Analog Dialogue

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

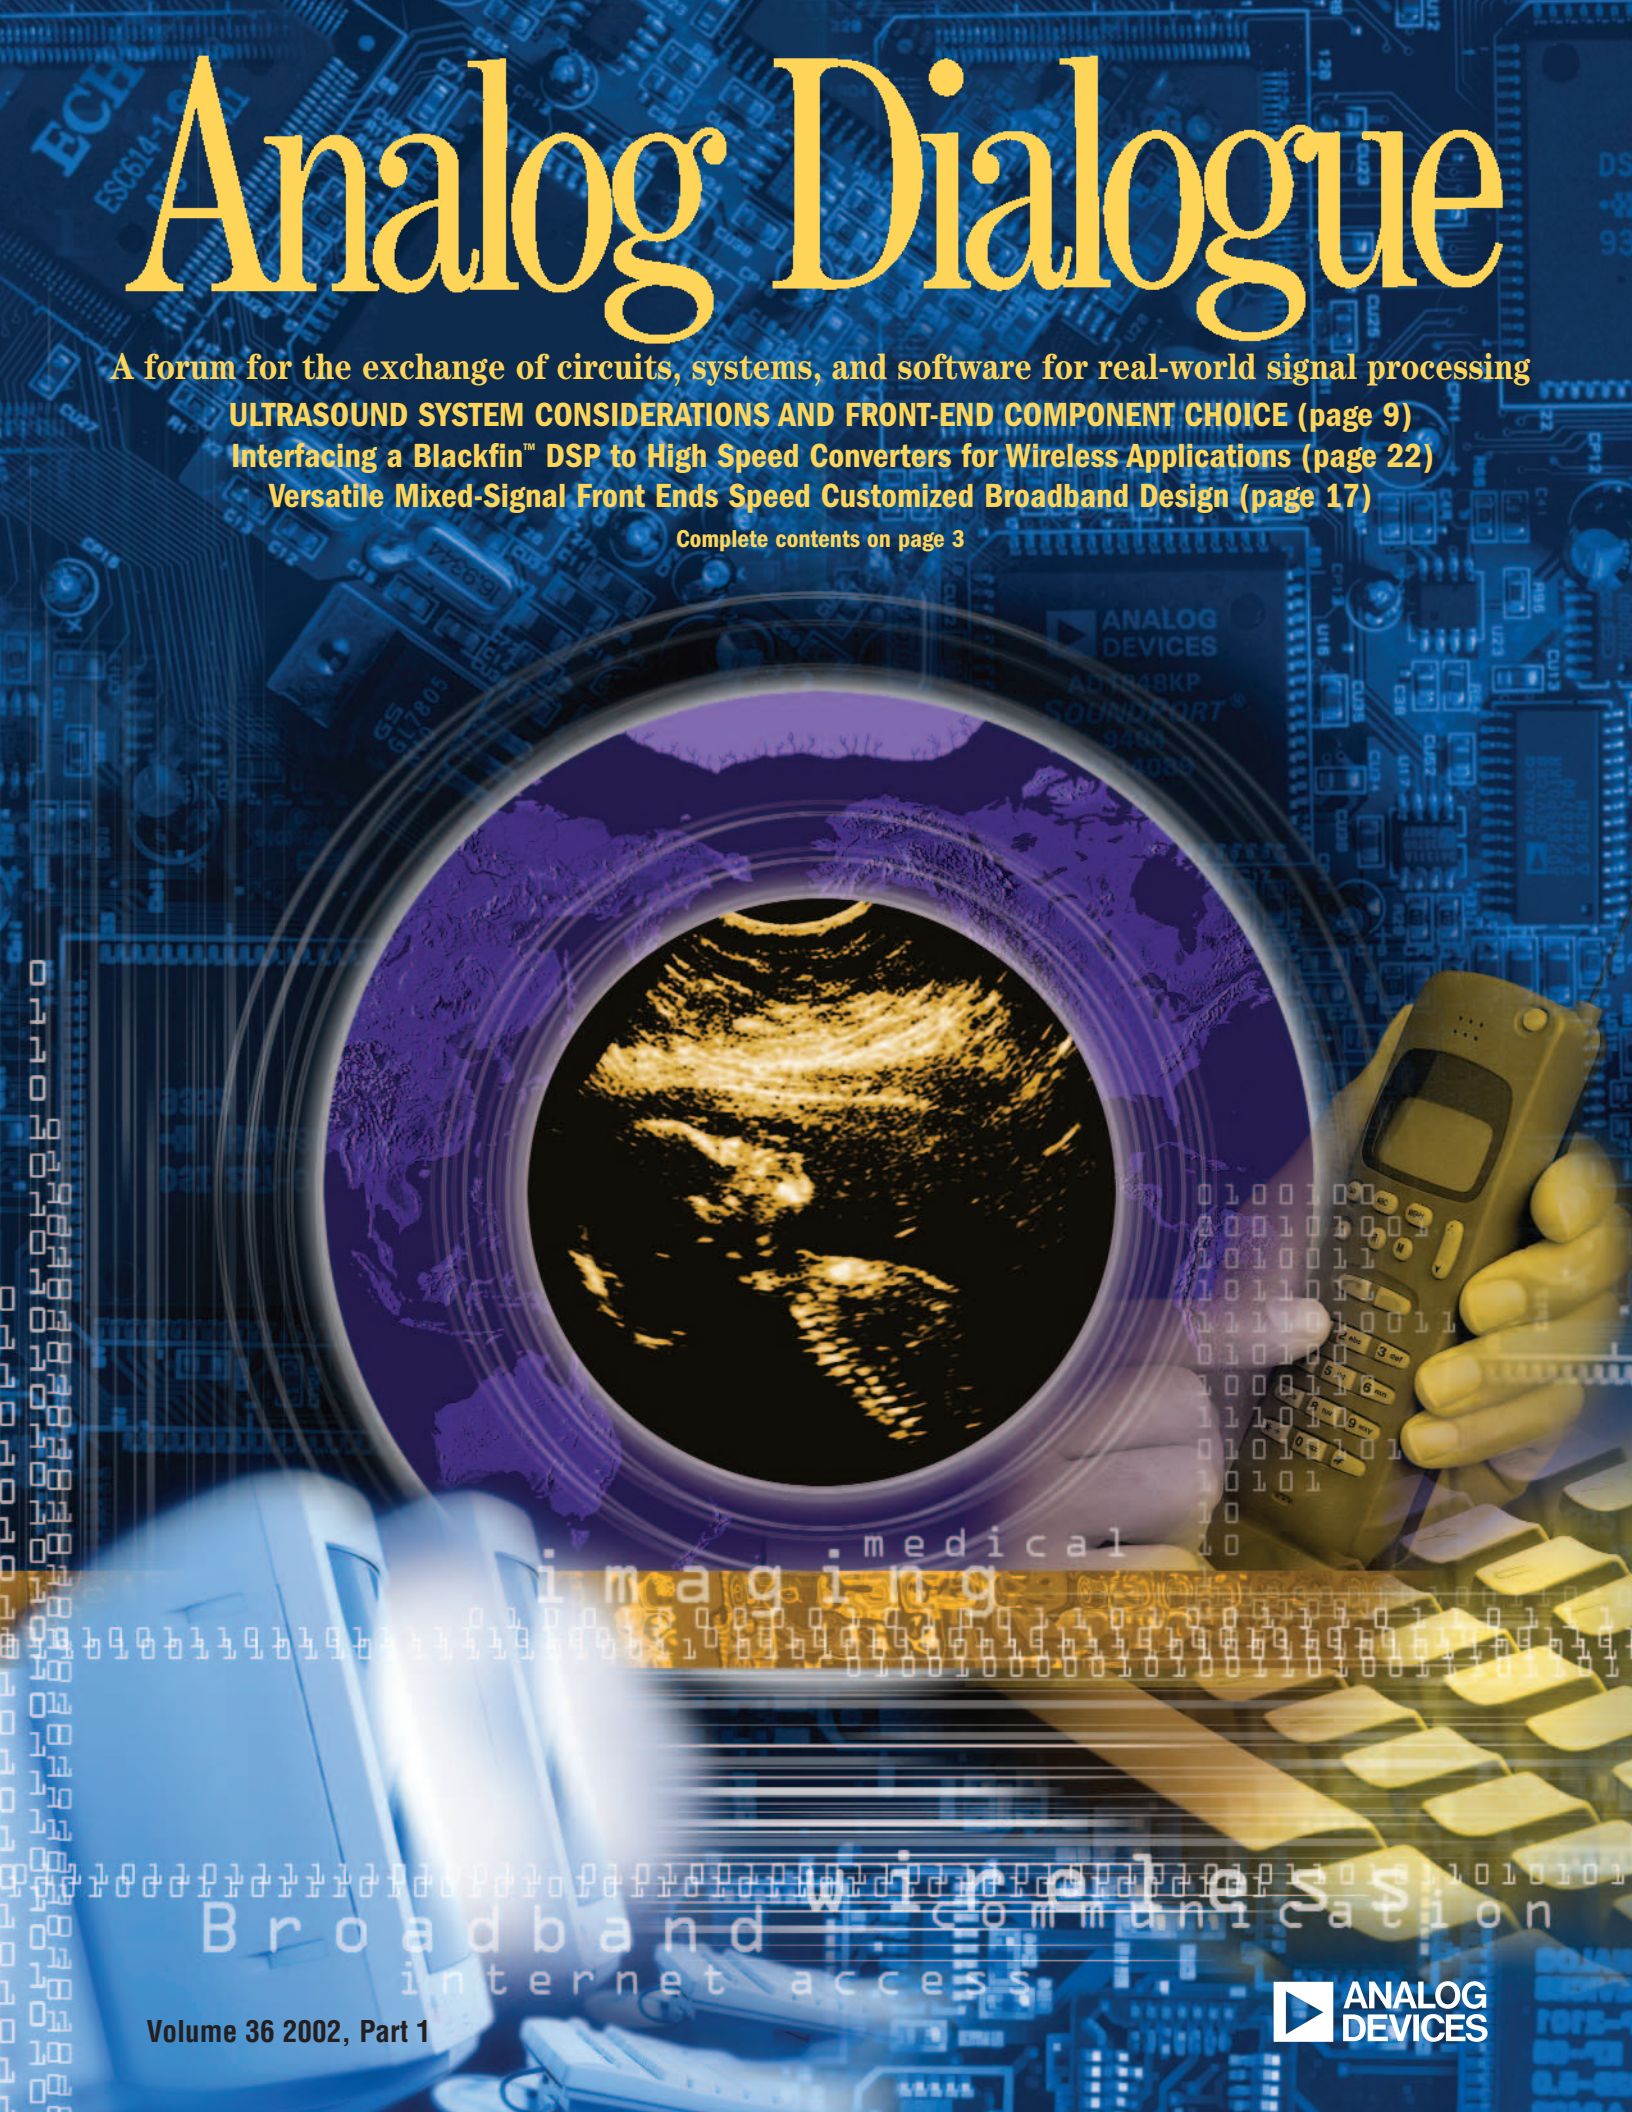

Volume 36 2002, Part 1

ANALOG DEVICES

Editor's Notes

RESTORING THE BALANCE

In the Fall of 1998, we decided to make *Analog Dialogue* more useful to our customers and our industry in the burgeoning on-line era. The Internet offered an opportunity to add timely new-product information to the core of timeless in-depth articles about Analog Devices applications, products, and technologies. We could archive older issues and have them readily accessible. And, most compellingly, we could link our writings with the many sources of information on the Analog Devices website: data sheets, application notes, press releases, white papers, contact information, etc.—and our colleagues could link their product sites to the cogent background stories in *Analog Dialogue*.

But there were questions: How do we plunge into a medium whose potential spoke loudly but that we understood only dimly? How should we present ourselves? What kind of schedule made sense? Could we still serve faithful readers, a few of whom boasted complete files of our publication, going back 36 years? Or should we pocket the Web's cost savings and abandon print?

We chose monthly publication, and—recognizing that we had an international audience, many of whose members had limited access-speed and memory—we decided on a simple format that permitted downloading with a minimum of time- and memory-consuming "bells and whistles." During the past three years, the on-line issues have been well accepted—and have garnered many, many thousands of eyeballs and mouse clicks.

Knowing that there was also an audience that could not—or would not—use the Internet, we resolved to publish large annual print editions containing all the articles that had appeared during each year. Unexpectedly, though, annual was not often enough. "Out of sight, out of mind!" A surprising number of inveterate print readers registered their dismay at the apparent *disappearance* of *Analog Dialogue*. Had it gone the way of the *Hewlett Packard Journal*? Our answer: "Let's publish more often!" So, this issue, Volume 36, Part 1, is the first of a series of two to four issues per year—returning to the goal of regular quarterly publication—with continuing strong ties to the more frequent Internet edition.

Dan Sheingold,

Editor

Analog Dialogue

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

Analog Dialogue is the free technical magazine of Analog Devices, Inc., published continuously for 36 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—*on-line*, monthly or bimonthly at the above URL, and—less frequently—*in print*, as periodic retrospective collections of articles that have appeared on-line. In addition to technical articles, the on-line 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 29, Number 2 (1995), plus two 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** and click on **<Contact ADI>** or type in this URL: **www.analog.com/request.html**. 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]**.

AUTHORS

Eberhard Brunner (page 9) has been with ADI since 1991, starting as a technician for Barrie Gilbert. He progressed to design and worked mostly on VGAs for ultrasound and on the first generation of GSM chipsets. In 1997 he moved back to his native Germany to provide field support for customers in the RF and medical ultrasound fields. He is now returning to NW Labs, where he enjoys the wonderfully rainy winters in the Northwest, as well as windsurfing and skiing.

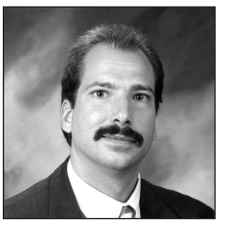

Joe DiPilato (page 17) is the product-line manager for communication-specific integrated products within ADI's High-Speed Converter (HSC) group. Since joining ADI in 1987, he has been involved in the start-up, early development, and marketing of industrial, instrumentation, and communication-focused product lines—and for the last eight years HSC Wilmington's standard- and communication-specific products, including TxDAC® and MxFE™ brands. He has authored numerous articles, many resulting in cover stories. Joe has a BSEE from Worcester Polytechnic Institute (1982) and an MBA from Anna Maria College (1984). He is active in his church, enjoys swimming and quality beach time with his wife Lisa and son Nicholas.

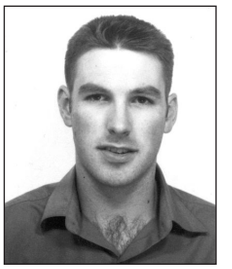

Adrian Fox (page 13), a graduate of the University of Limerick, with a BEng in Electronic Engineering, works in the RF Applications group in Limerick, Ireland, supporting PLL synthesizer products and the new ISM-band transmitter products. Hobbies include soccer, surfing (on the west coast of Ireland!), and playing the accordion.

Rick Gentile (page 22) leads the Blackfin DSP Applications Group. He joined ADI in 2000 as a Senior DSP Applications Engineer after working as a member of the Technical Staff at MIT Lincoln Laboratory, where he designed signal processors used in a wide range of radar sensors. He received a BS in 1987 from the University of Massachusetts at Amherst and an MS in 1994 from Northeastern University, both in Electrical and Computer Engineering.

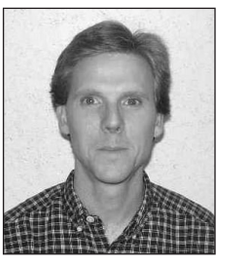

Brian Harrington (page 17) has recently become a Field Applications Engineer—supporting ADI's analog products in northeastern U.S. He had spent the prior six years developing reference designs for powerline, phoneline, ADSL, and other broadband modems. When the chores are finished at home, Brian enjoys playing with his three children, plus running, eating, and napping.

Jeritt Kent (page 22) is the senior Communications Field Applications Engineer for western North America, based in Bellevue, Washington. Jeritt received both a BSEE and an MEng from the University of Idaho. Earlier, he worked at Allegro Microsystems, and then at American Microsystems, where he spent six years as a mixed-signal Design Engineer.

X-Amp™, A New 45-dB, 500-MHz Variable-Gain Amplifier (VGA) Simplifies Adaptive Receiver Designs

by Eric J. Newman (eric.newman@analog.com)

INTRODUCTION

Wireless communications equipment design usually starts with strategic *signal-chain* definition and analysis. *Noise Figure* (NF), *linearity*, *distortion*, and *dynamic range* all need to be considered at an early stage in the product development cycle to properly identify component specifications for each element in the signal path. Signal-chain budget analysis allows designers to quickly select components, analyze, and compare the performance of design architectures being considered. The challenge is greater in mobile communications systems, where special attention needs to be focused on the *spectral selectivity*, *linearity*, and *noise mechanisms* associated with RF and IF signal blocks.

Receivers can be designed to provide adaptive sensitivity to incoming signal strength by employing variable gain at the lower IF frequencies, where it is easier to manipulate the signal of interest. Most spectral grooming (frequency shaping and filtering) tends to be implemented at the lower IF frequencies where very narrow-band pass filters can be easily realized through the use of SAW devices, crystals, and passive lumped-element RLC filter networks. After precise channel selection, automatic gain-control (AGC) circuitry can be employed to scale the received signal to a desired level. The use of AGC yields a receiver design whose sensitivity varies, based on received signal strength. Adaptive sensitivity reduces the effects of distance inherent in fading-channel mobile environments. High-performance variable-gain amplifiers are often necessary to provide the needed dynamic range and noise performance.

Background

Variable gain amplifiers (VGAs) have been used in a variety of remote sensing and communications equipment for more than a half century. Applications ranging from ultrasound, radar and lidar, to wireless communications—and even speech analysis—have utilized variable gain in an attempt to enhance dynamic performance. Early designs achieved gain selection by switching in fixed-gain amplifier stages to adjust receiver sensitivity in a binary fashion. Later implementations used step attenuators followed by fixed-gain amplifiers to achieve a wider range of discrete gain control. Modern designs achieve continuous voltage controlled gain, using analog techniques, by such means as voltage-variable attenuators (VVAs), analog multipliers, and gain interpolators.

