Amplifiers Power Management Processors DSP MEMS Converters

Analog Dialogue

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

Volume 38, Number 3, 2004

JPEG 2000 Image Compression ... *see page 3*

In This Issue

ANALOG DEVICES

www.analog.com/analogdialogue

Editors' Notes

THE HARMONY OF ANALOG AND DIGITAL

It's interesting to observe that the three feature articles in this issue are highly representative of the currents that are basic to our businesses. Each illustrates a facet of what real-world signal processing (and Analog Devices) is all about—digital, analog, and the interface. They exemplify our long-held conviction that, although we live in an analog physical world, analog and digital must cooperate to solve each other's problems.

Take, for example, *JPEG 2000 Image Compression* (page 3). The JPEG 2000 standard defines a new image-coding scheme that uses state-of-the-art compression techniques based on wavelet technology. Its architecture is useful for many applications, including Internet image distribution, security systems, digital photography, and medical imaging. The article highlights some of its benefits.

The technology involves applying the highly sophisticated *wavelet* transformation in a *purely digital* interpretation scheme for encoding, transmitting, receiving, storing, and selectively using pictorial material. Yet it starts with electrical samples that depend on light intensity (an analog quantity), and is intended for ultimate display in some form by modifications of light from a source to produce pixels of light whose intensity (an analog quantity) correspond more-or-less faithfully to an intended relationship to the original.

Consider now, if you will, *Adjustable Cable Equalizer Combines Wideband Differential Receiver with Analog Switches* (page 13). *Category-5 unshielded twisted-pair cable*, like any transmission medium, suffers from dispersion and high-frequency signal loss. This article presents an equalizer design that compensates Cat-5 cable at frequencies to 100 MHz and lengths to 1000 feet, making it suitable for KVM networking and high-resolution video transmission.

The subject of this article is, quite evidently, preserving the integrity of *analog* signals. But what do the initials, KVM, stand for? *Keyboard, video, mouse*! What could be more digital? Again, we have a scheme for preserving information, but this time, the purpose of the design is to preserve *digital* information—subjected to the tender mercies of the analog world in the cable. Since KVM suggests *computer*, the source and destination of the information could both be totally digital in nature—starting with symbols and ending with symbols.

Finally, we have: *All About Direct Digital Synthesis*—Ask the Applications Engineer—33 (page 8). *Direct digital synthesis* (DDS) is a method of producing an analog waveform—usually a sine wave—by generating a time-varying signal in digital form and then performing a digital-to-analog conversion. Because operations within DDS devices are primarily digital, they can offer fast switching between output frequencies, fine frequency resolution, and operation over a broad spectrum of frequencies.

The operation speaks for itself! The properties of the waveform to be generated enter in purely symbolic form as *numerical* information (btw, the French word for "digital" is *numérique*). And lo! the device's DAC—in cahoots with the clock that (along with the power supply) is essential to the device's operation—emits an *analog* waveform of the appropriate frequency and phase.

Dan Sheingold [dan.sheingold@analog.com]

THE ANALOG WORLD—FICTIONAL AND REAL

Analog gets a bad rap, even in popular culture. In movies, TV shows, and magazines, people are told that analog is dirty and old fashioned, and that digital is clean and modern. In an ad from one of the electronics superstores, for example, a guy gets dumped for being "too analog". In *The Teeth of the Tiger* by Tom Clancy[1], readers are told that

> "The world was not digital, after all—it was an analog reality, always untidy, always with loose ends that could never be tied up neatly like shoelaces, and so it was possible to trip and fall with every incautious step." (page 172);
>
> "The world, one had to remember, was analog, not digital, in the way it operated. And *analog* actually meant *sloppy*." (page 286);

and

> "I think that we can depend on that." "Yeah, unless he got an unexpected phone call, or he saw something in the morning paper that caught his interest, or his favorite shirt wasn't properly pressed. Reality is analog, Sam, not digital, remember?" (page 316).

Clancy is right—the world *is* analog. But that doesn't make it dirty, unpredictable, or imprecise. While digital signals are limited by finite resolution, analog signals can have infinite resolution, limited only by noise or quantum effects. *Analog* signal processing can respond nearly instantaneously, without the computational delays inherent to digital signal processing. Analog circuitry can often operate at far lower power levels than digital circuitry providing the same function. Yet it is difficult to maintain the speedy, pristine nature of an analog signal through further signal processing for communications or storage, especially over long distances, in hostile environments, or over extended periods of time. Thus, the world needs precision data converters, high-speed operational amplifiers, power management components—and the expertise to use them.

The images in a digital camera are stored as 1s and 0s, but they are acquired by a CCD analog imager. Processing the CCD signal requires analog functions such as sampling, variable-gain amplification, and A/D conversion. Displaying the image on the liquid-crystal display requires analog functions, such as D/A conversion, filtering, and gamma correction. Many digital cameras include audio functions, and therefore require audio codecs and amplifiers to drive the speakers and microphones.

Wireless communication is also made possible by analog technology. Cellular carriers brag about their *all-digital* networks, but humans don't speak in 1s and 0s, and don't hear that way either. Voices must be digitized by A/D converters and reconstructed by D/A converters. And, while the data being transmitted is digital, the transmission medium is analog. 1s and 0s can't be transmitted as-is—they must first be modulated onto high-frequency carriers. On the *receive* side, weak signals must be captured by low-noise amplifiers—and demodulated. Power for all of these functions must be supplied by a small battery that lasts for weeks between charges, can be recharged quickly, and can be used while it is being charged. This requires complex analog power-management techniques.

Analog Devices, with its technologies, products, application notes, data sheets, application seminars, design tools, web site, field application engineers—and publications such as *Analog Dialogue*—seeks to help foster both the technology and expertise that designers—trained in either analog or digital—can use to cope with the realities of a mixed-signal world.

Scott Wayne [scott.wayne@analog.com]

[1]Tom Clancy, *The Teeth of the Tiger*. New York: G. P. Putnam's Sons, hardcover (2003).

Analog Dialogue

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

Analog Dialogue is the free technical magazine of Analog Devices, Inc., published continuously for 38 years—starting in 1967. It discusses products, applications, technology, and techniques for analog, digital, and mixed-signal processing. It is currently published in two editions—*online*, monthly at the above URL, and quarterly *in print*, as periodic retrospective collections of articles that have appeared online. In addition to technical articles, the online edition has timely announcements, linking to data sheets of newly released and pre-release products, and "Potpourri"—a universe of links to important and rapidly proliferating sources of relevant information and activity on the Analog Devices website and elsewhere. The *Analog Dialogue* site is, in effect, a "high-pass-filtered" point of entry to the **www.analog.com** site—the virtual world of *Analog Devices*. In addition to all its current information, the *Analog Dialogue* site has archives with all recent editions, starting from Volume 29, Number 2 (1995), plus three special anniversary issues, containing useful articles extracted from earlier editions, going all the way back to Volume 1, Number 1.

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

JPEG 2000 Image Compression

By Christine Bako [christine.bako@analog.com]

INTRODUCTION

The JPEG (Joint Photographic Experts Group) 2000 standard, finalized in 2001, defines a new image-coding scheme using state-of-the-art compression techniques based on wavelet technology. Its architecture is useful for many diverse applications, including Internet image distribution, security systems, digital photography, and medical imaging.

A lot of confusion exists as to what JPEG 2000 is and how it compares with other compression standards such as MPEG (Moving-Picture Experts Group) -2, MPEG-4, and the earlier JPEG. With brief comparisons to other compression standards, this article is primarily intended to highlight some of the often misunderstood and rarely mentioned potential-become-actual benefits of JPEG 2000.

Applications

CCTV Security

When transmitting or storing picture information, *compression* must be employed to maintain picture resolution while making best use of limited channel bandwidth. Compression is defined as *lossless* if full recovery of the original is available from the channel without any loss of information; otherwise, it is *lossy*. Standards are required to ensure interoperability. JPEG 2000 is the only standard compression scheme that provides for both lossless and lossy compression. As such, it lends itself to applications that require high-quality images despite limitations on storage or transmission bandwidths.

An important feature of systems based on JPEG 2000 is the ability to extract a variety of resolutions, components, areas of interest, and compression ratios from a single JPEG 2000 code stream. This is not possible with any other compression standard because the image size, bit rate, and quality must be specified on the encode side and can not be determined or changed on the decode side.

For example, a closed-circuit TV (CCTV) security system can make use of this feature by sending a single JPEG 2000 code stream over a low bandwidth network. High-resolution images can be stored on a hard-disk drive (HDD), while several lower-resolution images are displayed on monitors. The operator on the *receive* side can decide what information to extract from the single code stream sent.

JPEG 2000 is *frame accurate*, in that every single frame of the input is contained in the compressed format. MPEG systems, on the other hand, reduce the amount of data through *temporal compression* (which does not encode each frame as a complete image), so MPEG compression is not *frame-accurate*. For this reason, legal issues restrict the use of MPEG compression in some security applications. To get around this problem, security system and equipment providers have had to develop their own compression schemes—or use the highly inefficient *motion JPEG* (M-JPEG) compression standard—in order to provide a compressed stream that contains every single field of the original. They can now use JPEG 2000 for new designs.

Internet Image Distribution

Progressive coding, another feature of the JPEG 2000 standard, means that the bit stream can be coded in such a way as to contain less-detailed information at the beginning of the stream and more detailed information as the stream progresses. This makes it ideal for Internet/network applications—especially with large images and low bandwidths—as the image can be seen instantly on the decoding side, even with low-speed networks or image databases. The lower subbands are shown first, and more detail is added as time progresses. The picture thus becomes sharper and more detailed over time, and the entire image does not have to be downloaded before it can be seen.

With the low-quality image instantly available, the user at the receiving end can decide whether to view the picture in its fully decoded version, or to pass it by and scan the next picture instead. Clients can view images at different resolutions or quality levels [*compression rates*] making them suitable for any transmission bandwidth, connection speed, or display device. In addition, JPEG 2000 coding provides the option to zoom in or out on a particular area of the image—or to display a particular region of the image at a different resolution or compression rate.

High Definition

At extreme compression levels, JPEG 2000 video starts to blur, but is still quite viewable. MPEG or JPEG artifacts are much more disturbing to the eye, with the picture visibly broken down into small blocks at high compression ratios. The high image quality of medium-to-high bit rate signals containing a lot of motion, the lack of block artifacts, and high efficiency make JPEG 2000 ideal for high-definition (HD) applications, such as digital cinema, HD recording systems, and HD camera equipment.

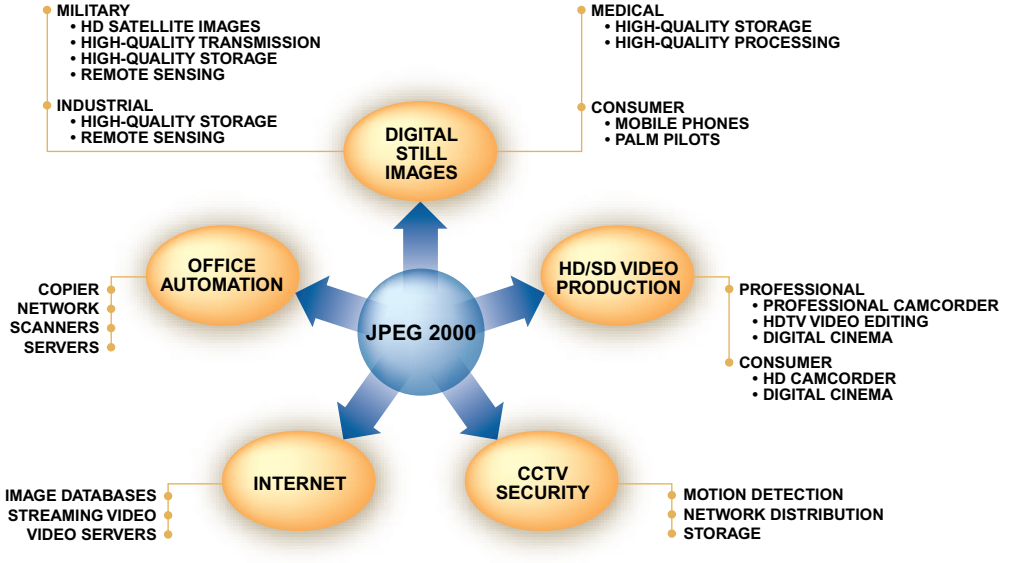

Figure 1. JPEG 2000 applications.

Many applications require exact bit-rate control, which only JPEG 2000 can provide. Exact bit-rate control is possible because an entire frame or field is transformed at once; it is then broken down into bit streams or code blocks that can be processed independently with the techniques described below. In systems using DCT, quantization is the only technique used, and this makes exact bit-rate control difficult. In order to control bit rate in DCT systems, the information must be repeatedly re-processed and re-quantized. The rate-control algorithm used in JPEG 2000 truncates each bit stream to meet a specific target bit rate, adjusting the truncation and re-quantization of each code block's data as required. In addition to programming the target bit rate, the standard allows the user to specify a particular quality metric. In this case, the target bit rate will vary to meet the specified quality factor, as long as the performance does not fall below a specific *peak signal-to-noise ratio.* The PSNR is an indication of picture quality comparable to *perceived picture quality.*

JPEG 2000 Code Stream

A given input image or part of the image [*tile*] is sent to a set of wavelet filters, which transform the pixel information into wavelet coefficients, which are then grouped into several *subbands* [the use of *wavelets* in encoding was first explained in Analog Dialogue 30-2 (1996)]. Each subband contains wavelet coefficients that describe a specific horizontal and vertical spatial frequency range of the entire original image. This means that lower-frequency, less-detailed information is contained in the first transform level, while more-detailed, higher-frequency information is contained in higher transform levels. For simplicity, only two levels of transform are shown here. The first transform level results in subbands LH1, HH1, HL1, and LL1. Only subband LL1 is passed on for further filtering, generating the next transform level and creating subbands LH2, HH2, HL2, and LL2.

