techtrends *By David Marsh, Contributing Technical Editor*

Software-defined radio TUNES IN

MILITARY PROJECTS HAVE LONG BEEN THE TARGETS FOR SOFTWARE-DEFINED RADIO. HOWEVER, HYBRID ARCHITECTURES MAY PAVE THE WAY FOR STRATEGIC DEVELOPMENT WITHIN THE COMMERCIAL SPHERE.

Until recently an expensive R&D exercise, SDR (software-defined radio) is finally breaking cover. The last few months have seen a flurry of product announcements—from application-specific semiconductors to the first software-driven radio to gain approval from the FCC (Federal Communications Commission). The US military has a big interest here, having earmarked as much as $25 billion for SDR development through the JTRS (Joint Tactical Radio System) initiative. Its objective is to support some 33 waveform profiles from 2 MHz to 55 GHz—one of which includes some of the cellular standards—with one platform, rather than requiring a truckload of transceivers and a patch system to enable cross-agency communications. As JTRS overseer, the US Department of Defense is now working with agencies in Canada, Japan, Sweden, and the United Kingdom to foster this development. In the wider world of commerce, the technology similarly promises to shrink operators' costs and increase service-provision flexibility by using a generic, reprogrammable hardware platform. So, what's the truth behind the banner headlines, and how soon will designers benefit from taking a software approach to a technology that's traditionally rooted in analog hardware?

First off, it's worth reviewing what differentiates an SDR system from contemporary telecommunications technologies, such as the CDMA (code-division multiple access) and GSM (global system for mobile) variants that serve most of today's cell phones. CDMA and GSM systems already carry extensive programmable hardware for tasks that range from managing base-station links to baseband processing within individual cell phones. According to the FCC, SDR's definition is

disarmingly simple: "In a software-defined radio, functions that were formerly carried out solely in hardware, such as the generation of the transmitted signal and the tuning and detection of the received radio signal, are performed by software that controls high-speed signal processors." Similarly, the SDR Forum defines an SDR device as one that functions independently of carrier frequencies and can operate within a range of transmission-protocol environments. Architecturally, these definitions suggest transceivers that perform upconversion and downconversion between baseband and RF exclusively in the digital domain, reducing the RF interface to a transmit-channel power amplifier, low-noise amplifier for the receive path, and minimal analog filtering (**Figure 1**).

The generic nature of the hardware appeals to military and commercial operators alike, because it prevents operators from being locked into any one system supplier. Crucially, SDR will make it possible to upgrade a network simply by loading new software. Given estimates of $1 billion to upgrade a 2G (second-generation) network to 3G, this move enables massive savings in new equipment purchases and shortens the 10-year average network lifetime that previous economic models dictate. Moreover, base stations will become protocol-aware and capable of bridging otherwise-incompatible networks—an increasingly desirable technical goal as global operators consolidate their operations. (The politics are currently something else.) Such bridges will blur today's distinction between networking and telecommunications to the point that, say, a W-CDMA (wideband-CDMA) handset will be able to tap into a local WiMax infrastructure to gain broadband data access. Ultimately, handset manufacturers will use a common global platform, enabling manufacturing economies of scale that will reduce hardware costs to make their products competitive with today's single-band and multiband phones. The revenue stream for mobile devices will move from connectivity to a true service-provision model, with bandwidth-on-demand serving subscribers' needs from voice to video and allowing operators to dynamically share bandwidth among network resources.

Today, this grand vision is some way off due to formidable hardware and software obstacles. In the meantime, industry insiders agree that SDR will evolve in phases that reflect increasing technical ability in areas such as DSP and converter ICs, power management, and network-infrastructure design. From the software side, developers desire a common framework that enables and promotes portability (see **sidebar** "SCA standardizes software development"). The SDR Forum's predictions show commercial 3G telecommunications maturing by 2008 and beginning to move into 4G services around 2010. It identifies SDR opportunities opening with 2.5G services, such as the EDGE (enhanced data for GSM evolution) packet-switched service, which is operational in North America and beginning to roll out in Europe. The technology's commercial uptake is mostly starting with base-station manufacturers, such as Vanu, which recently became the first company to receive FCC approval for its software radio.

Terminals will evolve more slowly, due primarily to power requirements but also to the low manufacturing cost of today's cell phones. Most observers agree that SDR cell phones won't appear until around 2010, when the increasing availability of mass-produced chip sets will stimulate growth beyond the cell-phone arena. In the meantime, the first consumer-level SDR terminals are beginning to appear in less power-challenged mobiles, such as laptops and PDAs, and within vehicles. Philips Semiconductors recently announced its SAF7730—a single-chip, dual-IF car radio and audio DSP that enables designers to cost-effectively implement a range of functions, such as adaptive ultrabass boost, to further product differentiation using a single platform. Elsewhere, equipment designers such as UK-based RadioScape are using software techniques at the core of their DAB (digital-audio-broadcast) and DRM (digital-radio-mondiale) products. RadioScape bases its consumer-level RS200 DAB/FM receiver module on Texas Instruments' DRE200 receiver chip.

