Volume 40, Number 2, 2006

Analog Dialogue

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

In This Issue

Editors' Notes 2
Oscar Stirs [The Wit and Wisdom of Dr. Leif—3] 3
Class D Audio Amplifiers: What, Why, and How 6
Power-Supply Management—Principles, Problems, and Parts 13
Design Your Own VoIP Solution with a Blackfin Processor—Add Enhancements Later 17
Product Introductions and Authors 23

www.analog.com/analogdialogue

ANALOG DEVICES

Editors' Notes

ANNIVERSARY

As faithful readers already know (and are perhaps weary of knowing!), **Analog Dialogue** *is celebrating our 40th anniversary in print. But there is another significant anniversary to celebrate: Exactly 20 years ago, in Volume 20, No. 2 (1986), we broke with our purely analog past and introduced the first Analog Devices* ***digital*** *signal processor (DSP), the ADSP-2100. You may be interested in recalling the theme of the Editor's Note welcoming our readers to that issue:*

ANALOG & DIGITAL

"Microprocessor?" we hear you ask. "Isn't it a bit unseemly for a nice 'Analog' IC company to be designing a microprocessor? (What could be more digital?)"

Good question.

Our objective has always been to design and manufacture cost-effective components that are key elements of the signal path for processing real-world (i.e., *analog*) data and for which performance is maximized and errors minimized.

The signal path? *Real-world data* almost always starts out as *analog* (i.e., parallel, nonnumeric) variables, which are measured by sensors that provide *analog* electrical signals—voltage and current. The signals must be accurately and speedily amplified, conditioned (almost always in parallel), and converted to digital for processing. Once in digital form, they must be processed rapidly. Often, they again wind up as analog signals.

Key elements of the signal path include preamplifiers, analog signal processors, data converters to and from digital, and—when the signal is in digital form—a digital processor. Inadequacy in any one of the key elements—amplifier, analog processor, data converter, or microprocessor—can cause poor performance of the overall system.

Obstacles in the signal path include noise, drift, nonlinearity, and measurement lag at the analog stages, similar obstacles in conversion—and throughput delays in digital processing, often because of the lack of parallelism in von Neumann architecture.

Throughout our history, our role in the signal path has been to initiate new products (or product lines) when dissatisfied with the cost-effectiveness of what's available (which is often limited to user-assembled kludges, when nothing else is available). At this point in time, we (and our competitors) have virtually eliminated the user-assembled amplifier, signal conditioner, and data converter by designing and marketing families of cost-effective products.

We have always been dissatisfied with the cost, power dissipation, and slow throughput in the digital domain; this concern led to our pioneering development of CMOS multipliers and other digital signal-processing ICs (note that because we were already familiar with analog multipliers, digital multipliers became just another analog signal-processing tool; note also our commitment to *signal processing*—not payroll, desktop publishing, or order-handling products). Our dissatisfaction with the complexity of systems using Bit-Slice parts led to the powerful and compact Wordslice™ microcoded system parts.

And finally, our dissatisfaction with insufficient throughput in DSP processors led to design of the ADSP-2100, which stresses the use of that *analog* characteristic, parallelism, to minimize instruction cycles, whether in processing, data transfer, or interrupt handling. It's neat! We invite you to read about it.

Dan Sheingold [dan.sheingold@analog.com]

RANDOM THOUGHTS ON THIS ISSUE

Dr. Leif's World—And Ours

In this issue, enjoy *Oscar Stirs*, part three of Niku's quest to understand the start-up phase of an ideal oscillator. In Dr. Leif's world, the effects of global warming have already been felt, yet there is hope for the future. Is our outlook as bright, or will some of the cataclysmic events experienced in Niku's lifetime become reality in ours? Everyone should see the intentionally biased movie, *An Inconvenient Truth*, regardless of their current views on global warming. At the very least it will make them wonder about how robust or fragile our environment is and question what they read in the popular press. Some may be motivated to act by reducing their own impact on our ecosystem or becoming advocates for change. For more information, and some simple suggestions for becoming carbon neutral, visit http://yosemite.epa.gov/oar/globalwarming.nsf.

VoIP—Pro and Con

After using VoIP for about two years, I finally cut my personal ties with Ma Bell. For less than 20 dollars a month I got everything that was previously provided by my land line, plus a lot more. Some of the cool features include voice mail messages sent via email, allowing me to access them from any PC, and the ability to reject anonymous calls without being forced to jump up and check the caller ID. In addition, I can make unlimited calls to the US and Canada, and limited free international calls. So, what's the downside? Well, even with 8 Mbps Internet access and 92% quality of service (QoS), the voice quality, noise level, and echo rejection are not as good as with the wired phone. Also, a call is occasionally dropped, something I *never* experienced with the land line. Furthermore, I'd lose my VoIP connection in the event of a power failure, in contrast to the always-up reliability of the plain old telephone system (POTS). With a cell phone as a backup, though, VoIP is a smart choice. In this issue, read about implementing VoIP on a Blackfin processor.

Power-Supply Management

With today's portable devices and green appliances, power consumption must be kept to a minimum, yet complex consumer electronics equipment often requires multiple power supplies for the analog circuitry, memory, microprocessors, DSPs, and ASICs. These supplies must be turned *on* and *off* in sequence to avoid damage to their associated components, with timeouts occurring in the event of a system fault. All of this increases the importance of integrated power-supply management. In this issue, learn the basics of the problem and of the available solutions.

Class D Amplifiers

Class D amplifiers provide the advantages of higher efficiency, smaller size, and lower cost, while still achieving low distortion, wide dynamic range, and clickless muting. Recently, everyone seems to be writing about Class D. After reading the article in this issue, you may also be interested in reading "How Class D Amplifiers Work" by Jun Honda and Jonathan Adams, "Design and Analysis of a Basic Class D Amplifier" by Chi Ho Li, and "Class D in Audio Hubs Extends Battery Life" by Robert Hatfield. All can be found at http://www.audiodesignline.com.

Scott Wayne [scott.wayne@analog.com]

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 17, Number 1 (1983), 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].**

ISSN 0161-3626 ©Analog Devices, Inc. 2006

OSCAR STIRS
[The Wit and Wisdom of Dr. Leif—3]

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

Back in her small apartment, Niku tossed two units of MycoPro into the hyperphase and settled to watching the top stories on the TransGlobal Network (TGN, Channel 28.20.20). She was feeling upbeat. After weeks of orientation at ADI, she'd just been given her first real assignment as a Product Originator. From the initial discussions, it didn't appear that her contributions to this mixed-signal ASIC, a new type of monitor of deep-ocean conditions for the International Oceanic Authority, would require much in the way of "original" design; but she was excited at the prospect of finally doing something productive, working under the guidance of her team leader, Jeff Rawlings, another protégé of Dr. Leif.

It was Leif who first alerted him to Niku's peculiarly dogged interest in some little-studied subtleties of HF oscillators. Until recently, Jeff had been too preoccupied to ask about this work. But when the opportunity arose, she enthusiastically shared her notes of the step-by-step questioning she'd applied to the start-up behavior of the basic two-transistor differential topology she was exploring. He was immediately impressed by her diligence in this self-appointed enterprise, and encouraged her to present a summary of her work at "Daedalus Days," the series of meetings that Dr. Leif had instituted, years ago, and had named after the mythical Greek inventor. After that chat, Jeff was even more certain that her penetrating curiosity and insatiable appetite for pursuing basic insights would be valuable assets for their new project. Asking questions is the beginning of wisdom.

The majority of the other team members were *Fusers*. While they were experienced and efficient, and the older ones were broadly informed about the IC business, their satisfaction came largely from providing a quick solution to a field application, using the large library of precharacterized, time-proven cells. Most of the Fusers had scant interest in topological exploration, save for one of them: Mauricio Pegna, hired several months before Niku. Rico showed great promise as an eager treasure seeker in the vast and rich forests of analog cell invention, and was clearly destined for an early promotion. A remote relative, Guido Pegna, once taught physics at the University of Caligari in Sardinia. Prof. Pegna, an exceptional teacher, offered a boldly different perspective on science—one that was, he said, "less worrying, less dramatic, more fun, more human, more delightful." Among his legacies was a fascinating site on the old Internet: www.pegna.com.

While picking away at her nutritional but totally uninspiring dinner, Niku watched the latest reports of the disappearing polar ice. The erosion of habitat for polar bears in the Arctic, and the confusing loss of landmarks for penguins and other migratory birds in Antarctica, who for eons had regarded these stable regions as their home and birthright, was troubling enough. But disruptions of that sort were eclipsed by the sheer magnitude of the threat to human life and property posed by the ever-rising sea level.

As a kid, during the overture to the 21st century, Niku had heard plenty of talking heads indulging in endless exchanges of blame: across national borders, between political parties, and between various agencies charged with caring for the environment. But little was done to reverse these trends. For all the bickering, pollutants and greenhouse gases were still on the rise as late as 2016. Now, in the mid '20s, at least they were stable, but at levels far too high for complacency. Only when low-lying coastal regions began to be inundated did world governments finally grasp the ominous reality of the science. After years of glory, the older areas of Venice were the first to succumb, and were soon abandoned.

These shocking events were followed by the asteroid strike in August of 2017. Fortunately for humanity, after a terrifyingly blazing journey on an uncertain and wobbly trajectory, it had burned itself down to the size of a one-seat commutobile when it hit a remote region of Antarctica at a fairly steep angle. The heat equivalent of its massive kinetic energy had greatly accelerated the melting of the ice and the rising sea levels. That event led to the formation of the Coalition of Terrestrial Governments, and galvanized its member-nations into drafting laws which earlier would have been regarded as quite untenable. Several smaller nations declared martial law, strictly rationed supplies, and then applied repressive penalties for any kind of wastage. There were sporadic demonstrations, but they soon fizzled out.

Now, only eight years later, the outrage and tensions had eased, in a seemingly miraculous—certainly *astonishing*—reversal of history. The disparate and entrenched ideological stances of the leading nations were dissolving into acts of genuine cooperation. The first fragile signs of a shift in attitude toward the arbitrariness and irrationality of nationalism were becoming evident. Another big shake-up: In 2018, NASA, ESA, and other space agencies were obliged to accept a new charter: to put space exploration on hold, until Earth, this most important and precious of planets, was clearly on its way to recovering from the ravages and excesses of more than a century of unmitigated abuse.

Understandably, most scientists and engineers in these agencies, and many academicians, were violently opposed to this directive; but the dire urgency of the prevailing circumstances soon forced them to come around. These reorganized agencies, and new ones such as the IOA, had, for about seven years now, applied their best minds to Project Milton—"Restoring Our Common Habitat and Heritage." The notion of freedom had taken on a new, more vital and self-disciplined meaning, and not a decade too soon.

Removing the visionizer headband, and putting aside her food tray, Niku settled down for the evening. Saying *RAISE*! to the SyntheDown pillow unit, it fine-tuned the personalized presets interactively to conform even more comfortably to her contours. She reached for the thin blue Actablet, which kept the details of her studies. Propping it up on her raised knees, it came alive at her touch. She struggled to refocus her mind for the evening's work: to be well prepared for next week's Daedalus Day presentation, she needed to get these notes, and the many results, into some semblance of order.

But almost immediately, Marcoloonie, her lithe and intelligent Abyssinian kitten, was vying for attention. Plunking himself on her chest, with no room for negotiation, he offered his wide eyes in place of the screen—an old trick of cats. Succumbing to the hypnosis, she planted her hands on his warm fur. Nose to nose, they shared air, and her racing mind dropped down to first gear. Half-waking sometime later, she found Marco asleep, curled in a contented ball at her feet. She sighed *LOWER*! to the pillow unit.

* * *

At mid-morning, Dr. Leif encountered her in the hall. "Hi, Niku! How are you? Y'know, I've been meaning for some time to ask you about the rest of your oscillator story. Do you have time for a good cup of coffee over at GalaxyBux?"

"Sure, Dr. Leif. I've been trying to put my notes into a logical order for the Daedelus talk next week, but so far I haven't made much progress. A chat with you would help a lot. Let me just pop back to the lab for my Actablet."

As they arrived at the coffee shop, the Greeter did its usual thing, except that now it also recognized Niku: "Good morning, Dr. Leif, Dr. Yeng. I trust your day is going well. We have a great special

Petitioners
Ex. 1062, p. 3

this morning. It's …" They didn't wait to hear about the day's special, nor the ad that was sure to follow. Finding a table for four, Niku sat beside her mentor, to more easily share the screen of her 'tablet. With coffee ordered, a touch brought up the usual gallery of icons. She selected one labeled *Oscar*.

"Oscar?" quizzed the amused senior. "Why Oscar?"

"Oh, I've become quite fond of this little circuit, so I gave him a name. Actually, the title I'll use for the talk is "Oscar Awakes," to emphasize that it's about the brief start-up phase, and how it's dominated by circuit noise in a well-designed and well-balanced circuit. But I'm afraid I got sidetracked into other issues."

"Ah ha. I see. So, where did we leave off?"

"Well, if you remember, I was taking each step cautiously, and the last time we talked I'd gotten as far as showing that when a fast step of current, I, is applied to a parallel LC tank, the circuit behaves as a pure resistance of $\sqrt{(L/C)}$, resulting in a voltage, $V = I\sqrt{(L/C)}$, which persists as the nondecaying amplitude of a sinusoid, while the magnetic and dielectric energies oscillate between LI^2 and CV^2. Of course, that never happens in a real circuit, mostly because of the finite series resistance of the inductor. But even if the elements were 'perfect,' some of the stored energy must be lost by radiation—mostly from the inductor, whose magnetic field, in the case of an open coil, generates an EM wave."

"Did you try calculating the radiation loss from first principles?"

"Hmm … I haven't given it any thought and I'm not sure I could. I'm trying to stay focused on the issue of the noise-excited start-up trajectory. I mean … there are so many other basic questions you could investigate. For example, intuition tells me that there's a dual—the series-connected LC tank, driven by an ideal voltage source. A fast step of V would then set up a circulating current of $V/\sqrt{(L/C)}$. Thinking about the Fundaments—if I may adopt your valuable idea—it couldn't be anything else, since L and C must determine the impedance, and, from a dimensional perspective, these have to be combined as $\sqrt{(L/C)}$. I don't think one needs to demonstrate that by separate simulations."

The coffee arrived. "Your intuitions serve you well, Niku. These kinds of *nonmathematical* analysis provide just one example of what I call the *What Must Be* philosophy. Understanding of crucial aspects of circuit design or behavior can often be gained simply by listing the known variables, then asking how the *dimensional constraints* can be satisfied. In the case of little Oscar's parallel-tuned tank, you knew there has to be a voltage, V, caused by the applied current, I, which demands a relationship via impedance, because $V = IZ$. So there's one *What Must Be* factor, right away.

"Then, when GE°E later showed you that the amplitude of the resulting sinusoid didn't change over time, it confirmed your intuition that 'Z' has to be a simple *resistance*. You realized that there's only one way of getting something with the dimension of resistance out of an L and a C, and that's from their ratio L/C—although this is actually R^2. After all, there is only one other simple dimensionally correct way of combining L and C—as the product, LC. This probably led you to think of two standard forms: $\sqrt{(L/C)}$, which has the dimension of resistance; and $\sqrt{(LC)}$, which has the dimension of time—or $1/\sqrt{(LC)}$, which has the dimension of angular frequency. One can easily arrive at these basic forms from first principles, considering the inherent dimensions of inductance and capacitance. This concern for *what has to be included and what can be excluded without significant loss* is philosophically related to the Principle of Occam's Razor."

"Mmm ...?" said Niku, sipping the dark brew. "Who is Occam, and what do his shaving habits have to do with a principle?"

