



COMPLEX ICs ARE NOT ONLY ABSORBING MORE OF THE SYSTEMS AROUND THEM, BUT ALSO SWALLOWING THE TEST EQUIPMENT DESIGNERS USE TO BRING UP, EVALUATE, AND CALIBRATE THE CHIPS.

AS SOCs GROW, TEST-AND-MEASUREMENT INSTRUMENTS MOVE ON-CHIP

BY RON WILSON • EXECUTIVE EDITOR

It's just geometry. As system-level ICs grow larger and more complex, they become impossible to observe and stimulate. Internal nodes aren't accessible to bonding pads or even to probes. Signal voltages are small, noise thresholds are tiny, and drive strengths are negligible. As critical circuits reach gigahertz frequencies, it becomes physically impossible to get an accurate representation of signals off the die, even if you can probe the circuit.

Yet the need remains. Chip designers must be able to observe and stimulate individual blocks in an SOC (system on chip) to bring up the silicon. Manufacturing-test engineers must be able to create fast test programs on affordable test equipment. Increasingly, chip designers must also create autocalibration routines

that can compensate critical circuits for process, voltage, temperature, impedance, and noise variations while the chip is in use. The only apparent option is to move test-and-measurement instruments—the racks of logic analyzers, bus analyzers, communications testers, and oscilloscopes that populate the bring-up lab—onto the chip itself.

And this option is now available. Beginning perhaps with debugging facilities built into CPU cores, extending through bus diagnostic blocks and logic built-in

self-test blocks, on-chip instrumentation is today extending into high-speed transceivers and RF circuits. In the future, you may see on-chip analog instrumentation for characterization and calibration routinely becoming part of analog design (see sidebar “MEMS accelerometers need instrumentation, too”).

BEGINNING IN THE CPU

The idea of building debugging hardware into a CPU dates back at least to the IBM 360 architecture. But the race

to fit more complex CPU cores into less die area meant the concept became lost in the early days of SOC design. It resurfaced because of necessity.

“As processor complexity and frequency increased, it became just too difficult to control the core through external circuitry,” says William Orme, general manager of the CoreInsight debugging program at ARM. “Designers resisted adding to the core area for debug, but it became a trade-off of die area versus pain. Eventually, in-core debug became cost-effective in the overall cost of the SOC. Then, it simply became obligatory.”

During that evolution, the basic techniques have not changed. Designers still need to put the CPU core in a known state, start it running, observe and record the resulting sequence of states, and stop the core again when something interesting happens. Hardware integrated into the core can accomplish all these things with essentially no impact on performance or energy consumption during normal operation.

But as SOCs evolved, the CPU core ceased to be the only problem. The growing web of buses—wide, fast, seg-

mented, and multilayer—that spread out from the CPU also became unobservable. So, ARM and other vendors of interconnect IP (intellectual property) built debugging circuitry into the interconnect architecture, just as they had built it into the CPU.

“In an AHB [advanced-high-speed-bus] multilevel interconnect, designers need to monitor what is going on at any level of the interconnect: the source, destination, and content of every transaction,” Orme says. “That [process] requires monitoring from inside the die.”

The situation grows even more complex as SOCs evolve from a single CPU core at the center of a bus network to multicore designs in which numerous processing sites are active at once. Now, an event may be not simply the state of a core or a bus but a complicated merg-

AT A GLANCE

- Chip designers are starting to design test-and-measurement instruments into their complex ICs.
- The movement toward designing instruments into ICs started with digital debugging hardware for CPU cores and buses.
- Now, designers are also building analysis instruments into high-speed-I/O blocks.
- Designers are integrating more complex analog and RF instruments into the inner workings of high-frequency chips, such as read-channel ICs.

ing of state sequences from a number of processors and interconnect structures in different clock and voltage domains. Just being able to start and stop such a

system and to capture some idea of its actual state becomes a challenging problem, Orme admits.