AT A GLANCE

- SDR (software-defined radio) promises operators massive cost reductions.
- Cognitive radio will soon furnish bandwidth on demand.
- SDR forces re-evaluation of superhet and direct-conversion technologies.
- Carrier-speed data converters are beginning to emerge.

ANALOG STILL DOMINATES RFEs

Digitally processing the transmission signal at carrier frequencies faces numerous difficulties that begin with data-converter requirements. Protocols such as GSM and CDMA use frequencies that can exceed 2 GHz, demanding converters that run at 5 GHz or more. Worse, converters require upward of 13 bits of resolution to preserve the dynamic range that allows subsequent processing to reliably extract the signal content. Microwave processes, such as SiGe (silicon-germanium) and GaAs (gallium-arsenide), are potentially suitable but neither cheap nor power-efficient. The demands on receiver-chain processing can

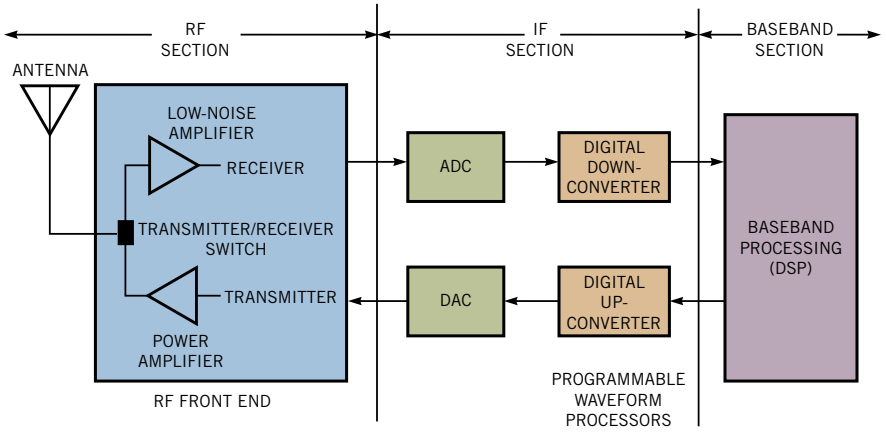

Figure 1 The ideal software-defined radio performs carrier-speed data conversion and operates on signals exclusively in the digital domain.

be stringent, too, depending heavily upon the modulation format and channel bandwidth. Although GSM shares 200 kHz among eight TDMA (time-division multiple-access) slots, 3G services have channel bandwidths as high as 5 MHz that the protocols share between multiple sessions. This shift from narrowband to wideband mandates DSP techniques for signal-content extraction but reduces the need for analog filters by moving filtering, channel selection, and baseband processing into the digital domain.

The paucity of carrier-speed data converters compels designers to consider analog RFEs (RF front ends) for upconversion and downconversion to an IF value that converters can handle. Although the receiver side is especially challenging, transmitters aren't without their problems, such as maintaining linearity conflicting with maximizing transmitter efficiency via switch-mode techniques. In general, the linearity requirements of ultrawideband communications make Class C amplifiers unusable without modification. GSM's simple signals allow power amplifiers to run at approximately 40% efficiency, but a 3G amplifier's linearity requirements can reduce efficiency to about 3%—requiring 700W for 20W of transmitting power. Techniques to improve linearity within nonlinear power amplifiers include predistortion, in which compensation circuitry distorts the incoming signal in the opposite direction of the transmitter's transfer characteristics. Such compensation requires digital techniques to calibrate the amplifier for sufficient accuracy but can improve efficiency by 20% or more.

But designers generally accept that receiving is an order of magnitude more difficult than transmitting; without the necessity of extracting rapidly changing information from signals buried within a sea of noise, transmitter algorithms are relatively simple. Equally, many receiver techniques similarly apply in the reverse direction, so most discussion centers on receiver architectures and how best to adapt schemes to fit within the SDR context. As a result, there's new focus on superheterodyne, direct-conversion, and hybrid techniques. Such radios aren't true SDR implementations in the sense that **Figure 1** conveys; rather, they extend the software-defined baseband processors that today's cell phones embody to become software-defined IF processors.

The demands on a generic wideband RFE are stringent, as research by the United Kingdom's DERA (Defence Evaluation and Research Agency) for a 1.6- to 2690-MHz receiver divulges (**Reference 1**). Key parameters for this multirole military radio include the ability to receive signals from −113 to +7 dBm; 100 dB of image rejection and 90 dB of IF rejection within the RF downconverter and first mixer stages; spurious rejection of 73 dB maintained through the RF and IF stages; a combined RF/IF-stage noise figure of −6 dB; and adjacent-channel rejection of 100 dB within a bandwidth that varies from 0.2 to 1728 kHz. Although it's possible to meet these selectivity metrics in a narrowband analog radio, such wideband frequency agility would demand multiple receiver paths in what's known as a dedicated-resources model. With contemporary airborne radios carrying as many as 100 channels, it's infinitely preferable to share resources as widely as possible. Examining current analog RFEs that work well in single- and limited multiband applications reveals significant performance barriers for effective SDR use.

