardware in an environment of changing standards. In order to keep up with ever changing air interfaces and their associated design challenges, the

solution is to use a software definable radio platform. The ideal SDR platform for rapid prototyping of OFDM and other MCM techniques incorporates a distributed network of signal processing nodes, combined with a high speed, scaleable and flexible embedded fabric. Such architecture makes it easy to channelize, transfer and process data. At the heart of the architecture is an embedded communications fabric.

What's in a Fabric?

The trend towards systems that support large numbers of logical channels within a communications structure has led developers to migrate from bussed and circuit switched architectures towards packet-based switch communications fabrics. Such a fabric consists of switches, which route data based on a destination address embedded in each transmitted packet. Parallel packet switched architectures such as Parallel Rapid IO and HyperTransport all support high bandwidth and low protocol overhead using multiple data lines running in parallel with a separate data clock line. Serial packet switched architectures such as Serial RapidIO and Switched Ethernet provide even more advantages over parallel implementations because they reduce signal skew, use less wires and are more ideal for board-to-board and chassis-to-chassis communications. Systems that make use of Serial RapidIO can take advantage of multiple point-to-point links each exceeding 1.25Gbps per links, thereby providing the bandwidth required in OFDM and other wideband air interfaces.

Using these packet-switched technologies we can create a modular system that can span multiple boards in a cPCI form factor. Such systems consist of an RF transceiver stage, channelizer, modem, and CODEC stage (Figure 3). The RF stage can be implemented using a combination of discrete analog components. Subsequent stages are implemented with a distributed network of processing nodes that can be FPGAs, DSPs, or general purpose processors. All of the processing elements will be held together by a packet-switch embedded fabric. The result is an architecture that can support multiple transceiver channels, whether it is OFDM, W-CDMA, or some other protocol. The packet-switch embedded fabric gives designers the ability to easily scale the system on a channel level, partition their algorithms across a variety of different processors, and dynamically reconfigure the dataflows as necessary when switching to a new air interface.

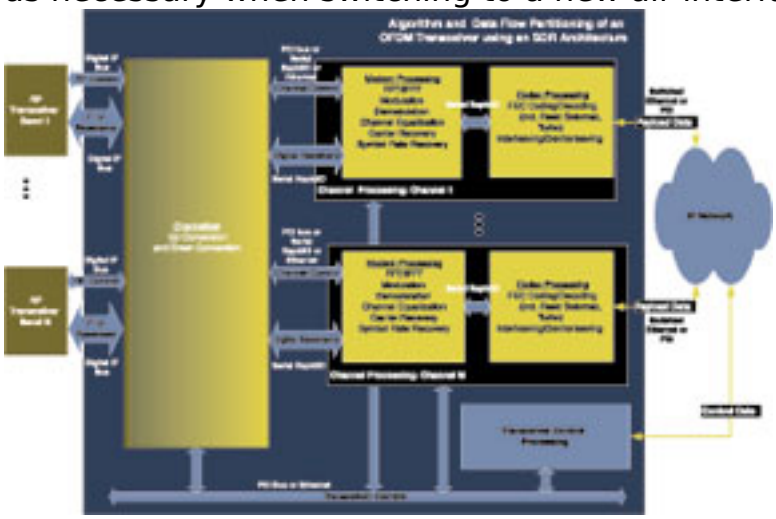

Figure 3. Distributed Processing SDR Architecture

A flexible and scaleable hardware architecture only gives designers one piece of the

puzzle. How quickly designers can rapidly prototype their OFDM or 4G subsystem is largely dependent on the firmware and software. FPGA vendors such as Xilinx provide IP cores for their FPGAs to speed time to market. Such vendors have several off-the-shelf IP cores that implement many of the elements of a communications system, from FEC coding/decoding algorithms, FFT engines, to digital down converters. Companies such as Spectrum Signal Processing also have developed a wide range of cores and algorithms for FPGAs and PowerPC processors, and also implement system-level solutions, by providing the necessary interfaces between the FPGA, PowerPC elements and the rest of the system. Along with a comprehensive inter-processor communications software library, developers no longer have to deal with the intricacies of board-level setup and control, and can focus on their core competencies in application development.

Figure 4: OFDM Transceiver Implementation in an SDR Platform

Figure 4 illustrates one implementation of an OFDM transceiver using an SDR platform. A PowerPC processor can act as the network interface, managing incoming and outgoing IP traffic into the transceiver. The CODEC stages are implemented in FPGA processing nodes, and are connected to the modem stages that are implemented in DSP or a general-purpose processor such as the PowerPC. Using the distributed transceiver architecture, the processing elements such as the PowerPC's are running multiple signal processing functions for one channel. The embedded communications fabric and inter-processor communications library allow the user to easily scale the system to handle multiple channels.

Conclusion

Individual technologies already exist which address many of the design challenges facing 4G developers. High performance processors from vendors such as Motorola and Texas Instruments, and Field Programmable gate arrays (FPGA) from Xilinx are all very capable of handling the intense algorithms needed for baseband processing. The emergence of switched-packet communications protocols like RapidIO provide the necessary fabric to allow high-speed, flexible and scaleable data communications between processing elements. The software defined radio platform presented here illustrates the marriage of high-performance processing elements with a switched-fabric interconnect creating a reconfigurable and scaleable platform designed to overcome the challenges of 4G systems.

J. Cyrus Sy is a Field Applications Engineer in the Wireless Systems Business Unit at Spectrum Signal Processing. He can be reached at cyrus_sy@spectrumsignal.com.

A list of references is available by request. Please email kpotts@reedbusiness.com.