"William of Ockham was a 14th-century … No, that can wait. Roughly speaking, he conjectured that when faced with two or more ways to explain some phenomenon, the most simple is usually the right one: the razor will pare away the rest. In the same spirit, faced with several ways to configure a basic circuit, in meeting some performance challenge, it often pays to use the simplest possible form. Of course, you can't *create* new circuit forms, from scratch, in this way, only *choose* from pre-existing possibilities. The *origination of topology* is a daunting aspect of analog design, and the chief reason why many find it Difficult.

"Still, Occam's Razor can nicely complement the philosophy of *intuitive design*, the most important aspect of which is the habit of ceaselessly asking *What If? How About? Why Is That? What Must Be?*—that kind of question—before plunging into a detailed and time-consuming analysis that may turn out to be irrelevant or fail to contribute any useful insights.

"A simulator is an indispensable partner in telling you *what* will happen— with high numerical accuracy—if that is all you need to know. But after seeing *what happens*, your eyes are invariably opened to *insights*; and these lead to even more deeply probing questions, as you're demonstrating. Sometime, I'd like to share some personal case histories of *What Must Be* reasoning. Many of them rely on elegance-driven heuristics like Occam's Razor. But I'm anxious to hear about Oscar. What *did* you do next?"

"Well, this actually *wasn't* the next thing I looked at, during my early studies; but I was curious and went back later to explore it. In fact, Dr. Leif, that's my dilemma. I don't know whether to include such details in my talk or not, because they are merely tangential to the original topic. But I will tell you about this one.

"I was curious about the effect of the *rise time*, t_R, of the current step. Only a pulse having a rise time of *zero* would generate the full $I\sqrt{(L/C)}$, so I wanted to discover the relationship between the rise time and the resulting amplitude of the tank voltage. What rise time would cause it to be reduced by, say, 10%? What if it's equal to a significant fraction of the period, up to or even beyond the full period? One thing was certain: for very slow rise times—when t_R is hundreds of times $\sqrt{(LC)}$—practically all of the current must flow in the inductor, generating a small voltage LdI/dt."

"Niku," interrupted Dr. Leif, touching her lightly on the wrist. "Y'know, these questions aren't all that far off the beaten track. After all, in a practical oscillator, the currents coming from the driving circuit *will* have finite rise times, and you might have uncovered an 'incidental' fact that would alter the direction of your studies—perhaps radically. The trick is knowing when to stop chasing every side thought. You can never suppress them. Even if one's thinking is clear and controlled, the mind will be nonetheless seething with questions continuously generated by our innate propensity for confabulation. Look, you have hardly touched your coffee. Please, take a moment to enjoy it."

But Niku rattled on as the coffee contributed to cosmic entropy.

"Yes, I really do find it difficult to cope with all the questions and ideas for new experiments that keep coming to mind! The question about the importance of the rise time led me to worry about the effect of the detailed *wave-shape* of the current step, since this obviously determines the frequency-domain spectrum of the voltage. How would my ideal, infinite-Q tank be affected by this? GE°E provides a rich assortment of forms, including the simple ramp [the PULSE source in SPICE—Ed.] and the raised-cosine form, both of which occupy a finite time; and it provides the Gaussian form, which must be truncated before $t = 0$, for any finite delay—though, it's only 0.148 ppm of the final value at a time four times earlier than that needed to get 10% up the edge.

Petitioners
Ex. 1062, p. 4

"Anyway, I thought it might be useful to explore these effects for an isolated tank, using a 20-nH inductor shunted by a 10-pF cap, which resonates at 355.88… MHz, this time using a single unit step of 1 A. As expected, the tank voltage for an extremely slow ramp is $dI(t)/dt$ times the inductance. But I didn't expect this …" Turning the screen of her Actablet to a better angle for Leif, she showed him the first of two sets of plots.

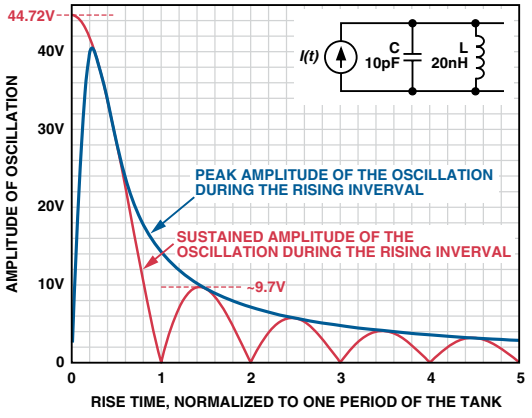

Figure 1.

"Let's see," he pondered. "The horizontal axis (Figure 1) shows the number of periods of the tank over which the current linearly rises to 1 A; and vertically, you're showing the peak amplitude of the oscillation both *during* and *after* the rising edge, right?" Niku nodded. "And I see that for fast-rising pulses, the stable *periodic-steady-state* amplitude of the sine wave is 44.72… V, which is consistent with your $\sqrt{(L/C)}$ value of 44.72… Ω. Then, as the rise time increases, the PSS amplitude steadily falls, becoming zero when it's exactly equal to one whole period of the tank. For slower rise times, it rises back up, to roughly 9.7 V, falling back to zero whenever the rise time is an integer number of periods. And your second set (Figure 2) shows the actual time-domain voltages for several rise times up to one full period. Okay. Now, what did you *expect* would happen?" Leif asked mischievously.

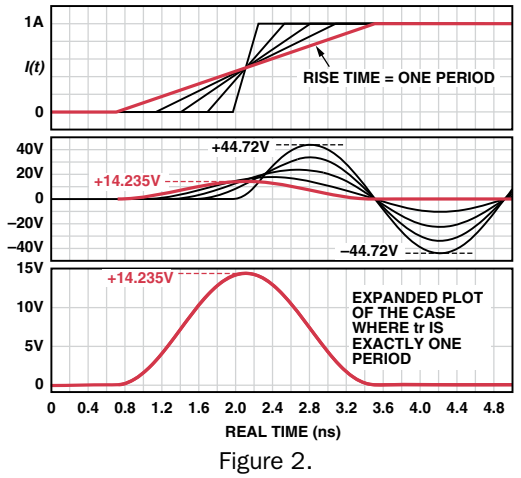

Figure 2.

"Well, until I saw this, I expected the amplitude to simply keep on falling as t_R increased. I assumed that, for a rise time as long as one tank period, it would be already slow enough to approach the point where the voltage would be just LdI/dt, and that a 1-A step with a rise rate of 355.88127 A/μs, when multiplied by the inductance of 20 nH, would make a rectangular voltage pulse of 7.117625 V. But the actual waveform seems to be a raised sine. Then I realized that if a step of voltage *did* appear, *immediately*, at the foot of the current ramp, there'd be a very large capacitor current, so my 'simple' expectation was just a bit too simplistic! On the other hand, the measured peak voltage of 14.235… V is *exactly twice* my erroneously predicted value, which happens to be exactly (1/π) × 44.72… V—isn't *that* intriguing! So until I have all my facts together, I won't touch on this in my talk."

"Well, yes; but it would be more intriguing if the factor weren't (1/π) but, say, exactly (1/7), because *then* you'd need a theory to explain where this new number comes from! Anyway, now that GE°E has shown you the error of your ways, I agree it wouldn't be wise to introduce this matter, unless you know what's going on," said Leif, his eyes twinkling. "I understand why you'd feel uncomfortable about not being able to get to the bottom of all these details. Still, at this point, I'm going to let you—how does it go?—stew in your own juices for a while!

"However, I *will* mention that this is one of those times when you will need to dust off your Laplace transforms, and use paper and pencil to find out *Why Is That*? Though simulators are excellent tools, and fine for showing *what* happens, they haven't a clue as to *why*. They aren't aware, as we are, of the Fundaments, and in spite of decades of promises from the AI folks, GE°E still can't emulate creativity in any useful way. That's why we hire people like you, Niku, and why we use the job description *Originator*. But I have a question for you to mull over: Do you recognize the *functional shape* of your plot of PSS amplitude vs. rise time?"

"Well, it *looks* like the absolute-sine-*x*-over-*x* form," Niku said, "but I don't know why. I suppose if I could trust Occam's Razor, that simple answer would probably be the right one."

"Mmm. It's not *really* a job for Occam, since you have only one suggestion! But it's a smart guess. Still, *why* would it have *that* form, and what causes all those nulls? Every effect has a cause."

His wink was the only encouragement she needed to strengthen her resolve to understand this by Daedalus Day. "Maybe I *will* include these analyses, after all," she beamed, confident of her practiced ability to wield the Laplacian sword. As they left, her coffee continued its spiraling descent to ambient. Walking back to the lab, under a yellowish sky, Niku said, "By the way, you still haven't told me what the Fourth Dee of analog design is."

"Ah, but I did," said Dr. Leif, "only a few moments ago."

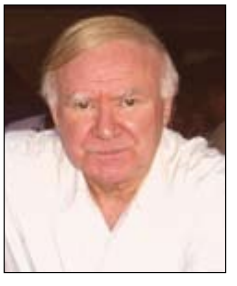

Barrie Gilbert, the first-appointed ADI Fellow, has "spent a lifetime in pursuit of analog excellence." Barrie was born in Bournemouth, England, in 1937. Before joining ADI, he worked with first-generation transistors at SRDE in 1954. At Mullard, Ltd. in the late '50s, he pioneered transistorized sampling oscilloscopes, and in 1964 became a leading 'scope designer at Tektronix. He spent two years as a group leader at Plessey Research Labs before joining Analog Devices in 1972, where he is now director of the Northwest Labs in Beaverton, Oregon. Barrie is a Life Fellow of the IEEE and has received numerous service awards. He has about 70 issued patents, has authored some 50 papers, is a reviewer for several professional journals, and is a co-author or co-editor of five books. In 1997, he was awarded an honorary doctorate of engineering from Oregon State University.

Petitioners
Ex. 1062, p. 5

Class D Audio Amplifiers: What, Why, and How

By Eric Gaalaas [eric.gaalaas@analog.com]

Class D amplifiers, first proposed in 1958, have become increasingly popular in recent years. What are Class D amplifiers? How do they compare with other kinds of amplifiers? Why is Class D of interest for audio? What is needed to make a "good" audio Class D amplifier? What are the features of ADI's Class D amplifier products? Find the answers to all these questions in the following pages.

Audio Amplifier Background

The goal of audio amplifiers is to reproduce input audio signals at sound-producing output elements, with desired volume and power levels—faithfully, efficiently, and at low distortion. Audio frequencies range from about 20 Hz to 20 kHz, so the amplifier must have good frequency response over this range (less when driving a band-limited speaker, such as a *woofer* or a *tweeter*). Power capabilities vary widely depending on the application, from milliwatts in headphones, to a few watts in TV or PC audio, to tens of watts for "mini" home stereos and automotive audio, to hundreds of watts and beyond for more powerful home and commercial sound systems—and to fill theaters or auditoriums with sound.

A straightforward analog implementation of an audio amplifier uses transistors in linear mode to create an output voltage that is a scaled copy of the input voltage. The forward voltage gain is usually high (at least 40 dB). If the forward gain is part of a feedback loop, the overall *loop gain* will also be high. Feedback is often used because high loop gain improves performance—suppressing distortion caused by nonlinearities in the forward path and reducing power-supply noise by increasing the power-supply rejection (PSR).

The Class D Amplifier Advantage

In a conventional transistor amplifier, the *output stage* contains transistors that supply the instantaneous continuous output current. The many possible implementations for audio systems include Classes A, AB, and B. Compared with *Class D* designs, the output-stage power dissipation is large in even the most efficient *linear* output stages. This difference gives Class D significant advantages in many applications because the lower power dissipation produces less heat, saves circuit board space and cost, and extends battery life in portable systems.

Linear Amplifiers, Class D Amplifiers, and Power Dissipation

Linear-amplifier output stages are directly connected to the speaker (in some cases via capacitors). If bipolar junction transistors (BJTs) are used in the output stage, they generally operate in the linear mode, with large collector-emitter voltages. The output stage could also be implemented with MOS transistors, as shown in Figure 1.

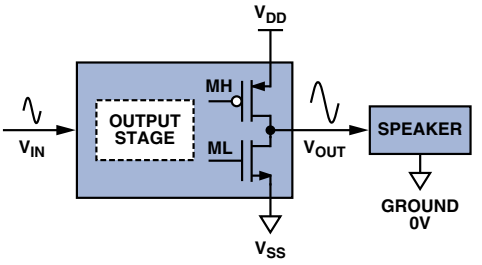

Figure 1. CMOS linear output stage.

Power is dissipated in all linear output stages, because the process of generating V_{OUT} unavoidably causes nonzero I_{DS} and V_{DS} in at least one output transistor. The amount of power dissipation strongly depends on the method used to bias the output transistors.

The *Class A* topology uses one of the transistors as a dc current source, capable of supplying the maximum audio current required by the speaker. Good sound quality is possible with the Class A output stage, but power dissipation is excessive because a large dc bias current usually flows in the output-stage *transistors* (where we do not want it), without being delivered to the *speaker* (where we do want it).

The *Class B* topology eliminates the dc bias current and dissipates significantly less power. Its output transistors are individually controlled in a push-pull manner, allowing the MH device to supply positive currents to the speaker, and ML to sink negative currents. This reduces output stage power dissipation, with only signal current conducted through the transistors. The Class B circuit has inferior sound quality, however, due to nonlinear behavior (*crossover distortion*) when the output current passes through 0 and the transistors are changing between the on and off conditions.

Class AB, a hybrid compromise of Classes A and B, uses some dc bias current, but much less than a pure Class A design. The small dc bias current is sufficient to prevent crossover distortion, enabling good sound quality. Power dissipation, although between Class A and Class B limits, is typically closer to Class B. Some control, similar to that of the Class B circuit, is needed to allow the Class AB circuit to supply or sink large output currents.

Unfortunately, even a well-designed class AB amplifier has significant power dissipation, because its midrange output voltages are generally far from either the positive or negative supply rails. The large drain-source voltage drops thus produce significant $I_{DS} \times V_{DS}$ instantaneous power dissipation.

Thanks to a different topology (Figure 2), the *Class D* amplifier dissipates much less power than any of the above. Its output stage switches between the positive and negative power supplies so as to produce a train of voltage pulses. This waveform is benign for power dissipation, because the output transistors have zero current when not switching, and have low V_{DS} when they are conducting current, thus giving smaller $I_{DS} \times V_{DS}$.

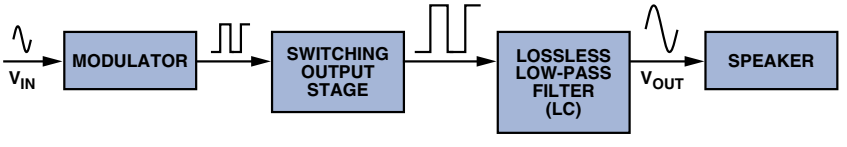

Figure 2. Class D open-loop-amplifier block diagram.

Petitioners
Ex. 1062, p. 6

Since most audio signals are not pulse trains, a modulator must be included to convert the audio input into pulses. The frequency content of the pulses includes both the desired audio signal and significant high-frequency energy related to the modulation process. A low-pass filter is often inserted between the output stage and the speaker to minimize electromagnetic interference (EMI) and avoid driving the speaker with too much high frequency energy. The filter (Figure 3) needs to be lossless (or nearly so) in order to retain the power-dissipation advantage of the switching output stage. The filter normally uses capacitors and inductors, with the only intentionally dissipative element being the speaker.

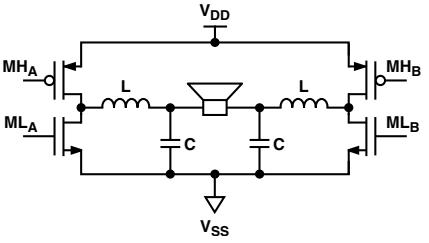

Figure 3. Differential switching output stage and LC low-pass filter.