With a channel bandwidth of just 30 kHz, the United States' first-generation AMPS (advanced mobile-phone-system) cellular service suits traditional superhets that still appear in many GSM cell phones (**Figure 2a**). The receiver translates incoming RF to a fixed IF value by multiplying the signal with a local-oscillator frequency in an analog mixer, easing gain and filter-stage design; further downconversion and demodulation follow. Using superhet-inventor and US Signal Corps officer Edwin Armstrong's original demonstration of transforming a 1-MHz signal into a 100-kHz IF, the first local-oscillator frequency is 1.1 MHz. This mixing process generates spurious images that differ in frequency from the wanted signal by twice the IF value, requiring image-rejection filters to quash noise that otherwise blocks information retrieval. These images are an intrinsic part of analog downconversion that leads designers to employ techniques such as using a high IF value to increase the distance between the signal and its image, hence easing filter requirements.

The first IF filter in a superhet cellphone receiver is almost invariably a SAW (surface-acoustic-wave) device that tunes the receiver to meet the protocol's blocking-signal specifications. An inevitable conflict exists between IF value and filter performance. Consider a GSM-900 receiver with a center frequency of 947.5 MHz and a 25-MHz bandwidth that uses an 80-MHz IF (a value that approaches the maximum conversion rate for commodity high-resolution ADCs). Assuming high-side mixing, its image frequency bands lie between 1095 and 1220 MHz. Filters can reduce these values, but not by the 110 dB that the blocking specification requires at this point.

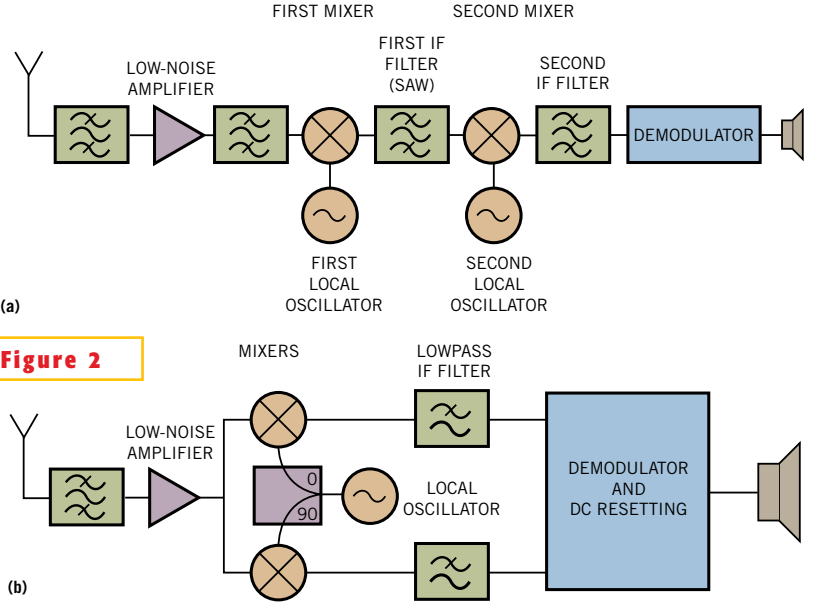

Figure 2

The traditional superhet mixes variable RF to a fixed IF before further downconversion and demodulation (a). A direct-conversion receiver transforms RF to baseband frequencies in one step (b).

Moving the IF up to 500 MHz creates images 1 GHz above signal frequency that are easy to filter but demands multiple downconversion steps to reach practical ADC sample rates. As well as replicating components, these steps create further signal degradation.

Because a true SDR receiver requires wide bandwidth to accommodate protocols such as CDMA or simply to search for signals of interest, the narrowband-superhet model is inappropriate. Future options that may extend its usefulness include filters that are programmable over very wide frequency ranges, possibly using MEMS (microelectromechanical-systems) technology. Such technology is currently in the research stage. But one approach that's gaining ground in contemporary SDR front ends uses variations of the direct-conversion-receiver architecture that's also popular in cell phones, in which it dispenses with SAW filters. Here, an I/Q (in-phase/quadrature) mixer transforms RF to complex baseband in one step, typically centered on dc, and a lowpass filter can then condition the signal before analog-to-digital conversion (**Figure 2b**). Mixing the signal down to dc—the so-called zero-IF approach—folds the spectrum around zero frequency, producing positive and negative frequencies and dividing signal bandwidth in two. The demodulator/DSP resolves the ambiguity in the received signal's instantaneous frequency by observing the I/Q phase relationship: If Q leads I, the frequency is positive; otherwise, it's negative.

Crucially, if the mixers are perfectly linear, operate within perfect quadrature, and have perfect amplitude balance, image signals self-cancel. In practice, imbalances in the real (I) and imaginary (Q) paths create a sideband around the signal of interest that translates to a wandering dc level, necessitating compensation schemes to retain complex signal components within their correct relationships. The I/Q components represent the coordinates of a vector in a constellation diagram of symbols that the protocol transmits; in turn, each symbol carries several bits of information. For example, EDGE encodes transmission data into 3-bit symbols that phase-modulate a carrier to produce eight equally spaced points around a circle (**Figure 3**). This 8PSK (eight-level phase-shift-keying) modulation scheme allows the carrier amplitude to fall to zero between transitions, which places impractical strain on receiver linearity requirements. The transmitter thus applies $3\pi/8$ rotation that produces 16 points, none of which then passes through zero. This setup eases linearity requirements, but the receiver must now keep track of vectors that lack unique positions. Worse, the receiver must disentangle these points from a sea of noise that blurs vectors before the decoder can map symbols back to data. As modulation schemes become more complex to pack ever more information into finite bandwidth, the distance between vectors decreases, and content becomes increasingly difficult to decode.

