Software running on general-purpose hardware means not having to throw away your cellphone when standards change

A Cellphone For All Standards

BY BENNY BING & NIKIL JAYANT
Georgia Institute of Technology

Can you imagine keeping the same cell phone as you switch from one service provider to another to take advantage of deals and new phone service options? A new technology called software-defined radio could make that possible—of course, assuming that service providers choose to adopt it.

As any reader of the business pages knows, the most appropriate single word for describing the world of mobile telephony is "uncertainty." No one can say with confidence what forms third- (and future) generation cellular systems will take, nor even when they will be deployed. Making decisions about what protocols to adopt (cdma2000, UMTS/W-CDMA, EDGE, GPRS, or even some yet-to-be-finalized 4G technology) and therefore about what equipment to invest in has been difficult for the Verizons and Nextels of this rapidly evolving world.

Fortunately for them, software-defined radio (SDR) is now becoming an alternative. It makes the main characteristics of a cellphone—or any other kind of radio, for that matter—reconfigurable with software rather than with hardware alterations. SDR thereby makes it possible to reprogram cellphones to operate on different radio interface standards like North America's IS-95 (CDMA) and IS-136 (TDMA).

But that's not all. Putting much of a radio's functionality in software opens up other benefits. A mobile SDR device can cope with the unpredictable dynamic characteristics of highly variable wireless links. It can allow efficient use of radio spectrum and power, and nimbly jump from one radio standard to another (say, from GSM to PDC, when a European subscriber travels to Japan, or from GSM to EDGE, when she decides to upgrade her handset).

The flexibility goes even further. Engineers are free to pick and choose among the software modules they created for earlier SDRs and use them in designing later products. This code reusability creates an open platform for both the evolutionary development of wireless systems and their rapid prototyping. Even hardware design becomes simpler because SDR relieves it of many complicated analog circuits by performing their functions in software.

The technology is being taken seriously: SDR is widely used in cellular base stations, mainly because of its ability to simplify hardware design, and it is just starting to be put to use in cellphone handsets.

It is also being widely studied by the U.S. military in its multi-service Joint Tactical Radio System (JTRS). Here the goal is to allow software-based transceivers to simultaneously carry voice, video, and data, using a variety of modulation schemes on carrier frequencies between 5 MHz and 2 GHz. The software is based on an open architecture, in which the operating system is made publicly available—and thus portable—although it has security features such as encryption.

Aircraft, tanks, trucks, and soldiers can have versions tailored to their individual needs. By building to a common standard and migrating existing systems to that standard, the U.S. military hopes to ensure that all forces at all levels can communicate, which they cannot always do today.

In a program that predates JTRS, General Dynamics Systems Inc. (Scottsdale, Ariz.) developed a software-defined radio known as the Digital Modular Radio (DMR), production quantities of which it is now delivering to the U.S. Navy. These devices can perform the functions of at least a dozen separate communications devices—up to four of them at a time.

NOAH WOODS

This commercial and military interest in SDR has grown in recent years primarily because high-performance digital signal processors (DSPs) are now available at reasonable prices. Radio functions in ever greater numbers are being implemented in programmable digital devices rather than hardwired analog components.

As processors optimized for high-speed arithmetic, DSPs are technically capable of doing the job on their own; but in today's marketplace, it is more economical for them to share the work with faster or more flexible chip types: application-specific integrated circuits (ASICs), which use hard-coded logic to perform the arithmetic; and field-programmable gate arrays (FPGAs), whose programmable interconnect and logic functions can be redefined after manufacture. Deciding how best to partition the signal-processing functions of an SDR among these devices is one of the challenges with which cellphone designers must deal.

Although ASICs provide better performance at lower cost, their programmability declines as their level of integration increases. Consequently, radios handling multiple radio interface standards often need multiple ASIC devices. In contrast, several interface standards can be easily integrated into a single DSP or FPGA with no loss of flexibility worth noting.

Architectural blueprint

SDR transceivers implement many functions by running software on general-purpose hardware. The analog hardware for functions like frequency tuning, filtering, modulation, and demodulation is replaced by software that implements those functions digitally. Such an arrangement enables a single radio to reprogram its mixers and filters to handle multiple modulation schemes and to work across many frequency bands.

The conventional dual-mode cellphone [see diagram, p. 36] is a typical second-generation device. In North America, it would likely work on two kinds of networks—one based on the old AMPS (analog) standard, and one based on the European digital standard known as GSM. To do that, the phone has two separate transmitters and two separate receivers.