Equally sized code blocks, which are essentially bit streams of data, are generated within each subband. This break-down is necessary for coefficient modeling and coding, and is done on a code-block-by-code-block basis. In essence, the actual compression is achieved by truncating and/or re-quantizing the bit streams contained in each code block. These bit streams are then optimally truncated using a technique knows as *post-compression-rate-control* (PCRC).

Code blocks can be accessed independently. Their bit streams are coded with three coding passes per bit plane. This process, called context modeling, is used to assign information about the importance of each individual coefficient bit. The code blocks can then be grouped according to their significance. On the decoding side it is then possible to extract information according to its significance, allowing the most significant information to be seen first.

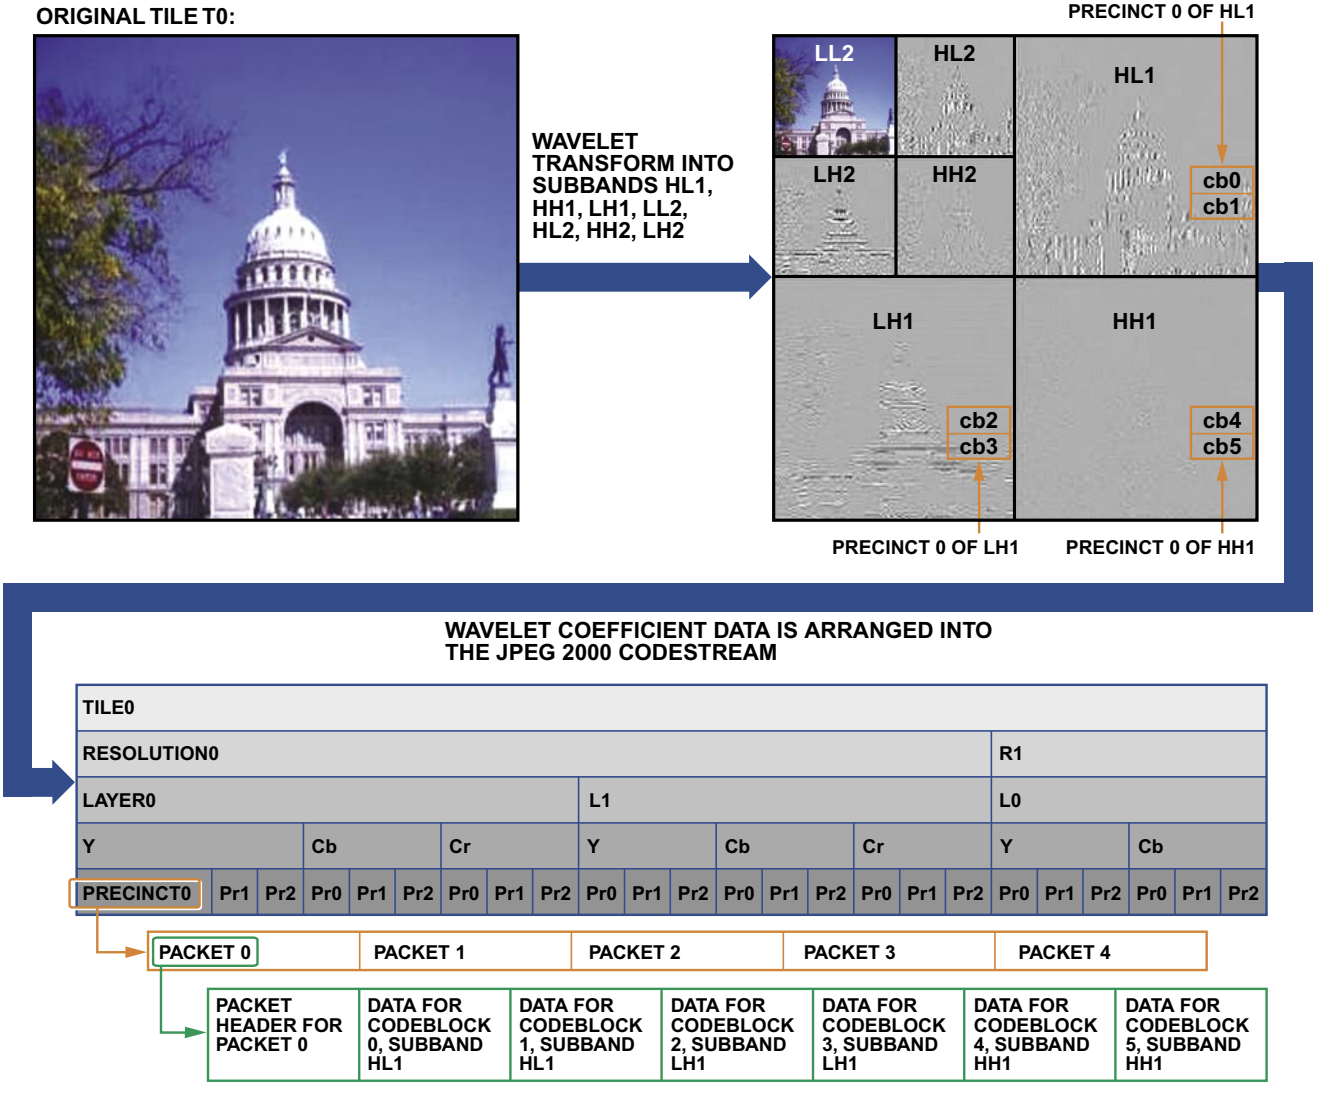

Figure 2. ENCODE—image over wavelet transform into subbands and resolutions.

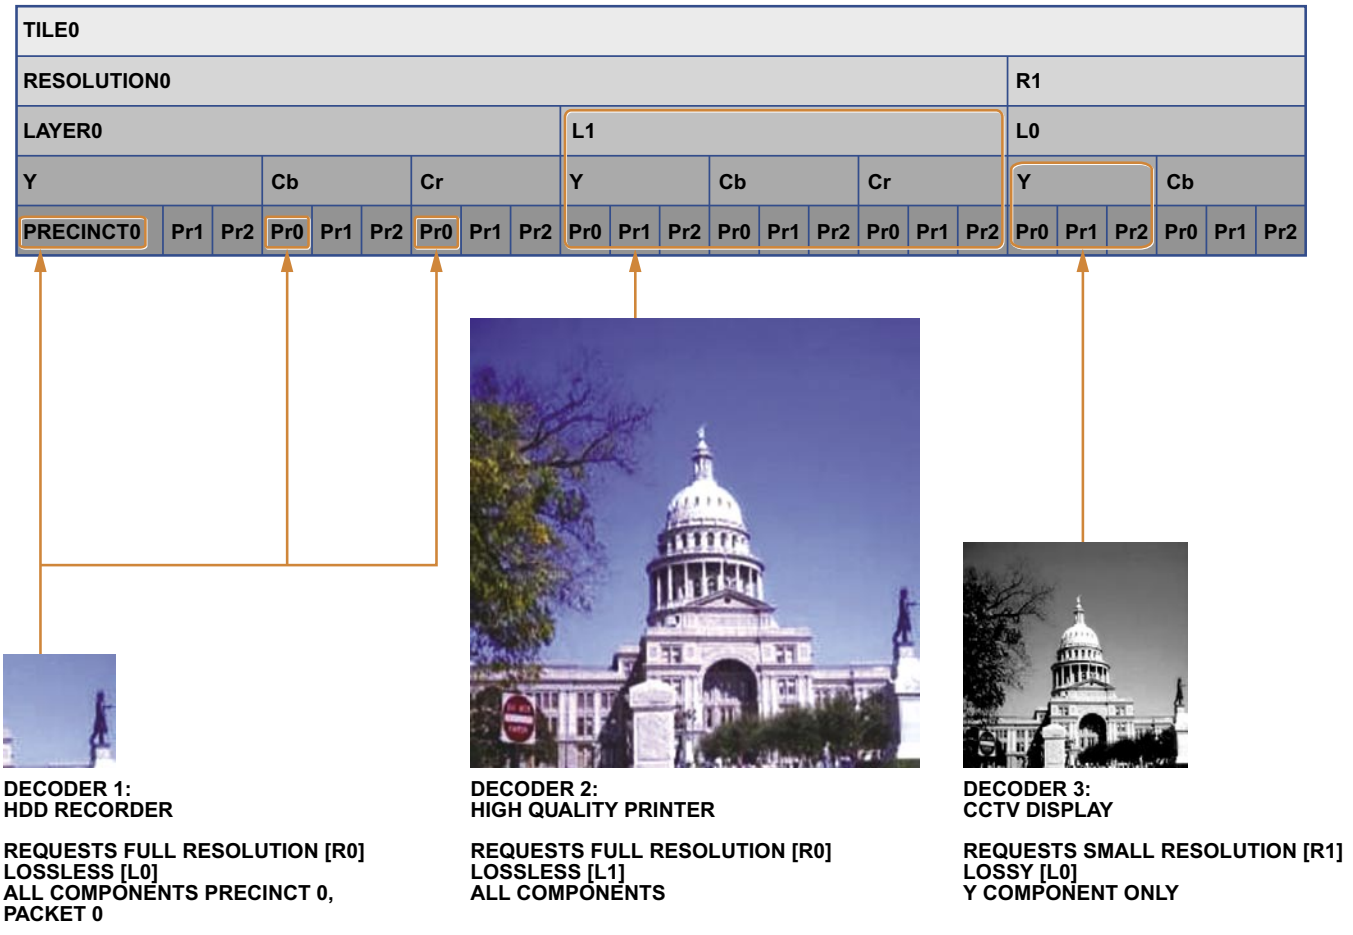

Figure 3. DECODE—one JPEG 2000 stream is received by several decoders.

JPEG 2000 can contain a user-defined number of layers, which are defined by PCRC and context modeling. Each layer stands for a particular compression rate, where the compression rate is achieved from the quantization-, rate-distortion-, and context modeling processes. Layer 0, for example, contains bit streams—from the lossy WT transform—that are heavily truncated, contain no coding passes, and thus provide the highest compression rate and the lowest quality. Layer 16 can then contain bit streams that are less truncated and use a higher number of coding passes, thus providing low compression and high quality.

Tiles or images are further partitioned into *precincts*. Precincts contain a number of code blocks, and are used to facilitate access to a specific area within an image in order to process this area in a different way, or to decode only a specific area of an image. The JPEG 2000 bit stream is generated by arranging code blocks or precincts into an array of packets with the lower subbands coming first.

The JPEG 2000 stream starts with a main header containing information such as: uncompressed image size, tile size, number of components, bit depth of components, coding style, transform levels, progression order, number of layers, code block size, wavelet filter type, quantization level, etc. The entire image data, grouped in code blocks of LL, HL, LH, and HH subbands, follows the header. Data is *not* contained in the header information. Also, a *table of contents* can be stored on the encode side, and allows a decoder to call up a certain resolution on demand, without first having to decode or download the entire JPEG 2000 code stream.

DCT versus WT

JPEG 2000 uses the *wavelet transform* (WT) to reduce the amount of information contained in a picture, while MPEG and JPEG systems use the *discrete cosine transform* (DCT). It is true that the WT requires more processing power than the DCT, but MPEG systems require more than just the DCT. The DCT, or any type of *Fourier* transform, expresses the signal in terms of frequency and amplitude—but only at a single instant in time. The WT transforms a signal into frequency and amplitude over time, and is therefore more efficient. The figures on the following page demonstrate this.

To obtain the same amount of information as with one WT pass, the DCT must be used for every frequency; and each of these frequencies must be transformed at each time instant for each 8×8 pixel block. In addition, MPEG systems use inter-frame compression [motion estimation] in order to reduce the amount of data further for motion estimation. This requires storage of at least two entire fields in external memory. The computation-intensive motion estimation process requires a very powerful processor. Temporal compression can be used in JPEG 2000 systems, but it is not inherent in the JPEG 2000 standard.

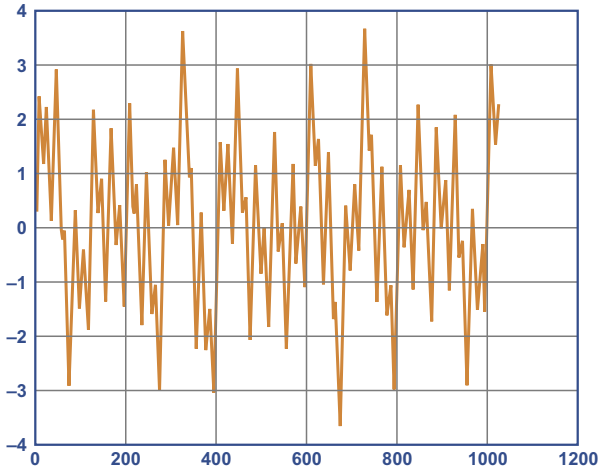

Figure 4. Input signal 1 containing frequencies A, B, C, and D.

