Next-generation networks

A software-defined radio architecture for wireless hubs

By Louis Luneau and François Luneau

In 1996, for the first time PCs outsold TVs. We are now entering a new era where a new form of PC will be dominant – the digital cellular phone. Digital cellular phones will be as portable as your wallet. They will recognize speech and convert it to text. They will interconnect with the electronic system in your car and help you navigate streets. They will consult electronic yellow pages and give directions to the nearest gas station, restaurant, police headquarters, or hotel. They will collect your news and your mail and, if you wish, they will read it to you. They will conduct transactions and load credit into a credit chip on a smart card, which can be used like cash. They will play music. They will take digital pictures and project them onto a wall or screen, or dispatch them to any other digital phone or computer. They will have an Internet address and a Java run-time engine that will execute any applet or program written in that increasingly universal language. They will dock in a more powerful machine to perform more demanding functions. They will link to any compatible display, monitor, keyboard, storage device, or other peripheral through infrared pulses or radio frequencies.

The problem

What is needed to realize the above is a family of affordable, high capacity radios that support rapid deployment of new mobile services to users within a specific geographical area. This family of radios needs to simultaneously support multiple frequency bands, channels, protocols, applications, and needs to be capable of transmitting voice, video, and data, all at once. However, one solution is not a one-size-fits-all system. Rather, it will build upon a common open system architecture to provide affordability and scalability. It will also provide the ability to reconfigure waveform/protocol software in the radios. Therefore, although the radio will be used for a variety of radio applications, for the purpose of illustration, in this article we will be referring mostly to the cellular network applications.

Underlying concepts

The technical approach used to develop a family of radios is driven by several underlying concepts. This section discusses these concepts and their impact on a solution.

Multi-band. The product is capable of receiving from and transmitting to all users inside a cell. The receiver/transmitter has the capability of handling all of the operating spectrums that users may use. For example, cellular users in North America may operate in the cellular and PCS bands. For each band of interest, there is a wideband transceiver that has an ADC to digitize the analog signal. The system then integrates many wideband receivers/transmitters to cover the complete operating spectrum of interest. For the cellular user, this would mean a system that can integrate a wideband transceiver that covers the whole cellular band and a wideband transceiver that covers the whole PCS band. The system has the capability of integrating many wideband transceivers to cover all the operating spectrums of interest.

Multi-channel. Every user within a cell is allocated a dedicated channel to communicate with the hub. When SDR is used, the broadband signals are digitized using a single ADC/DAC, then the desired individual channel is extracted from the digitized signals through a Digital Down Converter (DDC)/Digital Up Converter (DUC) (the software-defined bandwidth and channel selector) and the data is then sent to the Digital Signal Processing (DSP) devices. The system should also have the capability of increasing the number of DDCs/DUCs and the processing power when additional channels are needed.

Multi-protocol. The mobile users within a cell can use several air interfaces (protocols). For example, cellular users in an area may get services from four different digital cellular phone service providers and four different air interfaces might be used such as AMPS, TDMA (IS-136), GSM, and CDMA (IS-95). A flexible radio platform has been developed to support these protocols. For each channel, the platform is able to select, configure, and/or program the appropriate building blocks within the DSP devices to support the selected air interface.

Leveraging existing ASIC and protocol software. Many companies already have digital radios that perform the desired functionalities. These radios normally use application specific integrated circuits (ASICs) and protocol software written in C and C++. What is needed is a way to port an existing design. An ASIC design can be ported on Field Programmable Gate Array (FPGA) and C/C++ protocol software can be ported to a DSP microprocessor or a General Purpose Processor (GPP).

Leveraging existing computer telephony integration (CTI). Naturally, wireless users need to communicate with the wired world, mostly other users connected via their telephones and computers. A wireless hub must therefore be able to connect to the wired world. Companies such as Intel and NMS Communications, involved in computer telephony integration, have developed many applications for interconnecting telephone and computer systems inside a wired network (PSTN and TCP/IP LAN/WAN). CTI companies offer a comprehensive set of carrier-grade, standard-based CompactPCI building block solutions for connecting to the wired world (telecommunications and the Internet). These building blocks use the high-availability and hot-swap capabilities of the CompactPCI bus with the H.110 bus and the new PICMG 2.16 (Packet Switching Backplane) to provide a scalable path to the next-generation Internet and voice communications network. The benefits of CTI technology are:

- **Architectural scalability** – Backplane interconnect speeds are user-definable, scalable from 10

Reprinted from CompactPCI Systems / January-February 2003

Mbits/sec to 2000 Mbits/sec (dual-star Ethernet full duplex) per node slot. This allows customers to start with a lower-cost, lower-speed set of components, and upgrade performance by swapping nodes or fabric components, as needed.

