# Software-defined Radio Comes of Age

Jeffrey Steinheider
Mobile Radio Technology, Feb 1, 2003

In 1991, the term "software defined radio" was coined to describe radio devices implemented in software and running on generic hardware. This idea of software-programmable radio technology was a departure from traditional, hardware-specific radio architectures, and offered improvements in flexibility and upgradability.

Before these multi-band, multi-mode radios could be fully implemented, however, advancements were necessary in the processing, A/D and RF hardware. These radios were not targeted for development until the year 2000 or beyond.

## In the beginning

Early cellular basestations were analog, with a shift to basestations built around digital hardware in the late 1980s. The evolution of SDR continued in the 1990s mostly through military interoperability endeavors. Commercial applications of SDR remained neither practical nor cost effective until the past few years. Advances in A/D converters, RF technology and processing hardware have allowed SDR to finally achieve commercial viability.

SDR encompasses a wide range of communications systems, from reconfigurable hardware-based digital radios to fully programmable software radios, but can be grouped into three categories.

## The first tier

The simplest example of a SDR is a dual-mode cell phone. This is the modal SDR. A dual-mode cell phone has two hardware radios in it, one for each supported standard. Software determines which standard needs to run and activates the correct radio. Though the phone allows switching between the modes built into the radio, the user is limited only to those modes and lacks the ability to upgrade the system with new waveforms.

## The next generation

Reconfigurable SDR is the category of software radio that has been built for defense applications over the last decade, typically involving a combination of processing technologies such as application-specific integrated circuits (ASICs), fieldprogammable gate arrays (FPGAs), and digital signal processors (DSPs). Despite good performance on current systems, the software investment for these specialized systems is high, and they rapidly become obsolete as technology development accelerates.

One such example is the SpeakEasy system. It was built around a combination of FPGAs and 40 MHz TI C40 DSPs. By the time the first prototype was demonstrated, COTS DSPs were available at 166 MHz. As the SpeakEasy software was tied not only to the C40, but also to a specific layout of C40s and FPGAs, the new DSPs could not be exploited.

## The final frontier

The most advanced type of SDR, Software Radio (SWR), maximizes software reuse across platforms and hardware generations. SWR implements the signal processing software as an application-level program running on top of a standard operating system (OS) (whether on general purpse - GP, central processing units - CPUs, DSP, or other processing engines), giving it the flexibility lacking in other SDR types.

The use of application-level software and an OS both reduces software development costs and allows the underlying hardware components to be upgraded without incurring the high cost of redeveloping the software. As a result, SWR systems can track the Moore's Law performance curve over time at a much lower cost than other types of SDRs.

As SDR technology progresses, the flexibility and performance of SWR will give it a clear advantage over not only traditional radio architecture, but also other SDR types. Its unique ability to add features through software upgrades and to enable a single radio to support multiple standards has drawn interest to this technology from many markets, from cellular providers to public safety agencies.

**Hardware architecture**

The architecture of a software defined radio can be divided into three distinct elements: a digital signal processing section; a section responsible for the conversion between RF and digital; and the antenna. While the performance of the antenna and RF to digital conversion play a key part in determining the capabilities of an SDR platform, the flexible digital signal processing is what qualifies a radio as being a software defined radio. FPGAs, DSPs, and general purpose processors are the three leading technologies that can provide the flexibility and processing power needed for SDR systems.

The hardware architecture, groups the hardware components into three blocks representing the antenna, RF-to-digital and processing subsystems. No hardware component in the architecture is specialized to any particular waveform. While the architecture places no limitation on the achievable waveforms, any given implementation of the architecture can only support some waveforms.

Each implementation supports a limited range of RF frequencies, bandwidths, and amount of computational power. For example, in order for a platform to be software upgradeable from 2G to 3G cellular standards, the implementation must be able to receive a 5 MHz wide band in the appropriate frequency ranges and have enough computational power to perform the 3G processing.

The first block represents the antenna subsystem. The interfaces to the antenna block are RF transmit and receive analog lines and a digital control interface. With these interfaces, the architecture can accommodate traditional passive antennas (for which the digital interface has no function) as well as advanced systems such as electrically controllable antenna arrays. The architecture does not specify a particular type of digital connection (e.g. RS-232), as this is a detail of the implementation.