A variety of architectures are commonly used to provide both continuous and discrete variable-gain control. Applications such as automatic gain control often require *continuous* analog gain control. The most straightforward designs utilize analog multipliers followed by fixed-gain buffer amplifiers. Such designs often involve a

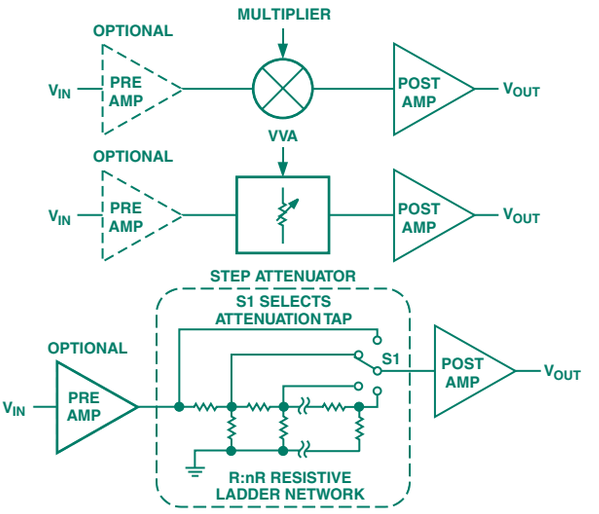

Figure 1. Typical variable-gain architectures.

nonlinear gain-control function that requires calibration. Additionally, the multiplier cores suffer from temperature and supply voltage dependencies that can result in poor gain-law accuracy and stability, as well as unacceptable high-frequency gain variation. Designs that use preamplifier/attenuator/post-amplifier architectures can provide low-noise operation and good bandwidth, but tend to have quite low input third-order intercepts (IIP3), limiting their ability to perform in high-dynamic-range receivers.

Another class of solutions utilizes voltage-variable attenuators, followed by fixed-gain post-amplification. VVAs can provide an accurate attenuation transfer function that is linear in dB, but it is often necessary to cascade multiple VVAs in order to provide adequate attenuation range. The cascading results in an increased sensitivity to variations of the attenuation transfer function. It is sometimes necessary to preamplify the signal to buffer the signal source from the loading effects of the VVA, as well as to decrease the attenuator's influence on noise figure. The high gain required to yield a low-noise figure results in a decreased input third-order intercept.

The AD8367 X-AMP VGA with AGC

The X-AMP architecture, originating ten years ago with the Analog Devices AD600 and AD602, (*Analog Dialogue* 26-2, 1992), permits a linear-in-dB gain-control function that is essentially independent of temperature. It comprises a resistive ladder network, along with a highly linear amplifier and interpolator stage, to provide a continuous linear-in-dB gain-control function. The AD8367

IN THIS ISSUE

Volume 36, Part 1, November 2002

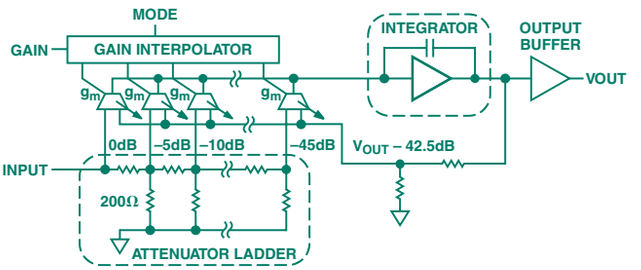

Figure 2. Architecture of the AD8367 X-Amp VGA.

(Figure 2) is the latest generation of X-AMP VGAs. Its design is implemented on a new eXtra Fast Complementary Bipolar process (XFCB2.0) that provides moderate gain out to hundreds of MHz and improved linearity at higher frequencies than heretofore available with conventional semiconductor processing.

As Figure 2 shows, the input signal is applied to a ground-referenced 9-stage *R-nR* resistive ladder network, designed to produce 5-dB steps of attenuation between tap points. Smooth gain control is achieved by sensing the tap points with variable-transconductance (g_m) stages. Depending on the gain-control voltage, an interpolator selects which stages are active. For example, if the first stage is active, the 0-dB tap point is sensed; if the last stage is active, the 45-dB point is sensed. Attenuation levels that fall between tap points are achieved by having neighboring g_m stages active simultaneously, creating a weighted average of the discrete tap-point attenuations. In this manner, a smooth, monotonic, linear-in-dB attenuation function with very precise scaling is synthesized. The ideal linear-in-dB transfer function can be expressed as:

$$Gain\,(dB) = M_Y \times V_{GAIN} + B_Z \qquad (1)$$

where

M_Y is the gain scale (slope) usually expressed in dB/V, typically 50 dB/V (or 20 mV/dB)

B_Z is the gain intercept in dB, typically –5 dB, the extrapolated gain for V_{GAIN} = 0 V

V_{GAIN} is the gain-control voltage

The AD8367's basic connection outline, gain transfer function, and typical gain-error pattern are illustrated in Figure 3, showing the gain transfer function's slope of 50 dB/V and –5-dB intercept over a gain-control voltage range of 50 mV ≤V_{GAIN} ≤950 mV. The device allows the gain slope to be reversed by a simple pin-strap of the MODE pin. The inverse gain mode is convenient in automatic gain control (AGC) applications, where the gain-control function is derived from an error integrator, which compares the detected output power to a predetermined set-point level. A square-law detector and the error integrator, integrated on-chip, allow the device to be used as a self-contained AGC subsystem.

A typical stand-alone AGC circuit is shown in Figure 4, along with its time-domain response to a 10-dB input voltage step. In this example the signal input is a 70-MHz sinusoid, and its input is step modulated from –17 dBm to –7 dBm (referred to 200 Ω). The output signal power is measured as a voltage by the internal square-law detector and compared to an internal 354-mV rms reference. The output of the detector is a current, which is integrated using an external capacitor, C_{AGC}. The voltage that is developed across the C_{AGC} capacitor drives the GAIN pin to reduce or increase the gain. The loop is stabilized when the output signal level's rms value becomes equal to the internal 354-mV reference. When the input signal is less than 354 mV rms, the DETO pin sinks current which reduces the voltage at the GAIN pin. As the input signal increases above 354 mV rms, the DETO pin sources current causing the voltage at the GAIN pin to increase. The inverse gain mode is required in this application to ensure that the gain decreases when the input signal's rms value exceeds the internal reference. The resulting voltage applied to the GAIN pin, V_{AGC}, can be used as a received-signal-strength indication (RSSI), representing the input signal strength as compared to a 354-mV rms reference. For a sinusoidal waveform this results in a 1-V p-p output signal for a 200-Ω load.

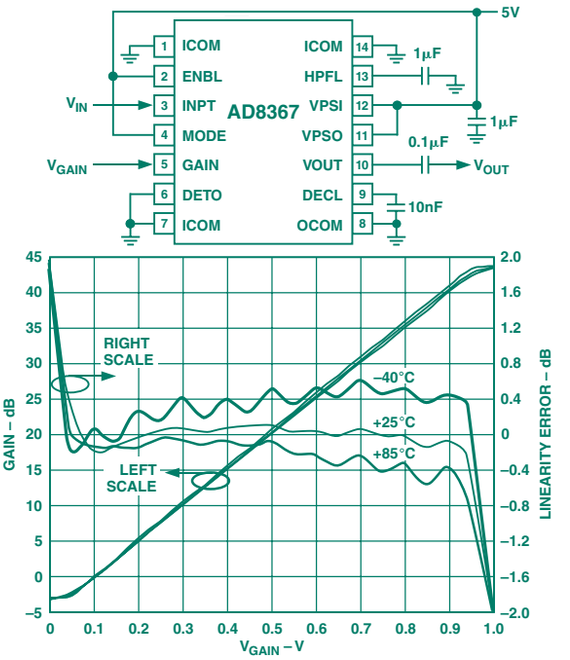

Figure 3. Basic AD8367 VGA application circuit and gain-control transfer function, showing typical errors at various temperatures.

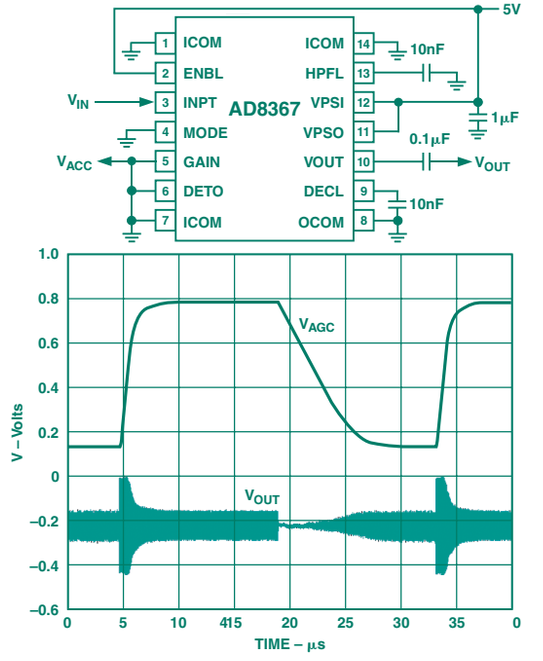

Figure 4. Basic AD8367 AGC application circuit and time-domain response at 70 MHz.

Signal Chain Analysis

A modern superheterodyne architecture is depicted in Figure 5. The AD8367 is used in the *receive* (Rx) path to adaptively adjust overall receiver gain as the RF signal level changes. In the *transmit* (Tx) path, the AD8367 is used in conjunction with an RF power detector to maintain a desired output power level.

Considering the receive path, the overall sensitivity and dynamic range can be assessed using signal-path budget analysis. For this example a PCS-CDMA signal was selected, using a 1-MHz noise bandwidth. Working backwards from the output of the AD8367 IF VGA, the input sensitivity and dynamic range can be analyzed. Figure 6 represents a detailed budget analysis from the receiver input to the output of the IF VGA.

In the example above, the AD8367 controls received signal levels prior to the I & Q demodulator. The AD8367 is an example of a VGA that uses variable attenuation followed by a post-gain amplifier. This style of VGA will exhibit essentially a constant OIP3 and a noise figure that varies with gain setting. The AD8367 provides minimum noise figure at maximum gain and maximum input third-order intercept at minimum gain. This unique combination allows for dynamic control of a receiver's sensitivity and input linearity, based upon received signal strength.

AD8367 is characterized over temperature from –40°C to +85°C and is packaged in a 14-lead thin-shrunk small-outline package (TSSOP). It operates on a single 3-V to 5-V supply. The device has a –3-dB operating bandwidth of 500 MHz; and its data sheet provides detailed specifications at common IF frequencies—such as 70 MHz, 140 MHz, 190 MHz, and 240 MHz. If you are reading the PDF or printed version of this article, please visit www.analog.com to download the data sheet or to request samples. The AD8367 is normally available from stock, and an evaluation board is also available.

Acknowledgements

The innovative AD8367 was designed by Barrie Gilbert and John Cowles. The author would like to recognize Leon Small, Dana Whitlow, and Pete Kearney for their characterization efforts.

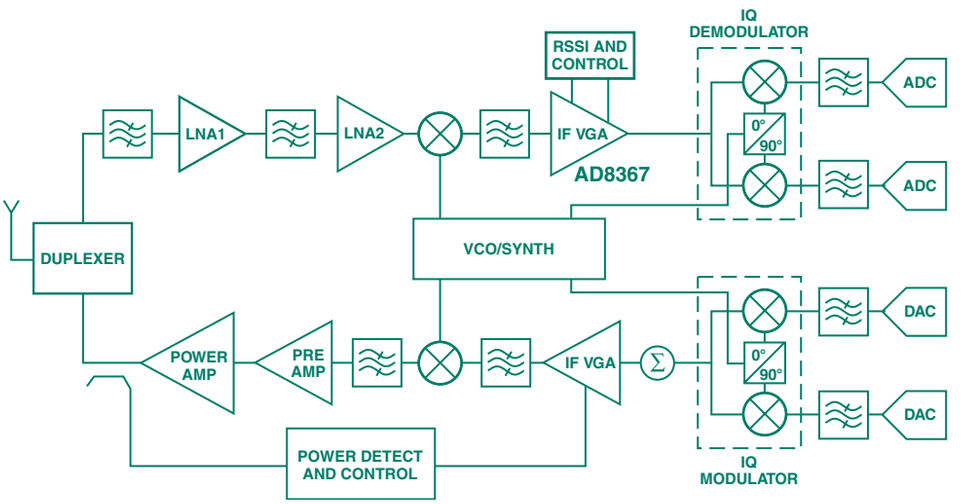

Figure 5. Superheterodyne architecture using VGAs for IF level control. VGAs are used in the intermediate frequency stages to adjust overall receiver sensitivity adaptively and to control transmitted power levels.

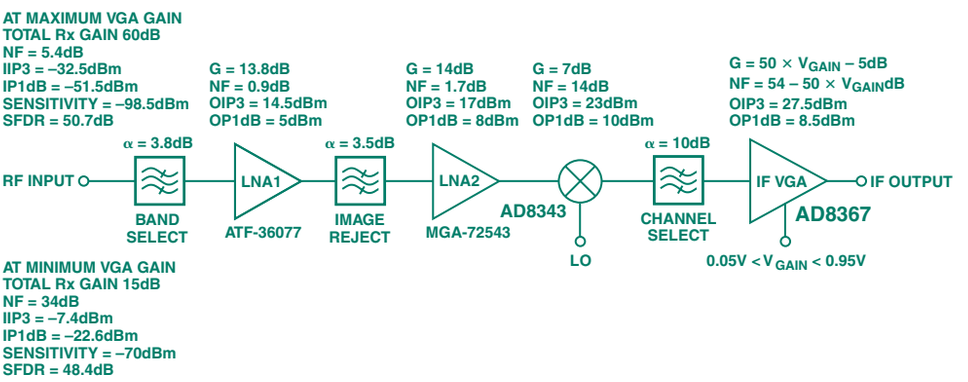

Figure 6. Rx Path Budget Analysis for 1900-MHz CDMA with a 70-MHz IF.

Simplify Audio Setups with a SigmaDSP™ Pre-Programmed, Fully Configurable digital Audio Processor

by Thomas Irrgang (thomas.irrgang@analog.com)

INTRODUCTION

While audio *sources* at present largely consist of digital media (CDs, DVDs, Internet), most of the audio processing has stayed in the analog domain. It is high time to migrate the audio *processing* to digital as well.

In the last two decades, audio technology has experienced many advances. Starting 20 years ago with the introduction of the CD, and progressing to MiniDisk and DAT, today we have a choice of high-resolution audio formats such as DVD-Audio, SACD, and the incredible flexibility and storage density of MP3. All these advances concentrate on *storage* media for the music. But how does the audio signal get processed once it comes off the storage media? How does it get to the output of a system? Are current "digital" systems truly digital? The vast majority of systems today are not.

In audio/video (A/V) receivers digital signal processing became popular due to the nature of Dolby Digital decoding, but just about all popular audio systems, such as mini-components, car stereos and PC add-on speakers, are still using analog signal-processing technology.

The reason for this is that earlier digital solutions, based on general-purpose digital signal processors (DSPs), with separate D/A and A/D conversion ICs, carry significant overhead in terms of hardware and programming. Accordingly, the implementation of digital solutions has been difficult and prohibitively expensive.

Now, at a time when major consumer systems are completing the crossover to all-digital media, Analog Devices introduces the AD1954, the first member of the new, cost-effective SigmaDSP family—the first solution to deal directly with the essential problem of integrating an audio-specific DSP, together with high-performance audio converters, on a single IC.

This family of digital sound processing devices, led by the AD1954, offers:

- Professional quality digital sound processing with integrated converters (112-dB SNR)
- ZERO-hassle programming, along with a highly user-friendly graphical configuration tool
- Very low price, which allows a majority of systems to provide the superior sonic qualities of digital technology.

What's inside the AD1954?

The AD1954 is a pre-programmed, fully configurable digital audio processor. The internal structure is shown in Figure 1. It is optimized for 2.0 (left/right) and 2.1 (left/right + subwoofer) configurations, and comprises the following processing blocks:

- 3-to-1 digital source selector
- Stereo 7-band equalizer (48-bit double-precision filters)
- Professional quality dynamic processor (dual-band structure)
- Phat™ Stereo spatial enhancement
- Delay for speaker location adjustment
- Crossover for independent subwoofer processing
- Volume control
- Three D/A converters (112-dB SNR) for left-, right-, and sub-woofer outputs.

All parameters are *fully configurable* by the system designer. This allows quick design cycle times while permitting full flexibility in adjusting the system to the specific requirements of each market and customer.

Which DSP core is used in a SigmaDSP?

The DSP core used in these products is entirely new—it has been optimized for the requirements of audio processing (Figure 2). Among the features that drastically reduce the cycle count required for it to perform a given audio algorithm are hardware accelerators for double precision and dynamic processing. The DSP core is based on a 26 × 22 multiply-accumulate engine with dual 48-bit accumulators. The input word length is 24 bits, but the core's internal resolution of 26 bits (3.23 format) provides two extra bits for up to +12 dB of gain. Since +12 dB of gain is common in many audio algorithms, no scaling is necessary in most applications.

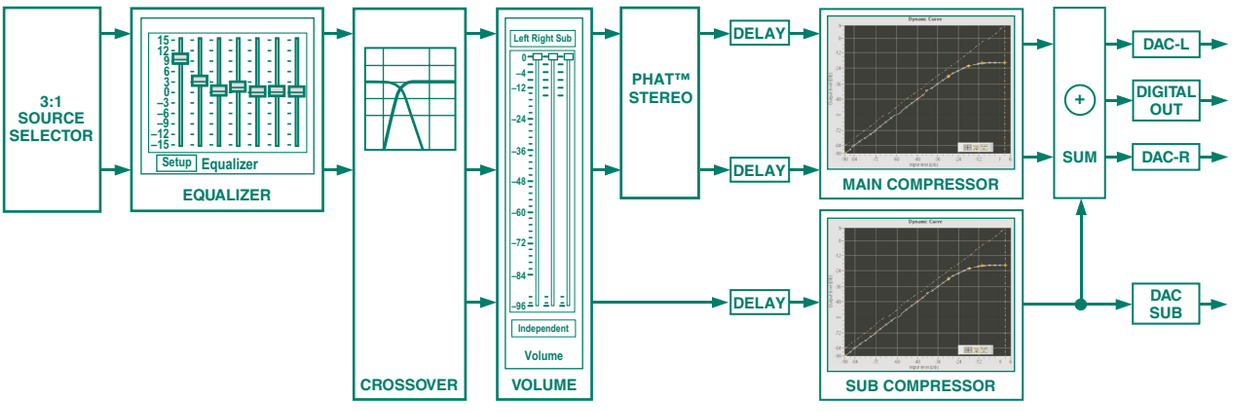

Figure 1. SigmaDSP architecture.

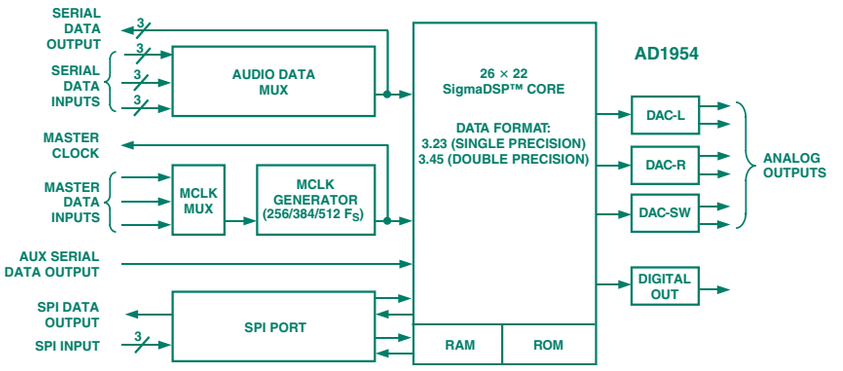

Figure 2. SigmaDSP processing core.

