Volume 40, Number 1, 2006

Analog Dialogue

A forum for the exchange of circuits, systems, and software for real-world signal processing

In This Issue

40 years of Analog Dialogue 1967 2006

www.analog.com/analogdialogue

ANALOG DEVICES

Editors' Notes

A READER COMMENTS

Last fall, we published an article about adding stereo audio to a satellite set-top box. The key idea was to apply a phase-locked loop, with the vestigial 15.734-kHz pilot signal in the composite spectrum as the reference, allowing one channel of the AD71028 BTSC encoder to derive the primary channel's master clock while the other channel provides the MTS stereo-encoded output for satellite set-top boxes and receivers.

We received the following email from Matt Laun, at NASA:

> "I found your design using the AD71028 as a STB (satellite set-top box) stereo synthesizer clever and interesting. However, I can suggest two corrections to the information in the article. One is the RF spectrum of a modulated NTSC video and audio signal. It appears in your article as the following:

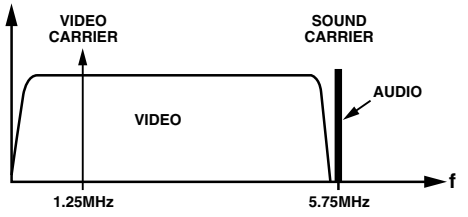

> The frequency information may be misleading to a reader not familiar to the standard. More commonly, the audio carrier frequency is measured relative to the channel carrier frequency or the video carrier frequency. The audio subcarrier would be 4.5 MHz. It is more intuitive to view the spectrum as seen below:

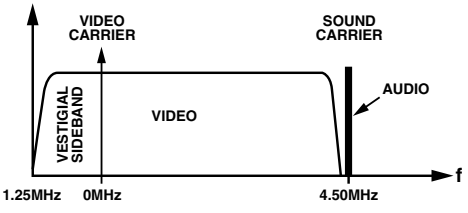

> The second correction is minor but may be significant if a reader is trying to understand the design by performing calculations. When describing the need for a precision N-divider, the value N is listed as 780.9838... The precision value is actually 780.9706666... { = 48 kHz × 256/15.734265734 kHz}; [fractionally: = 49201152/63000]. I understand the clever design eliminates the need for a separate precision divider, I just thought a more accurate number would help a reader understand how that number is derived.
>
> Thanks for a relevant, practical, and interesting article."

Lead author, Jeritt Kent, responded:

> "First of all, thank you very much for this kind letter. I am very pleased that you enjoyed the article. Victor and I worked hard to ensure that the paper was clear and easy to understand.
>
> Your two clarifications are very valid. The first one was likely an interpretation by the graphic designer, as the original intended figure would have matched Fig. 2 in this document http://www.sencore.com/custsup/pdf/TT213.pdf
>
> The second observation shows an example where rounded mathematics can generate problematic results, albeit this is not a high volume Intel Pentium processor design nor the Hubble telescope:) Nonetheless, a Google search on <15734.265734 Hz> provides some excellent history on the subject from the University of Victoria, Canada (UV)." (www.ece.uvic.ca/~yli/misc/discussion.htm)

That UV document, titled "44100," was a fascinating bit of history. An abbreviated version follows:

> The original appearance of 44.1 kHz was with the Sony F1 digital audio recorder. It recorded digitized stereo audio as an emulated video signal onto Betamax tape. **You might recall that the horizontal scan rate for NTSC video is 15750 Hz (actually 99.9% of that but we'll get to that later). This comes from 30 frames/sec × 525 lines/frame. Note that 44.1 kHz is exactly 2.8 times 15.75 kHz.**
>
> The vertical blanking normally done on video was not used, three stereo samples were recorded on each horizontal line, and every fifth line had a checksum or some kind of redundancy sample for error detection and correction. With 14 samples for every five lines, you get a ratio of 2.8. There was considerable resistance by the AES and the audio community to have 44.1 kHz adopted as any kind of standard. We wanted 48.000 kHz because it was nicely synced to movies and both U.S. and European TV and, being a little higher, gave us a little more transition band for the antialiasing and anti-imaging filters (and we really needed that transition band back before the day of sigma-delta converters), but Sony and Philips had pretty big guns and dictated the standard since they owned the technology.
>
> Why would they not let it be 48 kHz? Rumor has it that the president of Sony wanted *all* of Beethoven's 9th to fit uninterrupted on a single CD. The 12-cm diameter had already been carved into stone, and there was no way for it to fit with a 48-kHz sampling rate. I don't know how long Beethoven's 9th is, but a CD can hold 74 min, 33 sec of music.
>
> **The frame rate (and horizontal scan rate) was actually reduced by 0.1% due to the emergence of color TV. Originally, it was exactly 15750 Hz and the sound subcarrier was at exactly 4.5 MHz.** Color TVs encode R, G, & B into a B&W luminance signal called Y, and I&Q (in-phase and quadrature) chroma signals. Y is transmitted just like a normal B&W signal. I&Q are transmitted using quadrature carrier modulation, and are bumped up to exactly halfway between the 227th and 228th harmonic of the horizontal scan rate. There is not much B&W energy up there at the 227th harmonic, so they could stick the chroma up there without degrading the B&W image too much. In addition, putting the chroma exactly at 227.5 times the horizontal rate (and all of the harmonics of the chroma would be halfway between harmonics of the B&W signal) would cause any interference from the chroma signal to be exactly negated on the neighboring scan line. The same is true if you consider the effect that the B&W signal has on the chroma signal: the B&W signal is at the –227.5th harmonic of the chroma baseband, so its interference is negated with every other scan line. The sound subcarrier is a lot closer to the chroma signal than to the B&W signal, so the color TV guys had to worry about it.
>
> **The sound carrier is at exactly 285.71428571 times the horizontal. The color TV guys wanted it to be an exact integer multiple (so it would be 1/2 harmonic off from the chroma signal). So, they decided to make the sound carrier exactly 286 times the horizontal. Unfortunately, instead of bumping up the sound carrier to 4.5045 MHz (these guys in the '50s thought that this would be too far off for existing FM receivers to lock to), they bumped the horizontal scan rate down to 15734.265734 Hz. This caused the horrible 29.97 Hz drop-frame mess that today's technicians have to deal with and the 44055.944056 Hz sampling rate.**

By the way, 15734.265734 is a truncation of 15750/1.001, or 15734.2657342657..., while 15750 × 0.999 = 15734.25. Now, here is the last word from Matt Laun:

> "Thank you too for the fascinating Google '.doc' about the origins of 15734.265734! I am a recording engineer as well as an electronic hobbyist with a nostalgic love for analog NTSC and BTSC modulations. I never knew that audio 44100 kHz was derived from video (and not even standard NTSC sync at that)! Incredible!"

If only Beethoven had known! A few bars less, and he might have simplified life for designers in the Age of Media.

Dan Sheingold [dan.sheingold@analog.com]

SIMULATED REALITY

In addition to three outstanding feature articles, this issue includes the second installment of Barrie Gilbert's fantasy. In it, Barrie introduces an example showing the power of circuit simulation. In a short time, Niku was able to gain valuable insights regarding the behavior of an ideal oscillator and to debunk the commonly held belief that it would start up given a spike on a supply or bias line.

In this column, Dr. Leif has provided some clues regarding the fourth "Dee" of Analog. We invite our readers to guess what Barrie has in mind, and to provide anecdotes from your experiences with simulators or with analog design in general.

Scott Wayne [scott.wayne@analog.com]

ISSN 0161-3626

THE FOURTH DEE: TURNING OVER A NEW LEIF

By Barrie Gilbert [barrie.gilbert@analog.com]

At the time Niku Yeng received an official offer of employment from Analog Devices in 2025, shortly after her interview with Dr. Leif, she had briefly considered two others. All held the prospect of an exciting and rewarding career in advanced microelectronics. But she recalled that the interviewer at one of the companies seemed to be unusually concerned about her willingness to accept highly detailed directives and to rapidly produce solutions in response to specific market demands. She'd been trained to expect this, in an industry that had become just another provider of commodity items—much like hyperphase foods or disposable clothing—she nonetheless felt that a strong emphasis on products focused only on near-term needs was myopic, and it was bound to discourage originality and clash with any aberrant, singular vision, leading to mediocrity and a poor reputation for quality and service.

The third offer was quite different. During that interview, all the questions were concerned with very tentative technologies, facets of the ongoing struggle to make nanostructures not only as crude logical elements but in the vastly more complex arena of analog design—where device quality is of paramount importance. Very good progress had been made in applying hybrid neural networks (such as those in the greeter at Galaxybux) having quasi-analog circuits on silicon substrata acting as message concentrators for the layers of super-stressed fibers of dimethyl-3, 5-ribocarbon—which provide the fast-learning, slow-fading memory cells with massive parallelism for noise immunity and redundancy. But the broad promises of the nanodisciplines, peaking just after the turn of the century, had all but evaporated between 2012 and 2018, as more pragmatic concerns dictated which technologies could deal with the peculiar demands of analog signals. Furthermore, with atmospheric CO_2 rising at an imminently perilous rate, research in global climate control now held center stage in every country. Degree courses in this field had become the first choice of many a young scientist who aspired to a place in history.

But, more than any other factor, it was the sheer enthusiasm of Leif, his keen interest in exploring some still-unanswered, though seemingly rudimentary, questions about analog circuits that led Niku to accept ADI's offer for a position as *Entry-Level Product Originator*. At first, she was upset to discover that her work would not be under Leif—who was, it seemed, a sort of roving consultant to various groups in the company. (She later discovered that he got involved in entry interviews only in cases where exceptional talent was evidenced during prescreening, a revelation that had helped to assuage her initial disappointment.)

By now, though, she felt sure she'd made the right decision. Leif was no mystical hermit. He regularly wandered around the local design groups, asking penetrating questions about their projects, vigorously engaging in theoretical issues here or offering advice there, and applying a deft technique of asking leading questions that left no one feeling put down. It was on one of these routine visits, roughly two months after Niku had joined her team, that he sat down with her for a while. Before long, she was reminding Leif that he had never gotten around to telling her what the "Fourth Dee" of analog design was all about.

"Ah, yes, those 'Dees.' What do you remember about them?" he quizzed.

"Well, you told me that analog products are far more **Durable** than digital, often having generations measured in decades; and that the little circuits that go into them are highly **Diverse**—like the myriad musical tunes composed out of just a few notes; and that their constituent components, as well as the actual signals, have a crucial **Dimensional** aspect."

"Hearing you summarize them so well, Niku, they don't appear to be especially profound, do they? Are you finding these issues to be as important as I have suggested?"

"Not really. I believe you, of course. But any truths of substance have to be learned, from one's individually acquired knowledge and hands-on experience, rather than accepted simply on their face value," Niku said, sounding wiser than her years. "Are you going to tell me what the Fourth Dee is now?"

"Tell you what," said Leif mischievously. "I dropped by to see what you've been doing, and I'd first like to hear all about that. Perhaps the answer you're seeking will occur to you by the time you have finished telling me about your project."

Niku explained that she hadn't yet been assigned a development project. She'd been given time to familiarize herself with the lab environment, the many in-house and foundry IC processes that would be available to her, and the vast network of databases and design centers in every corner of the United World. She had also been familiarizing herself with the large suite of simulation tools called GE°E, pronounced "gee-oh," which she was told stood for *general electro°ptical emulator*. She had never been exposed to anything so powerful and all-encompassing, let alone so user-friendly and fast.

"So how are you getting along with your teammates?"

"Oh, they're okay." She briefly hesitated—but then mentioned that several days ago she'd heard a noisy argument in the halls about exactly how high-frequency oscillators start up. She was quite surprised that this was a matter for disagreement.

"It appears that Bob somebody—he has a CyberCyte" (Leif's wry smile acknowledged his recognition of this character) "had stated in no uncertain way that oscillators start only because of a sudden step on the supply, or on a bias line. A couple of the guys seemed to agree with him, but most didn't. They argued that if a circuit satisfies the conditions required for sustained oscillations, it should only depend on its internal noise to start, and the longer it takes, the better. The argument got pretty fierce at times!

"So I decided that studying this question in depth would be an excellent learning experience, as my first serious exercise in the use of GE°E. But I have told no one about this, because it might seem a waste of the company's time."

"Absolutely not, my good young lady! If you ever start to believe that the exploration of such fundamental questions is a waste of time—even *after* you have product responsibilities—you'll hear from me! Do you know what *genius* is?"

Niku blinked, startled by Leif's sudden intensity.

"Genius is nothing more than this: The curiosity of childhood constantly recaptured—every day of your life! It wouldn't be proper for me to advise you to allow your curiosity to rule your actions when you become faced with urgent deadlines. But if those deadlines should ever become the constant feature of your life, you'll feel frustrated by the pressure, for perhaps a year or two, and later miserable and irritable. Eventually, your precious flame of individualism will be fully extinguished. Years ago, this would happen to a fine product designer, who would turn into a designing robot by the pressures of the work. Today, we have a better awareness of these dangers. In fact, it is my job to ensure that creativity flourishes in this group, by defending flights of the imagination—such as the one you are starting to tell me about!"

"Okay, but I'm no genius! Just very curious, simple-minded, and terribly inexperienced," she said, blushing deeply. "Well, to start

with, it seemed obvious that if an oscillator actually did start up through being disturbed by a supply transient, it would be a pretty poor circuit!"

"How does that follow?" said Leif, knowing full well where this was going.

"Well, the guys in the hall were talking about RF oscillators, for which phase noise is a critical performance issue; and if such an oscillator is easily upset by supply noise—enough to induce it to start up from cold—then I reasoned it would not be in the league they were talking about. It was only a short step to conclude that oscillators for demanding applications, as in this TransInformer" —Niku put her PDA on the table— "must use a *fully balanced* cell topology, if for no other reason than to reject the common-mode noise voltages, but also to minimize the even harmonic terms."

"Splendid! These are quite remarkable leaps of the imagination! I'm beginning to think that the Fourth Dee may not be one you need to worry about! But—keep going!"

"Well, it just seems like common sense to me. Anyway, I wanted my test cell to be as simple as possible, to reduce the number of unknown influences, so I used this ..." Niku pulled up the circuit on the screen of her PDA (Figure 1). "I know it's not a practical oscillator. For example, I learned in one of my courses at Nova Terra that once it does start oscillating, the amplitude will build up until the transistors start to saturate, and then the frequency plunges. I was interested not only in whether it will start under disturbed conditions; I also wanted to verify that the circuit noise is important to startup, and to learn *exactly how* this process unfolds over time—the oscillator's start-up trajectory, I suppose you'd call it. And I wanted to discover the relationship between the tail current required to sustain oscillation and the size of its resistive load, and ..."

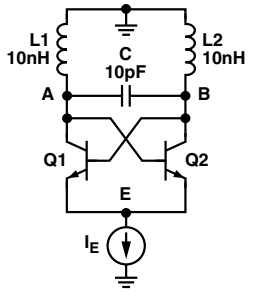

TO MINIMIZE THE INTRODUCTION OF DISTRACTING COMPLICATIONS OF THE KIND FREQUENTLY ENCOUNTERED IN PRACTICAL OSCILLATORS, Q1 AND Q2 ARE INITIALLY ALLOWED TO BE "IDEAL"; THAT IS, BF = BR = VAF = VAR = INF. AND THE DEVICE RESISTANCES AND CAPACITANCES ARE ZERO; τ_f = 10ps.

LIKEWISE, FOR THE INITIAL EXPERIMENTS, THE TANK IS ASSUMED TO BE LOSS-LESS AND WITHOUT A LOAD.

I_E IS TURNED ON VERY RAPIDLY, AND THE CONSEQUENCES ARE OBSERVED FOR A VARIETY OF CONDITIONS, WHICH WILL BE ELABORATED ON IN DETAIL DURING FURTHER DISCUSSIONS WITH DR. LEIF.

Figure 1. Niku's First Basic Experimental Oscillator. Note that the only "supply" is the current I_E, further minimizing sources of enigma.

"Niku! Whoa!" said Leif, again looking rather serious. "Are you aware that you have set forth a series of studies—solely for your own enlightenment and pleasure—on a complex topic that others have regarded as a sufficient basis for a thesis degree?"

"Oh, not really. I didn't expect these virtual experiments to take very long, using GE°E. As it turned out, these studies brought to my attention a long, connected sequence of questions, as I saw various effects coming into play—some quite puzzling at first—and I wanted to explain them all. I have written them up, in case anyone else might be interested," said Niku.

"I have no doubt of that! May I put your name into our schedule of Daedalus Days?"

"What's a 'Daedalus Day'?" she giggled.

"Oh, I don't want to break your train of thought right now, but we will certainly get back to that, sometime," said Leif.