Figure 4 compares ideal output-stage power dissipation (P_{DISS}) for Class A and Class B amplifiers with measured dissipation for the AD1994 Class D amplifier, plotted against power delivered to the speaker (P_{LOAD}), given an audio-frequency sine wave signal. The power numbers are normalized to the power level, P_{LOAD} *max*, at which the sine is clipped enough to cause 10% total harmonic distortion (THD). The vertical line indicates the P_{LOAD} at which clipping begins.

Significant differences in power dissipation are visible for a wide range of loads, especially at high and moderate values. At the onset of clipping, dissipation in the Class D output stage is about 2.5 times less than Class B, and 27 times less than Class A. Note that more power is consumed in the Class A output stage than is delivered to the speaker—a consequence of using the large dc bias current.

Output-stage power efficiency, *Eff*, is defined as

$$Eff = \frac{P_{LOAD}}{P_{LOAD} + P_{DIS}}$$

At the onset of clipping, *Eff* = 25% for the Class A amplifier, 78.5% for the Class B amplifier, and 90% for the Class D amplifier (see Figure 5). These best-case values for Class A and Class B are the ones often cited in textbooks.

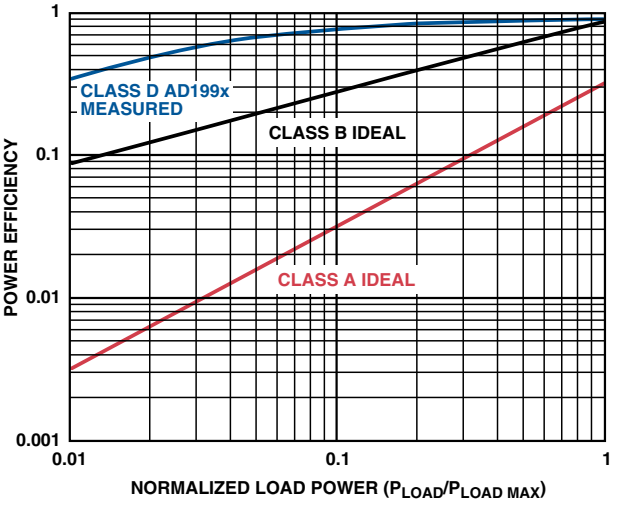

Figure 5. Power efficiency of Class A, Class B, and Class D output stages.

The differences in power dissipation and efficiency widen at moderate power levels. This is important for audio, because long-term average levels for loud music are much lower (by factors of five to 20, depending on the type of music) than the instantaneous peak levels, which approach P_{LOAD} *max*. Thus, for audio amplifiers, [$P_{LOAD} = 0.1 \times P_{LOAD}$ *max*] is a reasonable average power level at which to evaluate P_{DISS}. At this level, the Class D output-stage dissipation is nine times less than Class B, and 107 times less than Class A.

For an audio amplifier with 10-W P_{LOAD} *max*, an average P_{LOAD} of 1 W can be considered a realistic listening level. Under this condition, 282 mW is dissipated inside the Class D output stage, vs. 2.53 W for Class B and 30.2 W for Class A. In this case, the Class D efficiency is reduced to 78%—from 90% at higher power. But even 78% is much better than the Class B and Class A efficiencies—28% and 3%, respectively.

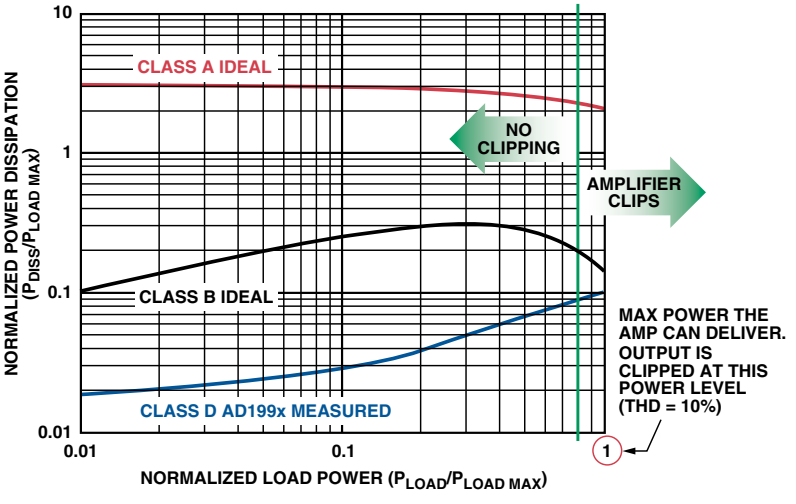

Figure 4. Power dissipation in Class A, Class B, and Class D output stages.

Petitioners
Ex. 1062, p. 7

These differences have important consequences for system design. For power levels above 1 W, the excessive dissipation of linear output stages requires significant cooling measures to avoid unacceptable heating—typically by using large slabs of metal as heat sinks, or fans to blow air over the amplifier. If the amplifier is implemented as an integrated circuit, a bulky and expensive thermally enhanced package may be needed to facilitate heat transfer. These considerations are onerous in consumer products such as flat-screen TVs, where space is at a premium—or automotive audio, where the trend is toward cramming higher channel counts into a fixed space.

For power levels below 1 W, wasted power can be more of a difficulty than heat generation. If powered from a battery, a linear output stage would drain battery charge faster than a Class D design. In the above example, the Class D output stage consumes 2.8 times less supply current than Class B and 23.6 times less than Class A—resulting in a big difference in the life of batteries used in products like cell phones, PDAs, and MP3 players.

For simplicity, the analysis thus far has focused exclusively on the amplifier *output* stages. However, when all sources of power dissipation in the amplifier system are considered, linear amplifiers can compare more favorably to Class D amplifiers at low output-power levels. The reason is that the power needed to generate and modulate the switching waveform can be significant at low levels. Thus, the system-wide quiescent dissipation of well-designed low-to-moderate-power Class AB amplifiers can make them competitive with Class D amplifiers. Class D power dissipation is unquestionably superior for the higher output power ranges, though.

Class D Amplifier Terminology, and Differential vs. Single-Ended Versions

Figure 3 shows a differential implementation of the output transistors and LC filter in a Class D amplifier. This *H-bridge* has two *half-bridge* switching circuits that supply pulses of opposite polarity to the filter, which comprises two inductors, two capacitors, and the speaker. Each half-bridge contains two output transistors—a high-side transistor (MH) connected to the positive power supply, and a low-side transistor (ML) connected to the negative supply. The diagrams here show high-side *p*MOS transistors. High-side *n*MOS transistors are often used to reduce size and capacitance, but special gate-drive techniques are required to control them (Further Reading 1).

Full H-bridge circuits generally run from a single supply (V_{DD}), with ground used for the negative supply terminal (V_{SS}). For a given V_{DD} and V_{SS}, the differential nature of the bridge means that it can deliver twice the output signal and four times the output power of single-ended implementations. Half-bridge circuits can be powered from bipolar power supplies or a single supply, but the single-supply version imposes a potentially harmful dc bias voltage, $V_{DD}/2$, across the speaker, unless a blocking capacitor is added.

The power-supply voltage buses of half-bridge circuits can be "pumped" beyond their nominal values by large inductor currents from the LC filter. The dV/dt of the pumping transient can be limited by adding large decoupling capacitors between V_{DD} and V_{SS}. Full-bridge circuits do not suffer from bus pumping, because inductor current flowing into one of the half-bridges flows out of the other one, creating a local current loop that minimally disturbs the power supplies.

Factors in Audio Class D Amplifier Design

The lower power dissipation provides a strong motivation to use Class D for audio applications, but there are important challenges for the designer. These include:

- Choice of output transistor size
- Output-stage protection
- Sound quality
- Modulation technique
- EMI
- LC filter design
- System cost

Choice of Output Transistor Size

The output transistor size is chosen to optimize power dissipation over a wide range of signal conditions. Ensuring that V_{DS} stays small when conducting large I_{DS} requires the on resistance (R_{ON}) of the output transistors to be small (typically 0.1 Ω to 0.2 Ω). But this requires large transistors with significant gate capacitance (C_G). The gate-drive circuitry that switches the capacitance consumes power—CV^2f, where C is the capacitance, V is the voltage change during charging, and f is the switching frequency. This "switching loss" becomes excessive if the capacitance or frequency is too high, so practical upper limits exist. The choice of transistor size is therefore a trade-off between minimizing $I_{DS} \times V_{DS}$ losses during conduction vs. minimizing *switching* losses. Conductive losses will dominate power dissipation and efficiency at high output power levels, while dissipation is dominated by switching losses at low output levels. Power transistor manufacturers try to minimize the $R_{ON} \times C_G$ product of their devices to reduce overall power dissipation in switching applications, and to provide flexibility in the choice of switching frequency.

Protecting the Output Stage

The output stage must be protected from a number of potentially hazardous conditions:

Overheating: Class D's output-stage power dissipation, though lower than that of linear amplifiers, can still reach levels that endanger the output transistors if the amplifier is forced to deliver very high power for a long time. To protect against dangerous overheating, temperature-monitoring control circuitry is needed. In simple protection schemes, the output stage is shut off when its temperature, as measured by an on-chip sensor, exceeds a *thermal-shutdown* safety threshold, and is kept off until it cools down. The sensor can provide additional temperature information, aside from the simple binary indication about whether temperature has exceeded the shutdown threshold. By measuring temperature, the control circuitry can gradually reduce the volume level, reducing power dissipation and keeping temperature well within limits—instead of forcing perceptible periods of silence during thermal-shutdown events.

Excessive current flow in the output transistors: The low on resistance of the output transistors is not a problem if the output stage and speaker terminals are properly connected, but enormous currents can result if these nodes are inadvertently short-circuited to one another, or to the positive or negative power supplies. If unchecked, such currents can damage the transistors or surrounding circuitry. Consequently, current-sensing output-transistor protection circuitry is needed. In simple protection schemes, the output stage is shut off if the output currents exceed a safety threshold. In more sophisticated schemes, the

Petitioners
Ex. 1062, p. 8

current-sensor output is fed back into the amplifier—seeking to limit the output current to a maximum safe level, while allowing the amplifier to run continuously without shutting down. In these schemes, shutdown can be forced as a last resort if the attempted limiting proves ineffective. Effective current limiters can also keep the amplifier running safely in the presence of momentarily large transient currents due to speaker resonances.

Undervoltage: Most switching output stage circuits work well only if the positive power-supply voltages are high enough. Problems result if there is an *undervoltage* condition, where the supplies are too low. This issue is commonly handled by an *undervoltage lockout* circuit, which permits the output stages to operate only if the power-supply voltages are above an undervoltage-lockout threshold.

Output transistor turn-on timing: The MH and ML output stage transistors (Figure 6) have very low *on* resistance. It is therefore important to avoid situations in which both MH and ML are on simultaneously, as this would create a low-resistance path from V_{DD} to V_{SS} through the transistors and a large *shoot-through* current. At best, the transistors will heat up and waste power; at worst, the transistors may be damaged. *Break-before-make* control of the transistors prevents the shoot-through condition by forcing both transistors off before turning one on. The time intervals in which both transistors are off are called *nonoverlap time* or *dead time*.

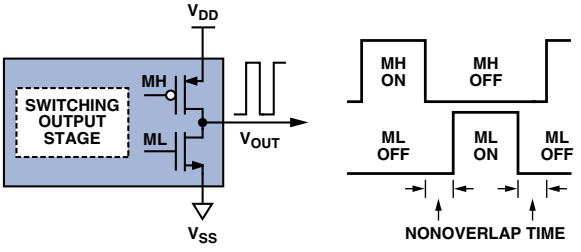

Figure 6. Break-before-make switching of output-stage transistors.

Sound Quality

Several issues must be addressed to achieve good overall sound quality in Class D amplifiers.

Clicks and pops, which occur when the amplifier is turning on or off, can be very annoying. Unfortunately, however, they are easy to introduce into a Class D amplifier unless careful attention is paid to modulator state, output-stage timing, and LC filter state when the amplifier is muted or unmuted.

Signal-to-noise ratio (SNR): To avoid audible hiss from the amplifier noise floor, SNR should typically exceed 90 dB in low-power amplifiers for portable applications, 100 dB for medium-power designs, and 110 dB for high-power designs. This is achievable for a wide variety of amplifier implementations, but individual noise sources must be tracked during amplifier design to ensure a satisfactory overall SNR.

Distortion mechanisms: These include nonlinearities in the modulation technique or modulator implementation—and the dead time used in the output stage to solve the shoot-through current problem.

Information about the audio signal level is generally encoded in the widths of the Class D modulator output pulses. Adding dead time to prevent output stage shoot-through currents introduces a nonlinear timing error, which creates distortion at the speaker in proportion to the timing error in relation to the ideal pulse width. The shortest dead time that avoids shoot-through is often best for minimizing distortion; see Further Reading 2 for a detailed design method to optimize distortion performance of switching output stages.

Other sources of distortion include: mismatch of rise and fall times in the output pulses, mismatch in the timing characteristics for the output transistor gate-drive circuits, and nonlinearities in the components of the LC low-pass filter.

Power-supply rejection (PSR): In the circuit of Figure 2, power-supply noise couples almost directly to the speaker with very little rejection. This occurs because the output-stage transistors connect the power supplies to the low-pass filter through a very low resistance. The filter rejects high-frequency noise, but is designed to pass all audio frequencies, including noise. See Further Reading 3 for a good description of the effect of power-supply noise in single-ended and differential switching output-stage circuits.

If neither distortion nor power-supply issues are addressed, it is difficult to achieve PSR better than 10 dB, or total harmonic distortion (THD) better than 0.1%. Even worse, the THD tends to be the bad-sounding high-order kind.

Fortunately, there are good solutions to these issues. Using feedback with high loop gain (as is done in many linear amplifier designs) helps a lot. Feedback from the LC filter input will greatly improve PSR and attenuate all non-LC filter distortion mechanisms. LC filter nonlinearities can be attenuated by including the speaker in the feedback loop. Audiophile-grade sound quality with PSR > 60 dB and THD < 0.01% is attainable in well-designed closed-loop Class D amplifiers.

Feedback complicates the amplifier design, however, because loop stability must be addressed (a nontrivial consideration for high-order design). Also, continuous-time analog feedback is necessary to capture important information about pulse timing errors, so the control loop must include analog circuitry to process the feedback signal. In integrated-circuit amplifier implementations, this can add to the die cost.

To minimize IC cost, some vendors prefer to minimize or eliminate analog circuit content. Some products use a digital open-loop modulator, plus an analog-to-digital converter to sense power-supply variations—and adjust the modulator's behavior to compensate, as proposed in Further Reading 3. This can improve PSR, but will not address any of the distortion problems. Other digital modulators attempt to precompensate for expected output stage timing errors, or correct for modulator nonidealities. This can at least partly address some distortion mechanisms, but not all. Applications that tolerate fairly relaxed sound-quality requirements can be handled by these kinds of open-loop Class D amplifiers, but some form of feedback seems necessary for best audio quality.

Modulation Technique

Class D modulators can be implemented in many ways, supported by a large quantity of related research and intellectual property. This article will only introduce fundamental concepts.

All Class D modulation techniques encode information about the audio signal into a stream of pulses. Generally, the pulse *widths* are linked to the amplitude of the audio signal, and the spectrum of the pulses includes the desired audio signal plus undesired (but unavoidable) high-frequency content. The total integrated high-frequency power in all schemes is roughly the same, since the total power in the time-domain waveforms is similar, and by Parseval's theorem, power in the time domain must equal power in the frequency domain. However, the distribution of energy varies widely: in some schemes, there are high energy tones atop a low noise floor, while in other schemes, the energy is shaped so that tones are eliminated but the noise floor is higher.

