Schematic shows a single-cell, 2100mAh NiMH battery being fast-charged through a high-power USB port. Current drawn from the port is 420mA. (See article inside, page 13.)

News Brief

■ MAXIM REPORTS 32% YEAR OVER YEAR EARNINGS PER SHARE GROWTH FOR FISCAL 2005

Maxim Integrated Products, Inc., (MXIM) reported net revenues of $400.4 million for its fiscal fourth quarter ended June 25, 2005. Net income for the fourth quarter was $126.1 million, a slight increase over the $124.7 million reported for the fourth quarter of last year. Diluted earnings per share were $0.37 for the fourth quarter, a 2.8% increase over the $0.36 reported for the same period a year ago. Net revenues, net income, and diluted earnings per share for the fourth quarter were similar to those reported for the third quarter of fiscal year 2005.

For the 2005 fiscal year, Maxim reported net revenues of $1.672 billion compared to $1.439 billion for last year, a 16.2% increase. Net income for the 2005 fiscal year was $540.8 million compared to $419.8 million reported for fiscal 2004, a 28.8% increase. Diluted earnings per share grew 31.7% from $1.20 per share reported in fiscal 2004 to $1.58 per share in fiscal 2005.

The Company's free cash flow was $145 million, or $0.43 per diluted share, for the fourth quarter of fiscal 2005, compared to $111 million, or $0.32 per diluted share, for the fourth quarter of fiscal year 2004. Excluding a one time payment of $40 million for the settlement of a license matter, free cash flow for the fourth quarter of fiscal 2005 would have been $185 million or $0.54 per diluted share.

Free cash flow is defined as cash from operating activities (after tax) less additions to property, plant and equipment as reported in the Company's statements of cash flows.

During the quarter, cash and short-term investments increased $71.5 million after the Company repurchased 1.8 million shares of its common stock for $71.7 million, paid dividends of $32.7million, and acquired $20.9 million in capital equipment. At year end, cash and short-term investments totaled $1.475 billion. Accounts receivable remained unchanged during the quarter at $192.3 million, and inventories increased $8.3 million to $167.8 million in the fourth quarter.

Research and development expense was $84.9 million or 21.2% of net revenues in the fourth quarter, compared to $83.1 million or 20.8% of net revenues in the third quarter of fiscal year 2005. The increase in research and development expense in the fourth quarter was primarily due to hiring additional engineers to support new product development. Selling, general and administrative expenses decreased slightly from $24.7 million in the third quarter or 6.2% of net revenues to $23.4 million or 5.8% of net revenues in the fourth quarter. The decrease in selling, general and administrative expenses was primarily due to lower litigation costs in the fourth quarter. Below-the-line spending was 27.0% of net revenues for both the third and fourth quarters of fiscal 2005.

Fourth quarter bookings were approximately $398 million, a 7% increase from the third quarter's level of $373 million. Maxim bookings were up 11% while Dallas bookings were down 7%. Turns orders received in the quarter were approximately $171 million or 43 % of net bookings, a 10% increase over the $156 million or 42% of net bookings received in the prior quarter (turns orders are customer orders that are for delivery within the same quarter and may result in revenue within the same quarter if the Company has available inventory that matches those orders). Bookings increased in all geographic locations except the United States. Fourth quarter ending backlog shippable within the next 12 months was approximately $313 million, including approximately $273 million requested for shipment in the first quarter of fiscal year 2006.

Jack Gifford, Chairman, President, and Chief Executive Officer, commented: "Even though our industry experienced a significant inventory correction in FY 2005, we were able to achieve significant growth in net income and diluted earnings per share. For the year, net income grew by about 29% while diluted earnings per share was up almost 32%."

Mr. Gifford continued: "We continue to be one of the most prolific inventors of innovative devices in the analog mixed signal space. But it is not a matter of numbers alone. We are proud to note that two of our new products received Product of the Year Awards. Electronic Products magazine gave this award to our digital video equalizer product. Our low cost, high accuracy analog output temperature sensor device won a similar accolade from AnalogZONE. These two products demonstrate the breadth of Maxim's analog expertise."

Mr. Gifford concluded: "The Company's Board of Directors has declared a quarterly cash dividend of $0.10 per share. Payment will be made on August 30, 2005 to stockholders of record on August 15, 2005."

For the complete Q405 press release, including safe harbor information, go to: www.maxim-ic.com/NewsBrief

Integrated DC Logarithmic Amplifiers

For over half a century, engineers have used log amps for compressing signals and for computation. Although digital ICs have mostly replaced the log amp in applications that require computing, engineers continue to use log amps to compress signals. Therefore, the log amp remains a key component in many video, fiber, medical, test and measurement, and wireless systems.

As implied by the name, a logarithmic amplifier expresses an output that is related to its input by the mathematical log function (the logarithmic base is not important, as the different log-based functions are related by constants). By utilizing the log function, you can compress the dynamic range of signals encountered by a system. Compressing wide-dynamic-range signals has several important benefits. The combination of a log amp and a low-bit-count ADC can often save board space and system cost, whereas a high-bit-count ADC might otherwise be required. Furthermore, low-bit-count ADCs are often already present in a given system or on a resident microcontroller. The conversion to a logarithmic parameter is also useful in many applications where measured quantities are evaluated in decibels, or where sensors exhibit exponential or near-exponential transfer characteristics.

During the 1990s, the fiber communications industry began using log-amp circuits to measure optical intensity in certain optical applications. Before that time, precision log-amp ICs were both costly and bulky; such expense was warranted in only a handful of electronic systems. The only alternative to these IC solutions were log amps constructed from discrete components. In addition to consuming even more board area, discrete-component log amps were frequently prone to temperature changes and required careful design and board layout. Highly matched components were also necessary to guarantee adequate performance over a wide range of input signals. Semiconductor manufacturers have since developed smaller and less costly integrated log-amp products with reduced temperature sensitivity and added functionality.

Classes of Logarithmic Amplifiers

There are three major classes of logarithmic amplifiers. The first class, the DC log amp, traditionally operates on slowly changing DC signals with bandwidths extending to about 1MHz. Without question, the most popular implementations use the logarithmic I-to-V transfer characteristic inherent in pn junctions. These DC log amps operate on unipolar inputs (current or voltage), and are frequently referred to as diode, transdiode, translinear, and transimpedance log amps. Due to their current inputs, DC log amps are commonly used to monitor wide-dynamic-range, unipolar photodiode currents—either absolute or ratiometric. Photodiode current monitoring is not only a common requirement in fiber communications equipment, but is also found in a wide range of chemical and biological sample-processing instrumentation. Other DC-based log amps exist, such as those based on the logarithmic time-voltage relationship of RC circuits. However, these circuits are typically prone to complications such as widely varying, signal-dependent resolution and conversion times, as well as high-temperature sensitivities.

The second class of log amp is known as the baseband log amp. This circuit class operates on rapidly changing baseband signals in applications where the compression of AC signals is required (common in certain audio and video circuits). The amplifier provides an output proportional to the logarithm of the instantaneous input signal. A special version of the baseband log amp is the "true log amp," which accepts bipolar inputs and provides a compressed output voltage that preserves the polarity of the input. True log amps are sometimes used in radio IF stages and medical ultrasound receiver circuits for dynamic range compression.

Finally, the third class of log amp is the demodulating log amp, or successive detection log amp. This class of log amp both compresses and demodulates RF signals, yielding the logarithm of the rectified signal's envelope. Demodulating log amps are prevalent in RF transceiver applications, where received RF signal strength is measured to control transmitter output power.

The Classic DC Logarithmic Amplifier

In the classic pn-junction-based implementation of the DC log amp, a bipolar transistor is used to generate the logarithmic I-to-V relationship. As shown in **Figure 1**, bipolar junction transistors (BJTs) are placed in the feedback path of an operational amplifier. Depending on the type of transistor chosen, npn or pnp, the log amp is either a current-sinking or current-sourcing circuit, respectively (Figures 1a and 1b). Through negative feedback, the op amp places enough output voltage on the base-emitter junction of the BJT to ensure that all available input current is drawn through the collector of the device. Note that a floating-diode implementation causes the op-amp output voltage to include input-referred offset; the grounded-base implementation does not possess this problem.

With the addition of an input series resistor, the DC log amp can also function as a voltage-input device. Input voltages are converted to a proportional current though the resistor, using the op amp's virtual ground as the reference. Clearly, op-amp input-referred offset must be
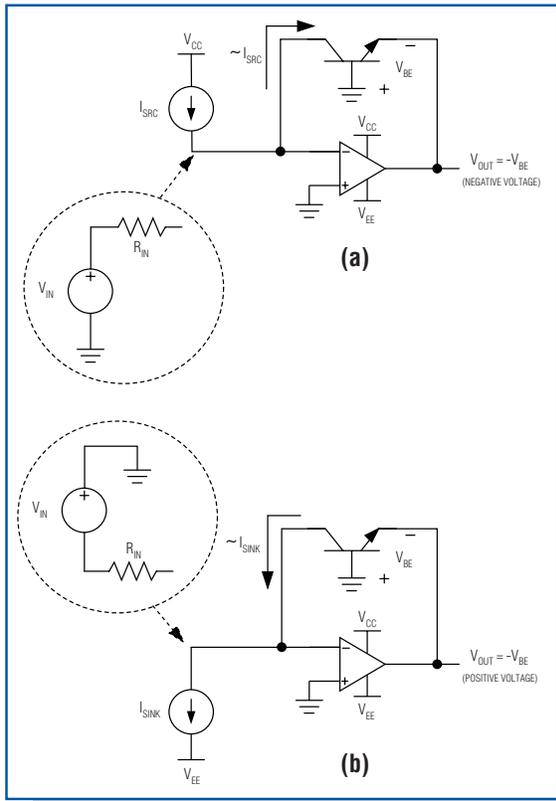
Figure 1. (a) The basic BJT implementation of a DC log amp has current-sinking inputs that generate negative output voltages. (b) By changing the BJT from npn to pnp, the log amp becomes a current-sourcing circuit and the output is positive.

minimized so that accurate voltage-to-current conversion can be achieved. The bipolar-transistor approach is prone to temperature variations but, as will be discussed, this sensitivity is drastically reduced by using a reference current and on-chip temperature compensation.

