Smart Partitioning in WiMAX Radios

By Noman Rangwala, Marketing Manager [noman.rangwala@analog.com]
and Tom Gratzek, Product Line Director [tom.gratzek@analog.com]

The digital revolution has changed the way we communicate, work, and travel by reshaping our relationship with the world around us. The digitization of electronics has transformed our world by enabling a vast network of portable, accessible, interconnected communications media. However, the promised advantages of digital technology are only as good as the ability of the analog technologies to faithfully translate the digital language of 1s and 0s into natural analog signals.

The advance of the digital revolution has been characterized by Moore's law—which states that the number of transistors on a chip doubles every 18 months. Analog technologies, on the other hand, are characterized by Murphy's Law—if anything can go wrong, it will. Analog technologies progress at a more measured pace dictated not by process enhancements but by innovations in circuits and physical transistor modeling. These technologies improve incrementally on multiple dimensions of performance, power, and integration.

Integration Trends and the Case for Partitioning

Integration trends are a function of volume and system maturity; in many cases system acceptance and unit volume production never grow to justify recurring generational development. In other applications, such as base stations, instrumentation, and military applications, stringent performance requirements lead to discrete implementations. In cases such as cellular and Wi-Fi, where consumer acceptance is universal, competitive forces drive the continual cost reduction. As technology becomes more expensive to deploy (such as mask, tool, and engineering costs), the return needed to justify these developments increases. At the same time, competitive forces drive companies to invest heavily early in a standard's life cycle. If a market takes off, and a company's chipset is not ready, the financial result can be dire.

In essence, companies are forced to invest to be ready when a market takes off, and this investment is increasingly expensive, while at the same time, customers are requiring more performance from their suppliers. Obtaining an acceptable return on the R&D investments required to build today's complex communication systems is a very tricky proposition. Depending on the complexity of the SOC—development costs can easily range from $10 to $20 million, and higher, for a 90 nm design. Thus, success of a new initiative depends on identifying a market where your IP is valuable and then lining up partners to meet customer needs. Fewer and fewer companies are able to handle all aspects of a system development. However, focus on performance cost, TTM, and financial payback is an absolute requirement.

For emerging communications applications like WiMAX, the first generation systems have typically been developed using multiple
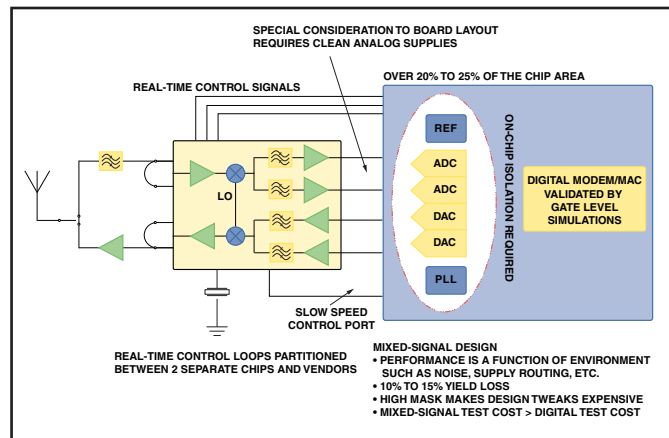

Figure 1. Traditional partitioning.

ICs. The MAC/modem section may use FPGAs and off-the-shelf DSPs; the RF sections often use discrete components such as LNAs, mixers, and synthesizers, with the ADCs and DACs bridging the gap. As volumes grow, the digital logic is often integrated together on a dedicated ASIC and, in some cases, the ADCs/DACs are included on this digital ASIC, for use with more integrated RF solutions. For other applications with size constraints, such as mobile phones or USB dongles, the analog and digital functionality can be integrated together, either in one system in a package using multichip modules, or on a single chip. There are many different ways to drive to lower size and cost, but the trend is that as volumes increase, size and cost decline. In some cases, cost is king and RF performance can be sacrificed (i.e., some WLAN consumer applications), although customers don't realize it. In other cases, size is paramount, and integration of functionality is the driver.

There is no one recipe for success. Companies have been successful with many different integration and cost reduction strategies. To be clear, development choices must be made that minimize electronic bill of materials (eBOM), size, and TTM. Intelligent design of system partitioning is instrumental in achieving success.

