Reinventing the Signal Processor

FPGAs are ideal for building high-performance, reconfigurable signal processing systems such as software defined radios.

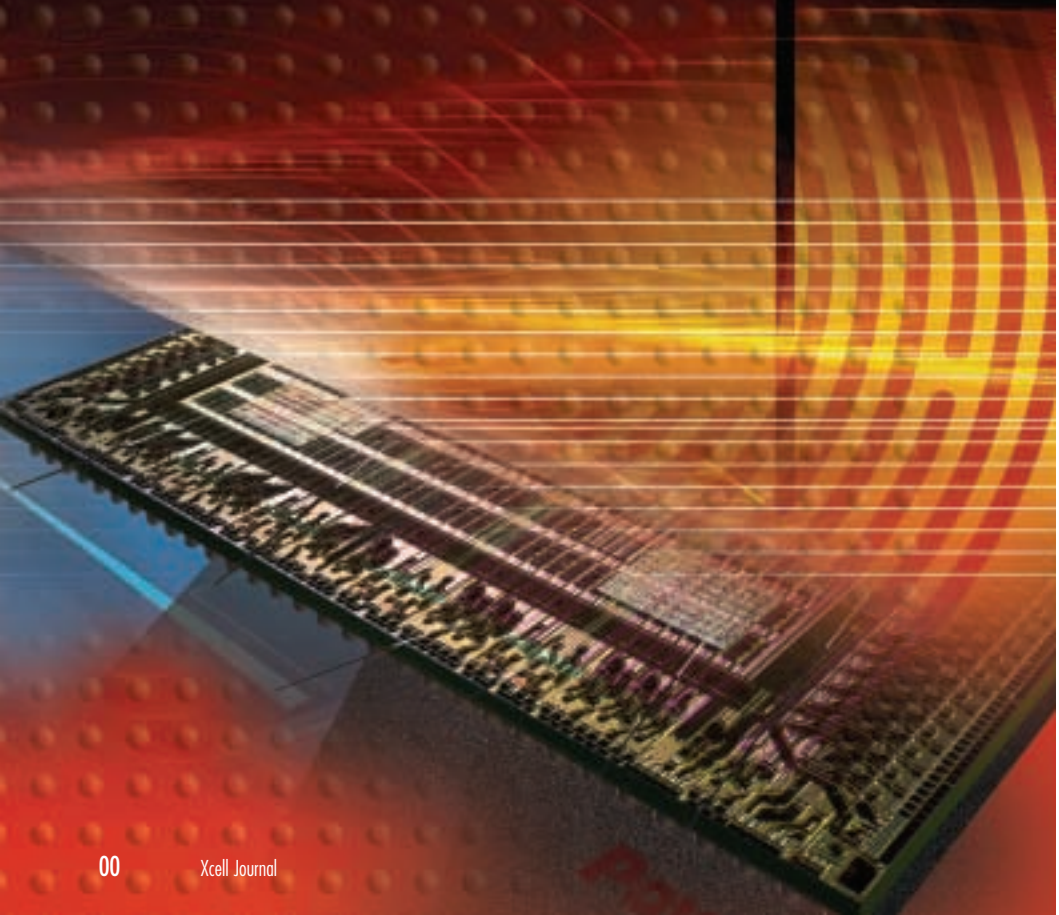

by Chris Dick, Ph.D.
Chief DSP Architect and Director
Signal Processing Engineering
Xilinx, Inc.
chris.dick@xilinx.com

The ultimate goal in software radio has been the realization of an agile radio that can transmit and receive at any carrier frequency using any protocol, all of which can be reprogrammed virtually instantaneously.

The Software Defined Radio Forum (SDRF) (*www.sdrforum.org*), an organization dedicated to supporting the development, deployment, and use of open architectures for advanced wireless systems, defines a software defined radio (originally coined by Joe Mitola in 1991 [1]) as radios that provide software control of a variety of modulation techniques. These include wide-band or narrow-band operation, communications security functions (such as hopping), and waveform requirements over a broad frequency range.

Figure 1 shows the architecture of a generic software radio. Smart antenna array technology is used for both the receive and transmit paths in the system. On the receive side, multiple high-bandwidth digitized antenna data is channelized, converted to baseband, and filtered – typically the sample rate is adjusted at this node. Other sections of the radio's physical layer (PHY) perform demodulation, synchronization, multiuser detection, adaptive interference cancellation, source decoding, forward error correction, beam forming, and adaptive equalization.

All of these computations present significant challenges for the radio PHY signal processing engine. Furthermore, much of the processing occurs at very high data rates.