Further Details

The circuit of **Figure 2** shows a BJT log amp with two inputs, I_{IN} and I_{REF}. As described in the previous section, current presented to I_{IN} causes op amp A_1 to develop a corresponding output voltage:

$$V_{OUT1} = \frac{kT}{q} \ln\left(\frac{I_C}{I_S}\right) \approx \frac{kT}{q} \ln\left(\frac{I_{IN}}{I_S}\right) \qquad \text{(Eq 1)}$$

where:

$k = 1.381 \times 10^{-23}$ J/°K

T = absolute temperature (°K)

$q = 1.602 \times 10^{-19}$°C

I_C = collector current (mA or same units as I_{IN} and I_S)

I_{IN} = log-amp input current (mA or same units as I_C and I_S)

I_S = reverse saturation current (mA or same units as I_{IN} and I_C)

(In Equation 1, "ln" is used to represent the natural logarithm function. "Log$_{10}$" is used to express the base-10 logarithm function in subsequent equations).

Although this expression clearly shows the logarithmic dependence of V_{OUT1} on I_{IN}, the terms I_S and kT/q depend on temperature and will introduce a high-degree of variation in the V_{BE} voltage. To remove the temperature dependence caused by Is, a second junction voltage is subtracted from V_{OUT1} through the differencing circuit, composed of A_3 and its surrounding resistors. The second junction voltage is created in the same way as V_{OUT1}, except that I_{REF} serves as the input current. The transistors used to create both junctions must possess nearly identical properties and be in close thermal contact for proper cancellation.

$$V_{OUT} = \frac{kT}{q} \ln\left(\frac{I_{LOG}}{I_S}\right) - \frac{kT}{q} \ln\left(\frac{I_{REF}}{I_S}\right) \qquad \text{(Eq 2)}$$

$$= \frac{kT}{q} \left[\ln\left(\frac{I_{LOG}}{I_S}\right) - \ln\left(\frac{I_{REF}}{I_S}\right) \right] \qquad \text{(Eq 3)}$$

$$= \frac{kT}{q} \ln\left(\frac{I_{LOG}}{I_{REF}}\right) \qquad \text{(Eq 4)}$$

$$= \frac{kT}{q} \ln(10) \log_{10}\left(\frac{I_{LOG}}{I_{REF}}\right) \qquad \text{(Eq 5)}$$

The presence of I_{REF} has two benefits. First, it sets the desired x-axis "log-intercept" current—the current that makes the log-amp output current theoretically equal to zero. Second, it allows the users to take ratiometric measurements, in addition to absolute measurements. Ratiometric measurements are frequently used in optical sensors and systems, where an attenuated light source must be compared to a reference light source.

The expression of Equation 5 is still subject to temperature effects, as V_{DIFF} is proportional to absolute temperature (PTAT). By adding subsequent temperature-compensation circuitry (normally, an additional op-amp amplifier stage with a resistive temperature detector [RTD], or similar device, incorporated as part of the gain), the remaining
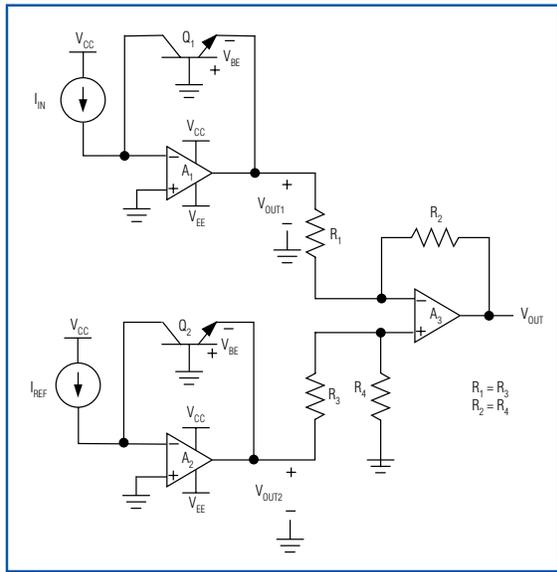
Figure 2. When two of the basic BJT-based input structures are used, subtracting V_{OUT2} from V_{OUT1} removes I_S temperature dependence from the output. The remaining "PTAT" dependence can be minimized by using a suitably selected RTD (resistive temperature detector) with the gain-setting resistors of the difference amplifier.

PTAT error can be virtually eliminated, yielding the ideal log-amp relationship:

$$V_{OUT} = K \log_{10}\left(\frac{I_{LOG}}{I_{REF}}\right) \qquad \text{(Eq 6)}$$

where K is the new scaling constant, also known as log-amp gain, given in V/decade. Because applying $\log_{10}$ to the ratio I_{LOG}/I_{REF} determines the number of decades I_{LOG} is above or below I_{REF}, multiplication by K yields the desired units of volts.

Integrated designs are well suited for DC log amps, because key temperature-dependent components can be co-located on the physical circuit, yielding excellent temperature tracking of those components. Furthermore, trimming of various remaining errors is also possible during production. Any remaining errors are generally well documented in the log amp's data sheet.

Modern-Day DC Logarithmic Amplifiers

The functional block diagram of **Figure 3** illustrates a typical contemporary DC log amp, the MAX4206. Like previous generations, today's DC log amp possesses op-amp input structures, BJT feedback, a differencing amplifier, and temperature compensation. To eliminate negative driving voltages at its emitter, the BJT transistor's circuit connections have been rearranged to facilitate

single-supply operation. An uncommitted op amp is still commonly used for subsequent gain, offset adjustment, or even PID control.

Unlike its ancestors, the contemporary log amp incorporates all its electronics within a small package (the MAX4206 is available in a 4mm x 4mm, 16-pin TQFN package). Prior to 2001, DC log amps were only commercially available in much larger DIP packages with pin counts ranging from 14 to 24. These older components carried a substantial price tag of $20 to $100; today's replacements are readily available in the $5 to $15 range.

Single-supply operation is a new improvement appearing in some modern-day DC log amps, making them desirable for use with single-supply ADCs/systems. The MAX4206 can operate from either a single +2.7V to +11V supply or a dual ±2.7 to ±5.5V supply. A consequence of single-supply operation is that these log amps generally hold a typical 0.5V common-mode voltage at their input terminals in order to maintain proper biasing on the logging BJTs. Because these log amps are current-input devices, this internally generated common-mode voltage is generally not a problem in most current-measurement applications.

The presence of an on-chip current reference has become quite popular in most contemporary DC log amps. This reference can be connected to the reference input of the log amp, thereby permitting an absolute, rather than ratiometric, measurement of the current presented to the log amp's main current input. In the case of the MAX4206, a reference current is obtained by means of a 0.5VDC voltage source, a voltage-to-current converter, and a 10:1 current mirror. An external resistor is required to program the desired reference current.

Also new to DC log amps, an on-chip voltage reference is sometimes provided to assist in adjusting amplifier offset at the uncommitted op amp. This reference can also be used for general purposes.

Application Examples

Without question, the most common applications of the DC logarithmic amplifier are those involving the measurement of light. Two implementations are generally used. In the first, a single photodiode is connected to the logging input, while a reference current is connected to the reference input. The second implementation uses two photodiodes, one connected to the logging input and the other to the reference input. The former implementation is used when absolution measurements of light intensity are desired, the latter for logarithmic-ratiometric ("log-ratio") light-intensity measurements.

