

SDR COMES OF AGE: TECHNOLOGY MEETS ECONOMICS

Manuel Uhm (Xilinx, Inc., San Jose, CA, United States; manuelu@xilinx.com); David Hawke (Xilinx, Inc., San Jose, CA, United States; dhawke@xilinx.com); Mark Quartermain (Xilinx, Inc., San Jose, CA, United States; markq@xilinx.com)

Software-defined radio (SDR) technology has “crossed the chasm” in many market segments, including tactical radios, satellite modems and commercial wireless infrastructure. However, in order to be truly successful and reach mainstream, technology has to not only initially meet the requirements of the visionaries, but also meet the requirements of the pragmatists and conservatives – both technical as well as economic. Technologies that fail to become mainstream often do so because they fail to meet the economic requirements.

SDR is a technology that has now become mainstream in many markets. By way of example, this article examines the commercial wireless infrastructure market to illustrate how SDR is meeting both the technical and economic requirements of telecommunications equipment manufacturers (TEMs) through cost reduced architectures that not only lower total cost of ownership, but also BOM (bill of materials) cost.

1. SDR TECHNOLOGY

The IEEE and SDR Forum define an SDR as “a radio in which some or all of the physical layer functions are software-defined.” Clearly, this is a broad definition. It is important to recognize that by this definition an SDR does not require a wideband radio frequency (RF) front end that is frequency agile. In fact, the definition is restricted solely to layer 1 of baseband processing. Stated another way, while not precluded, a wideband RF front end is not necessary for a radio to be defined as an SDR. Understanding this definition is critical to the context of this SDR discussion.

In addition, it is also necessary to have an accurate understanding of the SDR technology adoption curve. Many people in high technology companies are familiar with Geoffrey Moore’s technology adoption curve as described in his book “Crossing the Chasm”. The concept is illustrated below highlighting key market segments and where they reside on the SDR technology adoption curve. This shows that SDR adoption in MILCOM is the most advanced to the point where even the laggards and skeptics are adopting SDR. Mobile handsets and terminals adoption

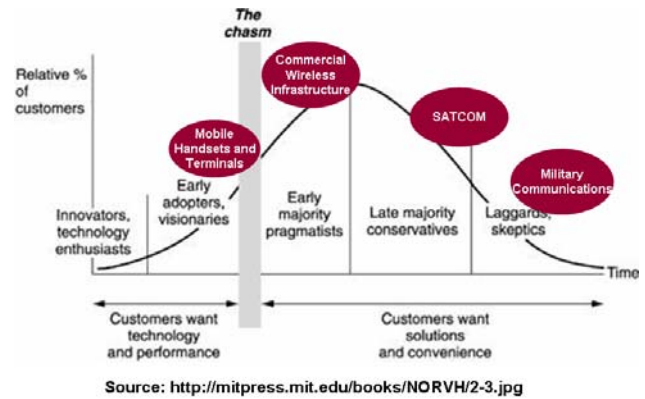


Figure 1: Market segments acceptance of SDR on the technology adoption curve

of SDR is the least advanced where some smartphones are adopting reprogrammable baseband processors but it’s not clear yet if SDR technology will cross the chasm and become a de facto standard in mobile handsets.

In order for a technology to cross the chasm, it must appeal to pragmatists that have both technical and economic/business criteria to satisfy. There are many examples of technology that have appealed to visionaries on their innovation and technical merits, but have failed to cross the chasm due to economic or business issues such as a lack of manufacturability, cost effectiveness, quality or profitable business models.

Essentially, all current generation radios (exclusive of legacy radios which are still in production) targeting military communications (MILCOM) and SATCOM applications use reprogrammable reconfigurable processors. As such, the MILCOM and SATCOM markets are in the late majority phase and have clearly crossed the chasm in adopting SDR technologies.

The rate of adoption of SDR technology in commercial wireless infrastructure is also most worthy of note. In the last couple of years, there have been public announcements of SDR basestations from major TEMs, including NSN, ALU, Huawei, ZTE, as well as smaller players such as Vanu. In this domain, SDR is rapidly becoming the de facto

standard for baseband processing in basestations that can be marketed as SDRs, multi-mode basestations, or common platforms. SDR is the fundamental enabling technology that makes it possible for these basestations to support multiple air interfaces and over-the-air upgrades as standards continue to evolve.