The growing complexity may be shifting the emphasis in on-chip-debugging circuitry away from the CPU core and toward a more system-level approach. ARM, for instance, offers a cross-triggering switch matrix that attempts to bring together state signals from different blocks in the SOC. And processor-independent debugging-hardware companies, such as Dafca, are also emerging.

“Design teams have been trying to create a comprehensive debug strategy by combining CPU-debug facilities with their own in-house-designed on-chip instrumentation,” says Paul Bradley, Dafca’s vice president of engineering. “But as the complexity of the SOC grows, voltage and clock domains proliferate,

MEMS ACCELEROMETERS NEED INSTRUMENTATION, TOO

MEMS (microelectromechanical-system) accelerometers, with their accuracy, low power, and low cost, are revolutionizing a number of applications. Fuses for automotive air bags, attitude and motion sensors for computer-input devices, and sensors for automotive attitude-control systems are some large-volume examples.

The acceleration-sensing elements are often differential capacitors in which silicon springs suspend one plate of the capacitor above the other two. Acceleration causes the suspended plate to shift, changing the relative capacitance of the two capacitors and, hence, the relative amplitude of signals passing through them. Properly designed and fabricated examples of such structures are sensitive, accurate, and durable. But they require calibration. Changes in applied voltage and ambient temperature can change the sensitivity of the device, and each sensor has an offset reading at 0g that the user must measure.

In high-volume applications, that requirement is not a problem, according to Bob Scannell, business-development manager at Analog Devices. If you use hundreds of thousands or millions of sensors, it’s no great cost to set up a calibration tool that connects to the sensors and spins them at known rates in a temperature chamber. But, in lower volume applications in the industrial world, it makes great sense to have on-chip instrumentation to calibrate the devices ahead of time (Figure A). “Precalibration removes barriers to use,” Scannell says.

By embedding a temperature sensor and a microcontroller in the MEMS device, for instance, the vendor can preload temperature-compensation tables and have the chip remove temperature-based variations on the fly. Similarly, firmware can take the chip through an auto-zeroing procedure to establish a reference horizontal

plane in-system, electrically stimulate the MEMS structure, and compare the output with a known reference for a quick pass/fail-operation test.

In examples such as these, on-chip instrumentation is worth the silicon investment not because there is no other way to make the measurements or even to reduce exorbitant test cost, but simply to make the device easier to use, opening new markets.




Figure A MEMS accelerometers, such as Analog Devices’ ADxl202, require significant on-chip instrumentation for bring-up, test, and in-circuit calibration.

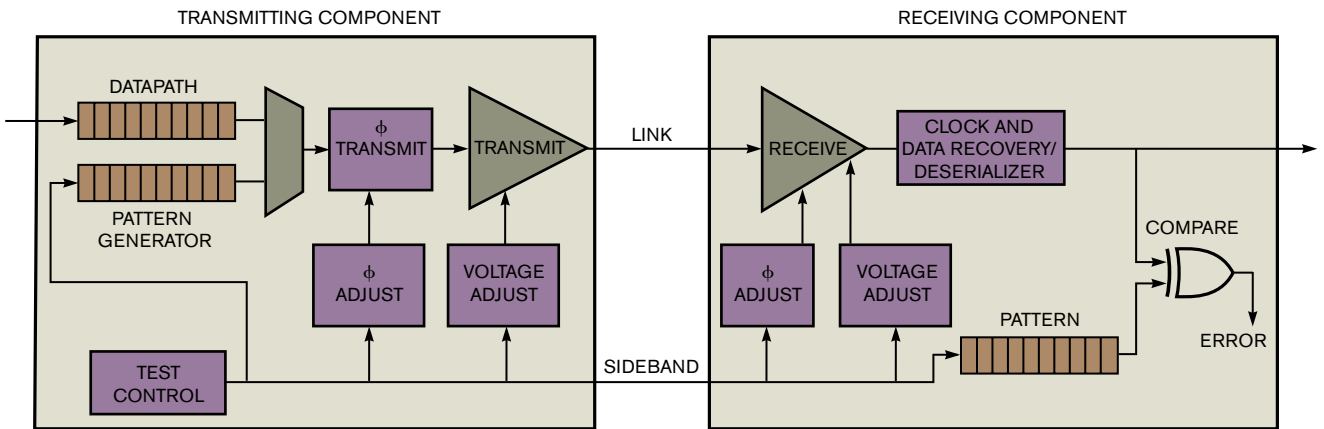


Figure 1 On-chip instruments permit a variety of test-and-measurement procedures in this Rambus serial-link architecture.

and reuse gets more important, designers need a comprehensive approach, not a collection of ad hoc designs.”

