

BY RON WILSON • EXECUTIVE EDITOR

DESIGNERS CAST A SKEPTICAL EYE ON MIXED-SIGNAL SOCs

THE FUNCTIONS
ARE NECESSARY,
BUT INTEGRATION
CHALLENGES KEEP
ANALOG IP OUT OF
THE MAINSTREAM
FOR SOC DESIGN.



The easiest lesson to draw from the consumer-electronics market is that integration is everything. Mobile beats portable. Palmtop beats mobile, and shirt pocket trumps palmtop. Even in desktop devices, such as game consoles, sleek packaging and low-manufacturing-cost targets depend on ever-higher levels of integration. This trend is not limited to low-cost consumer electronics. Perhaps inevitably, military, automotive, medical, and even industrial applications are demanding the capabilities that consumers want. The integration trend will likely spare only the largest and lowest volume system designs.

But the classic tool of integration, the SOC (system on chip) is running up against a technological barrier. SOCs advance by vacuuming up all the digital functions in a system until there is nothing of significant cost or function left outside the chip. But this

approach eventually undermines its own success: When nothing is left to vacuum up, integration is over. Many system designs are approaching that state today. The SOC has absorbed all of the significant digital blocks, leaving only commodity parts, such as mass memory and passive components, with one big exception. Most systems still contain significant cost and functions in precision analog or RF blocks that have remained outside the SOC. Accordingly, these blocks have become the next frontier for SOC integration.

You can already see this trend in such speculative designs as recent “single-chip” cell-phone-handset SOCs from Infineon and Texas Instruments. Such chips include not only the familiar baseband and application processors, but also the analog audio circuitry and the small-signal portion of the RF circuitry. They are not only leading-edge examples of digital integration, but also tours de force in analog and RF integration.

These chips also exemplify what an integrated-device manufacturer can achieve if it has enormous resources; design teams that include analog, RF, device modeling, and process engineers; and an intimate link between the chip-design and process-engineering teams.

Ordinary SOC design teams with ordinary skills and resources are not developing these designs. But can they? Can designers bring precision analog or RF blocks into the IP (intellectual-property) assembly-design flow that they have successfully used to create digital SOCs? Or are there barriers, such as inappropriate IP, missing tools, or fundamental issues that will prevent designers from handling analog blocks as black-box IP cores?

THE SOC FLOW

The existence of the digital SOC rests on the ability of its design flow to control complexity. That flow, IP-block assembly, in turn depends upon the existence of previously designed functional blocks that designers can treat almost as black boxes through the early stages of the design cycle. This approach allows behavior-level modeling of the entire chip before designers begin a detailed design, allowing the designers to synthesize, place, and route the individual blocks relatively independently of each other and vastly simplifies the chip-level-verification process. Any experienced SOC designer would state that these processes describe gross oversimplifications of how SOC design actually occurs. However, the IP-assembly flow

does significantly reduce design-process complexity compared with treating all of the nets in an SOC with equal attention.

Can designers extend this flow to precision analog or RF circuits? The initial answers from seasoned designers are decidedly mixed. In some simple cases, designers have already done so. For example, some ASIC vendors offer drop-in blocks with low-speed ADCs or moderate-performance PLLs. In other cases, experts say that such an extension to more complex circuits is impossible. “For us, analog design is a matter of constant evolution, not of design reuse,” explains Julian Hayes, vice president of marketing for consumer products at Wolfson Laboratories. “IP can provide the foundation for a new design, but it can never be used for cut and paste.”

Fundamentally, the digital-IP-integration flow depends on a number of assumptions. First, it assumes that such a thing as reusable IP exists—that a designer can use a block that functions properly in one design for the same function in another design without modification. Second, the methodology assumes that it is possible to model the behavior and timing of a block with acceptable accuracy without understanding the details of its internal function. Third, it assumes that the behavior of a block is independent of the placement and routing of the block and of the internals of signals not connected to it. The problem with precision analog or RF blocks in an SOC flow is simple: They violate each of the assumptions.

REUSABLE IP

An IP-assembly flow can’t work if there is no such thing as reusable IP. And many analog experts suggest just that: No analog design is reusable without the intervention of a skilled analog designer. Part of the issue is that nothing in the analog world is analogous to RTL or to synthesis. Designers have attempted to produce analog-synthesis tools, but these tools have for the most part failed, whether through simply not working; working, but only with the intervention of skilled analog engineers; or working only on a narrow range of functions and environment. Today, essentially no analog-synthesis



AT A GLANCE

Running out of digital functions to integrate, SOC (system-on-chip) designers are starting to eye complex analog and RF functions.