What may be surprising to some is that SDR technology is also gaining traction in the mobile handset market. The most prominent example is Apple’s 3G iPhone, which uses an Infineon baseband processor. The Infineon baseband processor includes a reprogrammable digital signal processor (DSP) for baseband processing and a general purpose processor for higher layers of the protocol stack. While SDR technology is making inroads in this market space, however, it has not yet crossed the chasm into mainstream.

2. COMMERCIAL WIRELESS INFRASTRUCTURE: BEYOND THE CHASM

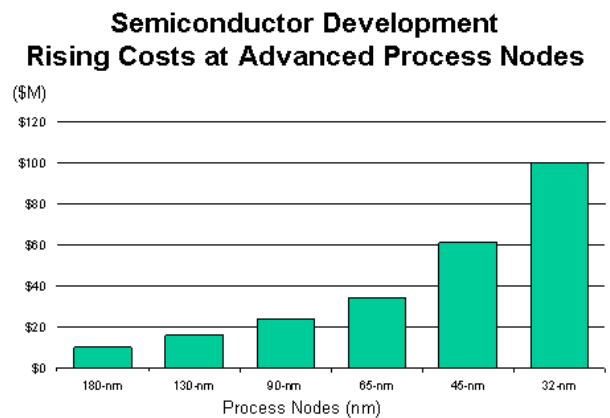
An issue that has traditionally hampered the adoption of SDR technology in commercial domains has been the BOM cost premium associated with the flexibility of SDRs over and above the cost of a fixed solution. Hence, much of the previous argument for the economic viability of SDRs has centered around total cost of ownership, as opposed to BOM cost. The logic is that once you factor in the costs of development (which could include multiple ASIC spins, multiple board designs, etc.), as well as logistics costs (such as higher inventory and supply chain costs), SDRs are lower cost than developing multiple fixed hardware platforms. While there is significant merit to this argument, the BOM cost premium has been a barrier to adoption of SDR technology in commercial wireless for at least a couple reasons:

- 1) Many TEMs lacked visibility into the cost structure around total cost of ownership, whereas BOM costs were highly visible.
- 2) Many of the pragmatists are measured on their ability to develop the most BOM cost effective basestations,

So then, why has SDR adoption in commercial wireless infrastructure now crossed the chasm to become mainstream? Let’s start with the first issue. There is no doubt that TEMs have much better visibility into total cost of ownership now than even just three to five years ago. However, the primary forcing function has been the harsh economic realities that basestation hardware has become a commodity. This is evidenced by the expectation of 10 to 15

percent lower prices year over year, and the significant reduction in engineering resources that has occurred as TEMs struggle to control costs and increase profitability. The combination of these factors has forced many TEMs to move towards a “common platform” engineering approach, whereby a single hardware platform is developed to support multiple operators and air interfaces. These common platforms are fundamentally software-defined radios that enable TEMs to address the diversity of requirements from multiple operators with different software loads even in the face of limited engineering resources.

With regard to the second issue, reconfigurable and reprogrammable processing technologies, the key enablers for SDR, are now able to compete on BOM cost for the volumes driven by the commercial wireless infrastructure market due to the increasingly prohibitive mask and development costs of ASSPs and ASICs. In fact, the business case for development of ASSPs and ASICs for wireless infrastructure is steadily declining over time. As illustrated in Figure 2, the costs of semiconductor development from 90 nanometer (nm) to 32nm are expected to increase 170 percent during the same timeframe that the wireless infrastructure market is projected to grow at a CAGR of only 3 percent. As a result, ASSPs and ASICs are both costly and risky given that the cost of servicing the market is growing orders of magnitude faster than the industry itself. By contrast, reconfigurable and reprogrammable processors are still able to aggressively move to successive process node geometries to drive down BOM cost, while addressing the wireless infrastructure industry’s needs, by amortizing the mask and development costs across multiple industries.



