

SINGLE-CHIP POSE PERPLEXITIES FOR SOC ARCHITECTS

INTEGRATING WIRELESS CAPABILITY INTO AN SOC
REQUIRES CAREFUL ATTENTION TO PARTITIONING
AND ARCHITECTURAL DECISIONS THAT OFTEN
HAVE NO CLEAR-CUT ANSWERS.

RON WILSON • EXECUTIVE EDITOR



Wireless links are becoming ubiquitous in electronic systems. At high data rates demanding wireless connections, WiMax and 3G cellular dominate, whereas, at the other extreme, standards such as ZigBee and Wibree move data at far lower rates and energy-consumption levels. But, in all cases, market economics and consumer expectations demand that system architects at least consider integrating the radio hardware into an SOC (system on chip).

This demand raises some critical questions. First, is a single chip the best approach, or does isolating the radio on a separate die better serve the application? Second, given the answers to these questions, what are the best architectures to employ for the receiver and the transmitter? Underlying these questions are others: Can architects make these decisions in a systematic way, or must they make them not only case by case, but also with a case-by-case decision process?

Some vendors, particularly those in the cellular-handset market, seem to sug-

gest that anything less than a radio fully integrated into the SOC is as obsolete as the vacuum-tube tuner. Determining whether to implement a radio on a separate die using an RF process or to implement at least some of the RF circuitry onto the SOC in a vanilla-CMOS-logic process is far from obvious, however. The system architect must consider issues such as market expectations, life-cycle cost, performance requirements, and project risk. The architect reaches an informed decision only by assessing all these factors.

The first item on the list is an emo-

tional, rather than an objective, judgment. Particularly in consumer mobile markets, in which cost and form-factor issues are huge, a perception often exists that only a single-die approach is acceptable. "The pressure is to do a one-chip solution, even if [a one-chip design] is hard, and that means using digital CMOS," says Maryam Rofougaran, senior director of radio-technology engineering at Broadcom. "Even an RF-CMOS process tends to cost too much, so we target digital-CMOS processes."

Other vendors that play in the cellular-handset business, at least in segments with large volumes, agree. "For us, single-die is a business consideration only: We can make a business case only for a single-die solution," explains Robert Aiello, chief technology officer of Staccato Communications. "So, we have to build a team with experience in integrating radios for volume production."

Texas Instruments, perhaps the most insistent supporter of the approach, has gone so far as to brand its CMOS-single-die-radio architecture as the DRP (digital-RF processor). "There are clear benefits to a single-chip approach," says

Bill Krenik, TI's chief technology officer of the wireless-terminals-business unit. "Today, radio chips for low-cost applications are overwhelmingly done in digital processes," Krenik says. "But the design complexity and development cost are large for a single-chip architecture, so such a design requires a large market and good expected market share. A smaller segment, such as a high-end feature phone, might not make sense as a single-chip device."

Another view that you cannot ignore comes not from customers but from the venture-capital community. "I think that if you are a start-up today, your investors are going to push you to do a single-chip design in digital CMOS, whether you think that is the right approach or not," observes Teddy O'Connell, client manager at IBM. "But there are still very significant issues you need to look at before making the decision."

Many of these issues involve the cost of the approach. O'Connell and others point out that, as multichip packaging has become mainstream, the economics of single-die versus multichip designs have

AT A GLANCE

- ▣ Market forces are pushing integration of RF circuits into SOCs (systems on chips).
- ▣ There are still some compelling reasons to keep RF on a separate die.
- ▣ The choice of radio architectures—especially for receivers—is far from settled.
- ▣ The architect must think across the whole system, not just one chip.

changed. The budgetary cost of a two-die assembly may be greater than that of a single packaged die, but that difference is just the beginning of the equation. "That single die may be 20% RF circuitry, with significantly lower yield than the digital portion," O'Connell points out. "And, as the design migrates to finer geometries, the RF section may grow rather than scaling down. Further, the need for RF performance may push the entire design into a more advanced digital-process node than would have

been necessary if the design used separate RF and digital dice."