"Okay. Let's see. Oh yes! I felt it would be a good idea to further minimize the unknowns by using *ideal* bipolar transistors. I knew that, as long as the fundamental shot noise was modeled—and of course the BJT's beautifully straightforward transconductance—then, including the realism of the complex full transistor model would add nothing to help me gain the insights I was looking for. So I set the junction resistances and capacitances to zero, and the forward and reverse betas, as well as the forward and reverse early voltages, to infinity. Everything else used default values; except that, even though I wasn't interested in exactly modeling the base charge terms, I included a τ_F of a few picoseconds."

"That leaves very little of the reality, Niku! Are you confident that these drastic simplifications can be justified?" asked Leif. But he was not frowning, only putting her to the test.

"Yes, I think so. Originally, these experiments were intended to demonstrate only that an exactly balanced, *noise-free* oscillator will not start up when the tail current switches on suddenly, even if its rise time is less than the tank period, and even without any load. I also had a hunch that it wouldn't start up if I introduced a deliberate imbalance, provided the rise rate of the tail current was below a critical value, which I wanted to quantify; and certainly not with a load resistance below a critical value across the tank. By the way, I could have used two equal and separate tanks as loads, but that would introduce one more capacitor and another degree of freedom in the behavior."

"Good thinking. So, how did you upset the perfect balance?" asked Dr. Leif.

"I just altered the relative size of one of the transistors, using the SIZE parameter," she explained. (Note: Although GE°E is a far cry from SPICE, a surprising number of its commands, variable names, and other parameters can be traced to that earlier era.) "And, Dr. Leif, I wanted to add that as these studies progressed, the circuit opened up its many secrets to me; and I was glad I'd chosen to use primitive models because, even with these, there were times when I had to think hard to explain what was going on. It's safer to add in the additional reality of the full transistor model in small steps. Then you can see precisely at what point some puzzling new phenomenon first appears."

"Yes, many of us appeal to that paradigm, particularly when we are exploring a novel cell topology. It was called 'Foundation Design,' about 50 years ago, by one of ADI's Fellows. Well, now that you have whetted my appetite, Niku, tell me: How did you start your journey, and what did you discover first?"

"The first thing I did was to demonstrate that the application of a 10-mA tail current, I_E, having a 1-ps rise time, would never start this perfectly balanced oscillator under any of the conditions I tested. Of course, such a shock probably *would* be the primary reason for startup in a real circuit, which is always unbalanced, to some degree. But bias currents don't appear this quickly!"

"Now, from what you are telling me, I gather you are running GE°E in its primitive mode, as a SPICE emulator; because none of today's circuit simulators will nicely leave such an oscillator circuit in its meta-stable condition. By the way—*why* is that?"

"Oh, I know what you're getting at! Yes, that's correct. I chose to run the initial simulations in the old SPICE mode because I wanted to temporarily eliminate real-world noise processes. The SPICE-based simulators of, say, 2005 could predict small-signal noise values quite well, provided the device models were correct. However, SPICE only 'knows about' noise in a numerical sense, and merely handles the math to add up all the numbers. It has no idea about noise as a *process in time*—it does not treat the noise mechanisms in the various elements as a set of *time-stochastic variables*, whereas GE°E does."

"Precisely! Good. I assume that you used tight convergence tolerances to ensure the simulator wasn't simply stuck inside a broad numerical tolerance range. And I guess you chose 10 mA simply as a representative tail current for this type of oscillator. Okay, so at this juncture, you felt justified that those 'start-by-spike' fellas were incorrect?"

"Oh, no; it was just the first step. I really wanted to demonstrate that, in practice, the noise voltages across the tank at resonance would be the more significant source of disturbance, and that in a real circuit, with or without mismatches, noise is the root cause of the start-up trajectory. In fact, during my CyberFind studies, I turned up an article in *Analog Dialogue* about this, going back to 2006. It was very helpful. But I had to find out for myself."

Leif smiled with a mixture of approval and growing affection for this unusually curious and perceptive young mind.

"My next step was to introduce a 20% mismatch, equivalent to about 5 mV of V_{BE} difference, by giving Q1 a SIZE factor of 1.2 and again pulsing the tail current. Clearly, if this current appears very rapidly, with a rise time similar to the oscillator's period, it is bound to generate a voltage change across the tank, and even the slightest disturbance will get things going. So, I thought it would be interesting to ask how large that voltage would be."

Her mentor struggled to be ready with a quick calculation, in case Niku asked, "Do you know what I found?" Alas, it wasn't immediately obvious to him how to figure it out. With his eyes loosely closed in concentration, he could have been asleep.

"Dr. Leif? I said, 'Do you know what I found?'"

"Well, let me see now," he replied, still not having the answer he had hoped would come to mind. "The amplitude of the initial step of differential voltage across the tank, labeled V_{OUT} in your sketch, must be proportional to the step in tail current, I_E, and to the 20% mismatch—which gives us a factor of 0.2 times I_E for the current step into the tank, which is 2 mA. But then, the load resistor complicates things ..."

"No, wait! The difference current is $[(1.2/2.2)I_E - (1/2.2)I_E]$, and that's only 0.909 mA," she corrected him. "And remember, these initial studies assumed an *open-circuit* tank. Also, my idealized circuit assumed no other losses in the tank inductors, the capacitor, or the transistors. But I reduced this to an even more basic question. Setting aside the active circuit, and the effect of its power gain, what happens if an almost instantaneous step of current is applied to an LC tank? What is the voltage waveform just after t-zero? And since there is no damping, that question is equivalent to asking: 'What is the amplitude of the undiminished ringing, a true oscillation at $1/2\pi\sqrt{(LC)}$ in the tank voltage?'"

The situation was suddenly reversed. Leif was likely to have said, "I dunno. What is it?" even if he *did* know; he understood that the art of teaching must necessarily involve a great deal of humility, and allow students to feel the full glow of their proud discoveries. But the fact was, today he didn't have a clue. He honestly replied, "Niku, you have a way of asking the darnedest questions! This one is simple, as are all the best questions; but it's not one I've thought about before. Classically, it would be solved by using Laplace transforms. But experience teaches us that there are often more direct ways of seeing **'What Must Be,'** just by thinking about the Fundaments."

There was that word again. "I have to tell you, Dr. Leif, that it was your passionate concern for what you call the Fundaments that most excited me during the interview! I want to spend my life thinking afresh about the Fundaments, as it is evident you have. Well, I confess that at first I cheated! I used the simulator, and this is what I found." Niku pulled up a waveform on her PDA, showing what happens when a current step of 0.909 mA with a 1-ps rise time is applied to a parallel tank of 20 nH (that is, L_1+L_2 in Figure 1) and 10 pF (C_T). The voltage immediately assumes a steady sinusoidal form, with a continuous amplitude of 40.65... mV (top panel, Figure 2). Furthermore, over many periods, the amplitude remains within much less than one part per billion, at 40.651715831... mV.

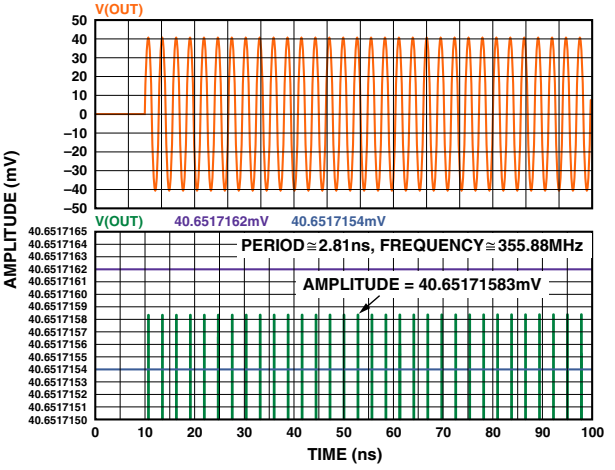

Figure 2. Result of applying a sudden step of current of 0.0909 × 10 mA directly to a tank of 2 × 10 nH and 10 pF. Niku's calculated amplitude is 40.65171583 mV. The vertically expanded trace (lower panel) showing the tips of the sine wave fully validates her theory. The added marker lines are at ±1 part per billion. This result incidentally illustrates the excellent conservation of charge provided by the GE°E simulator.

"GE°E gave me this result in less than a second," Niku enthused, "but my immediate instinct was to ask: 'Where does this funny number come from?' It implies that the tank presents a rather low impedance of (40.651... mV/0.909 mA), or 44.72135955... Ω, which is just another funny number. But doesn't a parallel-tuned tank exhibit an infinite impedance at resonance? And this tank was manifestly *resonating* in response to my stimulus!"

"Excuse me, Niku, but my TransInformer has just reminded me of a meeting, so I'm afraid I will need to leave in a few minutes. But before I go, I would like to say something about this notion that using a simulator is 'cheating.' Mathematicians once used to scorn 'numerical methods' as a way to gain insights, or to prove theorems. And any engineer who relied on 'computer-aided design'—in other words, simulation—to gain an understanding of circuit behavior was regarded by some as weak-minded and poorly equipped. But for decades we have viewed such methods in a very different light.

"Circuit designers once had to rely entirely on mathematics—and on their slide rule, pens and paper, and erasers—working through the night, fueled by endless cups of GalaxyBusters, because that was the *only way* of getting all the calculations done—like the way in which our transmobiles used to have four wheels and an engine that bravely managed to convert tens of thousands of explosions per minute into forward motion at 150 kilometers an hour on the old nonautomated MainWays. We simply didn't have anything better back in the 20th century.

"But we grew out of these things. Today, we no longer speak of computer-*aided* design, because so much of the old drudgery of calculation and optimization is managed by resourceful systems like GE°E. The equations in a modern simulator represent, in every important respect, an almost-perfect analog of the reality—whether a new molecule, a space elevator, or a clever circuit—and this allows us to examine numerous boundaries and optima to serve the *immediate* needs, while at the same time allowing us to gain

(continued on page 22)

Switching in USB Consumer Applications

By Eva Murphy [eva.murphy@analog.com]
Padraig Fitzgerald [padraig.fitzgerald@analog.com]

The *universal serial bus* (USB) has become a dominant interface to fulfill the ever increasing needs for rapid data transfer between end devices—for example, downloading and uploading data between PCs and portable devices such as cell phones, digital cameras, and personal media players.

CMOS switches can be used for connecting and routing data lines in USB systems. By selecting suitable switches, designers can significantly shorten design cycles by enhancing existing designs rather than developing new ones. In this article, we describe the USB, then go on to explore the crucial role of switches in improving performance in applications such as portable media players, cell phones, and wireless pen drives. We also show how key parameters of the switches affect the overall system design and discuss basic design challenges, such as the trade-offs between meeting bandwidth requirements and minimizing signal reflections. Additionally, we suggest how to maximize the opening in eye diagrams by careful board layout.

What Is USB and Why Has It Become So Popular?

USB has become the most popular standard for PC-to-peripheral communication in the world. Keyboards, printers, data-storage devices, and mobile phones are among the many peripherals that can be connected to a PC, employing the USB standard. Devices that previously used serial ports and parallel ports are migrating to USB, while designers of devices such as hard drives and digital cameras are often choosing USB over other standards, such as FireWire or serial-port communication. Connectivity to mobile phones, MP3 players, and game consoles is another recent development.

USB's main attraction is the ability to *plug and play*. The device is plugged into the PC, recognized by the PC; then, after the first installation of appropriate software, the device will always be recognized by the host PC—a user-friendly handshake.

The *USB Implementers Forum, Inc.*, an industry-standard-generating body sponsored by leading companies from the computer and electronics industry, lays down the standards for USB. Device designs can receive USB certification—and use the USB symbol on a product, but only after passing very strict software and hardware tests. This ensures that all USB-certified devices, whether PC or peripheral, will function correctly when interconnected, from the standpoints of both software and hardware. The standard ensures that all certified software routines, connectors, cables, signal drivers, and receivers comply, ensuring interconnnectability (Figure 1a).

Figure 1a. USB devices: a port expander, a pen drive, and a webcam.

USB is based on a serial master-slave architecture. In general, the PC is the master, known as the *host* (Figure 1b); it controls the transaction. The slave, known as the *peripheral*, tells the host its bandwidth requirement, and then a data transaction is initiated. A complete sequence of normal USB events includes these steps:

1. Peripheral plugged into host (initiating USB event)
2. Handshaking (peripheral identified, bandwidth allocated)
3. Bulk data transfer (e.g., to printer), or peripheral polled (mouse)
4. Peripheral disabled by host
5. Peripheral disconnected

Figure 1b. Typical host- and peripheral USB devices.

The hardware used in a USB system transmits data using a 2-wire (plus ground) differential bidirectional system. The data lines, D+ and D–, transmit the data as shown in Figure 2. Data can only be transmitted in one direction, so in one instance the host transmits while the peripheral receives, and then the peripheral transmits while the host receives. The USB standard also includes a 5-V power line. Generally used to power downstream devices, it obviates the need for batteries in low-power devices such as USB pen drives, webcams, and keyboards.

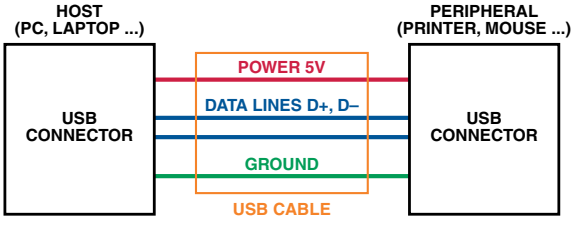

Figure 2. USB interconnections.

How Do USB 1.1 and USB 2.0 Compare?

The USB standard specifies three data rates: Low-Speed (1.5 Mbps), Full-Speed (12 Mbps), and Hi-Speed (480 Mbps). USB 1.1 devices have ±3.3-V signal levels and can operate at low- and full speeds. USB 2.0 devices have ±400-mV signal levels and can operate at low-, full-, and high speeds.

Table I. Comparison of USB 1.1 and USB 2.0

	USB 1.1	USB 2.0
Symbol		
Nomenclature	Low-/Full-Speed	Low-/Full-/Hi-Speed
Bit Rate (Mbps)	1.5/12	1.5/12/480
Single-Ended Amplitude	0 V to 3.3 V	0 V to 400 mV

What Is USB On-The-Go (USB OTG)?

Many consumer products—such as cell phones and digital cameras that connect to the PC as a USB peripheral—can also be connected to other USB devices. Since, in these circumstances, the PC cannot be the host, one of the peripherals needs to take on the responsibility. USB OTG defines a *dual-role device*, which can act as either a host or a peripheral—and can use the same connector for both PCs and other portable devices.

By enabling basic functions between digital devices, USB OTG makes these PC peripherals more capable, hence more valuable to consumers and corporate users. USB OTG devices will, of course, connect to PCs, as they comply with the USB 2.0 specification. In addition, they will have limited host capability to allow them to connect to a targeted set of other USB peripherals. When two dual-role devices get connected together via a cable, the cable sets a default host and default peripheral. If the application requires that the roles be reversed, the Host Negotiation Protocol (HNP) provides a handshake to perform that function, a reversal completely invisible to the user.

What Are the Switch Requirements for USB 1.1/USB 2.0?

The USB data lines, D+ and D–, can be connected and routed by internal CMOS switches. In the example of Figure 3, a switch is connected in series with each data line. Additional switch capacity is available by using multiplexers. Low-, Full-, and Hi-Speed USB use a 45-Ω system; the driver has a source impedance of 45 Ω, and the receiver has a termination of 45 Ω to ground. All USB cables and tracks should have a single-ended impedance of 45 Ω to preserve signal integrity. We will discuss transmission-line impedance and board layout later.

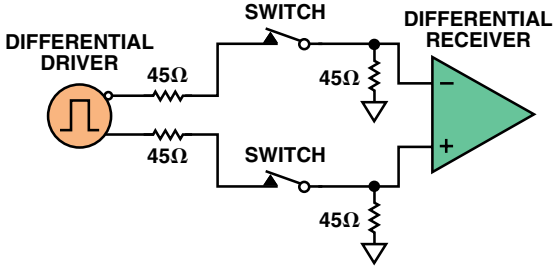

Figure 3. USB 45-Ω system.

USB standards call for stringent tests to ensure that signals are handled in conformance with their requirements. One of the key tests is an "eye" diagram. This is an intuitive visual test, which can tell a lot about the signal's quality. An eye diagram is generated by probing a randomly varying digital signal, plotting it vs. sweeps of one or more cycles, and setting the scope for long persistence. The result is that all possible bit permutations are overlaid on a single view, showing the range of deviations from an ideal "eye" pattern in amplitude, phase, and rise- and fall times. Hence any bit patterns that could cause problems may be seen on the plot.