The most common modulation technique is *pulse-width modulation* (PWM). Conceptually, PWM compares the input audio signal to a triangular or ramping waveform that runs at

Petitioners
Ex. 1062, p. 9

a fixed carrier frequency. This creates a stream of pulses at the *carrier* frequency. Within each period of the carrier, the duty ratio of the PWM pulse is proportional to the amplitude of the audio signal. In the example of Figure 7, the audio input and triangular wave are both centered around 0 V, so that for 0 input, the duty ratio of the output pulses is 50%. For large positive input, it is near 100%, and it is near 0% for large negative input. If the audio amplitude exceeds that of the triangle wave, *full modulation* occurs, where the pulse train stops switching and the duty ratio within individual periods is either 0% or 100%.

PWM is attractive because it allows 100-dB or better audio-band SNR at PWM carrier frequencies of a few hundred kilohertz—low enough to limit switching losses in the output stage. Also, many PWM modulators are stable up to nearly 100% modulation, in concept permitting high output power—up to the point of overloading. However, PWM has several problems. First, the PWM process inherently adds distortion in many implementations (Further Reading 4); next, harmonics of the PWM carrier frequency produce EMI within the AM radio band; and finally, PWM pulse widths become very small near full modulation. This causes problems in most switching output-stage gate-driver circuits—with their limited drive capability, they cannot switch properly at the excessive speeds needed to reproduce short pulses with widths of a few nanoseconds. Consequently, full modulation is often unattainable in PWM-based amplifiers, limiting maximum achievable output power to something less than the theoretical maximum—which considers only power-supply voltage, transistor *on* resistance, and speaker impedance.

An alternative to PWM is *pulse-density modulation* (PDM), in which the number of pulses in a given time window is proportional to the average value of the input audio signal. Individual pulse widths cannot be arbitrary as in PWM, but are instead "quantized" to multiples of the modulator clock period. 1-bit sigma-delta modulation is a form of PDM.

Much of the high-frequency energy in sigma-delta is distributed over a wide range of frequencies—not concentrated in tones at multiples of a carrier frequency, as in PWM—providing sigma-delta modulation with a potential EMI advantage over PWM. Energy still exists at images of the PDM sampling clock frequency; but with typical clock frequencies from 3 MHz to 6 MHz, the images are outside the audio frequency band—and are strongly attenuated by the LC low-pass filter.

Another advantage of sigma-delta is that the minimum pulse width is one sampling-clock period, even for signal conditions approaching full modulation. This eases gate-driver design and allows safe operation to theoretical full power. Nonetheless, 1-bit sigma-delta modulation is not often used in Class D amplifiers (Further Reading 4) because conventional 1-bit modulators are only stable to 50% modulation. Also, at least 64× oversampling is needed to achieve sufficient audio-band SNR, so typical output data rates are at least 1 MHz, and power efficiency is limited.

Recently, *self-oscillating* amplifiers have been developed, such as the one in Further Reading 5. This type of amplifier always includes a feedback loop, with properties of the loop determining the switching frequency of the modulator, instead of an externally provided clock. High-frequency energy is often more evenly distributed than in PWM. Excellent audio quality is possible, thanks to the feedback, but the loop is self-oscillating, so it's difficult to synchronize with any other switching circuits, or to connect to digital audio sources without first converting the digital to analog.

The full-bridge circuit (Figure 3) can use "3-state" modulation to reduce differential EMI. With conventional differential operation, the output polarity of Half-bridge A must be opposite to that of Half-bridge B. Only two differential operating states exist: Output A high with Output B low; and A low with B high. Two additional common-mode states exist, however, in which both half-bridge outputs are the same polarity (both high or both low). One of these *common-mode* states can be used in conjunction with the differential states to produce 3-state modulation where the differential input to the LC filter can be positive, 0, or negative. The 0 state can be used to represent low power levels, instead of switching between the positive and negative state as in a 2-state scheme. Very little differential activity occurs in the LC filter during the 0 state, reducing differential EMI, although actually increasing common-mode EMI. The differential benefit only applies at low power levels, because the positive and negative states must still be used to deliver significant power to the speaker. The varying common-mode voltage level in 3-state modulation schemes presents a design challenge for closed-loop amplifiers.

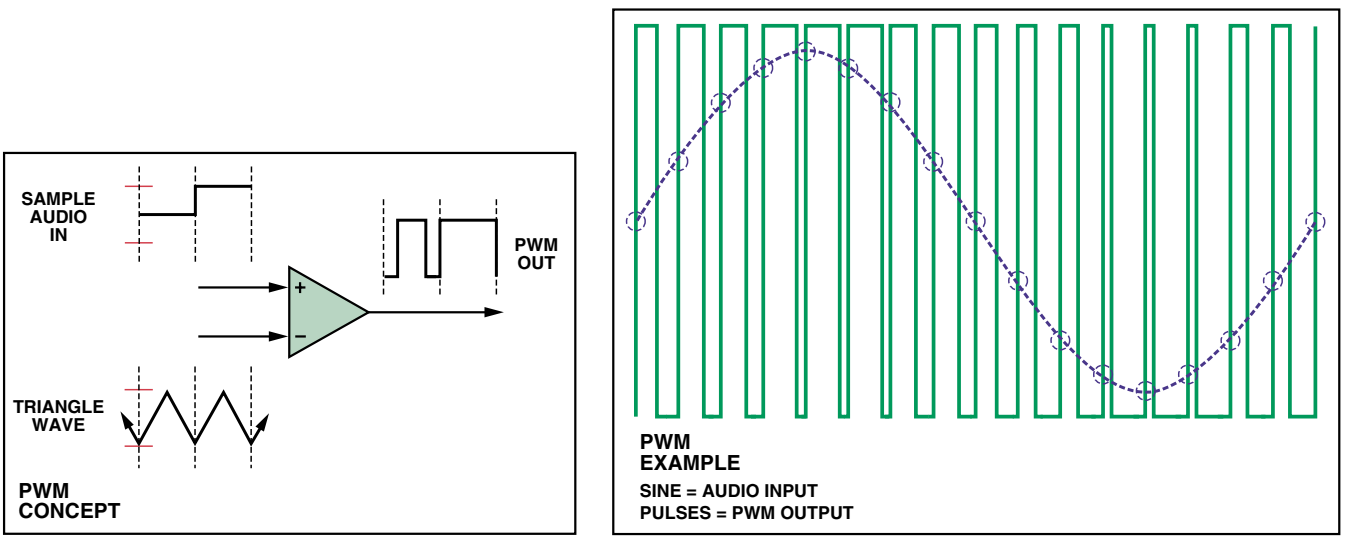

Figure 7. PWM concept and example.

Petitioners
Ex. 1062, p. 10

Taming EMI

The high-frequency components of Class D amplifier outputs merit serious consideration. If not properly understood and managed, these components can generate large amounts of EMI and disrupt operation of other equipment.

Two kinds of EMI are of concern: signals that are radiated into space and those that are conducted via speaker- and power-supply wires. The Class D modulation scheme determines a *baseline* spectrum of the components of conducted and radiated EMI. However, some board-level design techniques can be used to reduce the EMI emitted by a Class D amplifier, despite its baseline spectrum.

A useful principle is to minimize the area of loops that carry high-frequency currents, since strength of associated EMI is related to loop area and the proximity of loops to other circuits. For example, the entire LC filter (including the speaker wiring) should be laid out as compactly as possible, and kept close to the amplifier. Traces for current drive and return paths should be kept together to minimize loop areas (using twisted pairs for the speaker wires is helpful). Another place to focus is on the large charge transients that occur while switching the gate capacitance of the output-stage transistors. Generally this charge comes from a *reservoir* capacitance, forming a current loop containing both capacitances. The EMI impact of transients in this loop can be diminished by minimizing the loop area, which means placing the reservoir capacitance as closely as possible to the transistor(s) it charges.

It is sometimes helpful to insert RF chokes in series with the power supplies for the amplifier. Properly placed, they can confine high-frequency transient currents to local loops near the amplifier, instead of being conducted for long distances down the power-supply wires.

If gate-drive nonoverlap time is very long, inductive currents from the speaker or LC filter can forward-bias parasitic diodes at the terminals of the output-stage transistors. When the nonoverlap time ends, the bias on the diode is changed from forward to reverse. Large reverse-recovery current spikes can flow before the diode fully turns off, creating a troublesome source of EMI. This problem can be minimized by keeping the nonoverlap time very short (also recommended to minimize distortion of the audio). If the reverse-recovery behavior is still unacceptable, Schottky diodes can be paralleled with the transistor's parasitic diodes, in order to divert the currents and prevent the parasitic diode from ever turning on. This helps because the metal-semiconductor junctions of Schottky diodes are intrinsically immune to reverse-recovery effects.

LC filters with toroidal inductor cores can minimize stray field lines resulting from amplifier currents. The radiation from the cheaper *drum* cores can be reduced by shielding, a good compromise between cost and EMI performance—if care is taken to ensure that the shielding doesn't unacceptably degrade inductor linearity and sound quality at the speaker.

LC Filter Design

To save on cost and board space, most LC filters for Class D amplifiers are second-order, low-pass designs. Figure 3 depicts the differential version of a second-order LC filter. The speaker serves to damp the circuit's inherent resonance. Although the speaker impedance is sometimes approximated as a simple resistance, the actual impedance is more complex and may include significant reactive components. For best results in filter design, one should always seek to use an accurate speaker model.

A common filter design choice is to aim for the lowest bandwidth for which *droop* in the filter response at the highest audio frequency of interest is minimized. A typical filter has 40-kHz Butterworth response (to achieve a maximally flat pass band), if droop of less than 1 dB is desired for frequencies up to 20 kHz. The nominal component values in the table give approximate Butterworth response for common speaker impedances and standard L and C values:

Inductance L (μH)	Capacitance C (μF)	Speaker Resistance (Ω)	Bandwidth –3 dB (kHz)
10	1.2	4	50
15	1	6	41
22	0.68	8	41

If the design does not include feedback from the speaker, THD at the speaker will be sensitive to linearity of the LC filter components.

Inductor Design Factors: Important factors in designing or selecting the inductor include the core's current rating and shape, and the winding resistance.

Current rating: The core that is chosen should have a current rating above the highest expected amplifier current. The reason is that many inductor cores will magnetically saturate if current exceeds the current-rating threshold and flux density becomes too high—resulting in unwanted drastic reduction of inductance.

The inductance is formed by wrapping a wire around the core. If there are many turns, the resistance associated with the total wire length is significant. Since this resistance is in series between the half-bridge and the speaker, some of the output power will be dissipated in it. If the resistance is too high, use thicker wire or change the core to a different material that requires fewer turns of wire to give the desired inductance.

Finally, it should not be forgotten that the form of inductor used can affect EMI, as noted above.

System Cost

What are the important factors in the overall cost of an audio system that uses Class D amplifiers? How can we minimize the cost?

The *active* components of the Class D amplifier are the switching output stage and modulator. This circuitry can be built for roughly the same cost as an analog linear amplifier. The real trade-offs occur when considering other components of the system.

The lower dissipation of Class D saves the cost (and space) of cooling apparatus like heat sinks or fans. A Class D integrated-circuit amplifier may be able to use a smaller and cheaper package than is possible for the linear one. When driven from a digital audio source, analog linear amplifiers require D/A converters (DACs) to convert the audio into analog form. This is also true for analog-input Class D amplifiers, but digital-input types effectively integrate the DAC function.

On the other hand, the principal cost disadvantage of Class D is the LC filter. The components—especially the inductors—occupy board space and add expense. In high-power amplifiers, the overall system cost is still competitive, because LC filter cost is offset by large savings in cooling apparatus. But in cost-sensitive, low-power applications, the inductor expense becomes onerous. In extreme cases, such as cheap amplifiers for cell phones, an amplifier IC can be cheaper than the total LC filter cost. Also, even if the monetary cost is ignored, the board space occupied by the LC filter can be an issue in small form-factor applications.

Petitioners
Ex. 1062, p. 11

To address these concerns, the LC filter is sometimes eliminated entirely, to create a *filterless* amplifier. This saves cost and space, though losing the benefit of low-pass filtering. Without the filter, EMI and high-frequency power dissipation can increase unacceptably—unless the speaker is inductive and kept very close to the amplifier, current-loop areas are minimal, and power levels are kept low. Though often possible in portable applications like cell phones, it is not feasible for higher power systems such as home stereos.

Another approach is to minimize the number of LC filter components required per audio channel. This can be accomplished by using single-ended half-bridge output stages, which require half the number of Ls and Cs needed for differential, full-bridge circuits. But if the half-bridge requires bipolar power supplies, the expense associated with generating the negative supply may be prohibitive, unless a negative supply is already present for some other purpose—or the amplifier has enough audio channels, to amortize the cost of the negative supply. Alternatively, the half-bridge could be powered from a single supply, but this reduces output power and often requires a large dc blocking capacitor.

Analog Devices Class D Amplifiers

All of the design challenges just discussed can add up to a rather demanding project. To save time for the designer, Analog Devices offers a variety of Class D amplifier integrated circuits,[1] incorporating programmable-gain amplifiers, modulators, and power output stages. To simplify evaluation, demonstration boards are available for each amplifier type to simplify evaluation. The PCB layout and bill-of-materials for each of these boards serve as a workable reference design, helping customers quickly design working, cost-effective audio systems without having to "reinvent the wheel" to solve the major Class D amplifier design challenges.

Consider, for example, the AD1990,[2] AD1992,[3] and AD1994,[4]—a family of dual-amplifier ICs, targeted at moderate-power stereo or mono applications requiring two channels with output-per-channel of up to 5-, 10-, and 25 W, respectively. Here are some properties of these ICs:

> The AD1994 Class D audio power amplifier combines two programmable-gain amplifiers, two sigma-delta modulators, and two power-output stages to drive full H-bridge-tied loads in home theater-, automotive-, and PC audio applications. It generates switching waveforms that can drive stereo speakers at up to 25 W per speaker, or a single speaker to 50 W monophonic, with 90% efficiency. Its single-ended inputs are applied to a programmable-gain amplifier (PGA) with gains settable to 0-, 6-, 12-, and 18 dB, to handle low-level signals.
>
> The device has integrated protection against output-stage hazards of overheating, overcurrent, and shoot-through current. There are minimal clicks and pops associated with muting, thanks to special timing control, soft start, and dc offset calibration. Specifications include 0.001% THD, 105-dB dynamic range, and >60 dB PSR, using continuous-time analog feedback from the switching output stage and optimized output stage gate drive. Its 1-bit sigma-delta modulator is especially enhanced for the Class D application to achieve average data frequency of 500 kHz, with high loop gain to 90% modulation, and stability to full modulation. A *standalone* modulator mode allows it to drive external FETs for higher output power.
>
> It uses a 5-V supply for the PGA, modulator, and digital logic, and a high-voltage supply from 8 V to 20 V for the switching output stage. The associated reference design meets FCC Class B EMI requirements. When driving 6 Ω loads with 5-V and 12-V supplies, the AD1994 dissipates 487 mW quiescently, 710 mW at the 2 × 1-W output level, and 0.27 mW in *power-down* mode. Available in a 64-lead LFCSP package, it is specified from –40°C to +85°C

More technical information about Class D amplifiers—including implementations with Blackfin processors—can be found in the Further Reading section.

ACKNOWLEDGEMENTS

The author would like to thank Art Kalb and Rajeev Morajkar of Analog Devices for their thoughtful inputs to this article.