Source: Xilinx

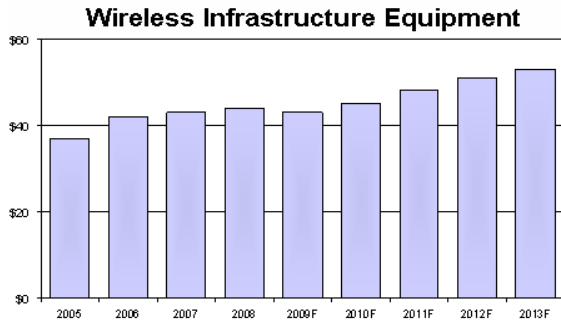


Figure 2: Rising semiconductor development costs juxtaposed against flat wireless infrastructure growth. Source: GSA and iSuppli.

These points are further illustrated with the following two examples. The first example is a digital radio subsystem. It comprises functionality such as digital up/down conversion, crest factor reduction, digital pre-distortion and CPRI, which would be found in a radio card or remote radio head. Figure 3 illustrates an ASSP implementation for a radio supporting 2x2 LTE (Long Term Evolution). It is a seven-chip digital solution with estimated power consumption at 10.6W and a list price of \$316.

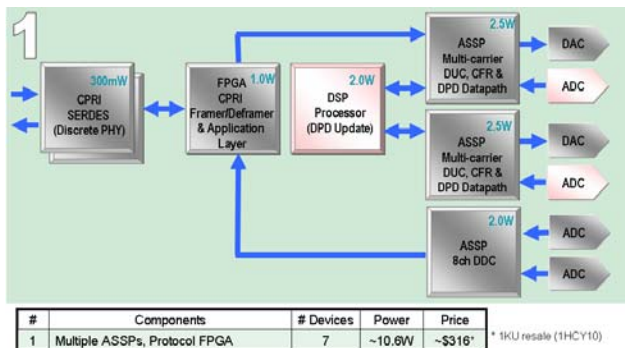


Figure 3: LTE 2x2 digital radio ASSP implementation

Figure 4 illustrates the same functionality implemented in a single Xilinx® Virtex®-6 FPGA that offers significant advantages over the ASSP implementation. This reprogrammable approach harnesses the true power of SDR with a cost competitive BOM and much lower total cost of ownership:

- 1) Roughly equivalent BOM cost
- 2) 45% reduction in power consumption
- 3) Significantly smaller footprint
- 4) Upgradeable to support evolutions in the LTE standard and updates with new features or standards.

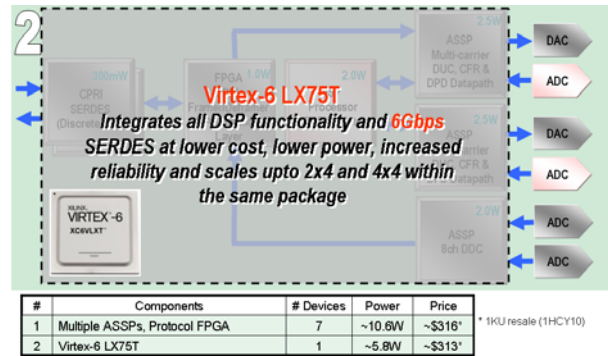


Figure 4: LTE 2x2 digital radio FPGA implementation

The second example is a baseband processing system. The complexity of baseband processing and the need for TEMs to differentiate themselves with proprietary algorithms or implementations does not bode well for fixed implementations. Indeed, there are no simple ASSP solutions, whereby the semiconductor vendor provides a complete solution that enables TEMs to easily integrate devices on board. Rather, typical solutions today are comprised of reprogrammable DSPs with some hardened functionality, such as turbo decoders, or FPGAs, ASICs, or some combination. In most scenarios, these implementations are still classified as SDRs as some or all of the physical layer processing is defined in software.