Analog blocks don't fit into the standard IP (intellectual-property) integration strategy for building SOCs; they have too many interactions with the rest of the design.

Both using conservative design and having analog experts participate in the integration process are proven solutions to the problem.

In the future, digitally controlled and self-adjusting analog blocks may serve as black-box functions for SOC designers.

tool is on the market that design teams report they are using.

That lack of availability still leaves open the possibility of hard IP; a designer can drop in a placed-and-routed design unaltered. This approach can work, according to many designers, if the requirements on the block are easy, if the user has sufficient design guidelines to ensure that the circuit stays within its original design space, and if the design has sufficient isolation. But those are big ifs.

A third possibility, which companies including Barcelona Design have explored, is to use a topology to define an analog block and then use automated tools to scale the devices within the topology to meet specifications in an environment. Unfortunately, this approach has not met with great success, either. So, to a first-order approximation, the only reusable analog IP is simple analog IP.

BEHAVIORAL MODELING

To listen to people in the EDA market, you'd think that the mixed-signal-behavioral-modeling problem is the one part of the puzzle that designers have solved. Unfortunately, that is not the story one hears from experienced analog-design teams. For example, Analog Devices may be one of the most advanced design groups in language-based modeling of

analog- and mixed-signal circuits. The company has its own internally developed hardware-description language, which can model both discrete- and continuous-time blocks in what engineers describe as C-level abstractions (Figure 1). David Robertson, product-line director for high-speed converters at the company, says that it takes a few weeks to a few months to get a current engineering graduate up to speed on the tool. But just knowing how to use the tool is only part of the problem, Robertson adds. "Knowing where to apply it is still an art," he says. In critical situations, modeling still goes right back to the Spice level.

Even vendors of analog IP are skeptical when it comes to accurate behavioral modeling of an analog block. Knowlent Chief Executive Officer Sandipan Bhanot says that, in the analog domain, Spice is still the ultimate court of appeals. "There really isn't anything equivalent to the testbenches, modeling, and measurement languages of the digital domain," he says.

THE DELICACY OF ANALOG

If there has been some progress in design reuse and in high-level modeling of analog or RF circuits, the great issue that remains is independence. Far from being little functional blocks a designer can drop into a design wherever necessary, analog blocks are exquisitely sensitive to their surroundings. One way to think about the problem is in terms of digital and analog netlists. In the digital

world, an IP block has inputs, outputs, and power and ground contacts. You assume that any interaction between the block and the rest of the chip takes place on the input- and output-contact points. Most digital flows, at least those below 180 nm, recognize the possibility of capacitive coupling to adjacent routing. However, rather than model these couplings as additional contacts on the block, designers generally ignore them until detailed postroute extraction and then model them as added delays rather than as signal sources or passive networks.

In the analog world, it would be wonderful if that simple model would work. And it can work, assuming that the analog circuit in question is robust to begin with, that it is operating well within its performance envelope, and that the designer has observed sufficiently demanding guidelines with regard to external noise sources, impedances, and isolation. Without those assumptions, the real model of the analog block becomes considerably more complex (Figure 2). Even at a behavioral-level view, clock, power, and ground pins become signal paths, not abstract concepts that don't go anywhere. Thus, clock-line cleanliness and supply coupling become fundamental issues that designers must model to determine how the circuit will behave.

This problem is becoming worse, not better. Increasingly, large SOCs are employing aggressive power-management techniques that include clock and

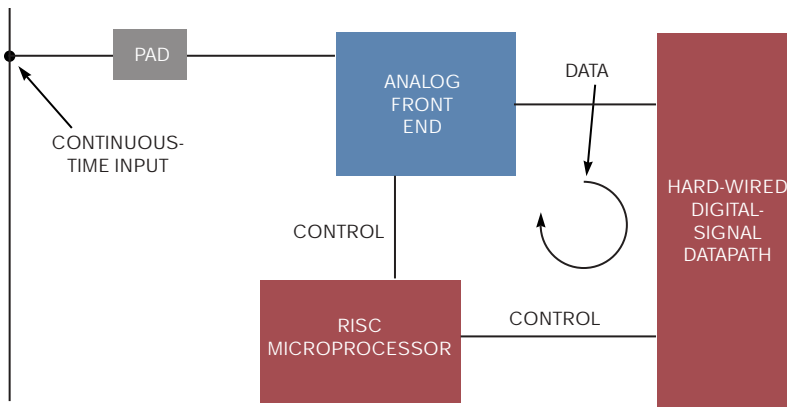


Figure 1 Increasingly, analog blocks depend on tight interaction with digital blocks and even processors, requiring a sophisticated mixed-mode simulation environment even for behavioral modeling (courtesy Analog Devices Inc).