FURTHER READING

1. International Rectifier, Application Note AN-978, "HV Floating MOS-Gate Driver ICs."
2. Nyboe, F., et al, "Time Domain Analysis of Open-Loop Distortion in Class D Amplifier Output Stages," presented at the AES 27th International Conference, Copenhagen, Denmark, September 2005.
3. Zhang, L., et al, "Real-Time Power Supply Compensation for Noise-Shaped Class D Amplifier," Presented at the 117th AES Convention, San Francisco, CA, October 2004.
4. Nielsen, K., "A Review and Comparison of Pulse-Width Modulation (PWM) Methods for Analog and Digital Input Switching Power Amplifiers," Presented at the 102nd AES Convention, Munich, Germany, March 1997.
5. Putzeys, B., "Simple Self-Oscillating Class D Amplifier with Full Output Filter Control," Presented at the 118th AES Convention, Barcelona, Spain, May 2005.
6. Gaalaas, E., et al, "Integrated Stereo Delta-Sigma Class D Amplifier," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, December 2005, pp. 2388-2397. About the AD199x Modulator.
7. Morrow, P., et al, "A 20-W Stereo Class D Audio Output Stage in 0.6 μm BCDMOS Technology," *IEEE J. Solid-State Circuits*, vol. 39, no. 11, November 2004, pp. 1948-1958. About the AD199x Switching Output Stage.
8. *PWM and Class-D Amplifiers with ADSP-BF535 Blackfin® Processors*, Analog Devices Engineer-to-Engineer Note EE-242. ADI website: www.analog.com (Search) EE-242 (Go)

REFERENCES—VALID AS OF JUNE 2006

[1] http://www.analog.com/en/content/0,2886,759_5F1075_5F57704,00.html
[2] ADI website: www.analog.com (Search) AD1990 (Go)
[3] ADI website: www.analog.com (Search) AD1992 (Go)
[4] ADI website: www.analog.com (Search) AD1994 (Go)

This article can be found at http://www.analog.com/library/analogdialogue/archives/40-06/class_d.html, with a link to a PDF.

Petitioners
Ex. 1062, p. 12

Power-Supply Management—Principles, Problems, and Parts

By Alan Moloney [alan.moloney@analog.com]

INTRODUCTION

Power-supply designers are using flexible supply monitoring, sequencing, and adjustment circuits to manage their systems.[1] This article discusses why and how.

The monitoring and control of a growing number of power-supply voltage rails has been vital for safety, economy, durability, and proper operation of electronic systems for many years—especially for systems employing microprocessors. Determining whether a voltage rail is above a threshold or within an operating window—and whether that voltage is powered on or off in the correct sequence with respect to the other rails—is crucial to operational reliability and safety.

Many methods exist to solve various aspects of this problem. For example, a simple circuit using a precision resistive divider, comparator, and reference can be used to determine whether the voltage on a rail is above or below a certain level. In *reset generators*, such as the ADM803,[2] these elements are combined with a delay element to hold devices—such as microprocessors, application-specific ICs (ASICs), and digital signal processors (DSPs)—in *reset* while powering up. This level of monitoring is adequate for many applications.

Where multiple rails need to be monitored, multiple devices (or multichannel comparators and their associated circuitry) are used in parallel, but increasing opportunities call for monitoring ICs that do more than simple threshold comparison.

For example, consider a common requirement for power-supply sequencing: an FPGA (field-programmable gate-array) manufacturer may specify that the 3.3-V core voltage must be applied 20 ms before the 5-V I/O (input/output) voltage to avoid possible damage when the device is powered up. Meeting such sequencing requirements may be as crucial for reliability as keeping the device's supply voltage and temperature within specified operating limits.

Also, the number of power rails in many applications has increased dramatically. Complex, expensive systems, such as LAN switches and cellular base stations, commonly have line cards with 10 or more voltage rails; but even cost-sensitive consumer systems, such as plasma TVs, can have as many as 15 separate voltage rails, many of which may require monitoring and sequencing.

Many high-performance ICs now require multiple voltages. For example, separate core- and I/O voltages are standard for many devices. At the high end, DSPs may require up to four separate supplies per device. In many cases, numerous multisupply devices can coexist in a single system that contains FPGAs, ASICs, DSPs, microprocessors, and microcontrollers (as well as analog components).

Many devices share standard voltage levels (such as 3.3 V), while others may require device-specific voltages. In addition, a particular standard voltage level may have to be independently furnished in numerous places. For example, separate analog- and digital supplies, such as 3.3 V_{ANALOG} and 3.3 $V_{DIGITAL}$, may be required. Generating the same voltage numerous times may be necessary to improve efficiency (e.g., memory rails running at hundreds of amperes) or to meet sequencing requirements (3.3 V_A and 3.3 V_B needed by separate devices at different times). All of these factors contribute to the proliferation of voltage sources.

Voltage monitoring and sequencing can become quite complex, especially if a system must be designed to support a power-up sequence, a power-down sequence, and multiple responses to all possible fault conditions on the various supply rails at different points during operation. A central power management controller is the best way to solve this problem.

As the number of supply voltages increases, there is a much higher probability of things going wrong. The risk increases in proportion to the number of supplies, number of elements, and complexity of the system. External factors also add risk. If, for example, the main ASIC is not completely characterized at the time of the initial design, the power-supply designer must commit to hardwiring voltage-monitoring thresholds and timing sequences that are subject to change as the ASIC specifications are developed. If the requirements change, the PC board may have to be revised—with obvious schedule- and cost implications. Furthermore, the supply voltage specifications for certain devices may change during their development. In such circumstances, a way to readily adjust supplies would be useful to any central power-system manager. In fact, the flexibility to monitor, sequence, and adjust the voltage rails of such systems is a vital necessity.

Evaluating the robustness of the chosen fault protection and timing sequence can be a sizable job, so a device that simplifies this process will speed up board evaluation and reduce time to market. Fault logging and digitized voltage and temperature data are useful features, both in the field and in all phases of design from early PCB development through prototype evaluation.

Basic Monitoring

Figure 1 shows a simple method for monitoring multiple voltage rails using the ADCMP354[3] comparator and reference IC. An individual circuit is used for each rail. Resistive dividers scale the voltage rails down, setting an undervoltage trip point for each supply. All outputs are tied together to generate a common *power-good* signal.

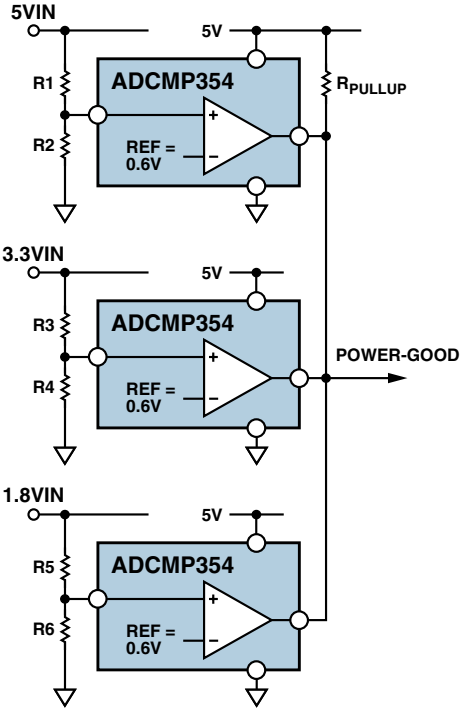

Figure 1. Comparator-based undervoltage detection with common power-good output for a three-supply system.

Petitioners
Ex. 1062, p. 13

Basic Sequencing

Figure 2 shows how basic sequencing can be implemented with discrete components, using logic thresholds instead of comparators. The 12-V and 5-V rails have been generated elsewhere. A time delay must be introduced to ensure that the system operates correctly. This is achieved by using a resistor-capacitor (RC) combination to slowly ramp the gate voltage on the *n*-channel FET in series with the 5-V supply. The RC values are chosen to ensure a sufficient time delay before the FET reaches its voltage threshold and begins to turn on. The 3.3-V and 1.8-V rails are generated with ADP3330[4] and ADP3333[5] low-dropout (LDO) regulators. The turn-on times of these voltages are also sequenced by RCs. No series FET is required, as the RC drives the shutdown ($\overline{SD}$) pin of each LDO. The RC values are chosen to ensure sufficient time delays (t_2, t_3) before the voltages on the $\overline{SD}$ pins climb above their thresholds.

This simple, low-cost approach to sequencing power supplies uses little board area and is perfectly acceptable in many applications. It is suitable for systems where cost is the primary driver, the sequencing requirements are simple, and the accuracy of the sequencing circuit is not critical.

But many situations require more accuracy than is available with RC lag circuits. In addition, this simple solution does not permit faults to be dealt with in a structured way (e.g., a 5-V supply failure will eventually bring down the other rails).

Sequencing with ICs

Figure 3 shows how the ADM6820[6] and ADM1086[7] power-supply-sequencing ICs can be used to accurately and reliably sequence power rails in a similar system. Internal comparators detect when a voltage rail goes above a precisely set level. The outputs are asserted after programmable turn-on delays, enabling the ADP3309[8] and ADP3335[9] regulators in the desired sequence. The thresholds are established by resistance ratios; the delay is established by a capacitor.

A wide variety of power-supply-sequencing ICs[10] is available. Some devices have outputs that can be used to enable power modules directly, and numerous output configurations are available. Some include on-board *charge-pump* voltage generators. This is especially useful for low-voltage systems that need to sequence rails that are generated upstream but lack a high-voltage source, such as a 12-V rail, to drive an *n*-channel FET gate. Many of these devices also have *enable* pins to allow an external signal—from a push-button switch or a controller—to restart the sequence or shut the controlled rails off when required.

Integrated Power System Management

Some systems have so many power-supply rails that discrete approaches that use a large number of ICs and set timing and threshold levels with resistors and capacitors become too complex and costly, and cannot provide adequate performance.

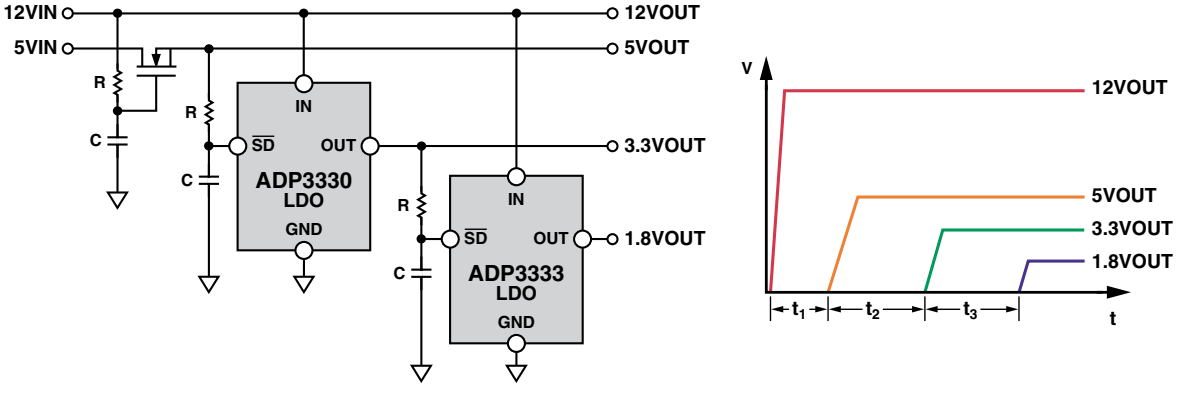

Figure 2. Basic discrete sequencing for a four-supply system.

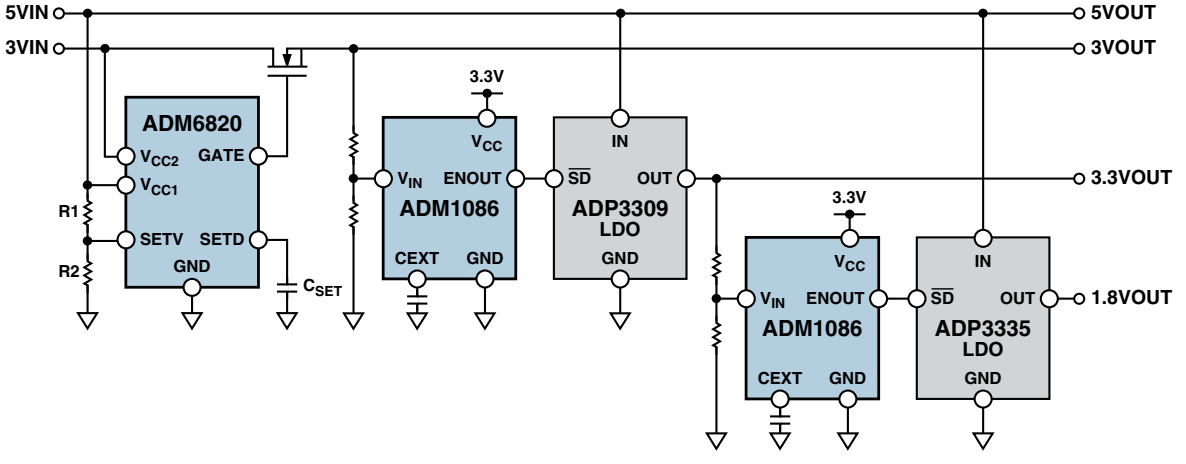

Figure 3. Sequencing a four-supply system with monitoring ICs.

Petitioners
Ex. 1062, p. 14

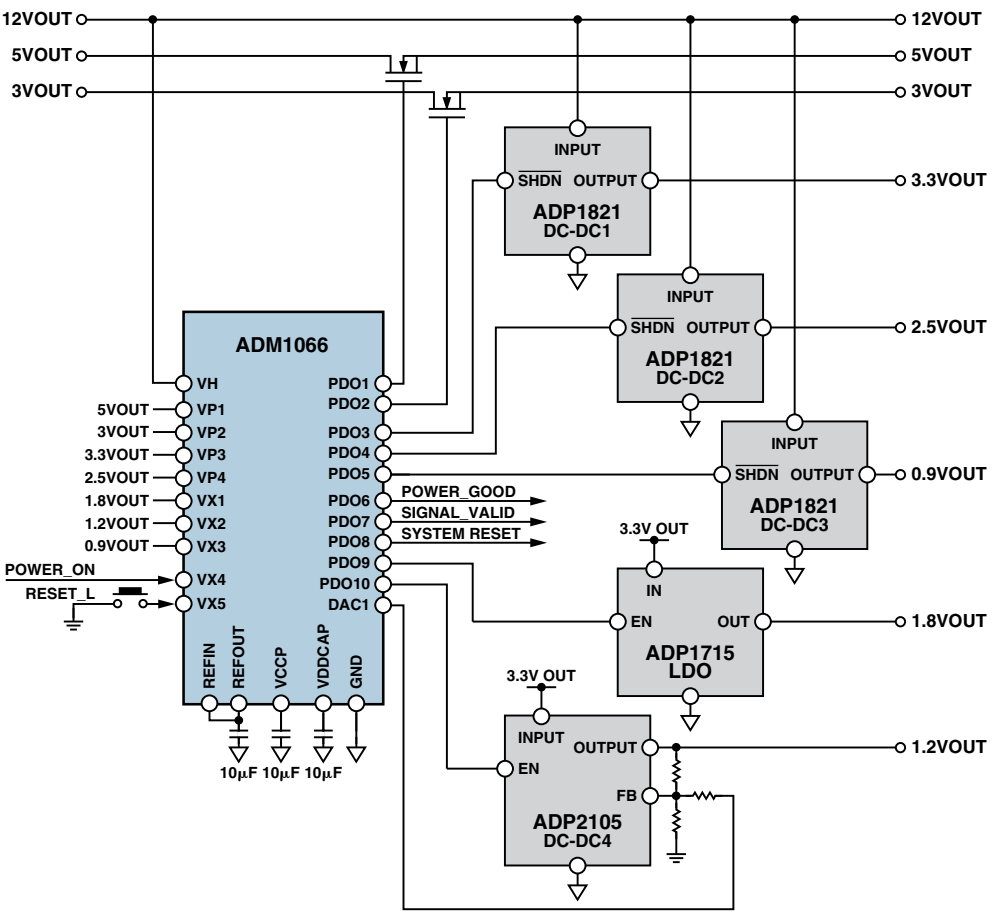

Figure 4. A centralized sequencing and monitoring solution for an eight-supply system.