Dafca Chief Executive Officer Peter Levin adds another dimension to the concern: “As this complexity grows, so does the complexity of the debug software you need to control the hardware and make sense of the data. Today, the debug-software effort is at least 10 times greater than the effort to design and integrate the debug hardware.” Again, reuse is vital, and licensing the whole thing from a specialist looks more and more reasonable.

Dafca’s architecture gives an idea of how complex full-chip debugging support has become. Dafca works with, rather than replaces, the debugging hardware in the design’s CPU or DSP cores. But it adds “analysis instruments”—elaborate but compact reprogrammable state machines that can stimulate a block by pre-empting the block’s inputs for some number of cycles, monitor outputs of a block, and report exceptions into a global network.

“Deployment of the instruments is very application-specific,” Bradley says. “Usually, a design team will use two to six analysis instruments on a chip, often associating one with each major design domain. In this way, they can observe key interfaces and control points—usually at the second level in the design hierarchy.” Pulling together the information from the instruments and the processor cores into a cohesive picture of the operating chip and providing a level of abstraction at which designers can define and impose a chipwide state are critical needs in any such architecture.

That requirement may explain the huge software component of the design.

Digital on-chip instrumentation, then, has evolved from fairly simple debugging cores within CPU cores to much more sophisticated hardware/software systems that designers base on distributed state machines and a triggering-and-control network on the chip. As SOCs become still more complex, the next step seems to be full logic-analysis capability on the die, with programmable triggers, large trace and vector buffers, and an ability to probe nodes within the chip.

THE ANALOG DOMAIN

Analogous to the trend in on-chip digital instrumentation, there is growing interest in generating and measuring

analog signals with on-chip instruments. And if it’s fair to trace the origins of digital instrumentation to CPU cores, the analog instruments are starting out in the cores of high-speed serial interfaces. But the techniques are branching out in the analog domain, as well, incorporating sensor-based measuring systems for auto-calibration, mixed-signal servo systems, and instruments for validation and test.

As in the digital world, the driving force behind on-chip instrumentation of high-speed I/O is the growing invisibility of key signals. “The mere existence of an external probe alters a high-speed link,” says Rich Perego, senior principal engineer at Rambus. “Yet, there is value in being able to look at a waveform.”

“Often, you can probe the transmitter directly,” adds Ken Chang, design-engi-

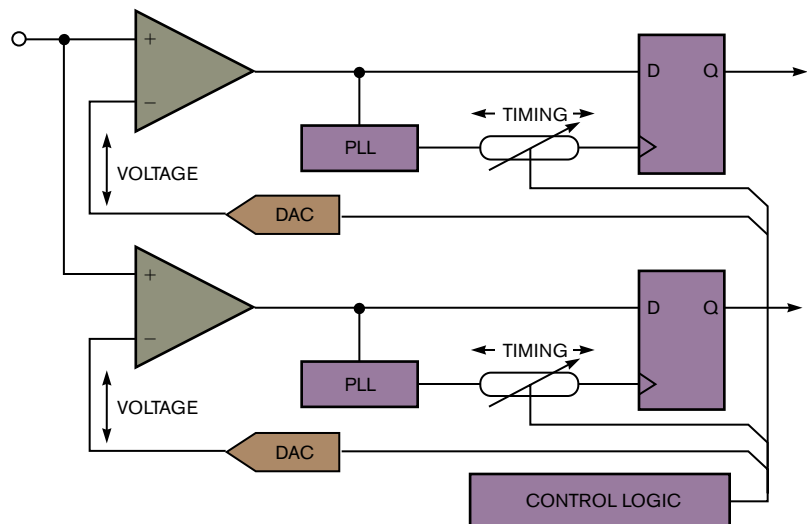


Figure 2 Vitesse has recently taken receiver instrumentation a step further, devising an interesting two-channel approach.

neering manager at Rambus. “But you can’t just probe the receiver. You have to look at the receiver indirectly by extracting data on-chip and inferring eye diagrams or bathtub diagrams.”