Generalized circuits for both implementations are shown in shown in **Figure 4**. In Figure 4(a), a single photodiode measures the light from a fiber optic channel by viewing
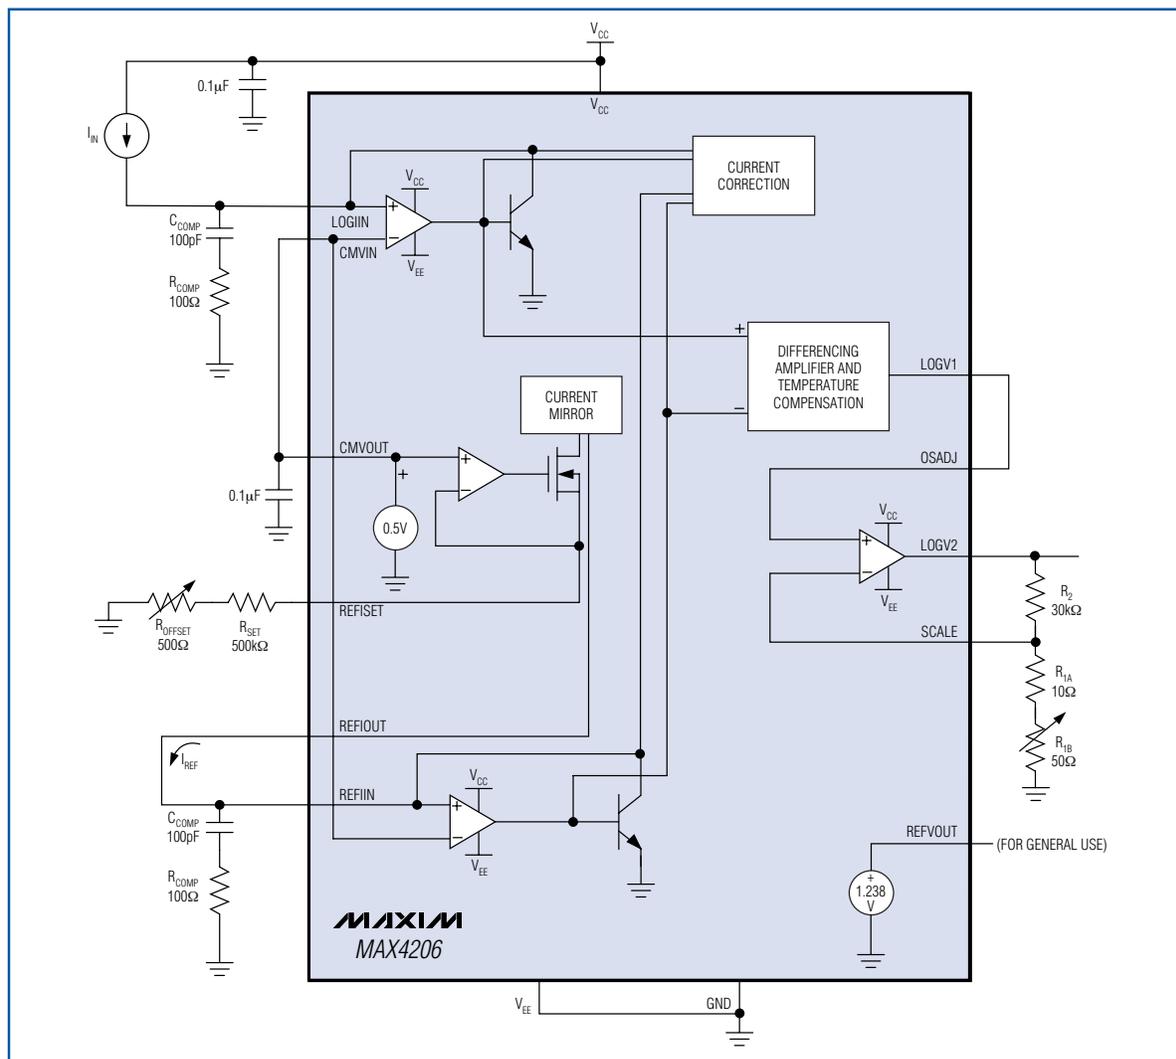
Figure 3. A typical DC log amp, like the MAX4206, integrates components such as trimmer pots and output amplifiers. Therefore, it requires few additional components for operation.

the light emanating from a fiber optic tap (1% transmission). A PIN photodiode is depicted in the diagram, although an avalanche photodiode could have been used for greater sensitivity (proper power-supply precautions should be taken if high voltages are used to bias the photodiode). Because photodiodes' output current is generally linear with incident optical power (0.1A/mW is a typical photodiode sensitivity) and the MAX4206 operates over five decades of dynamic range, a circuit like this can reliably measure the fiber's optical intensities from 10μW to 1W. Note that, although the MAX4206 is guaranteed to operate over a -40°C to +85°C temperature range, the effects of changing operating temperature and optical frequency can drastically affect the performance of the photodiode.

For cases where the photodiode's anode is reserved for other circuitry, such as the high-speed transimpedance amplifier (TIA) found in many fiber optic modules, a precision current mirror/monitor may be used at the photodiode's cathode. The MAX4007 series of products are well-suited for this application. Refer to the MAX4206 and MAX4007 data sheets for further details.

When two photodiodes are used in a logging application, the intent is to compare a reference light source against an attenuated light source that is derived from the reference. In this way, the attenuation caused by a given medium can be measured independent of light source intensity (or at least small changes in the intensity). This type of application is commonly found in many optically based gas-sensor applications. In Figure 4(b), a light source's output is split equally into two paths. The first is incident upon the reference PIN photodiode, whose anode feeds the MAX4206's REFIIN input. The other path reflects 90° off
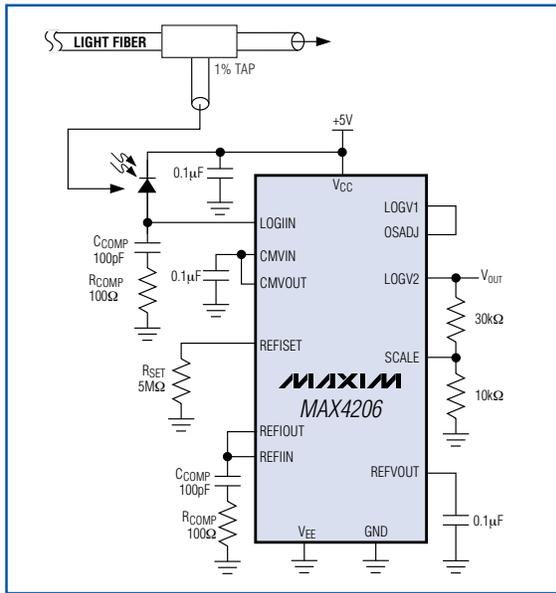
Figure 4a. Logging applications that measure optical intensity are easily created by placing a photodiode at the input of the log amp.
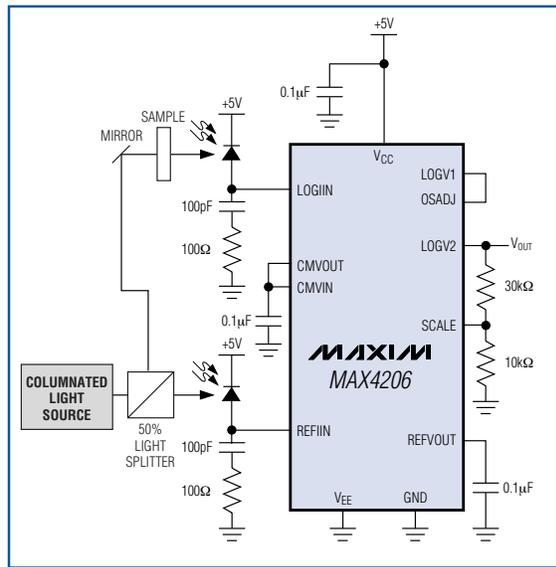
Figure 4b. Log-ratio applications use two photodiodes and generally measure optical attenuation.

a mirror, though the sample, and onto the other PIN photodiode (connected to the LOGIIN input). When calibrated so that the current from the reference photodiode measures 1mA, the current from the other photodiode will measure 1mA or less, depending on the attenuation encountered by the light. By anchoring input currents to 1mA or less, the MAX4206's wide five-decade dynamic range is fully utilized.

It should also be mentioned that, although the MAX4206 is not guaranteed to operate beyond in the 10nA to 1mA input current range, it is often possible to operate the device beyond this range while still maintaining monotonicity between inputs and output.

Sources of Error in DC Log Amps

Today's DC logarithmic amplifiers are still subject to the same limitations seen in previous generations of products. Equation 6 is an ideal approximation of the DC log amp. To obtain the most accurate expression possible, terms resulting from errors in gain, bias currents, offset, and linearity must be considered as well. This is especially important when such inaccuracies are worsened by temperature and time-dependent drift.

A more comprehensive expression representing the BJT-based DC log amp is given by:

$$V_{OUT} = K\left(1 \pm \Delta K\right)\left[\log_{10}\left(\frac{I_{LOG} \pm I_{BIAS1}}{I_{REF} \pm I_{BIAS2}}\right) \\ \pm \left(V_{CONF} \pm V_{OSOUT}\right)\right]$$

(Eq 7)

where ΔK is the gain variation; I_{BIAS1} and I_{BIAS2} are the bias currents associated with the LOGIIN and REFIIN inputs, respectively. V_{CONF} is the log-conformity error and V_{OSOUT} is the output-referred offset. K, I_{LOG}, I_{REF}, and V_{OUT} are defined previously. In many applications, errors associated with bias currents are quite small relative to the input and reference currents, and thus are usually omitted from the error expression. Log-conformity error is defined as the maximum deviation from the ideal log relationship of Equation 6 (assuming all other error sources have been nulled). This error is frequently presented in a difference format, so that the small deviations from the ideal plot line can be easily inspected (**Figure 5a**).

Though not immediately obvious, the reference current, I_{REF}, is potentially a large source of error, composed of initial inaccuracy, temperature drift, and age-related drift. Such errors should be included when assessing the overall error budget of the log amp's operation.