All filters are calculated in 48-bit double-precision resolution, utilizing special hardware accelerators. Double precision ensures that low-frequency IIR filters can be operated without limit-cycle problems, which cause audible artifacts.

The core memory comprises 2.5 KB of program RAM, 2.5 KB of program ROM, and 1 KB of parameter RAM. All memory is directly accessible through the SPI interface, which uses a self-addressing 32-bit format (8-bit address, 24-bit data) that allows single-cycle access to any memory location. The internal clock rate of the AD1954 is 25 MHz—equal to about 50 MIPS of a general-purpose DSP due to the hardware acceleration.

SigmaDSP Graphical User Interface (GUI) gives the designer *Total Setup* Control in real time.

SigmaDSP technology is targeted to both the experienced digital designer and the analog designer who knows his audio system well but does not want to dive into the intricacies of low-level DSP programming. Meeting both objectives requires a tool that can be operated intuitively, but provides control of the entire signal flow in real time.

The solution is the AD1954 graphical user interface (GUI), an example of which is shown in Figure 3. This graphical representation of the AD1954 signal flow makes the use of this combination (part plus GUI) truly intuitive. Every parameter in

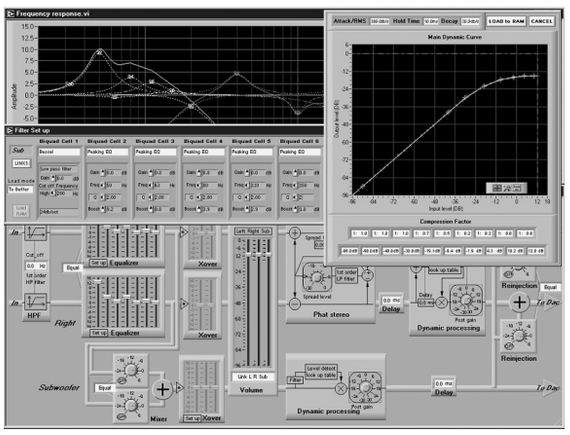

Figure 3. SigmaDSP graphical user interface.

the signal chain, including filter coefficients, volume settings, and dynamic processing functions, can be directly accessed and altered in real time. The GUI connects through the printer port of the PC to the AD1954 evaluation board. In this way, any parametric changes are sent to the AD1954 in SPI format and become immediately effective (in real time).

While the SigmaDSP GUI is intended as a dedicated tool for the system designer, it could also be made available (in a modified version) to the enthusiastic end user. With this PC interface, a user can have total control via a notebook PC.

Why is professional quality dynamic processing so essential?

Small and medium size systems, particularly car sound systems, are often limited by their amplifier and speaker power. There are several potential limiting factors: In *car systems*, the barrier is simply the 12-V supply, which limits the maximum output power to about 20 W rms into 4 Ω or 40 W rms into 2 Ω In *mini-components* there is a space constraint on transformer size, and there are thermal constraints. Another factor is speaker equalization: In small and medium systems, we typically see small cabinets containing small loudspeakers. A popular solution is to use relatively heavy equalization, particularly in the bass section (bass boost) calling for increased power at low frequencies, to compensate for this acoustically imperfect setup. Finally (driven by the younger generation), there is a common desire—if not a demand—that these systems have a high maximum volume.

This combination of limited amplifier power, heavy bass equalization, and significant total loudness of the system leads readily to a situation where the amplifiers are saturated and start to introduce heavy distortion—resulting in an unsatisfactory and annoying listening experience. In the past, attempts to solve this problem typically used primitive clipped-signal detectors, which avoided the clipping but led to artifacts that were nearly as bad as the clipping distortion itself. The professional quality, dual-band dynamic processors of the AD1954 SigmaDSP make it possible to control the system limits without artifacts.

Increase clarity and loudness of your system

Figures 4 and 5 show an example of a transfer function without any dynamic processing and with a soft knee compressor/limiter function. By using the function shown in Figure 5, the professional quality dynamic processing of the AD1954 allows natural handling of the clip level, resulting in lower distortion at high volumes. This effectively allows a user to turn up system volume by about 10 dB. A 10-dB increase in volume represents doubling the experienced sound pressure level; thus the user can run his system *twice as loud*.

This is particularly critical in smaller systems with limited amplifier and speaker power.

Real-World adjustments

The arbitrarily adjustable transfer function of the SigmaDSP dynamic processors permits many other applications, and the total flexibility allows a user to combine several applications into the same device.

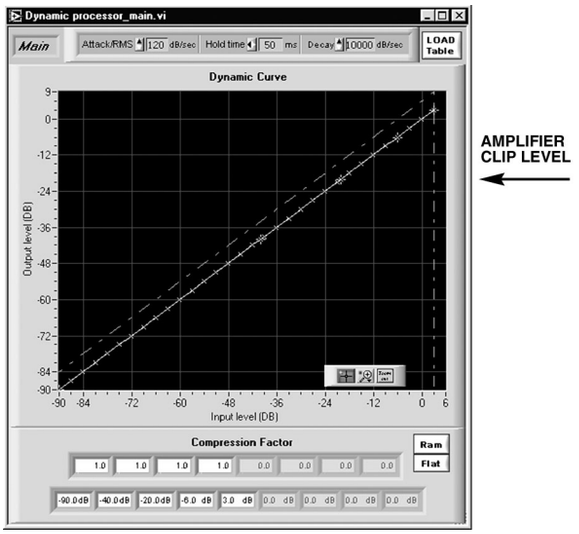

Figure 4. Without Dynamic Processing, the audio signal will get increasingly distorted if allowed to increase linearly above the clip level.

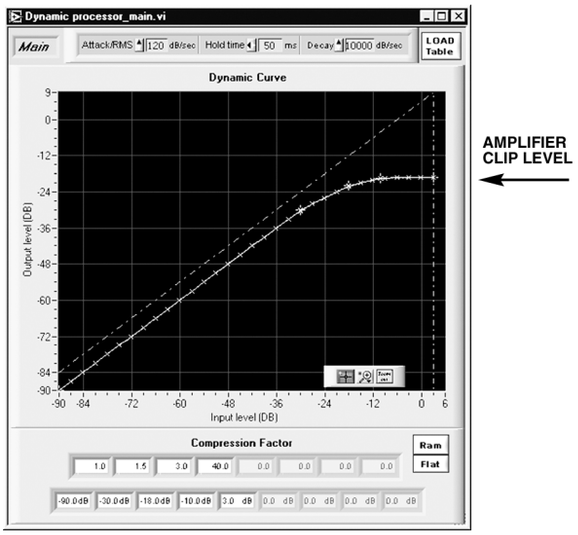

Figure 5. With Dynamic Processing, the gain is dynamically adjusted to produce accurate compression so that the peak amplitudes will not be driven beyond the limits of the amplifier/speaker system.

Road Noise Compensation

A powerful use of SigmaDSP dynamic processing is in car systems. Besides the sophisticated equalizing strategies and distortion handling, it is possible to compensate for the noisy environment of the car interior. Since the noise itself cannot be decreased (despite ongoing research projects in active noise cancellation), the only practical way to improve the listening experience is to make soft signals louder. This is especially critical with classical music, where a large dynamic range and many passages of quiet music are common.

With a proper transfer function programmed into the AD1954, signals below a certain threshold can be compressed and the music level can be amplified to maintain a level above the road noise. The ultimate use of this technology can be achieved in OEM car systems, where signals from the speedometer and RPM control signals make it possible for the compensation ratio to be altered depending on the speed (wind noise) and engine RPM (motor noise). Anyone who has experienced a smart road noise compensation strategy in a car sound system never wants to be without it again—with the possible exception of drivers of a mid-20th-century Rolls Royce. ("At 60 mph, the loudest sound you hear is the ticking of its electric clock.")

Midnight mode

A final example of a special use of the SigmaDSP dynamic processing technology is in the compensation of sudden variations of loudness in movie soundtracks. It seems that at home, the loudness of the soundtrack—especially in action movies—is just never right. (Your remote-control finger knows this best.) One reason is that soundtracks are typically mixed for movie theaters. In the cinema, the acoustical experience is an essential element of the entire movie experience, and large dynamic variations are a tool that the director uses to generate excitement. In theaters nobody gets disturbed. At home this is all different. While we still want the movie experience to be as good as possible, we have to avoid waking the children or nearby neighbors.

A "midnight" mode can handle this annoying problem automatically by reducing the dynamic range of the soundtrack. To implement this feature, a similar transfer function to that used for dynamic clip control is used, but at a much lower threshold. To avoid audible artifacts, a professional type dual-band dynamic processor like that implemented in the AD1954 is necessary.

The two pairs of sound tracks (Figure 6) are for voices followed by an action scene (bomb explosion). One can see that, while the dynamic range of the action scene gets reduced when the midnight mode is turned on, the voices stay at the same level.

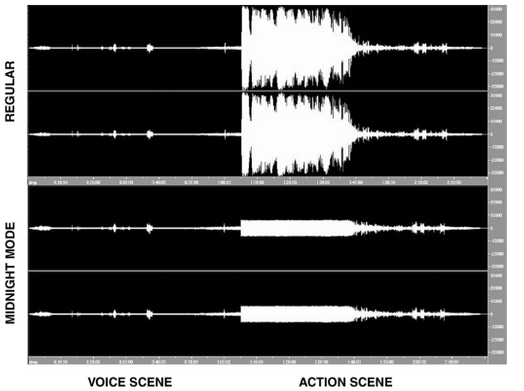

Figure 6. In midnight mode, high levels are suppressed without affecting normal levels.

SUMMARY

With the introduction of SigmaDSP technology, the experience of listening to music and movies migrates into a new era. Processing performance, converter technology, and sophisticated algorithms, which have formerly been known only to owners of professional recording studios, are becoming available for cost-sensitive consumer systems. By using SigmaDSP technology, one can develop stereo systems that remain digital until the final output in order to take full advantage of the superb quality of today's digital media.

AD1954 Availability

Samples of the AD1954, the first member of ADI's new SigmaDSP family, are available now, and evaluation boards and control software are available. The AD1954 comes in 44-lead MQFP (YS) and 48-lead TQFP (YST) packages and is rated for the extended temperature range of –40°C to +105°C. Pricing in 10,000 piece quantities is $5.88 for the AD1954YST.

How Ultrasound System Considerations Influence Front-End Component Choice

by Eberhard Brunner [eberhard.brunner@analog.com]

INTRODUCTION

There are major trade-offs to be considered when designing ultrasound front-end circuits. Performance parameters in the front-end circuit components affect diagnostic performance—and conversely, system configuration and objectives affect the choice of components.

It is essential for designers to understand the specifications that are of particular importance, their effect on system performance, and how they are affected by integrated-circuit (IC) design trade-offs—in terms of integration and semiconductor process technology—that will limit user design choices. Awareness of these considerations will help the designer to achieve the most advantageous system partitioning. We start with a high level system overview, followed by a more detailed description of how ultrasound systems work.

System Introduction

Medical ultrasound machines are among the most sophisticated signal processing machines in widespread use today. As in any complex machine, there are many trade-offs in implementation due to performance requirements, physics, and cost. Some system-level understanding is necessary to fully appreciate the desired front-end IC functions and performance levels, especially for: the low-noise amplifier (LNA); time gain compensation amplifier (TGC); and analog-to-digital converters (ADCs).

In ultrasound front-ends—as well as many other sophisticated electronic systems, these analog signal processing components are key elements in determining the overall system performance. The front-end component characteristics define the limits on system performance; once noise and distortion have been introduced, it is virtually impossible to remove them. This is, of course, a general problem in any receive signal-processing chain, be it ultrasound or wireless.

It is interesting to consider that ultrasound is basically a radar or sonar system, but it operates at speeds that differ from these by orders of magnitude. A typical ultrasound system is almost identical in concept to the phased-array radar systems on board commercial and military aircraft, and on military ships. Radar works in the GHz range, sonar in the kHz range, and ultrasound in the MHz range. Ultrasound designers adopted and expanded on the principle of steering beams using phased arrays, originated by radar system designers. Today those systems involve some of the most sophisticated signal processing equipment to be found.

Figure 1 shows a simplified diagram of an ultrasound system. In all such systems there is a multi-element transducer at the end of a relatively long (about 2-m) cable. Containing from 48 to 256 micro-coaxial cables, the cable is one of the most expensive parts of the system. In most systems, several different transducer probe heads (also called handles—a handle is the unit that contains the transducer elements and is attached to the system via cable) are available to be connected to the system, allowing the operator to select the appropriate transducer for optimal imaging. The handles are selected via high voltage (HV) relays, which add large parasitic capacitances to those of the cable.

A HV multiplexer/demultiplexer is used in some arrays to reduce the complexity of transmit and receive hardware, but at the expense of flexibility. The most flexible systems are phased-array digital beamformer systems—they also tend to be the most costly systems, due to the need for full electronic control of all channels. However, today's state-of-the-art front-end ICs, like the AD8332 variable-gain amplifier (VGA) and the AD9238 12-bit analog-to-digital converter (ADC) are pushing the cost-per-channel down continuously, so that full electronic control of all elements is now being introduced even in medium to low cost systems.

On the transmit (Tx) side, the Tx beamformer determines the delay pattern and pulse train that set the desired transmit focal point. The outputs of the beamformer are then amplified by high voltage

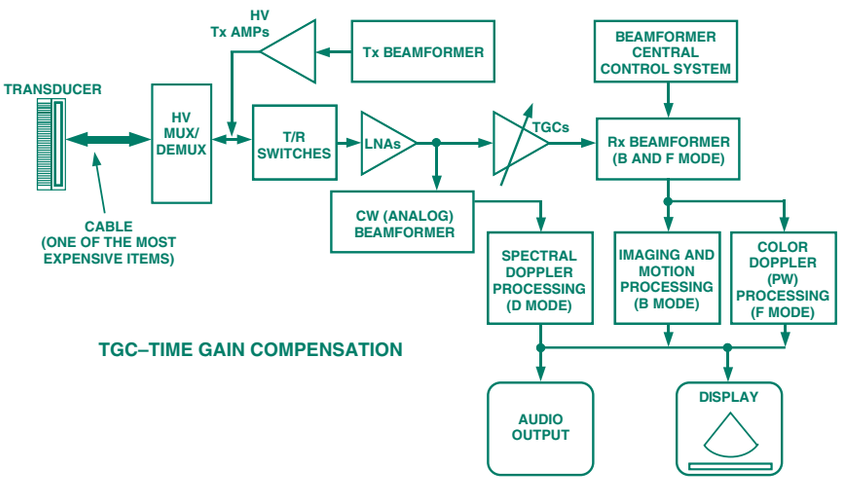

Figure 1. Ultrasound system block diagram.

transmit amplifiers that drive the transducers. These amplifiers might be controlled by digital-to-analog converters (DACs) to shape the transmit pulses for better energy delivery to the transducer elements. Typically, multiple transmit focal regions (zones) are used—that is, the field to be imaged is deepened by focusing the transmit energy at progressively deeper points in the body. The main reason for multiple zones is that the transmit energy needs to be greater for points that are deeper in the body, because of the signal's attenuation as it travels into the body (and as it returns).

On the receive (Rx) side, there is a T/R switch, generally a diode bridge, which blocks the high voltage Tx pulses. It is followed by a low-noise amplifier (LNA) and one or more variable-gain amplifiers (VGAs), which implement time gain compensation (TGC) and sometimes also apodization (spatial "windowing" to reduce sidelobes in beam) functions. Time gain control—which provides increased gain for signals from deeper in the body (and therefore arriving later)—is under operator control and used to maintain image uniformity.

After amplification, beamforming is performed, implemented in either analog (ABF) or digital (DBF) form. It is mostly digital in modern systems, except for continuous-wave (CW) Doppler processing, whose dynamic range is still too large to be processed through the same channel as the image. Finally, the Rx beams are processed to show either a gray-scale image, Colorflow overlay on the 2-D image, and/or a Doppler output.

Ultrasound System Challenges

To fully understand the challenges in ultrasound and their impact on the front-end components, it is important to remember what this imaging modality is trying to achieve. First, it is supposed to give an accurate representation of the internal organs of a human body, and second, through Doppler signal processing, it is to determine movement within the body (for example, blood flow). From this information a doctor can then make conclusions about the correct functioning of a heart valve or blood vessel.

Acquisition Modes

There are three main ultrasonic acquisition modes: B-mode (gray-scale imaging; 2D); F-mode (Colorflow or Doppler Imaging; blood flow); and D-mode (Spectral Doppler). B-mode creates the traditional gray-scale image; F-mode is a color overlay on the B-mode display that shows blood flow; D-mode is the Doppler display that might show blood flow velocities and their frequencies. (There is also an M-mode, which displays a single B-mode time line.)

Operating frequencies for medical ultrasound are in the 1-MHz to 40-MHz range, with external imaging machines typically using frequencies of 1 MHz to 15 MHz, while intravenous cardiovascular machines use frequencies as high as 40 MHz. Higher frequencies are in principle more desirable, since they provide higher resolution—but tissue attenuation limits how high the frequency can be for a given penetration distance. However, one cannot arbitrarily increase the ultrasound frequency to get finer resolution, since the signal experiences an attenuation of about 1 dB/cm/MHz; i.e., for a 10-MHz ultrasound signal and a penetration depth of 5 cm, the round-trip signal has been attenuated by $5 \times 2 \times 10 = 100$ dB! To handle an instantaneous dynamic range of about 60 dB at any location, the required dynamic range would be 160 dB (a voltage dynamic range of 100 million to 1)! Dynamic ranges of this magnitude are not directly achievable; therefore one has to pay the costs of a highly sophisticated system and trade off something at the front end—either penetration depth (limited by safety regulations due to maximum transmit power that is allowed) or image resolution (using a lower ultrasound frequency).

The large dynamic range of the received signals presents the most severe challenge. The front-end circuitry must have very low noise and large-signal handling capability simultaneously—requirements familiar to anyone experienced in the demands of communications. Cable mismatch and loss directly add to the noise figure of the system. For example, if the loss of the cable at a particular frequency is 2 dB, then the NF is degraded by 2 dB. This means that the first amplifier after the cable will have to have a noise figure that is 2 dB lower than that needed with a lossless cable. One potential way to get around this problem is to situate an amplifier in the transducer handle. However, there are serious size and power constraints; also, the need for protection from high voltage transmit pulses makes such a solution difficult to implement.

Another challenge is the large acoustic impedance mismatch between the transducer elements and the body. The acoustic impedance mismatch requires matching layers (analogous to electrical-impedance-matching RF circuitry) to transmit energy efficiently. This normally consists of a couple of matching layers in front of the transducer elements in the handle, followed by a lens, followed by coupling gel. The gel establishes good acoustic contact with the body—since air is a very good acoustic reflector.

Another important issue for the receive circuitry is fast overload recovery. Even though the T/R switch is supposed to protect the receiver from large pulses, a small fraction of these pulses leaking across the switches can be sufficient to overload the front-end circuitry. Poor overload recovery will make the receiver "blind" until it recovers, with a direct impact on how close to the surface of the skin an image can be generated.