In analog mode, outgoing signals emerge from the (analog) signal processor and are fed to a chain of functional blocks of the same general design as in most radio and TV stations—the venerable superheterodyne architecture. From the signal processor, the signal is modulated onto a carrier, translated up to an intermediate frequency (I-F), then up to a higher radio frequency (RF), and finally amplified and sent to the antenna.

In a reverse process, the analog receiver downconverts the received analog signal in two stages, selects the channel assigned to its particular conversation by means of the analog filter, and then passes it on to the signal processor for demodulation.

The digital transceiver is similar except that operations on the received signal are carried out in a DSP instead of in single-purpose analog circuitry. These operations may include decompression and even decryption as well as filtering.

In essence, the dual-mode cellphone employs configurable functions with multiple firmware cores that are activated and deactivated as needed. SDR offers a more elegant approach, using programmable DSPs that first download and then run the functions needed to implement a particular standard.

The first step in transforming a conventional cellphone into an SDR system is to make as much of the circuitry digital as possible. To start with, this means eliminating the baseband analog operations. These are carried out on the input (usually

Conventional Dual-Mode Cellphone

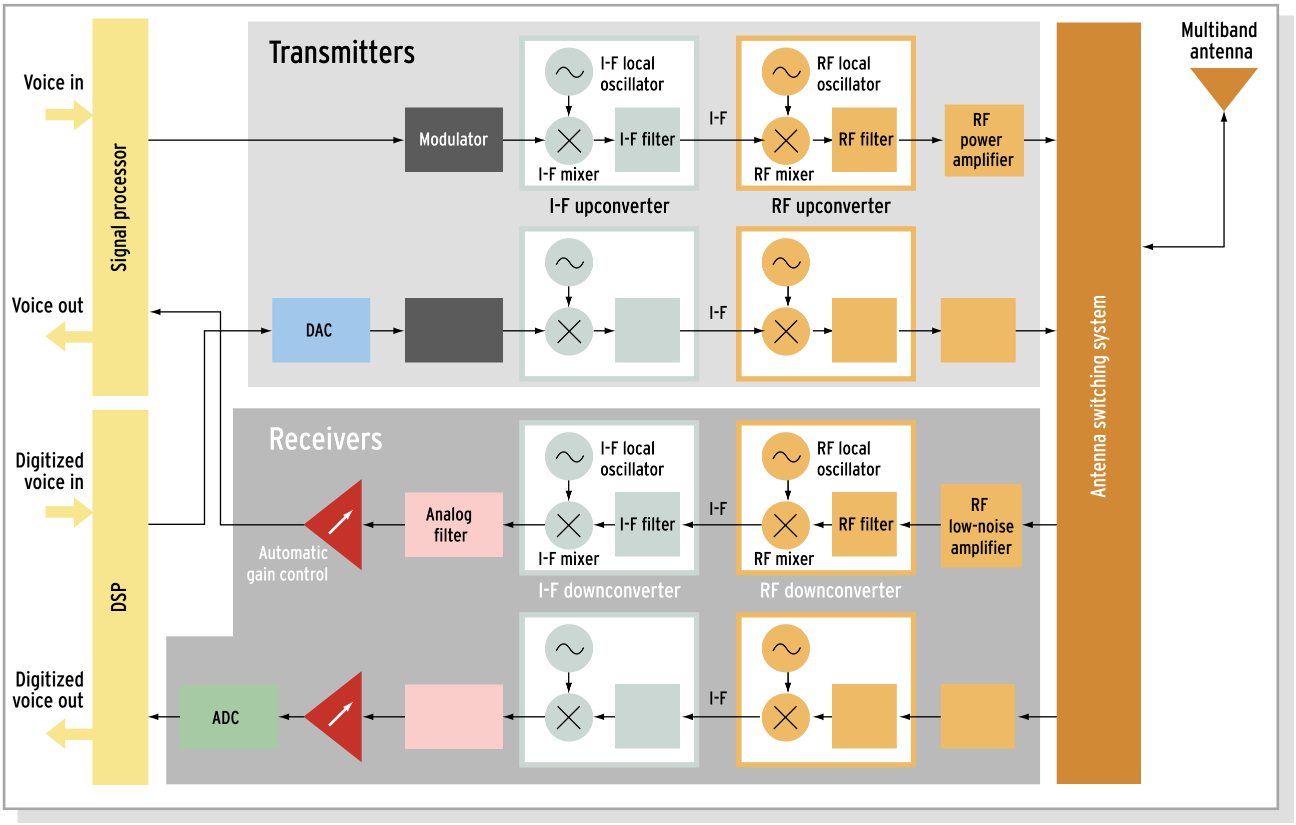

A conventional dual-mode handset works on two different networks—in this case, one analog and one digital—by simply switching between transmitters and receivers dedicated to each. This approach works, but it obviously doesn't scale very well to many more networks: each added network would require its own dedicated transmitter and receiver.