The cost question can be tricky, however. Choosing a single-die or multichip strategy can influence the choice of radio architecture, which in turn has significant influence over the number and quality of external passive components the radio will require. And those passives, after you pay for components, real estate, insertion, and test, can be significant factors in total system cost.

The questions of performance requirements and design risk also interrelate. On paper, it would appear that a 65-nm or even a 90-nm CMOS-logic process should work just fine for RF. "Just a few years ago, gigahertz-range RF meant GaAs [gallium-arsenide] processes with expensive transistors," observes Analog Devices' business-development director, Doug Grant. "But process shrinks have brought lower parasitic capacitances, shorter transit paths for carriers, and larger transistor budgets, all of which are benefits for RF designers."

Transistor cutoff frequencies in these processes can be greater than 40 GHz for

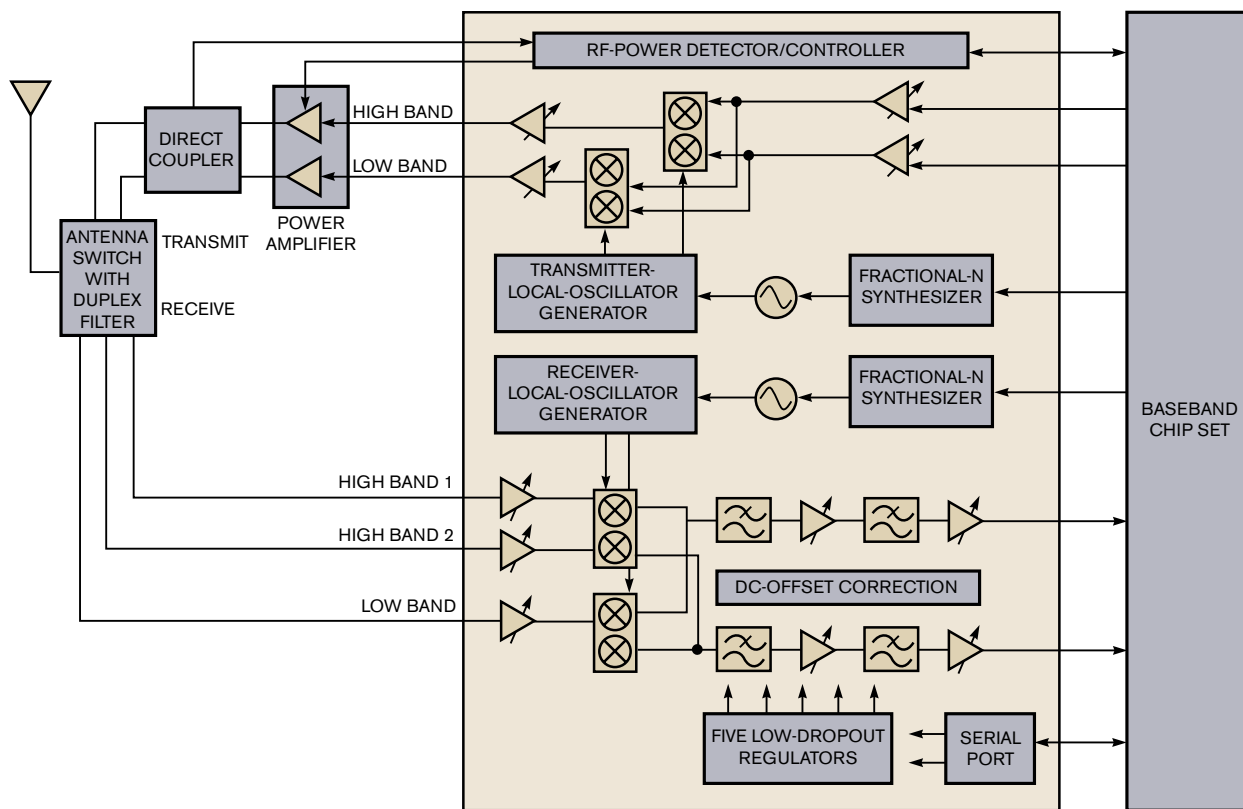


Figure 1 Analog Devices' single-chip radio for wideband-CDMA applications includes on-chip voltage regulators to stabilize delicate circuits.