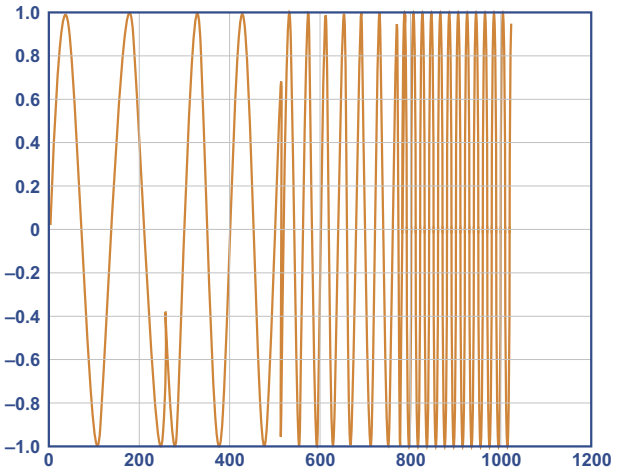

Figure 5. Input signal 2 containing frequencies A, B, C, and D.

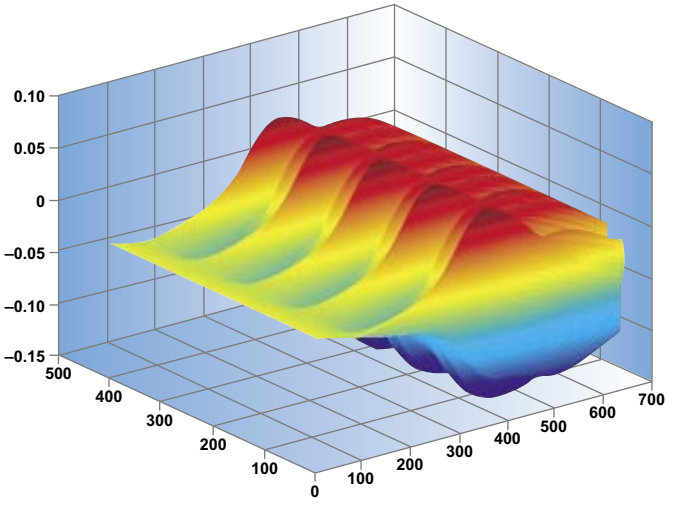

Figure 6. Wavelet transform of signal 1.

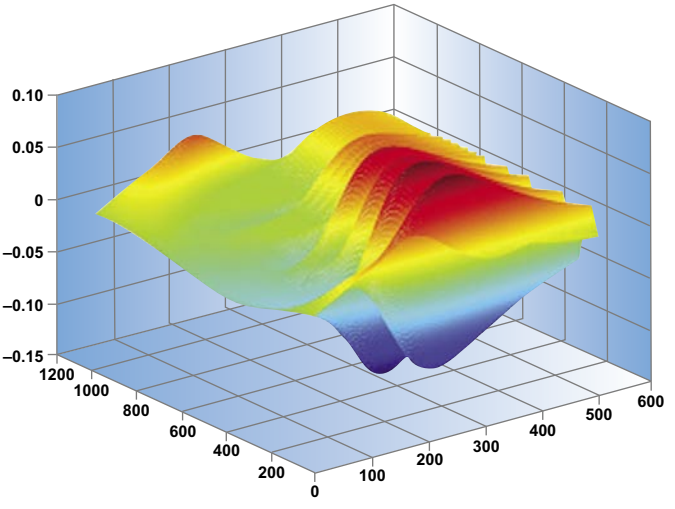

Figure 7. Wavelet transform of signal 2.

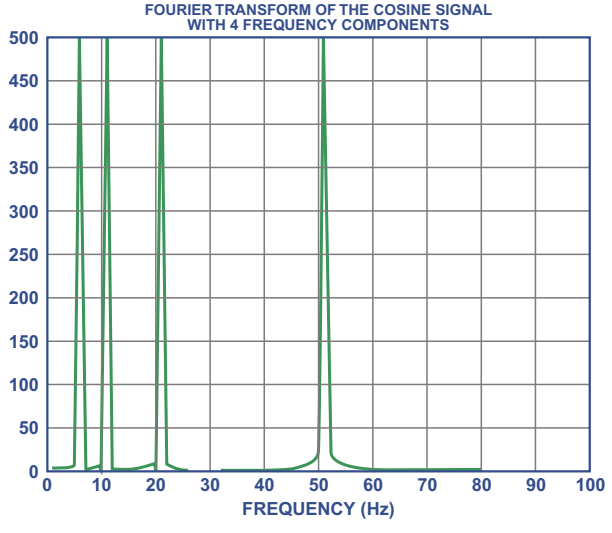

Figure 8. Fourier transform of signal 1.

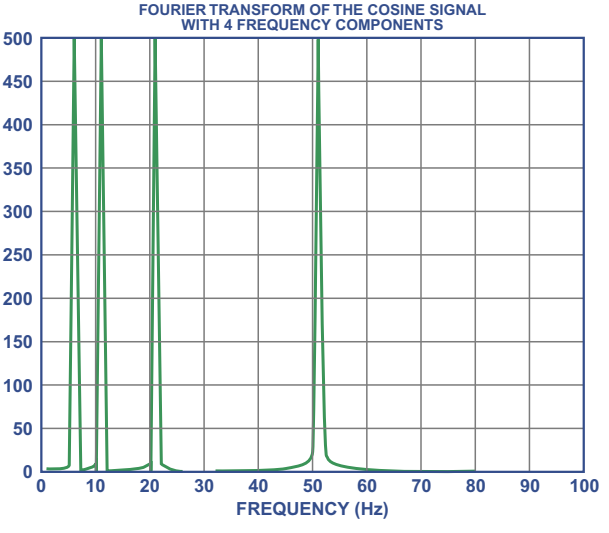

Figure 9. Fourier transform of signal 2.

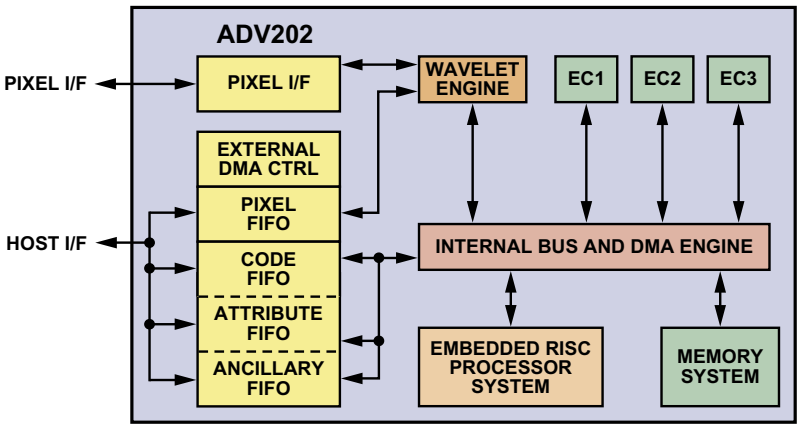

Figure 10. ADV202 block diagram.

JPEG 2000's Advantages Over Other Compression Standards

All MPEG standards are complex and computation intensive. This translates into extensive processing latency and memory requirements in *standard-definition* (SD) applications. These factors become even more of a problem when *high-definition* (HD) formats are considered, and JPEG 2000 becomes even more desirable. Another strength of JPEG 2000 is the standard itself, which allows immense flexibility and control in many different applications. There is also much versatility regarding formats: JPEG 2000 supports anything from 8-bits per sample to an unlimited amount of bits per sample, whereas MPEG only supports 8-bit data.

JPEG 2000 continues to gain popularity, even though MPEG-2 is the established standard for DVD and broadcast applications. JPEG 2000 is also very popular in HD applications that require high-quality storage or transmission of HD images over wireless or other links.

The ADV202

Since the early 1990s, Analog Devices has invested heavily in wavelet-compression R&D. We were the first to introduce a wavelet-compression hardware solution in 1996 with the ADV601. Now ADI's newest wavelet codec, the ADV202, released in July 2004, is so far the only dedicated JPEG 2000 IC on the market. A complete single-chip JPEG 2000 compression/decompression IC, the ADV202 works with high-definition video, standard-definition video, and still images. It supports all features of the ISO/IEC15444-1 [JPEG 2000] image-compression standard [except Maxshift ROI]. Its patented SURF™ (spatial ultra-efficient recursive filtering) technology enables low-power, low-cost wavelet-based compression. Containing a dedicated wavelet transform engine, three entropy codecs, a RISC processor, and on-board memory systems, the ADV202 provides a glueless interface to common video standards such as ITU.R.BT656, SMPTE274M, or SMPTE296M. It can create a fully compliant JPEG 2000 code stream [.j2c, .jp2]. It can also provide raw code-block and attribute data, allowing the host processor to have complete control over the generation- and compression processes.

Even though digital signal-processor (DSP) performance has improved significantly, a DSP would have to perform 20 billion instructions per second to match the performance of the ADV202 in a standard-definition encode application. Effectively serving as accelerators, the ADV202's three dedicated on-chip entropy codecs are responsible for the high throughput rate.

CONCLUSION—THE OUTLOOK FOR JPEG 2000

A major advantage of using a JPEG 2000 hardware solution is lower latency than any other compression scheme, a factor which is of particular importance in medical applications, for example.

Several major manufacturers of video or broadcast equipment have implemented JPEG 2000 into such future HD products as real-time encoding and decoding systems and video servers.

The Digital Cinema Initiative (DCI) has recently announced that it will use JPEG 2000 as the compression method in the delivery of digital motion pictures. The ADV202 has already found its way into many designs in the CCTV/security market in *video-over-network* applications. For more information, please visit http://www.itscj.ipsj.or.jp/sc29/29w02901.pdf

Because of its flexibility and image-compression quality, the ADV202—operating under JPEG 2000—could find its way into virtually every design that uses image or video compression.

This article can be found at http://www.analog.com/library/analogDialogue/archives/38-09/jpeg2000.html, with a link to a PDF.

Ask The Application Engineer—33

All About Direct Digital Synthesis

By Eva Murphy [eva.murphy@analog.com]
Colm Slattery [colm.slattery@analog.com]

What is Direct Digital Synthesis?

Direct digital synthesis (DDS) is a method of producing an analog waveform—usually a sine wave—by generating a time-varying signal in digital form and then performing a digital-to-analog conversion. Because operations within a DDS device are primarily digital, it can offer fast switching between output frequencies, fine frequency resolution, and operation over a broad spectrum of frequencies. With advances in design and process technology, today's DDS devices are very compact and draw little power.

Why would one use a direct digital synthesizer (DDS)? Aren't there other methods for easily generating frequencies?

The ability to accurately produce and control waveforms of various frequencies and profiles has become a key requirement common to a number of industries. Whether providing agile sources of low-phase-noise variable-frequencies with good spurious performance for communications, or simply generating a frequency stimulus in industrial or biomedical test equipment applications, convenience, compactness, and low cost are important design considerations.

Many possibilities for frequency generation are open to a designer, ranging from *phase-locked-loop* (PLL)-based techniques for very high-frequency synthesis, to dynamic programming of *digital-to-analog converter* (DAC) outputs to generate arbitrary waveforms at lower frequencies. But the DDS technique is rapidly gaining acceptance for solving frequency- (or waveform) generation requirements in both communications and industrial applications because single-chip IC devices can generate programmable analog, output waveforms simply and with high resolution and accuracy.

Furthermore, the continual improvements in both process technology and design have resulted in cost and power consumption levels that were previously unthinkably low. For example, the AD9833, a DDS-based programmable waveform generator (Figure 1), operating at 5.5 V with a 25-MHz clock, consumes a maximum power of 30 *milliwatts*.

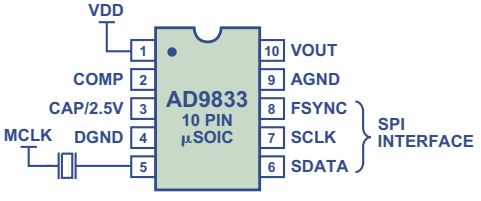

Figure 1. The AD9833—a one-chip waveform generator.

What are the main benefits of using a DDS?

DDS devices like the AD9833 are programmed through a high-speed *serial peripheral-interface* (SPI), and need only an external clock to generate simple sine waves. DDS devices are now available that can generate frequencies from less than 1 Hz up to 400 MHz (based on a 1-GHz clock). The benefits of their low power, low cost, and single small package, combined with their inherent excellent performance and the ability to digitally program (and re-program) the output waveform, make DDS devices an extremely attractive solution—preferable to less-flexible solutions comprising aggregations of discrete elements.

What kind of outputs can I generate with a typical DDS device?

DDS devices are not limited to purely sinusoidal outputs. Figure 2 shows the square-, triangular-, and sinusoidal outputs available from an AD9833.