voice) signal while it is still occupying its native region of the spectrum and before it modulates a carrier and is thereby translated up to a higher frequency band [see software-defined cellphone, next page].

On the transmit side, this means digitizing the input voice as close to the microphone as possible so that all subsequent signal processing (compression, filtering, modulation) can be done digitally. Now the processing can be made programmable. (Of course, as the diagrams show, at some point the digital signal must be converted back into analog form for transmission, preferably as close to the antenna as possible.)

Similar reasoning holds on the receive side. The goal there is to convert the incoming analog RF signal into digital form as close to the antenna as possible, to process it digitally in a programmable device or devices, and then to convert it back into analog form as close as possible to the earpiece [not shown].

The next step, as more powerful analog-to-digital and digital-to analog converters (ADCs and DACs) become commercially available, is to achieve programmability at higher frequencies, first I-F and later RF.

Right away, the transceiver needs only half as much hardware as the conventional configuration. Instead of multiple transmit and receive chains, programmable radios have just one of each, which can be programmed to handle whatever radio standard is used by the subscriber's network—and even others that do not yet exist!

The higher-frequency functions, like filters and mixers, are hard to make programmable in silicon, which is by far the most common and least expensive chip material. The speeds required to implement them digitally exceed silicon's reach. The I-F section of a GSM cellphone runs at a few hundred megahertz, no problem for any modern silicon IC process. But to implement that same function digitally in the same processor that runs the rest of the cellphone's functions means executing around 100 billion instructions per second, a much more difficult feat, and one that may require silicon-germanium chips. To put the number in perspective, the single-purpose silicon chips in present-day cellphones execute roughly 10–100 million instructions per second.

Quartet of key technologies

The component technologies that form the backbone of SDR systems—and set their performance limits—are ADCs, digital signal processors, filters, and RF amplifiers.

The ADC is the most critical element of an SDR since its speed determines how close to the antenna the analog-to-digital conversion can be done. Defining ADC performance is always difficult because it involves specifying both analog and digital parameters. In essence, three main areas must be characterized: speed (number of samples per second), resolution

ARMAND VENEZIANO

(how many bits each sample is coded into), and linearity (how accurately the digital output codes are related to the analog input values).

For SDR purposes, the situation is as follows: The fastest ADCs in commercial use today—the kind found in the fastest digital oscilloscopes—can acquire roughly 10 billion samples per second. They cost too much and consume too much power to succeed in the cellphone application.

An ADC priced low enough and with enough resolution for use in cellphones can acquire about 100 million samples per second. That's high enough to digitize the I-F section of the transceiver—sampling the entire I-F cellular band and extracting individual channels in the digital domain—but it is nowhere near up to the job of digitizing the RF portion of the radio.

A straightforward approach to implementing a software-defined radio is to use a direct-conversion architecture in which analog circuitry downconverts the RF signal directly to baseband, skipping the I-F stage completely [see bottom diagram on this page]. The signal is then digitized by an ADC after which the desired channel is selected using selection filters implemented in a DSP.

While it is elegant in concept, direct conversion faces some formidable challenges in practice, according to Zoran Zvonar, systems development manager at Analog Devices Inc. (Wilmington, Mass.). Among the barriers to IC realizations of direct-conversion radios, there are difficulties in reducing dc voltage offsets to the point at which compensation techniques can eliminate them completely so they do not degrade overall system performance.