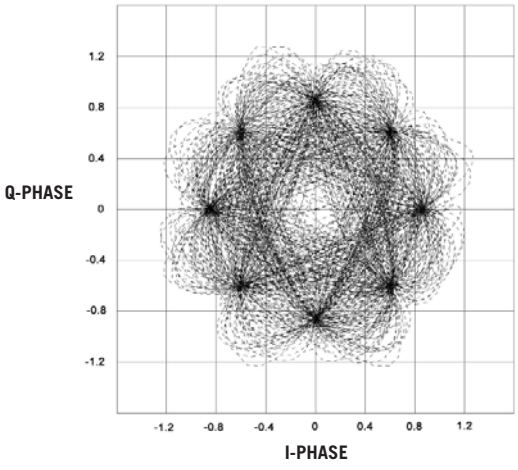

Figure 3 To pack more information within finite bandwidth, protocols such as EDGE phase-modulate symbols that appear as points in a vector diagram.

Much of the noise that receivers see originates externally, but, along with receiver-path imbalances, leakage from the local oscillator heavily degrades SNR (signal-to-noise-ratio) performance. Because any leakage in the local oscillator that couples into the mixer is at the same frequency as the wanted signal, the local-oscillator component mixes down as a dc term that appears as a baseband-blocking signal. Leakage originates from numerous sources, such as through the IC's bond wires and substrate and can couple back into the receiver's front end. The user can contribute to the leakage path as radiation couples through the body and back to the receiver's antenna, creating a varying dc offset that's difficult to compensate. This self-mixing interference can worsen with movement due to varying reflections between the antenna and the surrounding objects. One strategy that reduces these effects shifts the baseband frequency away from dc to create a low- or near-zero-IF receiver. An ac-coupling technique is now available to defeat dc offsets, but this approach requires ADCs with twice the bandwidth of a zero-IF design and complicates filter requirements.

Depending on how the designer partitions the radio, these considerations reappear with varying emphasis throughout SDR implementations. A conventional approach cascades a superhet RF-to-IF downconverter ahead of an ADC and I/Q transformation in digital mixers (**Figure 4**). In the absence of tunable antennas and filters that remain mostly in the research stage, this situation demands multiple RF channels to cover the bands of interest. The choice of IF value now adds ADC speed and precision to analog-downconverter considerations. Suitable 14-bit commodity ADCs that run at 70M to 125M samples/sec are available from vendors including Analog Devices, Linear Technology, Maxim, and Texas Instruments. Martin Streckfuss, European vertical business manager for wireless infrastructure products at Analog Devices, reports that most users are currently sampling IF at about 60 to 70M samples/sec: "The real target is about 170M samples/sec, which offers advantages in filtering and frequency planning to meet blocking-signal specifications and also greatly eases the first RF-downconversion stage."

Although engineers typically recognize oversampling for signal-definition preservation, subsampling also meets Nyquist's criterion. Rather than band-limiting the ADC input signal to half the sampling frequency in an antialiasing filter to prevent image generation, it's also possible to apply a bandpass filter that excludes signals below half the sample rate. Because the bandwidth is again half the sampling frequency, subsampling meets Nyquist's criterion, but the ADC's input bandwidth must be sufficiently high to avoid distorting the signal. One example is TI's all-CMOS, 125M-sample/sec ADS5500, which consumes just 780 mW from a 3.3V supply. Its input sample-

and-hold amplifier has an analog bandwidth of some 750 MHz that allows direct IF sampling beyond 300 MHz. For a faster conversion rate, TelASIC's monolithic TC1410 uses SiGe/BiCMOS to achieve 240M samples/sec with 500-MHz input bandwidth; the trade-off is as much as 14W power dissipation from ±5V supplies.

For carrier-speed performance, watch out for TechnoConcepts, a design house that specializes in SDR products. Its TSR (true-software-radio) technology boasts a 5-GHz delta-sigma digitizer that it originally developed with sponsorship from the US Department of Energy. Ron Hickling, co-founder and chief technology officer, says that the prototypes that will soon become available for sampling to key OEMs will be 10- or 11-bit devices built using GaAs (gallium-arsenide) MESFETs. But the company is working on 12-bit CMOS versions with SiGe front ends. "Our aim is 14+ bits with over 90 dB dynamic range, and it's coming real soon," Hickling predicts. He explains that, compared with sampling at IF rates, sampling at carrier frequencies reduces the resolution that's necessary to preserve the signal's phase and amplitude information. Thus, a carrier-speed converter of more than 10 bits can create redundant information. "The biggest challenge today is to move from superhet receivers that create large amounts of distortion due to their long mixer and amplifier chains," he says. Results from his early tests confirm that low-IF is preferable to zero-frequency downconversion: "Any amplifier has random noise, and, if that knee frequency is around, say, 10 kHz, it's relatively easy to filter. But, in a zero-frequency downconverter, noise in the 1-Hz region has a mean-square value that's virtually limitless."