Consider a system with eight voltage rails that requires a complex power-up sequence. Each rail must be monitored for undervoltage and overvoltage faults. In the event of a fault, all voltages could be turned off, or a power-down sequence could be initiated, depending on the failure mechanism. Actions must be taken depending upon the state of the control signals, and flags must be generated depending on the state of the power supplies. Implementing a circuit of this complexity with discrete devices and simple ICs may require hundreds of individual components, a huge amount of board space, and a significant combined cost.

In systems with four or more voltages, it may make sense to use a centralized device to manage the power supplies. An example of this approach can be seen in Figure 4.

Centralized Monitoring and Sequencing

The ADM106x Super Sequencer™[11] family continues to use comparators, but with some important differences. Two comparators are dedicated to each input so that undervoltage and overvoltage detection can be implemented, thus providing windowed monitoring for the rails created by the ADP1821[12] and ADP2105 dc-to-dc converters and the ADP1715 LDO. An undervoltage fault is the normal condition of a rail before it powers up, so this indication is used for sequencing. An overvoltage condition usually indicates a critical fault—such as a shorted FET or inductor—and calls for immediate action.

Systems with higher supply counts usually have greater complexity, and thus have tighter accuracy constraints. Also, setting accurate thresholds with resistors becomes challenging at lower voltages, such as 1.0 V and 0.9 V. Although a 10% tolerance may be acceptable on a 5-V rail, this tolerance is generally insufficient on a 1-V rail. The ADM1066 allows input detector comparator thresholds to be set within 1% worst case, independent of the voltage (as low as 0.6 V)—and across the entire temperature range of the device. It adds internal glitch filtering and hysteresis to each comparator. Its logic inputs can be used to start the power-up sequence, shut down all rails, or perform other functions.

The information from the bank of comparators, fed into a powerful and flexible stage machine core, can be used for various purposes:

Sequencing: When the output voltage of a recently enabled supply comes into a window, a time delay can be triggered to turn on the next rail in the power-up sequence. Complex sequencing, with multiple power-up and power-down sequences, or vastly different sequences for power-up and power-down are possible.

Timeout: If a rail that has been enabled does not come on as expected, a suitable course of action can be taken (such as generating an interrupt or shutting down the system). A purely analog solution would simply hang at that point in the sequence.

Monitoring: If the voltage on any rail moves out of the preset window, a suitable course of action can be taken—depending on the rail that faulted, the type of fault that occurred, and the current operating mode. Systems with more than five supplies are often expensive, so comprehensive fault protection is crucial.

Petitioners
Ex. 1062, p. 15

An on-board charge pump is used to generate approximately 12 V of gate drive, even if the highest available system voltage is as low as 3 V, allowing outputs to directly drive series *n*-channel FETs. Additional outputs enable or shut down dc-to-dc converters or regulators, allowing an output to internally pull up to one of the inputs or the on-board regulated voltage. The outputs can also be asserted open-drain. Outputs may also be used as status signals such as *power good* or *power-on* reset. Status LEDs can be directly driven from the outputs if required.

Supply Adjustment

In addition to monitoring multiple voltage rails and providing a solution for complex sequencing, integrated power management devices, such as the ADM1066, also provide the tools to temporarily or permanently adjust individual rail voltages. The voltage output of a dc-to-dc converter or regulator can be altered by adjusting the voltage at the trim or feedback node of that device. Typically, a resistive divider between the output and ground of the module sets a nominal voltage at the trim/feedback pin. This, in turn, sets a nominal output voltage. Simple schemes involving switching extra resistors or controlling variable resistances in the feedback loop will alter the trim/feedback voltage and hence adjust the output voltage.

The ADM1066 is equipped with digital-to-analog converters (DACs) to provide direct control over the trim/feedback node. For maximum efficiency, these DACs do not operate between ground and a maximum voltage; instead they operate across a relatively narrow window centered on the nominal trim/feedback level. The value of an attenuation resistor scales the incremental change in the output of the power module with each LSB change of the DAC. This open-loop adjustment provides margin-up and margin-down levels equivalent to those obtained by digital resistance switching in the reference circuit, and will adjust the output to a similar accuracy.

The ADM1066 also includes a 12-bit analog-to-digital converter (ADC) to measure the supply voltage, so a *closed-loop* supply adjustment scheme can be implemented. With a given DAC output setting, the voltage output of the power module is digitized by the ADC and compared with the target voltage in software. The DAC can then be adjusted to calibrate the voltage output as closely as possible to the target voltage. This closed-loop scheme provides a very accurate method for supply adjustment. With a closed-loop method, the accuracy of the external resistors is completely irrelevant. In Figure 4, the output voltage of DC-DC4 is adjusted by one of the on-chip DACs.

There are two primary uses for the supply adjustment scheme. The first is the concept of *margining* the supplies, i.e., testing the system's response to operating its supplies at the margins of the specified supply voltage range of the equipment. Manufacturers of datacom, telecom, cellular infrastructure, server, and storage area network equipment are required to test their systems strenuously before shipping to their end customers. All of the power supplies in the system will be specified to operate with a certain tolerance (e.g., ±5%, ±10%). Margining allows all the supplies on board to be adjusted to the high-end and low-end of the tolerance range, with tests performed to ensure correct operation. A centralized power management device with supply-adjustment capability can be used to do this margining test while minimizing the need for the extra components and PCB area required to perform a function that is only needed once—during the margin test at the manufacturer's test site.

Four-*corners* testing, i.e., testing across the operational voltage and temperature range of the equipment, is often required, so the ADM1062[13] integrates temperature-sensing and readback in addition to closed-loop power-supply margining circuitry.

The second use for the supply adjustment scheme is to compensate for system supply variations in the field. There are many causes for such variation. Short term, it is quite common for the voltages to change slightly as the temperature changes. Longer term, some of the component values may drift slightly over the life of the product, which can result in voltage drift. The ADC and DAC loops can be activated periodically (e.g., every 10, 30, or 60 seconds) in conjunction with a software calibration loop to keep the voltages where they should be.

Flexibility

The ADM1066 has on-board nonvolatile memory, allowing it to be reprogrammed as many times as necessary, while the sequencing and monitoring needs of the system evolve during the development process. This means that the *hardware* design can be completed early in the prototype process, and optimization of the monitoring and sequencing can be done as the project progresses.

Functions such as digital temperature- and voltage measurement simplify and speed up the evaluation process. Margining tools will allow adjustment of the voltage rails during the development cycle. So in a situation where a key ASIC, FPGA, or processor is also in development, and the supply-voltage levels or sequencing requirements are in a state of flux as new silicon revisions are shipped, a simple adjustment can be performed via the software[14] GUI. Thus, the power management device can be reprogrammed in a few minutes to take the changes into account, without the need to physically change components on the board or—worse still—redesign the hardware.

CONCLUSION

The increasing number of voltage rails and the emergence of power-supply sequencing have increased the demands on the power designer in all sorts of devices and systems—everything from notebook PCs, set-top boxes, and automotive systems to servers and storage, cellular base stations, and Internet-routing and switching systems. More stringent testing procedures, new levels of information gathering, and quick and simple programmability are also of interest, especially in mid- to high-end systems. For increased robustness and reliability, and the addition of these vital new features, there are many new power management integrated circuits available to help solve these problems safely, efficiently, and with minimum board area, while reducing time to market.

REFERENCES—VALID AS OF JULY 2006

[1]http://www.analog.com/en/cat/0,2878,766,00.html
[2]ADI website: www.analog.com (Search) ADM803 (GO)
[3]ADI website: www.analog.com (Search) ADCMP354 (GO)
[4]ADI website: www.analog.com (Search) ADP3330 (GO)
[5]ADI website: www.analog.com (Search) ADP3333 (GO)
[6]ADI website: www.analog.com (Search) ADM6820 (GO)
[7]ADI website: www.analog.com (Search) ADM1086 (GO)
[8]ADI website: www.analog.com (Search) ADP3309 (GO)
[9]ADI website: www.analog.com (Search) ADP3335 (GO)
[10]http://www.analog.com/Analog_Root/static/promotions/sequencing/index.html
[11]ADI website: www.analog.com (Search) ADM1066 (GO)
[12]ADI website: www.analog.com (Search) ADP1821 (GO)
[13]ADI website: www.analog.com (Search) ADM1062 (GO)
[14]http://www.analog.com/en/content/0,2886,766_821_70849,00.html

This article can be found at http://www.analog.com/library/analogdialogue/archives/40-05/power.html, with a link to a PDF.

Petitioners
Ex. 1062, p. 16

Design Your Own VoIP Solution with a Blackfin® Processor—Add Enhancements Later

By David Katz [david.katz@analog.com]
Tomasz Lukasiak [tomasz.lukasiak@analog.com]
Rick Gentile [richard.gentile@analog.com]
Wayne Meyer [wayne.meyer@analog.com]

INTRODUCTION

The age of *voice-over-Internet-protocol* (VoIP) is here, bringing together telephony and data communications to provide packetized voice and fax data streamed over low-cost Internet links. The transition from circuit-switched to packet-switched networking, continuing right now at breakneck speed, is encouraging applications that go far beyond simple voice transmission, embracing other forms of data and allowing them to all travel over the same infrastructure.

The VoIP challenge to the embedded-system designer is to choose a processing solution that is cost-effective, easy to deploy, and scalable in performance across market spaces. A "sweet-spot" embedded-solution approach is to design with a platform that can implement a low-channel-count basic VoIP solution, yet retain plenty of capacity for value-added capabilities and services—like video, music, imaging, and system control. The discussion below makes the case that the Blackfin[1] processor family from Analog Devices offers just such an attractive solution.

What Is VoIP?

Today's voice networks—such as the *public switched telephone network* (PSTN)—utilize digital switching technology to establish a dedicated link between the caller and the receiver. While this connection offers only limited bandwidth, it does provide an acceptable quality level without the burden of a complicated encoding algorithm.

The VoIP alternative uses *Internet protocol* (IP) to send digitized voice traffic over the Internet or private networks. An IP *packet* consists of a train of digits containing a control header and a data payload. The header provides network navigation information for the packet, and the payload contains the compressed voice data.

While circuit-switched telephony deals with the entire message, VoIP-based data transmission is packet-based, so that *chunks* of data are *packetized* (separated into units for transmission), compressed, and sent across the network—and eventually re-assembled at the designated receiving end. The key point is that *there is no need for a dedicated link between transmitter and receiver*.

Packetization is a good match for transporting data (for example, a JPEG file or email) across a network, because the delivery falls into a non-time-critical "best-effort" category. The network efficiently moves data from multiple sources across the same medium. For voice applications, however, "best-effort" is not adequate, because variable-length delays as the packets make their way across the network can degrade the quality of the decoded audio signal at the receiving end. For this reason, VoIP protocols, via QoS (quality-of-service) techniques, focus on managing network bandwidth to prevent delays from degrading voice quality.

Packetizing voice data involves adding header and trailer information to the data blocks. Packetization *overhead* (additional time and data introduced by this process) must be reduced to minimize added latencies (time delays through the system). Therefore, the process must achieve a balance between minimizing transmission delay and using network bandwidth most efficiently—smaller size allows packets to be sent more often, while larger packets take longer to compose. On the other hand, larger packets amortize the header and trailer information across a bigger chunk of voice data, so they use network bandwidth more efficiently than do smaller packets.

By their nature, networks cause the rate of data transmission to vary quite a bit. This variation, known as *jitter*, is removed by buffering the packets long enough to ensure that the slowest packets arrive in time to be decoded in the correct sequence. Naturally, a larger jitter buffer contributes to more overall system latency.

As mentioned above, *latency* represents the time delay through the IP system. A *one-way latency* is the time from when a word is spoken to when the person on the other end of the call hears it. *Round-trip latency* is simply the sum of the two one-way latencies. The lower the latency value, the more natural a conversation will sound. For the PSTN phone system in North America, the round-trip latency is less than 150 ms.

For VoIP systems, a one-way latency of up to 200 ms is considered acceptable. The largest contributors to latency in a VoIP system are the network and the gateways at either end of the call. The *voice codec* (coder-decoder) adds some latency—but this is usually small by comparison (<20 ms).

When the delay is large in a voice network application, the main challenges are to cancel echoes and eliminate overlap. *Echo cancellation* directly affects perceived quality; it becomes important when the round-trip delay exceeds 50 ms. Voice overlap becomes a concern when the one-way latency is more than 200 ms.

Because most of the time elapsed during a voice conversation is "dead time"—during which no speaker is talking—codecs take advantage of this silence by not transmitting any data during these intervals. Such "silence compression" techniques detect voice activity and stop transmitting data when there is no voice activity, instead generating "comfort" noise to ensure that the line does not appear dead when no one is talking.

In a standard PSTN telephone system, echoes that degrade perceived quality can happen for a variety of reasons. The two most common causes are impedance mismatches in the circuit-switched network ("line echo") and acoustic coupling between the microphone and speaker in a telephone ("acoustic echo"). Line echoes are common when there is a two-wire-to-four-wire conversion in the network (e.g., where analog signaling is converted into a T1 system).

Because VoIP systems can link to the PSTN, they must be able to deal with line echo, and IP phones can also fall victim to acoustic echo. Echo cancellers can be optimized to operate on line echo, acoustic echo, or both. The effectiveness of the cancellation depends directly on the quality of the algorithm used.

An important parameter for an echo canceller is the length of the packet on which it operates. Put simply, the echo canceller keeps a copy of the signal that was transmitted. For a given time after the signal is sent, it seeks to correlate and subtract the transmitted signal from the returning reflected signal—which is, of course, delayed and diminished in amplitude. To achieve effective cancellation, it usually suffices to use a standard correlation window size (e.g., 32 ms, 64 ms, or 128 ms), but larger sizes may be necessary.

Petitioners
Ex. 1062, p. 17

Emerging and Current VoIP-Based Applications

Because the high-speed network as a whole (rather than a dedicated channel) is used as the transport mechanism, a major advantage of VoIP systems is the lower cost per communication session. Moreover, VoIP calls allow network operators to avoid most interconnect charges associated with circuit-switched telephony networks; the additional infrastructure required to complete a VoIP phone call is minimal, because it uses the existing network already in place for the home or business personal computer (PC). Yet another reason for lower costs is that data-network operators often haven't used all the available bandwidth, so that the additional VoIP services currently incur an inconsequential additional cost-overhead burden.

VoIP users tend to think of their connection as being "free," since they can call anywhere in the world, as often as they want, for just pennies per minute. Although they are also paying a monthly fee to their Internet service provider, it can be amortized over both data and voice services.

Besides the low cost relative to the circuit-switched domain, many new features of IP services become available. For instance, incoming phone calls on the PSTN can be automatically rerouted to a user's VoIP phone, as long as it's connected to a network node. This arrangement has clear advantages over a global-enabled cellphone, since there are no roaming charges involved—from the VoIP standpoint, the end user's location is irrelevant; it is simply seen as just another network-connection point. This is especially useful where wireless local-area networks (LANs) are available; IEEE-Standard-802.11[2]-enabled VoIP handsets allow conversations at worldwide Wi-Fi[3] hotspots without the need to worry about mismatched communications infrastructure and transmission standards.

Everything discussed so far in relation to voice-over-IP extends to other forms of data-based communication as well. After all, once data is digitized and packetized, the nature of the content doesn't much matter, as long as it is appropriately encoded and decoded with adequate bandwidth. Because of this, the VoIP infrastructure facilitates an entirely new set of networked real-time applications, such as:

- Videoconferencing
- Remote video surveillance
- Analog telephone adapters
- Multicasting
- Instant messaging
- Gaming
- Electronic whiteboards

A CLOSER LOOK AT A VoIP SYSTEM

Figure 1 shows key components of a VoIP system: the signaling process, the encoder/decoder, the transport mechanism, and the switching gateway.