The next block, labeled RF-to-digital, is the only layer of the system that contains radio-specific analog components. On the receive side, its sole function is to generate a digitized representation of a downconverted slice of the radio spectrum. On the transmit side, it generates an upconverted radio signal from a digitized representation. This block does not perform waveform specific processing such as demodulation or equalization.

The name of the third block, motherboard, is borrowed from the PC world because software radios look much more like computers than like legacy radios. Like a PC motherboard, this layer contains memory and processor components, and provides I/O to a network, to the user, timing support, and similar functions.

**Applications**

SDR technology can be applied to wide range of markets. Fundamentally, SDR technology can be used in any device that uses RF for communication, which encompasses a wide range of products including cellular basestations, military communications systems and public safety radios.

**Technology in cellular basestations**

Cellular standards evolve slowly, from analog in the 1980s to digital in the 1990s, and possibly to 3G sometime this decade. While the underlying processing, communications and DSP technology evolves rapidly, cellular service is limited to once-a-decade upgrades because the high capital costs of infrastructure upgrades are prohibitive.

For example, AT&T and Cingular are upgrading their networks from time-division multiple access to (TDMA) to the global system for mobile communications (GSM). This "upgrade" actually involves building out a new GSM network in parallel to their existing TDMA network, an initiative that costs each carrier upwards of $4 billion and requires a 10-year deployment to achieve a reasonable return on investment.

| CATEGORY | EXAMPLES |
|---|---|
| **Modal SDR -** | |
| Software configures the radio ASIC or fixed hardware does processing | Dual-mode cell phones |
| **Reconfigurable SDR -** | |
| All signal processing reconfigurable Significant use of FPGA or assembly | SpeakEasy and AirNet |
| **Software Radio (SWR)** | |
| Exploits Moore's Law Supports software reuse across platforms | Vanu |

A wireless network infrastructure using software radio technology can be software upgraded to new standards, thus deploying new standards more quickly and at lower cost than today's approach. Carriers can then increase revenue by rapidly implementing new revenue generating services as well as new systems that use spectrum more efficiently. A further benefit of SDR is reduced operating expenses - many of the maintenance and upgrades today that require truck rolls to tower sites could be serviced as remote software changes in an SDR system.

The architecture for a SDR base station is essentially a basic SDR with an array of processing elements that can be scaled to handle more capacity, or more complex waveforms. Using current x86 general purpose processors as an example, it is now possible to provide one GSM channel with 8 time slots for every 1 GHz of processing. Standard networking equipment such as gigabit Ethernet now has the bandwidth to supply digitized spectrum, and allows the use of standard PC servers with x86 processors to act as the cluster of processing units.

The radio section of a software radio basestation is responsible for converting a wide band of radio spectrum to a digital IF. This equipment is available today in the form of multi-carrier power amplifiers, wide-band upconverters and downconverters, and high-speed A/D and D/A converters. This provides a digital interface that is completely independent of the air standard, and able to support multiple channels of different standards in a band. When coupled with a SDR backend, it is possible to change air standards simply through a software upgrade.

**SDR Eases standards woes**

SDR can also mitigate problems carriers face when switching to a new standard. Generally, capacity is moved to the new standard slowly, so that customers are not forced to immediately upgrade their phones. Limited spectrum availability means the carrier must decide at some point to take away capacity from the old standard in order to add capacity to the new standard. A SDR base station can run two different air standards simultaneously, operating a control channel for each standard and saving an operator from having to make this decision at each tower.

Additional capacity then can be added to each standard on an as-needed basis, changing the number of channels used by each standard dynamically, depending on the number of users requiring voice channels for each standard.

**SDR and frequency allocation**

Base station hotelling is a new architecture for deploying cellular systems that takes advantage of SDR's flexibility to lower capital costs and make more efficient use of the spectrum. Companies are now separating the base station from the antennas in order to improve coverage in urban areas and add coverage to tunnels, stadiums, and within buildings by putting the antennas where they are most needed. These remote antennas provide the RF spectrum over a fiber optic cable back to a central location where all of the basestation processing resides.

This method also better utilizes base station resources, as channels can be allocated to different locations to match the load as it varies over the course of a day. For example, at rush hour, more resources can be applied to towers on the highway, whereas these same processing resources could be allocated to the downtown office area at other times of the day. It is no longer necessary to outfit towers with capacity for the peak load - capacity that will sit idle during off-peak hours. Adding capacity to the entire system is now as easy as adding a server to a rack in the central location, eliminating a trip to the tower.

