

RF: Will it ever be plug-in IP?

AS SOCs FOR MOBILE DEVICES INTEGRATE RADIO CIRCUITS, THEY WILL NEED TO REUSE RF IP. BUT WILL THEY BE ABLE TO?

BY RON WILSON • EXECUTIVE EDITOR

For the next generation of smart phones, mobile video players, and roaming Web accessories, integration means not just placing multipurpose basebands onto the same SOC (system on chip) with the application processor, accelerators, and memory, but also integrating the small-signal RF circuits of many radios onto the SOC. And therein lies a challenge.

In today's SOCs, reusing silicon IP (intellectual property) is absolutely necessary to meet schedules and maintain reasonable design complexity. But, although IP reuse is well-understood in the digital world and has even become common for some very-high-frequency blocks, such as SERDES (serializer-deserializer)

functions for high-speed I/O and PLLs (phase-locked loops) in all sorts of applications, design reuse is virtually unknown for RF circuits in radio applications. Why? And can you do anything about it? Or will the next generation of SOCs for mobile devices have to be largely custom chips?

SIZING THE QUESTION

At stake in these questions is a large fraction of the die area in the mobile-system SOC. This situation occurs not simply because radios tend to be large—with larger-than-minimum transistors and virtually unshrinkable inductors and capacitors—but also because the advanced handset will require lots of radios on one SOC.

For example, consider a typical concept device that acts as a multiband phone and Internet terminal, reaching out to whatever best connection is available at the time. Such a handset would have an LTE (long-term-evolution) radio for connecting to the cellular-phone network; an 802.11n MIMO

(multiple-input/multiple-output) Wi-Fi-radio bank for making multiple-antenna connections to any available Wi-Fi access point; at least one UWB (ultrawideband) radio to implement Bluetooth, wireless-USB, or proprietary short-range, high-bandwidth protocols; and a GPS (global-positioning-system) receiver for obtaining location information. None of these components is trivial, and, in the case of the 802.11n MIMO radio, the area is substantial (Figure 1).

Initially, these radios may all be on separate dice. But design pressures will force many, if not all, of them to migrate onto the SOC. The exception will be the large-signal RF blocks, such as power amplifiers and antenna switches, which are likely to remain off the SOC.

The small-signal RF blocks in a typical radio are not hard to enumerate (Figure 2). They recur in about the same configuration in most modern radios. But that situation does not mean that it's easy to reuse either a whole radio block or the individual functions.



With so much complexity, so much die area, and so much at stake, it would seem obvious that SOC designers—particularly SOC teams that lack internal RF expertise—would want to obtain their radio blocks as third-party IP. But the reality is that, today, they will not. And even in the moderate-term future, some or most of the radio blocks in these SOC designs will be the work of the SOC-design team itself, not a third-party-IP developer. There are several important reasons for this predicament.

All of these reasons follow from the concept of reusable IP. To be reusable, an IP block must have a clearly understood function, upon which the creator and the user agree. It must have clearly defined ports, and the signals at these ports must be unambiguously defined both as to their relationship to the function of the block and as to their allowable signal characteristics. Without these basic constraints, IP is not so much reusable as redesignable. But, as you will see, each of these requirements is problematic for RF-radio IP.

The first problem is the matter of standards. Certainly, there are clear air-interface standards, with verification IP and interoperability tests, for virtually all nonproprietary air interfaces. But the problem is the implementation of the radios. “The transceiver architecture for these standard radios is still evolving,” warns Berkeley Design Automation Chief Executive Officer Ravi Subramanian. “How you want to build the radio still depends on the market you are going to serve with it.” For example, a

AT A GLANCE

- ▶ The next generation of mobile wireless devices will depend on integration of some radio functions into the SOC (system on chip).
- ▶ SOC development depends on IP (intellectual-property) reusability.
- ▶ RF circuitry is notoriously hard to reuse for some very compelling reasons.
- ▶ Reuse of third-party-radio blocks is likely to be difficult and slow to emerge.