Figure 4, taken from the USB-IF spec, shows the setup used to establish the eye diagram. The "SQiDD" (*signal quality drop/droop*) test board, which the USB-IF distributes, functions as a host; and the mouse (the device under test) is plugged into this board. The signals D+ and D– are probed and then overlaid on the scope, generating the eye diagram. The eye opening is then compared with a mask of the required shape to allow the viewer to see if the signal quality complies with the USB standard.

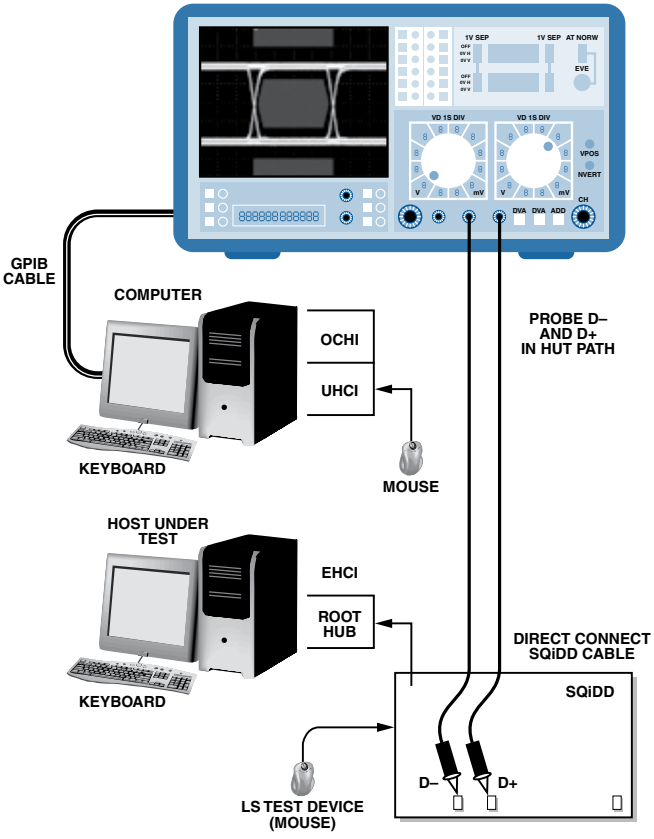

Figure 4. USB IF recommended eye-test setup.

In testing the suitability of CMOS switches for use in USB products, they cannot be tested by themselves as USB devices, since they are used within the device in the signal path. Therefore, a data generator could be used to generate the required signal, and this signal, passing through the switch, is terminated at the scope. The scope is triggered using an external clock, which is synchronized with the random digital signal. This will result in an eye diagram of the CMOS switch.

For example, a set of typical eye diagrams is shown in Figure 5, generated at USB Hi-Speed data rates (480 Mbps) and signal levels (0 to 400 mV). They compare the performance of ADG774A[1] (bandwidth >500 MHz) and ADG736[2] (200-MHz-bandwidth) CMOS switches, passing the same signals. Included in the plot is a USB-IF mask (red hexagon). According to the USB spec, if the signal crosses the boundaries of the mask, the device fails on signal integrity.

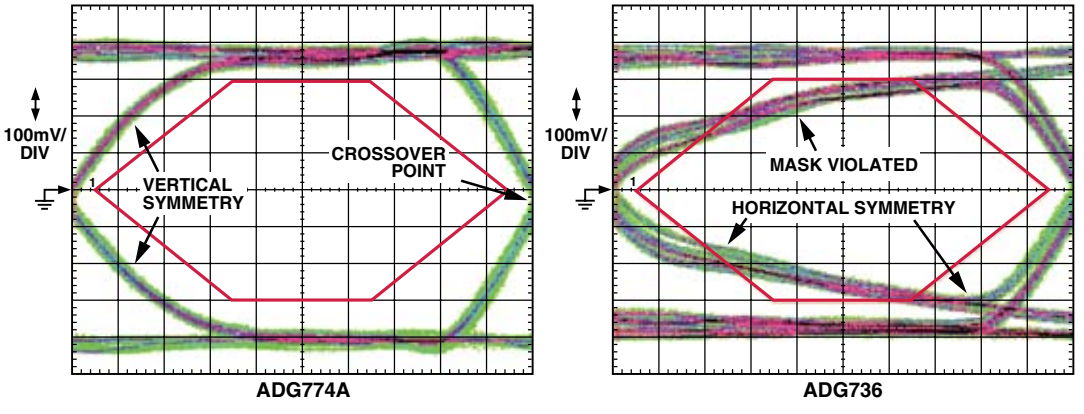

Figure 5. Comparison of the ADG774A and ADG736 at USB Hi-Speed.

The illustration shows that the ADG774A complies with the mask, displaying little ripple, even at these high data rates. The ADG736, however, with its higher capacitance and lower bandwidth, has slowed down the edges, thereby causing the signal to cross the mask on the left side—a clear violation, which disqualifies it from being used to pass Hi-Speed USB signals. Other noteworthy information is the lack of horizontal symmetry in the ADG736 eye, whereas the ADG774A is quite symmetrical, even at this high data rate. Both switches exhibit good symmetry vertically, however, which would indicate good matching of the two channels on both devices. *Channel matching* is a big concern when selecting a switch for USB applications. In a differential system, the D– signal must be the exact inverse of the D+ signal. Mismatches in cable length, capacitance, and resistance between the D+ and D– lines can cause serious skew in the eye, manifested as vertical asymmetry. The point where the signals cross (the *crossover point*) should be centered on ground. *Jitter* is also critical to USB qualification. The thicker the edges, the worse the jitter—not a problem with these CMOS switches. Actually, the jitter seen was also visible with the switch removed, suggesting that the jitter exists in the system.

Figure 6 is a typical plot for an ADG787,[3] using a USB Full-Speed signal (0 V to 3 V, 12 Mbps) in a setup similar to the one used for the above plots. The mask shown is taken from the USB-IF spec for USB Full Speed. The signal used had a rise- and fall time of six nanoseconds. As can be seen, the signal is free from the faults discussed above. No mask violation, good jitter, good crossover and symmetry, and little rippling can be observed. These plots demonstrate the value of an eye diagram, in that at a glance we can conclude that this ADG787 can easily pass a Full-Speed USB signal.

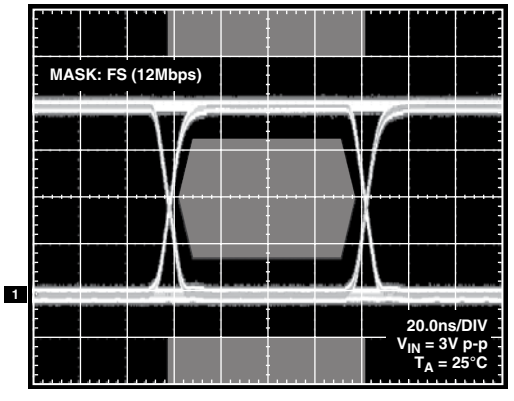

Figure 6. Eye diagram of the ADG787 at USB Full Speed.

How to Choose a CMOS Switch for USB Applications

We now illustrate the specific requirements of a switch and how they affect the signal. This section will consider the correlation between switch specifications and overall system signal integrity.

Switch requirements for both standards would call for as low an on resistance as possible, combined with low capacitance. The characteristics of the two switches need to be matched as accurately as possible to keep the data-line symmetry.

On Resistance

In a 45-Ω system, an on resistance of greater than 5 Ω is undesirable, as a 5-Ω on resistance will add to the source impedance, making it 50 Ω. In order for the receiver to receive a 3-V signal, the 45-Ω source transmits a 6-V signal, which is ideally halved by the divider formed by the 45-Ω source and the termination impedance. This is illustrated in Figure 7, which shows the switch as a resistance in series with the driver. With 5 Ω in series, the receiver sees a 50-Ω source and 45 Ω to ground.

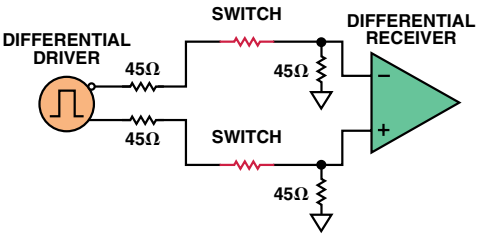

Figure 7. Model of the switch as a resistor

The equations compare performance of an ideal switch with one having 5 Ω of series resistance. A significant loss (>5%) is introduced by the switch. Therefore low R_{on} is critical.

$$6\,V \times \frac{45}{45+45} = 3\,V \qquad 6\,V \times \frac{45}{50+45} = 2.84\,V$$

ideal 0-Ω switch with 5-Ω switch

The source-to-drain resistance of a CMOS switch varies with both the supply voltage and the bias voltage, as illustrated in the R_{on} plot for an ADG787 switch. As the voltage on the source is varied, the resistance measured from source to drain changes.

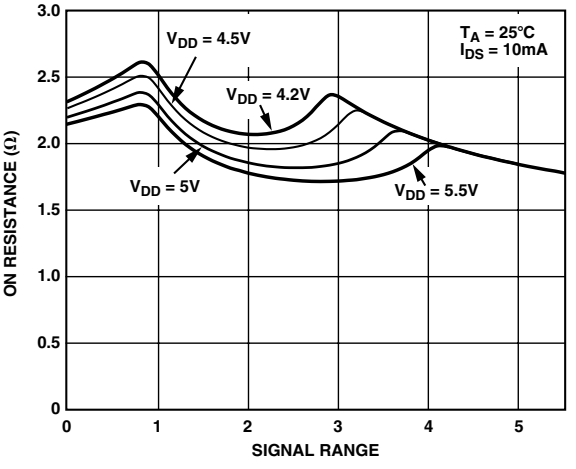

Figure 8. On-resistance variation over input source voltage of the ADG787.

If the resistance of the switch varies, with either bias voltage, temperature, or supply, the amplitude seen by the receiver will also vary, as can be seen for a varying R_{on} (i.e., $R_{on} + \Delta R_{on}$).

$$6\,V \times \frac{45}{45+45+R_{on}+\Delta R_{on}} = \text{a variable amplitude}$$

R_{on} flatness is also vital to ensuring that the rise- and fall times of the switch are as close as possible. If R_{on} varied significantly with bias, the rising and falling edges would see different impedances at different stages in their transition. Differences here would be seen as poor crossover in the eye diagram.

Therefore, R_{on} variability with supply voltage, temperature, and bias are big considerations when designing a switch for use in USB products. Variability of R_{on} over supply tolerances and temperature would be seen on the eye diagram as jitter. As a rule, lower R_{on} means lower flatness and distortion, as can be seen by comparing Figure 9 (ADG836)[4] with Figure 8. The ADG836, which is a dual-SPDT switch fabricated on a 0.35-µm geometry, has R_{on} of about 0.5 Ω and 0.05-Ω flatness, compared with 2 Ω and 0.25-Ω flatness for the ADG787. Keeping R_{on} low is the key to keeping R_{on} flatness low.

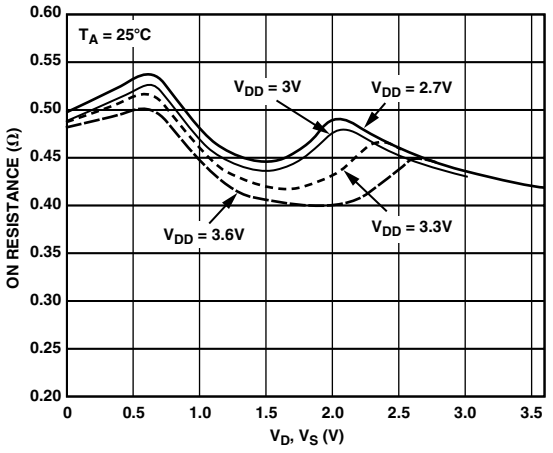

Figure 9. Ultralow on resistance of the ADG836 ensures excellent on-resistance flatness.

Channels should be matched as closely as possible when designing, to ensure R_{on} and ΔR_{on} are the same for the two switch channels passing the differential signals. An eye diagram would indicate poor matching of on resistances.

Capacitance

Capacitance of CMOS switches in the on state increases with size of the switch. However, since low on resistance is achieved by *increasing* the size of the switch, there is a direct trade-off between R_{on} and capacitance. This capacitance, which dictates the bandwidth of the switch, becomes more critical for Hi-Speed USB signals, where switch capacitance greater than 10 pF can significantly degrade the signal. The high capacitance slows the edges down, causing the eye to cross the mask. This was seen in the comparison of the ADG736 and the ADG774A USB Hi-Speed eye diagram of Figure 5. The ADG736 has a bandwidth of 200 MHz. The ADG774A has a much lower capacitance with a bandwidth of 400 MHz. A –3-dB switch bandwidth of greater than 6 MHz (12 Mbps) is required for USB Full Speed, with 240 MHz (480 Mbps) required for Hi-Speed USB. The layout engineer needs to ensure very close similarity of switch layouts to maintain symmetry capacitively.

Propagation Delay

By itself, a CMOS switch in the *closed* state adds negligible delay to a digital signal passing through it. The switch introduces no buffers in the path, and it can be modeled as a series resistance. The only real delay added by the switch is the time taken by the signal to get to the die, and out again. This value can be measured in picoseconds.

Supplies

For Low- and Full-Speed USB, the signal amplitude is 3.3 V ± 10%. Therefore, 3.6 V is the minimum supply allowable. The amplitude of the Hi-Speed signal is 400 mV ± 10%, which can easily be passed by a switch on a 3.3-V supply. It is possible for the CMOS switch to be powered using the USB cable's 5-V supply line. When passing Full-Speed signals (3 V, 12 Mbps), a full signal range is desirable.

Switch Protection

The USB spec states that the data lines of a USB device must be able to withstand being shorted to the 5-V supply line for a period of 24 hours. This has implications for using 3.3-V (0.35-μm geometry) switches in order to obtain the required R_{on} and capacitance. It also has implications for portable devices such as handsets, which use a 3.3-V supply.

Figure 10 shows a 0.35-μm switch being supplied by a 3.3-V regulator at the input to a USB transceiver. One channel is shown for simplicity. This is a typical circuit using a 0.35-μm-geometry switch in a USB application.

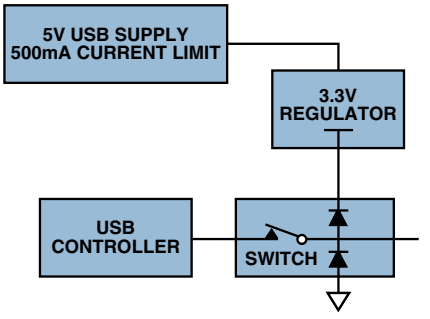

Figure 10. Switch in a normal USB situation with no stress.

Figure 11 introduces a short (in red) from the 5-V supply to the data line. This could happen if the device were plugged into a faulty port.

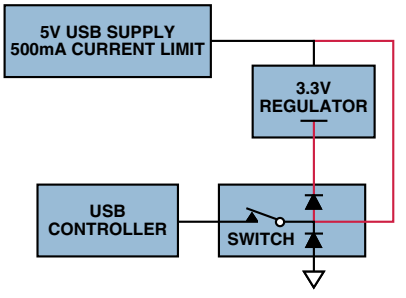

Figure 11. Positive supply of switch forward-biased.

The short circuit forward-biases the ESD (*electrostatic-discharge-protection*) diode to VDD, which means that 500 mA could flow continuously through the ESD diode—a circumstance potentially very damaging to the CMOS switch, which would not be likely to survive more than 24 hours. This is a limitation in implementing 0.35-μm parts. In systems that require this USB condition to be met, and a 3-V switch was to be used, the designer would need to provide adequate protection to prevent this failure mechanism. The easiest way of doing this is using a resistor to limit the current flow. However, the most common solution is to avoid this altogether by using a switch powered from 5-V supplies.

Consumer Applications

Having shown the basic ways in which switches are used in USB applications, we now survey some specific areas of application and discuss the ways they make use of switches. It will be noted that many of them have common topologies.

Portable Media Players (PMPs)

PMPs are rapidly becoming a must-have gadget throughout Asia; it is predicted that they will soon replace the MP3 market. PMPs can record directly from a TV, VCR, DVD player, cable box, or a satellite receiver, and can store up to 120 hours of video, 300,000 pictures, 16,500 songs or 30 GB of data. A portable device that can store this amount of data must have a fast, easy-to-use interface. The interface of choice, usually USB Hi Speed, is one that can be used with a USB camera, a USB card reader, or USB hard drives.

Consumer demand for this type of product also dictates a slim, portable device, so traditional bulky headphone connectors could not even be considered. Instead, the headphone connector is replaced by a mini USB connector, which is shared by the USB data stream and audio outputs.