How an Ultrasound Image is Generated—B-Mode

Figure 2 shows how the different scan images are generated. In all four scans, the pictures with the scan lines bounded by a rectangle are an actual representation of the image, as it will be seen on the display monitor. Mechanical motion of a single transducer (in the directions indicated by the arrows) is shown here to facilitate understanding of the image generation; but the same kinds of images can be generated by a linear array without mechanical motion. In the example of a linear scan, the transducer element is moved in a horizontal direction; for every scan line (the lines shown in the images), a Tx pulse is sent and the reflected signals from different depths are recorded and scan-converted to be shown on a video display. How the single transducer is moved during image acquisition determines the shape of the image. This directly translates into the shape of a linear array transducer, i.e., for the linear scan, the array would be straight, while for the arc scan, the array would be concave.

The step that is needed to go from a mechanical single transducer system to an electronic system can also be easily explained by examining the linear scan in Figure 2. If the single transducer element is divided into many small pieces, then if one excites one element at a time and records the reflections from the body, one also gets the rectangular image as shown, only now one does not need to move the transducer elements. From this one can see that the arc scan can be made of a linear array that has a concave shape; and the sector scan would be made of a linear array that has a convex shape.

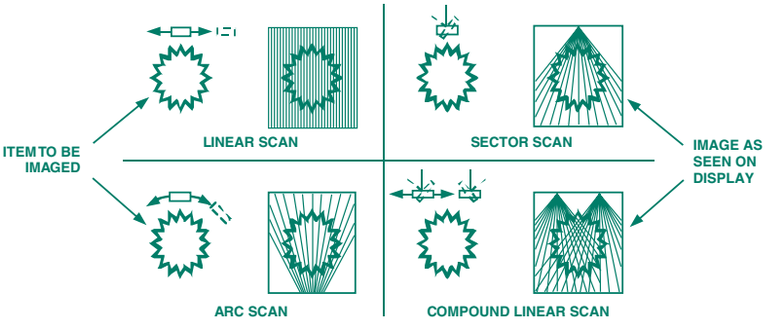

Figure 2. Single-transducer image generation.

Even though the example above explains the basics for B-mode ultrasound image generation, in a modern system more than one element at a time is used to generate a scan line because it allows the aperture of the system to be changed. Changing the aperture is like changing the location of the focal point in optics—it helps create clearer images. Figure 3 shows how this is done for a linear array and a phased array; the main difference is that in a phased array all elements are used simultaneously, while in a linear array only a subset of the total array elements is used. Using a smaller number of elements has the advantage of saving electronic hardware; but it increases the time to image a given field of view. A phased array is different; because of its pie shape a very small transducer can image a large area in the far field. That is why phased array transducers are the transducers of choice in applications like cardiac imaging where one has to deal with the small spaces between the ribs through which the much larger heart needs to be imaged.

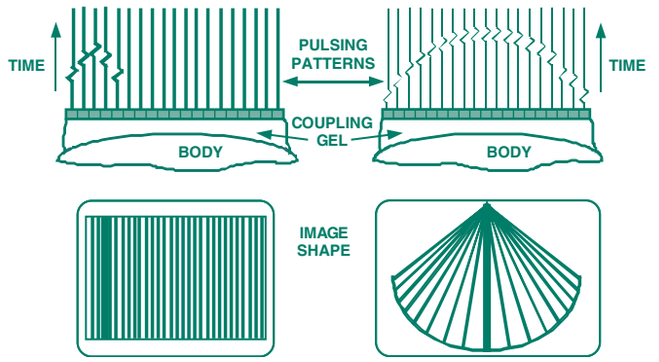

Figure 3. Linear vs. phased-array imaging.

Excitation in arrays is directed along scan lines, determined by the delay profile of a set of pulses intended to arrive simultaneously at a focal point. The pulses (Figure 3) are represented by the "squiggles" on the vertical time lines above the array (shaded color)—with time increasing vertically from the array surface. The linear stepped array, in Figure 3, will deliver shaped excitation to a group of elements (aperture), then step the aperture by adding a leading element and dropping a trailing one. On each step one scan line (beam) is formed by the simultaneous arrival of the pulses. In the phased array, all transducers are active at the same time. In the examples shown, the darkened lines are the scan lines imaging the reflection data produced by the representative pulsing patterns.

Analog vs. Digital Beamforming

In analog beamforming (ABF) and digital beamforming (DBF) ultrasound systems, the received pulses reflected from a particular focal point along a beam are stored for each channel, then aligned in time, and coherently summed—this provides spatial processing gain because the noise of the channels is uncorrelated. Images may be formed as either a sequence of analog levels that are delayed with analog delay lines, summed, and converted to digital after summation (ABF)—or digitally by sampling the analog levels as close as possible to the transducer elements, storing them in a memory (FIFO), and then summing them digitally (DBF). Figures 4 and 5 show basic respective block diagrams of ABF and DBF systems. Both types of systems require perfect channel-to-channel matching. Note that the variable-gain amplifiers (VGAs) are needed in both implementations—and will continue to be in the digital case until ADCs with a large enough dynamic range become available at reasonable cost and low enough power. Note that an

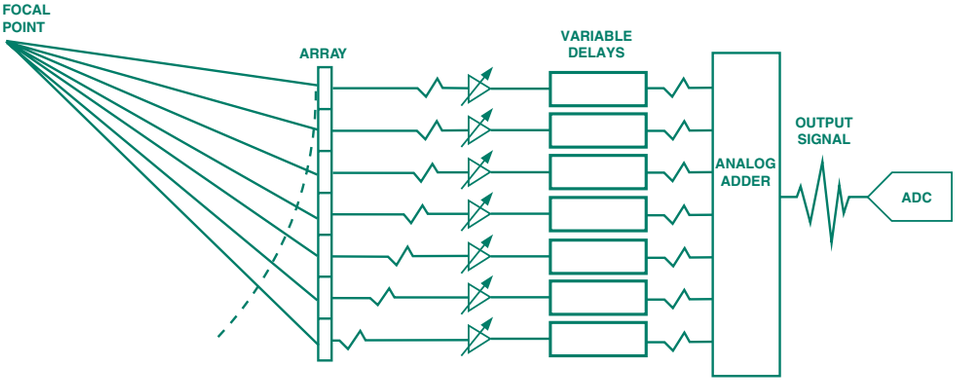

Figure 4. Simplified block diagram of ABF system.

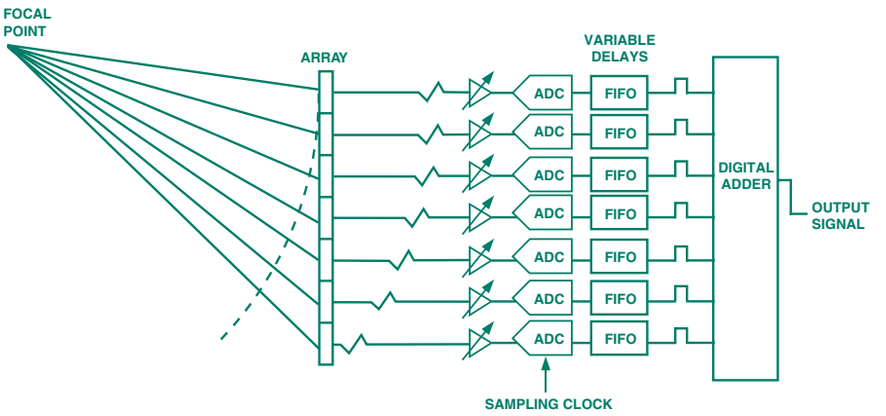

Figure 5. Simplified block diagram of DBF system.

ABF imaging system needs only one very high resolution and high speed ADC, but a DBF system requires many high speed, high resolution ADCs. Sometimes a logarithmic amplifier is used in the ABF systems to compress the dynamic range before the ADC.

Dynamic Range

In the front-end circuitry, the noise floor of the LNA determines how weak a signal can be received. But at the same time—especially during CW Doppler signal processing—the LNA must also be able to handle very large signals. So it is crucial to maximize the dynamic range of the LNA (in general, it is impossible to implement any filtering before the LNA due to noise constraints). Note that these same conditions apply for any receiver—in communications applications, the circuitry closest to the antenna does not have the advantage of a lot of filtering either; accordingly, it needs to cope with the largest dynamic range.

CW Doppler has the largest dynamic range of all signals in an ultrasound system—during CW, a sine wave is transmitted continuously with half of the transducer array, while the other half is receiving. There is a strong tendency for the Tx signal to leak into the Rx side; and there are also strong reflections coming from stationary body parts that are close to the surface. This tends to interfere with examination of, for example, blood flow in a vein deep in the body with concomitant very weak Doppler signals.

At the current state of the art, CW Doppler signals cannot be processed through the main imaging (B-mode) and PW Doppler (F-mode) path in a digital beamforming (DBF) system; for this reason, an analog beamformer (ABF) is indicated for CW Doppler processing in Figure 1. The ABF has larger dynamic range. Naturally, the "Holy Grail" in DBF ultrasound is for all modes to be processed through the DBF chain (at realistic cost), and there is a great deal of ongoing research as to how to get there.

Power

Since ultrasound systems require many channels, power consumption of all the front-end components—from T/R switch, through LNA, VGA, and ADC, to the digital circuitry of the beamformer—is a very critical specification. As has been pointed out above, there will always be a push to increase the front-end dynamic range in order to arrive at eventual integration of all ultrasound modes into one beamformer—a tendency that will lead towards increasing the power in the system. However, there is a corresponding need to make the ultrasound systems forever smaller—with a tendency towards reducing power. Power in digital circuits usually decreases with supply voltage; but this is not necessarily true for analog and mixed signal circuitry. Furthermore, taking into account the fact that reduced analog "headroom" tends to reduce dynamic range, there will be a limit to how low the supply voltage can go and still achieve a desired dynamic range.

CONCLUSION

We have sought to show here the trade-offs required in front-end ICs for ultrasound by explaining the basic operation of such a system first, and then pointing out what particular performance parameters are needed to ensure optimal system operation. A more complete version of this paper[1] is available to provide additional details.

[1]Brunner, Eberhard, "*Ultrasound System Considerations and their Impact on Front-End Components*," Analog Devices, Inc., 2002 (http://www.analog.com/library/analogDialogue/archives/36-03/ultrasound/UltrasoundFrontend.pdf).

Ask the Applications Engineer—30

by Adrian Fox [adrian.fox@analog.com]

PLL SYNTHESIZERS

Q. What is a PLL Synthesizer?

A. A *frequency synthesizer* allows the designer to generate a variety of output frequencies as multiples of a single reference frequency. The main application is in generating local oscillator (LO) signals for the up- and down-conversion of RF signals.

The synthesizer works in a *phase-locked loop* (PLL), where a phase/frequency detector (PFD) compares a fed back frequency with a divided-down version of the reference frequency (Figure 1). The PFD's output current pulses are filtered and integrated to generate a voltage. This voltage drives an external voltage-controlled oscillator (VCO) to increase or decrease the output frequency so as to drive the PFD's average output towards zero.

Frequency is scaled by the use of counters. In the example shown, an ADF4xxx synthesizer is used with an external filter and VCO. An input reference (R) counter reduces the reference input frequency (13 MHz in this example) to PFD frequency ($F_{PFD} = F_{REF}/R$); and a feedback (N) counter reduces the output frequency for comparison with the scaled reference frequency at the PFD. At equilibrium, the two frequencies are equal, and the output frequency is $N \times F_{PFD}$. The feedback counter is a *dual-modulus prescaler type*, with A and B counters ($N = BP + A$, where P is the prescale value).

Figure 2 shows a typical application in a superheterodyne receiver. Base station and handset LOs are the most common application, but synthesizers are also found in low frequency clock generators (ADF4001), wireless LANs (5.8 GHz), radar systems, and collision-avoidance systems (ADF4106).

Q. What are the key performance parameters to be considered in selecting a PLL synthesizer?

A. The major ones are: phase noise, reference spurs, and lock time.

Phase Noise: For a carrier frequency at a given power level, the phase noise of a synthesizer is the ratio of the carrier power

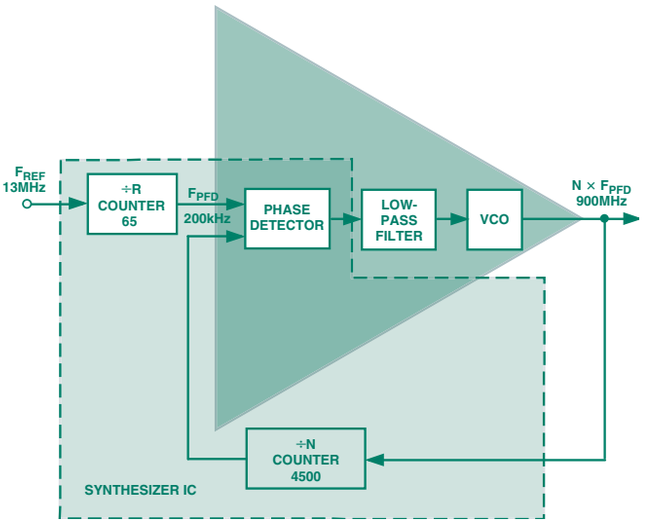

Figure 1. Block diagram of a PLL.

to the power found in a 1-Hz bandwidth at a defined frequency offset (usually 1 kHz for a synthesizer). Expressed in dBc/Hz, the in-band (or close-in) phase noise is dominated by the synthesizer; the VCO noise contribution is high-pass filtered in the closed loop.

Reference Spurs: These are artifacts at discrete offset frequencies generated by the internal counters and charge pump operation at the PFD frequency. These spurs will be increased by mismatched up and down currents from the charge pump, charge-pump leakage, and inadequate decoupling of supplies. The spurious tones will get mixed down on top of the wanted signal and decrease receiver sensitivity.

Lock Time: The lock time of a PLL is the time it takes to jump from one specified frequency to another specified frequency within a given frequency tolerance. The jump size is normally determined by the maximum jump the PLL will have to accomplish when operating in its allocated frequency band. The step-size for GSM-900 is 45 MHz and for GSM-1800 is 95 MHz. The required frequency tolerances are 90 Hz and 180 Hz, respectively. The PLL must complete the required frequency step in less than 1.5 time slots, where each time slot is 577 µs.

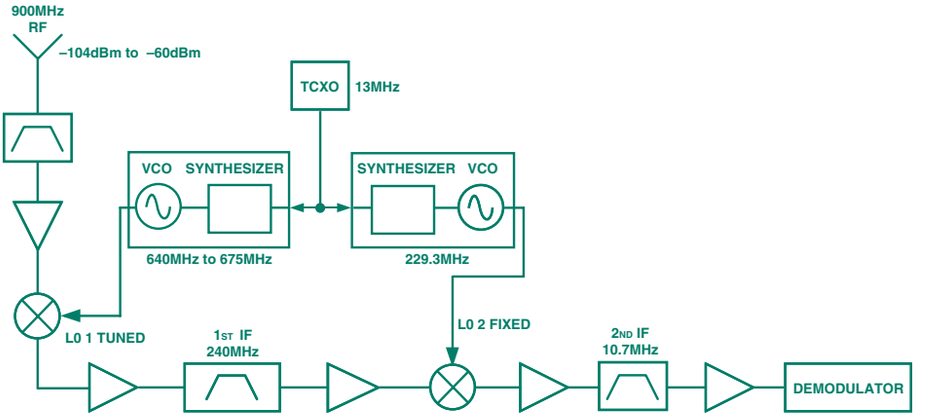

Figure 2. Dual PLL used to mix down from GSM RF to baseband.

Q. *I've selected my synthesizer based on the output frequency required. What about choosing the other elements in the PLL?*

A. **Frequency Reference**: A good, high quality, low-phase-noise reference is crucial to a stable low-phase-noise RF output. A square wave or clipped sine wave available from a TCXO crystal offers excellent performance, because the sharper clocking edge results in less phase jitter at the R-counter output. The ADF4206 family features on-board oscillator circuitry allowing low cost AT-cut crystals to be used as the reference. While predictable AT crystals cost one third as much as TCXOs, their temperature stability is poor unless a compensation scheme with a varactor is implemented.

VCO: The VCO will convert the applied tuning voltage to an output frequency. The sensitivity can vary drastically over the full frequency range of the VCO. This may make the loop unstable (see *loop filter*). In general, the lower the tuning sensitivity (Kv) of the VCO, the better the VCO phase noise will be. The synthesizer phase noise will dominate at smaller offsets from the carrier. Farther away from the carrier, the high-pass-filtered noise of the VCO will begin to dominate. The GSM specification for out-of-band phase noise is –130 dBc/Hz at a 1-MHz offset.

Loop Filter: There are many different types of loop filter. The most common is the third-order integrator shown in Figure 3. In general, the loop filter bandwidth should be 1/10 of the PFD frequency (channel spacing). Increasing the loop bandwidth will reduce the lock time, but the filter bandwidth should never be more than PFD/5, to avoid significantly increasing the risk of instability.

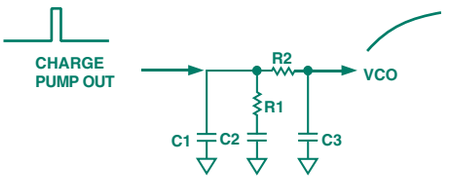

Figure 3. A third-order loop filter. The R2C3 pole provides extra attenuation for spurious products.

A loop filter's bandwidth can be doubled by doubling either the PFD frequency or the charge-pump current. If the actual Kv of the VCO is significantly higher than the nominal Kv used to design the loop filter, the loop bandwidth will be significantly wider than expected. The variation of loop bandwidth with Kv presents a major design challenge in wideband PLL designs, where the Kv can vary by more than 300%. Increasing or decreasing the programmable charge-pump current is the easiest way to compensate for changes in the loop bandwidth caused by the variation in Kv.

Q. How do I optimize PLL design for phase noise?

A. *Use low N-value*: Since phase noise is multiplied up from the PFD (reference frequency) at a rate of 20 log$N$, reducing N by a factor of 2 will improve system phase noise by 3 dB (i.e., doubling the PFD frequency reduces phase noise by 10 log2). Therefore the highest feasible PFD frequency should always be used.

Choose a higher frequency synthesizer than is required: Operating under the same conditions at 900 MHz, the ADF4106 will give 6-dB better phase noise than the ADF4111 (see Table 1).

Use the lowest Rset resistor specified for operation: Reducing the Rset increases the charge-pump current, which reduces phase noise.

Table 1. The integrated phase jitter depends heavily on the in-band phase noise of the synthesizer. System parameters: [900-MHz RF, 200-kHz PFD, 20-kHz loop filter]

Synthesizer Model	In-Band Phase Noise (dB)	Integration Range (Hz)	Integrated Phase Error Degrees rms
ADF4111	–86	100 to 1 M	0.86
ADF4112	–89	100 to 1 M	0.62
ADF4113	–91	100 to 1 M	0.56
ADF4106	–92.5	100 to 1 M	0.45

Q. Why is phase noise important?

A. Phase noise is probably the most crucial specification in PLL selection. In a *transmit* chain, the linear power amplifier (PA) is the most difficult block to design. A low-phase-noise LO will give the designer greater margin for non-linearity in the PA by reducing the phase error in the up-conversion of the baseband signal.