UWB radio, although structurally much like an 802.11 radio, can look different in detail (Figure 3).

“Even standard radios are less well-defined than you might think,” says Navraj Nandra, mixed-signal-product-marketing director at Synopsys. For instance, he notes, different countries might implement the same standard radio differently. WiMedia radios in the United States use Band Group 1. Other countries use higher-frequency band groups requiring different radio designs.

“And, putting these radios into CMOS on an SOC is the crown jewel for the companies that can do it,” Subramanian says. “The ability to integrate the radios with high performance is essential to their products, and it differentiates them. They are not going to buy the radios as IP, even if they could.”

But what about licensing the sub-blocks within the radios? After all, you can license third-party PLLs that run as fast as some of these blocks. That sce-

nario is not in the cards either, Nandra says. “There are clear functional blocks in a radio: the RF front end, the mixers, the data converters, and so on. But the problem here is that the exact partitioning into blocks, and the interfaces between them, is not defined. For example, in the PCI [peripheral-component-interconnect] Express world, there is an industry-standard pipe interface between the PHY [physical] and the MAC [media-access-control] layers. But the interfaces between the functions in radio hardware are not defined.”

“These are very difficult interfaces,” says Carlos Leme, executive vice president of the wireless-systems group at MIPS. “It’s not a matter of just plugging them together. You have to observe all the loading and impedance-matching requirements of the RF signals between blocks.”

The functions within the blocks also are not clearly defined. “The partitioning of RF circuits is very complex,” Leme continues. “The specifications of the blocks interact with each other.” Leme explains that, for example, skilled RF designers may choose to accept more noise in one portion of a circuit and make up for it in another. “[For this reason] there has never been a significant market for RF-building-block IP,” Leme adds. “You always end up having to work closely with the design team; it’s not really IP then.”

SIGNALS AND NOISE

Even if a design team would accept previously defined building blocks, the integration process would prove difficult. The problem involves another component of the definition of reusable IP: The block should have well-defined signals on clearly defined pins. The difficulty with this concept for RF design is that interactions between an RF circuit and the rest of the chip don’t involve only defined signals or even intentional paths between the IP block and the rest of the chip. This situation can lead to interesting problems.

First, there is the connection problem. You can’t simply tell a digital-routing tool to connect the pins on a hard-IP block to their destinations and end up with a working radio. RF-signal paths must have impedance matching. They care about parasitics. And they care—




Figure 1 The Broadcom BCM2055 is a 4×4 MIMO radio supporting the 802.11n draft specification. Note that both complexity and die area are not insignificant.

sometimes a great deal—about everything that is going on at the node to which you attach them.

Even with impedance-matching connections to quiet, well-behaved nodes, a piece of hard IP validated in the process you are using may not behave in your chip as it behaved in the test chip. Often, this unpredictability is due to interactions between the RF circuitry and the rest of the die that don't involve the defined signal paths.

"Models are problematic in RF design," says Hany El Hak, senior product-marketing manager at Cadence. "It's not because the RF models you get from the foundry or the IP vendor are inaccurate, necessarily. It's because the assumptions the IP designers made while constructing the models don't always get communicated to the IP user."

He explains that, for instance, if the IP designer assumed a maximum supply-noise figure, you need to know it to verify that the supply pins on the IP in your design don't exceed that figure. "In general," he notes, "the problem is that, in the RF domain, there are couplings and interactions that don't follow the signal path."

Supply noise is just one example El Hak offers. Another is substrate coupling. Until recently, it was all but impossible to get accurate substrate models for even the best CMOS-logic processes. Now, those models exist, and foundries are happy to share them, El Hak reports. "But substrate-coupling models are very complex. If you include them in your circuit models, the complexity of the overall problem explodes. You must have some formal way of reducing the complexity of the model—removing parasitic paths to which the circuit isn't particularly sensitive—to make simulation feasible. There are tools to [accomplish this task], in Spectre, for example. But it's not fully automatic. The accuracy of the reduced model still depends on the designer's guidance in pruning the circuit."