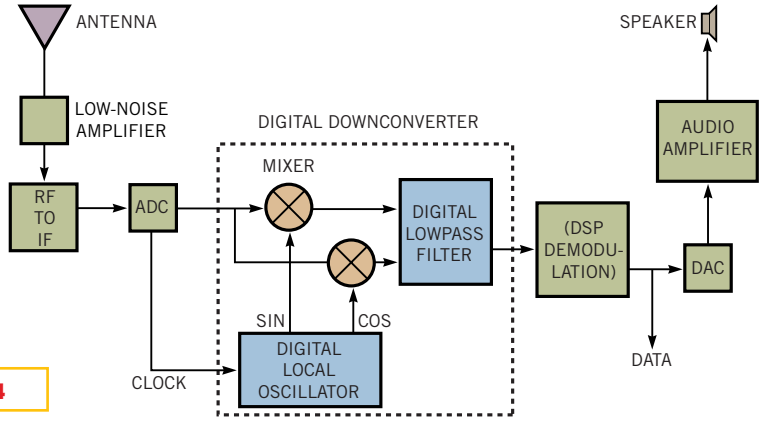

Figure 4

Many contemporary software-defined radios combine superhet front ends with digital IF and baseband processing.

FPGAs ACCELERATE DSP

Once digitized, a mixer built from two digital multipliers performs signal-I/Q conversion with precision that's virtually independent of local-oscillator frequency. The local oscillator is a digital synthesizer that uses sine/cosine look-up tables and a phase accumulator to generate samples of two sine waves precisely offset by 90°. Because the mixer runs at the ADC's native sample rate, the local oscillator varies the mix-down frequency by adjusting the phase advance between sine and cosine. The phase advance per sample is directly proportional to output frequency; hence, increasing angles represent higher frequencies. Using this FSK (frequency-shift-keying) approach and assuming 32-bit precision, the mixer's resolution betters 1 Hz from dc to 35 MHz. Crucially, the digital mixer's mathematical precision alleviates the unwanted artifacts of its analog counterpart, which far better suits SDR's wideband nature. Digital mixers can downconvert to zero frequency without dc-offset or significant image problems, thus allowing a lowpass filter to extract signal content. This filter is normally a decimating FIR (finite-impulse-response) design, in which the decimation factor sets filter bandwidth and dictates the output-sample rate to the baseband processor. This composite local oscillator/mixer/filter block is available as DDC (direct-downconverter) ICs from vendors including Analog Devices, Freescale, Intersil, National Semiconductor, TelASIC, and TI.

Alternatively, many designers prefer to include DDC functions within other logic, such as FPGAs or DSP arrays that can also serve baseband-processing duties. The shift from narrowband voice to wideband data-oriented services has a huge impact on processing power that conventional DSPs find hard to address. For example, estimates of the processing power that's necessary to accommodate GSM run at about 10 MIPS. Successive enhancements incur order-of-magnitude computational increases: GSM's first-generation GPRS (general-packet-radio-service) consumes about 100 MIPS, and EDGE demands approximately 1000 MIPS to deliver a theoretical best data rate of 384 kbps. True 3G protocols, such as UMTS (universal mobile-telecommunications system) and W-CDMA, require yet another order of magnitude—or some 10,000 MIPS—to process signal-channel bandwidths of about 5 MHz. And wireless data rates continue to evolve at a challenging pace, with developments such as the recent release 5 of the 3GPP's HSDPA (Third-Generation Partnership Project's high-speed downlink packet-access) specification. This W-CDMA enhancement carries as much as 14 Mbps, easily enough to deliver streaming video.

At the top end, such dataflow rates make it impractical to meet these requirements using conventional DSPs to process data straight from an ADC. Joel Seely, technical marketing manager at Altera's military business unit, observes that users then have options that reduce to designing an ASIC, wiring the logic in an FPGA, or using multiple DSPs. With development costs of approximately $10 million for a basically fixed-function device, designers are increasingly eschew the ASIC route. The availability of Altera's configurable Nios-II embedded processor to complement the company's Stratix-family FPGAs allows infrastructure designers at companies such as Panasonic Mobile Communications to design 3.5G base stations that support HSDPA. Seely notes that the key lies in the FPGA's ability to support multiple parallel-data streams: "Conventional processors are limited by fixed datapaths and finite clock speed, and their generalized nature often puts far too much processing power in the pathway for the task in hand. An FPGA allows you to break down the problem into multiple

parts to perform relatively simple operations in parallel at very high speed."