The system maximum phase error specification for GSM receivers/transmitters (Rx/Tx) is 5° rms. As one can see in Table 1, the allowable PA phase-error contribution can be significantly greater when the phase noise contributed by the PLL is reduced.

On the receive side, low phase noise is crucial to obtaining good receiver selectivity (the ability of the receiver to demodulate signals in the presence of interferers). In the example of Figure 4, on the left the desired low level signal is swamped by a nearby undesired signal mixing with the LO noise (enclosed dashed area). In this case the filters will be unable to block these unwanted interferers. In order to demodulate the desired RF signal, either the transmit side will require higher output power, or the LO phase noise will need to be improved.

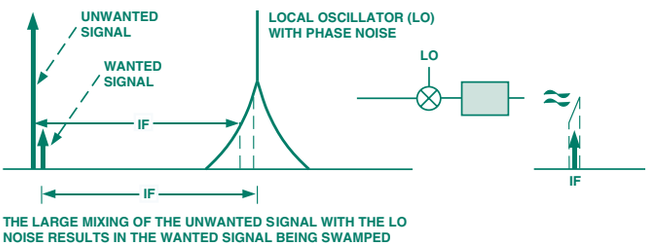

Figure 4. A large unwanted signal mixing with LO noise swamps the wanted signal. Increased phase noise will reduce the sensitivity of the receiver, since the demodulator will not be able to resolve the signal from the noise.

Q. Why are spur levels important?

A. Most communication standards will have stringent maximum specifications on the level of spurious frequency components (*spurs*) that the LO can generate. In *transmit* mode, the spur levels must be limited to ensure that they do not interfere with users in the same or a nearby system. In a *receiver*, the LO spurs can significantly reduce the ability to demodulate the mixed-down signal. Figure 4 shows the effect of reciprocal mixing

where the desired signal is swamped with noise due to a large undesired signal mixing with noise on the oscillator. The same effect will occur for spurious noise components.

A high level of spurs can indirectly affect lock time by forcing the designer to narrow the loop bandwidth—slowing response—in order to provide sufficient attenuation of these unwanted components. The key synthesizer specifications to ensure low reference spurs are *low charge-pump leakage* and *matching of the charge pump currents*.

Q. Why is lock time important?

A. Many systems use frequency hopping as a means to protect data security, avoid muti-path fading, and avoid interference. The time spent by the PLL in achieving frequency lock is valuable time that cannot be used for transmitting or receiving data; this reduces the effective data rate achievable. Currently there is no PLL available than can frequency-hop quickly enough to meet the timing requirements of the GSM protocol. In base-station applications, two separate PLL devices are used in parallel to reduce the number of wasted slots. While the first is generating the LO for the transmitter, the second PLL is moving to the next allocated channel. In this case a super-fast (<10-µs) settling PLL would significantly reduce the bill of materials (BOM) and layout complexity.

Q. How do I minimize lock time?

A. By increasing the *PFD frequency*. The PFD frequency determines the rate at which a comparison is made between the VCO/N and the reference signal. Increasing the PFD frequency increases the update of the charge pump and reduces lock time. It also allows the *loop bandwidth* to be widened.

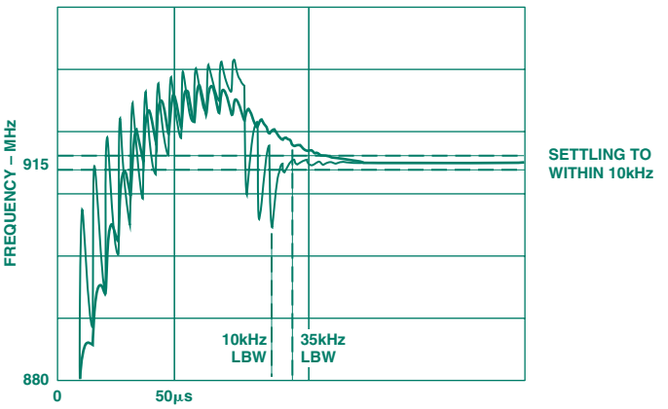

Figure 5. Loop bandwidth has a significant effect on the lock time. The wider the loop bandwidth, the faster the lock time, but also the greater the level of spurious components. Lock time to 1 kHz is 142 µs with a 35-kHz LBW—and 248 µs with a 10-kHz LBW.

Loop Bandwidth. The wider the loop bandwidth, the faster the lock time. The trade-off is that a wider loop bandwidth will reduce attenuation of spurious products and increase the integrated phase noise. Increasing the loop bandwidth significantly (>PFD/5) may cause the loop to become unstable and permanently lose lock. A phase margin of 45 degrees produces the optimum settling transient.

Avoid tuning voltages nearing ground or Vp. When the tuning voltage is within a volt of the rails of the charge pump supply (Vp), the charge pump begins to operate in a saturation region. Operation in this region will degrade settling time significantly; it may also result in mismatch between jumping-up in frequency and jumping down. Operation in this saturation region can be avoided by using the maximum Vp available or using an active loop filter. Using a VCO with a higher Kv will allow Vtune to remain closer to Vp/2 while still tuning over the required frequency range.

Choose plastic capacitors. Some capacitors exhibit a dielectric memory effect, which can impede lock time. For fast phase locking applications 'plastic-film' Panasonic ECHU capacitors are recommended.

Q. What factors determine the maximum PFD frequency I can use?

A. In order to obtain contiguous output frequencies in steps of the PFD frequency

$$F_{PFD} < \frac{V_{CO}\ Output\ Frequency}{(P^2 - P)}$$

where P is the prescaler value.

The ADF4xxx offers prescaler selections as low as 8/9. This permits a higher PFD frequency than many competitive parts, without violating the above rule—enabling lower phase noise PLL design. Even if this condition is not met, the PLL will lock if $B > A$ and $B > 2$ in the programming registers.

Q. Fractional-N has been around since 1970. What are its advantages to PLL designers?

A. The resolution at the output of an integer-N PLL is limited to steps of the PFD frequency. Fractional-N allows the resolution at the PLL output to be reduced to small fractions of the PFD frequency. It is possible to generate output frequencies with resolutions of 100s of Hz, while maintaining a high PFD frequency. As a result the N-value is significantly less than for integer-N. Since noise at the charge pump is multiplied up to the output at a rate of 20 logN, significant improvements in phase noise are possible. For a GSM900 system, the fractional-N *ADF4252* offers phase noise performance of –103 dBc/Hz, compared with –93 dBc/Hz for the *ADF4106* integer-N PLL.

Also offering a significant advantage is the lock-time improvement made possible by fractional-N. The PFD frequency set to 20 MHz and loop bandwidth of 150 kHz will allow the synthesizer jump 30 MHz in <30 µs. Current base stations require 2 PLL blocks to ensure that LOs can meet the timing requirements for transmissions. With the super-fast lock times of fractional-N, future synthesizers will have lock time specs that allow the 2 "ping-pong" PLLs to be replaced with a single fractional-N PLL block.

Q. If fractional-N offers all these advantages, why are integer-N PLLs still so popular?

A. Spurious levels! A fractional-N divide by 19.1 consists of the N-divider dividing by nineteen 90% of the time, and by twenty 10% of the time. The average division is correct, but the instantaneous division is incorrect. Because of this, the PFD and charge pump are constantly trying to correct for instantaneous phase errors. The heavy digital activity of the sigma-delta modulator, which provides the averaging function, creates spurious components at the output. The digital noise,

combined with inaccuracies in matching the hard-working charge pump, results in spurious levels greater than those allowable by most communications standards. Only recently have fractional-N parts, such as the ADF4252, made the necessary improvements in spurious performance to allow designers to consider their use in traditional integer-N markets.

Q. What PLL devices have you released recently, how do they differ, and where would I use them?

A. **ADF4001** is a <200-MHz PLL, pin-compatible with the popular ADF4110 series, but with the prescaler removed. Applications are stable reference clock generators, in cases where all clocks must be synchronized with a single reference source. They are generally used with VCXOs (voltage-controlled crystal oscillators), which have lower gain (Kv) and better phase noise than VCOs.

ADF4252 is a dual fractional-N device with <70 dBc spurious. It offers <20-µs lock times vs. 250 µs for integer-N, with <100 dBc/Hz phase noise due to the high PFD frequency—a ground-breaking product with a software-programmable trade-off between phase noise and spurs.

ADF4217L/ADF4218L/ADF4219L are low-phase-noise upgrades for the LMX2331L/LMX2330L/LMX2370. They consume only 7.1 mA, with a 4-dB improvement in phase noise over competitive devices. Great news for handset designers!

ADF4106 is a 6-GHz PLL synthesizer. Ideal for WLAN equipment in the 5.4-to-5.8-GHz frequency band, it is the lowest-noise integer-N PLL on the market.

Q. What tools are available to simulate loop behavior?

A. ADIsimPLL is a simulation tool developed with Applied Radio Labs. It consists of extensive models for the ADI synthesizers as well as popular VCOs and TCXOs. It allows the user to design passive and active loop filters in many configurations, simulate VCO, PLL, and reference noise, and model spurious and settling behavior. Once a design is completed, a custom evaluation board may be ordered based on the design using an internal weblink to Avnet.

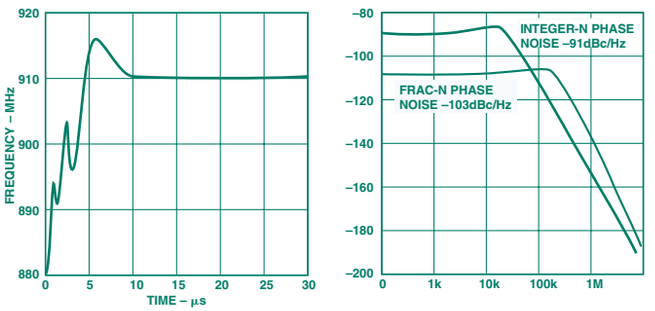

(a) (b)

Figure 6. Lock time and phase noise are just two parameters that can be modeled by ADIsimPLL. While phase noise is reduced by >8 dB, the wider loop bandwidths and high PFD frequency allowed by fractional-N reduce the lock time to <30 µs for 30-MHz jumps (as shown).

The tool is free and may be downloaded from www.analog.com/pll. Also widely used are the commercially available Eagleware and MATLAB tools.

Q. Do ADI proprietary parts have specific advantages over comparable competitive parts?

A. Phase noise is the critical specification for many system designers. Phase-noise performance in the ADF4113 family is typically 6 dB better than the National equivalent and >10 dB better than Fujitsu or Philips equivalents. The extended choice of prescaler settings protects the designer from being compromised in selecting a higher PFD frequency by the '$P^2 - P$' rule. Another major advantage is the choice of eight programmable charge-pump currents; in wideband designs where the gain of the VCO changes dramatically, the programmable currents can be adjusted to ensure loop stability and bandwidth consistency across the entire band.

Q. What is the future direction of the PLL industry?

A. While chipset solutions are prominent in the headlines, particularly for GSM, the new generation of cellular phone and base stations are still likely to initially favor discrete solutions. Discrete PLL and VCO modules offer improved noise performance and isolation, and are already in high volume production at the start of the design cycle.

The demand for reduced size and current consumption in handsets has driven the development of the ADI L-series of dual synthesizers on 0.35-µm Bi-CMOS in miniature CSP packages. Integrated VCO and PLL modules will be a major growth in newer system designs, where board area and cost reduction of an initial design is crucial.

However the most exciting developments are likely to be in fractional-N technology. Recent improvements in spur performance have allowed the release of the ADF4252 and created unprecedented interest. The phase-noise improvement, super-fast lock times, and versatility inherent in the architecture are likely to dominate LO blocks of future multi-standard high-data-rate wireless systems.

Acknowledgements

The author would like to thank Mike Curtin, Brendan Daly, and Ian Collins for their valuable contributions.

REFERENCES

(1) "Fractional-N Synthesizers," (Design Feature), *Microwaves and RF*, August 1999.

(2) *Microwave and RF Wireless Systems*, by David M. Pozar. Wiley (2000).

(3) "Phase locked loops for high frequency receivers and transmitters," by Mike Curtin and Paul O'Brien. *Analog Dialogue* Volume 33, 1999.

(4) *Phase-locked Loops*, by Roland E. Best. McGraw-Hill (1993).

(5) "Phase Noise Reference" (Application Note), Applied Radio Labs.

Versatile Mixed-Signal Front Ends Speed Customized Design of Wireline Broadband Modems and Home Networks

by Joseph DiPilato [joe.dipilato@analog.com],
Iuri Mehr [iuri.mehr@analog.com],
Brian Harrington [brian.harrington@analog.com]

INTRODUCTION

The widespread growth of broadband modems (predominantly based on cable and DSL technology) has made broadband communications available to residential households throughout the world for use in applications such as Internet access, interactive gaming, and telecommuting. The number of households with multiple personal computers (PCs) is also on the rise. The ready availability of these facilities has led to the increasing popularity of home networking for sharing Internet access and printer resources for work, academics, and entertainment. Figure 1 shows a typical domestic network connected to the broadband gateway and to various devices inside the residence/office.

A variety of technologies have been developed to address high-speed communication between home computers and peripheral devices—and to interface to the *hub* (or gateway) for broadband access. Until recently, *Ethernet* has been the most viable method for providing networking within the home. The attractiveness of Ethernet has been the availability of inexpensive network interface cards based on proven technology. However, Ethernet suffers from one main disadvantage—the requirement of having *Category 5* (CAT5) cable wired throughout the home. This almost always means that the homeowner must run new wires—a cumbersome and potentially expensive solution.

Wireless LAN provides an alternative to Ethernet that does not require installation of new wires. Despite the convenience of a wireless solution, it has achieved limited success due to higher cost, potential insecurity, multiple competing standards, lack of interoperability, and interference/robustness concerns. Recent advancements in the wireless standards within IEEE 802.11 and the emergence of the WiFi consortium have increased the prospects for wireless home networking. Nevertheless, the scarcity of frequency spectrum and the ubiquity of potent interferers—such as microwave ovens, garage door openers, etc.—will continue to pose numerous challenges to a wireless home network, including increased cost.

The newer preferred *wireline* technologies are those that require "no new wire"; they offer the homeowner a potentially better way of establishing a home network by avoiding the burden of installing new cabling, while providing comparable performance to that of Ethernet at an affordable price. The two choices for operating over the installed infrastructure involve either telephone or power-line wiring. Until recently, utilizing existing home wiring for broadband networking was impossible due to very poor channel quality. However, cost-effective broadband home networking is now a reality because of advances in signal-processing techniques, lowered costs resulting from perennial reduction of silicon fabrication geometries, and performance improvements in high-speed mixed-signal circuitry. The Home Phone Networking Alliance (HPNA) standard provides the framework for *phone-line* networking and allows data rates of up to 10 Mbits/s. And the HomePlug™ standard describes the specifications for implementing a *power-line* network system with data rates comparable to those of HPNA.

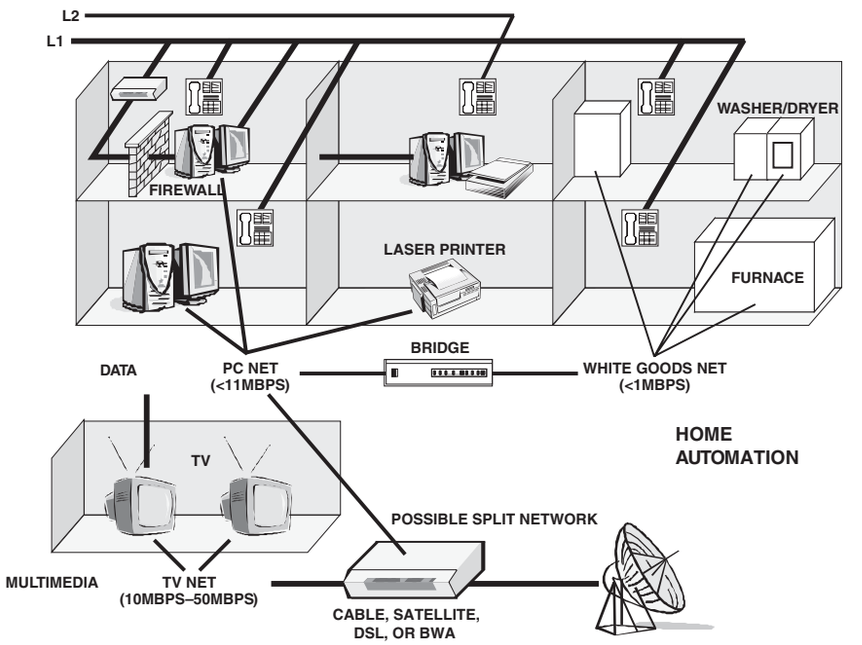

Figure 1. Home-networking/broadband access configuration.

These wireline home networking technologies, together with high-speed access technologies like DSL and cable, require mixed-signal interfacing between the transmission medium (power line, cable, twisted pair) and the digital baseband processors and controllers. Analog Devices has developed a family of monolithic mixed-signal front-end (MxFE™) integrated circuits (ICs) to bridge this gap (see *Analog Dialogue* Volume 35, Number 1, January–February, 2001 article, "AD9873 Mixed-Signal (MxFE) Front End for Broadband Digital Set-Top Boxes").

The recently introduced AD9875 and AD9876 MxFE devices, to be discussed in this article, were developed and aimed specifically at broadband home-networking and broadband access applications, both of which require high data rates or signal bandwidth up to 25 MHz and impose similar demands in functionality, performance, and cost. Given their flexibility, these parts can be used in modems for both home networking (HomePlug, HPNA) and high-speed data access (VDSL, power line). Developed for large volume, cost-sensitive, consumer-class applications, they offer exceptional value to providers of system solutions.

Figure 2 depicts the versatile role played by the MxFE chip. Its *receive* (Rx) circuitry accepts the (analog-domain) signals from the transmission medium, once they have been appropriately interfaced, and provides analog signal conditioning and A/D conversion to produce multiplexed digital signals that can be dealt with by the digital *physical layer* (PHY) and/or *media access controller* (MAC). It also accepts digital data from these entities, processes it and converts it to analog—and outputs this transmit (Tx) signal to the media interface. The device design is based on the objective of optimizing *system* performance—irrespective of the analog or digital nature of its I/O—as implemented in ADI's "smart-partitioning methodology" (see Appendix).

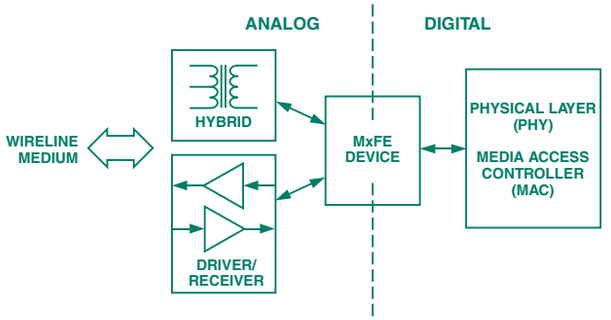

Figure 2. Typical wireline network node.