Demands for Configurability and Agility

One of the driving objectives underlying SDR concepts is the desire to have a single hardware platform capable of servicing a number of radio environments. This type of reconfigurability could be used in several ways. For example, manufacturers developing infrastructure equipment or network operators building out a network could deploy a software radio system in Europe configured to support Universal Mobile

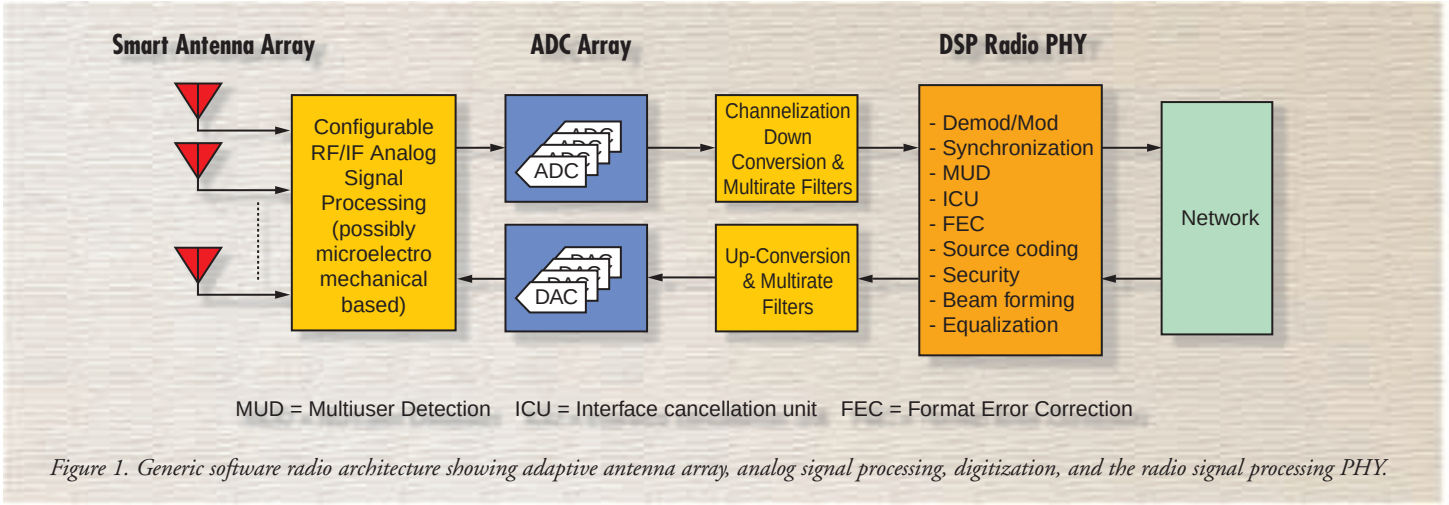

Figure 1. Generic software radio architecture showing adaptive antenna array, analog signal processing, digitization, and the radio signal processing PHY.

Telecommunications System (UMTS) or Global System for Mobile Communication (GSM) standards, or operate the system in the U.S. with a Code Division Multiple Access (CDMA)2000 radio personality profile. This one system could also be operated as a multimode radio in an environment that employs both wideband and narrowband CDMA communications.

Radio agility is important in situations where standards are fluid. For example, consider the evolution of the 3GPP standard and the length of time required for that standard to stabilize. Agility is also important during transition periods. As we move from 2G to 3G mobile cellular systems, multiple standards such as Personal Digital Communication System (PCS), (GSM), IS-95, Personal Handyphone System (PHS), DECT, EDGE, GPRS, IMT-2000, and CDMA2000 must all coexist.

Support will also be necessary. When the 4G wireless network build-out is completed, multimode operation will be required to support third-generation wireless direct sequence spread spectrum (DSSS) and orthogonal frequency division multiplexing (OFDM), the modulation scheme most likely to be deployed in 4G systems. From a network operator perspective, base tranceiver station (BTS) configurability could be used to dynamically allocate radio resources. This might occur on the time-scale of hours in order to provide the highest quality of service to the subscriber base at any given time.

Both manufacturers and network operators could also use a configurable BTS to permit field upgrades or bug fixes to equipment already deployed, by supplying a new BTS profile using the Internet or a microwave link from a radio network controller to the BTS. Soft radios can also be viewed as a means to protect infrastructure investments by keeping radio hardware from becoming obsolete as new standards and techniques become available.