The effects of these nonideal variations are shown in the transfer curves of **Figure 5b** (these effects have been exaggerated for demonstration purposes). The ideal/desired expression is represented by the black solid line, shown with a log intercept of 100nA and a gain of 1V/decade. Output offset error shifts the black solid line up or down, as indicated by the blue dashed line. Gain error rotates the offset-shifted transfer characteristic, and is represented by the black dashed line. The blue dotted line shows the effects of incorporating of nonlinearity and output-margin errors.
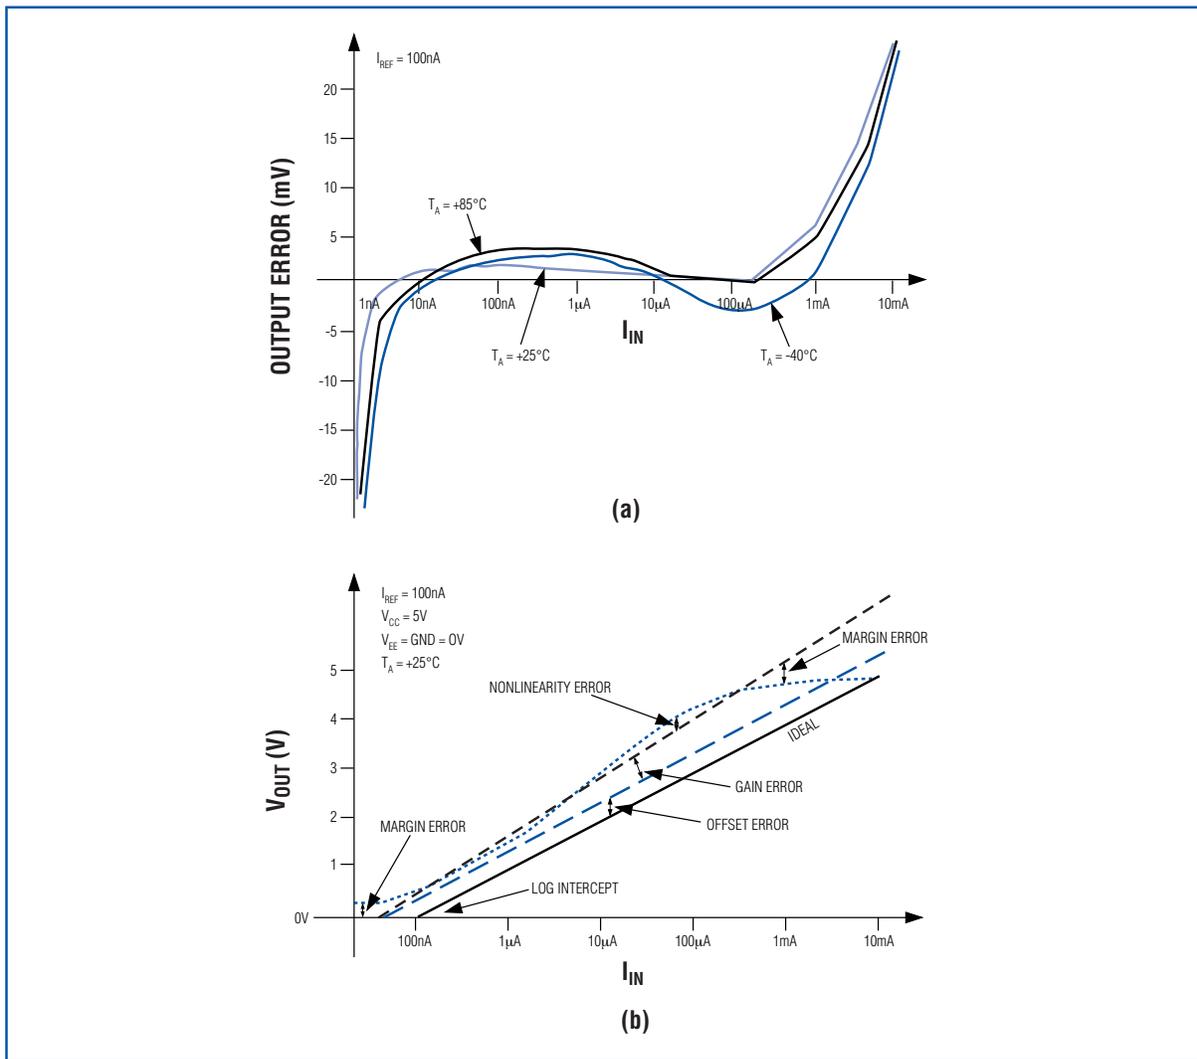
Figure 5. (a) A typical log-conformity error plot is commonly shown as a function of input current(s) and operating temperature. (b) The effect of the different errors, presented in Equation 7, on the log transfer function is shown. Errors have been exaggerated for clarity.

In practice, the log-amp manufacturers minimize many of the errors presented in this section. With additional calibration and temperature monitoring, designers can reduce the effects of these errors further. The designers usually use calibration tables to perform these corrections after the log-amp output has been digitized.

DC Log-Amp Implementation

The performance of a DC log amp is only as good as the circuit within which it resides. Good design and layout techniques minimize input leakage currents and cross-component temperature variations. However, good design and layout are rarely sufficient to ensure the performance required by most log-amp applications, especially when operated over a range of input currents and temperatures.

Depending upon an application's requirements and operating conditions, a suitable calibration procedure should be implemented to minimize cumulative errors.

Here are some tips to consider when implementing a DC log amp.

One-Point Calibration This is a "bare-minimum" technique that effectively shifts the raw performance line (blue dotted line) of Figure 5b vertically, such that it intersects the ideal-performance line (black solid line) at a single point. At the typical operating temperature, nominal input and reference currents are applied to the respective log-amp inputs. The resulting deviation from the desired result is subtracted from the actual log-amp output during normal operation.

Advantages: Calibration is quick, can be performed during final product testing, and does not require much computation. Calibration can also be performed in the analog domain with a single trimming resistor.

Disadvantages: Gain and offset errors are combined into a single, over-generalized correction. Correction value loses validity as inputs and temperatures are moved from the calibration conditions.

Two-Point Calibration Slightly more complex than the previous technique, this technique yields much better results. It effectively rotates and vertically shifts the blue dotted line in Figure 5b to approximate the desired black solid line. Again, a typical operating temperature should be selected. Input currents should span the desired operating range. The process is greatly simplified if only one reference current is used in both calibration and operation.

Advantages: Calibration is fairly quick and greatly reduces gain and offset errors. Calibration can be performed in the digital domain by applying gain and offset computations, or in the analog domain with gain and offset trimming resistors.

Disadvantages: Correction values lose validity as inputs and temperatures are changed.

Multipoint Calibration This technique creates a table of calibration data from key sample points. Samples are taken at a single operating temperature. Corrections are performed by interpolating between sampled points.

Advantages: With a sufficient number and strategic selection of input conditions, gain, offset, and nonlinearity errors can be greatly minimized.

Disadvantages: Some form of interpolation is required, which increases the amount of computation involved. Calibration loses validity as inputs and temperatures are changed.

Calibration with Temperature Adjustment Similar to multipoint calibration, this technique also considers test temperature, creating an extra degree of freedom.

Advantages: This technique greatly reduces gain, offset, nonlinearity, and temperature-imposed variations affecting the total error. This is a good option for high-performance, low-volume products.

Disadvantages: Calibration times during final product testing are much longer due to spanning of temperatures. Multidimensional interpolation of sampled data significantly increases the required computational resources. An additional temperature monitoring circuit is also required.

Maintaining Proper Input Margin Log-amp outputs should not operate near the power supply rails, as log-amp outputs have limited sourcing and sinking abilities near these rails. This suggestion is easily overlooked when trying to measure currents near or below the reference current, or near the maximum input current. Choose a reference current that is lower than the lowest expected input current. Gain should be set so that the maximum log-amp output voltage is not attained when the maximum input current is applied. A dual-supply log amp can also help, as identical input and reference currents put the output at midscale in most designs.

Advantages: Accuracy and response time increase under extreme input conditions.

Disadvantages: Usable output range decreases slightly.

Component Selection Use external resistors that are the same type and possess low-temperature coefficients. This is especially important for resistors whose absolute value influences performance (e.g., a reference current generator circuit). Parameters affected by resistance ratios, like gain and offset, are generally less affected by temperature changes. The temperature stability of compensation components is generally not critical. To avoid leakage issues when measuring small currents, low-leakage PC material should also be considered.

Advantages: Minimizes additional performance degradation created by external components.

Disadvantages: Low-temperature-coefficient components are generally slightly more expensive but are well worth the cost, given the increased performance they can provide.

Uniform Temperature Exposure No part of the log-amp circuit should be exposed to a temperature that is significantly different from any other part of the circuit. This precaution ensures that all circuit changes caused by temperature will track each other more closely.

Advantages: Further degrees of freedom are eliminated in the calibration process.

Disadvantages: This may cause inconvenience in layout routing or overall circuit size.

Conclusion

In summary, DC log amps have evolved into small, easy-to-use, cost-effective circuits nicely suited for certain analog designs. The logarithmic function conveniently compresses wide dynamic range signals and linearizes sensors with (semi-)exponential transfer functions. The compression created by the log function reduces the ADC bit count that would be required to digitize a wide-dynamic-range signal. Circuit implementations of DC log amp ICs are straightforward, and can be optimized with minimal effort. Calibration procedures can enhance log-amp performance, but are not necessary in all applications.