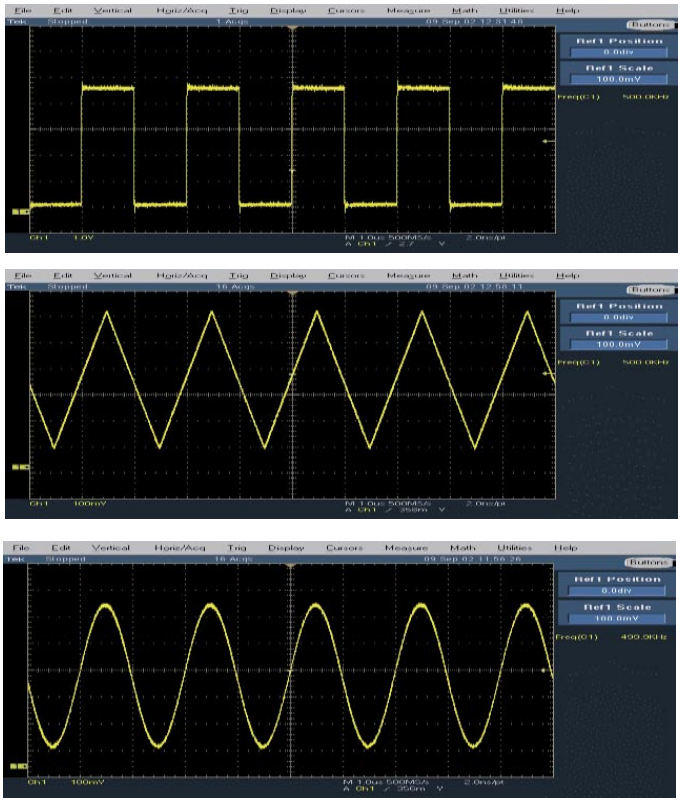

Figure 2. Square-, triangular-, and sinusoidal outputs from a DDS.

How does a DDS device create a sine wave?

Here's a breakdown of the internal circuitry of a DDS device: its main components are a *phase accumulator*, a means of *phase-to-amplitude conversion* (often a sine look-up table), and a DAC. These blocks are represented in Figure 3.

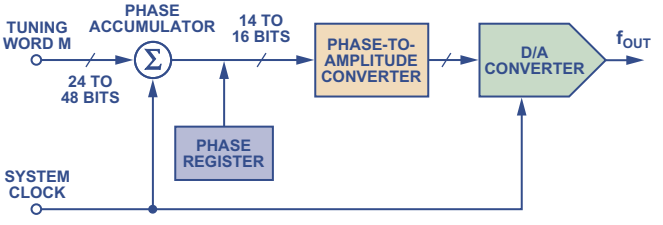

Figure 3. Components of a direct digital synthesizer.

A DDS produces a sine wave at a given frequency. The frequency depends on two variables, the *reference-clock* frequency and the binary number programmed into the frequency register (*tuning word*).

The binary number in the frequency register provides the main input to the phase accumulator. If a sine look-up table is used, the phase accumulator computes a phase (angle) address for the look-up table, which outputs the digital value of amplitude—corresponding to the sine of that phase angle—to the DAC. The DAC, in turn, converts that number to a corresponding value of analog voltage or current. To generate a fixed-frequency sine wave, a constant value (the phase increment—which is determined by the binary number) is added to the phase accumulator with each clock cycle. If the phase increment is large, the phase accumulator will step quickly through the sine look-up table and thus generate a high-frequency sine wave. If the phase increment is small, the phase accumulator will take many more steps, accordingly generating a slower waveform.

What do you mean by a *complete* DDS?

The integration of a D/A converter and a DDS onto a single chip is commonly known as a *complete* DDS solution, a property common to all DDS devices from ADI.

Let's talk some more about the phase accumulator. How does it work?

Continuous-time sinusoidal signals have a repetitive angular phase range of 0 to 2π. The digital implementation is no different. The counter's carry function allows the phase accumulator to act as a *phase wheel* in the DDS implementation.

To understand this basic function, visualize the sine-wave oscillation as a vector rotating around a phase circle (see Figure 4). Each designated point on the phase wheel corresponds to the equivalent point on a cycle of a sine wave. As the vector rotates around the wheel, visualize that the sine of the angle generates a corresponding output sine wave. One revolution of the vector around the phase wheel, at a constant speed, results in one complete cycle of the output sine wave. The phase accumulator provides the equally spaced angular values accompanying the vector's linear rotation around the phase wheel. The contents of the phase accumulator correspond to the points on the cycle of the output sine wave.

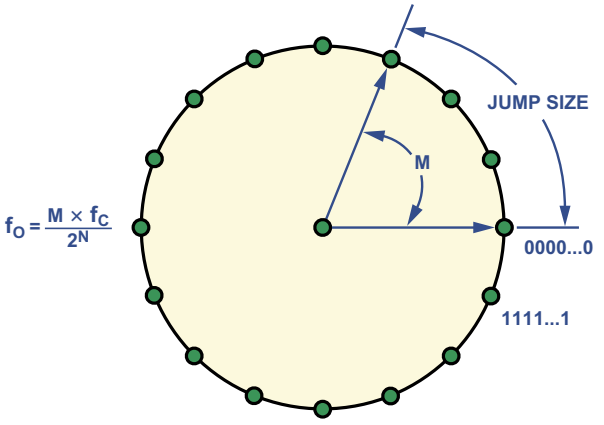

n	NUMBER OF POINTS
8	256
12	4096
16	65535
20	1048576
24	16777216
28	268435456
32	4294967296
48	281474976710656

Figure 4. Digital phase wheel.

The phase accumulator is actually a modulo-M counter that increments its stored number each time it receives a clock pulse. The magnitude of the increment is determined by the binary-coded input word (M). This word forms the phase step size between reference-clock updates; it effectively sets how many points to skip around the phase wheel. The larger the jump size, the faster the phase accumulator overflows and completes its equivalent of a sine-wave cycle. The number of discrete phase points contained in the *wheel* is determined by the resolution of the phase accumulator (n), which determines the tuning resolution of the DDS. For an n = 28-bit phase accumulator, an M value of 0000...0001 would result in the phase accumulator overflowing after 2^{28} reference-clock cycles (increments). If the M value is changed to 0111...1111, the phase accumulator will overflow after only 2 reference-clock cycles (the minimum required by Nyquist). This relationship is found in the basic tuning equation for DDS architecture:

$$f_{OUT} = \frac{M \times f_C}{2^n}$$

where:

f_{OUT} = output frequency of the DDS

M = binary tuning word

f_C = internal reference clock frequency (system clock)

n = length of the phase accumulator, in bits

Changes to the value of M result in immediate and *phase-continuous* changes in the output frequency. No loop settling time is incurred as in the case of a phase-locked loop.

As the output frequency is increased, the number of samples per cycle decreases. Since sampling theory dictates that at least two samples per cycle are required to reconstruct the output waveform, the maximum fundamental output frequency of a DDS is $f_C/2$. However, for practical applications, the output frequency is limited to somewhat less than that, improving the quality of the reconstructed waveform and permitting filtering on the output.

When generating a constant frequency, the output of the phase accumulator increases linearly, so the analog waveform it generates is inherently a ramp.

Then how is that linear output translated into a sine wave?

A phase-to-amplitude lookup table is used to convert the phase-accumulator's instantaneous output value (28 bits for AD9833)—with unneeded less-significant bits eliminated by truncation—into the sine-wave amplitude information that is presented to the (10-bit) D/A converter. The DDS architecture exploits the symmetrical nature of a sine wave and utilizes mapping logic to synthesize a complete sine wave from one-quarter-cycle of data from the phase accumulator. The phase-to- amplitude lookup table generates the remaining data by reading forward then back through the lookup table. This is shown pictorially in Figure 5.

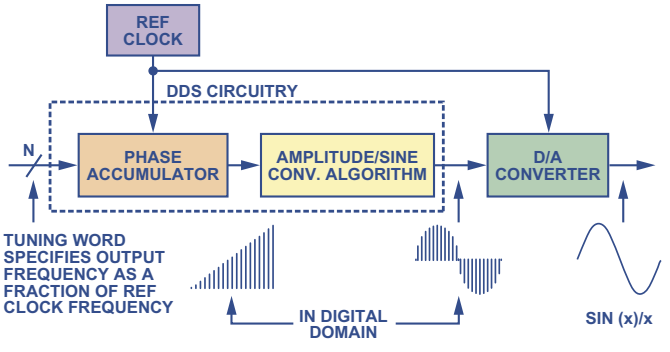

Figure 5. Signal flow through the DDS architecture.

What are popular uses for DDS?

Applications currently using DDS-based waveform generation fall into two principal categories: Designers of communications systems requiring agile (i.e., immediately responding) frequency sources with excellent phase noise and low spurious performance often choose DDS for its combination of spectral performance and frequency-tuning resolution. Such applications include using a DDS for modulation, as a reference for a PLL to enhance overall frequency tunability, as a local oscillator (LO), or even for direct RF transmission.

Alternatively, many industrial and biomedical applications use a DDS as a programmable waveform generator. Because a DDS is digitally programmable, the phase and frequency of a waveform can be easily adjusted without the need to change the external components that would normally need to be changed when using traditional analog-programmed waveform generators. DDS permits simple adjustments of frequency in real time to locate resonant frequencies or compensate for temperature drift. Such

applications include using a DDS in adjustable frequency sources to measure impedance (for example in an impedance-based sensor), to generate pulse-wave modulated signals for micro-actuation, or to examine attenuation in LANs or telephone cables.

What do you consider to be the key advantages of DDS to designers of real-world equipment and systems?

Today's cost-competitive, high-performance, functionally integrated DDS ICs are becoming common in both communication systems and sensor applications. The advantages that make them attractive to design engineers include:

- digitally controlled micro-hertz frequency-tuning and sub-degree phase-tuning capability,
- extremely fast *hopping speed* in tuning output frequency (or phase); phase-continuous frequency hops with no overshoot/undershoot or analog-related loop settling-time anomalies,
- the digital architecture of DDS eliminates the need for the manual tuning and tweaking related to component aging and temperature drift in analog synthesizer solutions, and
- the digital control interface of the DDS architecture facilitates an environment where systems can be remotely controlled and optimized with high resolution under processor control.

How would I use a DDS device for FSK encoding?

Binary *frequency-shift keying* (usually referred to simply as FSK) is one of the simplest forms of data encoding. The data is transmitted by shifting the frequency of a continuous carrier to one of two discrete frequencies (hence *binary*). One frequency, f_1, (perhaps the higher) is designated as the mark frequency (binary one) and the other, f_0, as the *space* frequency (binary zero). Figure 6 shows an example of the relationship between the mark-space data and the transmitted signal.

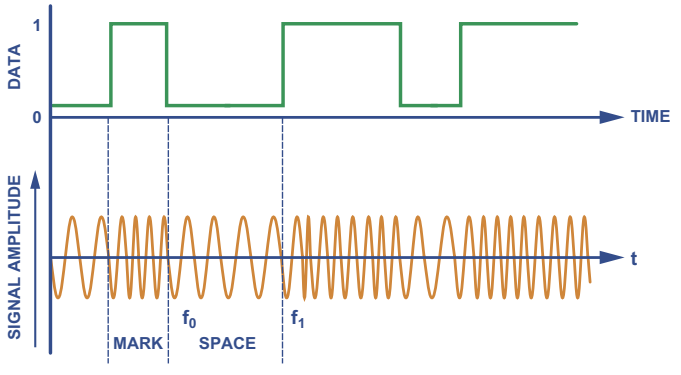

Figure 6. FSK modulation.

This encoding scheme is easily implemented using a DDS. The DDS frequency tuning word, representing the output frequencies, is set to the appropriate values to generate f_0 and f_1 as they occur in the pattern of 0s and 1s to be transmitted. The user programs the two required tuning words into the device before transmission. In the case of the AD9834, two frequency registers are available to facilitate convenient FSK encoding. A dedicated pin on the device (FSELECT) accepts the modulating signal and selects the appropriate tuning word (or frequency register). The block diagram in Figure 7 demonstrates a simple implementation of FSK encoding.

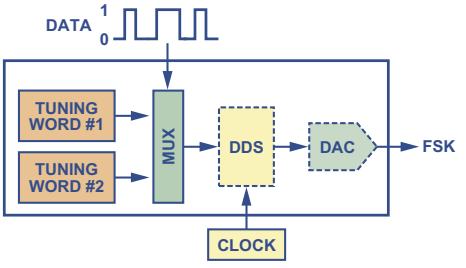

Figure 7. A DDS-based FSK encoder.

And how about PSK coding?

Phase-shift keying (PSK) is another simple form of data encoding. In PSK, the frequency of the carrier remains constant and the *phase* of the transmitted signal is varied to convey the information.

Of the schemes to accomplish PSK, the simplest-known as binary PSK (BPSK)—uses just two signal phases, 0 degrees and 180 degrees. BPSK encodes 0° phase shift for a logic 1 input and 180° phase shift for a logic 0 input. The state of each bit is determined according to the state of the preceding bit. If the phase of the wave does not change, the signal state stays the same (low or high). If the phase of the wave reverses (changes by 180 degrees), then the signal state changes (from low to high, or from high to low).