Figure 2 represents a typical wireline network node. The MxFE function may be connected to the wireline medium passively through a hybrid transformer block, or actively, with an amplifier on the receive side and a driver on the transmit side. On the digital side, the MxFE device needs to interface with the PHY and the MAC, which reside on a separate chip.

Critically important factors for the designer of MxFE circuitry for the home networking and high-speed access markets are performance, time to market, and low cost. In this type of emerging broadband communication application, discrete analog components can be prohibitively expensive, power-hungry, and demanding of board space. The combination of cost, performance, size, and power dissipation made possible by integrating the difficult mixed-signal, analog, digital, and signal processing functions on a single chip—such as the AD9875/AD9876—makes complex consumer-market communication products feasible. In fact, by implementing their intellectual property and design expertise in a digital gate array or ASIC, in conjunction with these devices, design engineers can develop new products and prototypes faster than ever to capitalize on rapidly changing markets.

Figure 3 shows a block diagram of the AD9875/AD9876 mixed-signal front end converter for broadband modems. On the analog side, the AD9875 and AD9876 combine the ADC, DAC, clock generation, programmable-gain amplification, and analog and digital filtering circuitry to provide a level of integration and performance usually implemented with discrete solutions costing significantly more. The ADC uses a pipelined multistage architecture to achieve high sample rates while consuming low power. In the AD9875, the *receive* path provides 9.5 ENOB @ 32 MSPS and 8.6 ENOB @ 50 MSPS. Comparable numbers for the AD9876 are 10.2 ENOB @ 32 MSPS and 9.3 ENOB @ 50 MSPS. The DACs in each device are, respectively, 10- and 12-bit interpolating TxDAC™ circuits.

On the digital side, the DAC inputs and ADC outputs are available on separate ports to accommodate both full-duplex and half-duplex operations. Each converter's port is multiplexed into high- and low nybbles to reduce the number of package pins. [A 10-bit half-duplex device, the AD9875-HD, available in summer 2002, will connect seamlessly to currently available HomePlug-compliant physical-layer digital ASICs that support multiplexed *transmit* and *receive* data ports.]

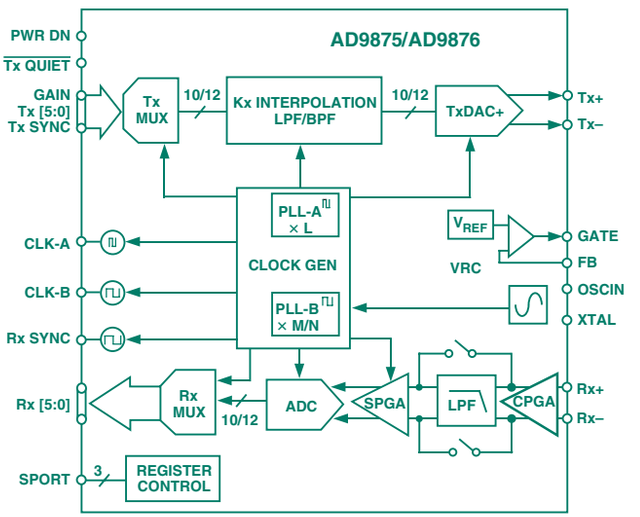

Figure 3. AD9875/AD9876 block diagram.

The availability of low-cost, flexible, high-performance, off-the-shelf mixed-signal front ends, such as these, with optimized functionality for broadband modem design using a variety of modulation formats—including OFDM—simplifies the ASIC design, specification, and testing process for both vendor and OEM and can substantially cut time to market. The unique features of the AD9875/AD9876 make these parts ideal for phone and power line networking as well as some xDSL applications.

Broadband wireline modems

Wideband signals with high peak-to-average ratios, commonly found in broadband modems—irrespective of the modulation scheme employed—put great demands on the system's analog- and mixed-signal processing components. Figure 2 and the AD9875/AD9876

block diagram in Figure 3 together show details of a generic wireline broadband modem. Not every modem requires every block shown, although most of them do employ the same or similar functions, while others may even require additional functional blocks like additional analog filtering. The block-level components of a wireline modem can be grouped into three main functions; a transmit path, a line coupler or hybrid, and a receive path. Most of the transmit and receive blocks are addressed directly by the AD9875/AD9876 component.

The main function of the *transmit* path is to send the signal onto the line with sufficient fidelity to reach the far end of the wire with a high enough *signal-to-noise ratio* (SNR) to allow faithful decoding at the receiver end. Invariably, the transmitter must also do this while conforming to a spectral mask to ensure that the modem does not cause excessive noise outside its channel bandwidth. Meeting the requirements of a *power spectral-density* (PSD) mask usually drives the design and performance requirements of the transmit path components. The two converter parameters that get the most attention are the number of bits and the sampling rate.

For the DAC, the number of bits required will depend on the desired SNR, the signal's peak-to-average ratio (PAR) and the ratio of the signal bandwidth to the sample rate. For a given SNR requirement, higher-precision converters are required as signal bandwidths and PARs increase. Often the most demanding DAC performance parameter turns out to be its *spurious-free dynamic range* (SFDR). The spurs generated by the non-ideal DAC transfer function may show up anywhere in the spectrum. If the magnitudes of the spurs are high and fall close to the signal band, they may be impossible to filter adequately. The DAC's SFDR performance must be able to meet the system linearity requirements.

The AD9875/AD9876 incorporates a 10-/12-bit DAC and makes use of interpolation filters to oversample the input data. Oversampling, because it moves DAC image frequencies away from the frequency of the desired signal, may result in substantially simpler external analog filter requirements, which translate into lower complexity and cost. Ideally, the driver will be able to provide the transmit-path gain and deliver the required output power while maintaining the DAC linearity performance. In cases where the DAC's available peak output power is not sufficient, AD832x cable drivers or AD8xxx DSL drivers could provide the interface to the wire. To deliver high peak-output signal with low distortion requires high voltage rails and high bias currents in the amplifier's output stage, which conflicts with the need for low power consumption and CMOS integration.

The method of coupling the modem's analog front-end to the line depends on whether the type of modulation/demodulation (modem) is *time-domain duplex* (TDD) or *frequency-domain duplex* (FDD). A TDD modem will normally employ a simple transformer coupling to the line and a switch that connects either the transmitter or receiver to the transformer. The main concern is that the switching and settling times, when the connections are made, meet the system requirements. An FDD modem will normally employ a *hybrid* to connect the modem's analog front-end to the line. A hybrid is required because the modem can be transmitting large signals while the receiver is listening to greatly attenuated signals—which can be orders of magnitude smaller. In order to limit the amount of signal coupling from the transmitter to the receiver, some type of line matching, cancellation circuitry, and filtering is used. The AD9875/AD9876 integrates a low-pass analog receive filter (LPF in Figure 3) with variable cutoff to provide for various signal bandwidth requirements.

The effectiveness with which the receiver maintains the SNR of the incoming signal within the *receive* channel bandwidth will be the most important determinant of the modem's raw data rate. The SNR of the receive signal over the signal bandwidth, determined by the channel, puts a fundamental limit on the amount of data the channel will carry. Any degradation beyond this in the circuitry between the line interface and the digital output samples is considered *implementation noise* and is determined by the quality of the receiver. It is the receiver's job to eliminate out-of-band channel noise and compensate for signal attenuation, then digitize the analog signal for further digital signal processing.

The out-of-band noise and interference on the line reduce the receiver SNR in two ways. First, the noise present may be folded back into the signal band of interest by the sampling process, raising the existing noise floor. Second, if the noise and interference are orders of magnitude greater than the desired signal, this will reduce the gain that can be applied to the signal to compensate for signal attenuation—again resulting in lower SNR. Effective filtering is essential to reduce these effects. In order to optimize noise and distortion performance, a variable-gain amplifier is implemented on-chip. This function is shared by a *continuous-time programmable gain stage* (CPGA in Figure 3) and a discrete-time *switched programmable gain stage* (SPGA), sandwiching the LPF. Following the SPGA, a 12-/10-bit ADC digitizes the signal with sampling rates up to 50 MHz.

Several auxiliary function are also present on-chip. Two *phase-locked-loop* (PLL) blocks, a voltage-regulator control circuit, and a *serial-port interface* help reduce external component count and optimize performance.

Phone-line networking

Figures 4 and 5a show the AD9875 in a phone-line networking application that is capable of data rates up to 32 Mbps using QAM modulation. The separation of mixed-signal and digital circuitry allows the digital ASIC to be implemented on the most cost effective geometry possible. Because the modulation coding is located in the digital ASIC, designers can maximize their 'value added' while they minimize time to market.

Power-line and VDSL modems

The AD9876 provides the optimum partitioning for a power-line or a VDSL modem, as illustrated in Figures 4 , 5b, and 5c. With its integrated programmable-gain amplifier, low-pass filter, and 12-bit ADC, in combination with 2×/4× interpolation filter and 12-bit TxDAC® D/A converter, the AD9876 provides nearly the entire analog portion of the signal chain-in a 48-lead LQFP package. By adding analog filters and line drivers, designers can bring a power-line or VDSL modem product to market quickly with an integrated single-chip solution. Compared with other, less-integrated, solutions in this type of application, component count can be reduced by as much as 50%; and the overall bill of materials can be reduced by more than the purchase price of the included integrated MxFE component. System costs can be reduced further by using the AD9875/AD9876's integrated 4× PLL clock multiplier and system clock outputs. They allow the entire system clock to be implemented with an inexpensive low-frequency crystal.

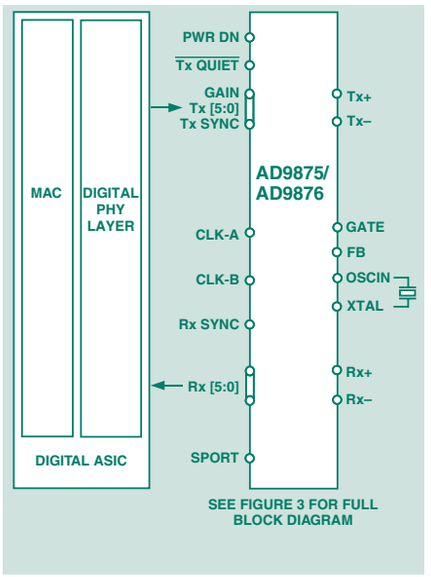

Figure 4. Mixed-signal and digital interface.

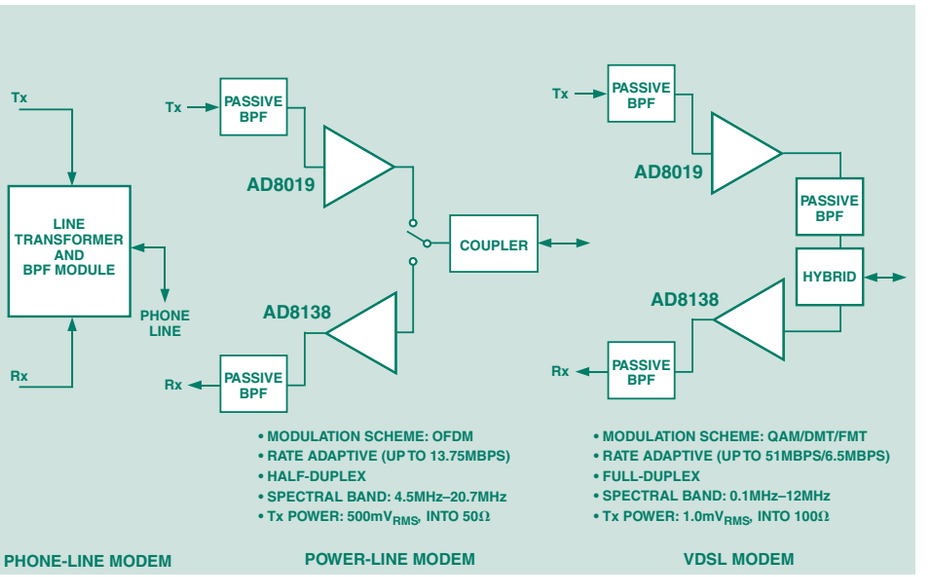

(a) (b) (c)

Figure 5. Analog interfaces.

AD9875/AD9876 Key Features, Specifications, and Performance

- Low-cost 3.3V-CMOS mixed-signal front-end converter for broadband networking/modems
- 10-/12-bit 128-MSPS TxDAC+® D/A converter
- 64-/32-MSPS input word rate
- *2×/4× interpolation transmit LPF or BPF transmit*
- Flexible power-down modes
- 10-/12-bit, 50-MSPS ADC
- 4th-order low-pass filter 12- or 26-MHz with bypass
- –6-dB to 36-dB programmable-gain amplifier
- *Internal 4× clock multiplier (PLL) clock outputs*
- Voltage regulator controller
- 48-lead LQFP package

AD9875/AD9876 performance was characterized over the –40° to +85°C extended industrial temperature range.

The following two graphs (Figures 6 and 7) show the transmit-path performance in multi-tone applications of these parts and illustrate their exceptional linearity. Figure 6 shows intermodulation distortion of the AD9876; under the indicated conditions, it is less than –80 dB. Figure 7 shows the multi-tone power ratio of the transmit path: about 55 dB when transmitting 70 tones in the 4.5-MHz to 20.7-MHz frequency range, corresponding to the HomePlug frequency band.

The AD9876 12-bit ADC's performance meets the requirements for two-band VDSL and has been designed into system solutions for that application and for power line access modems. Figure 8 is a plot of THD as a function of ADC sampling rate for a 5-MHz sinusoidal input.

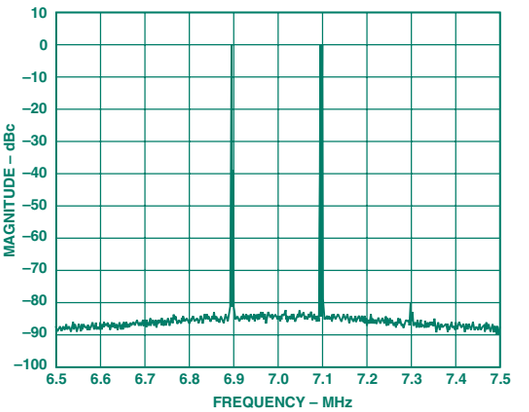

Figure 6. Dual-tone spectral plot of AD9876's 12-bit DAC @ f_{DATA} = 50 MSPS, f_{OUT} = 6.9 MHz, and 7.1 MHz.

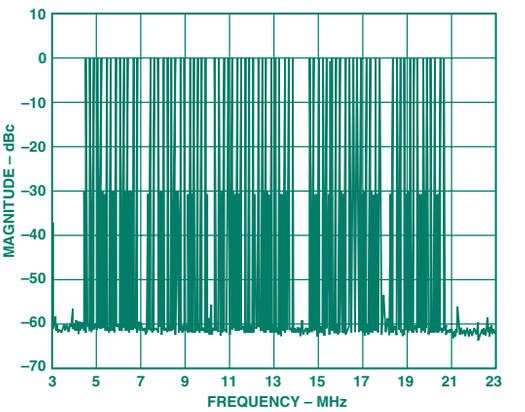

Figure 7. "In-band" multi-tone spectral plot of AD9876's 12-bit DAC @ f_{DATA} = 50 MSPS, f_{OUT} = k× 195 kHz, 2× LPF.

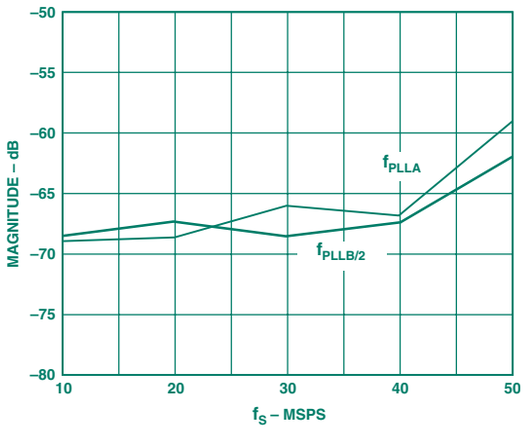

Figure 8. AD9876 12-bit ADC THD performance vs. f_{ADC} @ f_{IN} = 5 MHz.

The Rx LPF of the MxFE device has two frequency settings, one with a center frequency of about 10.8 MHz, the other at about 26 MHz. The filter transfer function is similar to that of a fourth-order Butterworth. The filters have a self-tuning feature that corrects for variation in device elements from part to part and for temperature drift. The center frequency can be adjusted over a 20% range if different cutoff frequencies are desired. Figure 9 shows LPF performance with the lower cutoff frequency selected.

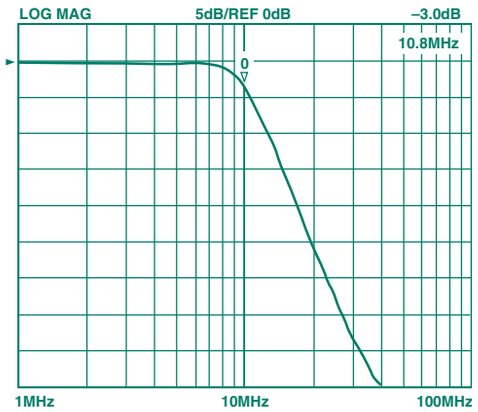

Figure 9. AD9876 Rx LPF frequency response, low f_C.

Evaluation Board and Software

The AD9875/AD9876 *evaluation board* (Figure 10), with its software, allows users to easily program and quickly evaluate the device for a specific modem application. The evaluation board provides connector access to the device's digital transmit and receive ports—which are buffered for reliable data transmission. The on-chip configuration registers can be programmed through the use of a PC serial port, which connects directly to a connector on the evaluation board; and the PC-resident evaluation-board software provides for a simple means of configuring the MxFE operation.

The analog output from the DAC, and the analog input to the ADC circuitry can be accessed through SMA connectors or pin headers on the evaluation board. The jumper-programmable interface configurations allow several different circuit options, to suit different testing methods. An on-board *digital loopback* feature allows the ADC and DAC to be tested simultaneously with the use of a signal source and spectrum analyzer. Also available is a *daughter card*, which can provide a phone-line compatible interface to the MxFE device.

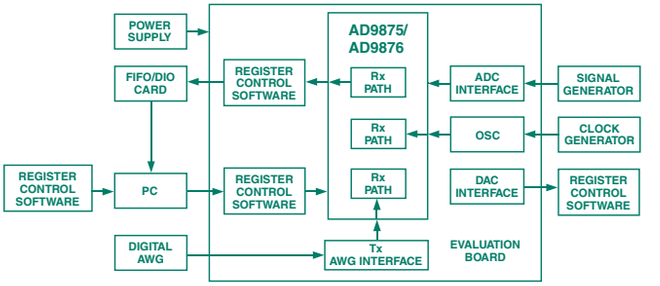

Figure 10. AD9875/AD9876 evaluation setup.