Timekeeping Accuracy, Automatic and Affordable

Electronic timekeeping has always lacked a high level of accuracy due to the inferior characteristics of the quartz crystal over temperature. Many different techniques have been applied to improve the accuracy provided by a 32.768kHz quartz crystal. This article describes a highly integrated device that provides unparalleled timekeeping accuracy at a price point comparable to an uncalibrated stand-alone real-time clock (RTC). This device makes current accuracy-improvement techniques obsolete, and makes accurate timekeeping the standard rather than a luxury.

<div align="center">

"**You may delay, but time will not.**"

- Benjamin Franklin

</div>

If Benjamin Franklin had to use a quartz crystal and an RTC to maintain the time of day, he may have rethought his statement. The inaccuracy of the crystal over temperature usually makes time appear to delay (or, occasionally, to move quicker).

An RTC with a 32.768kHz quartz tuning-fork crystal oscillator is the standard timekeeping reference for most electronic applications. The RTC maintains the time and date by counting seconds, which requires a 1Hz clock signal derived from the 32.768kHz crystal oscillator. The current time and date information is stored in a set of registers, which is accessed through a communication interface.

The Problem

There is nothing inherently wrong with using an RTC for timekeeping. However, the time is only as accurate as the reference used. Unfortunately, the typical 32.768kHz tuning-fork crystal does not provide much accuracy over a wide temperature range. Due to its parabolic nature over temperature (Figure 1), this accuracy is typically ±20ppm at room temperature (+25°C). This is the equivalent of gaining or losing 1.7 seconds of time each day, or 10.34 minutes per year. As **Figure 1** shows, accuracy decreases even further at more extreme high and low temperatures. The typical accuracy at these temperatures is much worse than 150ppm, the equivalent of losing almost 13.0 seconds of time each day, or over 1.3 hours per year.

The frequency deviation (Δf) of a typical crystal at a specific frequency and temperature is:

$$\Delta f/f = k\,(T - T_o)^2 + f_o$$

where f is the nominal crystal frequency, k is the curvature constant, T is the temperature, T_o is the turnover temperature, and f_o is the frequency deviation at room temperature.

An analysis of this equation reveals only three variables that control each crystal's frequency response over temperature. These are the curvature constant, turnover temperature, and room-temperature frequency deviation. The curvature constant has the most effect on the parabolic nature of the frequency deviation over temperature, but this constant has a very small deviation. Different turnover temperatures shift the deviation curve left or right; different frequency deviations at room temperature shift the curve up or down.

Various Solutions

For applications that demand accuracy in timekeeping, there have been limited options available to improve upon the crystal inaccuracies. Applications can improve timekeeping accuracy through crystal screening, integrated crystals, calibration registers, or temperature-compensated crystal oscillators.

Crystal Screening

One option to improve timekeeping is to have a supplier provide crystals that fall within a specified range of room-temperature accuracy. This requires that the supplier analyze each crystal's frequency deviation at room temperature before shipment, a screening process that obviously adds to the cost of the crystal. This method has no effect on the parabolic nature of the crystal's accuracy curve.

By using such a screening process, a crystal manufacturer could provide a subset of crystals that improve the room temperature accuracy from ±20ppm to ±10ppm or ±5ppm. These "improved" crystals would still suffer from large inaccuracies at high and low temperatures. Depending on
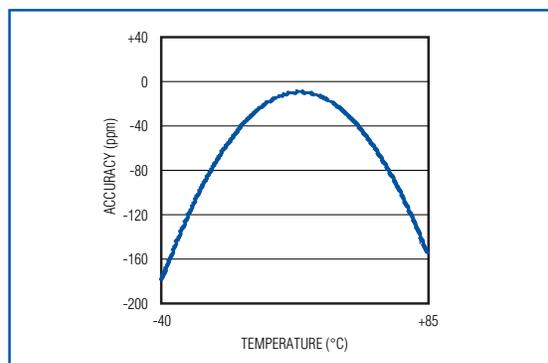
Figure 1. Temperature vs. accuracy is shown for a typical 32.768kHz tuning-fork crystal.

the level of accuracy and load capacitance required, there would also be a yield loss. This could result in an insufficient quantity of acceptable crystals.

A manufacturer can also control the crystal turnover temperature by the angle at which the crystal is originally cut, but this is impractical and costly. Crystal manufacturers use many automated processes, but still struggle to keep up with demand. The probability of inducing a manufacturer to interrupt a manufacturing sequence for a nonstandard part is low.

Integrated Crystals

Taking the crystal-screening process one step further, some companies include the tuning-fork crystal in the same package as the timekeeping device, which shifts the burden of providing crystals to the device manufacturer. Providing an integrated crystal reduces the designer's workload by eliminating crystal procurement issues. This also alleviates concerns about matching the crystal parameters with the timekeeping device requirements, and reduces printed circuit board (PC board) layout issues.

Companies that are not vertically integrated do not have the ability to measure or trim the crystal parameters. These companies purchase crystals from a supplier and assemble the die and crystal into a single package. No accuracy improvement is expected with this option. Dallas Semiconductor provides this type of integrated device with the DS1337C, DS1338C, DS1339C, DS1340C, and DS1374C. These are excellent RTCs for applications that do not require a high level of accuracy.

Other companies that manufacture their own crystals have the ability to place a crystal blank, which is the unpackaged quartz, in a smaller, hermetically sealed package and trim the blank to meet certain accuracy requirements. This method does not change the parabolic curve, and only provides a small accuracy improvement at room temperature. The improvements at high and low temperatures are negligible. The downside of this method is that the ceramic package and crystal trimming add cost to the total solution.

Temperature Compensation

To achieve timekeeping accuracy over a wide temperature range, some form of temperature compensation is required. Temperature compensation requires periodic measurement of temperature and subsequent adjustment of either the crystal loading or the clock source according to the measured temperature.

Temperature compensation can be accomplished in one of two ways. The first option is to develop a temperature-compensation algorithm by using a temperature sensor with a timekeeping device that provides some form of analog or digital clock calibration. This method usually requires an extensive development and calibration invest-

ment. The second option is to use an off-the-shelf, temperature-compensated crystal oscillator (TCXO) as the clock source for an RTC.

Calibration Register

Some RTCs, like the DS1340, provide a digital calibration register that can be used to periodically adjust the time of day in discrete amounts. This method does not attempt to alter the crystal behavior at all, but instead, periodically adjusts time according to the expected frequency deviation at a specified temperature. The effect is to move the 32.768kHz parabolic curve up or down to approach 0.0ppm accuracy at a desired temperature. This is accomplished by adding or subtracting clock cycles from the oscillator's divider chain. The number of clock pulses removed (subtracted for negative calibration) or inserted (added for positive calibration) is set by the value in the calibration register. By adding clock pulses, time is sped up (the crystal curve moves up); by subtracting clock pulses, time is slowed down (the crystal curve moves down). **Figure 2** shows how the typical crystal curve is shifted upward until the accuracy approaches 0.0ppm. In this example, the measured temperature is +55°C.

An RTC with a calibration register can be combined with a temperature sensor to achieve accuracy levels from -2.034ppm to +4.068ppm at one specific temperature. The total adjustment range is from -126ppm to +63ppm; so, at extreme high and low temperatures, the curve cannot be adjusted enough to achieve 0.0ppm. Processor overhead is required to periodically measure the temperature, calculate the new calibration register value, and adjust the appropriate RTC register.

The major difficulty with this method is the required factory-calibration effort. Because each crystal behaves differently, a custom calibration table for the desired temperature range is essential for each timekeeping device. The amount of manpower and time needed for such an effort could become cumbersome. Some amount of nonvolatile memory is also necessary to store the calibration data, adding to the overall cost. In addition, the compensation values do not compensate for the inevitable crystal aging, which can approach ±3ppm for the first year alone.

While the calibration register method does not provide automatic adjustments as temperature changes, it still provides an incremental improvement in accuracy. But is this method worth the cost?

Temperature-Compensated Crystal Oscillator

Another option that greatly improves timekeeping accuracy uses a 32.768kHz temperature-compensated crystal oscillator (TCXO), like the DS32kHz, as the clock source for a stand-alone RTC. These TCXOs are factory calibrated, and can provide accuracy up to ±7.5ppm over
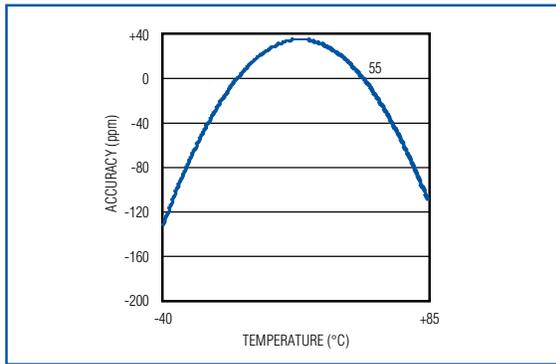
Figure 2. A typical crystal curve is shifted upward until the accuracy approaches 0.0ppm.

Figure 3. Using a TCXO will flatten a crystal curve.

the industrial temperature range (-40°C to +85°C). In effect, a TCXO flattens the parabolic nature of the crystal curve over temperature (**Figure 3**).