VERIFICATION ISSUES

With all the potential for interactions between a radio-IP block and the substrate, the supply pins, the signal pins, and even nearby but unrelated traces, it should be no surprise that even experienced RF designers approach verification of the integrated block with a cer-

tain amount of respect. There is no easy way out.

"You really want to do verification of the whole system, not just the IP block," warns Irshad Rasheed, president and chief executive officer of RF-design-services house Tahoe RF Semiconductor. "Just defining the system from the top level can be 15 to 25% of the design cycle. Once you [define it], many design teams begin to analyze the system from a behavioral level with Verilog models and with enough extracted data to allow an analog/mixed-signal functional simulation." Moving directly to integration of the IP and producing GDS (graphic-data system)-II for the whole chip is possible but extremely risky, he cautions. "Models for substrate coupling and for the noise spurs from the digital circuitry are never that good. The VCOs [voltage-controlled oscillators] are never centered. The risks are very high."

Instead, Rasheed recommends, the design team can implement the radio circuitry on test chips. These components can start out with small structures simply to validate the circuit models and progress to an entire radio block surrounded by digital-noise generators

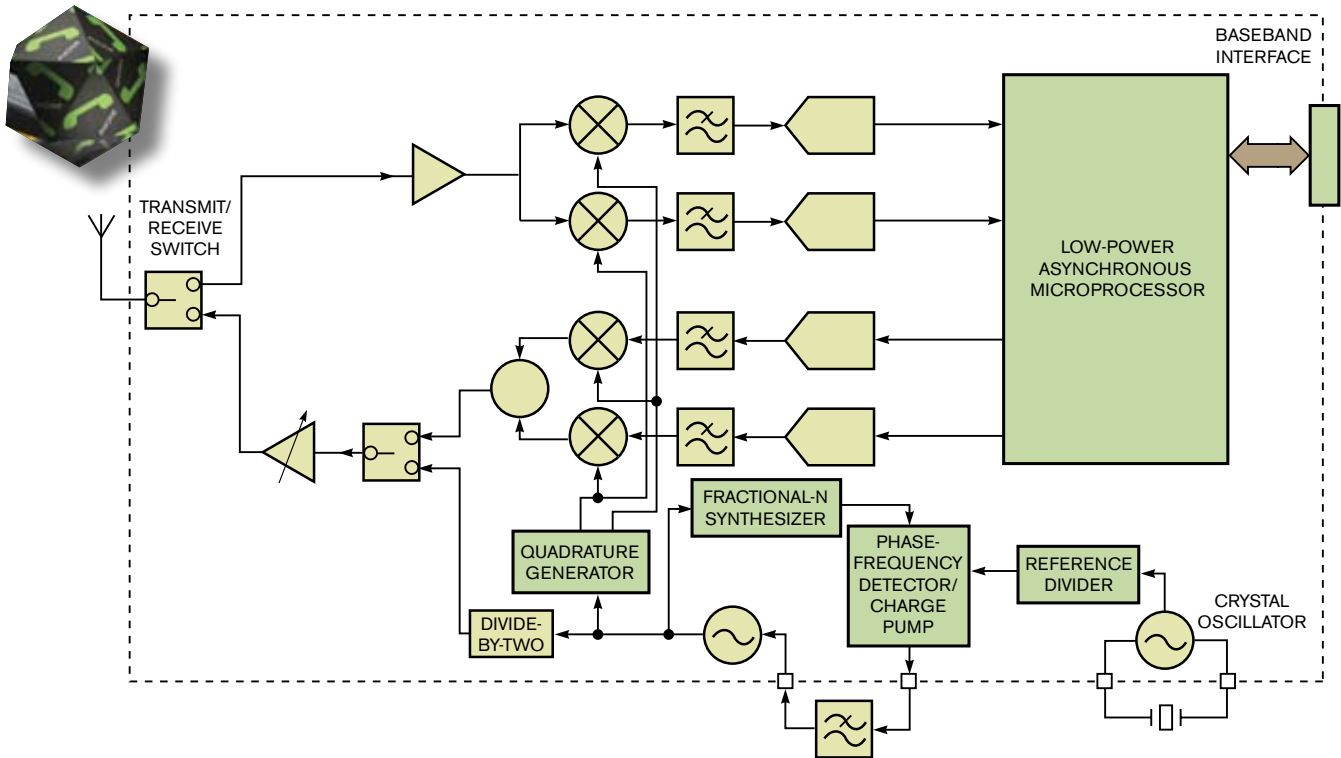


Figure 2 A typical two-way radio block shows the major functional subblocks (courtesy Berkeley Design Automation).

to simulate the environment on the final SOC. "With test chips, you can verify a great deal of the radio behavior before you find yourself signing off on final masks for an entire SOC," he says.

In any case, Rasheed emphasizes the importance of top-level simulation that is both abstract enough to view the behavior of the design as a working radio and accurate enough to predict problems. "You need Verilog-A models that reflect circuit-level reality," he says. "[Achieving that goal] takes a lot of experience in RF synthesis. It takes moving back and forth easily between Spectre, RF-Spice, and Verilog-A models. And it takes knowing where the 'gotchas' will be so you can capture them in the higher-level models and not be surprised by them down the road. Realistically, the RF designers must be involved in the chip-verification process."