The AD9875-EB software provides a graphical user interface for easy programming and querying of the AD9875 registers. Three programming windows are available. The *direct-register-access* window allows AD9875 write- and readback in decimal, binary or hexadecimal data formats. The *register-map* window provides easy, function-oriented programming of the AD9875 registers. This window displays graphically all of the MxFE™ functions on the screen. The *advanced-register-access* window allows programming of sequences of register accesses.

Availability

The AD9875 and AD9876 were released to production in summer 2001. They are available in an economical, space-saving 48-lead LQFP and are priced at $9.24 and $14.45 (1000s), respectively. The AD9875 is priced at less than $5.00 (AD9876 <$7.00) in high volumes.

APPENDIX

These members of the Analog Devices MxFE family, the AD9875/AD9876, support *smart partitioning*, a mixed-signal design technique that partitions the signal path along lines that optimize system performance, rather than at analog/digital boundaries. This methodology complements the digital PHY (physical layer) and MAC (media access controller), which consist of several hundred thousand to well over a million gates fabricated on state-of-the-art 0.18-µm or finer geometry CMOS processes to take full advantage of the speed, power, and cost advantages that those processes offer. Meanwhile, at the heart of the AD9875/AD9876 are high-performance data-converter cores, fabricated on a well-established and cost-effective 0.35-µm CMOS process. To these cores are added both *analog and digital* signal-processing circuitry, to provide the necessary interfacing to both worlds, thus reducing complexity and cost of circuitry on the system board, without reducing flexibility. This combination of mixed-signal and digital functions provides overall system benefits by reducing component count and footprint without compromising performance.

Besides its performance and cost advantages, "smart partitioning" provides a logical place for the division of labor in developing complete physical-layer solutions. Companies whose core competencies are in the development of communications systems and algorithms can concentrate on those aspects of the digital design and capitalize on their differentiating technology. They can rest assured that their combined analog and mixed-signal solution is at the state of the art by partnering with a leader in the field, such as Analog Devices.

Interfacing a Blackfin DSP to High-Speed Converters for Wireless Applications

by Jeff Sondermeyer [jeff.sondermeyer@analog.com],
Jeritt Kent [jeritt.kent@analog.com],
Martin Kessler [martin.kessler@analog.com], and
Rick Gentile [richard.gentile@analog.com]

INTRODUCTION

In the 1970s and '80s high-speed mixed-signal designs were most frequently constrained by the limitations of digital circuitry, not analog. As an example, high-speed parallel converters (>10 MSPS) have been available from industry leaders like Analog Devices, Inc. (ADI), since the 1970s. Now, high-resolution data is being handled at higher sample rates (for example, 14 bits at >50 MSPS) by both analog-to-digital and digital-to-analog converters (ADCs and DACs). In addition, more and more applications are demanding intensive real-time algorithms. These factors mandate faster programmable general-purpose (GP) digital signal processors (DSPs) to handle the challenges presented by high-speed data rates.

Until recently, most designers had to interface high-speed parallel converters to application-specific ICs (ASICs) or fast *field-programmable gate arrays* (FPGAs). Devices like these are capable of resolving the many required simultaneous parallel digital operations; but they are often inflexible and can be prohibitively expensive. Now, with the recent introduction of Blackfin™ DSPs, such as the ADSP-21535, users have available a programmable general-purpose 16-bit fixed-point vector DSP—with a 300-MHz-capable core—that can handle the sustained input/output (I/O) and core throughputs required to process data from the many available high-speed converters. Depending on the core clock frequency, a maximum *system clock* (SCLK) of 133 MHz can be achieved. [This SCLK should not be confused with the serial clock for the *serial peripheral interface* (SPI).]

Why Choose a General-Purpose DSP?

GP DSPs typically cost much less than their closest digital-processing counterparts—FPGAs and ASICs—and they are easily programmed. In addition, since GP DSP design cycles are much shorter, time to market can be faster. With FPGAs/ASICs, users must often hire or consult professionals with specialized design skills. They may even be required to send their intellectual property (IP) out-of-house, incurring risks to the confidentiality of hardware, firmware, and software. On the other hand, GP DSP code can be stored in read-only memory (ROM) or masked into a DSP (such as members of the ADSP-2153x family), which further protects IP. Finally, GP DSPs are fully programmable, in contrast to ASIC implementations, where every change requires a costly redesign (time and money). These factors easily motivate many engineers to consider GP DSP as the solution of choice, especially when core rates can approach those of "Pentium®-class" chips.

The ADSP-21535, the first member of ADI's Blackfin family, was designed to work optimally in a computer-bus environment, while newer designs available soon (within the year) will have a parallel peripheral interface (PPI), specifically designed to work with I/O data. In the interim, however, the ADSP-21535's power can be made available for use in urgently needed designs, such as wireless applications, by using it with a small amount of readily available external circuitry.

What are the issues? In general, to guarantee sufficient data-processing bandwidth, the DSP needs a minimum clock speed an order of magnitude (10×) faster than the converter's sample rate. In turn, the amount of processing bandwidth needed depends upon the DSP's interface capabilities, which are in turn influenced by several other factors. These considerations include: block processing versus sample processing, the existence of a *direct memory-access* (DMA) controller, multi-ported memory, and whether external FIFOs are used. Fortunately, the ADSP-21535 has a full DMA controller that operates independently of the core, with multi-ported level-1 (L1) and level-2 (L2) memories. The combination of core speed, an independent DMA controller, and a large multi-ported on-board memory (308 Kbytes), allows the ADSP-21535 to perform efficient block processing at high data rates. For example, if the Revision 2.2-compliant, 33-MHz, 32-bit (4 bytes) peripheral component interconnect (PCI) interface is used (not shown in this application), transfer bandwidths can be achieved that approach 132 MB/s.

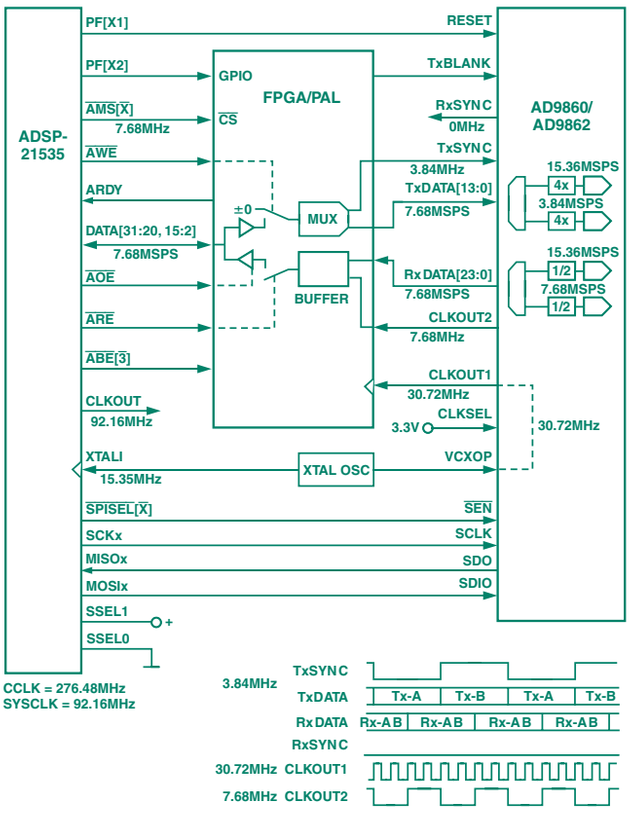

Figure 1. External logic connections between the ADSP-21535 and the AD9860/AD9862.

The ADSP-21535's *external bus interface unit* (EBIU) provides interfaces to asynchronous (ASYNC) external memories. If the PCI bus must be used for other system communications, the EBIU is the only available parallel interface to connect the ADSP-21535 to a high-speed converter. To combine the DSP-mastered, asynchronous control of this port with the synchronous, continuous data stream of converters may pose somewhat of a challenge for a system designer.

This article describes one particular hardware implementation, utilizing low-pin-count, low-cost, commonly available "glue logic" devices, such as a *programmable array logic* chip (PAL), a *complex programmable logic device* (CPLD), or an FPGA. This logic performs the control functions between the AD9860/AD9862 Mixed Signal Front End (MxFE™) and the ASYNC external memory bus of the ADSP-21535. The application depicted in Figure 1 is for an *orthogonal frequency-division multiplexed* (OFDM) wireless portable terminal. The ADC and DAC are time-shared (*time-division multiplexed*, or TDM) over the ASYNC interface of the DSP. (The information given here applies equally to other parallel high-speed ADCs and DACs.)

An Engineering Note* is available describing the details of the interconnection scheme. It assumes that the reader has information on hand about both the ADSP-21535 and the AD9860/AD9862, including the "ADSP-2153x/ADSP21535 Blackfin™ DSP Hardware Reference" and the data sheet for the AD9860/AD9862. These can be found at http://www.analog.com.

Design Goals

One of the early design goals for this project was to minimize the amount of external control logic necessary to interface the DSP and the converter(s). Driven by cost, Engineering wanted to eliminate any FIFOs or memory within the external logic device. An additional constraint was to avoid routing the data buses through the logic, thereby reducing the number of pins, package size, and cost of the logic device. The initial design shown in Figure 1 combines all functions (including data latching) into a single logic device. However, production models of this design will utilize inexpensive tri-state-able latches driven by a logic device. These latches or buffers will multiplex (pack) the samples from the DSP memory interface to the 12-/14-bit DAC as well as buffer or demultiplex (unpack) the 10-/12-bit ADC samples to the DSP memory interface.

Design Challenges

One of the key factors in any mixed-signal/DSP design is a solid understanding of the constraints and consequent trade-offs between the devices. The following discussion will illustrate the various trade-offs that must be considered when interfacing ADCs/DACs to the ADSP-21535.

Some of the major design constraints are:

- The OFDM modulation scheme for this design drove a required converter sample rate of 15.36 MSPS.
- The AD9860/AD9862 has a dual 10-/12-bit, 64-MSPS ADC and a dual 12-/14-bit, 128-MSPS DAC.
- Unlike SHARC® processors, which have a DMA request and DMA grant (i.e., DMA can be mastered from external devices), the ADSP-21535 has only one set of internal memory DMA channels (memDMA), which must be mastered from the DSP.

In addition, when the ADSP-21535 ASYNC interface is connected to devices that do not contain FIFOs or memory, all latencies must be thoroughly understood. For example, every time the memDMA relinquishes the bus after a burst of eight transfers, it requires ten SCLK cycles to begin the next transfer.

Future Blackfin family members will have programmable priority levels for the DMA controller, as well as a dedicated high-speed parallel interface—with DMA request and DMA grant signaling. With a dedicated PPI, these future Blackfin products will not require the ASYNC memory interface to connect with parallel converters.

The approach used here assumes that the memory interface is dedicated to the converters. Multiplexing external SRAM/SDRAM memory with the converter(s) would be difficult and is not recommended, especially considering that there is only one memDMA, and it would need to be shared. The existence of a large on-board L2 memory (256 Kbytes) minimizes the need for any external memory. However, it is permissible to multiplex the parallel converter(s) with a Flash or EPROM for the initial booting process.

This design uses a TDM time-slice approach for sharing the external bus between the ADCs and the DACs, because simultaneous access is not possible here, since the single memory interface does either a read or a write—and there is only one set of memDMA channels (source and destination).

The ADSP-21535 will support a maximum SCLK of 133 MHz (peak DMA bandwidth). At this rate, and with no external FIFO, the memDMA could sustain a transfer (32-bit word) rate of 133 MSPS/10 (nine cycles are required for bus acquisition and one for next transfer), or 13.3 M words/s. However, the SCLK of the ADSP-21535 is derived from the core clock (CCLK). CCLK in turn is generated via the PLL divider, whose available ratios are 1 to 31—and there are only four available divide ratios: 2.0, 2.5, 3.0, and 4.0. So one possible combination of CCLK and divisor that will allow a 133-MHz SCLK is CCLK = 266 MHz and CCLK/SCLK = 2. But if the core must run at 300 MHz, as in this application, the highest SCLK that can be obtained is 120 MHz (divisor = 2.5) to stay under the maximum 133 MHz.

Now, since the ASYNC memory interface is 32 bits wide, up to two 16-bit samples (in this case I and Q) can be packed into each word. This effectively halves the word rate that the DSP must process. (With a 15.36-MSPS converter sample rate, the DSP will "see" 7.68 MSPS.) The highest external converter sample rate that the memDMA will support under these conditions is $2 \times 120/10 = 24$ MSPS. Furthermore, the SCLK must be an integer multiple of the converter sample rate to ensure proper phase alignment between converter timing and DSP timing and eliminate the need for any external FIFOs. Therefore, the highest converter sample rate that the ADSP-21535 will support at a 300-MHz core rate is $2 \times 120/10$ M = 24 MSPS—or twice the memDMA rate, as discussed in Commandment #10 (at the end of this article). Since the DSP will only process the packed data at half this rate, 12 MSPS is the maximum rate that the memDMA can sustain, i.e., 12 M words/second. Higher sample rates can be processed by the ADSP-21535 if small external FIFOs are included between the converter(s) and the EBIU.

*"Interfacing the ADSP-21535 to High-Speed Converters (like those on the AD9860/AD9862) over the External Memory Bus." A PDF of this file and a ZIP file of software code for EE-162 can be found on the DSP Application Notes Page under EE-162.

<http://www.analog.com/library/applicationNotes/dsp/applicationNotes.html>.

Table 1. Possible parameter scenarios for the ADSP-21535.

Converter Sample Rate (MSPS)	CCLK (MHz)	CCLK/SCLK Divide Ratio	SCLK (MHz)	memDMA (Mwrites/s)	SCLK/Converter Sample Rate
15.360*	276.48	3.0	92.16	9.216	6
24.000	300	2.5	120	12.0	5
26.600	266	2.0	133*	13.3	5
20.000	300*	3.0	100	10.0	5
15.000	300*	4.0	75	7.5	5
10.000	300	2.5	120	12.0	12
8.8670	266	2.0	133*	13.3	15
0.8867	26.6	2.0	13.30	1.33	15*
CSR <= 26.6	CCLK < 300	2.0, 2.5, 3.0, or 4.0	SCLK < 133	memDMA < 13.3	5 <= integer <= 15

*Denotes driving parameter

Recall now that OFDM requirements dictated a 15.36-MSPS converter sample rate. To obtain a SCLK that is an integer multiple of this converter sample rate, one must choose a *phase-locked-loop* (PLL) multiplier that is an integer multiple of one of the four available divisor ratios (2.0, 2.5, 3.0, or 4.0). With a PLL multiplier of 18, the maximum CCLK allowed is 276.48 MHz. This in turn limits the SCLK to an integer multiple of 3, because 276.48/3 = 92.16 MHz. (A divide ratio of 2 would give an SCLK over the 133-MHz maximum.) Under these constraints, the maximum sustained rate that the memDMA can support is 92.16/10 = 9.21 M words/s.

DMA Considerations

Careful consideration must be given to the combined, required, "sustained" DMA performance. Since the memDMA is a shared resource over the DMA bus (DAB), other DMA activity is arbitrated on this bus. This application requires a 10-Mbit/s serial channel on a *serial port* (SPORT) that also must arbitrate for the DAB. This will consume an additional 625 K words/s at 16 bits/word of DMA bandwidth. The ADSP-21535 can support a maximum of 133 M words/s (peak) DMA bandwidth, and the SPORT has higher arbitration priority over the memDMA (see Table 1). So the SPORT DMA should effectively utilize the above-mentioned ten-cycle delay and allow most, if not all, of the 9.21 M words/s to be used by the memDMA. There are 9.21 M – 15.36 M/2 (= 1.53 M) words/second of additional bandwidth, which should provide enough margin to sustain a 7.68 MSPS rate.

Table 2. Arbitration Priority

DAB Master	Arbitration Priority
SPORT0 RCV DMA Controller	0—highest
SPORT1 RCV DMA Controller	1
SPORT0 XMT DMA Controller	2
SPORT1 XMT DMA Controller	3
USB DMA Controller	4
SPI0 DMA Controller	5
SPI1 DMA Controller	6
UART0 RCV Controller	7
UART1 RCV Controller	8
UART0 XMT Controller	9
UART1 XMT Controller	10
Memory DMA Controller	11—lowest

Analysis of the DMA engine within the ADSP-21535 reveals a few other considerations. While the DMA engine supports two types of DMA transfers—descriptor-based and autobuffer-based—the memDMA controller does not support autobuffer-based DMA. Therefore, descriptor-based transfers must be used. The descriptor fetch from L1/L2 memory involves two 5-word block moves, one for the source descriptor and another for the destination descriptor. In addition, the memDMA has a 16-entry 32-bit FIFO that is filled from the source and emptied from the destination. If both descriptors are loaded simultaneously, 39 SCLK cycles (worst case) are required from L2. The destination descriptor load has priority over the source load to avoid overrunning the FIFO. Thus, in this example, the amount of time required to load both descriptors simultaneously is (1/92.16 M) × 39 = 423 ns. The DMA engine descriptor load performs best when the descriptors are loaded from L2 memory. If the descriptors are located in L1 memory, there are additional delays. The worst-case source-plus-destination descriptor load time from L1 is 65 SCLK cycles. To process data effectively at these sample rates, ping-pong buffers are normally used. (In this design, two 1024-word buffers are utilized.) This technique allows data to be filled into one buffer while the core processes the other buffer. As a reference, the complete VisualDSP++™ 2.0 project program is available from ADI.

There are two phases of operation that must be analyzed: Samples must be received by the DSP from the ADC (receiver TDM phase); and samples must be transmitted from the DSP to the DAC (transmitter TDM phase).

Receiver TDM Phase

During the receiver phase, data moves thus: ADC➔EBIU➔(source)➔memDMA➔FIFO➔L1/L2 (destination). At the 15.36-MHz converter sample rate, a new 32-bit sample arrives at the DSP every 1/7.68 M = 130.2 ns. As seen from the descriptor load-time latency, 423 ns, something must be done to avoid overrunning the DSP and losing samples. Fortunately, the converters are attached to an external bus, and the address bus is not being used. Thus, when moving samples into the DSP, one can set up the source descriptor with the maximum transfer count, 65536 words, and destination descriptor with intended ping-pong buffer transfer size, 1024 words. In this way, upon receiving an interrupt from the core every 1024 words, only the destination descriptor is reloaded, and the load time is reduced to 20 SCLKs × 1/92.16 M = 217 ns. As noted, this design uses a TDM scheme in which the ADC and

DAC occupy individual time slices. The multiplex rate is a variable 5 ms to 8 ms. Since the ADC and DAC data are interleaved, at a worst case, the interface changes from receiver to transmitter and back every 8 ms. Therefore, 65536 words × 130.2 ns, or 8.5 ms, is sufficient time, and the source descriptor only needs to be set up once at the beginning of each receiver TDM phase. Finally, the 16-entry memDMA FIFO "hides" the destination descriptor load time, because the source is still filling the FIFO while the destination descriptor is being loaded from memory. In a worst-case scenario, the memDMA FIFO will accumulate only a few samples of data before the descriptor is reloaded. Then, these samples are burst into memory. So the need for an external FIFO on the receiver side is eliminated, and no samples are lost.