A TCXO includes an integrated sensor to periodically measure device temperature. The measurement is used to access a lookup table, whose output is used to calculate and apply a load-capacitance value for the integrated 32.768kHz crystal to achieve 0.0ppm accuracy. The lookup table exists on the TCXO and requires no external inputs.

When a crystal is manufactured, it is optimized for a particular load capacitance, which is specified in the data sheet. If the actual load capacitance does not match the specification, the result is a deviation from the nominal frequency. This fact is utilized by a TCXO to improve accuracy. If the amount of frequency deviation at each temperature is known for a specific crystal, the TCXO can adjust the load capacitance to offset the temperature-dependent frequency deviation.

The advantage of using an off-the-shelf TCXO is that no algorithm development or factory calibration is required. The drawbacks are the additional cost and PC-board space required by a multichip solution.

The Most Accurate Solution—
RTC/TCXO/Crystal Integration

The ideal accurate timekeeping device would integrate an RTC, a TCXO, and a quartz crystal into a single package. The DS3231S is such a device. It provides unparalleled accuracy of ±2.0ppm from 0°C to +40°C, which is the equivalent of just over ±1.0 minute each year. Accuracy from -40°C to 0°C and +40°C to +85°C is ±3.5ppm, which is the equivalent of ±1.8 minutes each year. The worst-case accuracy of this device is displayed in **Figure 4**. As mentioned in the previous section, the integrated TCXO flattens the parabolic nature of the crystal curve over temperature.
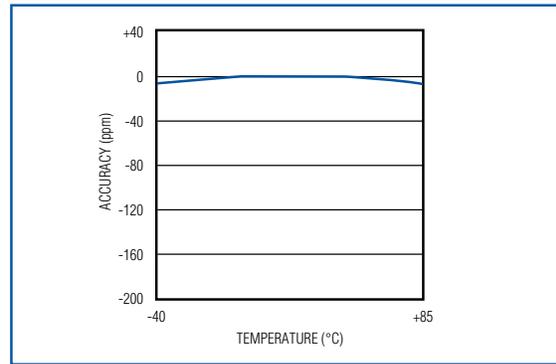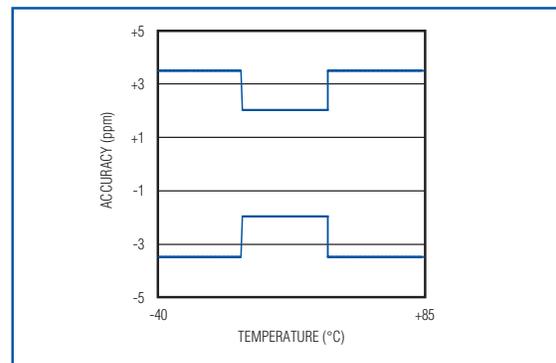
Figure 4. Worst-case accuracy is shown for an integrated RTC/TCXO/ quartz crystal, the DS3231S.

Like the previous TCXO solution in the DS32kHz, the integrated DS3231S is factory calibrated and requires no customer calibration or development effort. The single-package solution combines the same functionality in a smaller area and at a reduced cost.

Unlike a stand-alone TCXO, the integrated device register is accessible through the serial-interface port. An on-chip aging register adjusts for load capacitance and temperature compensation. This allows an application to also compensate for accuracy lost due to crystal aging.

Summary

Before the integration of a TCXO, an RTC, and a 32.768kHz crystal into one device, applications that required timekeeping accuracy had limited options. All options available required some combination of development effort, factory calibration, and additional cost. With the advent of single-package TCXO/RTC/crystal integration in the DS3231S, ±2.0 timekeeping accuracy is no longer a luxury, but available for all applications!

Charging Batteries from USB

The Universal Serial Bus (USB) port is a bidirectional data port with power and ground. Peripherals of all types can be connected to the USB, including external drives, memory devices, keyboards, mice, wireless interfaces, video and still cameras, MP3 players, and countless other electronics. Many of these devices are battery powered, some with internal batteries. The widespread availability of USB presents unique opportunities, as well as challenges, for battery-charging designs. This article describes how to interface a simple battery charger to a USB power source. This review of USB power bus characteristics includes voltage, current limits, inrush current, connectors, and cabling. An overview of nickel metal hydride (NiMH) and lithium battery technologies, charging methods, and charge-termination techniques is given. A complete example circuit for smart-charging NiMH cells from a USB port is presented, along with charging data.

USB Characteristics

The ubiquitous USB bus offers great opportunities as a power source for all types of low-power electronics. The bus' power source is isolated from power mains and is relatively well regulated. However, there are limitations on available current and potential interactions between the load and the host or power source.

The USB port consists of a 90Ω bidirectional differential shielded twisted pair, V_{BUS} (+5V power), and ground. These four wires are shielded with an inner shield of solid aluminum and a stranded outer shield. The USB specification is currently at revision 2.0, and copies are available free of charge from the USB organization (www.USB.org). Full compliance with the specification requires bidirectional communication between the device and the host through a function controller. The specification defines a unit load as 100mA (max). The maximum current that any device is allowed to draw is five unit loads.

USB ports are classified as either low-power ports, which supply up to one unit load, or high-power ports which supply up to five unit loads. When devices are first connected to the USB port, an enumeration process identifies the device to determine its load requirements. During this time, the device is only allowed to draw one unit load from the host. After the enumeration process, higher powered devices are permitted to draw higher current if the power-management software in the host allows it.

Some host systems (including downstream USB hubs) have current limiting either through fuses or active current sensors. If a USB device presents a high current (over one unit) load to the USB port without enumerating, it can cause a detectable overcurrent condition that could shut down one or more of the USB ports in use. Many commercially available USB devices, including stand-alone battery chargers, draw over 100mA without a function controller to handle the enumeration process; they run the risk of causing problems for the host under the wrong circumstances. For instance, if a device drawing 500mA is plugged into a bus-powered USB hub, it could overload both the hub port and the host port if it is not properly enumerated.

Further complications arise when the host operating system is using advanced power management, especially for notebook computers, and is expecting the port current to be extremely low. In some power-saving modes, the computer issues suspend commands to USB devices, which are then expected to go to a low-power mode. It is always a good idea to include a function controller to communicate with the host even with low-power devices.

The USB 2.0 specification is quite thorough and specifies power quality, connector construction, cable materials, allowable voltage drops, and inrush current. Low-current and high-current ports have different power-quality specifications. These are determined primarily by the voltage

Table 1. USB 2.0 Specification Power-Quality Standards

Parameter	Requirement
DC voltage, high-power port	4.75V to 5.25V
DC voltage, low-power port	4.40V to 5.25V
Maximum quiescent current (low power, suspend mode)	500μA
Maximum quiescent current (high power, suspend mode)	2500μA
Maximum allowable input capacitance (load side)	10μF
Maximum required output capacitance (host side)	120μF
Maximum allowable inrush charge into load	50μC

*These specifications apply to the pins of the host or hub port connector on the upstream side. Additional I x R drops due to cables and connectors must be counted separately.

drop in the connectors and cabling between the host and the load, including voltage drop across a USB-powered hub. A host, such as a computer or self-powered USB hub, has high-current ports capable of supporting up to 500mA. Lower current ports are found on passive, bus-powered USB hubs. Table 1 gives the allowable tolerances for the voltage at the pins on the upstream (source) side of the USB port for high- and low-current ports.

In hosts that are compatible with the USB 2.0 specification, the upstream side of a high-power port is provided with 120µF of low-ESR capacitance. The input capacitance of attached USB devices is limited to 10µF, and the total allowable charge drawn from the host (or powered hub) during an initial load connection is 50µC. Thus, when a new device is connected to a USB port, the transient voltage drop at the upstream port is less than half a volt. If more capacitance is required for correct operation of the load, it must be provided with an inrush current limiter to charge the larger capacitance at no greater than 100mA.

The allowable DC voltage drops for a USB port having a bus-powered USB hub with low-powered functions attached are shown in **Figure 1**. A high-power load connected to an unpowered hub has larger voltage drops than shown in Figure 1 and can overload the bus.

Battery-Charging Requirements

Single-Cell Lithium Ion and Lithium Polymer

Present-day lithium chemistries are typically 4.1V to 4.2V when the cells are charged to their maximum-rated capacity. Newer, higher capacity cells are being marketed with voltages in the 4.3V to 4.4V range. Typical prismatic lithium ion (Li+) and lithium polymer (Li-Poly) have capacities of 600mA to 1400mA.

The preferred charge profile for both Li+ and Li-Poly cells is to start the charge with a constant charge current until the cell voltage reaches the rated voltage. When this occurs, the charger then regulates the voltage across the cell. These two regulation states are called constant-current (CC) and constant-voltage (CV) charging; therefore, this type of charger is usually referred to as a CCCV charger. When the CCCV charger is in CV mode, the current into the cell begins to drop. For typical charge rates, between 0.5C and 1.5C, the transition between CC and CV mode occurs when the cell has accepted about 80% to 90% of its full-charge capacity. Once the charger is in CV charging mode, it then monitors the cell current; when it reaches a low threshold (milliamps or tens of milliamps), the charger terminates charge. A typical charge profile for a lithium chemistry battery is shown in **Figure 2**.