PSK encoding is easily implemented with DDS ICs. Most of the devices have a separate input register (a *phase register*) that can be loaded with a phase value. This value is directly added to the phase of the carrier without changing its frequency. Changing the contents of this register modulates the phase of the carrier, thus generating a PSK output signal. For applications that require high-speed modulation, the AD9834 allows the preloaded phase registers to be selected using a dedicated toggling input pin (PSELECT), which alternates between the registers and modulates the carrier as required.

More sophisticated forms of PSK employ four- or eight- wave phases. This allows binary data to be transmitted at a faster rate per phase change than is possible with BPSK modulation. In four-phase modulation (*quadrature* PSK or QPSK), the possible phase angles are 0, +90, –90, and 180 degrees; each phase shift can represent two signal elements. The AD9830, AD9831, AD9832, and AD9835 provide four phase registers to allow complex phase modulation schemes to be implemented by continuously updating different phase offsets to the registers.

Can multiple DDS devices be synchronized for, say, I-Q capability?

It is possible to use two single DDS devices that operate on the same master clock to output two signals whose phase relationship can then be directly controlled. In Figure 8, two AD9834s are programmed using one reference clock, with the same reset pin being used to update both parts. Using this setup, it is possible to do I-Q modulation.

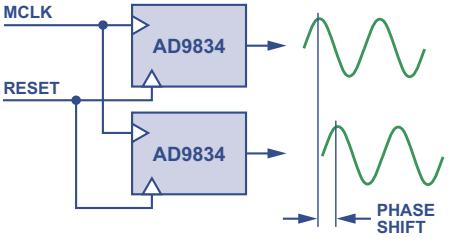

Figure 8. Multiple DDS ICs in synchronous mode.

A *reset* must be asserted after power-up and prior to transferring any data to the DDS. This sets the DDS output to a known phase, which serves as the common reference point that allows synchronization of multiple DDS devices. When new data is sent simultaneously to multiple DDS units, a coherent phase relationship can be maintained, and their relative phase offset can be predictably shifted by means of the phase-offset register. The AD9833 and AD9834 have 12 bits of phase resolution, with an effective resolution of 0.1 degree. [For further details on synchronizing multiple DDS units, please see Application Note AN-605.]

What are the key performance specs of a DDS based system?
Phase noise, jitter, and spurious-free dynamic range (SFDR).

Phase noise is a measure (dBc/Hz) of the short-term frequency instability of the oscillator. It is measured as the single-sideband noise resulting from changes in frequency (in decibels below the amplitude at the operating frequency of the oscillator using a 1-Hz bandwidth) at two or more frequency displacements from the operating frequency of the oscillator. This measurement has particular application to performance in the analog communications industry.

Do DDS devices have good phase noise?
Noise in a sampled system depends on many factors. Reference-clock jitter can be seen as phase noise on the fundamental signal in a DDS system; *and phase truncation* may introduce an error level into the system, depending on the code word chosen. For a ratio that can be exactly expressed by a truncated binary-coded word, there is no truncation error. For ratios requiring more bits than are available, the resulting phase noise truncation error results in spurs in a spectral plot. Their magnitudes and distribution depends on the code word chosen. The DAC also contributes to noise in the system. DAC quantization or linearity errors will result in both noise and harmonics. Figure 9 shows a phase noise plot for a typical DDS device—in this case an AD9834.

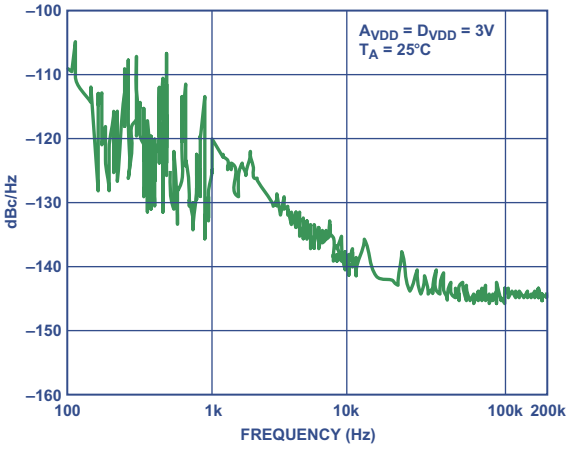

Figure 9. Typical output phase noise plot for the AD9834. Output frequency is 2 MHz and M clock is 50 MHz.

What about jitter?
Jitter is the dynamic displacement of digital signal edges from their long-term average positions, measured in degrees rms. A perfect oscillator would have rising and falling edges occurring at precisely regular moments in time and would never vary. This, of course, is impossible, as even the best oscillators are constructed from real components with sources of noise and other imperfections. A high-quality, low-phase-noise crystal oscillator will have jitter of less than 35 picoseconds (ps) of period jitter, accumulated over many millions of clock edges

Jitter in oscillators is caused by thermal noise, instabilities in the oscillator electronics, external interference through the power rails, ground, and even the output connections. Other influences include external magnetic or electric fields, such as RF interference from nearby transmitters, which can contribute jitter affecting the oscillator's output. Even a simple amplifier, inverter, or buffer will contribute jitter to a signal.

Thus, the output of a DDS device will add a certain amount of jitter. Since every clock will already have an intrinsic level of jitter, choosing an oscillator with low jitter is critical to begin with. Dividing down the frequency of a high-frequency clock is one way to reduce jitter. With frequency division, the same amount of jitter occurs within a longer period, reducing its percentage of system time.

In general, to reduce essential sources of jitter and avoid introducing additional sources, one should use a stable reference clock, avoid using signals and circuits that slew slowly, and use the highest feasible reference frequency to allow increased oversampling.

Spurious-Free Dynamic Range (SFDR) refers to the ratio (measured in decibels) between the highest level of the fundamental signal and the highest level of any spurious signal—including aliases and harmonically related frequency components—in the spectrum. For the very best SFDR, it is essential to begin with a high-quality oscillator.

SFDR is an important specification in an application where the frequency spectrum is being shared with other communication channels and applications. If a transmitter's output sends spurious signals into other frequency bands, they can corrupt or interrupt neighboring signals.

Typical output plots taken from an AD9834 (10-bit DDS) with a 50-MHz master clock are shown in Figure 10. In (a), the output frequency is exactly 1/3 of the master clock frequency (MCLK). Because of the judicious choice of frequencies, there are no harmonic frequencies in the 25-MHz window, aliases are minimized, and the spurious behavior appears excellent, with all spurs at least 80 dB below the signal (SFDR = 80 dB). The lower frequency setting in (b) has more points to shape the waveform (but not enough for a really clean waveform), and gives a more realistic picture; the largest spur, at the second-harmonic frequency, is about 50 dB below the signal (SFDR = 50 dB).

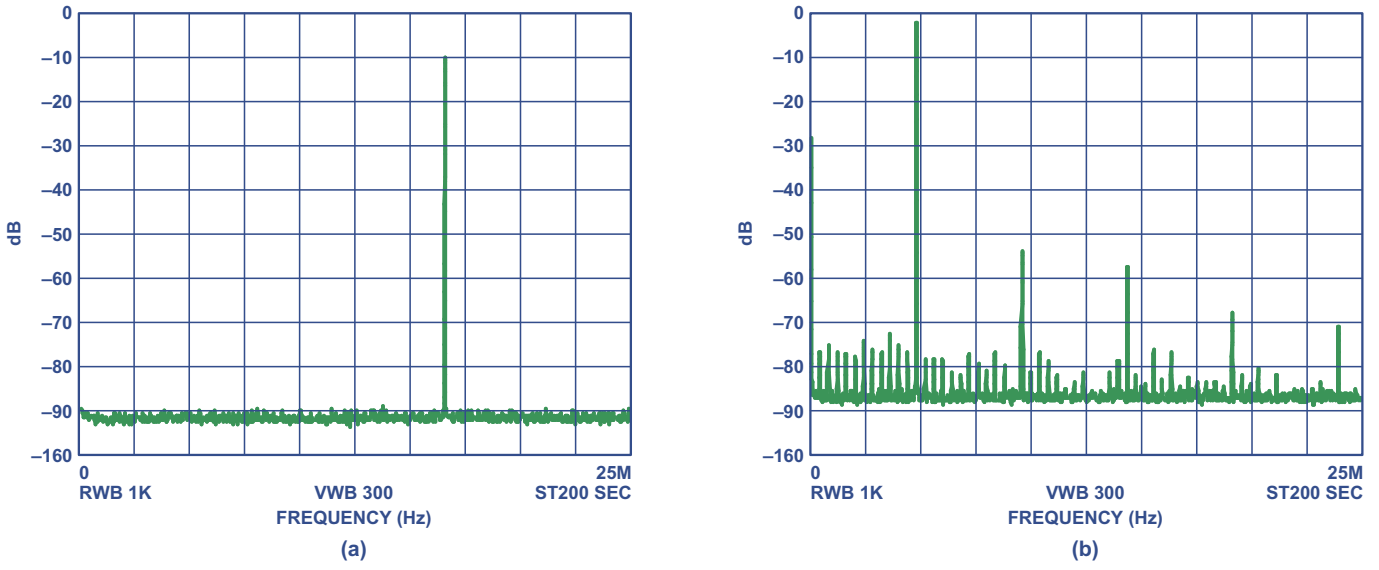

Figure 10. Output of an AD9834 with a 50-MHz master clock and (a) f_{OUT} = 16.667 MHz (i.e., MCLK/3); (b) f_{OUT} = 4.8 MHz.

Do you have tools that make it easier to program and predict the performance of the DDS?

The on-line *interactive design tool* is an assistant for selecting tuning words, given a reference clock and desired output frequencies and/or phases. The required frequency is chosen, and idealized output harmonics are shown after an external reconstruction filter has been applied. An example is shown in Figure 11. Tabular data is also provided for the major images and harmonics.

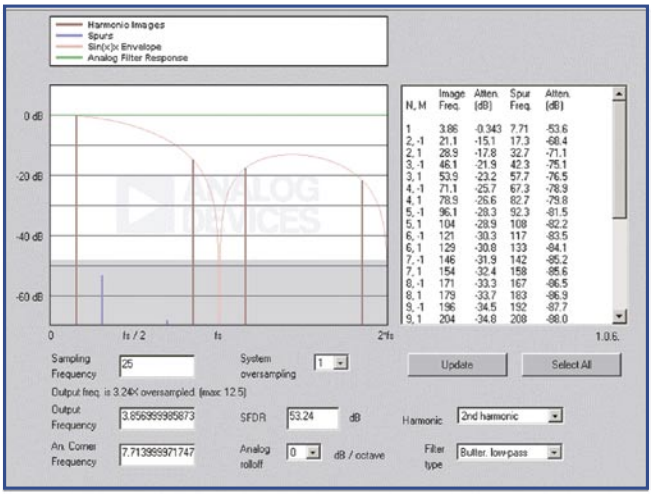

Figure 11. Screen presentation provided by an interactive design tool. A sin*x*/*x* presentation of a typical device output.

How will these tools help me program the DDS?

All that's needed is the required frequency output and the system's reference clock frequency. The design tool will output the full programming sequence required to program the part. In the example in Figure 12, the MCLK is set to 25 MHz and the desired output frequency is set to 10 MHz. Once the update button is pressed, the full programming sequence to program the part is contained in the Init Sequence register.

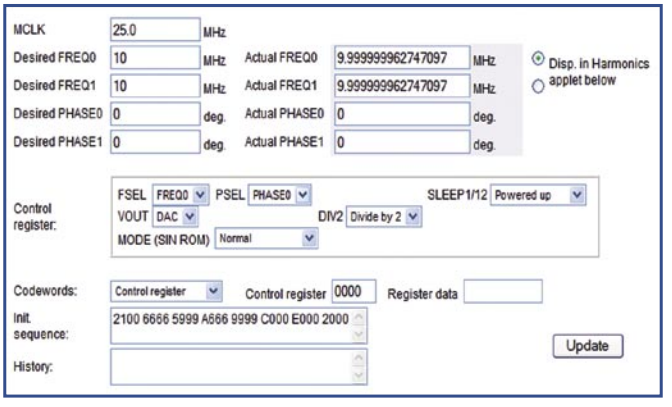

Figure 12. Typical display of programming sequence.

How can I evaluate your DDS devices?

All DDS devices have an evaluation board available for purchase. They come with dedicated software, allowing the user to test/evaluate the part easily within minutes of receiving the board. A technical note accompanying each evaluation board contains schematic information and shows best recommended board-design and layout practice.

Where can I find more information on DDS devices?

The main DDS homepage is located at www.analog.com/dds

Links to design tools are provided at http://www.analog.com/Analog_Root/static/techSupport/interactiveTools/#dds

An in-depth tutorial on DDS technology can be found at http://www.analog.com/UploadedFiles/Tutorials/450968421DDS_Tutorial_rev12-2-99.pdf

AN-605 can be found at http://www.analog.com/UploadedFiles/Application_Notes/371092853519044414816 8447035AN605_0.pdf

The latest DDS selection guide can be found at http://www.analog.com/IST/SelectionTable/?selection_table_id=27

This article can be found at http://www.analog.com/library/analogDialogue/archives/38-08/dds.html, with a link to a PDF.