As an indicator, Seely points to the company's 2S180 FPGA, which comprises 180,000 logic elements—enough to accommodate about 100 fast Nios processor instances and capable of some 20,000 MIPS, whereas a representative wideband application might consume about 100,000 logic elements. With tens of millions of transistors, the device's 90-nm process geometry now leaks to the extent that static power dissipation is an issue. Seely says that, to combat this problem, designers are applying power only as required in systems that require multiple FPGAs, such as airborne military radios. Here, one FPGA may mop up as many as four channels to enable substantial power savings over DSP/GPP techniques. Toolchain support includes links into The MathWorks' Matlab/Simulink graphical development and simulation environment, as well as Altera's SOPC (system-on-programmable-chip) Builder, which adds block-diagram-level composition to its Quartus-II suite. "The progress in graphical development tools such as SOPC Builder is now exerting the same dynamics in the hardware space as the move from assembly language to C++ had for software," Seely concludes.

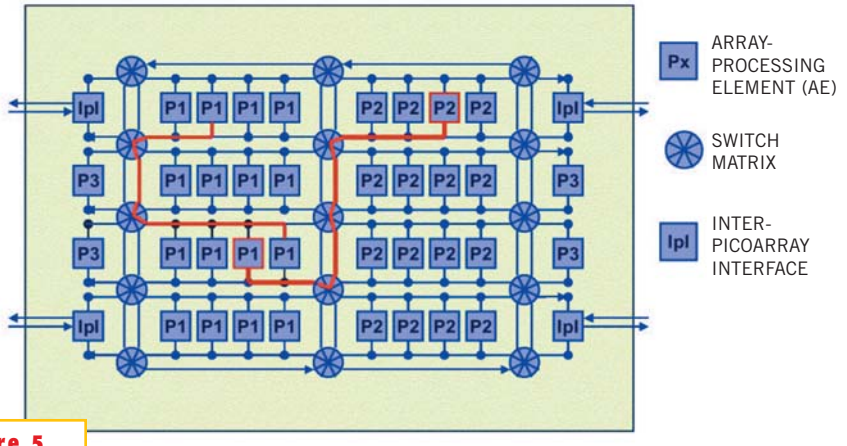

Figure 5

PicoChip's DSP array applies FPGA-like organization to more than 300 DSPs on a single chip.

Manuel Uhm, DSP marketing manager at Xilinx, agrees: "While VHDL [very-high-speed-IC hardware-description-language] allows users to handcraft hardware and is almost universally used by the military, user-friendly tools, such as our System Generator, greatly accelerate system-level development of signal-processing-intensive designs, from advanced military waveforms to commercial communications protocols." This object-oriented and model-based approach has several advantages over language-based approaches such as VHDL and Verilog, which Uhm says are more common in commercial circles. Links between Matlab/Simulink and Xilinx's System Generator enable block-diagram design that includes automatic FPGA configuration and real-time evaluation through hardware cosimulation. On-chip hardware enables DMA burst transfers between Simulink and the target device, accelerating simulation speed by several orders of magnitude. These tools support Xilinx's Spartan and Virtex families and include libraries of common DSP functions. Depending upon library-routine efficiency, Uhm estimates that, compared with an expert handcrafting VHDL, au-