The *signaling* process involves creating, maintaining, and terminating connections between nodes.

In order to reduce network bandwidth requirements, audio and video are encoded before transmission and decoded during reception. This compression and conversion process is governed by various codec standards for both audio and video streams.

The compressed packets move through the network governed by one or more *transport* protocols. A *switching gateway* ensures that the packet set is interoperable at the destination with another IP-based system or a PSTN system. At its final destination, the packet set is decoded and converted back to an audio/video signal, at which point it is played through the receiver's speakers and/or display unit.

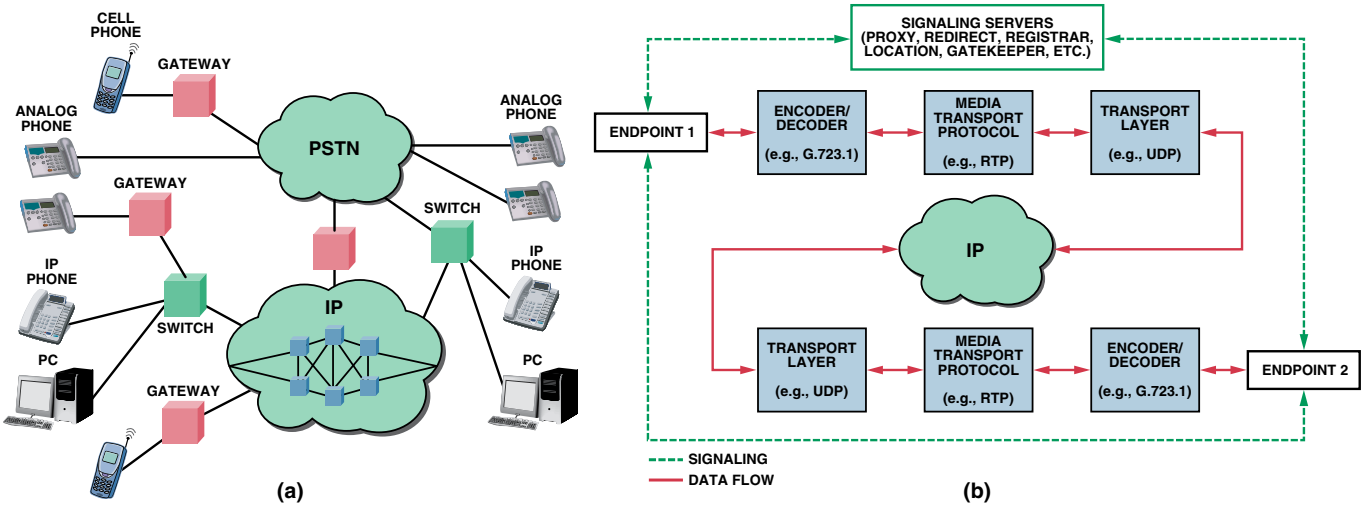

Figure 1. (a) Simplified representation of possible IP telephony network connections. (b) Signaling and transport flows between endpoints.

Petitioners
Ex. 1062, p. 18

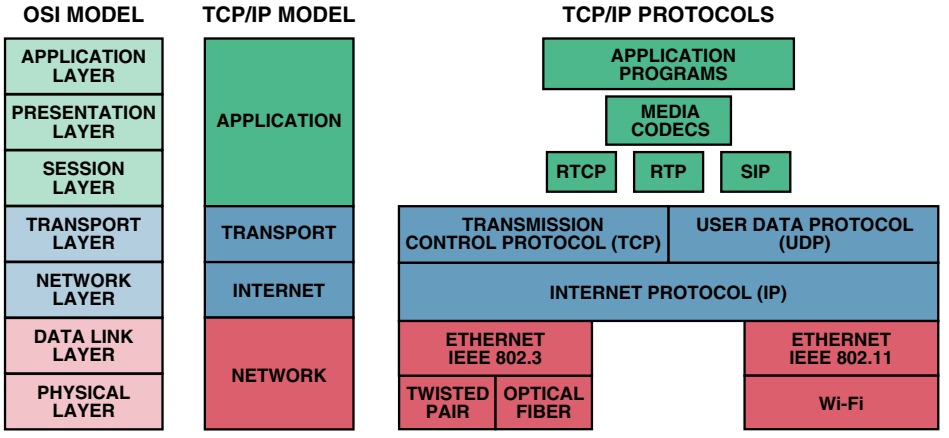

Figure 2. Open Systems Interconnection and TCP/IP models.

The OSI (Open Systems Interconnection) seven-layer model (Figure 2) specifies a framework for networking. If there are two parties to a communication session, data generated by each starts at the top, undergoing any required configuration and processing through the layers, and is finally delivered to the physical layer for transmission across the medium. At the destination, processing occurs in the reverse direction, until the packets are finally reassembled and the data is provided to the second user.

Session Control: H.323 vs. SIP

The first requirement in a VoIP system is a *session-control protocol* to establish presence and locate users, as well as to set up, modify, and terminate sessions. There are two protocols in wide use today. Historically, the first of these protocols was H.323*, but SIP (Session Initiation Protocol) is rapidly becoming the main standard. Let's take a look at the role played by each.

International Telecommunication Union (ITU) H.323

H.323[4] is an ITU standard originally developed for real-time multimedia (voice and video) conferencing and supplementary data transfer. It has rapidly evolved to meet the requirements of VoIP networks. It is technically a container for a number of required and optional network and media codec standards. The connection signaling part of H.323 is handled by the H.225 protocol, while feature negotiation is supported by H.245.

SIP (Session Initiation Protocol)

SIP[5] is defined by the IETF[6] (Internet Engineering Task Force) under RFC 3261. It was developed specifically for IP telephony and other Internet services—and though it overlaps H.323 in many ways, it is usually considered a more streamlined solution.

SIP is used with SDP[7] (*Session Description Protocol*) for user discovery; it provides feature negotiation and call management. SDP is essentially a format for describing initialization parameters for streaming media during session announcement and invitation. The SIP/SDP pair is somewhat analogous to the H.225/H.245 protocol set in the H.323 standard.

SIP can be used in a system with only two endpoints and no server infrastructure. However, in a public network, special proxy and registrar servers are utilized for establishing connections. In such a setup, each client registers itself with a server, in order to allow callers to find it from anywhere on the Internet.

*To be exact, the task of session control and initiation lies in the domain of H.225.0 and H.245, which are part of the H.323 umbrella protocol.

TRANSPORT LAYER PROTOCOLS

The signaling protocols above are responsible for configuring multimedia sessions across a network. Once the connection is set up, media flows between network nodes are established by utilizing one or more data-transport protocols, such as UDP or TCP.

UDP (User Datagram Protocol)

UDP[8] is a network protocol covering only packets that are broadcast out. There is no acknowledgement that a packet has been received at the other end. Since delivery is not guaranteed, voice transmission will not work very well with UDP alone when there are peak loads on a network. That is why a media transport protocol, like RTP,[9] usually runs on top of UDP.

TCP (Transmission Control Protocol)

TCP[10] uses a client/server communication model. The client requests (and is provided) a service by another computer (a server) in the network. Each client request is handled individually, unrelated to any previous one. This ensures that "free" network paths are available for other channels to use.

TCP creates smaller packets that can be transmitted over the Internet and received by a TCP layer at the other end of the call, such that the packets are "reassembled" back into the original message. The IP layer interprets the address field of each packet so that it arrives at the correct destination.

Unlike UDP, TCP does guarantee complete receipt of packets at the receiving end. However, it does this by allowing packet retransmission, which adds latencies that are not helpful for real-time data. For voice, a late packet due to retransmission is as bad as a lost packet. Because of this characteristic, TCP is usually not considered an appropriate transport for real-time streaming media transmission.

Figure 2 shows how the TCP/IP Internet model, and its associated protocols, compares with and utilizes various layers of the OSI model.

MEDIA TRANSPORT

As noted above, sending media data directly over a transport protocol is not very efficient for real-time communication. Because of this, a *media transport layer* is usually responsible for handling this data in an efficient manner.

Petitioners
Ex. 1062, p. 19

RTP (Real-Time Transport Protocol)

RTP provides delivery services for real-time packetized audio and video data. It is the standard way to transport real-time data over IP networks. The protocol resides on top of UDP to minimize packet header overhead—but at a cost; there is no guarantee of reliability or packet ordering. Compared to TCP, RTP is less reliable—but it has less latency in packet transmission, since its packet header overhead is much smaller than for TCP (Figure 3).

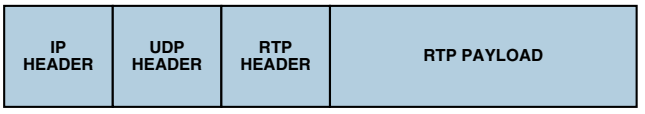

Figure 3. Header structure and payload of an RTP frame.

In order to maintain a given QoS level, RTP utilizes timestamps, sequence numbering, and delivery confirmation for each packet sent. It also supports a number of error-correction schemes for increased robustness, as well as some basic security options for encrypting packets.

Figure 4 compares performance and reliability of UDP, RTP, and TCP.

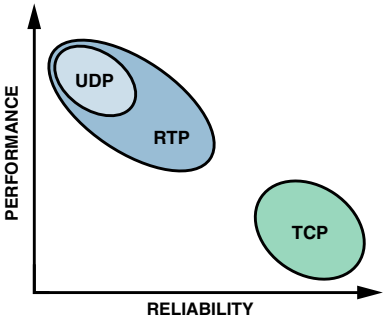

Figure 4. Performance vs. reliability.

RTCP (RTP Control Protocol)

RTCP[11] is a complementary protocol used to communicate control information, such as number of packets sent and lost, jitter, delay, and endpoint descriptions. It is most useful for managing session time bases and for analyzing QoS of an RTP stream. It also can provide a backchannel for limited retransmission of RTP packets.

MEDIA CODECS

At the top of the VoIP stack are protocols to handle the actual media being transported. There are potentially quite a few audio and video codecs that can feed into the media transport layer. A sampling of the most common ones can be found in the sidebar on the last page of this article.

A number of factors help determine how desirable a codec is—including how efficiently it makes use of available system bandwidth, how it handles packet loss, and what costs are associated with it, including intellectual-property royalties.

BLACKFIN VoIP COLLATERAL

Unlike traditional VoIP embedded solutions that utilize two processor cores to provide VoIP functionality, Blackfin processors provide a convergent solution in a unified core architecture that allows voice and video signal processing concurrent with RISC MCU processing to handle network- and user-interface demands. This unique ability to offer full VoIP functionality on a single convergent processor provides for a unified software development environment, faster system debugging and deployment, and lower overall system cost.

As an example, the ADSP-BF537[12] Blackfin processor family provides the necessary degree of integration and performance, with low power consumption, for VoIP deployment. It features multiple *integrated serial ports* (for glueless connection to audio analog-to-digital (A/D) and digital-to-analog (D/A) converters), *an external memory controller*, *a parallel peripheral interface* (PPI) for LCD or video encoder/decoder connectivity, and a *10/100BaseT Ethernet MAC*. If necessary, a second Ethernet MAC can be accommodated via the external memory interface.

A complete communication channel—including voice and networking stack—uses less than 75 MIPS of the processing bandwidth. With ADSP-BF537 performance at up to 600 MHz, there is plenty of available processor "horsepower" to spread across a VoIP product portfolio, as features such as multimedia compression or decompression become necessary. In contrast, competing dedicated VoIP choices are typically performance-limited and offer little or no ability to add features or differentiation.

For VoIP applications, Blackfin-based designs target high-quality, low-channel-count VoIP solutions—with processing headroom to accommodate added features such as music, video, and image transport, as well as overall system control. Here is a sampling of available VoIP offerings, ranging from open-source solutions to high-volume OEM reference designs:

Blackfin/Linphone

A Blackfin VoIP system can be designed using open-source software[13] based on μClinux, the embedded version of the popular GNU/Linux OS. One such *General Public License* (GPL-licensed) IP-phone package, called Linphone[14]—based on the *SIP* suite—has been ported to μClinux for Blackfin processors, allowing the Blackfin reference design to communicate with any SIP-compatible endpoint. In a public network with the proper SIP servers and gateway infrastructure, this system can even be used to connect to phones on a PSTN node. For voice encoding and decoding, the current Blackfin implementation of Linphone supports: G.711 (A-law and μ-law), *GSM* (Global System for Mobile Communications), and the *Speex* audio compression format.

The main components used in the Blackfin Linphone reference design are:

Linux TCP/IP networking stack: includes necessary transport and control protocols, such as TCP and UDP.

Linphone: the main VoIP application, which includes Blackfin-based G.711 and GSM codec implementations. It comprises both a *graphical user interface* (GUI) for desktop PCs and a simple command-line application for nongraphical embedded systems.

oRTP: an implementation of an RTP stack developed for Linphone and released under the LGPL license.

oSIP: a thread-safe implementation of the SIP protocol released under the LGPL license.

Speex: the open-source reference implementation of the Speex codec. Blackfin-specific optimizations to the fixed-point Speex implementation have been contributed back to the mainline code branch.

Unicoi Systems Blackfin-Based Fusion Voice Gateway

The Fusion Voice Gateway (Figure 5) is a complete *Voice Gateway/Terminal Adapter Reference Design*[15] from Unicoi Systems.[16] With router functionality and full-featured SIP telephony running on a single-core Blackfin Processor, the Fusion Voice Gateway allows for quick time-to-market for terminal adapters.

Petitioners
Ex. 1062, p. 20

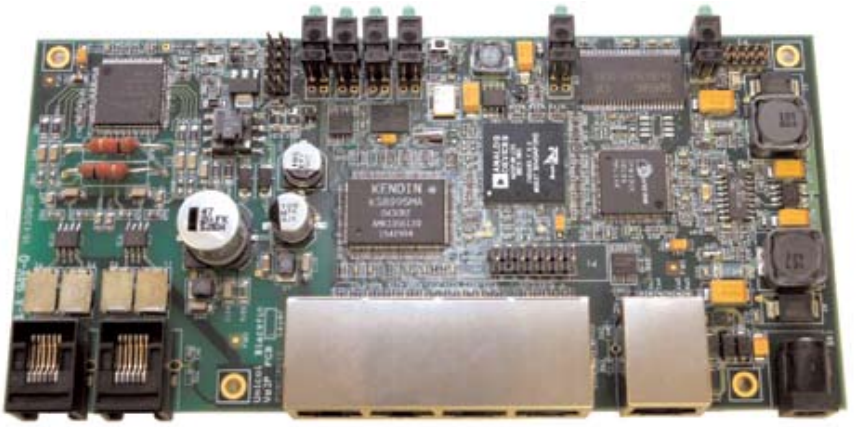

Figure 5. Blackfin-based Fusion Voice Gateway from Unicoi Systems.

The Fusion Voice Gateway features robust functionality, including G.168 echo cancellation and multiple G.7xx voice codecs. The Fusion reference design also includes full-featured telephony and router functionality by combining an Internet router, a 4-port Ethernet switch and VoIP gateway functionality.

Unicoi Systems Blackfin-Based Fusion IP Phone

The Fusion IP Phone from Unicoi Systems is a complete software/silicon solution that offers a full-featured platform supporting current and emerging IP phone standards, and has expansion capabilities for product differentiation.