Transmitter TDM Phase

During the transmit phase (data to the DAC), data movement is in the opposite direction: L2/L1 (source)➔memDMA➔FIFO➔EBIU (destination)➔DAC. Unlike the receiver mode, the source descriptor must be updated every 1024 words. This will require 20 SCLK cycles, or 217 ns. However, since the memDMA (9.21 M words/s) is running slightly faster than the sample rate (7.68 M words/s), this should maintain 16 samples in the memDMA FIFO, which will feed the DAC while the descriptor loads. The destination descriptor transfer count can be fixed at 65536 words. Again, no external FIFO is required, and no samples are lost.

LOGIC OVERVIEW AND TIMING

In avoiding the need for FIFOs in the external logic, it is still important to synchronize the converter clocks to the DSP system clock, SCLK. This limits the available ADSP-21535 clocking options in accordance with the sample rate. Minimally, SCLK must be evenly divisible by the converter sample rate, and CCLK may need to be evenly divisible by the converter sample rate as well. (There is only one non-integer divisor, 2.5—it may not be usable in some cases). External latches or buffers must be used to align the data from the converters with the timing of the DSP (see Figures 2 and 3 for sample skew and delay). The four-wire DSP SPI port is directly connected to the AD9860/AD9862 SPI port. To ensure proper power sequencing and initialization, the DSP should reset the converter(s). To further reduce the pin count of the external logic, another option available on the AD9860/AD9862 (not shown here) allows two 10-/12-bit ADC values to be time-multiplexed onto a single 10-/12-bit RXDATA bus. While this would eliminate one of the two 10-/12-bit buses, it requires the external logic to demultiplex the data before it is transmitted to the DSP.

All data movement is controlled or mastered by the memDMA within the DSP. When the ADC data is read (see Figure 2), the external logic must drive the data and the ARDY signal. The external logic must sample the /AOE pin to check when data can be driven to the ADSP-21535. The /AOE signal indicates to the external logic that the DMA controller is ready to take data. The receiver three-state machine is shown at the bottom of the figure.

When data is being sent out to the DAC (see Figure 3), the external logic has to sample the /AWE signal and then drive ARDY. /AWE indicates to the external logic when the DMA controller is ready with new data. The transmitter four-state machine is shown at the bottom of the figure.

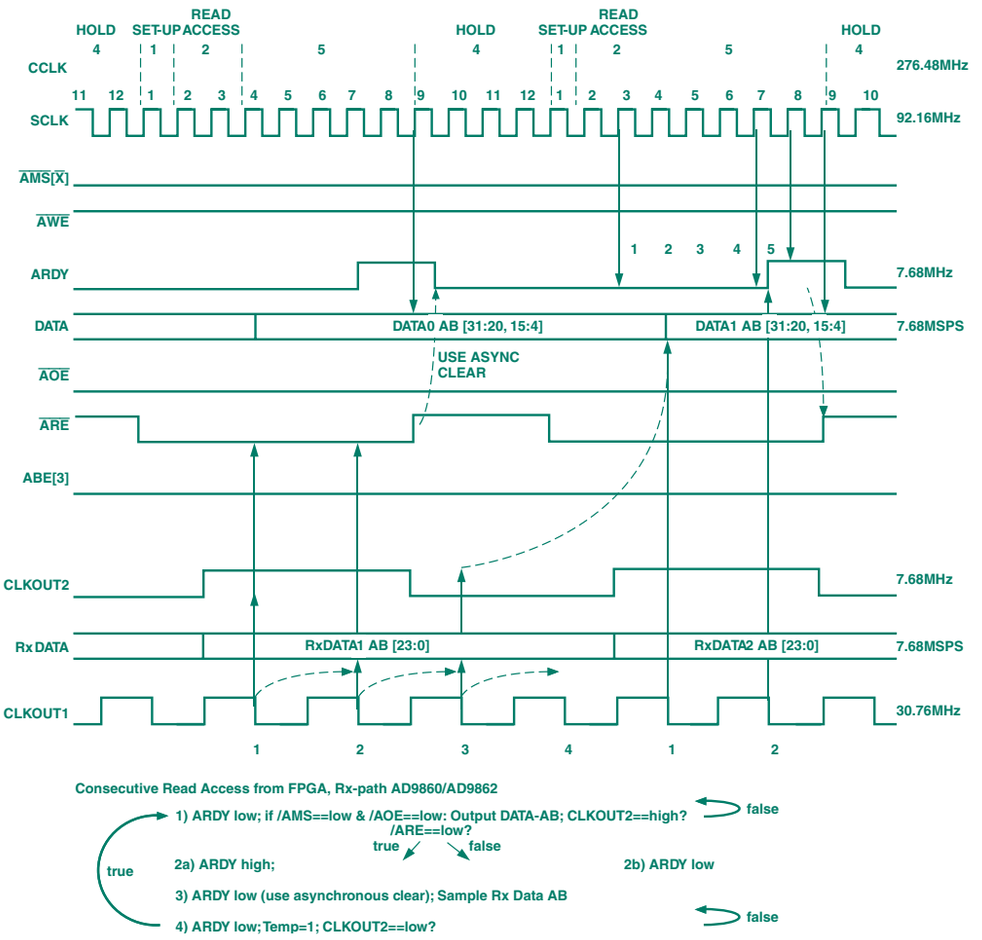

Figure 2. Receive timing and state machine.

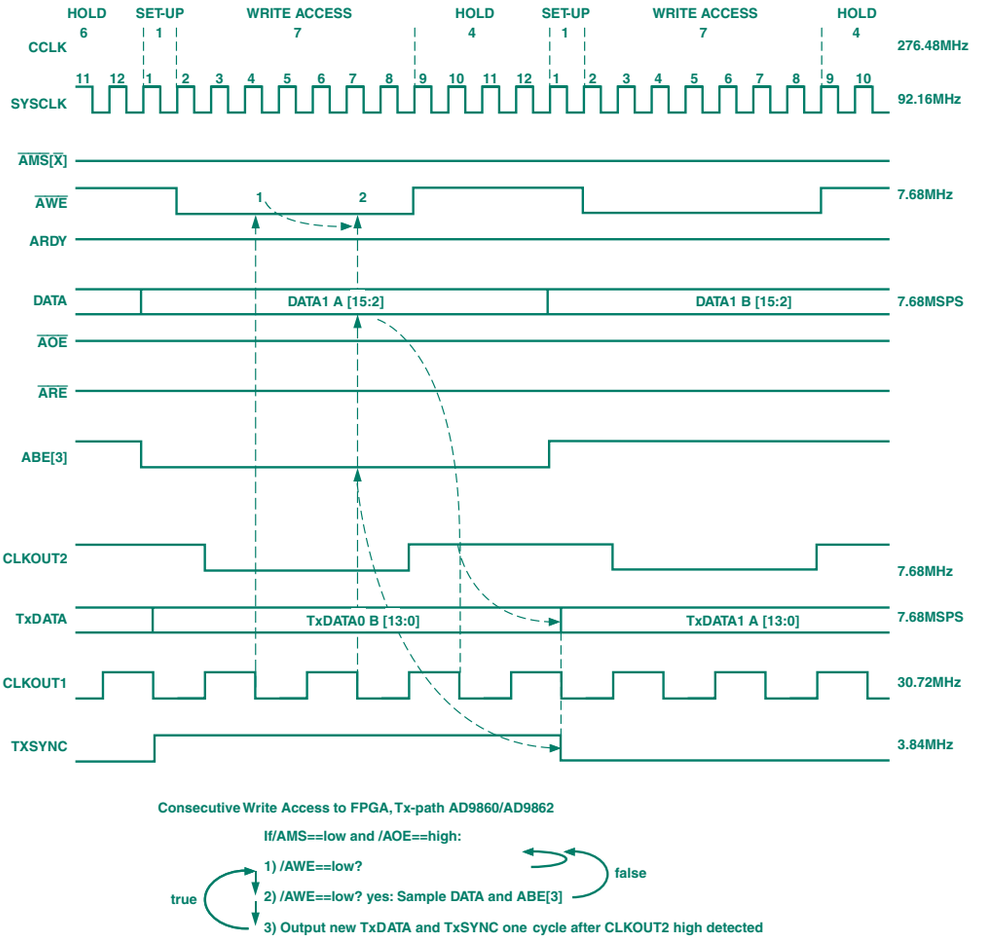

Figure 3. Transmit timing and state machine.

CONCLUSIONS

Even though the ADSP-21535 was not specifically designed to interface to high-speed parallel converters, it is available now, with all its many other advantages, for designs requiring rapid time-to-market. New-generation devices, such as the ADSP-21532, which has a dedicated parallel peripheral interface (PPI), will arrive soon to provide more-definitive solutions in such applications at lower cost. Until it is here-and available in production quantities-in order to fill an urgent need, we've suggested here a low-cost "FIFO-free" solution for interfacing today's ADSP-21535 to both ADCs and DACs with sample rates up to 24 MSPS-if the core is constrained to run at 300 MHz. If the ADSP-21535 core can be clocked specifically at 266 MHz, the highest converter sample rate is limited only by the maximum SCLK that the ADSP-21535 can support (133 MHz) and the inter-burst 10-cycle memDMA latency: 2 × 133 M/10 = 26.6 MHz.

The following set of interfacing rules, or "ten commandments," will help in optimizing performance of the ADSP-21535 when used with high-speed converters:

The ADSP-21535's Ten Commandments

1. The ADSP-21535's maximum allowed core frequency is 300 MHz.
2. ADSP-21535's maximum allowed system clock (SCLK) is 133 MHz.
3. ADSP-21535's memDMA has a worst-case 10-cycle reacquisition latency every time the DMA bus is relinquished.
4. Derive the ADSP-21535's SCLK from CCLK; there are four available divide ratios between CCLK and SCLK (2.0, 2.5, 3.0, and 4.0).
5. If the ADSP-21535's core runs at the maximum (300 MHz), the maximum SCLK is 120 MHz. See Commandments 1, 2, and 4.
6. The maximum ADSP-21535 memDMA rate is 133 M/10 = 13.3 M words/s. See Commandments 2 and 3.
7. In order to obtain the fastest ADSP-21535 SCLK (133 MHz) and memDMA transfer rate (13.3 M words/s), the core must not run at maximum, but at 266 MHz, with a CCLK/SCLK divide ratio of 2. This utilizes all but 34 MIPS.
8. When not using external FIFOs, the ADSP-21535's core should operate at a minimum rate an order of magnitude (10×) greater than the greatest external interfacing converter sample rate to provide sufficient processing bandwidth.
9. To eliminate the need for external FIFOs, the ADSP-21535's SCLK should be an integer multiple of the external interfacing converter sample rates to ensure proper phase alignment between converter timing and DSP timing.
10. To halve the sample rate that the DSP has to process, the external logic should pack up to two 16-bit samples into each 32-bit word. In fact, the maximum converter sample rate is twice the memDMA rate if two 16-bit words are packed into one 32-bit word.

NEW PRODUCT BRIEFS VOLUME 36, PART 1

Volume 36-01

16-bit, differential input, successive approximation **A/D Converters** **AD7675/76/77**
AC'97 SoundMAX® **Codec** **AD1887**
Variable gain **CATV Line Driver** **AD8327**
0.8-GHz-2.7-GHz direct-conversion **Quadrature Demodulator** **AD8347**
500-MHz, linear-in-dB VGA with **AGC detector** **AD8367**
Six-input **Analog Front-End plus DSP** **AD73460**
5-bit programmable **4-phase Synchronous Buck Controller** **ADP3164**
High-accuracy anyCAP® 1.5-A **Low Dropout Regulator** **ADP3339**
IMVP-II-compliant **Core Power Controller** for mobile CPUs ... **ADP3422**
High-frequency switch-mode **Li-Ion Battery Charger** **ADP3806**
Precision low-noise **Voltage References** **ADR423/25**
Remote/local **System Temperature Monitor** **ADM1032**
Low-resistance ±5-V **SPST/SPDT CMOS Switches** in small packages **ADG6xx**
CCD Signal Processors with integrated timing driver **AD9848/49**

Volume 36-02

16-channel, 1-MSPS, 12-bit **ADC with Sequencer** **AD7490**
16-bit, 570-ksps, low-cost **CMOS ADC** **AD7650**
16-bit sigma-delta ADC with PGA and switchable current sources **AD7709**
12-bit, 105-/125-MSPS **IF-sampling ADC** **AD9433**
3-V to 12-V **xDSL Line Driver** with power-down **AD8391**
Precision low-noise, low-bias-current ***FastFET***™ **Op Amp** **AD8610**
AC'97 SoundMAX® **Audio Codec** **AD1886A**
High-performance, single-chip stereo **Digital Audio Playback System** **AD1958**
6-GHz **PLL Frequency Synthesizer** **ADF4106**
12-bit mixed-signal **front-end for Set-top box and Cable Modem** **AD9877**
Serial Port Transceiver with 15-kV ESD protection and Green Idle™ **ADM3312E**
32-channel, 14-bit **DAC with infinite Sample-and-Hold mode** **AD5532B/33B**
1024-position **Digital Potentiometer with nonvolatile memory** **AD5231**
Active **Energy Metering IC** with di/dt sensor interface **ADE7759**
Ultralow I_Q, **anyCAP® Low Dropout Regulator** **ADP3342**
Dual bootstrapped **MOSFET Driver** **ADP3417**

Volume 36-03

24-bit read-only **sigma-delta ADC** **AD7782/83**
Dual 2-channel, 12-bit, 1-MSPS successive-approximation **ADC** **AD7866**
Single/dual low-power 350-MHz **voltage-feedback Op Amps** . **AD8038/39**
Single/dual low-cost 110-MHz **rail-to-rail Op Amps** **AD8091/92**
Precision gain-of-5 **Instrumentation Amplifier** **AD8225**
100 MHz to 2.7 GHz **RF Gain Block** **AD8353/54**
Dual precision **JFET-input Op Amp** **AD8512**
Quad precision **Op Amp** **OP4117**
96-kHz, 24-bit sigma-delta **Codecs** **AD1835/37/38/39**
Quad ultrahigh-speed **Pin Driver** **AD53513**
Dual high-speed **Pin Electronics IC: Driver, Comparator and Load** **AD53522**
4-channel, 104-MSPS **Digital Transmit Signal Processor** **AD6623**
4-channel, 100-MSPS **Digital Receive Signal Processor** **AD6624A**
160-dB **Logarithmic I-to-V Converter** **AD8304**
10-/12-/14-bit, 165-MSPS **TxDAC™ Family** **AD9740/42/44**
Serial Port Transceiver with 15-kV ESD protection and Green Idle™ **ADM3307E**
50-MBPS to 2.7-GBPS dual-loop **Laser Diode Driver** **ADN2841**
16-channel, 12-bit, **voltage-output DAC** with 14-bit increment mode **AD5516**
Mixed-signal **DashDSP™ with enhanced Analog Front-end** ... **ADMCF341**
16-bit fixed-point **DSP Microcomputer** **ADSP-2191M**
Precision **10.0-V/5.0-V References** **ADR01/02**
Multiplexers have 1-pC charge injection, 100-pA leakage, operate to 125°C **ADG604/36**
Dual **Interface for Flat Panel Displays** **AD9887**

MORE AUTHORS *(continued from page 2)*

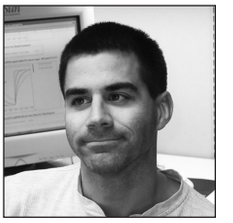

Martin Kessler (page 22) first encountered Analog Devices in 1990, as a coop student in Munich, Germany. He joined the company as a full-time employee in 1994, providing application support for all product lines. Subsequently, Martin moved to ADI Wilmington, where he focuses his efforts on digital-communication-specific ICs and mixed-signal front ends. Hobbies include skiing, hiking, backpacking, and playing the guitar.

Thomas Irrgang (page 6) is a Product Manager for SigmaDSP digital audio processors. Joining ADI in 1992, he has functioned in Central and Field Applications Engineering, and in consumer-market business development. In 2000 Thomas moved (from Munich) into the role of Technical Marketing for the digital audio group in Wilmington, MA. Earlier in his career, he had worked for professional audio companies Klotz Digital and AKG. Thomas received his MEE from Fachhochschule Muenchen in 1988. In his spare time, he enjoys his kids and likes playing drums.

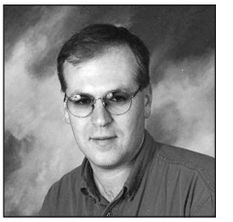

Iuri Mehr (page 17), a Design Engineer in ADI's High-Speed Converter Group, Wilmington, MA, designs mixed-signal circuits for communications. He received the MSEE degree from Washington State University in 1994, following BS and MS degrees in Applied Electronics from The Polytechnic Institute of Cluj Napoca, Cluj, Romania, in 1990. After a 3-year stint as Design Engineer for Crystal Semiconductor, he joined Analog Devices in 1996, in the Standard Products Group. He serves as an Associate Editor for *IEEE Transactions on Circuits and Systems.*

Eric Newman (page 3) is a Senior Applications Engineer in ADI's Advanced Linear Products Division. Earlier, he worked on adaptive receiver designs, and before that he worked as a board-level designer at Innovative Imaging Systems, Inc. In 1993 he received an MSEE from the University of Massachusetts at Lowell, with a focus on electromagnetics and communication systems.

Jeff Sondermeyer (page 22) is a Senior DSP Field Applications Engineer. He received his MSEE in Electrical Engineering from the University of Colorado in 1993. Before joining ADI in 2000, he worked in the professional audio field for ten years, designing award-winning audio-effects processors for Mackie Designs and Peavey Electronics. He has authored numerous articles and has a patent pending on "Automatic volume control in an audio system." He is an avid glacier/rock climber and guide in the Pacific Northwest, and he enjoys quality time with his wife Beth and daughter Rebecca.

Worldwide Headquarters
One Technology Way
P.O. Box 9106
Norwood, MA
02062-9106 U.S.A.
Tel: 781.329.4700,
(1.800.262.5643,
U.S.A. only)
Fax: 781.326.8703

Analog Devices Europe
c/o Analog Devices SA
17–19, rue Georges Besse
Parc de Haute
Technologie d'Antony
F-92182 Antony, France
Tel: 33.1.46.74.45.00
Fax: 33.1.46.74.45.01

Japan Headquarters
New Pier Takeshiba
South Tower Building
1-16-1 Kaigan,
Minato-ku, Tokyo
105-6891, Japan
Tel: 813.5402.8200
Fax: 813.5402.1063

Southeast Asia Headquarters
4501 Nat West Tower
Times Square
1 Matheson Street
Causeway Bay
Hong Kong, PRC
Tel: 852.2.506.9336
Fax: 852.2.506.4755

 AD02000-110-11/02(O)

ANALOG DEVICES

www.analog.com