Economics and The Effect on Software Radio Development

Although commercial technology and economics have always been inextricably linked, significant changes are occurring in both of these domains that will alter the way electronic equipment, including software radios, is developed and deployed. From a purely technological perspective, *The International Technology Roadmap for Semiconductors (ITRS)* shows that Moore's Law will remain in effect for at least another 15 years, and that in the year 2016 devices will be produced on a 22 nanometer node.

Yet Patrick Gelsinger, Intel's chief technology officer, announced at last year's International Solid State Circuit's Conference (ISSCC 2001) plans for a 20 or 30 nanometer process in 2010, delivering a device consisting of 2 billion transistors operating at a clock frequency of 30 GHz. The estimated power consumption of the device would be 3 to 5 kW, or a power density of 1 kW/square centimeter, about the same as a rocket nozzle. This has obvious thermal implications that must be dealt with using techniques radically different from today's methods.

Instruction set architecture (ISA) signal processors share many similarities to general purpose processors. Architectural differentiates such as very long instruction word (VLIW), super-scalar extensions, and various types of predictive enhancements are really micro-architecture evolutions of the basic architecture credited to von Neumann and his colleagues in the 1940s and 1950s. As such, signal microprocessors have leveraged most of their performance via a raw increase in clock frequencies. For example, in the early 1980s the first fixed-point signal processors supported clock frequencies in the 5- to 10-MHz region. Current-generation high-end ISA Digital Signal Processors (DSPs) use 600-MHz clocks and are on a trajectory to the Giga-Hertz region. Obviously, this curve has the same thermal pitfalls described above.

Although FPGAs take advantage of Moore's Law (and other advanced process technology such as all-copper interconnect and low-K dielectric substrates) to provide increased clock frequencies over time, their primary mechanism for supplying performance is completely different than the ISA approach. FPGAs exploit the large amount of parallelism inherent in most signal processing algorithms. With as many as 556 embedded multipliers and 125,136 logic cells in the Xilinx Virtex-II Pro™ Platform FPGA, we can readily see how these devices can be viewed as a naturally parallel processing engine that can take advantage of the rich parallelism in a software radio PHY.

The software radio PHY is a complex signal processing system in which algorith-

mic and functional-level parallelism can be leveraged to realize a high-performance system that does not rely on raw speed for its performance. The multiplier array could be used to implement space-time processing in a receiver, while at the functional level multiple turbo convolutional decoders could be operating concurrently to support multiple users, each with a 2 Mbps data rate in a 3G environment.

Manufacturers have made 60% more transistors available to circuit designers per area of silicon compared with what was available a year earlier. In contrast, the ratio at which designers are able to utilize transistors in circuits of any given tier of complexity has only been increasing at a rate of 20% per year [3].

This "design gap" is associated with performance supply and demand, but another aspect is methodology related. It is becoming an increasingly complex, time-consuming, and error-prone procedure to develop and verify a sophisticated ASIC. Furthermore, at a cost of $1-$2 million for mask set costs, it is becoming prohibitively expensive. ASIC development timelines are now spanning years, and may even extend beyond the window of opportunity for the intended product.

Harvard Business School Professor Clayton Christensen highlights that while price and performance are still important, there are signs that a seismic shift is taking place, leading to a new era where other factors – such as customization – matter more [4]. This is precisely where the FPGA fits in: it is the ultimate in customization.

FPGAs address the technical as well as business perspectives outlined above. Because they are off-the-shelf commodity items, companies can access state-of-the-art device technology with minimal NRE, and quickly build and deploy customized systems, achieving very short time-to-market while simultaneously maximizing first-to-market revenue streams. The Virtex-II Pro platform FPGA shown in Figure 2 is the cornerstone technology for building high-performance reconfigurable signal processing systems – which includes the PHY in a SDR. In conjunction with the logic fabric and active interconnect, this device has an array of embedded multipliers for supporting the most demanding of arithmetic tasks in a radio PHY. This particular FPGA family also offers integrated Power PC 405 technology, multi-Gigabit transceivers and dynamic impedance matching capability on the device I/O ports that can be used to simplify printed circuit board design and manufacturing. Using a platform-based approach to system implementation, system designers can create product differentiates by implementing signal processing functions in the logic fabric as well as through embedded software running on a Power PC.