Traditional Partitioning—a Time to Market Risk

The integration of mixed-signal circuits on a digital ASIC opens doors to many implementation challenges and hence introduces a time-to-market, and more importantly, time-to-revenue risk to the product. Even though the mixed core has been verified on a standalone basis, the performance of the core is a function of the environment in which it is integrated. Issues of power supply routing, parasitic capacitances, and process variations that are not important for a digital-only chip, now have a greater significance.
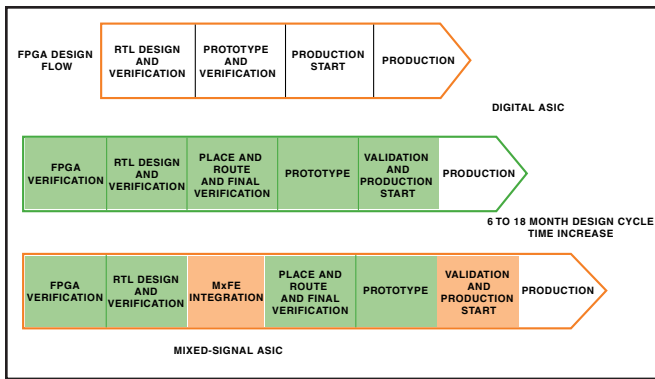
analog is everywhere.™

ANALOG DEVICES

Figure 2. Design cycle time.

The time from an FPGA-validated, digital-only design to silicon ranges from two to six months based on complexity, design flow, and automation tools. On the other hand, the cycle time to get a mixed-signal design to first silicon could take up to three times as long—assuming that the analog cores are available and verified in the appropriate process of choice. The sensitivity of analog circuitry to noise generated by the switching of millions of transistors in the presence of signals in the range of microvolts requires greater attention and multiple design and layout reviews, thereby increasing the time to silicon and working samples.

The problem is not insurmountable. Multiple techniques are available to mitigate the interaction, but these require careful attention to custom layout of the mask, which takes engineering time and resources. It certainly requires an entirely new set of core competencies in what may already be an overloaded engineering team.

The evaluation board design and layout also has a critical impact on the performance of the mixed-signal portion of the device. The analog I/O on the reference board is sensitive to external noise, and the supply routes to the mixed-signal portion of the design require high isolation. Eliminating analog I/O reduces the noise coupling issues to a minimum. In addition, it solves the problem of interfacing analog cores from different vendors (i.e., RF chip and mixed-signal converter cores). For example, some of the available ADC cores recommend that, to obtain data sheet specified performance, a discrete 5 V op amp driver buffer is required. For modems using a smaller process, such as 130 nm or 90 nm, the signal swing and common-mode level must be reduced and matched when using different vendor RF chips. These additional considerations require valuable engineering resources.

Being second to market often means steeply discounting product pricing in order to capture market share. Choosing a pure digital or an FPGA design flow can shorten the time to bring a product to volume manufacturing by six to 12 months.

Getting to functional silicon is only the first step—getting to production with a mixed-signal IC offers its own challenges. Mixed-signal circuits are sensitive to process variations such as thresholds, leakage, resistance of material, and other process parameters. Often, as the performance of the mixed-signal degrades, so does the system.

In high volume markets, the ability to manufacture at multiple fabrication sites is essential to ensure timely delivery and optimize costs. Digital design can be relatively fabrication-site agnostic while porting mixed-signal circuits to different fabs is time consuming and can require extensive redesign and optimization skills. The resources for targeting different manufacturing flows are usually very difficult to put together, and often better spent elsewhere.

Another important issue with traditional partitioning is that it requires a matched pair approach. That is, since the ADCs and DACs are separated from the RF, the real-time loops, such as automatic gain control and transmit power control, are forced to be shared between two chips and multiple parties. Significant up-front work is required to optimize a reference design from discrete devices.

These challenges of analog and mixed-signal design lessen the focus from the core competency of the system level design team and can delay the introduction of new products to market.

Smart Partitioning

With the availability of mature RF CMOS processes and advances in analog and RF modeling capabilities, it is now possible to move the data converters and other mixed-signal blocks to the RF IC. The next section will show why replacing the traditional analog baseband interface with a digital interface offers a "smarter" system partitioning for some communications systems.

The proposed change includes the appropriate partitioning of functionality such that the RF system on a chip (SOC) provides a complete RF to bits solution, which includes all the required control loops such as automatic gain control, transmit power control, and RF calibration loops. The inclusion of control loops on the radio front end results in ease of use and easier mix and match capability with different digital baseband PHY modems. A standard format, the ADI/Q™ digital I/Q interface, is available for the interface between the RF front end and the digital baseband. This interface format consists of bidirectional control and data lines and it supports interchangeability and ease of application. The reduction of real-time software control results in simpler system design. All the analog and RF specific controls are partitioned to the RF front end.
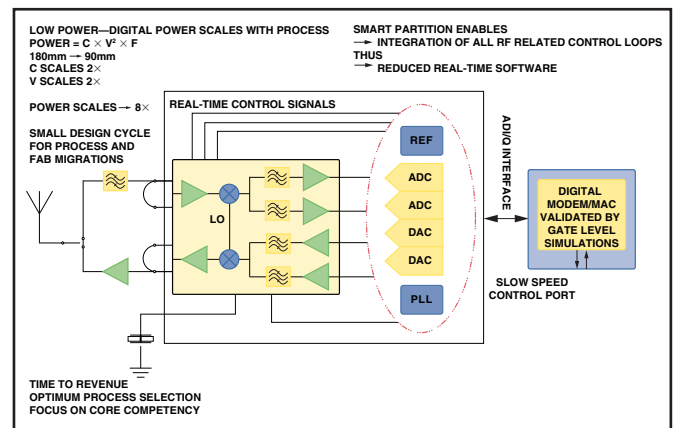

Figure 3. Smart partitioning.