Reflecting the difficulty of the verification task, Berkeley Design Automation's Subramanian describes five phases of verification for RF IP: functional simulation, performance analysis, noise analysis, investigation of the interac-

tions with package design, and analysis of sensitivities to process variations. Unfortunately, although you need to complete the first two of these steps on the block preintegration, all five are necessary after integration and layout of the SOC.

VARIABILITY

And then there is the matter of variability—not just process, voltage, and temperature variations, but also package and board variations. A radio you drop into an SOC must work in the digital-CMOS process in which you fabricated the SOC. But it must function within all the process corners for that process. And it must work in all the package variations marketing may dream up for the chip. And it must work in the customer's board design.

RF designers have two fundamental weapons in this unequal battle: robust circuit design and digital configuration. The first of these factors, circuit robustness, has been a tool of RF designers since the days of vacuum tubes. It is a matter of solid circuit-design experience, adequate

simulation, and, many designers argue, enough test chips. But with the integration of RF circuitry into digital-CMOS processes, designers have a new weapon that has changed the nature of radio design: digital configurability.

"Radios are very sensitive to parasitics, and parasitics are neither stable across process variations nor accurately modeled," argues Leme from MIPS. "We need improved design kits with more parasitic data. But even so, in the end, you will ... [be] using digital trimming to center the circuit to the process. We try to add as much configurability as possible to the IP."

Leme says that advanced CMOS processes help with adding configurability. "Once you get to 90 or 65 nm, analog switches are very good. You can use them without badly degrading signals." This ability has opened the door to a design style in which digital signals can open and close switches not only to adjust bias currents or match impedances, but also to switch active elements into or out of the signal path.