- **Improved density –** With the switch included in the chassis, no extra rack height is lost. Customers have the flexibly to connect from one to 21 node slots in the same chassis.
- **Increased system reliability –** All slots are served by redundant fabric connections, allowing for full system failover in the event of a fabric failure. Both switches use the same highly reliable power and cooling in the chassis. As part of Intel's fully IPMI-managed (Intelligent Platform Management Interface) architecture, power, and cooling systems may be remotely monitored and can be proactively serviced, thus avoiding possible system failure.
- **IP/Ethernet ecosystem –** IP/Ethernet is synonymous with the Internet. Use of PICMG 2.16-enabled products can accelerate convergence of telecom and IP-based Internet applications.
- **Accelerated system development –** By using standard-based modules rather than custom components, customers can focus on their core competence of building the next-generation application set.

Leveraging United States Department of Defence (U.S. DoD) research in software communication architecture. The U.S. DoD has developed a software architecture for software-defined radio called the Software Communication Architecture (SCA). The SCA requirements include the development of an architecture that will ease scalability, portability of waveforms, and extensibility. To develop an appropriate common open architecture, the U.S. military's Joint Tactical Radio System (JTRS) Joint Program Office took a two-step approach. The first step was to develop a baseline framework from which to proceed. The second step was to mature the baseline framework into a complete SCA, and validate that the completed SCA meets program goals. The result of this process is an SCA architecture that has been validated and matured. The OMG (Object Management Group) is currently advocating the SCA as a software-defined radio standard for use in commercial applications to:

- Promote portability, reusability, scalability, and interoperability of software radio-based platforms and applications
- Provide a standard platform-independent architecture to support software radio-based applications
- Promote the development of standard radio-based services for use by applications
- Promote the development of standard radio-based interface definitions for use in application development
- Promote, as required, the infrastructure and interface definitions needed to support:
  - Transparent security solutions
  - Safety critical solutions
  - Fault tolerant solutions
  - Real-time and embedded solutions
  - Ensure compatibility and consistency of resultant specifications related to software radios

The solution: Enabling technologies and software-defined radio

Software-defined radios are being regarded by many as the key solution for improving system performance in mobile communication. Software-defined radio is the next breakthrough in radio technology evolution and it will revolutionize the way we design and manufacture radios. The initial radio used analog-radio hardware. Today, most radios are digital. In addition to analog circuit, these radios use digital Integrated Circuits (IC) as well as computers and software to perform some digital filtering. Marketers readily describe these currently available digital products as software-defined radio components because software is used and can offer some of the advantages of true software-defined radios. But do not get fooled, these radios do not meet the criteria of a software-defined radio. With a few exceptions most software-defined radios are only laboratory curiosities being investigated by telecommunication and defense companies, and universities in the U.S., Europe, and Japan.

Recent technological improvements in Analog to Digital (A/D) and Digital to Analog (D/A) conversion, Digital Signal Processors (DSP), software architectures and tools are fundamental enabling factors for true software-defined radio implementations. The block diagram in Figure 1 illustrates current SDR technologies.

Antenna

Software-defined radio design begins at the antenna. The antenna is used to capture electromagnetic waves and transform them into electrical signals or vice-versa. In order to achieve improved service quality, SDR systems must have the capability to perform beam forming, diversity, and sectorization. Wideband antennas are also needed to access multiple RF bands dynamically (e.g. for military applications), sequentially or in parallel (e.g. in field re-configurable or multi-band PCS/wireless infrastructure). This technology, which has been known for decades in military practice, is now becoming affordable as the commercial demand for SDR expands.

Mixer

A mixer transforms the wideband RF analog radio frequency into an intermediate frequency (IF). This is done using local oscillators (LOs) to multiply the carrier frequency with a lower frequency. The key result is a frequency that represents the difference between the LO frequency and the carrier. This frequency is an Intermediate Frequency (IF) that holds all the information born by the carrier but at a level that can be processed by existing AD converters.