The USB voltage drops shown in Figure 1 show that a low-powered port, on the downstream side of a port-powered hub, would barely have enough headroom to charge a cell to 4.2V. Small amounts of additional resistance in the charge path could prevent proper charge.

Li+ and Li-Poly cells should be charged at moderate temperatures. Manufacturers' recommended maximum charging temperatures are typically in the +45°C to +55°C range, with allowable discharge temperatures about 10°C higher. The materials used in these cells are highly reactive and can ignite if the cell temperature exceeds +70°C. Chargers for lithium chemistry cells should be designed with a thermal cutoff circuit that monitors cell temperature and terminates charging if the cell temperature exceeds manufacturer-recommended maximum charging temperatures.
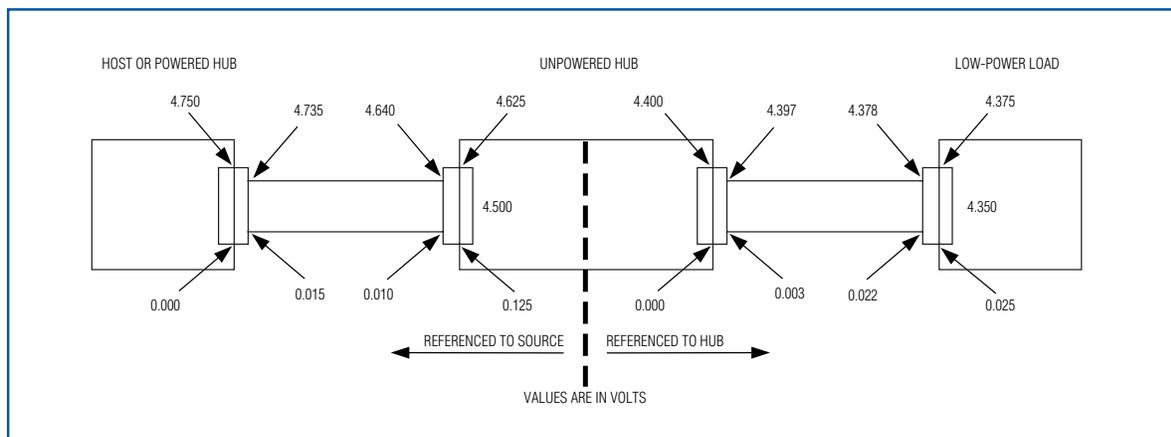
Figure 1. Drops larger than these allowable DC voltage drops from host to low-power load can overload the bus.
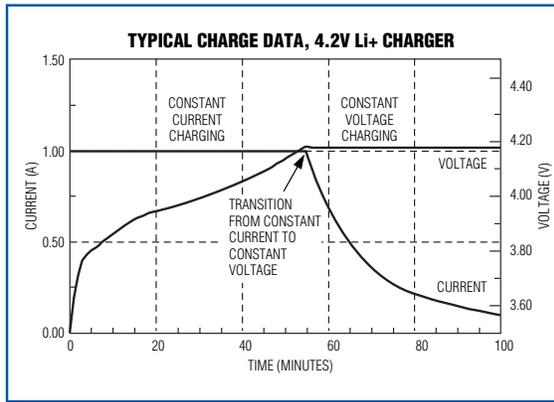
Figure 2. Typical results are shown for a Li+ battery charged using a CCCV charger.

Figure 3. A DS2712 charge controller is used to charge a NiMH cell.

Nickel Metal Hydride Cells

NiMH cells are heavier and have less energy density than lithium-based cells. Historically, they have been less expensive than lithium, but the price gap has recently been shrinking. NiMH cells are available in standard sizes and are a direct replacement for alkaline cells in most applications. Nominally, each cell is 1.2V, and when fully charged they are as high as 1.5V.

NiMH batteries are usually charged with a constant-current source. As they reach the fully charged state, an exothermic chemical reaction occurs that causes the battery temperature to increase and the terminal voltage to decrease. Either the rate of rise of the battery temperature or a negative voltage change can be detected and used to terminate the charge. These termination methods are known as dT/dt and -ΔV, respectively. At very low charge rates, the dT/dt and -ΔV effects become less pronounced and can be difficult to detect accurately. The dT/dt and -ΔV responses begin as the cell starts to become overcharged. Continued charging beyond this point can damage the cell.

Termination detection at charge rates above C/3 is much easier than at low charge rates. The temperature rise is about 1°C/min, and the -ΔV response is stronger than at lower rates. After fast-charge terminates, an additional charging period at a reduced current is recommended to top off the cell (top-off charge). When the top-off charge cycle is done, a trickle-charge current of C/20 or C/30 counters the effects of self-discharge and maintains the battery in a fully charged state. A graph showing the cell voltage during a charging cycle of a partially charged NiMH cell using a DS2712 NiMH charger is shown in Figure 3. In this graph, the upper curve is data taken while current is being driven into the battery; the lower curve shows data taken with the charge current source turned off. In the DS2712, this difference in voltage is used to distinguish between NiMH cells and alkaline cells. If an alkaline cell is detected, the DS2712 will not charge it.
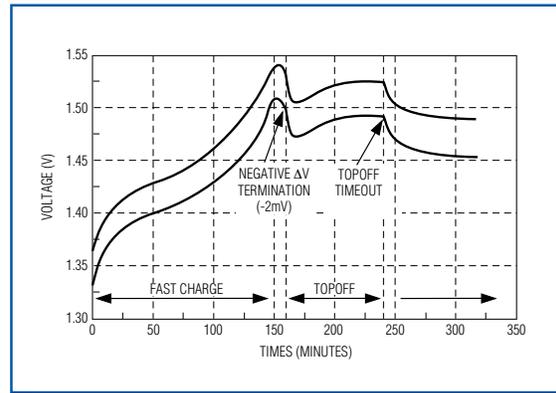
Switching vs. Linear

The USB 2.0 specification allows up to 100mA from a low-power port and 500mA from a high-power port. If a linear pass element is used to regulate charge current to the battery, these are the maximum allowable charge currents. The power dissipated in the linear pass element (**Figure 4**) is $P = V_Q \times I_{BATT}$. This causes power dissipation in the pass transistor and can require the use of a heat sink to prevent overheating.

For a nominal input voltage of 5V, the pass element dissipates a varying amount of power based on the type and number of cells, as well as the battery voltage.

Figure 5 shows the calculated power dissipation for a linear USB charger with NiMH cells at a nominal input voltage of 5.0V. With single-cell charging, a linear charger is only about 30% efficient; a two-cell charger is 60% efficient. Charging a single cell at 500mA results in up to 2W of power dissipation. This amount of power typically requires a heatsink. At 2W of dissipation, a +20°C/W heatsink heats up to about +65°C from an
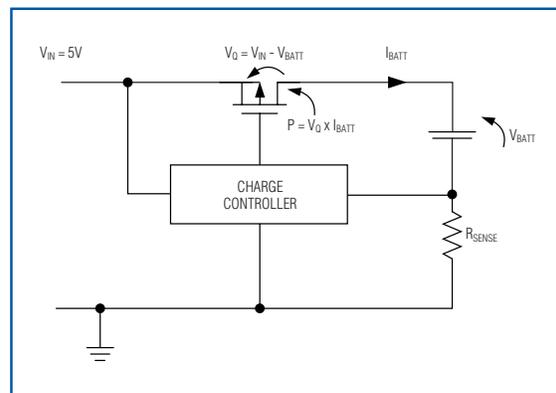
Figure 4. Power dissipation is the battery current times the voltage across the pass transistor.

POWER DISSIPATION vs. CELL VOLTAGE FOR NiMH LINEAR CHARGER
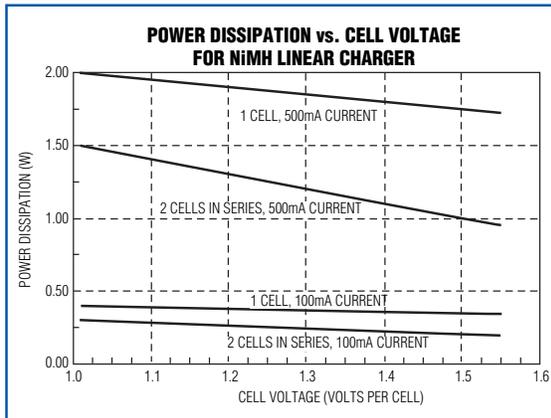
Figure 5. Power dissipates in a linear pass element, with NiMH cells being dissipated from a USB port at 5.0V input voltage.

CHARGE TIME FOR LINEAR vs. SWITCHING CHARGER FOR SINGLE-CELL NiMH

Figure 6. Charge times differ for a linear charger vs. a switching charger when charging a single-cell NiMH battery.

ambient temperature of +25°C, and it needs exposure to free air for full performance. Inside an enclosure with still air, it will get much hotter.

Several problems are solved by using a switching-regulator-based charger. First, the cells can be charged faster and at a higher current than with a linear charger (**Figure 6**). Because less power is lost as heat, the thermal management problem is reduced. Also, the charger is more reliable, as it runs cooler.

The calculated values in Figure 6 are based on charging from a high-power USB port at about 90% of the 500mA maximum allowable current. The switching regulator in this example is assumed to be a nonsynchronous buck converter with 77% efficiency.