Software-Defined Cellphone

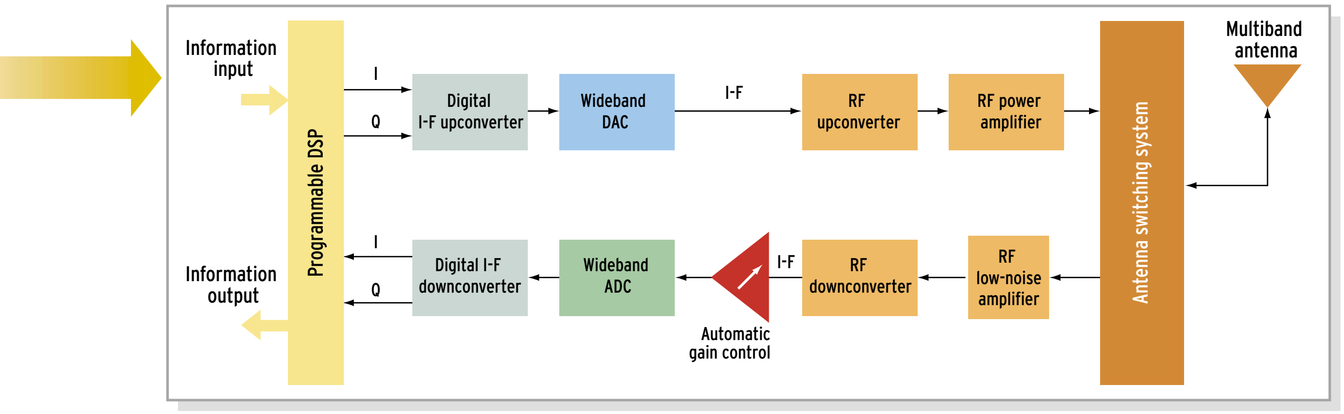

With its functions defined by software, an SDR cellphone handles different standards with but a single transmitter and receiver. The receiver has an analog RF front-end whose I-F output is converted to digital form by a single ADC and then downconverted to baseband digitally. Any channel is then selected from the baseband signal by digital filters implemented in the programmable DSP.

Direct-Conversion Cellphone

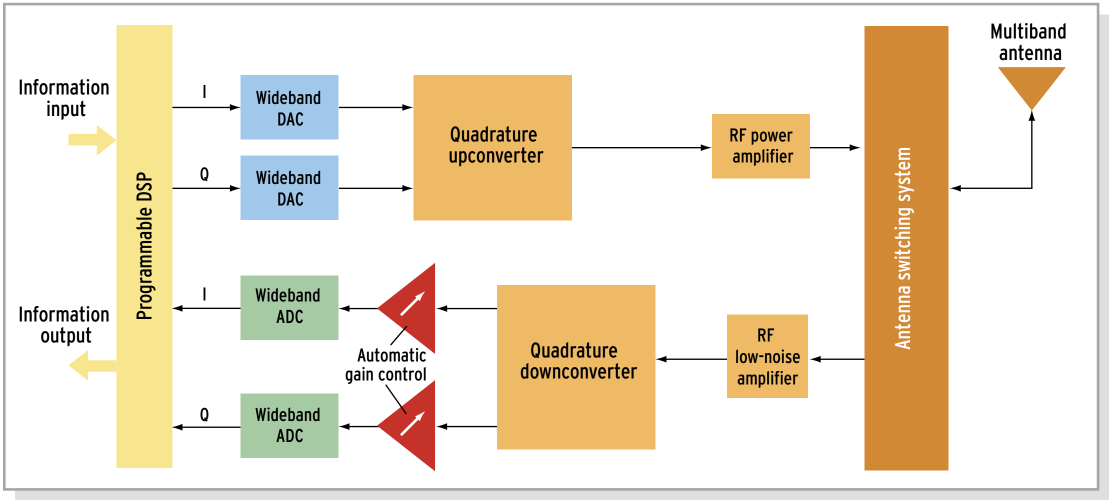

ADC=analog-to-digital converter DAC = digital-to-analog converter DSP = digital signal processor I-F = intermediate frequency RF = radio frequency

A direct-conversion software-defined cellphone can handle a wider range of frequencies and bandwidths because it goes directly from RF to baseband (or vice versa) without an I-F section and its inflexible components. Like many wireless devices, the unit saves bandwidth by using quadrature RF converters that split signals into in-phase and quadrature (I and Q) components, which can be transmitted over the same channel without interfering with one another.

The DSP is the fundamental building block of SDR. It can implement at least two radio interfaces at the same time, as is necessary when switching over from one radio standard to another. This sort of thing will become more common in the future, when cellphones need to link to other personal electronic devices over Bluetooth connections. In such cases, DSPs may have to run several parallel radio interfaces, much as some computers run multiple applications simultaneously. In general, functions assigned to a DSP include compressing and decompressing speech, modulation and demodulation, and filtering.