The Fusion IP Phone Reference Design reduces BOM cost as well as the time and complexity often associated with developing an IP phone. Designed around the ADSP-BF536, the reference design software delivers the critical processing (e.g., real-time operating system, call manager, voice algorithms, acoustic echo cancellation for full-duplex speakerphone), communication protocols (TCP/IPv4/v6, SIP, RTP, etc.), and peripheral functions (LCD and keypad controllers, etc.) required to build a basic or advanced IP phone.

Blackfin BRAVO VoIP Reference Designs

The Analog Devices Blackfin BRAVO™ VoIP[17] and Videophone reference designs are complete system solutions for OEMs building feature-rich, high-performance, low-cost VoIP desktop phones, videophones, and telephone adapters. The designs include the complete suite of software for VoIP applications, all controlled by a comprehensive set of *application program interfaces* (APIs) for customization and control of core system functions.

For audio, the designs support multiple G.7xx audio codecs, G.168-compliant network echo cancellation, and acoustic echo cancellation for enhanced audio clarity. Optionally, RF transceivers can be included in the design to provide wireless audio capability. The designs support both H.323- and SIP-compliant software stacks.

On the video front, the BRAVO Broadband Audio/Video Communications[18] reference design (Figure 6) provides up to 30 frames per second of *common intermediate format* (CIF) color video, including support for ITU-standard H.263 and H.264 video codecs, picture-in-picture, high-resolution graphics with overlay, alpha and chroma keying, and antiflicker filtering.

CONCLUSION

Clearly, VoIP technology has the potential to revolutionize the way people communicate—whether they're at home or at work, plugged-in or untethered, video-enabled or just plain audio-minded. The power and versatility of Blackfin processors, working with a wide variety of standards, will make VoIP increasingly pervasive in embedded environments, creating value-added features in many markets that are not yet experiencing the benefits of this exciting technology.

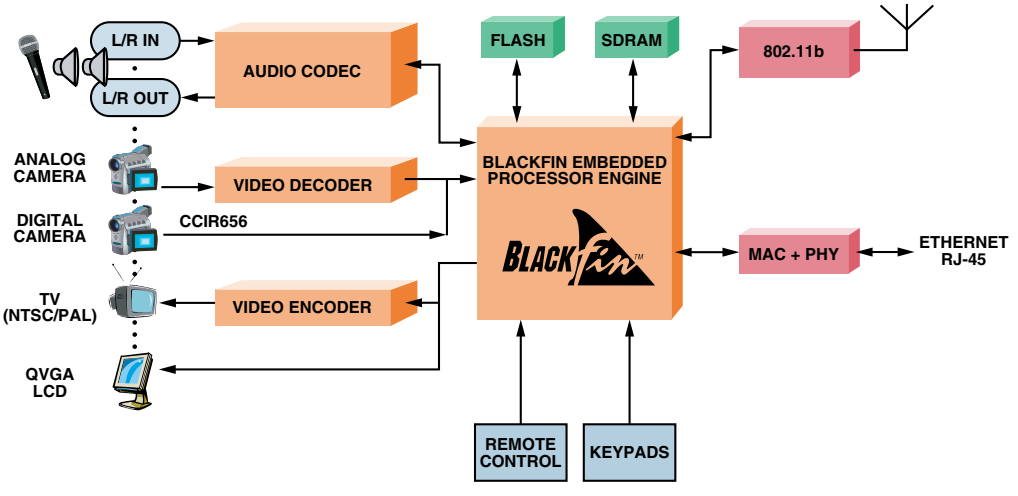

Figure 6. Blackfin BRAVO Broadband Audio/Video Communications reference design, functional diagram.

Petitioners
Ex. 1062, p. 21

REFERENCES—VALID AS OF JULY 2006

[1] http://www.analog.com/processors/processors/blackfin
[2] http://www.ieee802.org/11
[3] http://www.wi-fi.org
[4] http://www.itu.int/rec/T-REC-H.323/en
[5] http://www.ietf.org/rfc/rfc3261.txt
[6] http://www.ietf.org
[7] http://www.ietf.org/rfc/rfc2327.txt
[8] http://www.ietf.org/rfc/rfc768.txt
[9] http://www.ietf.org/rfc/rfc3550.txt
[10] http://www.ietf.org/rfc/rfc0793.txt
[11] http://www.ietf.org/rfc/rfc3550.txt (defined in the same RFC document as RTP)
[12] ADI website: www.analog.com (Search) ADSP-BF537 (GO)
[13] http://blackfin.uclinux.org
[14] http://docs.blackfin.uclinux.org/doku.php?id=linphone
[15] http://www.unicoi.com/Documentation/ProductBriefs/voicegateway_unicoi_blackfin.pdf
[16] http://www.unicoi.com
[17] http://www.analog.com/processors/platforms/blackfin_BRAVO/Blackfin_BRAVO_VoIP_final.pdf
[18] http://www.analog.com/processors/platforms/blackfin_BRAVO/Blackfin_BRAVO_Videophone(D)_final.pdf

MEDIA CODEC STANDARDS

Audio Codecs

G.711

Introduced in 1988, G.711—the international standard for encoding telephone audio on a 64-kbps channel—is the simplest standard among the options presented here. The only compression used in G.711 is *companding* (using either the μ-law or A-law standards), which compresses each data sample to an 8-bit word, yielding an output bit rate of 64 kbps. The H.323 standard specifies that G.711 must be present as a baseline for voice communication.

G.723.1

G.723.1 is an *algebraic code-excited linear-prediction* (ACELP)-based dual-bit-rate codec, released in 1996 to target VoIP applications. The encoding time frame for G.723.1 is 30 ms. Each frame can be encoded in 20 bytes or 24 bytes, thus translating to 5.3-kbps or 6.3-kbps streams, respectively. The bit rates can be effectively reduced through voice-activity detection and comfort-noise generation. The codec offers good immunity against network imperfections—like lost frames and bit errors. G.723.1 is suitable for video-conferencing applications, as described by the H.324 family of international standards for multimedia communication.

G.729

Another speech codec, released in 1996, is the low-latency G.729 audio data-compression algorithm, which partitions speech into 10-ms frames. It uses an algorithm called *conjugate-structure ACELP* (CS-ACELP). G.729 compresses 16-bit signals sampled at 8 kHz via 10-ms frames into a standard bit rate of 8 kbps, but it also supports 6.4-kbps and 11.8-kbps rates. In addition, it supports voice-activity detection and comfort-noise generation.

GSM

The GSM speech codecs find use in cell phone systems around the world. The governing body for these standards is the *European Telecommunications Standards Institute* (ETSI). Standards in this domain have evolved since the first one, GSM Full Rate (GSM-FR). This standard uses a CELP variant called *regular pulse-excited linear predictive coder* (RPELPC). The input speech signal is divided into 20-ms frames. Each frame is encoded as 260 bits, thereby producing a total bit rate of 13 kbps. Free GSM-FR implementations are available for use under certain restrictions.

Speex

Speex, an Open Source/Free Software audio compression format designed for speech codecs, was released by Xiph.org, with the goal of being a totally patent-free speech solution. Like many other speech codecs, Speex is based on CELP with residue coding. It can code 8-kHz, 16-kHz, and 32-kHz linear PCM signals into bit rates ranging from 2 kbps to 44 kbps. Speex is resilient to network errors, and it supports voice-activity detection. Besides allowing variable bit rates, Speex also has the unique feature of *stereo encoding*. Source code is available from Speex.org that includes assembly-level optimization for narrow-band compression, as well as a Blackfin-based echo canceller.

Video Codecs

H.261

This standard, developed in 1990, was the first widely used video codec. It introduced the idea of segmenting a frame into 16 × 16 "macroblocks" that are tracked between frames to establish motion-compensation vectors. It is mainly targeted at videoconferencing applications over ISDN lines ($p \times 64$ kbps, where p ranges from 1 to 30). Input frames are typically CIF (352 × 288) at 30 frames-per-second (fps), and output compressed frames occupy 64 kbps to 128 kbps for 10-fps resolution. Although still used today, it has been largely superseded by H.263. Nevertheless, H.323 specifies that H.261 must be present as a baseline for video communication.

H.263

This codec is ubiquitous in videoconferencing, outperforming H.261 at all bit rates. Input sources are usually *quarter-common intermediate format* (QCIF) (176 × 144) or CIF at 30 fps, and output bit rates can be less than 28.8 kbps at 10 fps, for the same performance as H.261. So whereas H.261 needed an ISDN (*integrated-services-digital-network*) line, H.263 can use ordinary phone lines. H.263 finds use in end markets such as video telephony and networked surveillance, and it is popular in IP-based applications.

This article can be found at http://www.analog.com/library/analogdialogue/archives/40-04/blackfin_voip.html, with a link to a PDF.

Petitioners
Ex. 1062, p. 22

PRODUCT INTRODUCTIONS: VOLUME 40, NUMBER 2

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

April

Amplifier, Power, Class D Audio **AD1994**
Clock and Data Recovery, 622-Mbps, integrated limiting amplifier **ADN2804**
Clock and Data Recovery, 622-Mbps **ADN2806**
Controller, Synchronous Buck, 2-/3-/4-phase **ADP3190**
Driver/Comparator/Active Load, dual, 250-MHz/500-MHz **ADATE205/ADATE206**
Driver, MOSFET, dual, 12-V, high-side bootstrap, output disable **ADP3120A**
Inclinometer/Accelerometer, dual-axis, programmable **ADIS16201**
Modulator, Sigma-Delta, isolated **AD7401**

May

Accelerometer, dual-axis, SPI interface, ±5-*g* range **ADIS16006**
Accelerometer, three-axis, ±3-*g* range **ADXL330**
ADC, Successive-Approximation, 8-channel, 12-bit-plus-sign, 1-MSPS **AD7329**
ADC, Successive-Approximation, 6-channel, 16-bit, 25-kSPS **AD7656**
Amplifier, Operational, dual, low-noise, high-precision, CMOS **AD8662**
DAC, Bipolar Voltage Output, quad, 16-bit, serial-input **AD5764**
Filter, Video, ultralow-power **ADA4430-1**
Isolators, Digital, 3-channel, enhanced ESD protection **ADuM3300/ADuM3301**
Isolators, Digital, 4-channel, enhanced ESD protection **ADuM3400/ADuM3401/ADuM3402**
Receiver, Diversity, IF-to-baseband, narrow-band GSM/EDGE **AD6650**
Processor, Embedded, Blackfin **ADSP-BF534**
Processor, Embedded, SHARC® **ADSP-21261**
Processor, Signal, Video, CCD, 2-channel, 14-bit, *Precision Timing*™ generator **AD9972**
Switches, Crosspoint 16 × 5, buffered, 260-MHz, video **AD8106/AD8107**

June

ADC, Pipelined, quad, 12-bit, 40-MSPS/65-MSPS, LVDS outputs **AD9228**
ADC, Pipelined, quad 10-bit, 40-MSPS/65-MSPS, LVDS outputs **AD9219**
ADC, Pipelined, 14-bit, 105-MSPS/125-MSPS, 1.8-V supply ... **AD9246**
ADC, Pipelined, 13-bit, 130-MSPS, IF-sampling **AD9461**
ADC, Pipelined, 12-bit, 105-MSPS/125-MSPS, 1.8-V supply ... **AD9233**
ADC, Successive-Approximation, 18-bit, 1.25-MSPS **AD7643**
ADC, Successive-Approximation, 18-bit, 400-kSPS, 1.5-LSB INL **AD7690**
Amplifier, Instrumentation, JFET-input, rail-to-rail output **AD8220**
Amplifier, Operational, CMOS, 1-μA max supply current **AD8500**
Amplifier, Operational, low-cost, high-speed **ADA4860-1**
Amplifier, Variable-Gain, ultralow-noise, quad **AD8334**
Amplifier, Variable-Gain, 800-MHz, automatic gain control **AD8368**
Amplifiers, Power, Class D Audio **AD1990/AD1992**
Codec, SoundMAX®, HD Audio **AD1983**
Controller, DC-to-DC, dual, step-down **ADP1823**
DAC, Voltage Output, 16-bit, dual, low-power, rail-to-rail outputs **AD5663**
DACs, Voltage Output, 12-/14-/16-bit, quad, 5-ppm/°C reference **AD5624R/AD5644R/AD5664R**
Driver, ADC, differential, 18-bit, single-supply **ADA4941-1**
Generator, Waveform, programmable frequency scan **AD5932**
Monitor, Temperature, ±1°C accuracy, series resistance cancellation **ADT7461A**
Monitor, Current Shunt, high-voltage, bidirectional **AD8210**
Multiplexers, *i*CMOS, 4-channel/8-channel, low capacitance, ±15 V **ADG1208/ADG1209**
Processor, Signal, Video, CCD, 14-bit, dual, *Precision Timing* core **AD9973**
Processor, Signal, Video, CCD, 12-bit *Precision Timing* core **AD9971**
Processors, Embedded, Blackfin **ADSP-BF536/ADSP-BF537**
Synthesizer, Frequency, 8 GHz, integer-N **ADF4108**
Synthesizer, Frequency, phase detector, 400 MHz **ADF4002**

AUTHORS

Eric Gaalaas (page 6) received Bachelor's and Master's degrees in EE from Cornell University in 1993 and 1994, and has 12 years experience as an IC design engineer, mostly on audio products. He is an inventor, with several patents related to audio DACs, sample-rate conversion, and Class D amplifiers, and an author of related publications. From 1999 to 2006, he worked in the Digital Audio group of Analog Devices in Wilmington, MA, and recently moved to the *i*Coupler Components group.

Rick Gentile (page 17), who joined ADI in 2000 as a senior DSP applications engineer, currently leads the Blackfin Processor Applications group. Additionally, he is co-author of *Embedded Media Processing* (Newnes 2005). Before joining ADI, Rick was a member of the technical staff at MIT Lincoln Laboratory, where he designed several signal processors used in a wide range of radar sensors. He received a BS from the University of Massachusetts at Amherst and a Master of Science degree from Northeastern University—both in electrical and computer engineering.

David Katz (page 17) joined ADI in 2001. He is a senior applications engineer responsible for new product development on Blackfin processors. He has published internationally dozens of embedded processing articles, and he has presented several conference papers in the field. Additionally, he is co-author of *Embedded Media Processing* (Newnes 2005). Previously, he worked at Motorola, Inc. as a senior design engineer in cable modem and factory automation groups. David holds both a BS and an MEng in electrical engineering from Cornell University.

Tom Lukasiak (page 17) has been a DSP applications engineer at Analog Devices since 2000. He is currently focusing on Blackfin processor products. Tom earned his ScB (2000) and ScM (2002) degrees in electrical engineering from Brown University.

Wayne Meyer (page 17) is the product marketing manager for Blackfin processors. He has over 15 years of experience in the semiconductor industry working for both Analog Devices and Motorola (Freescale) Semiconductor. Previous positions include product engineering and product manager roles for 8-bit and 16-bit microcontrollers and business development manager for 24-bit DSP and 32-bit RISC microprocessors. Wayne holds a BSEE degree from Southern Methodist University in Dallas, TX, and an MBA from St. Edward's University in Austin, TX.

Alan Moloney (page 13) is an applications engineer in the Power Systems Management group based in Limerick, Ireland. The PSM group provides voltage supervision, power-supply sequencing, and hot-swap controller solutions. Alan graduated from the National University of Ireland Galway in 2001 with a BE in electronic engineering.

Petitioners
Ex. 1062, p. 23

Analog Dialogue

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

Purchase of licensed I²C components of Analog Devices or one of its sublicensed Associated Companies conveys a license for the purchaser under the Philips I²C Patent Rights to use these components in an I²C system, provided that the system conforms to the I²C Standard Specification as defined by Philips.

©2006 Analog Devices, Inc. All rights reserved. Trademarks and registered trademarks are the property of their respective owners.
Printed in the U.S.A. M02000402-54-8/06

www.analog.com/analogdialogue

ANALOG DEVICES