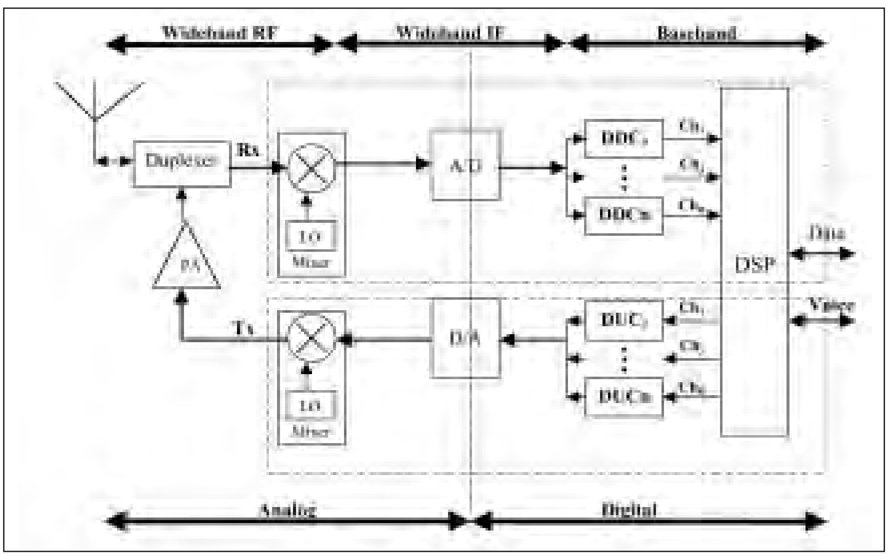

Figure 1

A/D and D/A converter

The A/Ds are considered critical components of SDR because the current technologies limit the bandwidth of the wideband receiver. The analog-to-digital converter takes the intermediate frequency, samples it at a rate of at least twice the frequency bandwidth, and translates each sample into a series of digital numbers. The ADC must be selected carefully. Challenges include accommodating high power in-band signals in some applications while in other applications what is important is achieving sufficient near/far performance. The following are key ADC characteristics:

- Sampling rate. The sampling rate limits the bandwidth of the wideband analog RF. The bandwidth must be smaller than half the maximum sampling rate because the Nyquist principle requires that the sampling rate be at least double the frequency bandwidth.
- Resolution. The resolution of each sample depends on the number of bits used.
- Spurious-free dynamic range (SFDR). Spurious-free dynamic range is defined as the range of signal amplitudes that can simultaneously be processed without distortion or be resolved by a receiver without the emergence of spurious signals above the noise floor. As a very rough analogy, imagine cranking the volume of your radio as high as possible without marring the desired signal with static and distortion. The spurious-free dynamic range of your radio would measure the distance between the lowest and the highest volumes with a clear signal. The SFDR is measured in dB.

For cellular applications, the ADC commands a sampling rate of at least 60 MHz, 14 bits and a spurious-free dynamic range of more than 90 decibels. Pushed by demands from software radio design companies, ADC manufacturers are introducing products that need to meet those requirements.

In the future, it is expected that ADC manufacturers will use super conducting technology that will operate at frequencies of several tens of GHz with SDFR between 120 and 160 dB. These ADCs will be able to translate microwave frequencies directly from the antenna into a digital bitstream. The high speed will enable the design of broadband SDR receivers that will cover bandwidths greater than 100 MHz. The high sensitivity will also enable the removal of any amplifier from the receiver chain (refer to Fujimaki et al. *Broadband Software-Defined Radio Receivers Based on Superconductive Devices.* Department of Quantum Engineering, Nagoya University, Japan).

D/A converters currently meet required performance and are not considered critical items.

DDC/DUC

Traditionally, analog radio tuners select a specific carrier frequency (channel). With software-defined radio, the DDC extracts from the A/D converter the samples that correspond to a specific channel (carrier frequency and its associated bandwidth).

The DDC has three major components: a local oscillator (LO), a complex mixer, and a decimating low pass FIR filter. The local oscillator is really a direct digital frequency synthesizer or DDS. It is sometimes called a numerically controlled oscillator, or NCO. It generates digital samples of two sine waves with a phase offset of 90 degrees, for sine and cosine signals. The sample from the LO is sent into a mixer.

The mixer actually consists of two digital multipliers. Digital input samples from the A/D are mathematically multiplied by the digital sine and cosine samples from the LO. The sine and cosine signals from the LO create a complex representation of the A/D input signal. The mixer translates the wideband input signal band to DC. The samples from the mixer are sent to a decimating low pass filter.

The decimating low pass filter selects only the carrier frequency of interest and the bandwidth of that signal. To set the bandwidth of the filter you need to program a parameter called the decimation factor.

For example, let's suppose that a 30 MHz wideband RF input was sampled by the A/D converter at 60 MHz. Let's also suppose that the signal of interest is at 20 MHz and that the bandwidth is 30 KHz. By setting the local oscillator at 20 MHz and the bandwidth of the low pass filter at 30 KHz, we can extract only a 30 KHz signal.

The DDC and DUC can be implemented in an FPGA or a commercial-off-the-shelf DDC or DUC chip.