power gating and dynamic voltage and frequency scaling. The use of these techniques means that the clock and supply networks on an SOC can change their effective topology, as well as their frequencies and voltages, on the fly, enormously complicating analysis. Ideally, no paths would exist between the supply pins on analog and digital blocks. But, with the increasing use of digital feedback and control into analog blocks and with the inevitable parasitic coupling of large transients across nets, you cannot assume that analog clock, power, and ground signals are clean. Many teams now perform detailed extractions of these nets and use Spice to model them before tape-out.

Perhaps more problematic still in SOCs with high-speed digital blocks, noise can couple into any node in an analog circuit through the substrate. This problem is especially worrisome for a number of reasons. First, the designer may not understand the location of the noise sources even after performing detailed floorplanning. The bad news may not arrive until the designer completes detailed placement. Second, digital designers are usually unaware of the noise their circuits are injecting into the substrate, so they may be no help in identifying signal sources, let alone locating or quantifying them. Digital designers

don't tend to think in terms of digital signals having frequency spectra. Third, neither design teams nor, frankly, foundries are likely to have adequate substrate-electrical models unless they have been previously involved in RF design. Fourth, in advanced processes, the high-frequency components of digital transients—and even harmonics from digital clocks—are so far into the RF region that simple substrate models may be highly misleading. So, even if the designers could locate and accurately model the noise sources, it might be anyone's guess what effect the sources would have on a particular node.

More widely known coupling problems between chunks of metal in the interconnect stack are just as problematic. Some design disasters on processes as large as 180 nm forced digital designers to recognize the role that capacitive coupling could play in signal integrity. But the tools that designers built to deal with the problem in the digital domain generally use static guidelines to screen layouts for possible capacitive coupling. The tools then either just flag those problems or degrade the delay parameters on the net on the assumption that, because everything is synchronous to the same clock, if another signal is coupling onto a victim net, you just have to wait for the aggressor to set-

tle and the noise will go away.

This model is less than helpful for analog blocks, of course. Some low-performance analog blocks hold during clock transitions of surrounding digital circuitry or simply shut down while big digital clock trees are operating. But, most of the time this approach is impractical, and the designer must protect the continuous-time analog circuits from aggressors. No tools in the digital domain can examine the surroundings of an analog block with that level of detail. So, it again comes back to the analog-block integrator: Extract accurate models of the parasitic capacitances, add them into the Spice model, and perform detailed simulations.

THE INDUCTANCE NIGHTMARE

This approach is not fun, but it is manageable. Another problem looming on the horizon—inductive coupling—is less manageable, however. A couple of years ago, researchers in digital tools were concerned that inductive coupling would finally defy Moore's Law, because inductive coupling gets more efficient with increasing frequency, and it strongly depends on the 3-D geometry of the aggressor and victim structures. Worse, the victims need not be nearby. All these issues appeared to present an uncomputable problem. Subsequent evidence