As shown in Figure 12, a switch is typically needed to isolate the USB signal from the analog audio output. This minimizes reflections by isolating the audio signals from the connector D– and D+ pins when in *data* mode. Reflections during fast signal logic state transitions can potentially cause higher bit-error rates and violate the 500-ppm accuracy requirement of the USB Hi-Speed connection.

For applications of this kind, switches with wide bandwidth and good on-resistance matching help minimize USB signal-edge distortion, while in *audio* mode (output connected to a headphone), low on resistance (about 2.5 Ω), and low total harmonic distortion (about 0.1%) are critical to minimize audio distortion.

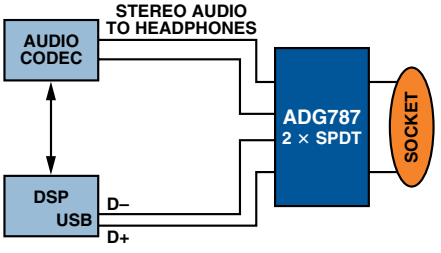

Figure 12. Sharing a mini USB connector between audio and USB.

Handsets/Cell Phones

As handsets acquire additional features, the challenges to a designer also increase. Many of the currently available handsets have cable connections to link to a PC. These connections are used to transfer data, such as emails, calendar, phone book, alarm clock, voice memos, and calculators. If the handset has an integrated camera, the ability to download pictures is also an attractive feature.

So, a handset may have many features that generate a need for USB-compatible switches. One of the most common requirements is in switching between different data standards, for example, between UART and USB. Handset manufacturers like to retain the capability of offering a choice of data-transmission standards to their customers, but they cannot afford the area needed for a separate connector for each interface. The easiest solution is to multiplex a number of pins on a common connector. An example of this is shown in Figure 13.

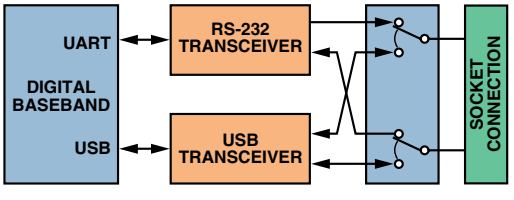

Figure 13. Switching between UART and USB using the switches on an ADG787.

Increasing resolutions of LCD panel displays and cameras in high-end phone designs generate a requirement for larger storage devices, such as embedded hard drives or external, reduced-size memory cards. Most cell phones use standalone hard-drive controllers with a USB interface to communicate with a PC host. When a full-speed I/O port for the baseband processor is also used for synchronizing address books or other data, sharing a single USB port becomes a challenge. The design is simplified by multiplexing the phone's USB connection, as shown in the example of Figure 14.

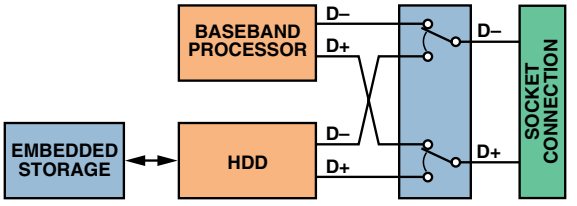

Figure 14. Multiplexing the USB port of a handset

For both functions in a handset as described above, the specifications that need to be considered by the designer are:

- the bandwidth requirements of the chosen USB standard
- on-resistance matching and/or matched propagation delays
- low on-resistance flatness/minimal additive jitter
- power and package size

A further function in a handset is port/bus isolation. This function, not limited to handsets, is also used in other portable designs such as digital still cameras (DSC), PMPs, and pen drives.

Switches are commonly used to protect internal ASICs that could be interfered with by external noise. Of greater importance: For high-end portable design with USB OTG interfaces, isolation between the USB PHY (USB physical layer transceivers) and the external world can further reduce the potential risk of triggering false session-request-protocol (SRP) pulses between dual-role devices—such as two cell phones. The specification of choice for a switch in this application is off isolation, needed when the switch is open and the USB port is not in use (Figure 15). On the other hand, when the USB bus is activated, wide switch bandwidth is needed for minimal deterministic jitter. Many Analog Devices switches are suitable for this application; the tables at the end of this article are a compact source of useful information.

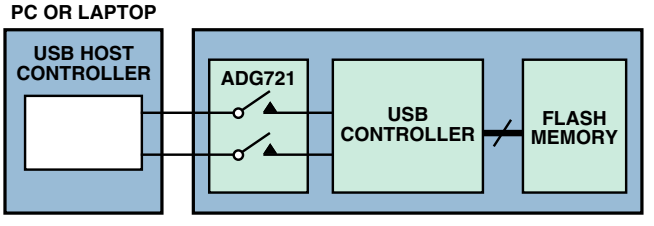

Figure 15. Dual-SPST switch used for isolation of USB.

Wireless Pen Drives and Wireless Adaptors

USB flash drives (pen drives) have become valuable tools for data sharing in both office and home applications because of their mobility, wireless capability, and scalable memory size. Another popular device is the USB wireless adapter; by simply plugging it into your PC, you can connect to the Internet wirelessly without the need for a Centrino™ chip, for example. Wireless USB adapters with memory storage capacity offer a convenient way for business travelers to switch between wireless Internet functions and storage and retrieval functions.

Most storage devices such as hard-disk drives or compact flash memory controllers have Hi-Speed USB interfaces, which are not integrated into a wireless LAN PHY. A USB-compatible switch can easily solve this design challenge by switching between flash memory storage and wireless functions (Figure 16). Low power consumption is desirable, since most of the power consumed by wireless USB adapters comes from the bus of the host application. Small and thin packaging is critical for such applications with very limited PCB space available inside the pen drives.

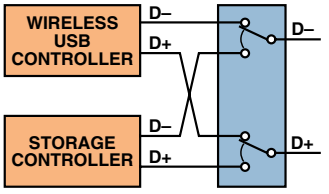

Figure 16. Using the ADG787 for memory/wireless switching in a USB wireless adaptor.

Personal Computers

The PC is at the hub of most USB systems. In all but USB OTG systems, the PC acts as the host of the system. The many traditional USB 1.1 peripherals, such as handsets, digital still cameras, modems, keyboards, mice, some CD-ROM drives, tape and floppy drives, digital scanners, and specialty printers—all interconnect with the PC. USB 2.0 Hi-Speed now accommodates a whole new generation of peripherals, including MPEG-2 video-based products, data gloves, and digitizers. USB has become a built-in feature of most PC chipsets, as well as operating-system- and other system software, without significantly affecting PC prices. By eliminating add-in cards and separate power supplies, USB can help make PC peripheral devices more affordable than they would be otherwise. In addition, USB's hot-swap capability allows users to easily attach and detach peripherals.

As with handsets, CMOS switches can be used to expand the USB bus internally. A further function of switches is in peripheral multiplexing. Figure 17 shows a printer being shared by multiple PCs.

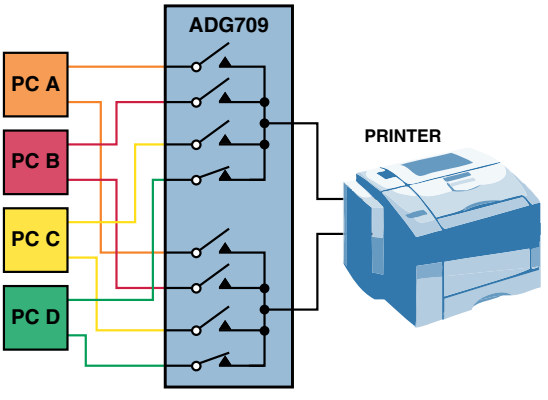

Figure 17. Connecting multiple PCs to a printer with an ADG709.[5]

Board Layout Considerations

Signal routing is key to performance in USB systems. This section details recommended USB PCB signal routing. These comments are based on the system chosen by the USB-IF in order that board and cable designers could design boards that have as little effect on the USB signals as possible at all USB speeds.

As noted earlier, USB is based on a 45-Ω single-ended transmission-line system. It requires that the D+ and D– tracks have impedance to the ground plane of 45 Ω for optimum signal integrity, i.e., to help prevent reflections and signal loss for high-speed signals.

Differentially, the D+ and D– lines should have a mutual impedance of 90 Ω, that is, the impedance between D+ and D– should be 90 Ω.

An impedance calculator should be used to come up with the differential trace spacing. In the example of Figure 18, a useful spacing for 1-ounce copper was calculated to be as shown.

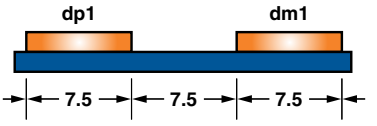

Figure 18. 1-ounce copper spacing to give 90-Ω differential impedance. Dimensions are in mils (inches × 10^{-3}).

Continuing with the above example, and knowing the track width to be 7.5 mils (0.1905 mm), a layer stackup depth can be calculated for 45-Ω single-ended impedance. Prepreg is the dielectric, commonly Rogers material or FR4. The dielectric constant of this material is critical in calculating the depth to the ground plane. The prepreg in the example of Figure 19 has a dielectric thickness of 4.5 mils, with 1-ounce copper traces and planes.

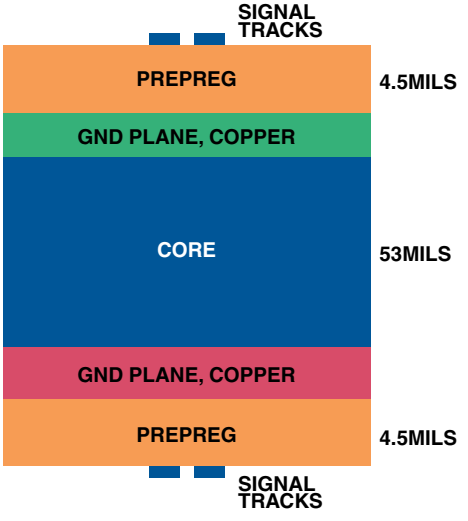

Figure 19. Layer stackup for 45-Ω single-ended impedance.

In the stackup shown, two internal planes were used, enabling signals to be routed along the top and bottom of the four-layer board. For a two-layer board, the stackup would only include the signal tracks, prepreg, and ground plane, with the same spacings and dielectric.

The board shown in Figure 20, using the above stackup and spacings, was employed to evaluate Analog Devices CMOS switches for USB-style measurements. USB connectors were also incorporated on the board.

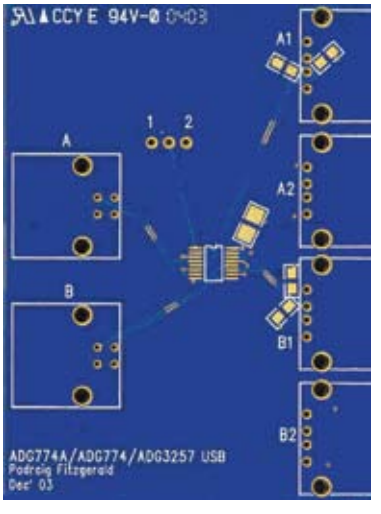

Figure 20. Board used by Analog Devices for USB verification, showing 90-Ω differential impedance matching, 45-Ω single-ended.

With rise- and fall-time edges of Hi-Speed USB as fast as 500 picoseconds, an impedance mismatch can result in transmission-line reflections. In order to avoid reflections, a switch should ideally be placed as close as possible to the USB driver output. The switch would then be seen as a lumped load at the driver output, and there would be minimal signal reflections. In addition, this placement helps improve EMI performance.

The difference between the trace lengths carrying the differential signals should be minimized to optimize the skew between channels; this helps to decrease the *deterministic* jitter (nonrandom, repeatable, or predictable jitter). For best signal integrity, minimal trace length between the USB driver and the connector is recommended. A lower bandwidth results in edge roll-off of the USB signals and may contribute to increased phase jitter and noise.

In addition to the natural decoupling capacitor between the power and ground planes inherent in the four-layer designs, additional paralleled decoupling capacitors (1 μF and 0.1 μF) should be attached close to the V_{dd} pin of the switch.

If the application requires higher ESD performance than is already available in the switch (for example, ADG787 has 2-kV HBM), you may add external ESD devices to the bus. However, it is recommended that the input/output capacitance of external ESD devices be less than 1 pF, and that they be placed close to the USB connector port to minimize bus loading.

For minimal static power consumption, the switch-control signal should swing as closely as possible between 0 V and V_{dd}.

Finally, if the USB controller's output-signal eye diagram has little passing margin or already fails the USB eye mask requirement, adding a switch will not result in successful eye compliance. To improve the eye, the output drive of the controller should be increased, or board-layout issues should be resolved, before the switch is incorporated.

CONCLUSION

With the increasing prevalence of USB functions in both portable and hand-held applications, high-quality switches, using ultralow power, are playing a key role. Analog Devices switches, with their small-footprint packaging, allow designers to add high-speed functionality to existing full-speed platforms at low cost and shorten the time-to-market for their applications. These factors, driven by the consumer demand for continual innovation, accelerated design, and reduced manufacturing cycles, are critical considerations for designers.

REFERENCES—VALID AS OF MAY 2006

[1] ADI website:www.analog.com (Search) ADG774A (Go)
[2] ADI website:www.analog.com (Search) ADG736 (Go)
[3] ADI website:www.analog.com (Search) ADG787 (Go)
[4] ADI website:www.analog.com (Search) ADG836 (Go)
[5] ADI website:www.analog.com (Search) ADG709 (Go)

Selection Table for 12 Mbps

Generic	Configuration	Supply	Package	Specifications
ADG711/ADG712/ADG713	Quad SPST	5 V	16-lead TSSOP/SOIC	2 Ω R_{on}, >200 MHz B/W
ADG781/ADG782/ADG783	Quad SPST	5 V	20-lead LFCSP	2 Ω R_{on}, >200 MHz B/W
ADG721/ADG722	Dual SPST	5 V	8-lead MSOP	2 Ω R_{on}, >200 MHz B/W
ADG736	Dual SPDT	5 V	10-lead MSOP	2 Ω R_{on}, >200 MHz B/W
ADG774	Quad SPDT	5 V	16-lead SOIC	2 Ω R_{on}, >200 MHz B/W
ADG784	Quad SPDT	5 V	20-lead LFCSP	2 Ω R_{on}, >200 MHz B/W
ADG788	Quad SPDT	5 V	20-lead LFCSP	2 Ω R_{on}, >200 MHz B/W
ADG821/ADG822	Dual SPST	5 V	8-lead MSOP	<1 Ω R_{on}, 24 MHz B/W
ADG709	Dual 4:1 Mux	5 V	16-lead TSSOP	3 Ω R_{on}, 100 MHz B/W
ADG729	Dual 4:1 Mux	5 V	16-lead TSSOP	I^2C®, 3 Ω R_{on}, 100 MHz B/W
ADG739	Dual 4:1 Mux	5 V	16-lead TSSOP	SPI®, Ω R_{on}, 100 MHz B/W
ADG759	Dual 4:1 Mux	5 V	20-lead LFCSP	3 Ω R_{on}, 100 MHz B/W

Selection Table for 480 Mbps

Generic	Configuration	Supply	Package	Specifications
ADG3241	SPST	3.3 V	SC70, SOT-66	2 Ω R_{on}, >480 MHz B/W
ADG3242	2 × SPST	3.3 V	SOT-23	2 Ω R_{on}, >480 MHz B/W
ADG3243	2 × SPS	3.3 V	SOT-23	2 Ω R_{on}, >480 MHz B/W
ADG3245	8 × SPST	3.3 V	TSSOP, LFCSP	2 Ω R_{on}, >480 MHz B/W
ADG3246	10 × SPST	3.3 V	TSSOP, LFCSP	2 Ω R_{on}, >480 MHz B/W
ADG3247	16 × SPST	3.3 V	TSSOP, LFCSP	2 Ω R_{on}, >480 MHz B/W
ADG3248	SPDT	3.3 V	SC70	2 Ω R_{on}, >480 MHz B/W
ADG3249	SPDT	3.3 V	SOT-23	<1 Ω R_{on}, >480 MHz B/W
ADG774A	4 × SPDT	5 V	QSOP, LFCSP	2.2 Ω R_{on}, >400 MHz B/W
ADG3257	4 × SPDT	5 V	QSOP	2.2 Ω R_{on}, >400 MHz B/W

This article can be found at http://www.analog.com/library/analogdialogue/archives/40-01/usb_switch.html, with a link to a PDF.