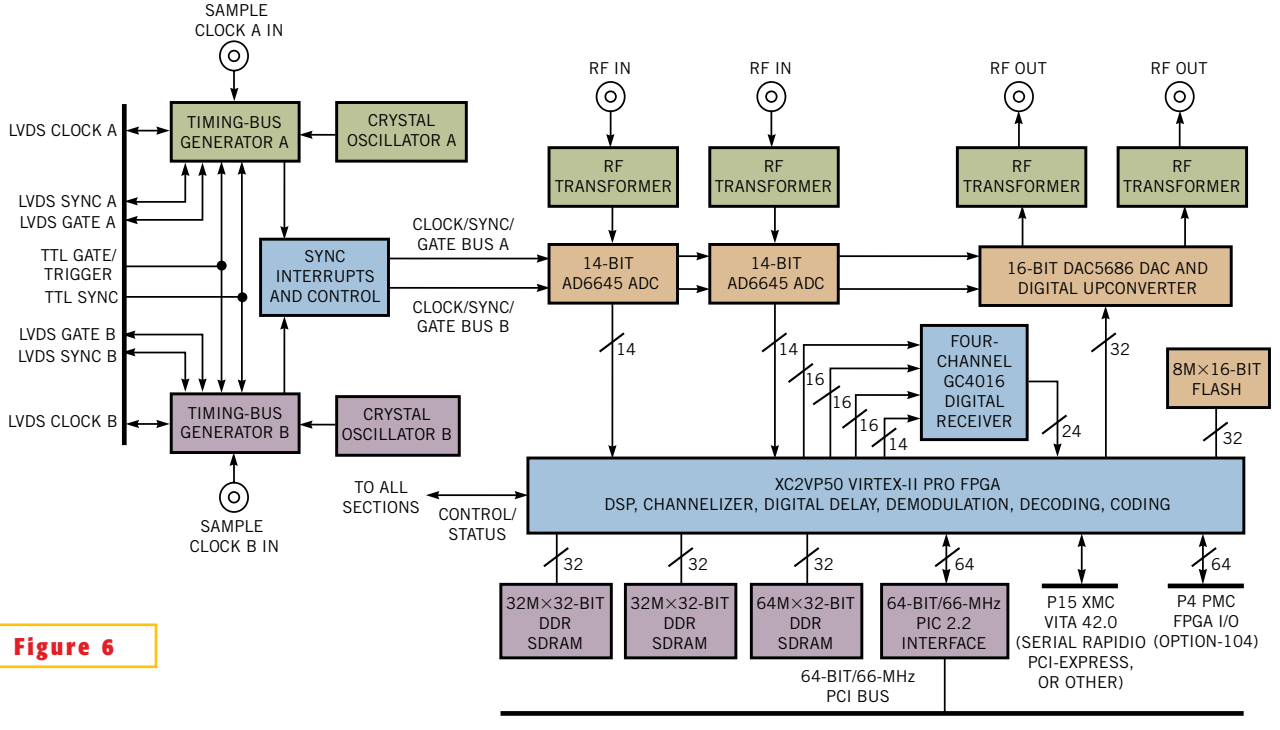

Figure 6

Pentek's 7140 transceiver module serves as an architectural model and reference platform for software-defined-radio development.

tomation provides approximately 80% optimization but saves at least 50% development time. Hardware efficiency is high, too; Uhm notes a recent benchmark test by independent test house BDTI (Berkeley Design Technology Inc): "Where a high-end DSP managed about 0.7 channels of OFDM (orthogonal-frequency-division-multiplex) decoding, the parallelism that FPGAs provide enabled about 60 channels."

DSP ARRAYS CHALLENGE FPGAs

The obvious approach to providing more bandwidth within a DSP-based architecture connects multiple DSPs in parallel, which is exactly what vendors such as picoChip offer. Rupert Baines, the company's vice president of marketing, says that its products apply an FPGA-like approach to DSP, in which each picoChip array houses 308 16-bit DSPs, 14 coprocessors, and programmable switches that interconnect the array processors via a 32-bit internal bus and out to adjacent array-processor blocks (**Figure 5**). Baines notes that the key is in making efficient DSP blocks: "The sheer scale of the cellphone industry has accelerated the knowledge base necessary to minimize DSP size," he says. He estimates that every 10% performance improvement in conventional DSPs requires another 30% in silicon area: "But 300 of our DSPs occupy less silicon than most Pentiums and are smaller than equivalent FPGA-based solutions." Depending upon configuration, resources within one PC102 device can furnish 240 general-purpose DSPs, 64 memory-control processors, and four supervisory processors that yield an aggregate performance approaching 200,000 MIPS from a 160-MHz clock. Such performance accommodates HSDPA and WiMax, which Baines reckons are today's toughest test. The picoTools design suite resolves synchronization issues by fixing the configuration at compilation time, resulting in predictable design operation and performance. Automatic internal routing optimization results in average resource use of around 90%.

Baines recognizes that picoChip isn't alone in believing that massively parallel DSP is the way forward, citing work at Freescale and Philips, among others. Freescale recently announced its MRC-6011 RCF (reconfigurable-compute-fabric) chip, an array of 96 DSPs. Robert Gach, design manager at Freescale's networking and computing systems group in Toulouse, France, says that each MRC-6011 houses six RCF cores, each of which comprises a 2×8 array of 16-bit DSPs. Total transistor count is more than 60 million using 130-nm process geometry. With 3G-baseband as its principal target, this first of a family of RCF devices includes resources optimized for CDMA tasks, such as correlation; initial 225-MHz devices can perform 21,600 million complex correlations/sec on 8-bit I/Q data. Again, tool-set design has been a major challenge, but full support is now available via extensions of the Metrowerks environment. Gach says that the MRC6011 is now qualified for production, with a future road map planning interfaces such as Ethernet and RapidIO, as well as control via embedded PowerPC cores.

Serving as an SDR reference platform that also showcases devices from Analog Devices, TI, and Xilinx, Pentek's 7140

SCA STANDARDIZES SOFTWARE DEVELOPMENT

One of the major barriers to a software-defined radio implementation has been the lack of a common software framework to stimulate co-development and facilitate interoperability. This situation has now changed with the release of the SCA (Software Communications Architecture), which is mandatory for all US military JTRS (Joint Tactical Radio System) development.

Researchers at Virginia Tech have developed a version known as Ossie (open-source SCA implementation embedded), that's available in the public domain to stimulate commercial SCA adoption. Because the military expects SCA to become a commercial standard, the architecture is designed from the outset to support COTS (commercial-off-the-shelf) hardware and software platforms.

SCA is now in Draft Version 3.0 of its requirements specification. Its key documentation, which appears at jtrs.army.mil, describes a basic architecture specification, defines the necessary application-programming interfaces, and considers security requirements that are unique to the military. To make the development environment as open and as interoperable as possible, the software structure builds upon CORBA (common-object-request-broker architecture) middleware and PoSIX-compliant operating systems that support COTS hardware, such as PCI and VMEbus boards.

At its heart lies the core framework that defines software interfaces and profiles that deploy, interconnect, manage, and communicate between software applications in SCA-compliant embedded systems. Other components include a platform-independent IDL (interface-definition language) that follows OMG (Object Management Group) specifications, and XLM (extensible markup language), which SCA's domain profiles use to identify the attributes of a system's hardware and software components. All hardware falls into a variety of classes of known behavior that communicate via defined interfaces, thus abstracting hardware considerations from higher level software layers.