Although DSP speed is improving every year, single-chip performance is still a limiting factor for SDR applications, which—remember—must execute 100 billion instructions per second if they are to take on I-F circuitry functions. One path to high speed would be to optimize large arrays of DSPs for parallel processing, but the increases in size, weight, power consumption, and price would be unacceptable.

The best way to build an SDR is with a programmable DSP, claims Panos Papamichalis, director of the Imaging and Audio Laboratory in the DSP R&D Center at Texas Instruments, and also 2000–2001 President of the IEEE Signal Processing Society. But given the speed and power constraints of today's DSPs, an ASIC-based solution will be a more realistic approach in the short term, he says. Later, a mixture of DSPs and ASICs (probably in the form of one or more ASIC coprocessors integrated onto a DSP chip) will take over. According to Papamichalis, the fastest processors currently available from his company are in the C6000 family, which offers speeds of up to 600 MHz and an execution rate of eight instructions per cycle (or 4800 million instructions per second).


NOAH WOODS

Fast as they are, those devices, as we have seen, are too slow by a factor of about 20 to implement even I-F functions, not to mention RF. In practice, slower DSPs are favored today for telecommunications because of their lower cost and power consumption. The relatively low power requirement of TI's C5000 family makes up for speeds of up to 300 MHz and the only 600 million instructions executed per second.

Dedicated filters

Because filters are critical to the performance of an SDR, gains can be achieved through the use of dedicated digital filtering chips. These devices can perform the filtering function at a small fraction of the complexity and cost of a programmable DSP. True, they are not programmable, but since the frequency bands in cellphone standards are fixed, they are nevertheless useful in that application.

Filters affect an SDR's signal-processing speed, sensitivity, dynamic range, and ability to avoid interference. Their importance is reflected in their physical presence—filters constitute a third of the volume of a conventional dual-mode cellphone, being used in all three sections: RF, I-F, and baseband. Existing SDRs do not eliminate analog filters altogether. In fact, because the systems operate over wide frequency ranges, they require filters made of new combinations of materials with electrical properties far beyond those of conventional inductors and capacitors. The novel materials and modern filter fabrication techniques will lead to smaller, yet more flexible and adaptive wideband filters. In this area, superconducting and microelectromechanical devices (MEMS) may play important roles.

Regarding output amplifiers, silicon is the material of choice for moderate-performance RF amplifiers in the cellular and PCS bands, all of which fall below 2 GHz. A new technology involving silicon-germanium transistors promises an extremely low-cost, low-voltage approach for analog RF power amplifiers as well as receiver front-ends operating well into the millimeter-wave region (up to at least 40 GHz). Another development in RF technology is the design of ultralinear power amplifiers to process multiband signals from multiple transmitters simultaneously and add them coherently with high fidelity. SDR designs based on these amplifiers will improve power efficiency and consume less space than traditional amplifiers. But such multiband amplifiers may still be very expensive to produce.

Adapting to change

So far we have stressed the value of embodying radio functionality in software as a means of protecting cellphone subscribers and network operators against unpredictable changes in the technological landscape. The concern is with insurance against obsolescence, if you will.

But SDR can do more. It can make all kinds of radios, not just cellphones, perform better by helping them adapt in real time to the rapidly changing characteristics of the wireless environment, which is an especially severe problem for mobile users. Specifically, software can add new functions to a radio without affecting its original functions or correctness. For example, it can implement RF power control to make a radio work better in a fringe reception area, or it may include additional code to mitigate interference in a congested radio environment. Conversely, specific functions and unnecessary code can be removed from existing transceivers to make them more efficient without affecting the remaining functions.

A major drawback of such adaptivity is that it typically introduces extra latency, or delay, into transmissions, since it takes some finite time to recognize and possibly react to a changed link characteristic. That latency can affect delay-sensitive traffic like voice and video.

Because they can monitor, identify, and make use of unused or underutilized frequency channels, multiband SDRs can help network operators use bandwidth more efficiently. For example, if the paging channel is not being used, these radios can use it to transmit other kinds of information like user data.

At this point, the question arises as to how functions and channels of mobile devices can be changed without violating regulatory rules. For example, the Federal Communications Commission currently authorizes each piece of equipment for a type of use and a specific frequency channel. Both national and international regulatory bodies will need to cooperate and resolve this issue before frequency- and service-agile SDRs can become widespread.