Analog Dialogue

www.analog.com/analogdialogue **dialogue.editor@analog.com**

Analog Dialogue is the free technical magazine of Analog Devices, Inc., published continuously for 40 years—starting in 1967. It discusses products, applications, technology, and techniques for analog, digital, and mixed-signal processing. It is currently published in two editions—*online*, monthly at the above URL, and quarterly *in print*, as periodic retrospective collections of articles that have appeared online. In addition to technical articles, the online edition has timely announcements, linking to data sheets of newly released and pre-release products, and "Potpourri"—a universe of links to important and rapidly proliferating sources of relevant information and activity on the Analog Devices website and elsewhere. The *Analog Dialogue* site is, in effect, a "high-pass-filtered" point of entry to the **www.analog.com** site—the virtual world of *Analog Devices*. In addition to all its current information, the *Analog Dialogue* site has archives with all recent editions, starting from Volume 19, Number 1 (1985), plus three special anniversary issues, containing useful articles extracted from earlier editions, going all the way back to Volume 1, Number 1.

If you wish to subscribe to—or receive copies of—the print edition, please go to **www.analog.com/analogdialogue** and click on **<subscribe>**. Your comments are always welcome; please send messages to **dialogue.editor@analog.com** or to these individuals: *Dan Sheingold*, Editor **[dan.sheingold@analog.com]** or *Scott Wayne*, Managing Editor and Publisher **[scott.wayne@analog.com].**

ADC Input Noise: The Good, The Bad, and The Ugly. Is No Noise Good Noise?

By Walt Kester [walt.kester@analog.com]

INTRODUCTION

All analog-to-digital converters (ADCs) have a certain amount of *input-referred noise*—modeled as a noise source connected in series with the input of a noise-free ADC. Input-referred noise is not to be confused with *quantization noise*, which is only of interest when an ADC is processing time-varying signals. In most cases, less input noise is better; however, there are some instances in which input noise can actually be helpful in achieving higher resolution. If this doesn't seem to make sense right now, read on to find out how *some* noise can be *good* noise.

Input-Referred Noise (Code-Transition Noise)

Practical ADCs deviate from ideal ADCs in many ways. Input-referred noise is certainly a departure from the ideal, and its effect on the overall ADC transfer function is shown in Figure 1. As the analog input voltage is increased, the "ideal" ADC (shown in Figure 1a) maintains a constant output code until a transition region is reached, at which point it instantly jumps to the next value, remaining there until the next transition region is reached. A theoretically perfect ADC has zero *code-transition noise*, and a transition region width equal to zero. A practical ADC has a certain amount of code transition noise, and therefore a finite transition region width. Figure 1b shows a situation where the width of the code transition noise is approximately one *least-significant bit* (LSB) peak-to-peak.

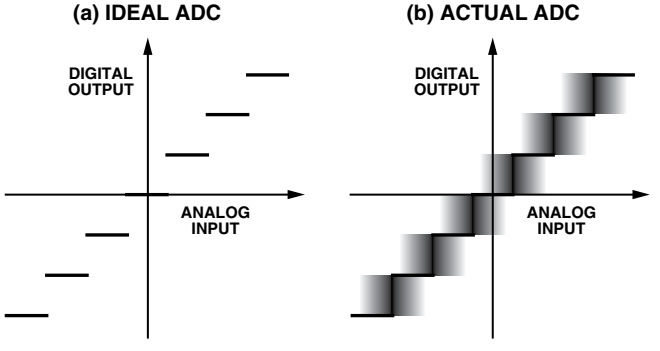

Figure 1. Code-transition noise (input-referred noise) and its effect on ADC transfer function.

Internally, all ADC circuits produce a certain amount of rms noise due to resistor noise and "kT/C" noise. This noise, present even for dc input signals, accounts for the code-transition noise, now generally referred to as *input-referred noise*. Input-referred noise is most often characterized by examining the histogram of a number of output samples, while the input to the ADC is held constant at a dc value. The output of most high speed or high resolution ADCs is a distribution of codes, typically centered around the nominal value of the dc input (see Figure 2).

To measure the amount of input-referred noise, the input of the ADC is either grounded or connected to a heavily decoupled voltage source, and a large number of output samples are collected and plotted as a histogram (referred to as a *grounded-input* histogram if the input is nominally at zero volts). Since the noise is approximately Gaussian, the standard deviation of the histogram, σ, which can be calculated, corresponds to the effective input rms noise. See Further Reading 6 for a detailed description of how to calculate the value of σ from the histogram data. It is common practice to express this rms noise in terms of LSBs rms, corresponding to an rms voltage referenced to the ADC full-scale input range. If the analog input range is expressed as digital numbers, or *counts*, input values, such as σ, can be expressed as a count of the number of LSBs.

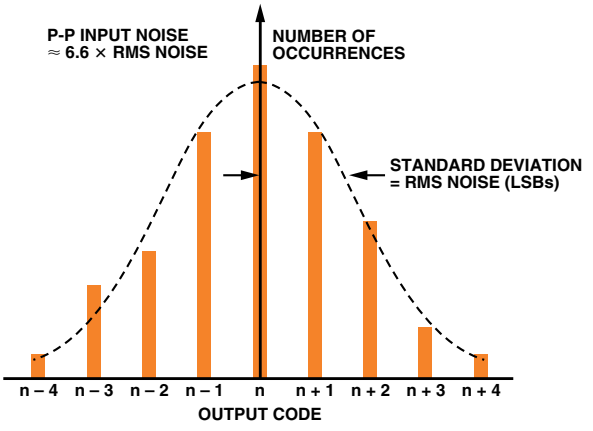

Figure 2. Effect of input-referred noise on ADC grounded-input histogram for an ADC with a small amount of DNL.

Although the inherent *differential nonlinearity* (DNL) of the ADC will cause deviations from an ideal Gaussian distribution (for instance, some DNL is evident in Figure 2), it should be at least approximately Gaussian. If there is significant DNL, the value of σ should be calculated for several different dc input voltages and the results averaged. If the code distribution is significantly non-Gaussian, as exemplified by large and distinct peaks and valleys, for instance, this could indicate either a poorly designed ADC or—more likely—a bad PC board layout, poor grounding techniques, or improper power supply decoupling (see Figure 3). Another indication of trouble is when the width of the distribution changes drastically as the dc input is swept over the ADC input voltage range.

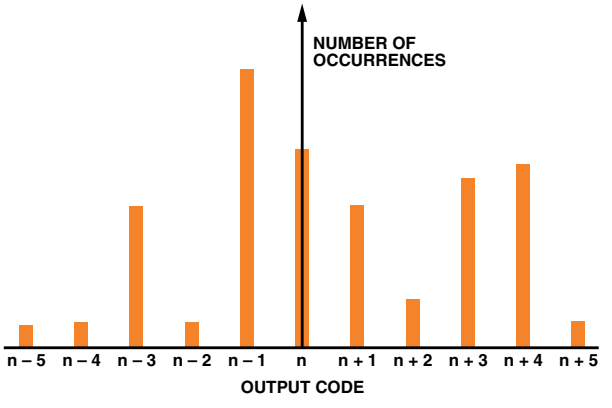

Figure 3. Grounded-input histogram for poorly designed ADC and/or poor layout, grounding, or decoupling.

Noise-Free (Flicker-Free) Code Resolution

The *noise-free code resolution* of an ADC is the number of bits of resolution beyond which it is impossible to distinctly resolve individual codes. This limitation is due to the effective input noise (or input-referred noise) associated with all ADCs and described above, usually expressed as an rms quantity with the units of *LSBs rms*. Multiplying by a factor of 6.6 converts the rms noise into a useful measure of peak-to-peak noise—the actual uncertainty with which a code can be identified—expressed in *LSBs peak-to-peak*.

Peak-to-Peak Resolution vs. Input Range and Update Rate (CHP = 1)

			Peak-to-Peak Resolution in Counts (Bits)					
Output Data Rate	**–3 dB Frequency**	**SF Word**	**Settling Time Normal Mode**	**Settling Time Fast Mode**	**Input Range = ±80 mV**	**Input Range = ±40 mV**	**Input Range = ±20 mV**	**Input Range = ±10 mV**
50 Hz	1.97 Hz	2048	460 ms	60 ms	230k (18)	175k (17.5)	120k (17)	80k (16.5)
100 Hz	3.95 Hz	1024	230 ms	30 ms	170k (17.5)	125k (17)	90k (16.5)	55k (16)
150 Hz	5.92 Hz	683	153 ms	20 ms	130k (17)	100k (16.5)	70k (16)	45k (15.5)
200 Hz*	7.9 Hz	512	115 ms	15 ms	120k (17)	90k (16.5)	65k (16)	40k (15.5)
400 Hz	15.8 Hz	256	57.5 ms	7.5 ms	80k (16.5)	55k (16)	40k (15.5)	30k (15)

*Power-On Default

Figure 4. Noise-free code resolution for the AD7730 sigma-delta ADC.

Since the total range (or span) of an N-bit ADC is 2^N LSBs, the total number of *noise-free counts* is therefore equal to:

$$Noise\text{-}free\ counts = \frac{2^N}{peak\text{-}to\text{-}peak\ input\ noise\ (LSBs)} \quad (1)$$

The number of noise-free counts can be converted into *noise-free (binary) code resolution* by calculating the base-2 logarithm as follows:

$$Noise\text{-}free\ code\ resolution = \log_2\left(\frac{2^N}{peak\text{-}to\text{-}peak\ input\ noise\ (LSBs)}\right) \quad (2)$$

The noise-free code resolution specification is generally associated with high-resolution sigma-delta measurement ADCs. It is most often a function of sampling rate, digital-filter bandwidth, and programmable-gain-amplifier (PGA) gain (hence input range). Figure 4 shows a typical table—taken from the data sheet of the AD7730 sigma-delta ADC[1].

Note that for an output data rate of 50 Hz and an input range of ±10 mV, the noise-free code resolution is 16.5 bits (80,000 noise-free counts). The settling time under these conditions is 460 ms, making this ADC an ideal candidate for a precision weigh-scale application. Data of this kind is available on most data sheets for high-resolution sigma-delta ADCs suitable for precision measurement applications.

The ratio of the full-scale range to the *rms* input noise (rather than peak-to-peak noise) is sometimes used to calculate resolution. In this case, the term *effective resolution* is used. Note that under identical conditions, *effective resolution* is larger than *noise-free code resolution* by $\log_2(6.6)$, or approximately 2.7 bits.

$$Effective\ resolution = \log_2\left(\frac{2^N}{rms\ input\ noise\ (LSBs)}\right) \quad (3)$$

$$Effective\ resolution = Noise\text{-}free\ code\ resolution + 2.7\ \text{bits} \quad (4)$$

Some manufacturers prefer to specify effective resolution rather than noise-free code resolution because it results in a higher number of bits—the user should check the data sheet closely to make sure which is actually specified.

Digital Averaging Increases Resolution and Reduces Noise

The effects of input-referred noise can be reduced by digital averaging. Consider a 16-bit ADC that has 15 noise-free bits at a sampling rate of 100 kSPS. Averaging two measurements of an unchanging signal for each output sample reduces the effective sampling rate to 50 kSPS—and increases the SNR by 3 dB and the number of noise-free bits to 15.5. Averaging four measurements per output sample reduces the sampling rate to 25 kSPS—and increases the SNR by 6 dB and the number of noise-free bits to 16.

We can go even further and average 16 measurements per output; the output sampling rate is reduced to 6.25 kSPS, the SNR increases by another 6 dB, and the number of noise-free bits increases to 17. The arithmetic precision in the averaging must be carried out to the larger number of significant bits in order to gain the extra "resolution."

The averaging process also helps smooth out the DNL errors in the ADC transfer function. This can be illustrated for the simple case where the ADC has a missing code at quantization level k. Even though code k is missing because of the large DNL error, the average of the two adjacent codes, $k - 1$ and $k + 1$, is equal to k.

This technique can therefore be used effectively to increase the dynamic range of the ADC at the expense of overall output sampling rate and extra digital hardware. It should also be noted that averaging will not correct the inherent integral nonlinearity of the ADC.

Now, consider the case of an ADC that has extremely low input-referred noise, and the histogram shows a single code no matter how many samples are taken. What will digital averaging do for this ADC? This answer is simple—it will do nothing! No matter how many samples are averaged, the answer will be the same. However, as soon as enough noise is added to the input signal, so that there is more than one code in the histogram, the averaging method starts working again. Thus—interestingly—some small amount of noise is good (at least with respect to the averaging method); however, the more noise present at the input, the more averaging is required to achieve the same resolution.

Don't Confuse Effective Number of Bits (ENOB) with *Effective Resolution* or *Noise-Free Code Resolution*

Because of the similarity of the terms, *effective number of bits* and *effective resolution* are often assumed to be equal. This is not the case.

Effective number of bits (ENOB) is derived from an FFT analysis of the ADC output when the ADC is stimulated with a full-scale sine-wave input signal. The root-sum-of-squares (RSS) value of all noise and distortion terms is computed, and the ratio of the signal to the noise and distortion is defined as SINAD, or $S/(N+D)$. The theoretical SNR of a perfect N-bit ADC is given by:

$$SNR = 6.02N + 1.76\ \text{dB} \quad (5)$$

ENOB is calculated by substituting the ADC's computed SINAD for SNR in Equation 5 and solving equation for N.

$$ENOB = \frac{SINAD - 1.76\ \text{dB}}{6.02} \quad (6)$$

The noise and distortion used to calculate SINAD and ENOB include not only the input-referred noise but also the quantization noise and the distortion terms. SINAD and ENOB are used to measure the dynamic performance of an ADC, while effective resolution and noise-free code resolution are used to measure the noise of the ADC under essentially dc input conditions, where quantization noise is not an issue.

Using Noise Dither to Increase an ADC's Spurious-Free Dynamic Range

Spurious-free dynamic range (SFDR) is the ratio of the rms signal amplitude to the rms value of the peak spurious spectral component. Two fundamental limitations to maximizing SFDR in a high-speed ADC are the distortion produced by the front-end amplifier and the sample-and-hold circuit; and that produced by nonlinearity in the transfer function of the encoder portion of the ADC. The key to achieving high SFDR is to minimize both sources of nonlinearity.

Nothing can be done externally to the ADC to significantly reduce the inherent distortion caused by its front end. However, the differential nonlinearity in the ADC's encoder transfer function can be reduced by the proper use of *dither* (external noise that is intentionally summed with the analog input signal).

Dithering can be used to improve SFDR of an ADC under certain conditions (see Further Reading 2 to 5). For example, even in a perfect ADC, some correlation exists between the quantization noise and the input signal. This correlation can reduce the SFDR of the ADC, especially if the input signal is an exact sub-multiple of the sampling frequency. Summing about ½-LSB rms of broadband noise with the input signal tends to randomize the quantization noise and minimize this effect (see Figure 5a). In most systems, however, the noise already riding on top of the signal (including the input-referred noise of the ADC) obviates the need for additional dither noise. Increasing the wideband rms noise level beyond approximately 1 LSB will proportionally reduce the SNR and result in no additional improvement.

Other schemes have been developed using larger amounts of dither noise to randomize the transfer function of the ADC. Figure 5b shows a dither-noise source comprising a pseudo-random-number generator driving a DAC. This signal is subtracted from the ADC input signal and then digitally added to the ADC output, thereby causing no significant degradation SNR. An inherent disadvantage of this technique, however, is that the input signal swing must be reduced to prevent overdriving the ADC as the amplitude of the dither signal is increased. Note that although this scheme improves distortion produced by the ADC's encoder nonlinearity, it does not significantly improve distortion created by its front end.

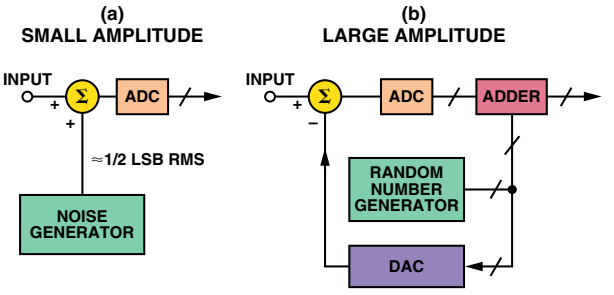

Figure 5. Using dither to randomize ADC transfer function.

Another method, one that is easier to implement—especially in wideband receivers—is to inject a narrow-band dither signal *outside the signal band of interest*, as shown in Figure 6. Usually, no signal components are located in the frequency range near dc, so this low-frequency region is often used for such a dither signal. Another possible location for the dither signal is slightly below $f_S/2$. The dither signal occupies only a small bandwidth relative to the signal bandwidth (usually a bandwidth of a few hundred kHz is sufficient), so no significant degradation in SNR occurs—as it would if the dither was broadband.