a 90-nm process and significantly higher for a 65-nm process. The lower series resistance of the small devices means that noise floors are lower, somewhat compensating for lower operating voltages. Linearity is not bad: “The third-order intercept is actually better in CMOS than it is in SiGe [silicon germanium],” says IBM’s O’Connell. And you can fabricate reasonably good spiral inductors, as long as the circuit does not require a very high Q.

So, why would anyone cite performance as a reason to use a specialized RF process on a separate die? There are plenty of reasons.

PROCESS AND PERFORMANCE

One good reason is frequency. “Regulatory requirements are pushing ultra-wideband radio links up to higher frequency bands. This [pressure] made operation at 10 GHz one of our design requirements,” says David Shoemaker, Alereon’s vice president of engineering and operations. “At that frequency, CMOS is definitely getting closer to viability; it can be fast enough. But, at those frequencies, it doesn’t take much variation to have a significant problem. For instance, it’s easy for VCOs [voltage-controlled oscillators] to become unstable, and there are issues with gain flatness.” After looking at these issues, Alereon chose to implement its RF front end in a SiGe process, with a promised migration path to CMOS. “We have the option to migrate the circuitry,” Shoemaker says, “but, right now, doing so wouldn’t buy us much in improved cost.”

Aside from gain, architects considering CMOS for RF have to think about linearity. The sensitivity of a radio to non-linearity depends on the application and modulation scheme. “Standards can have a big impact on whether you can implement a radio in a given technology,” says Staccato’s Aiello. “For instance, in our world, we use QPSK (quadrature-phase-shift keying), not 64-QAM (quadrature-amplitude modulation), and that substantially reduces the linearity requirements on the RF section.”

Another serious issue with digital CMOS is models. When high-frequency models are available at all for digital-CMOS devices, they are usually digital models. There is often little or nothing in the way of small-signal models.

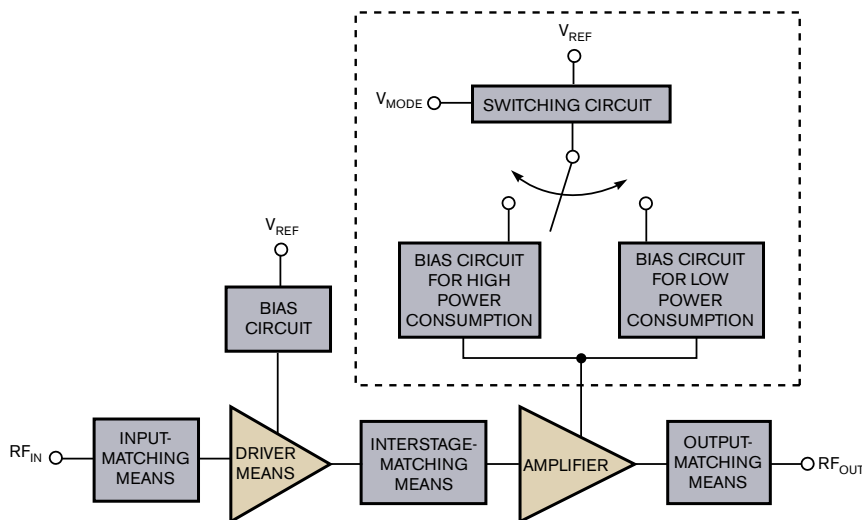


Figure 2 Even a simple RF amplifier must use a high level of digital intervention.

“Models are always an issue if you are an early RF user on a process, whether it’s CMOS or not,” says Dave Wright, lead systems architect at Cypress Semiconductor. “Even the BiCMOS process we use for our wireless-USB products was originally designed for high-speed-communications circuits, not RF. It was well-characterized but at different operating points than the ones we used.”