Adjustable Cable Equalizer Combines Wideband Differential Receiver with Analog Switches

By Jonathan Pearson [jonathan.pearson@analog.com]

Originally intended to carry LAN traffic, *category-5* (Cat-5) *unshielded twisted-pair* (UTP) cable has become an economical solution in many other signal-transmission applications, owing to its respectable performance and low cost. For instance, an application that has become popular is *keyboard-video-mouse* (KVM) networking, in which three of the four twisted pairs carry the *red, green, and blue* (RGB) video signals.

Like any transmission medium, Cat-5 imposes transmission losses on the signals it carries, manifested as signal dispersion and loss of high-frequency content. Unless something is done to compensate for these losses, they can render the cables useless for transmitting high-resolution video signals over reasonable distances. Presented here is a practical technique to compensate for Cat-5 losses by introducing an *equalizer* (EQ), with eleven (11) switchable cable-range settings, at the receiving end of the cable. Because each setting of the EQ provides the proper amount of frequency-dependent gain to make up for the cable losses, the EQ-cable combination becomes suitable for high-resolution video transmission.

The first step in the EQ design is to derive a model for the Cat-5 frequency response. It is well known that the frequency response of metallic cable follows a low-pass characteristic, with an exponential roll-off that depends on the square root of frequency. Figure 1 depicts this relationship for lengths of Cat-5 from 100 feet (30.48 m) through 1000 feet (304.8 m), in 100-foot increments. In this illustration, it should be evident that the *power* loss at a given frequency is characterized by a constant attenuation rate (expressed in dB/ft).

Table I presents the Cat-5 equivalent *voltage*-attenuation magnitudes as a function of frequency for the same cable lengths as shown in Figure 1.

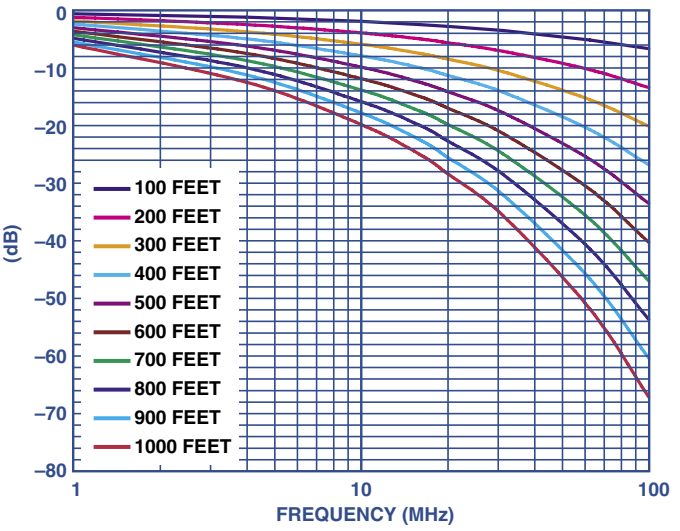

Figure 1. Frequency responses for various lengths of Cat-5 cable.

Using the data in Table I, the frequency response for each cable length can be approximated by a mathematical model based on a negative-real-axis pole-zero transfer function. Any one of the many available math software packages with the capability of least-squares polynomial curve fitting can be used to perform the approximation. Figure 1 suggests that, for long cables at high frequencies—because of the steepening slope, exceeding 20 dB/decade—consecutive negative-real-axis poles are required to obtain a close fit, while at low frequencies—to fit the nearly linear slope—alternating poles and zeros are required. As an extreme example, the frequency response for 1000 feet of cable at 100 MHz is rolling off approximately as $1/f^4$, which can only be attained by a model having multiple consecutive poles.

Equalization is achieved by passing the signal received over the cable through an equalizer whose transfer function is the reciprocal of the cable pole-zero model's transfer function. To neutralize the cable's frequency-dependency, the EQ has poles that are coincident with the zeros of the cable model and zeros that are coincident with the poles of the cable model.

Table I. Voltage-attenuation magnitude ratios of Cat-5 cable. For example, 500 feet of cable attenuates a 10-MHz, 1-V signal to 0.32 V, which corresponds to about –9.90 dB (Figure 1).

Frequency	100 ft	200 ft	300 ft	400 ft	500 ft	600 ft	700 ft	800 ft	900 ft	1000 ft
1 MHz	0.932	0.869	0.8100	0.7550	0.7040	0.65600	0.6120	0.57000	0.53200	0.496000
4 MHz	0.866	0.750	0.6490	0.5620	0.4870	0.42200	0.3650	0.31600	0.27400	0.237000
10 MHz	0.796	0.634	0.5040	0.4020	0.3200	0.25400	0.2030	0.16100	0.12800	0.102000
16 MHz	0.750	0.562	0.4220	0.3160	0.2370	0.17800	0.1330	0.10000	0.07500	0.056300
20 MHz	0.722	0.521	0.3760	0.2710	0.1960	0.14100	0.1020	0.07350	0.05300	0.038300
31 MHz	0.663	0.440	0.2920	0.1940	0.1280	0.08510	0.0565	0.03750	0.02480	0.016500
63 MHz	0.551	0.303	0.1670	0.0920	0.0507	0.02790	0.0154	0.00846	0.00466	0.002570
100 MHz	0.462	0.214	0.0987	0.0456	0.0211	0.00973	0.0045	0.00208	0.00096	0.000444

One of the properties of passive RC networks is that the alternating poles and zeros of their driving-point impedances are restricted to the negative real axis. This property also holds for those operational-amplifier circuits having a transfer function determined by the simple ratio of feedback-impedance to gain-impedance (Z_f/Z_g), where these impedances are RC networks. (The property does not hold for other cases, such as *active* RC filter sections that synthesize conjugate pole-pairs.)

For a practical equalizer design, we prefer that an EQ be based on a single amplifier stage in order to keep its adjustability manageable and to minimize cost and complexity. The equalizer to be discussed here uses RC networks of the former type, described by Budak, with alternating poles and zeros; but such a design precludes the use of a single amplifier stage to realize the consecutive zeros required to compensate for consecutive poles in the cable model at all frequencies. As a compromise that will provide good equalization for all but long cables at high frequencies, the design chosen uses a single amplifier to realize two zeros and two poles, alternating on the negative real axis.

Because equalization requires increased gain at the high end of the band, a low-noise amplifier is required. To avoid introducing significant errors due to amplifier dynamics, a large gain-bandwidth product is needed. For the specific design requirements of this application, the amplifier must have the capacity to perform frequency-dependent differential-to-single-ended transformations with voltage gain. The Analog Devices AD8129, just such an amplifier, is the heart of the basic frequency-dependent gain stage in the EQ. Figure 2 shows the dual-differential-input architecture of the AD8129, and its standard closed-loop configuration for applications requiring voltage gain.

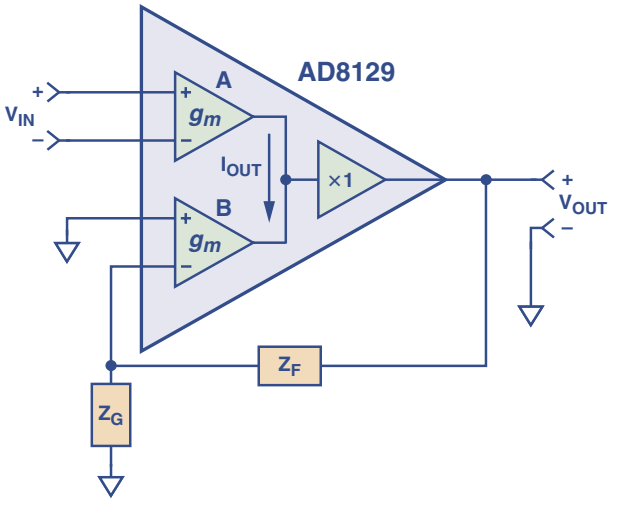

Figure 2. The AD8129 in a standard closed-loop gain configuration.

As can be seen, the AD8129 circuitry and operation differ from those of the traditional op amp; principally, it provides the designer with a beneficial separation of circuitry between the differential input and the feedback network. The two input stages are high-impedance, high-common-mode-rejection (CMR), wideband, high-gain transconductance amplifiers with closely matched g_m. The output currents of the two transconductance amplifiers are summed (at high impedance), and the voltage at the summing node is buffered to provide a low-impedance output. The output current of amplifier A equals the negative of the output current of amplifier B, and their transconductances are closely matched, so negative feedback applied around amplifier B drives v_{out} to the level that forces the input voltage of amplifier B to equal the negative of the input voltage of amplifier A. From the above discussion, the closed-loop voltage gain for the ideal case can be expressed as:

$$\frac{v_{out}}{v_{in}} = 1 + \frac{Z_f}{Z_g} \equiv A_V \qquad (1)$$

The EQ is designed using this gain equation, with RC networks for Z_f and Z_g. Its canonical circuit is depicted in Figure 3, which represents an EQ designed to compensate for a given length of cable.

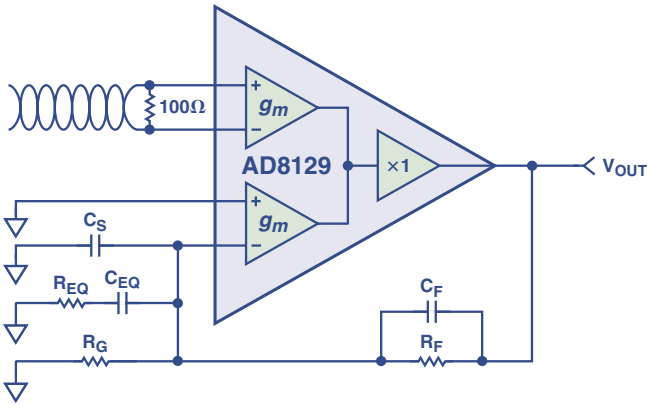

Figure 3. Canonical circuit of the equalizer.

In Figure 3, the high-differential input impedance of the upper transconductance amplifier facilitates provision of a good impedance match for the signal to be received over the Cat-5 cable; the lower amplifier provides the negative feedback circuit that implements the frequency-dependent gain. The Bode plot for the circuit has a high-pass characteristic, as shown in Figure 4. Z_n and P_n are the respective zeros and poles of the equalizer.

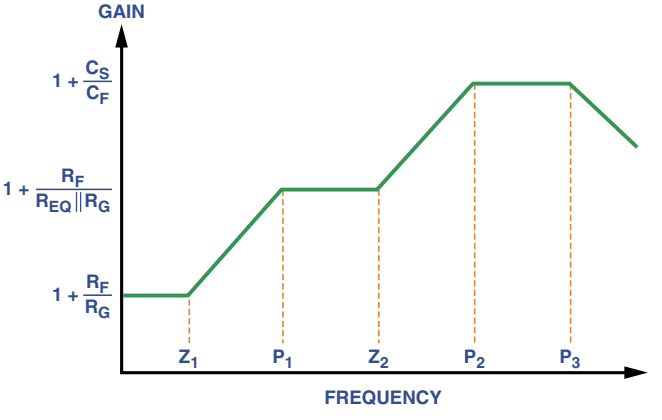

Figure 4. Bode plot of the canonical equalizer circuit.

In the following analysis, where the pole-zero pairs in Figure 4 are sufficiently separated, the capacitors can be approximated as short- or open circuits. The pole- and zero frequencies are expressed in radians per second. At low frequencies, all capacitors are open circuits, and the gain is simply

$$1 + \frac{R_f}{R_g}$$

This gain, set to compensate for flat (i.e., dc) losses, includes any loss due to matching and the cable's low-frequency flat loss. It also provides the flat gain required to stabilize the AD8129 when equalizing short cables (to be covered in greater depth below).

Moving up in frequency, the lowest-frequency pole-zero EQ section, comprising the series-connected R_{EQ} and C_{EQ}, starts to take effect, producing Z_1 and P_1. By approximating C_f and C_S as open circuits, the following equations can be written:

$$Z_1 = \frac{1}{\left(R_g + R_{EQ}\right) C_{EQ}} \tag{2}$$

$$P_1 = \frac{1}{R_{EQ}\, C_{EQ}} \tag{3}$$

The magnitude of the frequency response asymptotically approaches

$$1 + \frac{R_f}{R_{EQ} \parallel R_g}$$

as C_{EQ} approaches a short circuit.

As frequency increases, C_S begins to take effect, introducing another zero, Z_2. The primary function of C_f is to keep the amplifier stable by compensating for C_S. By initially approximating C_f as an open circuit ($C_f << C_S$), and presuming that Z_2 is sufficiently far in frequency from P_1 that C_{EQ} can be considered as having negligible impedance compared to R_{EQ}, the approximate expression for Z_2 can be written:

$$Z_2 = \frac{1}{\left(R_f \parallel R_{EQ} \parallel R_g\right)\left(C_S + C_f\right)} \approx \frac{1}{\left(R_f \parallel R_{EQ} \parallel R_g\right) C_S} \tag{4}$$

Finally, P_2 can be expressed as:

$$P_2 = \frac{1}{R_f\, C_f} \tag{5}$$

Between P_2 and P_3, the magnitude of the frequency response asymptotically approaches the closed-loop gain produced by the capacitive divider formed by C_f and C_S,

$$1 + \frac{C_S}{C_f}$$

This is the closed-loop gain at frequencies leading up to P_3, so P_3, which is due to the amplifier's dominant-pole roll-off, can be approximated as:

$$P_3 = \left[1 + A_O\left(\frac{C_f}{C_f + C_S}\right)\right]\omega_p \tag{6}$$

where A_O is the amplifier's dc open-loop gain, and ω_p is the amplifier's dominant pole. This result follows directly from standard op-amp gain-bandwidth analysis. P_3 is imposed by the gain-bandwidth product of the amplifier, and sets the approximate upper frequency limit of the equalizer. Using the above results, along with the cable's pole-zero model, an EQ can be designed for any practical length of cable that can be modeled by two alternating pole-zero pairs, provided that the amplifier has a sufficiently high-gain-bandwidth product.