Low Unit Cost and Lower Development Cost

Market segments which are characterized by high demand and production volumes attract more market entrants. To be successful in defending a lead and increasing market share, the solution providers need to pay attention to full factory cost for the chipset. Smart partitioning can offer significant device cost reduction.

For communication systems, such as WiMAX and broadband wireless access, consumer price points less than $100 are essential. CPE equipment for ADSL and 802.11g Wi-Fi ($20 to $30) are examples of where volumes increased dramatically as prices declined. An emerging market such as WiMAX will also experience similar price pressures. It is expected that the end user CPE prices will be under $100 by mid-2007. To achieve these targets, the chipset pricing will be required to fall into the range of $20 to $25. This is probably much lower than the current costs, and will require quantum improvements so that market prices yield an acceptable profit.

RF to bits radio ICs can help enable this transition.

For a given process, mixed-signal ASIC design is more expensive than a digital-only ASIC design, with the increased cost adders having four main components:

1. For a particular process, mixed-signal devices are inherently more expensive. The mixed-signal features require additional processing steps such as thicker oxides, low threshold devices, and additional implants. In general, mixed-signal wafer costs can be 20% higher than the digital-only wafer.

2. The fabrication plants invest heavily in the reduction of defect density, resulting in high yields, close to 97% to 98%, depending on die size. On the other hand, analog circuit IC yield is a function of the design itself. To achieve specified performance while making power dissipation trade-offs, analog circuits are designed to perform to specifications over a narrow window of process variations when compared to digital design, resulting in parametric-limited yield, thus increasing the costs for mixed-signal designs. This adds over a 10% increase in costs for mixed-signal designs.

3. The elimination of analog functions from the digital modem results in simplification of production test development and is instrumental in reducing production test time. Enabling test on a generic digital tester rather than an expensive mixed-signal tester can reduce tester cost by 15% to 20%.

Test coverage tools allow a digital designer to create fault coverage scan chains, simplifying production test. Whereas mixed-signal testing requires measuring various analog specifications in the range of a few microvolts. A mixed-signal test design could take at least five times longer than a digital-only test. The time can be reduced using parallel processing on the testers. Assuming an aggressive test program methodology—the test cost for mixed-signal devices can be in the range of two to three times greater.

4. The integrated converter core is usually intellectual property that is developed by a third party and/or an internal group with associated royalties, and/or NRE. The design and support tools used in a mixed-signal design flow are an added investment when compared with a design toolkit for a digital-only ASIC solution. A suite of tools
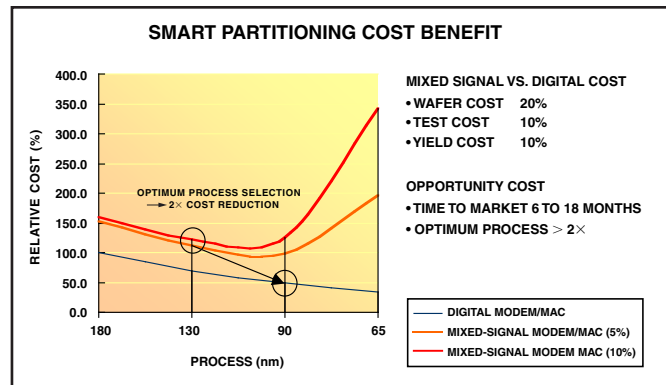

Figure 4. Cost benefit of smart partitioning.

required to design a new mixed-signal ASIC when compared to a digital-only ASIC can easily exceed $500k.

Additionally, analog circuits do not scale with process shrinks in the same way digital circuits do. Figure 4 illustrates the rising costs of mixed-signal ICs as a function of feature size. The cost curves are normalized to the cost of a digital-only ASIC in 180 nm. Historically, the digital ASIC cost tends to reduce by a third when migrating from one feature size to the next. In contrast, the mixed-signal IC cost increases as a function of the percentage of mixed-signal die area. This comes from the fact that the noise limited analog circuitry does not scale with lithography, while the digital circuitry tends to scale as quadratically with process.