The transistor models are not the only issues. Noise models—especially good models of noise propagation through the power supplies and substrate—are essential. This situation is true for any RF design in which low operating voltages make noise an issue. It is also true in systems such as transceivers with on-chip power amplifiers and designs using multiple antennas in which multiple uncorrelated RF signals are present. Even if the models appear good, you’ll have the nagging worry that process engineers get paid to get good yields on big digital designs. They aren’t necessarily keeping careful watch over small-signal RF-model parameters while they cen-

ter their processes for gates and SRAM cells. This situation makes it almost incumbent upon a design team using digital CMOS to have its own modeling group and a strong enough relationship with its foundry to calibrate its own models.

When uncertainties with the models exist, the design has added risk, which is difficult to quantify. Despite that difficulty, designers must include that risk in their decisions. If a plausible risk of a design spin—or three—is unacceptable, they have two alternatives. One is to rely on a proven design team and a mature RF process. The other is to rely on digital calibration of the RF circuits to compensate for the uncertainties in models or variations. “Basically, we put compensation where we are least confident in our models,” says Analog Devices’ Grant. He says that it is not unusual for ADI designers to use 20 or 30 additional transistors in what would have been a simple bias generator, just to stabilize the CMOS-signal path over the operating-temperature range. To control voltage variations within the CMOS die, the company sometimes implements low-dropout regulators adjacent to delicate RF circuits (Figure 1). Alternatively, designers may employ digital techniques: using an ADC to make on-the-fly measurements of a circuit’s operating point, computing compensation parameters, and load-

(continued on pg 36)

MORE AT EDN.COM

⊕ For a Web-exclusive sidebar to this article on the implications of high-integration RF, visit www.edn.com/070802df.

⊕ For a profile of a ZigBee-chip design, go to www.edn.com/article/CA6325290.

(continued from pg 32)

ing a DAC to adjust a bias or offset (Figure 2). Some designers joke that they are moving to a design style that has the output of a DAC attached to every terminal on every RF transistor.

“There’s no question that making a radio scalable is a digitally intensive process,” says TI’s Krenik. “Linearity, for instance, is one of the fundamental challenges. Old-school design said you had to have a very-high-linearity, high-voltage front end. But, with today’s processes, you have to let go of that notion. Now, we use precise digital control over bias and a few other parameters to linearize a less-well-behaved circuit.”

But this level of control in itself introduces another issue into the single/die controversy. “Power is one of our most important parameters,” says Adam Gould, senior vice president of product development at NextWave Wireless. “Normally, you would assume that an older process would require more power. But, in CMOS, you spend power consumption to get linearity. In WiMax, for instance, achieving sufficient linearity in the power amplifier is one of the biggest issues in the power budget. Right now, an optimized RF process is giving us better power consumption.”

So, the question of one chip versus two is not clear-cut for everyone. In some markets, there is no practical choice: The market expects a single-die approach, and the cost ceiling demands it. In these markets, a typical implementation might be a baseband digital-CMOS SOC integrating a radio. In other markets, a multichip package may allow the design team to go either way, and the decision depends on performance and power requirements, the choice of radio architecture, the availability of stable models, and the experience of the team. “You see both digital CMOS and SiGe RF in wideband-CDMA today,” IBM’s O’Connell says. In yet other areas, particularly at 10-GHz and higher frequencies, no practical alternative