This style is new to RF design, accord-

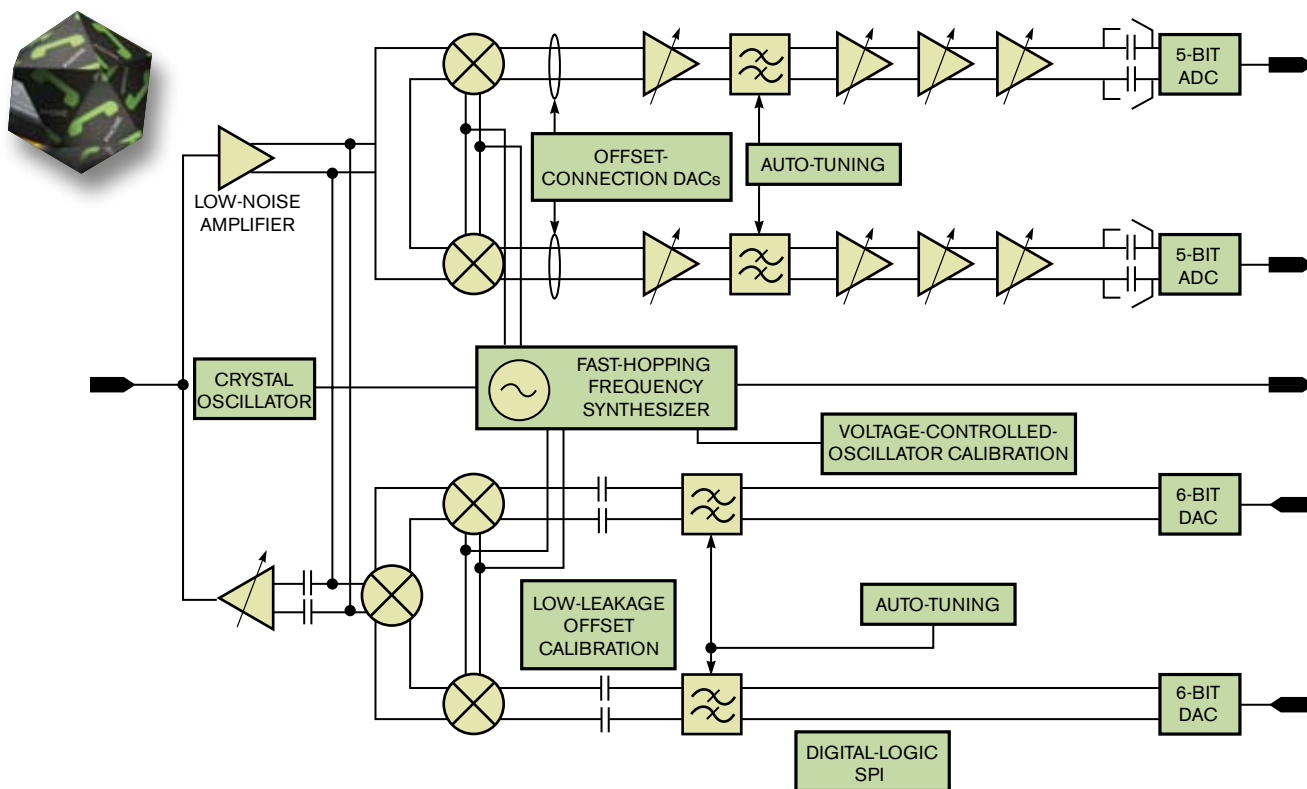


Figure 3 A UWB transceiver follows the basic pattern of the design in Figure 2 but adds functional blocks to support the UWB air interface (courtesy MIPS).



IT SHOULD BE POSSIBLE TO MAKE A RADIO CONFIGURABLE ENOUGH TO REUSE ACROSS A VARIETY OF SOC DESIGNS.

ing to Subramanian. “Advanced CMOS processes have their limitations for RF, but they give you a huge number of transistors to play with,” he says. “That [abundance] has led designers into the practice of throwing transistors at any inability to reach specifications. Consequently, RF-CMOS circuits on an SOC tend to be much larger than traditional RF designs: You can see 100,000 transistors in one of these things.”

In this new design style, configurability takes up the slack for imprecise knowledge of critical parameters at design time. For noise, in particular, Synopsys’ Nandra says, models are a problem. “Gate noise is the primary issue. If you are starting a design early in the process life—as aggressive SOC designers tend to—the transistor-level noise models may not be exactly stable. It’s a good idea to drop some process-monitoring devices into the scribe lines on test chips to help calibrate your models. Then, you can include the test structures in your IP further down the road to aid in calibration.” And the data from those structures can set digital parameters that center the circuit.