Digital Signal Processing (DSP)

The fundamental building block in a SDR is the DSP. In most SDR implementations, the DSP performs functions like channel equalization, filtering, synchronization, modulation, demodulation, frequency spreading, frequency de-spreading, channel coding, source coding, etc. Many different DSP devices are available to provide digital signal processing. The most commonly used is the DSP microprocessor. In addition to the DSP microprocessors, the following re-programmable devices are available for the radio designer: field programmable gate arrays (FPGA) and general-purpose microprocessors (GPP).

The signal processing requirements for military and commercial radio systems employing high data rate signals or spread spectrum modulation easily exceeds the processing speeds currently available in off-the-shelf DSP microprocessors and GPPs. Modern field programmable gate arrays are the only re-programmable components that can implement functions beyond the capabilities of today's DSP microprocessors. In fact, they have the potential to provide performance increases of an order of magnitude or better over traditional DSP microprocessors, but with the same flexibility. These devices can provide the programmability of software, the high speed of hardware and can be re-configured on a real-time basis with no physical change to the hardware. In fact, FPGAs are really "soft" hardware, in that they are a good compromise between flexible all-software approaches that unfortunately limit throughput, and custom hardware implementations, which are not as flexible. FPGAs offer a powerful approach – an architecture tailored to the specific application. As the logic in an FPGA is flexible and amorphous, a DSP function can be mapped directly to the resources available on the device. Modern FPGAs have sufficient capacity to fit multiple algorithms into a single device along with the interface circuitry required by the application – a single chip solution. However, since FPGA is a recent technology, the implementation of DSP algorithms in FPGAs presents several technical challenges over DSP microprocessors. Until these challenges are solved, we believe that a heterogeneous computing environment consisting of FPGAs, DSP microprocessors, and GPPs is the best solution.

DSP processing challenges can be dealt with by using a mixed approach in which an FPGA is used in the earliest stages, doing much of the filtering, modulation, and demodulation in fast digital logic and when the signal reaches the post-modulability refers to the scaling of content and implementation. For software-defined radio, the ability to increase and decrease the number of channels for a band of interest is very important. With *Flex*Cell, this can be done by increasing or decreasing the number of DSP boards, because the DDCs and DUCs that provide the channelization functions are on the DSP boards.

Single Board Computer (SBC)
The SBC executes the application software. This includes:

- **Human-machine interface (HMI) server application.** The HMI server application provides the system operator and maintenance operator with the capability to operate the system and get status information.
- **Multi-protocol air interface application**. The multi-protocol air interface application dynamically reconfigures the FPGA and the DSP microprocessors with an instance of required air interface (waveform) netlist/software for each DDC output (channel) and DUC input. This is the application that gives our system the multi-protocol capability.
- **Re-configurable application.** This application allows us to remotely add new waveforms to the system or modify existing waveforms in the system. This application gives us an important characteristic of future radio systems: re-configurability. Reconfiguration is required because of the large variety of air interfaces (waveforms) in use and their rapid evolution.

Software architecture

*Flex*Cell software architecture is based on the Software Communication Architecture (SCA), as defined by the U.S. DoD JTRS project. The JTRS objectives are:

- To increase flexibility and portability
- To reduce support costs
- To facilitate system upgrades
- To reduce acquisition and maintenance costs

Accordingly, the SCA provides the following features to support its objectives:

- Open common Architecture
- Multiple bands and protocols
- Multiple domains (aircraft, naval, cellular)
- Software reuse

The SCA is an open, non-proprietary specification. It specifies the software, hardware, security and networking architectures. It also specifies a core framework (common base of services). It is a modern architecture that is modular, component based and distributed. It uses industry standards, which consist of:

- A set of components that interact together to form a complete system
- Corba – The "pipes"
- XML – Rules that define how to use and connect the different components
- Core Framework – Set of common services
- Posix – The OS on which the SCA is built

Components are the basis for the SCA. A component is small piece of software that does a well-defined task. Each component also has simple and well-known interfaces. Components enable you to build more complex assemblies.

The SCA gives us the following advantages:

- Reuse of SCA architecture work already done
- Component architecture guaranties software reuse
- Standards-based, and thus easy to integrate COTS products
- Open specification makes it easy for partners to integrate their systems with our own

By conforming to the SCA, we can more easily work with partners, concentrate on our expertise, and build a flexible platform for the future.

Possible uses for *Flex*Cell – Bandwidth on demand

*Flex*Cell can be used for many different applications. The RF transceivers and waveforms (air interface protocols) loaded on the *Flex*Cell platform create the application. The following paragraphs illustrate a bandwidth on demand application.