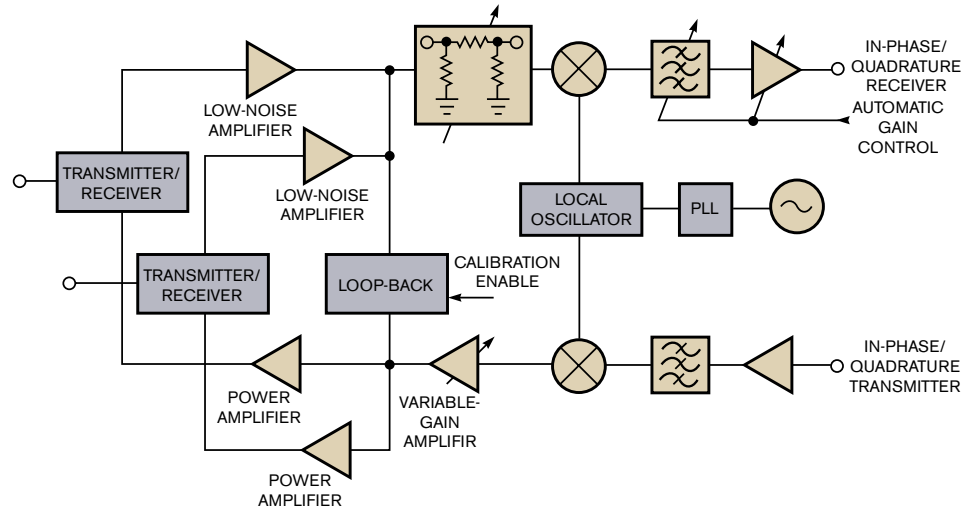


Figure 3 The 10-GHz Alereon ultrawideband-radio chip exemplifies the most common receiver architecture: the directly coupled receiver.

yet exists to RF processes for receivers. “WiMax, 5-GHz networks, and radars are still clearly SiGe applications today,” O’Connell says. Note, however, that the frequency at which a specialized process such as SiGe or GaAs becomes necessary can be considerably lower for power amplifiers and antenna switches, which involve significant power, than for receivers.

For power amplifiers, other issues arise beyond just the choice of process

ON THE RECEIVER SIDE, THERE IS A GROWING PREJUDICE THAT, THE SOONER YOU CAN DIGITIZE THE SIGNAL, THE BETTER.

technology. “As your effective channel length gets shorter, the problem of an integrated power amp gets harder and harder,” says Rofougaran. “The very small transistors can’t tolerate large signals for a number of reasons, including reliability and local heating. There can be really huge currents involved in power amps under some operating conditions.” Despite that fact, Rofougaran says, Broadcom does integrate the power amp in some low-power appli-

cations, such as some of the company’s Bluetooth chips.

“Technically, it’s feasible to integrate any power amplifier up to about 10 mW. But about 4 mW is the practical limit today,” says Cypress’ Wright. “Beyond that level, you have structures in the power amp that start consuming significant amounts of static power even in low-power modes, eating into battery life. Voltage swing can also be an issue, especially with more advanced processes.”

RADIO ARCHITECTURES

Decisions about one die or two interact with decisions about the radio architecture designers will implement. These decisions include choices for both the receiver and the transmitter sides of the radio. On the receiver side, there is a growing prejudice that, the sooner you can digitize the signal, the better. Some vendors have experimented with placing an ADC on the output of the low-noise amplifier and directly digitizing the incoming RF with only a band-pass filter between the converter input and the white-noise world. At the other extreme, some designers continue to use superheterodyne-receiver architectures that date back to the early days of vacuum-tube radio. As usual, the most frequently chosen architecture—the direct-conversion receiver—lies somewhere in between.

In many ways, just putting an ADC on the low-noise amplifier—or, for that matter, the antenna—would be ideal. This approach largely eliminates the problems and pains of RF design and gives rise to new possibilities. For instance, a cellular base station that simply digitizes the entire cellular spectrum could simultaneously filter bands and select channels for many modulation schemes. Designers would then not need to assign an RF front end to each active session, and the quality of filtering in the digital domain could be superior to anything a physical filter could achieve.

“But the converter for such a design would be very difficult and, in any case, would take plenty of current,” says O’Connell. To capture enough information for the digital processing to succeed, the converter would have to have an enormous sample rate; a large dynamic range, so that the DSP could extract small incoming signals from powerful neighboring-channel signals and out-of-band noise; and, for similar reasons, superb linearity. Although such a converter design is now possible, O’Connell believes that it would currently be feasible only for high-end designs, such as base stations and multistandard “world phones.”