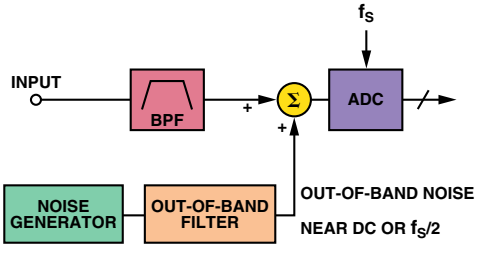

Figure 6. Injecting out-of-band dither to improve ADC SFDR.

A subranging, pipelined ADC, such as the AD6645 14-bit, 105-MSPS ADC[2] (see Figure 7), has very small differential nonlinearity errors that occur at specific code transition points across the ADC range. The AD6645 includes a 5-bit ADC (ADC1), followed by a 5-bit ADC2 and a 6-bit ADC3. The only significant DNL errors occur at the ADC1 transition points—the second- and third-stage DNL errors are minimal. There are $2^5 = 32$ decision points associated with ADC1, which occur every 68.75 mV ($2^9 = 512$ LSBs) for a 2.2-V full-scale input range.

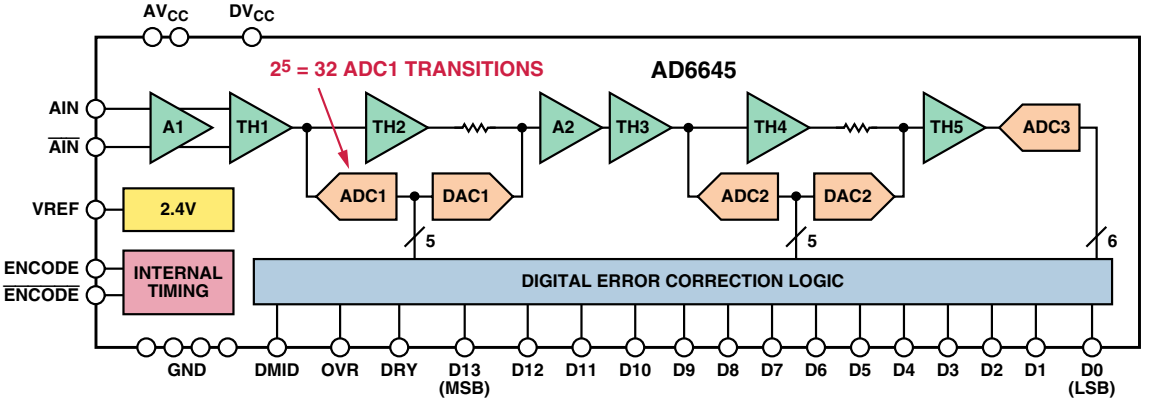

Figure 7. AD6645 14-bit, 105 MSPS ADC simplified block diagram.

Figure 8 shows a greatly exaggerated representation of these nonlinearities.

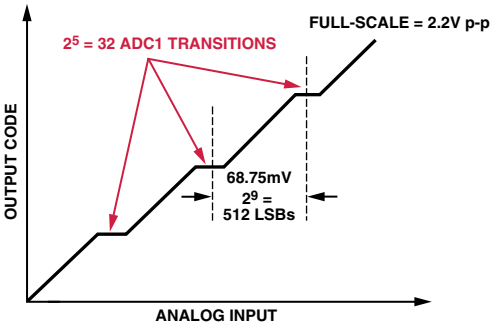

Figure 8. AD6645 subranging point DNL errors (exaggerated).

With an analog input up to about 200 MHz, the distortion components produced by the front end of the AD6645 are negligible compared to those produced by the encoder. That is, the static nonlinearity of the AD6645 transfer function is the chief limitation to SFDR.

The goal is to select the proper amount of out-of-band dither so that the effects of these small DNL errors are *randomized* across the ADC input range, thereby reducing the average DNL error. Experimentally, it was determined that making the peak-to-peak dither noise cover about two ADC1 transitions gives the best improvement in DNL. The DNL is not significantly improved with higher levels of noise. Two ADC1 transitions cover 1024 LSBs peak-to-peak, or approximately 155 (= 1024/6.6) LSBs rms.

The first plot shown in Figure 9 shows the undithered DNL over a small portion of the input signal range, including two of the subranging points, which are spaced 68.75 mV (512 LSBs) apart. The second plot shows the DNL after adding (and later filtering out) 155 LSBs of rms dither. This amount of dither corresponds to approximately –20.6 dBm. Note the dramatic improvement in the DNL.

Dither noise can be generated in a number of ways. For example, noise diodes can be used, but simply amplifying the input voltage noise of a wideband bipolar op amp provides a more economical solution. This approach, described in detail elsewhere (Further Reading 3, 4, and 5) will not be discussed here.

The dramatic improvement in SFDR obtainable with out-of-band dither is shown in Figure 10, using a deep (1,048,576-point) FFT, where the AD6645 is sampling a –35-dBm, 30.5-MHz signal at 80 MSPS. Note that the SFDR without dither is approximately 92 dBFS, compared to 108 dBFS with dither—a substantial 16-dB improvement!

The AD6645 ADC, introduced by Analog Devices in 2000, has until recently represented the ultimate in SFDR performance. In the few years since its introduction, improvements in both process technology and circuit design have resulted in even higher performance ADCs, such as the AD9444 (14 bits at 80 MSPS)[3], the AD9445 (14 bits at 105 MSPS/125 MSPS)[4],

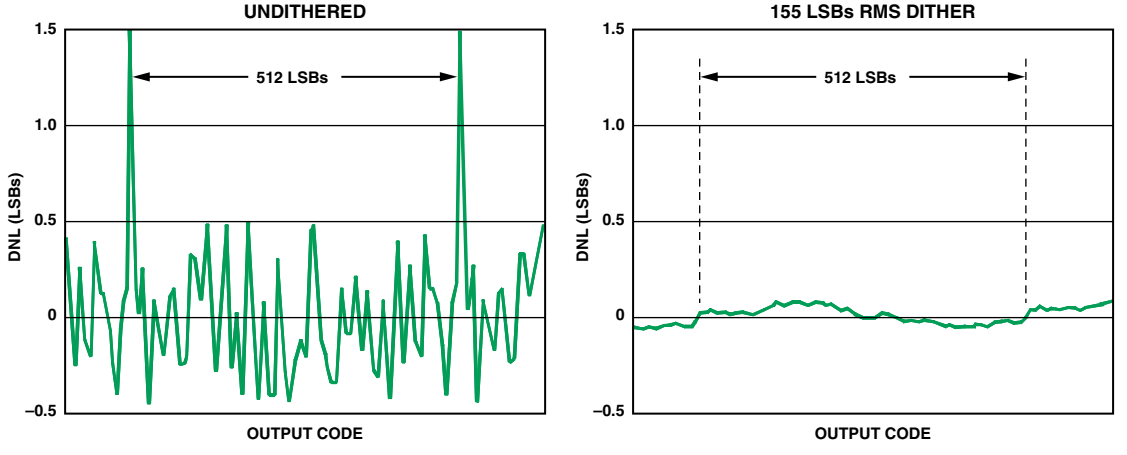

Figure 9. AD6645 DNL plot, without and with dither.

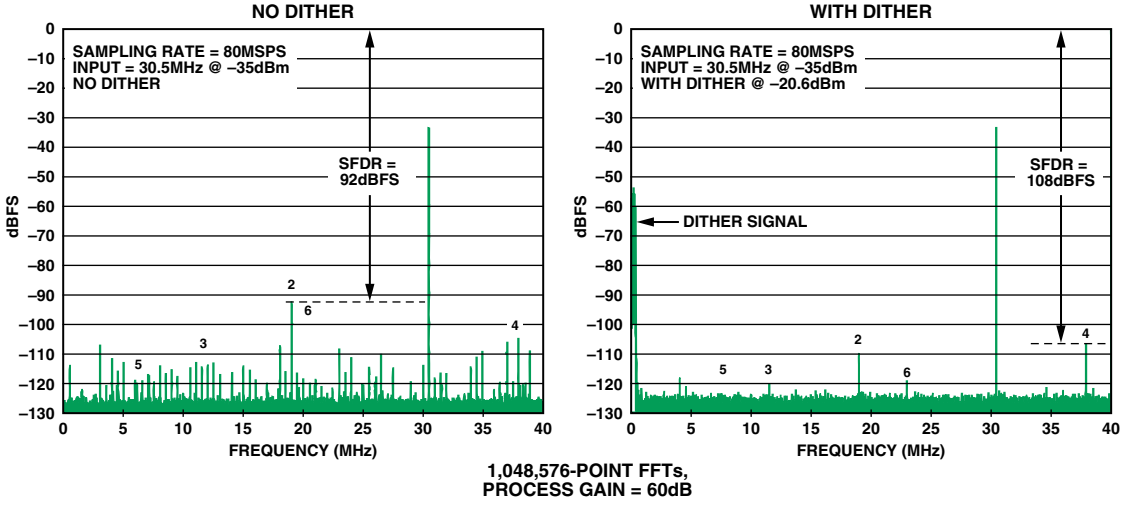

Figure 10. FFT plots showing AD6645 SFDR, without and with the use of dither.

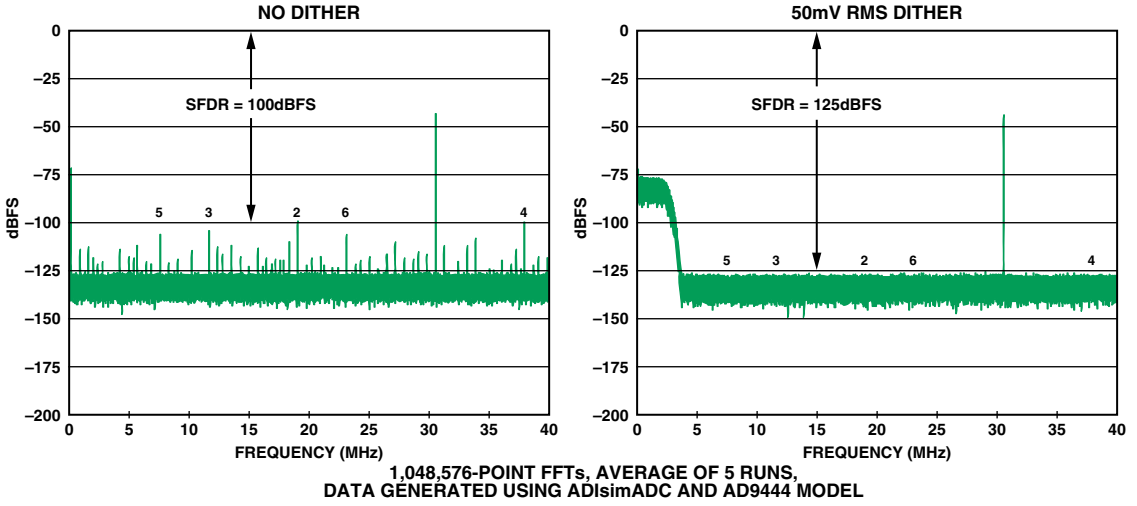

Figure 11. AD9444, a 14-bit, 80-MSPS ADC; f_S = 80 MSPS, f_{IN} = 30.5 MHz, signal amplitude = –40 dBFS.

and the AD9446 (16 bits at 80 MSPS/100 MSPS)[5]. These ADCs have very high SFDR (typically greater than 90 dBc for a 70-MHz, full-scale input signal) and low DNL. Still, the addition of an appropriate out-of-band dither signal can improve the SFDR under certain input signal conditions.

Figure 11 models FFT plots of the AD9444, with and without dither. It can be seen that, under the given input conditions, the addition of dither improves the SFDR by 25 dB. The data was taken using the ADIsimADC™ program[6] and the AD9444 model.

Even though the results shown in Figures 10 and 11 are fairly dramatic, it should not be assumed that the addition of out-of-band noise dither will always improve the SFDR of the ADC under all conditions. We reiterate that dither will not improve the linearity of the front-end circuits of the ADC. Even with a nearly ideal front end, the effects of dither will be highly dependent upon both the amplitude of the input signal and the amplitude of the dither signal itself. For example, when signals are near the full-scale input range of the ADC, the integral nonlinearity of the transfer function may become the limiting factor in determining SFDR, and dither will not help. In any event, the data sheet should be studied carefully—in some cases dithered and undithered data may be shown, along with suggestions for the amplitude and bandwidth. Dither may be a built-in feature of newer IF-sampling ADCs.

SUMMARY

In this discussion we have considered the *input-referred noise*, common to all ADCs. In precision, low-frequency measurement applications, effects of this noise can be reduced by digitally averaging the ADC output data, using lower sampling rates and additional hardware. While the resolution of the ADC can actually be increased by this averaging process, integral-nonlinearity errors are not reduced. Only a small amount of input-referred noise is needed to increase the resolution by the averaging technique; however, use of increased noise requires a larger number of samples in the average, so a point of diminishing returns is reached.

In certain high speed ADC applications, the addition of the proper amount of out-of-band noise dither can improve the DNL of the ADC and increase its SFDR. However, the effectiveness of dither in improving SFDR is highly dependent upon the characteristics of the ADC being considered.

ACKNOWLEDGEMENTS

The author would like to thank Bonnie Baker of Microchip Technology and Alain Guery of Analog Devices for their thoughtful inputs to this article.

FURTHER READING

1. Baker, Bonnie, "Sometimes, Noise Can Be Good," *EDN*, February 17, 2005, p. 26.
2. Brannon, Brad, "Overcoming Converter Nonlinearities with Dither," Application Note AN-410, Analog Devices, 1995.
3. Jung, Walt, *Op Amp Applications*, Analog Devices, 2002, ISBN 0-916550-26-5, p. 6.165, "A Simple Wideband Noise Generator." Also available as *Op Amp Applications Handbook*, Newnes, 2005, ISBN 0-7506-7844-5, p. 568.
4. Jung, Walt, "Wideband Noise Generator," Ideas for Design, *Electronic Design*, October 1, 1996.
5. Kester, Walt, "Add Noise Dither to Blow Out ADCs' Dynamic Range," *Electronic Design, Analog Applications Supplement*, November 22, 1999, pp. 20-26.
6. Ruscak, Steve and Larry Singer, "Using Histogram Techniques to Measure A/D Converter Noise," *Analog Dialogue*, Vol. 29-2, 1995.

REFERENCES—VALID AS OF MAY 2006

[1] ADI website: www.analog.com (Search) AD7730 (Go)
[2] ADI website: www.analog.com (Search) AD6645 (Go)
[3] ADI website: www.analog.com (Search) AD9444 (Go)
[4] ADI website: www.analog.com (Search) AD9445 (Go)
[5] ADI website: www.analog.com (Search) AD9446 (Go)
[6] ADI website: www.analog.com/ADIsimADC

This article can be found at http://www.analog.com/library/analogdialogue/archives/40-02/adc_noise.html, with a link to a PDF.

Wireless Short-Range Devices: Designing a Global License-Free System for Frequencies <1 GHz

By Austin Harney [austin.harney@analog.com]
Conor O'Mahony [conor.omahony@analog.com]

INTRODUCTION

The term *short-range device* (SRD) is intended to cover radio transmitters that provide either unidirectional or bidirectional communication and have little capability of causing interference to other radio equipment. One cannot list all the applications of SRDs because they provide many different services. Among their more popular applications are:

- Telecontrol for home- or other building-automation systems
- Wireless sensor systems
- Alarms
- Automotive, including remote keyless entry and remote car-starting
- Wireless speech and video

Designers of SRD wireless systems need to use great care in choosing the radio's communication frequency. In most cases, the choice is limited to those portions of the spectrum that allow license-free operation given that certain specifications and conditions on usage are met. Table I lists the frequency bands available globally.

Table I. Global SRD Frequency Allocations

Global Frequency Allocations	Comments
13.56 MHz	Used for near-field communications
40 MHz	Not often used, long range possible
433 MHz	Need to reduce power for U.S.
2.4 GHz	Popular global band
5.8 GHz	Some systems upbanding from 2.4 GHz
Multiregion Allocations	**Comments**
868 MHz/915 MHz	Available in Europe/U.S./Canada/ Australia/NZ

The 2.4-GHz band is widely used by designers who want to build systems that can operate worldwide. In fact, it has become the frequency band of choice for such standards as Bluetooth,[1] WLAN,[2] and ZigBee.[3] The 5.8-GHz band has also attracted some attention—in cordless phones or in the 802.11a version of WLAN, for example.

For systems that require both wider range and lower power, however, the sub-1-GHz bands remain compelling due to reduced co-existence issues and greater transmission range, as both of these affect power consumption—an important consideration in battery-powered applications.