It is worth noting that the same business case issue described earlier facing ASSPs and ASICs also applies to DSPs that are dedicated to the wireless infrastructure industry, such as the C6487 from Texas Instruments. While these devices use general purpose DSP cores, they are supplemented with hardened co-processors, like turbo decoders, that are specifically for wireless signal processing and hence are not really suitable to be applied to other markets, such as medical image processing. Hence, are they essentially reprogrammable ASSPs, but they still enable SDRs.

Consider the FPGA and DSP ASSP implementation of a 4G baseband channel card supporting three sectors of LTE as shown in Figure 5. This architecture is a functional-based design in which functions for all three sectors are partitioned across several devices. This is an SDR by definition, but has some fundamental technology limitations which restrict its capabilities.

First, there is a bottleneck between the FPGA and DSP ASSPs. For lower data rates, such as those from a 5 MHz bandwidth LTE implementation without MIMO, this may not be an issue. However, for higher data rates, such as a 10-

20 MHz bandwidth LTE implementation with 2x2 or 4x4 MIMO, this is likely to result in lost data.

Secondly, the hardened co-processors in the DSP ASSPs are fundamentally constrained by the data rates and air interfaces supported by their fixed architecture implementation. So, while this is an SDR implementation, its flexibility is limited. The architecture may not be capable of accommodating new air interfaces and higher data rates. Plus, the larger footprint of the design and the higher power consumption caused by the additional I/O required to move data between the seven discrete devices are major drawbacks.

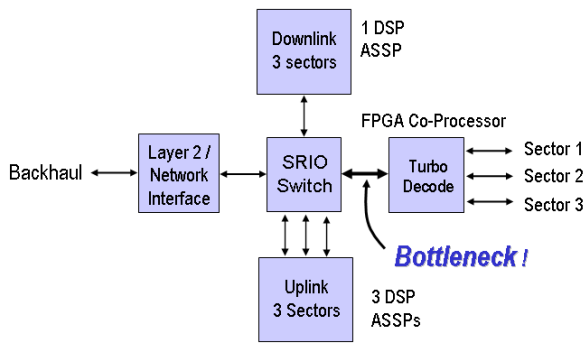


Figure 5: DSP ASSP Centric 4G SDR Channel Card Implementation

Figure 6 shows an alternative baseband implementation for LTE. This example uses a sector-based partitioning where all the functionality for a given sector is handled by a specific device. This is a much more flexible SDR approach than the fixed architecture implementation. There is no bottleneck because all the sector processing is done in a single device, rather than across multiple devices. The use of reconfigurable hardware also means that the highly computationally intensive portions (which required co-processors in the fixed implementation) do not need to be hardened and can be continuously evolved and optimized over time. Depending on the features of the PHY and how they are implemented, this approach can also result in lower BOM cost, lower power consumption, and smaller footprint.

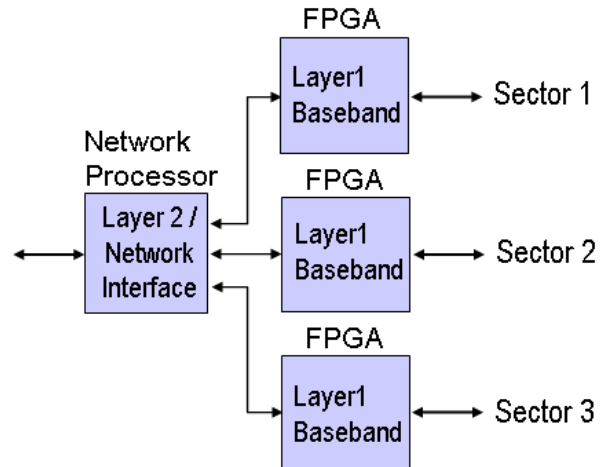


Figure 6: FPGA Centric 4G SDR Channel Card Implementation

3. SDR: TECHNOLOGY FOR PRAGMATISTS

SDR has clearly crossed the chasm of the technology adoption curve for commercial wireless infrastructure. It has taken more than a decade for the technology to achieve this milestone. While the promise has been a long time in coming to fruition, it has finally been achieved by overcoming the economic barriers to adoption. This has allowed pragmatists to embrace and fully utilize SDR technology to solve several key market problems in commercial wireless infrastructure.