TOWARD IP REUSE

That application of massive numbers of transistors, performance monitors, and calibration circuits into RF design has changed the nature of radios from elegantly small circuits with only a few active devices to incredibly complex digital circuits with only a few RF devices on the signal path. This evolution has moved us toward the goal of reusable RF IP.

Arya Behzad, distinguished engineer and director of engineering at Broadcom, bluntly sums up the current situation: “It would be impossible to produce our product lines if we could not reuse IP. But, typically, RF IP requires more changes during reuse than digital or other analog IP.” Behzad says that this reality has led Broadcom RF-design teams to consciously design their radio blocks for reuse—depending on the application. “If we are designing a radio for a completely new market area, and we just want to get some experience, we may do a one-off design,” Behzad says. “But, if we are developing a full product line for a market in which we have some experience, we will aim for reuse, even though

this costs some die area. The concept is to take advantage of all those short-channel devices, and use lots of them to make the radio flexible. You end up with a ton of circuitry around the core.”

Area and power are obvious costs of this flexibility. So design for reuse is not a dogma but yet another engineering trade-off. “For instance,” Behzad offers, “if you do the blocks for a radio with reuse in mind, you may find that 70% of the blocks you designed for a single-in, single-out radio you can reuse in a MIMO radio, even though the requirements for the MIMO radio are more severe.”

Assuming that you have designed the radio for reuse, integration becomes a process of matching up the configurability of the IP to the requirements of the new SOC, Behzad explains. But this process in itself can be complex. He points out one 802.11n transceiver that has more than 2 kbits of digital control. “And many of those bits interact with other parts of the chip in real time,” he explains. “To verify operation in a new environment, you have to move between Verilog-A models; digital simulations; and, in some interactions, transistor-level simulations. We’ve found that this [requirement] becomes particularly an issue with initialization sequences.”

Behzad says that validating coupling between blocks is another hard part. “It’s impossible to capture everything—for instance, substrate or package coupling. The problem is that you really need to model the die, package, and circuit board together. It becomes a monster model, and the simulations won’t run. So you make some simplifying assumptions, [which] open the door for mistakes.

“You can’t get the simulations right to the last decibel,” Behzad warns. “So, you work at the beginning to make the circuit less sensitive to the things—[such as] noise spurs—that are difficult to pre-

dict or just impossible to stop.” For example, Behzad says, there is substrate coupling. “There are more claims than realities on the tools for modeling substrates,” he laments. “So, you use your experience [to determine] when to add substrate capacitors, put in guard rings, and use N wells. Where you can’t predict, make the design robust.”

Will all of these steps make an RF-IP block reusable? “Within a company, definitely, yes,” Behzad says. “But with third-party IP, I think not.”

The question at last comes back to Subramanian’s point about differentiating versus commodity radios. It seems likely that, with the relatively standard operating voltages, extremely high cut-off frequencies, and enormous transistor counts of 65- and 45-nm processes, it should be possible to make a radio configurable enough to reuse across a variety of SOC designs. It might even be feasible to make the IP relatively portable across foundry processes, although some designers are pessimistic that this approach will ever be practical.

But Subramanian emphasizes that it will never be possible to make a radio at once configurable enough for reuse, small enough and low enough in power for a critical application, and offering high enough performance for the radio to help differentiate the finished SOC. “Over time, I think we will see Bluetooth, GPS, and maybe television-tuner blocks become commoditized enough to be third-party IP,” Subramanian speculates. “But for applications where the radio’s performance helps differentiate the end product, I think third-party IP will always be unlikely.” **EDN**

FOR MORE INFORMATION

Berkeley Design Automation
www.berkeley-da.com

Broadcom
www.broadcom.com

Cadence Design
www.cadence.com

MIPS
www.mips.com

Synopsys
www.synopsys.com

Tahoe RF Semiconductor
www.tahoerf.com

You can reach Executive Editor **Ron Wilson** at 1-408-345-4427 and ronald.wilson@reedbusiness.com.