The improvement in propagation range for lower-frequency radiators can be shown by a simplified version of the *Friis transmission equation*, which relates the power available in a receiving antenna, P_r, to the power delivered to the transmitting antenna, P_t:

$$P_r = \frac{P_t \lambda^2}{(4\pi d)^2} = \frac{P_t c^2}{(4\pi d)^2 f^2}$$

This equation assumes that both antennas have unity gain. It shows that, for a fixed transmit power, P_t, the received power will decrease with the square of the distance, d, and the square of the frequency, f (or increase with the square of the wavelength, λ). If the received power goes below the minimum power needed to demodulate the signal correctly (called the *sensitivity point*), the link will break down.

Worldwide Frequency Allocations Below 1 GHz

A more detailed description of the various sub-1-GHz standards is given by Table II. It is not an exhaustive list, but more detail can be found by following the links provided in the table.

The 433-MHz band is one option for global usage, with a slight frequency modification required for Japan (easily handled by most modern frequency-flexible transceivers, such as the

Table II. Some Common Regional SRD Bands

Region	Relevant Standards	Frequency Bands (MHz)	Relevant Links
Europe	ERC REC 70-03		
	EN 300 220 (Sept. '00)	433.05 to 434.79 868.0 to 870	http://www.ero.dk/ http://www.etsi.org
	EN 300 220 (Feb. '06)	863.0 to 870	
U.S.	FCC Title 47		
	Part 15.231	260 to 470	http://www.access.gpo.gov/nara/cfr/waisidx_04/47cfr15_04.html
	Part 15.247	902 to 928	
Canada	RSS-210	260 to 470 902 to 928	http://strategis.ic.gc.ca/epic/internet/insmt-gst.nsf/en/sf01320e.html
Japan	ARIB STD-T67	426.0375 to 426.1125 429.175 to 429.7375	http://www.arib.or.jp/english/
China	RADIO REGULATIONS OF THE PEOPLE'S REPUBLIC OF CHINA	315.0 to 316.0 430.0 to 432.0	http://ce.cei.gov.cn/elaw/law/lb93i1e.txt
Australia	AS/NZS 4268:2003	433.05 to 434.79 915 to 928	http://www.acma.gov.au/ACMAINTER.131180

ADF7020[4]—shown in Figure 1 below—or the ADF7021). Less than 2 MHz of bandwidth is available, however, and applications such as voice, video, audio and continuous data transmission are typically not allowed in this band, thus restricting its use. As a result, it is more commonly used for keyless entry systems and basic telecontrol.

The bands around 868 MHz (Europe) and 902 MHz to 928 MHz (U.S.) are more useful; they do not restrict applications, and they allow more compact antenna implementations. Other regions, such as Australia and Canada, have adopted versions of these specifications—thus making the band multiregional, albeit not quite fully global.

Prior to the latest EN 300 220 specification, however, the U.S. and European bodies took vastly different regulatory approaches. The U.S. adopted a frequency-hopping approach, while Europe applied duty-cycle limits in each of the subbands as described in the ERC REC-70 document. Both of these implementations are useful in minimizing interference, but manufacturers who were designing systems for both regions needed to completely rewrite the *media-access layer* (MAC) in the system's communication protocol.

Fortunately, the latest European EN 300 220 regulations (due for release in mid-2006) have extended the frequency bands to allow for *frequency-hopping spread-spectrum* (FHSS) or *direct-sequence spread-spectrum* (DSSS). This makes the MAC implementations more similar to those designed for the U.S, but some fine-tuning will still be required. The following sections describe some aspects of the new specification, covering areas that the SRD system designer needs to consider.

Frequency-Hopping Systems

The *frequency-hopping spread-spectrum* (FHSS) transmission technology spreads energy in the time domain by dividing the spectrum into a number of channels, then switching between them in a pseudo-random sequence, or "hopping code," that is known by both the receiver and transmitter. To welcome new nodes joining the network, the controller node periodically sends out a beacon signal to which the new node can synchronize. The synchronization time depends on the beacon interval and the number of hopping channels. U.S. and European standards both specify a similar number of hopping channels and a maximum dwell time (the time spent at a particular frequency during any single hop) of 400 ms.

Table III shows the number of channels, effective radiated power (ERP), and duty-cycle requirement for the extended frequency band in Europe (below 870 MHz) when FHSS is used. Bandwidths of up to 7 MHz are available once either the *listen-before-talk* (LBT) or duty-cycle limits are met, as compared to the 2-MHz range available previously.

Listen-before-talk, a "polite" communication protocol, scans the channel for activity before initiating a transmission. Also called *clear-channel assessment* (CCA), systems using it with frequency hopping have no duty-cycle limitations.

Wideband Modulation: DSSS

Besides FHSS, DSSS is also addressed in the new European regulations. In a DSSS system, a narrow-band signal is multiplied by a high-speed pseudo-random-number (PRN) sequence to generate a spread signal. Each PRN pulse is called a "chip," and

Table III. European Channel Requirements

Subband	Number of Hop Channels	Power/Magnetic Field	Other Requirements
865 MHz to 868 MHz	≥60	≤25 mW ERP	LBT or <1% Tx duty cycle
863 MHz to 870 MHz	≥47	≤25 mW ERP	LBT or <0.1% Tx duty cycle

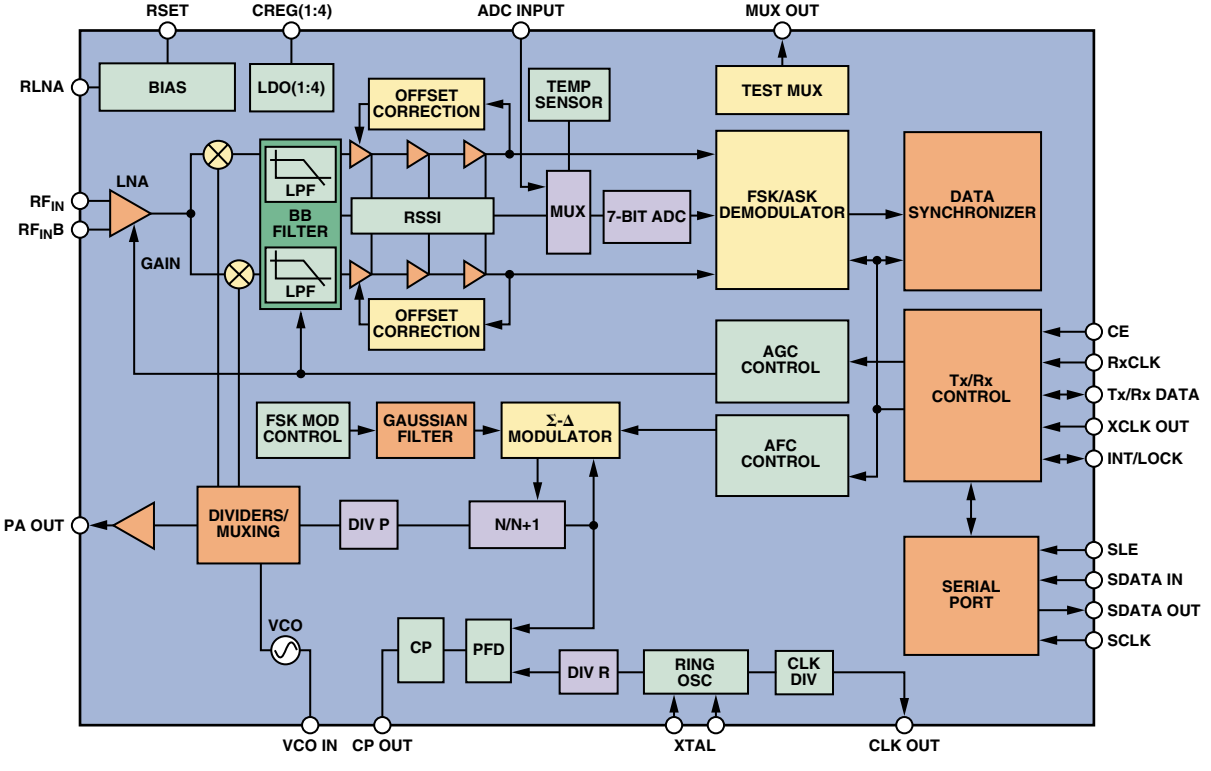

Figure 1. Block diagram of the ADF7020 SRD transceiver.

the rate of the sequence is called the "chip rate." The extent to which the original narrow-band signal is spread is referred to as the *processing gain*; it is the ratio of the chip rate (R_c) to the narrow-band data symbol rate. Frequency spectra of FHSS and DSSS are compared in Figure 2.

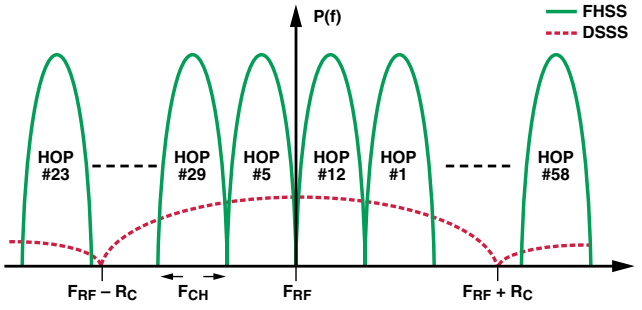

Figure 2. Frequency spectra for FHSS and DSSS.

At the receiver, the incoming spread-spectrum signal is multiplied with the same PRN code to despread the signal, allowing the original narrow-band signal to be extracted. At the same time, any narrow-band interferers at the receiver are spread and appear to the demodulator as wideband noise. The allocation of different PRN codes to each user in the system allows isolation between users in the same frequency band. This is known as *code-division multiple access* (CDMA).

A few examples of systems using DSSS modulation include IEEE 802.15.4 (WPAN), IEEE 802.11 (WLAN), and GPS. The main advantages of DSSS are:

1) Interference resilience—The essence of the interference-rejection capability of DSSS is that the useful signal gets multiplied twice (spread and despread) by the PRN code, while any interferers are multiplied just once (spread).
2) Low power spectral density—Introducing minimal interference with existing narrow-band systems.
3) Security—Very resistant to jamming because of spreading/despreading.
4) Mitigation of multipath effects

Wideband Modulation Other than DSSS or FHSS

An interesting aspect of the new European regulations is that they provide for other wideband spread-spectrum modulation schemes in addition to FHSS and DSSS. FSK/GFSK (*Gaussian frequency-shift-keying*) modulation, with an occupied bandwidth greater than 200 kHz, is considered wideband modulation under the European regulations. Table IV highlights the main specifications, which apply to wideband modulation schemes (including DSSS) in Europe:

Table IV. Maximum Radiated Power Density, Bandwidth, and Duty-Cycle Limits for Spread-Spectrum Modulation (Other than FHSS) and Wideband Modulation

Subband	Occupied Bandwidth (99%)	Maximum Radiated Power Density ERP	Requirements
865 MHz to 868 MHz	0.6 MHz	6.2 dBm/100 kHz	1% Tx duty cycle
865 MHz to 870 MHz	3.0 MHz	–0.8 dBm/100 kHz	0.1% Tx duty cycle
863 MHz to 870 MHz	7.0 MHz	–4.5 dBm/100 kHz	0.1% Tx duty cycle

An example of a device that can take advantage of this wideband standard using FSK modulation is the ADF7025[5] ISM-band transceiver IC. To operate in the 865-MHz to 870-MHz subband, the design must comply with the maximum occupied bandwidth (99%) and maximum power density limits. An edge-of-channel (or band) maximum power limit of –36 dBm is also specified.

With the ADF7025 set up as shown in Table V, all three of these limits were met. Figure 3 shows the occupied bandwidth as 1.7569 MHz and the peak spectral density as –1.41 dBm/100 kHz.

Table V. ADF7025 Parameters for Wideband Modulation Experiment

Frequency	867.5 MHz
Modulation	FSK
Frequency Deviation	±250 kHz
Data Rate	384 kbps

The ADF7025, using wideband modulation, has the possibility of a high data rate (in this case 384 kbps), allowing transmission of audio and medium quality video (several frames per second) in the sub-1-GHz European ISM frequency bands.

The U.S. regulation (FCC Part 15.247) has an allocation similar to that of Europe, which provides for frequency-hopping systems operating in the 902 MHz to 928 MHz, 2400 MHz to 2483.5 MHz, and 5725 MHz to 5850 MHz bands, while also

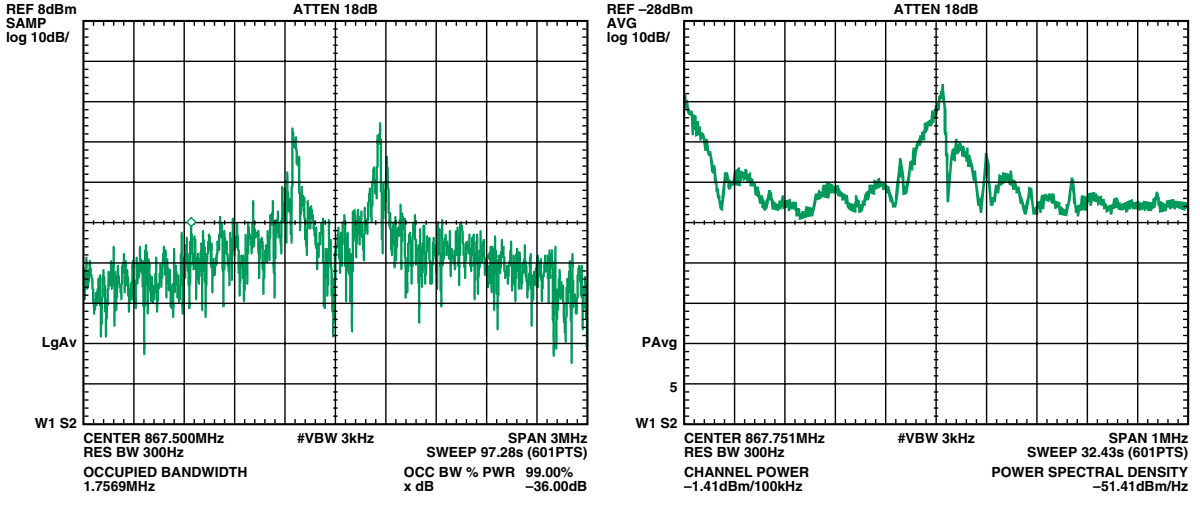

Figure 3. Results for wideband-modulation experiment on ADF7025: (a) FSK modulated signal, 99%-occupied bandwidth measurement, (b) zoomed-in view of (a) to measure maximum power spectral density.

providing for "digitally modulated" signals. This is a loose term that covers both spread-spectrum (DSSS) and other simpler forms of modulation (such as FSK, GFSK), and is thus similar to the "wideband modulation" specification in the European regulations. The two main requirements are:

1. The minimum 6-dB bandwidth shall be at least 500 kHz.
2. For digitally modulated systems, the power spectral density conducted from the intentional radiator to the antenna shall not be greater than 8 dBm in any 3-kHz band during any time interval of continuous transmission.

Anyone wishing to employ a system other than FHSS would normally have to limit the field strength to 50 mV/m (–1.5 dBm ERP). But in the case of "digital modulation," the maximum output power is 1 W, once the maximum power spectral density limit is met. Therefore, using the ADF7025 with an FSK frequency deviation wide enough to ensure that the 6-dB bandwidth is greater than 500 kHz permits a 1-W ERP. Also, with the wide signal bandwidths, higher data rates are possible (maximum 384 kbps for the ADF7025).

The ADF7025's co-channel rejection varies in the co-channel from –2 dB (worst case) to +24 dB, depending on the bandwidth of the interferer. This can be compared with a commercially available 802.15.4 DSSS transceiver, which has a co-channel rejection of –4 dB, where the jamming signal is an IEEE 802.15.4-modulated signal.

Using these methods, similar wideband modulated systems can now be employed in both the U.S. and Europe, thus simplifying the engineering of products intended for worldwide markets. The ADF7025 transceiver architecture lends itself to operation in both the "digital modulation" mode as defined in the U.S. standards and the "wideband modulation" mode as defined in the new European regulations.

Transient Power Requirements

Engineers should also be aware of a new specification in the European regulations that imposes restrictions on *transient power*, which is defined as the power falling into the adjacent spectrum when the transmitter is switched on and off during normal operation. This limit has been added to the latest regulations to prevent spectral splatter.

As the current to the power amplifier (PA) increases (turning on) or decreases (turning off), the load seen by the *voltage-controlled oscillator* (VCO) changes, causing the phase-locked loop (PLL) to unlock for an instant and produce spurious emissions—or *spectral splatter*—while the loop seeks to reacquire lock. In systems where a unit only transmits at intervals, the splatter can significantly increase the power falling into neighboring channels.