Additionally, the benefits of a SDR base station still apply. It is possible to run multiple standards simultaneously from a single hotel site using the same hardware, even supporting multiple wireless services providers from the same infrastructure base. This ability to share infrastructure between standards or carriers greatly reduces capital costs for the providers.

**Military**

The U.S. military has a significant radio interoperability problem. The story is often told of army troops calling in air support to Grenada using their personal calling cards and using Fort Bragg as an intermediary to communicate. Such problems are avoided with SDR.

Interoperability problems are also an obstacle in joint operations, where each nation typically has its own radio systems. Recently, emphasis on peacekeeping, disaster relief, homeland security and other non-combat military operations has created further problems. In these roles, military units must communicate with public safety agencies, humanitarian organizations, and the civilian population. A single SDR device with the ability to support multiple waveforms significantly reduces the number of devices needed in the field. For military users, who must maintain, transport, power, and manage each device under challenging operational conditions, the benefit of a streamlined system is substantial.

SDR also promises to reduce military radio development and acquisition costs. Without SDR, new device development requires investing anew in the implementation of each supported communication standard. With SDR, the bulk of the implementation knowledge for a communication standard is captured in portable software, which can then be reused at low cost in new or different platforms. This software reuse holds the potential to revolutionize radio procurement economics by significantly increasing competition among platform vendors, leading to reduced per-unit costs.

The United States Department of Defense (DoD) recognizes the potential cost reduction of SDR, and has established the Joint Tactical Radio System Joint Program Office (JTRS JPO) to achieve that goal. The JPO has begun to acquire software implementations of a first set of 33 communication standards. The linchpin of the JTRS effort is a software standard called the software communications architecture (SCA), intended to ensure portability of the implementations across platforms from many vendors. The SCA standardizes the software's operating environment, and the control and communication mechanisms for both the hardware and the external interfaces of the radio. Many NATO allies have signed agreements to apply the SCA in their future acquisitions, and the JPO intends the SCA to become the basis for commercial SDR software standards as well.

**Public safety interoperability**

Public safety agencies in the United States struggle with interoperability problems when collaborating with other public safety agencies.

In major emergency situations, from floods to plane crashes, a large number of local, state and federal agencies respond to the scene. Their incompatible radios, ranging from legacy analog FM systems, to digital trunked radios and even commercial cell phones, leave them unable to efficiently communicate with each other. Decentralized purchasing decisions lead to different agencies within different municipalities acquiring the systems that best meet their particular needs. As a result of this acquisition process, their voice and data communications systems cannot interoperate, and their databases cannot share information.

In a mutual aid scenario, multiple agencies (some which may work together infrequently or may never have worked together previously) must work together and communicate with little opportunity for prior planning, frequently outside the range of fixed communications infrastructure and in difficult terrain. It is unknown when and where the response will be necessary, and who will be involved. Ensuring interoperability in this context requires an extremely flexible and rapidly deployable solution. The Federal Emergency Management Agency has identified radio interoperability as the one item that could have made the most significant difference in the rescue and cleanup effort after the 9/11 disaster.

**SDR's position**

SDR is currently the best solution to solve the communications interoperability problem by providing an immediate, cost-effective solution that does not require the agencies to purchase new radios. A portable SDR device brought to an emergency scene can enable interoperability between selected members of different agencies by creating communication links between different radios and establishing infrastructure where none exists, or supplementing inadequate existing infrastructure by serving as a basestation.

The unplanned nature of an emergency requires extremely flexible radio systems that are able to adapt to the situation's communications needs, making SDR the ideal technology for public safety radio systems.

**Conclusion**

Software Defined Radio has been in development for many years, but is only now achieving commercial viability thanks to advances in integrated circuits. Initial SDR systems will likely appear in specialized areas such as military and public safety applications; visibility in consumer markets will follow.

Eventually, as SDR is deployed in infrastructure and consumer devices, consumers will benefit from improved wireless coverage, and new services offerings.

The extent of SDR progress will be decided largely by RF technology. Digital processing exists today with enough capacity to handle many different waveforms, but SDR is limited by specialized RF chipsets that are optimized for particular frequency bands and waveforms. Once economic RF designs are developed that are tunable over broad frequency ranges and can handle several bandwidths, SDR will be able to truly demonstrate all of its advantages.