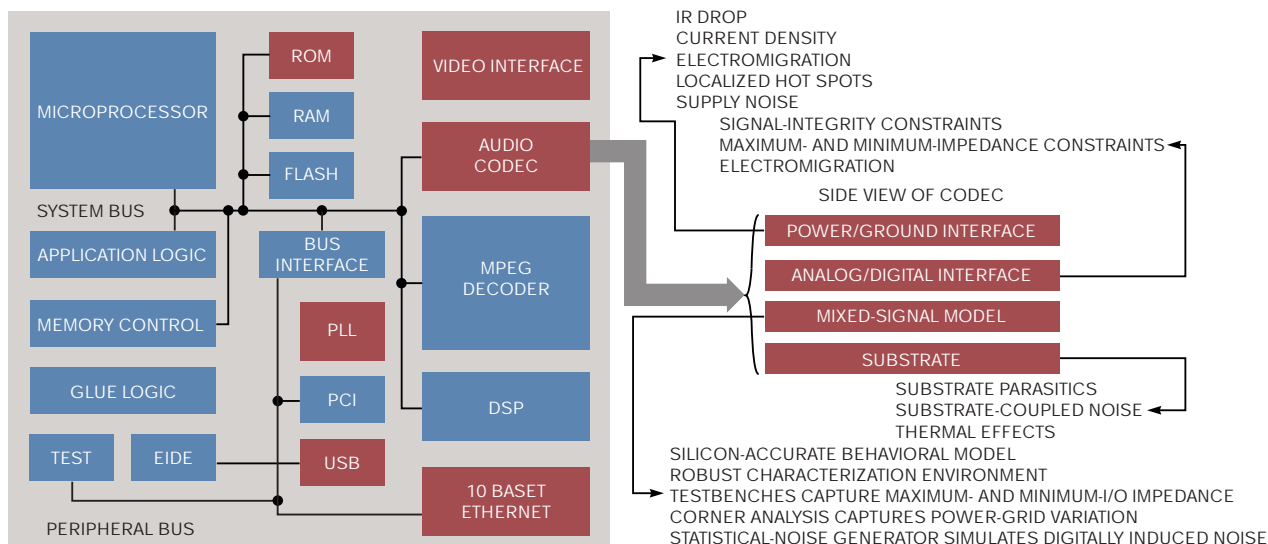


Figure 2 At every level of abstraction, a precision analog block breaks the simplifying assumptions on which digital-IP reuse depends. This approach complicates block integration (courtesy Cadence Design Systems).

suggests that, even for multigigahertz designs, the problem for digital SOCs may be much less than designers originally feared. However, no such reason for reassurance exists about the interaction between high-speed digital circuits and—in particular— analog RF blocks with on-chip inductors. Designers must consider inductive coupling for blocks operating at these frequencies.

Even connecting the signal pins to those of surrounding blocks involves a different skill for analog blocks from that for digital blocks. In the digital world, the only problem is to connect the right components with a trace that routes. If loading issues exist, the timing-analysis tools and buffers will handle them. In the analog world, every input and output requires more thorough specification (Figure 3). Voltage levels are not approximate and understood; engineers must specify them. Signal spectra are important. Noise that couples into the block from inputs or through outputs is important. Impedances matter. You could insert an analog block into a design, connect its pins in a logically correct manner, and inadvertently subject the block to impedances that prevent its operation. So, the noise and impedance characteristics of the pins on other blocks, as well as the impedances of the interconnect lines themselves, become issues.

The situation becomes even more interesting in designs using 130-nm or smaller geometries, in which design-for-manufacturing tools may change the physical layout of a metal layer after routing is complete. Tools that move metal lines to comply with line-pitch or -density rules can substantially change the electrical characteristics of a metal run. Tools that insert “dummy” metal— islands of metal that do not physically connect to the rest of the circuit, but keep the density of metal features constant across the surface of a layer—can also change the metal run’s electrical characteristics. Dummy metal beside or over an analog path can provide a new set of parasitic capacitors to many other circuit nodes.

GETTING IT DONE

All of these problems may sound formidable. Yet engineers say that solving

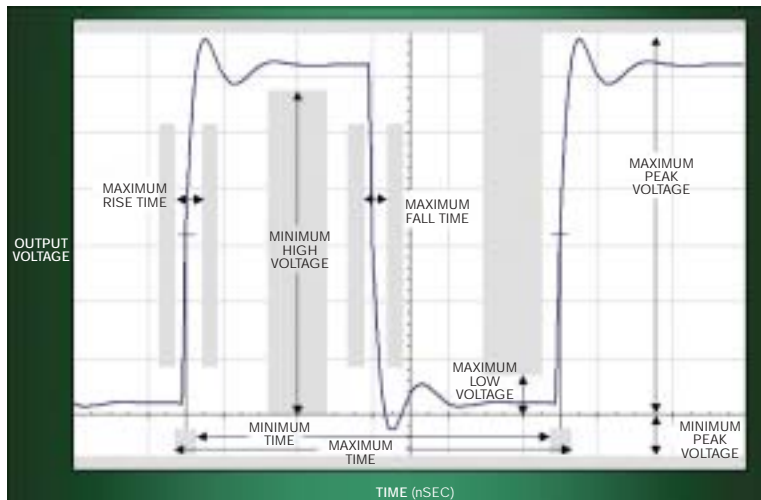


Figure 3 Just writing the specifications for an input to an analog block can become a significant design undertaking, requiring detailed knowledge of the internal circuitry (courtesy AMI Semiconductor).