In order for the EQ to be useful over a wide range of cable lengths, it must be adjustable. A simple means of adding adjustability is to switch different RC networks between the feedback pin of the AD8129 and ground. This scheme is illustrated in Figure 5.

Each EQ section in Figure 5 is appropriate for a range of cable lengths. Section EQ0 covers 0 to 50 feet, and Section EQ10 covers 950 to 1000 feet. The other sections are centered on 100 feet, 200 feet, etc., and cover ±50 feet from their centers. This resolution is sufficient for most RGB applications.

Practical Matters

The AD8129 is stable for gains greater than 10 V/V, where it has a nominal phase margin of 56°, but if care is taken with regard to layout and parasitic capacitance, it can be successfully operated with a gain of 8, where it has approximately 45° of phase margin.

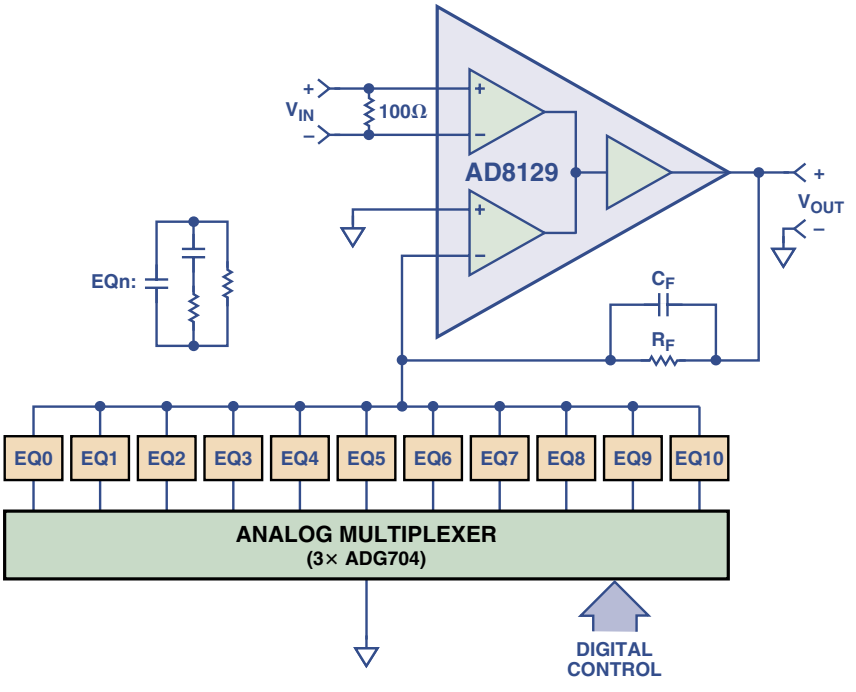

Figure 5. Equalizer with switchable sections.

This gain is required at high frequencies. For the longer cable lengths, sufficient high-frequency gain is provided by the high-pass nature of the equalizer. For cable lengths between 0 and 300 feet, however, excess flat gain is required in order to keep the AD8129 stable. Because the excess gain is flat, it can be easily inserted by adjusting the R_f/R_g ratio, and removed by switching in the same amount of flat attenuation after the equalizer.

The AD8129's input stage has a limited linear dynamic range (±0.5-V operating range). For optimum performance, it is best to attenuate the 700-mV RGB video signals by a factor of *four* before applying them to the AD8129 inputs. Sometimes the video signals are already attenuated by a factor of two before transmission over the cable. (This is not the matching loss—which is normally accounted for by using a cable *driver* with a gain of 2.) In this case, an additional factor-of-two attenuation can be inserted at the input to the AD8129 to produce an end-to-end flat attenuation factor of four. A buffer with a flat gain of four, placed after the EQ, is used to compensate for this attenuation (the AD8001 is an excellent choice for this stage). The buffer also simplifies the switched attenuator at the EQ output, which can be a simple L-pad.

The parasitic capacitance of each *off* channel in the ADG704 analog multiplexers used to select the EQ sections is 9 pF. The parasitic capacitance of the sum all of the unselected EQ sections is therefore quite large; it adds to the C_S value of the selected EQ section. For the EQ sections from 400 to 1000 feet, this parasitic capacitance can usually be absorbed into C_S. For the shorter sections, the excess closed-loop gain described above is used to compensate for the peaking caused by the parasitic capacitance. As a general rule, it is best to scale the impedances used in the EQ sections in such a way as to maximize the capacitance values, thus allowing absorption of as much parasitic capacitance into C_S as possible. This can't be carried too far however, since it reduces the associated resistances. The scaling is also limited by the parasitic inductance in the traces that connect the EQ sections. Small resistances provide little damping; if the resistance levels are too small, a moderate-Q tank circuit, resulting from the parasitic trace inductances and switch capacitances, can cause instability in the AD8129.

Optimizing the EQ PCB layout is of paramount importance. The major part of all power- and ground-plane copper must be removed from all layers under the traces that connect to the AD8129 summing node. Small ground-plane strips can be strategically placed as needed in these areas to provide low-Z return current paths, while minimizing stray capacitance at the summing node. The AD8129s and ADG704s should be in μSOIC packages, and the AD8001 should be in the SOT-5 package. Trace inductance in the EQ sections must be kept to an absolute minimum to avoid instability in the AD8129, so 0402 packages should be used for the resistors and capacitors, and the EQ sections should be laid out in such a way as to minimize trace lengths.

After the RC values that are based on the cable model have been determined, and the parasitic effects have been taken into account, a final tuning process in the time domain is required for RGB video applications. This is because one of the most important performance metrics for RGB video circuits is the step response; the *step response* of the cable and EQ combination must be tuned so as to exhibit fast rise time, minimum overshoot and ringing, and short settling time. C_S has the greatest effect on overshoot and ringing, and the series connection of R_{EQ} and C_{EQ} has the greatest effect on the long-term settling time. The positions of the pole and the zero produced by the series connection of R_{EQ} and C_{EQ} can be altered somewhat without changing the frequency response a great deal, because they are placed where the cable's frequency response has a rather gradual roll-off. This means that the equalized frequency response can appear to be quite good, while the positions of the pole and zero can be suboptimal from a step-response standpoint. It is therefore best to fine-tune the values of C_S, R_{EQ}, and C_{EQ} in the time domain by adjusting their values to produce a step response with the shortest settling time.

Since the equalizer must interface with long differential cables with no ground reference, the received signal may contain large common-mode voltage swings with respect to the power supply voltages at the receiver. It is therefore best to use dual power supplies of at least ±5 V. This also allows the output signal to swing to 0 V, which is generally required for video signals.

CONCLUSION

The equalizer presented here can stably compensate for lengths of Cat-5 cable from 0 to 1000 feet at frequencies to greater than 100 MHz at short cable lengths and to 25 MHz at 1000 feet, making it suitable for KVM networking and other high-resolution video transmission applications.

REFERENCE

Budak, *Passive and Active Network Analysis and Synthesis*, Houghton Mifflin, 1974.

This article can be found at http://www.analog.com/library/analogDialogue/archives/38-07/equalizer.html, with a link to a PDF.

A READER NOTES—REACTANT FLOW SENSOR

Jason Dugas, an electrical engineer at NASA's Johnson Space Center recently contacted the editors about an article published in *Analog Dialogue* in 1971. The article, entitled *Measuring Air Flow Using a Self-Balancing Bridge* (see below), details how an operational amplifier can be applied in the design of hot wire anemometers used to monitor flow rates. An excerpt of Jason's letter follows.

> The article was of interest because I wanted to learn more about the history of a similarly-designed flow meter that uses this architecture and is still flying on the space shuttle fleet today! The flow meters are used to monitor oxidizer and fuel flow rates at the inlets of the shuttle orbiter's fuel cell stacks.
>
> Analyzing the flow meter has been a lesson in history, as much as in electrical engineering, as the circuit for which I'm responsible was designed in 1971 (I wasn't born until 1978). The reactant flow sensors use the old "Royer oscillator" for the internal power supply. Both +24 and –12-V dc supplies are required for the circuit's operation; the internal supply produces these voltages from a 28-V bus. The power supply consists of two sections: (1) a dc-dc converter that takes the 28 V from the shuttle's power bus and produces regulated +20 V, and (2) a self-oscillating power oscillator that converts the +20 V into a square wave, which is then coupled via two transformer windings to bridge rectifier circuits that produce the isolated +24-V and –12-V supply voltages.
>
> The dc-dc converter is based upon a LM105 voltage regulator. The LM105—developed in the late 1960's and now obsoleted by newer designs/topologies—can be used as both a linear and a switching regulator. Refer to National Semiconductor application notes AN-1, AN-2, AN-8, and AN-23 (if you can find them by now—Ha!).
>
> Once the output of the dc/dc converter has ramped up, the oscillator starts to function. Start-up of the oscillator depends on a mismatch between two BJTs, which results in one of them reaching saturation first, while the other one is driven to cutoff.
>
> The other main functional area of the RFS is the circuitry that interfaces with the temperature sensors. The flow sensor works on thermal convection. The basic topology is a hot wire anemometer. Hot wire anemometers determine flow by measuring the change in temperature of an element due to convective heat loss. Traditionally, the sensor operates from a constant current source. This increases sensitivity but presents a problem when there is a wide range of reactant flow. At low flow rates, the element will overheat and potentially fail. The RFS avoids this problem by using a slightly different approach, called a self-balancing bridge. In this design, the temperature of the sensor is held constant by varying the excitation voltage of the bridge. The flow rate is inferred from the change in excitation voltage.
>
> In order to become smarter and familiar with the older architectures, I consulted the now historical literature, including early volumes of *Analog Dialogue* and the *Switching and Linear Power Supply, Power Converter Design* by Pressman, written in 1977—one year before my birth.

APPLICATION BRIEF

Reprinted from *Analog Dialogue* Volume 5, Number 1, January 1971

Measuring Air Flow Using a Self-Balancing Bridge

The design of a hot wire anemometer presents an interesting application for operational amplifiers.

The purpose of the instrument is to measure air speed by its cooling effect upon an electrically heated platinum filament that exhibits a high positive temperature coefficient of resistivity.

The filament characteristic is shown in Figure 1, for two values of air speed. Two classical ways of operating it would be at constant voltage or constant current.

With constant current applied (say, 0.6 amperes), the sensitivity is reasonably good—about 0.4V change for $\Delta s = s_2 - s_1$.

However, there is every prospect that in still air the filament has the potential of burning itself up (resistance increases as temperature increases, due to lack of air flow, which causes the voltage to increase [constant current], increasing the dissipation, the temperature, and again the resistance, etc.).

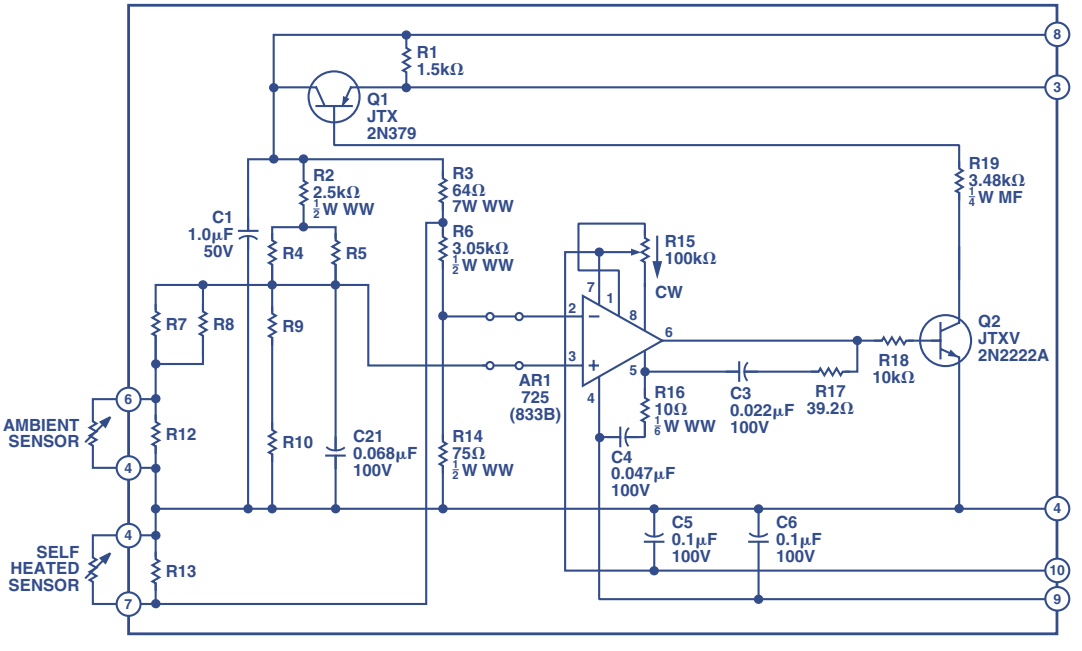

This bridge amplifier, used to measure oxidizer- and fuel flow rates on the space shuttle, was derived from the circuit described in a 1971 *Analog Dialogue* article.