This approach has become increasingly important in high-speed receivers, in which it is otherwise impossible to see what is going on inside the block. You can’t see from the outside what the signal looks like just before or just after equalization, according to Eric Sweetman, senior applications engineer at Vitesse. You have to make those observations from inside the receiver.

Rambus approaches this problem by using the receiver as, in effect, its own measuring instrument. The process starts with integrating programmable pseudorandom-pattern generators and bit-stream comparators into the I/O blocks. Designers can then add enough circuitry to digitally adjust the transmitter and receiver (Figure 1). In the case of the transmitter, for instance, this method means putting digital inputs and DACs on controls for the current swing and phase and being able to modify the equalization-tap coefficients. “That [approach] allows us to create shmoo plots of transmitter performance by sweeping timing and voltage,” Perego explains.

In the case of a receiver, it might mean adjusting the receiving amplifier’s phase and gain. That process allows designers to extract timing margins, for instance, by sweeping phase and watching the bit-error-rate accumulator on the output of the bit-stream compare block.

The art of such instrumentation is to do as much as possible in the digital domain. You can digitally control receiver phase by manipulating the digital-feedback path in a delay-locked loop, for instance. Or you can adjust gain by providing a digital value to a small DAC, which in turn adds an offset to a critical node. These insertions into the circuitry are not trivial, particularly in the case of analog nodes, in which any change has the potential to break a circuit. But they are achievable.

Vitesse has recently taken receiver instrumentation a step further, devising an interesting two-channel approach (Figure 2). This idea sets one channel—usually the receiver’s primary data channel—at the center of the eye. Designers can then sweep a second, essentially identical channel across the space

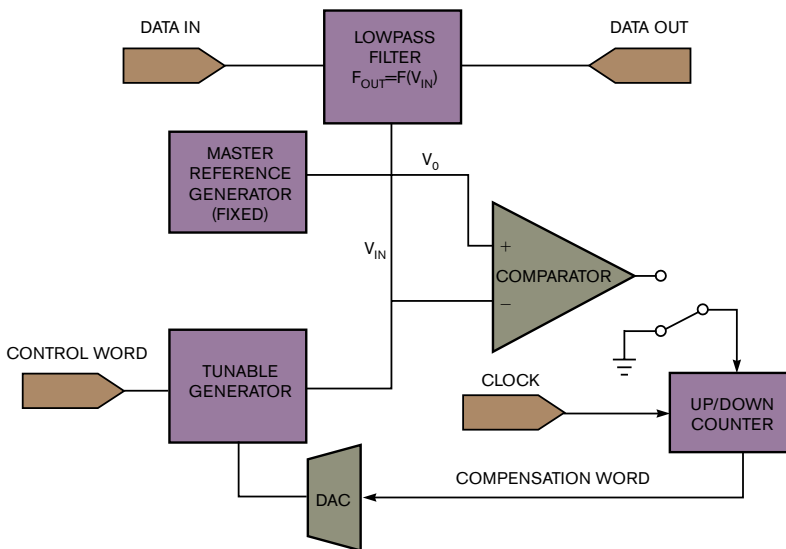


Figure 3 A simple comparator-based loop can remove drift from the voltage source that determines a filter cutoff frequency (courtesy STMicroelectronics).

of phase and amplitude points by altering the receiver’s phase and gain, stopping at each point, accumulating enough data to accurately estimate the bit-error rate, and moving on. The result after software postprocessing is either an eye diagram or a bathtub plot of data that the receiver itself collects. Because the second path is physically identical to the read path that is taking the measurements, the data are more accurate than a probe and scope could ever achieve, Sweetman says. The

THE ART OF SUCH INSTRUMENTATION IS TO DO AS MUCH AS POSSIBLE IN THE DIGITAL DOMAIN.

designer can see the exact signal that’s going into the receiver’s sampler from the equalizer, not a distorted approximation.