them is difficult but not impossible. It’s worth examining how a few shops are dealing with these issues to see what the future might be for analog integration into SOCs for the rest of us. One strategy is to encapsulate the analog IP in a purely digital wrapper. This technique allows you to model the block at the system level as a digital block and to treat it as conventional hard IP during physical design. This approach requires that the entire analog content of the design fit into one block or at least that the analog blocks have only digital links to each other. It also requires that you put enough thought into the isolation requirements and that you can integrate the block using a set of guidelines that eliminate any possibility of injected noise’s becoming a factor in the performance of the circuitry. That requirement in turn means using conservative design techniques. “The concept of ana-

log building blocks is viable, but it implies higher power, greater area, and lower performance,” says Wolfgang Meier, senior manager of business development at Infineon Technologies’ ASIC operation.

Such conditions may not be entirely practical in the ASIC world, but they can exist in another important technology: microcontrollers. Here, analog performance may not be a differentiating factor, and long experience with a single analog library may refine the integration process until it works smoothly. This scenario has been true at Silicon Laboratories. “The tradition was 100% hand-routing for analog,” says Silicon Labs’ vice president of technology, Douglas Holberg. “But our principle has been to exploit reuse. Today, we can effectively reuse 8- and even 16-bit data-converter blocks with a drop-in model. The integrator has to understand the block through the interface stages, and he has to understand the top-level process for assembling blocks in our methodology—practices about how you take a signal across a power-supply boundary or between clock zones, for instance. And we use techniques, such as deep N-wells, to make the block look as digital as possible.”

But is the result turnkey reuse? “In the beginning, it worked only because we had the analog experts involved to keep the integrators from breaking the block,”

FOR MORE INFORMATION

Actel www.actel.com	Infineon Technologies www.infineon.com
AMI Semiconductor Inc www.amis.com	Knowlent www.knowlent.com
Analog Devices www.analog.com	Silicon Laboratories www.silabs.com
Barcelona Design www.barcelonadesign.com	Wolfson Microelectronics www.wolfson.co.uk
Cadence www.cadence.com	ZMD www.zmd.com

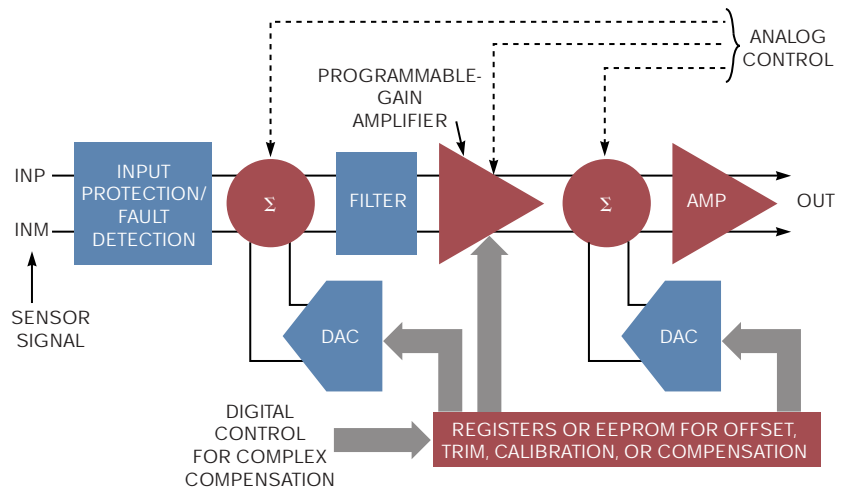


Figure 4 As transistors get smaller and analog-signal processing gets more complex, digital circuits must assist in trimming, compensating, and sometimes enhancing the functions of analog-signal paths (courtesy AMI Semiconductor).

Holberg says. Such a technique may result in turnkey reuse now, however.

AN ASIC APPROACH


What happens when you can't detune the analog blocks to improve integration? In that case, the model often becomes one of heavily assisted ASIC design. In a conventional digital-ASIC relationship, the customer would hand off a verified netlist, and the ASIC house would do the physical design and extraction. Similarly, in a mixed-signal ASIC, the customer may obtain relatively simply functional models of analog blocks from the ASIC vendor's library. The customer would use mixed-signal-simulation tools to approve the behavior of the blocks, and then the vendor design team—strong in analog-design skills—would integrate the analog functions into the design.