New processing equipment investments and the increased complexity of the manufacturing process result in a net increase in the die cost per sq. mm from one generation to the next. The digital circuitry scales proportionally to result in a lower cost per transistor. Since analog circuits do not scale with process, the total mixed-signal product cost tends to remain flat initially and increases with subsequent process shrinks.

In high volume markets, companies must remain cost competitive while meeting market pricing and providing a fair return to investors. If a company's cost structure is double the best-in-class competitors, new tactics or new strategies will soon become necessary. Although all the challenges associated with mixed-signal design continue to exist, the benefits of smart partitioning include dramatically lowering the systems cost by taking full advantage of Moore's law—not always available to analog/RF circuits.

In addition to the increased cost per device, the opportunity cost of not selecting an optimum process and longer time to market can doom the financial return on a project. The availability of ready-to-use analog and mixed-signal cores lag behind the digital process by approximately two years, or about one generation. With the availability of production ready cores being close to four years out, the smart partitioning approach enables the system vendors to choose an optimum process based on their needs and not be constrained by availability of a validated analog core. The opportunity cost associated with the selection of a nonoptimum process is high. For example, in the broadband wireless space, manufacturers have announced a 90 nm core design. The difference in product cost between a 90 nm

digital SOC design and 130 nm can be greater than 200 percent! At 65 nm, the multiplier can be even higher.

The proposed change offers an opportunity to use the additional time and resources to focus on developing the next generation product—potentially putting it one product generation ahead of competitors who are spending valuable resources fighting issues inherent in a mixed-signal ASIC design.

Performance Advantages from the Shift to a Digital Radio Baseband Interface

Along with the cost advantage in development, support, and per unit cost, smart partitioning enables a high performing system solution.

For advanced OFDM systems with high peak-to-average ratio, the high linearity achieved on the RF device, as well as the advanced synchronization and channel estimation algorithms on the DBB, must not be compromised by the dynamic range of the ADCs and DACs. Careful management of the headroom must be considered to enable robust performance in the presence of noise, fading channels, and interferers.

With the integration of an autonomous AGC loop, the dynamic range of the ADCs can be matched with the capability of the RF front end, thus enabling high data rates such as 64 QAM. There are many vendors that have struggled with bringing up their reference designs because of the complex interactions between the DBB and the RF IC. In addition, advanced techniques, such as symbol-to-symbol AGC, can be utilized to improve the performance of the system in fading channels which are common in mobile environments. Unlike a distributed AGC (i.e., AGC algorithm implemented on two separate devices), the proposed partitioning enables a fast convergence of the AGC, thus allowing the DBB to spend more time on channel estimation and synchronization, thus improving the system performance by many decibels, which translates into greater range and rate.

Filtering is required to eliminate undesired signals from adjacent or alternate channels. To address this issue, careful trade-offs must be made between linearity and filtering complexity. For low cost ZIF architectures, the final channel selectivity is performed by using digital filters. Filtering like gain must be distributed between the RF and subsequent digital filters. Smart partitioning enables the optimization of the filtering requirements between the analog and digital filtering, utilizing the converter dynamic range to the maximum.

Power dissipation is also an important parameter for mobile systems. Power dissipated on a digital chip is directly proportional to the square of the supply voltage and directly proportional to the gate capacitance. Thus, for a process migration from 130 nm to 90 nm, the result could be a power savings of $8\times$. With a smart partitioning philosophy, the DBB, when implemented in 0.13 μm, dissipating in the range of 1 W to 1.5 W, can be reduced aggressively down to 200 mW, when moved to a 90 nm process.

Summary

The digital revolution has resulted in solutions with millions of gates put together on fine line processes. These SOC solutions are expensive to develop and put tremendous pressure on return on investment. To succeed, one must choose the appropriate market segment, apply focus on a core competency to deliver a differentiated product at low cost in a timely manner. Partnering to minimize risk and executing to a schedule is an attractive option.

Partitioning with an "RF to bits" radio offers the four key ingredients for success—high performance solution, focus on core competency, lowest power cost, and fastest time to market.

The appropriate partitioning of analog and digital functionality solves many of the issues related to integration of analog circuits on digital ASICs and results in faster time to market and longer time-in-market. It enables the optimization of the system to achieve high performance.

For digital baseband vendors, with expertise in digital modems and media access controllers, smart partitioning offers the advantage of focusing critical resources on tasks and projects that further enhance their value proposition.

In high volume applications, the choice of process is critical. The ability to migrate to newer processes quickly results in new cost and performance points which will provide competitive advantage. The smart partitioning philosophy is being adopted by multiple standards bodies such as the Digi-RF group in mobile handsets, the JC-61 group targeting WLAN and WiMAX, as well as in various proprietary systems. Analog Devices offers the ADI/Q interface which allows easy implementation of this cost- and performance-optimized strategy.

www.analog.com

ANALOG
DEVICES