Radio systems have been marked by a perpetual shortage of spectrum frequencies. If you use a spectrum analyzer across the 300 to 3,000 megahertz (UHF) band, you discover that much of the radio band is empty most of the time. This unused spectrum might be available for transmission if we could take measurements and know exactly when and where to send the signal.

The behavior of cellular users in your geographic areas shows that the cellular networks are used up to 80 percent of their capacity between 0700H to 0900H and 1600H to 1800H. The rest of the day only 20 percent of the system capacity is used. On the other hand, the local news stations need bandwidth equivalent to 20 percent of the cellular network capacity throughout the day. The protocol needed by the news network is QAM 16 and OFDM – the bandwidth is variable.

A *Flex*Cell system that incorporates a transceiver for the cellular band and PCS band is produced. The system also includes the following protocol software: AMPS, TDMA, GSM, CDMAOne (IS-95), CDMA2000, UMTS, QAM 16, and OFDM. Finally, this is a software application that can supervise the bandwidth volume, detect trouble spots in the system, and then automatically distribute a variable amount of bandwidth (depending on the application), at a moment's notice, to where it is needed most is loaded on the platform. The *Flex*Cell platform configuration can meet the needs of digital phone users and local news stations and drastically increase efficiency in the use of frequency spectrums and revenues.

Conclusion

*flex*Cell software-defined radio can support multiple RF bands including not only 2G and 3G mobile systems, but also HF band, VHF/UHF military band, radio navigation aid band, low-powered ISM (industrial, scientific, and medical) bands, GPS band, IEEE-802.11x wireless LANs and Bluetooth, and the Local Multipoint Distribution Service (LMDS) band, etc. This is accomplished by including in the *flex*Cell platform an RF board for each band of interest. The RF board receives the wideband analog RF signal from one of the antennas, converts the signal through analog circuits to a wideband analog IF signal, and then converts the signal to a wideband digital IF signal using an ADC. The wideband digital IF signals from all RF boards are sent to an IF bus. The IF bus provides the wideband interconnect needed to achieve sufficient fan-out from the ADC to the digital signal processing elements on the DSP boards which includes DDCs, FPGAs, DSPs, and GPPs. This architecture simultaneously supports multi-band, multichannel, and multi-protocol. It provides channel scalability because to decrease/increase the number of channels, we only need to remove or add DSP boards since the channelization functions and baseband processing functions are on the DSP boards.

The *flex*Cell hardware architecture is designed around the CTI high-availability open architecture, which consists of a CompactPCI bus with the H.110 bus and the new

PICMG 2.16 (Packet Switching Backplane). This architecture allows *flex*Cell to use COTS boards to interconnect the system to PSTN and Internet.

The *flex*Cell design makes reuse of the U.S. DoD Software Communication Architecture (SCA). It is a modern architecture based on an object-oriented and component-based approach that includes message passing and encapsulation. It uses industry standards consisting of: a set of components and their interface, Corba, XML, Core framework (set of common services), and Posix. In the near term, to take advantage of existing DSP and GPP compilers, and to meet real-time performance requirements, the software is implemented using C++ and C. This open software architecture and implementation language provides the opportunity of reusing future partners' existing software (air interfaces and application software) and ease the integration of our solution in their systems.

An FPGA architecture for the *flex*Cell SDR radio has been developed to provide the methodology for incorporating changes and updates into the system. This architecture ensures that the FPGA design is appropriate for the problem at hand and allows for reusability, flexibility, and scalability.

In summary, *flex*Cell is a flexible market-ready platform, built to ensure current and future needs of wireless hubs.

Louis Luneau *is a founding partner in Radical Horizon and its Chief Technical Officer. Previously, Louis was the president of Isaac Newton Technologies, which he founded in 1989. Louis is the architect and visionary behind the flexCell software defined radio system. Louis obtained a Bachelor's degree in engineering and a diploma in telecommunications in 1991 from École de Technologie Supérieure in Montreal. E-mail Louis at lluneau@radicalhorizon.com*

François Luneau *is the Director of Project Management for Radical Horizon. During a 20-year military career, François spent 12 years in project management. He was selected by the government of Canada to become a risk management expert and was trained to perform risk evaluation of software projects using a technique called S: PRIME. François obtained a Bachelor's degree in engineering and management in 1983 from the Canadian Royal Military College. He did post graduate studies in software engineering from 1992 to 1994. E-mail François at fluneau@radicalhorizon.com*

RADICAL HORIZON

I N F O

For more information, contact Louis or François at:
Radical Horizon
1751 Richardson Street, Suite 8153
Montreal (Quebec), Canada H3K 1G6
Tel: 514-937-1757 • Fax: 513-937-1548
Web site: www.radicalhorizon.com