This model is the one that AMI Semiconductor uses, for example. The company keeps an extensive library of analog functional blocks and provides models to customers. When the customer is happy with the netlist, including the analog functions, the AMI team takes over and does the integration. That process may not involve an exact drop-in assembly. "We can leverage an extensive analog-block library," explains AMI's director of mixed-signal products, Ryan Cameron. "So, we rarely have to start with a clean sheet of paper. But the

result is rarely a drop-in, either. Except for simple blocks, it's probable that the analog circuitry will receive modification in some way during the integration process." That situation means that senior analog designers must take part in the process.

AMI tries to use a multitrack design flow, in which behavioral-modeling engineers work with customers to establish architecture and block specifications that they base on chip-level simulations. They must also define waveforms, envelopes, and corner-case behaviors for the analog pins. Meanwhile, circuit designers pull blocks from the library and tweak them to fit the chip-level specifications. Ideally, the two meet in the middle; designers then adjust the blocks to the needs of the chip design and adjust the chip architecture to require a minimum of new analog design. Two areas that most frequently become issues in adapting a block, according to Cameron, are initialization sequences—especially for high-voltage blocks—and details in the board-level environment.

Designers from another quarter equally emphasize the importance of communication between chip- and circuit-level designers during integration. Actel Corp acquired much of the analog IP it integrated into its family of mixed-signal FPGAs. That acquisition required a significant effort on Actel's part to assemble the analog blocks into a chip along with



the flash-based logic array and I/O circuitry, keeping the blocks configurable and not breaking them.


A POSSIBLE FUTURE

So, does a block-assembly methodology for mixed-signal SOCs depend on having a staff of analog gurus in the back room? Today, the answer is probably “yes.” But some changes on the horizon might make a difference. One is the increasing power of mixed-signal-design environments. As it becomes increasingly possible to accurately characterize analog blocks—including sensitivity analysis, not just functions and parameters—with a few test chips, it becomes more possible to harden analog blocks for the misfortunes that can happen during integration. It also becomes possible to characterize and simulate more of the important interactions between the block and its neighbors—moving this analysis from the Spice domain and into the realm of faster and less user-intensive mixed-signal-simulation tools. But tools for sign-off-quality noise modeling and analysis are still problematic. “It’s an area in which we still have some work to do,” admits a marketing director at one EDA vendor. “The work is going on only at the university level today.”


Another important change is also under way: the increasing degree to which analog circuits are coming under digital control. Engineers at Actel, for instance, compensate the on-chip RC oscillators in the company’s devices for voltage and temperature using a table in flash memory. AMI often puts a fair amount of digital programmability into analog blocks early in the design process and then backs much of it out as the device heads from early silicon to full production (**Figure 4**). In that way, AMI can use nonvolatile memory to trim early devices to meet specifications, and, as the company gains a better understanding of the design, it can pull out the trimming circuitry to improve die area.

This trimming capability may be as simple as controlling MOSFET switches in a resistor or capacitor network. It may be large-scale yet delicate, such as switching in and out stages in an RF-power amplifier, as ZMD does in the linear power amplifier on its single-chip ZigBee

MORE AT EDN.COM

 We encourage your comments! Go to www.edn.com/060511cs and click on Feedback Loop to post a comment on this article.

 To see how little has changed in two years, read an article from 2004 at www.edn.com/article/CA409045.

 To read more about ZMD’s ZigBee project, try www.edn.com/article/CA6325290.

sensor-interface devices. Or it may mean actually driving a DAC from the digital trimming signal to bias a node.

Today, such techniques find use primarily in early silicon versions and test chips, as at AMI, or in inherently messy analog problems, such as 90-nm RF amplifiers in single-chip cellular handsets. But engineers could conceivably use these trimming techniques as aids to integration. By identifying the points in an analog block that prove most sensitive to changes in the external environment and making them digitally tunable, designers could make trade-offs of performance, power, and noise immunity after the integration process or even after manufacturing. This approach would allow a relatively naïve design team to drop in a configurable analog block, perform tape-out, and then experiment with the configuration bits on the engineering samples to get the block working. This inelegant approach might extend the range of blocks that designers can use in block assembly. Such tactics may be necessary in any case simply because of the process variations at geometries smaller than 130 nm.

So, analog blocks can be black boxes in SOC design only if they are simple or if your ASIC vendor has a back room full of analog designers. In the future, however, things may change dramatically. EDN

You can reach Executive Editor [Ron Wilson](mailto:Ron.Wilson@reedbusiness.com) at 1-408-345-4427 and ronald.wilson@reedbusiness.com.