Figure 4 highlights the problem of spectral splatter. The yellow trace shows the PA output spectrum from an ADF7020 transmitter when the PA is turned on and off once every 100 ms while the spectrum analyzer is kept on maximum hold. It is evident that significant power is falling into channels on either side of the carrier. The blue trace shows the PA output being ramped on and off in 64 steps every 100 ms, and indicates a considerable reduction of the power falling into the neighboring channels. Specification 8.5 of the latest EN 300 220 regulation establishes a limit on the amount of power falling into these adjacent channels.

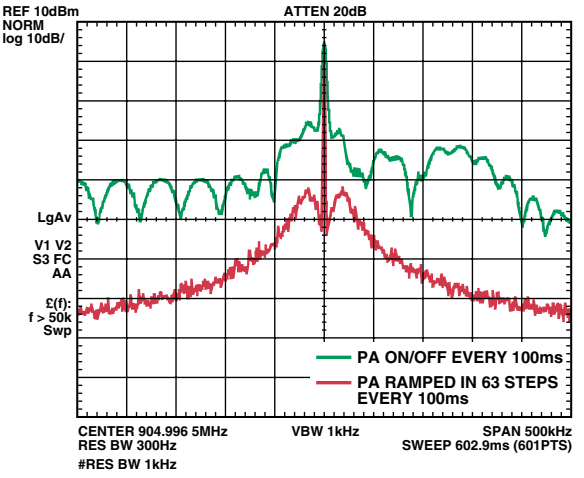

Figure 4. ADF7020 output spectrum for the test described in the text.

The measurement procedure requires that the transmitter be turned on and off five times at maximum output power; and that the power falling into the second, fourth, and 10th channels on either side of the carrier is measured.

The simplest way to ensure compliance with this specification is to ramp the PA gradually off-to-on or on-to-off. This is normally accomplished by using the microcontroller to turn the PA on/off in stages. With the ADF7020 transceiver, it is possible to step the PA from off to +14 dBm in a maximum of 63 steps. A faster and simpler approach is to use a transceiver with an automatic PA ramp. The ADF7021 has a programmable ramp for which both the number of steps and duration of each step can be set by the user.

Communication Protocol Considerations

Analog Devices is currently in the process of updating the ADIismLINK (Version 2.0) protocol software, which can be used with any of the ADF702x transceivers. This protocol, intended for use in the worldwide sub-1-GHz bands, incorporates the new European regulations. It is based on a star-type network (with up to 255 endpoints) illustrated in Figure 5.

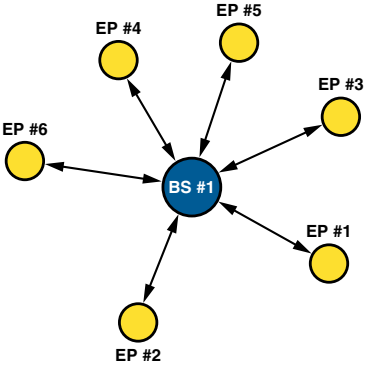

Figure 5. Star network topology.

At the heart of the protocol is a nonslotted, nonpersistent *carrier-sense-multiple-access* scheme with *collision avoidance* (CSMA-CA). The endpoint (EP) listens to the channel before transmitting (LBT), thereby avoiding collisions.

The *nonslotted* aspect of the protocol means that EPs can transmit as soon as they have data, subject to first performing a *listen-before-talk* operation. This approach also ensures that no synchronization is required. If an EP senses the channel is busy, it backs off for a random period before performing another *LBT*. The number of times this back-off can occur is limited, hence the *nonpersistent* nature of the protocol. In FHSS mode, the protocol uses this CSMA-CA system on each hopping channel, thus fulfilling the LBT requirement for the new European regulations.

The *physical-layer* (PHY) and *media-access-layer* (MAC) parameters of the ADIismLINK protocol are highly configurable, thus allowing thorough device- and system evaluation. Source code is also provided, simplifying the system-development procedure. The protocol comes as part of the ADF702x Development Kit (ADF70xxMB2). A system overview of ADIismLINK is shown in Figure 6. More information on this is available through the ADI website (ADF702x Development Kit).[6]

CONCLUSION

The new European regulations impose very specific requirements for over-the-air protocols in the 863-MHz-to-870-MHz band. Whether a system uses a single-channel protocol, FHSS, or DSSS, there are specific rules that must be observed. This of course complicates the protocol design. However, the upside of these new ETSI regulations is that they mirror the FCC Part 15.247 regulations in many aspects, thus simplifying the design of a protocol intended for multiregion use. In addition, the Analog Devices development kit includes protocol examples to simplify the challenges involved in designing short-range wireless networks.

REFERENCES—VALID AS OF MAY 2006

1 http://www.bluetooth.com/bluetooth/
2 http://grouper.ieee.org/groups/802/11/
3 http://www.zigbee.org/en/index.asp
4 ADI website: www.analog.com (Search) ADF7020 (Go)
5 ADI website: www.analog.com (Search) ADF7025 (Go)
6 ADI website: www.analog.com (Search) EVAL-ADF70xx (Go)

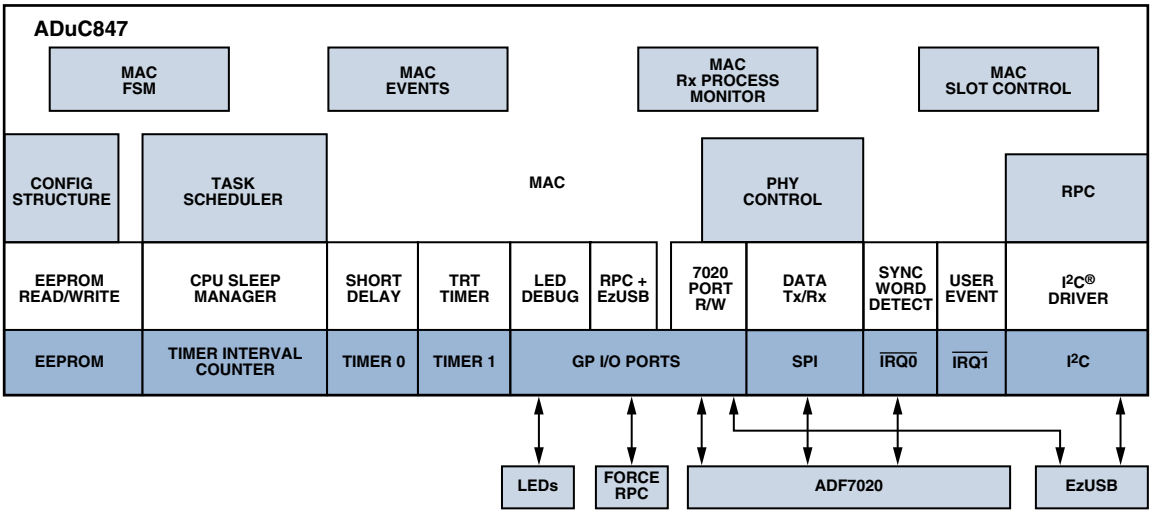

Figure 6. ADIismLINK system overview.

This article can be found at http://www.analog.com/library/analogdialogue/archives/40-03/wireless_srd.html, with a link to a PDF.

THE FOURTH DEE *(continued from page 5)*

immensely *deep insights*—far more quickly and clearly and accurately than classical mathematics ever could, with total control of all the conditions. For example, I expect you studied your oscillator's start-up trajectory (to use your nice descriptive word) at –50°C and +150°C, as well as somewhere in between, and had complete results for these conditions in seconds, right?"

"Yes, that's right, Dr. Leif. But the insights don't come directly from the machine. Even today's AI-rich simulators don't offer ideas! And the so-called genetic algorithms for circuit creation proved to be a great flop. *People* have the insights, and *people* create. I'd put it this way: *simulators stimulate solutions*. So, in this case, I found the tank appeared to have a small impedance, and GE°E gave me its value out to 20 decimal places. But now, with my curiosity stimulated, it was up to me to find the general solution."

"Exactly! You propose **'What If?'** and GE°E replies **'This will happen.'** Then you ask **'Why** does it do that? **How** did you get that value?' And these answers must come from thinking about the Fundaments—the **'What Must Be'** aspects of any circuit. GE°E is like a very willing and reliable, but junior, apprentice, who can calculate well but who cannot yet create. So, Niku, do you have the analytic solution, and how did you get it?"

"Oh, sure, that was one of the first things I took care of. Much more interesting insights came along later. Briefly, I thought of two ways of getting a general value for the impedance. Both rely on the **'What Must Be?'** philosophy. A simple tank only has two elements: an inductor and a capacitor. And there is only one way to get an impedance from these dimensional quantities—namely, by noting that $\sqrt{(L/C)}$ has the dimension of resistance. For our particular example of L = 20 nH and C = 10 pF, this evaluates to 44.721359549995796 Ω, which is precisely GE°E's answer! My second method is more roundabout."

"Niku, I really want to hear all the rest of your discoveries about this oscillator's behavior, but that will have to wait until our next visit. In the meantime, why don't you put your study notes up on ADI's CyberCyte? Hey, I'm proud of you already!"

As he was about to leave the table, he felt it appropriate to give his new protégée a fatherly pat on the shoulder.

(To be continued)

PRODUCT INTRODUCTIONS: VOLUME 40, NUMBER 1

Data sheets for all ADI products can be found by entering the model number in the Search Box at www.analog.com

January

ADC, Pipelined, 16-bit, 80-MSPS/100-MSPS, IF-sampling **AD9446**
ADC, Pipelined, 12-bit, 20-/40-/65-MSPS, low-power, 3-V **AD9237**
ADC, Sigma-Delta, 24-bit, 625-kSPS, serial interface, 109-dB DR **AD7763**
Amplifier, Instrumentation, auto-zero, shutdown, 1.8-V **AD8553**
Amplifier, Operational, CMOS, low-noise, high-precision, 16-V **AD8661**
Amplifier, Operational, low-noise, high-precision, 36-V **AD8675**
Amplifier, Operational, precision, low offset, TSOT package .. **AD8677**
Amplifiers, Operational, CMOS, low-noise, micropower, rail-to-rail I/O **AD8617/AD8619**
Current Mirror, high-side, precision, 3-nA to 3-mA range **ADL5315**
DAC, Voltage-Output, 16-bit, SPI interface, 1-LSB INL **AD5060**
DAC, Voltage-Output, 14-bit, SPI interface, 1-LSB INL **AD5040**
Driver, LCD, 12-bit, 12-channel, decimating *Dec*Driver **AD8387**
Driver, LCD, 10-bit, 12-channel, decimating *Dec*Driver **AD8386**
Driver, VCSEL-Diode, differential, 11.3-Gbps, active back termination **ADN2530**

February

Amplifier, Operational, precision, low-power, auto-zero, rail-to-rail I/O **AD8538**
Digital Isolator, single-channel, enhanced ESD protection .. **ADuM3100**
Digital Isolators, 3-channel and 4-channel **ADuM1310/ADuM1410**
Digital Potentiometer, 256-position, ±15-V **AD5290**
Driver, Half-Bridge, Isolated, 100-mA output **ADuM1230**
Interface, Display, dual analog/DVI **AD9396**
Interface, Display, dual analog/HDMI **AD9380**
Interface, Display, DVI **AD9397**
Interface, Display, HDMI, internal HDCP key storage **AD9381**
Interface, Display, HDMI **AD9398**
Modulator, Sigma-Delta, Isolated **AD7400**
Signal Processor, CCD, 14-bit, *Precision Timing*™ generator **AD9970**

March

ADC, Successive-Approximation, 18-bit, 2-MSPS, 3-LSB INL **AD7641**
ADCs, Successive-Approximation, 8-channel, 12-bit-plus-sign **AD7327/AD7328**
ADCs, Successive-Approximation, 4-channel, 12-bit-plus-sign **AD7323/AD7324**
ADCs, Successive-Approximation, 2-channel, 12-bit-plus-sign **AD7321/AD7322**
Amplifier, Differential, 2-GHz, ultralow distortion **AD8352**
Amplifier, Operational, CMOS, low-noise, micropower, rail-to-rail I/O **AD8613**
Amplifier, Operational, precision, 150-µV offset voltage **OP07D**
Amplifier, Operational, quad, 24-MHz, low-cost, rail-to-rail I/O **AD8648**
Amplifier, Operational, quad, precision, 1.8-nV/√Hz noise, 36-V **ADA4004-4**
Amplifier, Video, low-power, 1-µA disable **ADA4853-1**
Codec, Audio, embedded SigmaDSP processor **ADAV400**
Controller, Synchronous Buck, 7-bit, 1-, 2-, or 3-phase **ADP3207**
DAC, Current-Output, 14-bit, wideband, multiplying, serial interface **AD5453**
Generator, Waveform, programmable frequency sweep, output burst **AD5930**
Microcontrollers, Precision Analog, 12-bit I/O, ARM7 TDMI® MCUW **ADuC70xx**
Monitor/Controller, System, voltage/remote temperature .. **ADT7476A**
Monitor, System **ADT7462**
Reference, Voltage, 1.2875-V, micropower, shunt **ADR1500**
SHARC Processors, 32-bit/40-bit floating-point, high-performance audio **ADSP-2136x**
Switch, 8 × 12 Analog Array **ADG2128**
Switch, CMOS, low-cost, high-voltage, quad SPDT **ADG1334**
Switches, *i*CMOS, triple/quad SPDT, low-capacitance, high-voltage **ADG1233/ADG1234**
Transceiver, FSK/ASK, high-performance, ISM band **ADF7020-1**
Transceiver, FSK/GFSK, high-performance, ISM band **ADF7025**
Translator, Logic-Level, single-channel, bidirectional, 1.15 V to 5.5 V **ADG3301**

AUTHORS

Padraig Fitzgerald (page 6) earned his degree in electronic engineering from the University of Limerick, Ireland, in 2002. He joined the Analog Devices MPX group in Limerick that year, and has been in design evaluation for 3½ years, working on various switch and multiplexer developments.

Austin Harney (page 18) graduated in 1999 with a BEng from University College, Dublin, Ireland, and joined Analog Devices following graduation. He is currently an applications engineer for the ISM-band wireless product line, based in Limerick. In his spare time, Austin enjoys football, music, and spending time with his daughter.

Walt Kester (page 13) is a corporate staff applications engineer at Analog Devices, Greensboro, North Carolina. In his 35 years at ADI, he has designed, developed, and given applications support for high-speed ADCs, DACs, SHAs, op amps, and analog multiplexers. An author of many papers and articles, he has prepared and edited 11 major applications books for ADI's global technical seminar series; topics include op amps, data conversion, power management, sensor signal conditioning, mixed-signal circuits, and practical analog design techniques. His latest book, *Data Conversion Handbook* (Newnes), is a nearly 1000-page comprehensive guide to data conversion. Walt has a BSEE from NC State University and an MSEE from Duke University. He enjoys carpentry and travel.

Eva Murphy (page 6) is an applications engineer at ADI's facility in Limerick, Ireland. She provides applications support for general-purpose, digital-to-analog converters (DACs), switches and multiplexers, and direct-digital-synthesis (DDS) products. Joining Analog Devices in 2000, Eva holds a BEng from Cork Institute of Technology. Her interests include reading, traveling, and walking.

Conor O'Mahony (page 18) graduated in 2004 from University College, Cork, with a bachelor's degree in electrical and electronic engineering. After graduating, he joined ADI as an applications engineer for ISM Band Wireless Products. He's an avid sportsperson, playing soccer, Gaelic football, and hurling. In his spare time, he enjoys reading, music, and traveling.

Analog Devices, Inc.
Worldwide Headquarters
Analog Devices, Inc.
One Technology Way
P.O. Box 9106
Norwood, MA
02062-9106 U.S.A.
Tel: 781.329.4700
(800.262.5643,
U.S.A. only)
Fax: 781.461.3113

Analog Devices, Inc.
Europe Headquarters
Analog Devices, Inc.
Wilhelm-Wagenfeld-Str.6
80807 Munich
Germany
Tel: 49.89.76903.0
Fax: 49.89.76903.157

Analog Devices, Inc.
Japan Headquarters
Analog Devices, KK
New Pier Takeshiba
South Tower Building
1-16-1 Kaigan, Minato-ku,
Tokyo, 105-6891
Japan
Tel: 813.5402.8200
Fax: 813.5402.1064

Analog Devices, Inc.
Southeast Asia
Headquarters
Analog Devices
22/F One Corporate Avenue
222 Hu Bin Road
Shanghai, 200021
China
Tel: 86.21.5150.3000
Fax: 86.21.5150.3222

Printed in the U.S.A. M02000401-54-4/06

www.analog.com/analogdialogue

ANALOG DEVICES

Analog Dialogue