High-speed I/O is not the only application for such ideas, however. STMicroelectronics, for example, puts considerable instrumentation into its high-end disk-read-channel ICs. Again, the emphasis is on using the circuitry as much as possible to perform measurements under external control. But, in this case, the added circuitry can be complex.

Angelo Dati, architecture manager for the data-storage division at STMicroelectronics, says that the internal

instruments within a read-channel chip assist with initial troubleshooting and calibration of the chip; help characterize the complete media/electronics system once the chip is installed; and provide continuous compensation for voltage, temperature, head-height, and other variables during operation.

Accordingly, the read channel has a variety of measuring instruments. Simple sensors for temperature and voltage, for instance, run to comparators and linear filters, which the controlling processor can interrogate to determine whether to launch a calibration procedure on a local finite-state machine. The state machine runs a closed-loop calibration procedure in which it can alter voltage offsets, gains, and bias points to correct problems, such as the tendency for the channel’s lowpass-filter-cutoff point to drift with temperature.

More sophisticated measurements not only compensate for variations during operation, but also help designers and production engineers tune the chip to a particular head/media combination. One of these measurements sees use in

MORE AT EDN.COM

For more on on-chip-debugging hardware, go to www.edn.com/article/CA6451246.

Go to www.edn.com/080221df and click on Feedback Loop to post a comment on this article.

an adaptive-loop system based on pattern matching, Dati says. Instrumentation within the chip generates an analog waveform that the system feeds into the read circuitry and compares with the incoming signal from the read head, giving a frequency-error signal. This signal provides feedback for adapting to the noise and nonlinearity characteristics of the head/media combination. The process is closed-loop. But, Dati says, "There is no

proof of convergence. We have to seed the adaptive loop with a good starting value for it to work."

Still another analysis is useful in silicon bring-up and in manufacturing for performing harmonic analysis on the drive assembly. The chip actually contains what Dati describes as "a coarse spectrum analyzer—like an FFT engine but simplified to look for some specific frequencies, so it is simpler than a gen-

FOR MORE INFORMATION

Analog Devices
www.analog.com

ARM
www.arm.com

Dafca
www.dafca.com

IBM Corp
www.ibm.com

Rambus
www.rambus.com

STMicroelectronics
www.st.com

Vitesse
www.vitesse.com

eral-purpose spectrum analyzer." The instrument's harmonic analysis has a number of valuable purposes. By looking for specific frequencies that shouldn't be there, for instance, you can detect and measure nonlinearity in the head/media subsystem. And, by looking at the frequency envelope, you can estimate the magnetic spacing—the height of the head above the media. "With today's drives, you have to do that," Dati explains, "because if you don't measure and adjust the head height, ordinary differences in atmospheric pressure could cause a head crash."

A spectrum analyzer might seem like a large instrument to tuck inconspicuously into a highly cost-sensitive chip. But in reality, "receivers are getting so big that it's easy to hide some instrumentation circuitry in the design," Dati says.

That observation may be a good introduction to the future of on-chip instrumentation. As chips get more complex, there will be not only greater need for instruments on the die, but also more opportunities to exploit the functional circuitry to quickly move analog measurements into the digital domain. And there will be more space in which to stash even a fairly complex block, such as a modified FFT engine.

The next step is probably to follow what has been happening on the digital side: Pull together analog measurements from both the time and the frequency domains from different portions of the chip to provide a picture of the full state of the mixed-signal system. Only with such on-chip tools may it prove possible to debug the next generation of complex chips and to keep the generation after that running at all. **EDN**

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