The freeware Version B Ossie release is available at www.mprg.org/research/ossie. Virginia Tech operates the Web page and is home to the development team striving to meet the military's specifications. Ossie's framework hinges on Ace/Tao, an open-source version of OMG's CORBA that enables platform-independent data exchanges between compliant programs.

Another essential component is the Xerces-C++ parser, which adds XML read/write capabilities to user applications. These software components run on Fedora Linux and Windows 2000/XP platforms; programming-language environments include GNU's freeware gcc, or Version 6 of Microsoft's Visual C++ suite. Other open-source SDR software frameworks include OrcaCF from L-3 Communications and Scari from Canada's Communications Research Center. OrcaCF was developed with sponsorship from the US Air Force's research laboratory and implements SCA Version 2.2 on Linux platforms, also using Ace/Tao and Xerces components. Scari enjoys the SDR Forum's sponsorship, and is currently available to only its membership.

transceiver comes in several form factors, including PCI and CompactPCI modules (**Figure 6**). On the receiver side, the model 7140 accepts two 4-dBm full-scale analog RF inputs and transformer-couples them into two ADCs, typically 105M-sample/sec AD6645 parts from Analog Devices. Samples then pass through a Xilinx Virtex-II Pro FPGA for signal processing or for routing to other resources, such as TI's GC4016 quad-channel digital downconverter. Bob Sgandurra, Pentek's DSP product manager, explains that the company ships the board with the FPGA preconfigured to perform channel-selection, data-multiplexing, data-packing, gating, triggering, and SDRAM-control functions. Developers can take advantage of the spare capacity in the FPGA and its two embedded 405-PowerPC processors to develop custom algorithms. Pentek's GateFlow FPGA design kit complements the Xilinx tool set and includes VHDL source code and device-configuration information for the factory-installed functions.

You can reach Contributing Editor David Marsh at forncett@btinternet.com.

Various IP (intellectual-property) cores, such as FFTs, pulse-compression algorithms, and wideband digital receivers are also available, along with libraries of device drivers and functions.

According to Sgandurra, "[Pentek's] biggest design challenge was to include everything that a developer might need—partitioned in such a way as to compartmentalize essential housekeeping and data-management functions—without sacrificing the flexibility that FPGAs offer." Users can store the demodulated output from the FPGA for replay in onboard memory or stream it to either the PCI bus or a VITA-42 XMC (VMEbus International Trade Association extended-mezzanine-compatible) interface that furnishes switched-fabric backplane connectivity. The upconverter channel features a 16-bit DAC and DUC (digital upconverter), the 500M-sample/sec dual-channel DAC5686 from TI. Users can program the module to output single real streams or complex I/Q streams, as external hardware dictates. Sgandurra notes that his customers are partnering his company's hardware with COTS (commercial-off-the-shelf) RF modules from suppliers such as DRS Systems and Interad. Similarly, the Web site of SDR-system supplier Vanu confirms partnerships with RF specialists such as ADC and Raytheon. "COTS hardware gets customers up and running fast," Sgandurra says.□

References

1. Sharp, BA, PA Warr, RJ Wilkinson, and JP McGeehan, "The design of an analogue RF front-end for a multi-role radio," DERA 1998, www.argreenhouse.com/society/TacCom/papers98/09_01b.pdf.

For more information

For more information on the products and manufacturers this article covers, see the ***sidebar*** *"For more information" on the Web version of this article at www.edn.com.*

Talk to us

Post comments via TalkBack at the online version of this article at www.edn.com.

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, contact any of the following manufacturers directly, and please let them know you read about their products in *EDN*.

ADC
www.adc.com

Analog Devices
www.analog.com

Altera
www.altera.com

BDTI (Berkeley Design Technology Inc)
www.bdti.com

CRC (Communications Research Center, Canada)
www.crc.ca

DRS Systems
www.drs.com

Federal Communications Commission
www.fcc.gov

Freescale Semiconductor
www.freescale.com

Interad
www.microg.com/interad

Intersil
www.intersil.com

JTRS (Joint Tactical Radio System)
http://jtrs.army.mil

L-3 Communications
www.l-3com.com

Linear Technology
www.linear.com

The MathWorks
www.mathworks.com

Metrowerks
www.metrowerks.com

Maxim
www.maxim-ic.com

OMG (Object Management Group)
www.omg.org

Pentek
www.pentek.com

Panasonic Mobile Communications
www.panasonicmobile.com

Philips Semiconductors
www.philips.semiconductors.com

picoChip
www.picochip.com

RadioScape
www.radioscape.com

Raytheon
www.raytheon.com

SDR Forum
www.sdrforum.org

TechnoConcepts
www.technoconcepts.com

TelASIC
www.telasic.com

Texas Instruments
www.ti.com

Vanu
www.vanu.com

Virginia Tech's Mobile and Portable Radio Research Group
www.mprg.org

VITA (VMEbus International Trade Association)
www.vita.com

Xilinx
www.xilinx.com