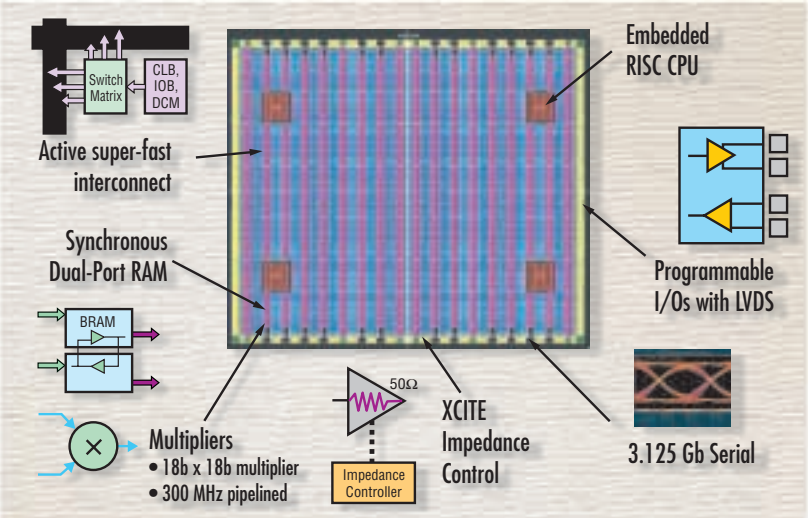

Figure 2. Virtex-II Pro platform FPGA showing the multiplier array for supporting parallel signal processing, multi-Gigabit transceivers for inter-chip and inter-system connectivity, and embedded RISC processor technology for performing decision-oriented tasks and running a real-time operating system.

The DSP Dilemma

One approach to BTS implementation has been to employ a combination of ASIC and ISA DSPs. The ASIC technology is typically used to address the significant arithmetic requirements of the radio front-end, such as digital down conversion and channelization filters to support multicarrier W-CDMA or CDMA2000 standards. These functions are beyond the capabilities of even state-of-the-art ISA DSPs. Even though the ASICs used in this part of the system may offer some programmability, it is generally limited in nature and is certainly a departure from the intended philosophy of the fully configurable soft radio. DSP processors might be used for certain baseband functions such as source (de-) encoding, as with CELP codec. Reduced Instruction Set Code (RISC) processing resources in the system could also support the requirements of the higher levels in the protocol stack.

From a soft radio perspective, the ASIC/processor combination is poor partitioning from both a flexibility and efficiency standpoint. In recent years FPGAs have experienced hyper-growth in both arithmetic complexity and compute density (number of operations/unit area of FPGA) that can be achieved by current generation devices. What types of signal processing functions can be usefully realized by an FPGA?

Radio designers working with FPGA technology implement infrared sampled receivers, channelizers of different varieties including classical digital down (and up) conversion (DDC and DUC) architectures, FFT-based polyphase transforms, multistage multirate polyphase decimators and interpolators, adaptive interference cancellers for DSSS channels, multiuser detection (MUD), and rake receivers (including acquisition and tracking). More recently, FPGAs have been used to construct space-time processors for advanced smart antenna systems. FPGAs are extremely adept and flexible at implementing FFTs, and this functionality has been used to construct OFDM modulators and demodulators.

FPGAs have also found extensive use in narrowband bandwidth-efficient Quadrature Amplitude Modulation (QAM) systems. In this environment they have been used to implement adaptive channel equalizers, digital timing recovery circuits, carrier recovery loops, frequency locked loops, and fractional

rate change filters.

FPGAs are also extensively used for forward error correction in communication systems. For example, OC-3 155 Mbps Viterbi decoders, Reed-Solomon decoding at OC-192 10 Gbps data rates, and (de-)interleavers operating at clock frequencies greater than 200 MHz are all achievable with current generation Virtex™-II FPGAs.

FPGAs are the ultimate device technology in terms of user customization. They allow system architects to perform area-performance tradeoffs and to therefore "right-size" the functional components in the system. FFTs with execution times in the microsecond to tens-of-microseconds are possible. In the context of an OFDM communication system, a small number of FPGA resources could be used to realize a (de-)modulator that supports a moderate data rate, or by using more resources an extremely high-performance high data-rate link could be realized.

With FPGA technology, control of the silicon is put back into the hands of the system developer rather than the chip architect – as is the case with an ISA signal processor. In fact, one way to view an FPGA is as a miniature silicon foundry with turnaround times of hours rather than months.

Leveraging these types of tradeoffs does not always mean that the engineering team has to construct the functional units from first principles. To facilitate rapid product development, many signal processing functions are available from the FPGA manufacturers themselves and from third party intellectual property (IP) suppliers. FFTs, multirate filters, Viterbi decoders, and Reed-Solomon encoders and decoders are all available as pre-verified IP from Xilinx (*www.xilinx.com/xlnx/xil_prodcat_landingpage.jsp?title=Xilinx+DSP*).