With constant voltage applied, the increase of resistance with temperature causes operation to be quite safe, but also relatively insensitive, especially at low air speeds.

Another factor that makes both constant-current and constant-voltage operation unsatisfactory is the necessity for the temperature of the filament to change to detect a change in air speed; this necessarily causes delay. If the air speed indicator is part of a control loop, the measurement delay could cause slow response or instability of the control loop.

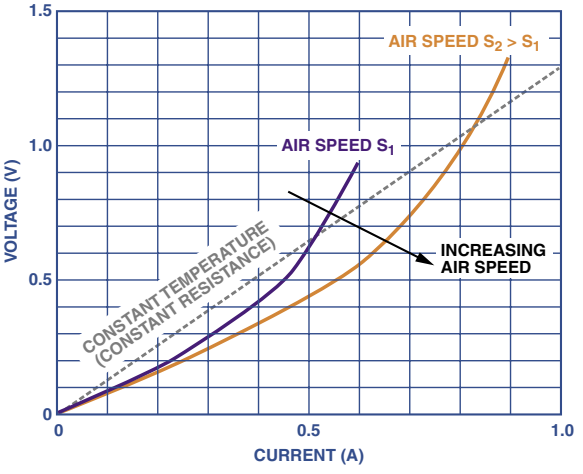

Figure 1. Volt-ampere characteristics of 7/16"L × 0.002"D straight filament of pure platinum in the presence of moving air, for two values of air speed.

An approach that answers all of these objections satisfactorily is to operate the filament as though it had constant resistance (i.e., constant temperature).

If, for example, the resistance were maintained at 1.3 Ω, as indicated by the dashed line, one could obtain a current change of 0.3 A and a voltage change of 0.4 V, with no danger of overheating in normal operation. Response would be quite speedy, since temperature changes are momentary and small.

The basic circuit for achieving constant temperature operation is the feedback circuit of Figure 2, consisting of a bridge, an op amp, and a power amplifier. The operational amplifier continuously adjusts the flow of current (through the power transistor) to maintain its two inputs equal. This can be done only by keeping the voltage across the filament equal to that across R2, and the filament current equal to the current through R1. However, since the current through R1 is proportional to the current through R0, (which has the same voltage drop as R1) and the current through R0 is determined by the voltage drop across R2, it can be seen that the resistance of the filament, RF, must be equal to that of R2, multiplied by the ratio of R_1 to R_0.

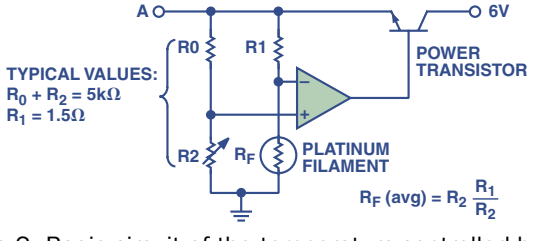

Figure 2. Basic circuit of the temperature-controlled bridge.

THE AUTHOR

In 1971, Ing. José Miyara was a partner in the firm TELEGUARD, in Rosario, Argentina. Their principal activities are in Industrial Electronics and Automatic Control.

Suppose now, that, starting from a given operating equilibrium point, the air flow increases. This will take heat away from R_F, causing its voltage to tend to drop. The amplifier's output voltage increases, which increases the current through the power transistor, and thus makes more power available for the filament to dissipate to maintain its temperature (and hence its resistance) constant.

The output voltage is measured at terminal "A", which provides an amplified version of the filament voltage, at an impedance level low enough to operate even the crudest of meter movements.

The zero-air-speed voltage is backed off by means of an auxiliary constant voltage, and the readings can be displayed with a moving coil meter. The scale is a nonlinear function of air speed, actually expanding toward the lowest values. LOW air speeds can be read with high sensitivity; in fact, the device can virtually detect a whisper several feet away.

CIRCUIT NOTES

For practical realization, the following points must be considered:

1. A voltage offset must be deliberately introduced into the operational amplifier (or elsewhere) to insure that the output goes positive with zero differential input; otherwise, the circuit might remain dead when turned on.
2. The power transistor must have ample current-handling capacity; the filament requires several hundred mA
3. Depending on the physical layout, especially when the filament is away at the end of a twisted pair, wild high-frequency oscillations are possible. Though not visible with low-frequency readout devices (however, rectification can cause voltage offsets), they look nasty on an oscilloscope screen. A 0.1 μF capacitor between the base and collector of the power transistor can often serve to stabilize these oscillations. A small resistor in series with the base may also be helpful.
4. The filament is a physical device with thermal lag. Although the circuit is fast enough to prevent loop oscillations when used in a larger control loop, it is itself a process control loop and may require the usual compensation techniques to maintain its own internal stability.
5. R0 and R2 form a trim potentiometer to set the operating temperature (e.g. resistance) of the filament. If R2 is a variable resistance, you start with $R_2 = 0$, and increase it until the filament just starts to glow, then back down a little. This will give optimal sensitivity.

APPLICATIONS

The device has been used in a commercially-produced apparatus to trip out equipment when the air speed in a forced draft duct falls below a preset value, but the approach is suggestive of a number of other applications in instrumentation.

In gas chromatographs, it could be used to monitor the minute gas flows required, and also to assure optimum sensitivity from thermal-conductivity filaments, while preventing their burnout in faulty operation.

Another application could be for constant-temperature ovens for crystals, differential pairs, etc., using copper or a thermally sensitive alloy such as "Balco" for the combined heater/temperature-sensor function. In such an arrangement, the control- led temperature could be slaved to another (arbitrarily) variable temperature (e.g., to insure a constant temperature difference) if R2 were a platinum temperature bulb with the circuit so dimensioned as to avoid causing it to introduce errors due to its own self-heating.

PRODUCT INTRODUCTIONS: VOLUME 38, NUMBER 3

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

July

ADC, Successive-Approximation, 16-bit, 250-ksps **AD7685**
Amplifier, Differential ADC Driver, low-cost, low-power **AD8137**
Amplifier, Differential ADC Driver, low-noise, rail-to-rail outputs **AD8139**
Amplifier, Difference, single-supply, 42-V systems **AD8205**
Amplifier, Limiting, 2.7-Gbps **ADN2890**
Amplifier, Operational, dual, precision, rail-to-rail outputs **AD8698**
Amplifier, Sensor Signal Conditioning, digitally programmable **AD8555**
Analog Interfaces, 110-/140-Msps for flat-panel displays ... **AD9985**
CCD Signal Processor includes vertical drivers, timing generator **AD9925**
Controller, Secondary-Side **ADM1041**
Controller, Synchronous Buck, for desktop PCs **ADP3188**
Controller, Synchronous Buck, for mobile PCs **ADP3206**
DACs, Voltage-Output, 12- and 14-bit, 32- and 40-channel **AD538x**
Digital Potentiometer, 256-position, ultralow-power, 1.8-V logic..... **AD5165**
Digital Signal Processor, 32-bit SHARC, low cost ... **ADSP-21262**
Digital Signal Processor, 32-bit SHARC **ADSP-21266**
Laser-Diode Driver, 4-channel, dual-output **AD9664**
LED Driver/Monitor for automotive applications **AD8240**
MicroConverter, system front-ends—include ADC, DACs, and flash MCU **ADuC84x**
MOSFET Driver, high-voltage **ADP3419**
Multiplexer, CMOS, 4-channel, 0.5-ohm *on* resistance ... **ADG804**
Multiplexer, CMOS, 4-channel, wideband, 37-dB isolation at 1 GHz **ADG904**
Switch/Multiplexer, CMOS, dual 2-channel, 0.5-ohm on resistance **ADG836L**
Video Decoder, multi-format SDTV, 9-bit **ADV7181**
Video Decoder, multi-format SDTV, 10-bit **ADV7183A**
Video Decoder, multi-format SDTV, 12-bit **ADV7189**
Video Decoder, multi-format SDTV/HDTV, RGB graphics digitizer **ADV7402**

August

Active Mixer, 400-MHz to 1200-MHz, *receive*-channel **AD8344**
ADC, Successive-Approximation, 14-bit, 100-ksps, 6-lead SOT-23 package **AD7940**
DACs, Voltage-Output, 8- and 16-channel, 12- and 14-bit **AD539x**
Digital Potentiometers, dual, 64-/256-position, nonvolatile memory, I^2C interface **AD5251/AD5252**
Frequency Synthesizer, fractional-N, 4-GHz **ADF4154**
RF Vector Modulator, 700-MHz to 1000-MHz **AD8340**

September

ADC, Pipelined, 8-bit, 250-Msps, 3.3-V supply **AD9480**
Amplifier, Differential Driver, triple **AD8133**
Amplifier, Operational, low-noise, ultralow-distortion, high-speed **AD8045**
Amplifier, Operational, quad, low-power, high-current, ideal for ADSL **AD8392**
Audio Codecs, stereo **ADAV801/ADAV803**
Clock Generator, 655-MHz, low-jitter, PLL-based **AD9540**
DAC, Wideband, dual 12-bit, includes resistors for 4-quadrant multiplication, with parallel interface **AD5405**
DAC, Wideband, dual 12-bit, includes resistors for 4-quadrant multiplication, with serial interface **AD5415**
DACs, Multiplying, dual 8-/10-/12-bit **AD5429/AD8439/AD5449**
DC-to-DC Converter, step-down, synchronous rectifier, 500-mA drive **ADP3051**
Detector/Controller, Logarithmic, 1-MHz to 8-GHz, 60-dB range **AD8318**
Digital Potentiometer, 3-channel, nonvolatile memory ... **AD5255**
Digital Potentiometer, 3-channel, nonvolatile memory .. **ADN2860**
Diversity Receiver, IF-to-baseband, 12-bit, 65-Msps **AD6652**
Frequency Synthesizer, PLL, includes voltage-controlled oscillator **ADF4360-0**
Frequency Synthesizer, DDS, 2.7-GHz **AD9956**
Gamma Buffer plus Vcom Driver, 10-channel, for flat-panel LCD displays **ADD8710**
Level Shifters for LCD-panel timing **ADSY8401**
Power-Supply Sequencers **ADM108x**
RF Vector Modulator, 1.5-GHz to 2.4-GHz **AD8341**
RS-485 Transceiver, isolated, high-speed, half-duplex .. **ADM2486**
Switch/Multiplexer, CMOS SPDT, 0.5-ohm, low-voltage, SC70 package **ADG849**
Thermal Monitor/Controllers, fan-speed .. **ADM1033/ADM1034**
Video Codec, JPEG 2000 ..... **ADV202**
Video Decoder, 10-bit, multi-format SDTV/HDTV, RGB graphics digitizer **ADV7400**

AUTHORS

Christine Bako (page 3) has been an applications engineer at our Austin, Texas, facility since 2002, supporting ADI's latest video/still image-compression devices. She joined ADI in 1998, in Limerick, Ireland, working in support of our video encoder products. She is a graduate—with a BEng—of the University of Salford, Manchester, UK.

Eva Murphy (page 8) is an applications engineer at ADI's facility in Limerick, Ireland. She provides applications support for general-purpose, digital-to-analog converters (DACs), switches/multiplexers, and direct-digital-synthesis (DDS) products. Joining Analog Devices in 2000, Eva holds a BEng from Cork Institute of Technology. Her interests include reading, traveling, and walking.

Jonathan Pearson (page 13) has been an applications engineer in the High-Speed Amplifier group since August 2002. Prior to working at ADI (and the telecom crash of 2001), he worked as an analog circuit and systems designer in the telecom industry. He holds a BSEE from Northeastern University, an MSEE from WPI, and two patents. Besides spending time with his family, he enjoys playing a variety of guitars, recording music, and collecting vacuum-tube guitar amplifiers and antique radios.

Colm Slattery (page 8) graduated in 1995 from the University of Limerick, Ireland, with a bachelor's degree in Electronic Engineering. After working at Microsemi in test-development engineering, he joined ADI in 1998 as a test-development engineer. He later (2001) moved to Applications in the Precision Data-Converter product line.

Analog Dialogue

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

Analog Devices, Inc. Japan Headquarters
New Pier Takeshiba
South Tower Building
1-16-1 Kaigan,
Minato-ku, Tokyo
105-6891, Japan
Tel: 813.5402.8210
Fax: 813.5402.1063

Analog Devices, Inc. Southeast Asia Headquarters
39/F One Pacific Place
88 Queensway
Admiralty,
Hong Kong, PRC
Tel: 852.2506.9336
Fax: 852.2506.4755

Printed in the U.S.A. M02000383-72-11/04

ANALOG DEVICES

www.analog.com/analogdialogue