A better choice, in the view of most radio architects these days, is a direct-conversion receiver (Figure 3). In this architecture, the RF signal passes from the low-noise amplifier into a mixer stage that multiplies the RF signal by I (in-phase) and Q (quadrature) signals from a local oscillator. The resulting products are two signals—I and Q—already at baseband frequency. Normally, the baseband signal goes through an antialiasing filter into a pair of ADCs.

This approach doesn’t make things particularly easy. The direct-conversion architecture emerged in the early days of vacuum-tube radio soon after the superheterodyne, but designers abandoned it because of its challenges. The low-noise amp and mixer still must be linear. The local oscillator still must be clean and stable, and the suitability of the direct-conversion architecture depends to some extent on the modulation scheme. “It depends on the data,” says Broad-

com’s Rofougaran. “In our 2046 Bluetooth chip, we are dealing with GFSK (Gaussian-frequency-shift keying) that has a relatively narrow data bandwidth but puts a lot of the information spectrum close to dc. A direct-conversion architecture would be a problem, because of the issues with dc-offset cancellation. So, we use a low-IF superheterodyne arrangement. In contrast, wireless LANs have relatively wide data bandwidths and little information near dc, and we use direct-conversion architectures for them.”

Cypress also opted for a low-IF superheterodyne architecture for its wireless-USB device. “It was optimal for our situation,” Wright explains. “We have a good circuit design and are getting better spurious rejection than others using direct-conversion architectures. And, with the low IF, we are not burning a lot of power in logic circuits. We could have moved to direct conversion at some point, but it’s sometimes best to just reuse an existing architecture.”

One issue that works against the superheterodyne is the difficulty of system integration of a receiver employing the approach. Alereon’s Shoemaker points out that IF sections tend to generate frequency spurs that aren’t necessarily a problem for the superheterodyne receiver itself but can wreak havoc with other receivers in the system. Because many mobile devices these days contain multiple radios, system architects must carefully plan the frequency spectrum for all the radios in the package—before designing the devices.

Similar decisions arise on the transmitter side. Here, the most common architecture appears to be the polar-loop power amplifier. This design, incorporating feedback loops for both amplitude and phase signals, provides good linearity in both amplitude and phase—important for modulation schemes such as QAM that encode information using both quantities. However, designers have discussed other schemes, including futuristic ones, such as directly driving the antenna with a power DAC.

So, in decision-making for radios on SOCs, growing pressure from end markets and investors urges manufacturers to go to single-chip designs, at least

with small-signal RF circuitry. Designers using this approach must overcome formidable problems in integration, however. Many of these problems respond to the aggressive use of digital circuitry for calibration and control and to design experience. But the industry is a long way from the day when high-frequency receiver-RF stages will as a matter of course go onto SOCs and even further from the day in which any but the lowest power transmitter circuits will dwell there.

Because CMOS-process technology is not moving in the direction of greater friendliness to RF design, increasing integration will likely mean—as it did several years ago at TI—significant innovation in circuit design and in radio architecture. Analog Devices’ Grant makes an important observation about how designers achieve progress in this direction. “On a project a few years ago, we stood back and looked at the whole design,” he says. “We realized that the RF, digital, and software designers really didn’t talk to each other much. They were all working on their own separate noise budgets, and that [situation] was leading to some real inefficiencies in the system. We realized at that point that you have to optimize across the whole signal chain. You can’t separate the blocks from each other. And you have to look at the whole radio, including the filters, the digital controls, and the operating modes the software creates; you can’t just look at the silicon.” EDN

FOR MORE INFORMATION

Alereon
www.alereon.com

Analog Devices
www.analog.com/en

Broadcom
www.broadcom.com

Cypress Semiconductor
www.cypress.com

IBM
www.ibm.com/us

NextWave Wireless
www.nextwave.com

Staccato Communications
www.staccato.com
communications.com

Texas Instruments
www.ti.com

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