SDRs can also work with smart antennas on combatting interference, poor connection quality, and limited system capacity. The antennas are arrays of devices capable of sending and receiving energy in precisely controlled, changeable directions. They form beams that receive (or radiate) energy only in particular directions, providing gain in those directions and attenuation at other angles. Placed at a cell tower base station, they can track the signals of several users while suppressing other interfering signals. In effect, they create multiple noninterfering channels where there had previously been just one, enlarging the traffic-carrying capacity of the cell.

Because it is difficult to build an antenna—smart or otherwise—that provides high gain over a broad frequency range, SDR designers have tended to focus on concatenating multiple narrowband antennas, thus leading to multiband SDR systems. The number of antennas required to cover a specified part of the RF spectrum depends on how well amplitude and phase errors across the individual bands can be compensated for, which in turn depends in part on how wide those bands are. In the future, advanced RF techniques may allow the development of reconfigurable multiband antennas.

TO PROBE FURTHER, SEE PAGE 79

Dynamic software download

Having made the case for SDR, it remains to explain how, exactly, a radio's software can be changed. Upgrading application features by software download over the Internet is now widely accepted, even though it can be a slow and frustrating experience when performed over a dial-up modem. The concept can be extended to the wireless realm by having subscribers download complete radio standards, or even custom applications over the air whenever it suits them. The download process can be made simple and entirely transparent to the user by having the service provider control the complete download process.

To arrive at this happy state of affairs, an open application programming interface needs to be defined to provide a mechanism through which different vendors can develop compatible software and hardware interfaces. To that end, object-oriented technologies like the Common Object Request Broker Architecture (Corba) and Java can be used to support the download mechanism.

But downloading software is hardly without problems. For example, it is not clear how alterations in equipment functionality will affect the equipment's FCC certification. It is also crucial to verify the integrity of software download to protect users against fraud.

Interested parties

SDR research is being pursued not only in industry and the military, but also in academia. At The Georgia Institute of Technology, through funding provided by the state of Georgia's economic development program known as the Yamacraw Initiative, researchers are developing a high–data-rate wireless system using SDR. The project has two main objectives:

- To develop fully adaptive mobile wireless communications. This will include the coordination of changes in the application, network, access, and physical layers of the network as it changes from one wireless channel or standard to another.
- To develop a high-performance wireless system. This will include the use of adaptive antennas and coding algorithms to combat the effects of channel variation in time and space.

The SDR testbed includes wideband RF front-ends, ADCs, DACs, DSPs, and other commercial off-the-shelf systems. It serves as a functional prototype for the implementation, test, and integration of cutting-edge research developed by faculty members.

Meanwhile, farther north, the Laboratory for Computer Science at the Massachusetts Institute of Technology (MIT) has developed a powerful, multipurpose, handheld computer that combines the functions of a cellular phone, a wireless connection to the Internet, a pager, an AM/FM radio, and a television set. Probably one of the most innovative applications of SDR, the Handy 21 replaces many current communications gadgets with a single portable device. Selecting each function is similar to selecting applications using a mouse and a computer. The hardware is essentially a PC equipped with an antenna and a wideband ADC. The hardware allows the selection of any 10-MHz region of the spectrum, converts it to I-F, and then relays the signal to the main memory of the handheld personal computer.

MIT is also developing a universal logic chip, called a Raw chip, that will deliver excellent performance and energy efficiency at low cost. The chip may be customized to suit a very wide variety of applications. Intended to replace both general-purpose microprocessors and special-purpose ICs, the Raw chip has more than 1000 I/O pins that can be dedicated to data streams—about 10 times more than in today's microprocessors. The faster data input will yield much better performance and energy efficiency than is possible today.

To sum up, software-defined radio concepts span all types of wireless computing devices, offering the possibilities of enhanced flexibility, easier upgradability, improved performance, and even customizing radio devices to individual needs. The concept will really have a chance to prove its worth when and if it is accepted as a platform for innovative new services that can be downloaded from the network. It will be interesting to see whether incumbent service providers have the foresight to seize this opportunity or whether it will be a flock of as-yet-unheard-of new entrepreneurs who provide the killer applications and reap the killer rewards. ●

Michael J. Riezenman, *Editor*

The lead author is grateful to Zoran Zvonar, who reviewed the original version of this article and provided many useful insights.