One of the roadblocks to the widespread deployment of FPGA-based signal processing has been design methodology related. In the past, FPGA-DSP design has required signal processing and communication engineers to use tool flows and languages with which they are typically unfamiliar. The introduction of tools like System Generator for DSP (*www.xilinx.com/xlnx/xil_prodcat_product.jsp?title=system_generator*) has gone a long way to let engineers work in the language of the problem. In this case the system is developed using a visual dataflow paradigm in The Mathworks Simulink environment. The approach not only allows the design to be specified, simulated, and parameterized, but it also enables design reuse through the use of IP cores.

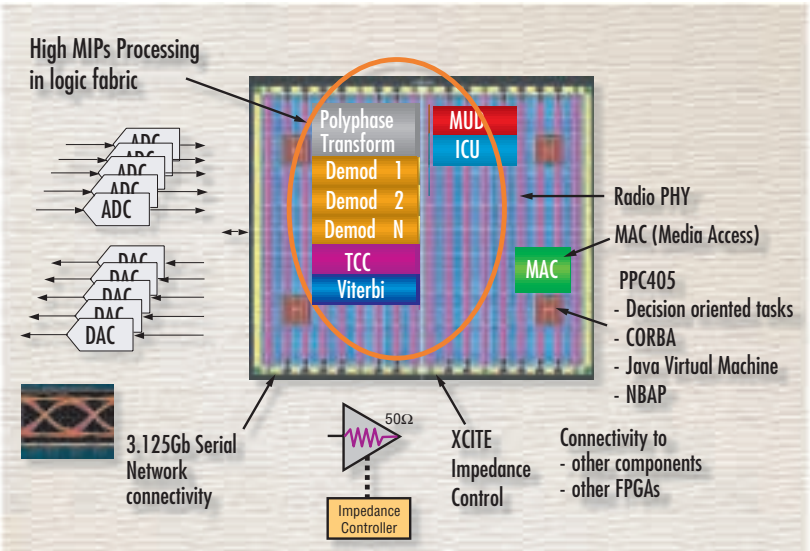

Figure 3. Platform FPGA approach to software-defined radio realization. The high MIPs processing is implemented in the logic fabric, while decision-oriented and non real-time tasks are provided as embedded software running on the Power PC. The multi-Gigabit transceivers could be used for providing connectivity to the broader network.

The Reconstruction of the Software Radio

The platform FPGA provides an opportunity for the radio architect to re-invent the system. Instead of having a radio card that is responsible for the DSP heavy lifting at the front-end of a soft radio system, and then passing this partially processed data over a VME or PCI-X bus to a baseband processor, multiple functions could be integrated into one or a small number of platform FPGAs.

As shown in Figure 3, compute-intensive tasks in the radio PHY could be implemented in the FPGA logic fabric, while more decision- and control-oriented tasks are run as embedded software on the Power PC. This embedded processor could even be used to run a Node B application protocol (NBAP) for a BTS, as a java virtual machine, or even provide CORBA support.

Advances in analog-to-digital converter technology are still required to support the high dynamic range requirements of wideband radio front-ends that offer true multimode global operability. Advances are also required in the area of configurable high-bandwidth analog signal processing for realizing the RF and IF stages of a radio. Micro-electro-mechanical systems (MEMs) appear to be a promising technology for addressing in-system configurable analog signal processing. As this technology matures and is combined on a single platform with digital functionality, the ideal of a completely configurable radio will move closer to reality.

The significant computation demands of the SDR PHY have been largely satisfied by highly parallel signal processing platforms realized using recent generation FPGA technology from companies like Xilinx. To complement the device technology, an increased emphasis on Signal Processing IP libraries and design methodologies such as System Generator for DSP are taking on renewed roles to provide a solution to the challenges presented by the software radio.

(Note: A useful primer for engineers and executives interested in developing products in the SDR application space can be found on the SDRF's web site, *www.sdrforum.org/sdr_primer.html.*)

References

[1] Mitola, J., III. 2000. *Software Radio Architecture: Object Oriented Approaches to Wireless Systems Engineering.* [Insert Location Here]: John Wiley and Sons.

[2] Allan, A., D. Edenfeld, W. H. Joyner Jr., A.B. Kahng, M. Rodgers, and Y. Zorian. 2002. 2001 Technology Roadmap for Semiconductors. IEEE *Computer Magazine* 35 (1):42-53.

[3] Verlinden, M.C., S.M. King, and C.M. Christensen. 2002. Seeing Beyond Moore's Law. *Semiconductor International* 25 (8):50-60.

[4] Bass, M.J., and C.M. Christensen. The Future of the Microprocessor. 2002. *IEEE Spectrum Magazine.* [Insert Volume and Issue Number Here]:34-39.