Circuit Example

The circuit shown in **Figure 7** is a switch-mode buck regulator for charging a single NiMH cell. It uses a DS2712 charge controller to regulate charging current and terminate charge. The charge controller monitors temperature, battery voltage, and battery current. If the temperature is above +45°C or below 0°C, the controller does not begin charging the battery.

In Figure 7, Q1 is the switching power transistor for the step-down charger; L1 is the smoothing inductor; and D1 is the freewheeling, or catch, diode. Input C1 is a 10μF, very low-ESR ceramic filter capacitor. Substituting tantalum or other electrolytic capacitors for C1 can adversely influence charger performance. R7 is the current-sense resistor for the sense amplifier of the current regulator. The reference voltage for the DS2712 is 0.125V with an hysteretic component of 24mV. Closed-loop, switch-mode current control is provided through CSOUT. Gate drive to Q1 is enabled when the gate of Q2 is pulled low by the charge control pin, CC1. Both Q1 and Q2 are low-V_t (gate-source threshold voltage) pMOSFETs. When CC1 and CSOUT are both low, the drain to source voltage of Q2 is slightly over one V_t. This voltage, plus the forward voltage drop of CSOUT, establishes the available switching-gate drive voltage for Q1.
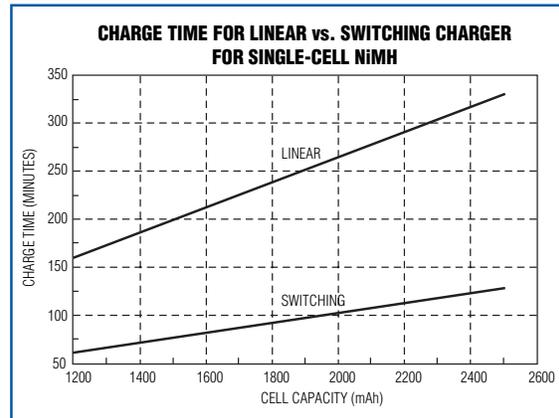
When CC1 is low, it enables closed-loop control of the current. Startup switching waveforms are shown in **Figure 8**. The top waveform is the voltage across the 0.125ø current-sense resistor, and the bottom waveform is the Q1 drain to GND voltage. Initially, current ramps up in the inductor while Q1 is ON (CC1 and CSOUT both low). When the current reaches 0.125V, CSOUT goes high. Also, current ramps down until the voltage on the current-sense resistor reaches about 0.1V, whereupon CSOUT goes low again. This process continues as long as CC1 is low.

The DS2712's internal state machine controls the gating action of CC1. At the beginning of charge, the DS2712 performs cell qualification tests to make sure that the cell voltage is between 1.0V to 1.65V, and to verify that the temperature is between 0°C and +45°C. If the voltage is below 1.0V, the DS2712 gates CC1 low at a duty factor of 0.125, which slow-charges the cell to prevent damage. Once the cell voltage exceeds 1.0V, the state machine transitions into fast-charge. The fast-charge duty factor is 31/32, or about 97%. The "skipped" pulse is used to perform impedance tests on the cell to make sure that a high-impedance cell, such as an alkaline cell, has not been installed in the charger. Fast-charge continues until a -ΔV of -2mV is detected. If no -ΔV is detected, fast-charge continues until the fast-charge timer expires or until an overtemperature or overvoltage fault condition (including impedance failure) is detected. When the fast-charge is done (either due to -¢V or expiration of the fast-charge timer) the DS2712 enters its timed top-off mode, with a duty factor of
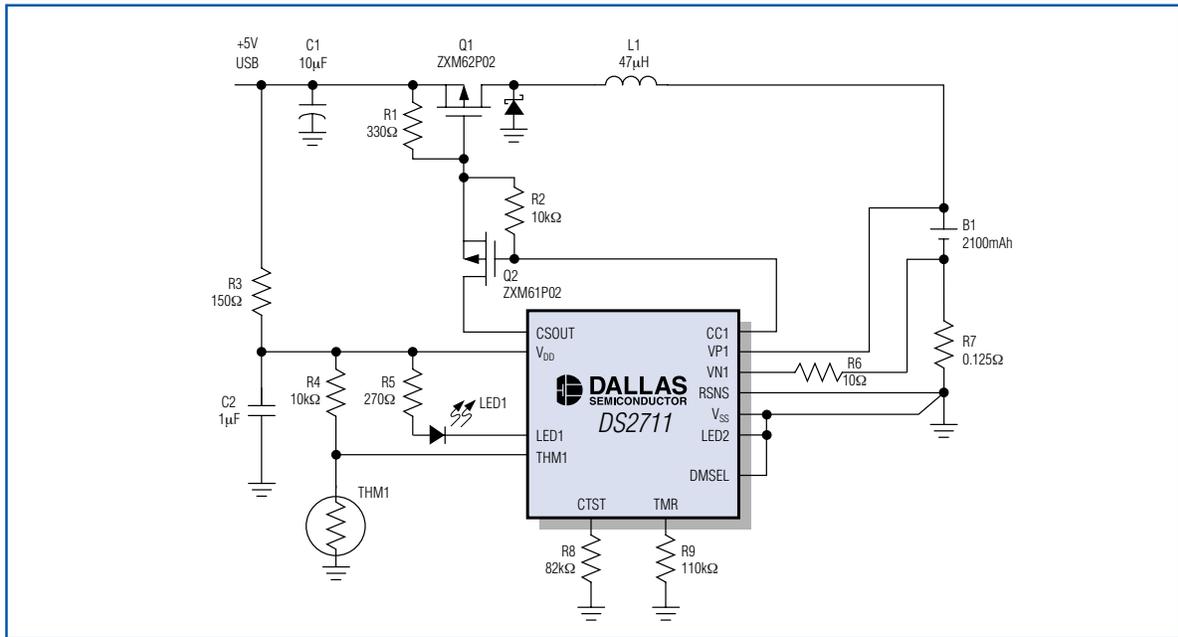
Figure 7. Schematic shows a single-cell NiMH battery being fast-charged from a USB port.

12.5% and a duration of half of the programmed fast-charge timeout. After top-off is complete, the charger goes into maintenance mode with a duty factor of 1/64, and remains in maintenance until the cell is removed or power is cycled.

The charger shown in Figure 7 fast-charges a 2100mAh NiMH cell from a high-power USB port in a little over two hours, with a full top-off charge achieved in about three hours. Current drawn from the port is 420mA. If enumeration with the host and high-current enable is required, an open-drain nMOSFET can be inserted in series between R9 and ground. If the MOSFET is off, TMR floats and the DS2712 is in its suspend state.

Summary

The USB port is an economical and practical power source for charging batteries for small consumer electronics. To be fully compliant with the USB 2.0 specification, loads connected to the USB port must be capable of bidirectional communication with the host. Loads must also comply with power-management requirements, including low-power modes and a means of allowing the host to determine when high power is drawn from a port. While partially compliant systems may operate compatibly with most USB hosts, they occasionally give unexpected results. A good understanding of USB requirements and expectations of the load are needed to make the right trade-off between full compliance and load complexity.
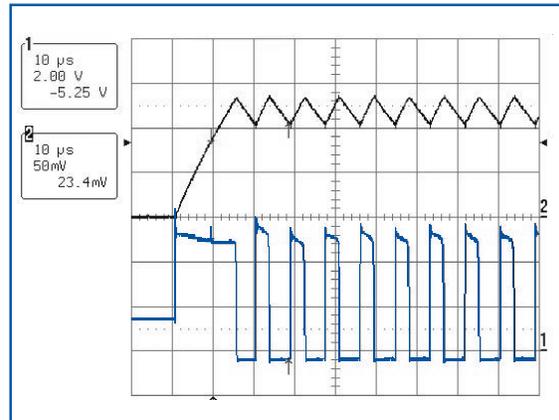
Figure 8. Startup waveforms are shown for a USB NiMH charger.

Digital Pots Achieve Voltage-to-Resistance Conversion

Voltage-to-resistance converters, sometimes needed in industrial controls and variable-bias circuits, can be difficult to implement. The simple approach of **Figure 1** builds such a converter using two digital potentiometers.

One digital potentiometer (U1) and an op amp (U3) form a digital track-and-hold circuit in which U1 adjusts its internal voltage-divider to make V_{WIPER} track V_{IN}. Wiper resistance is therefore proportional to V_{IN}. Because the digital inputs of the digital potentiometers (U1 and U2) are connected, U2's wiper position is the same as that of U1, and resistances between corresponding terminals are the same. Thus, the output resistance is proportional to V_{IN}, as required for voltage-to-resistance conversion.

Digital Track-and-Hold Operation

To track V_{IN}, the wiper (center tap) of U1 moves up or down as each clock pulse arrives. U3 compares the analog input (V_{IN}) with the wiper voltage (V_{WIPER}). If $V_{IN} > V_{WIPER}$, the comparator asserts logic-high, causing the wiper position to move up and increase the value of V_{WIPER}. V_{WIPER} keeps increasing until it is greater than V_{IN}; the comparator then toggles and drives the wiper position downward. On each clock cycle, the wiper moves up or down as required to track V_{IN}. Reference inputs for the voltage divider (V_H and V_L) set the input voltage range; if the V_{IN} range is 0V to 5VDC, set V_L = GND and V_H = 5VDC.

Due to the connected digital inputs and identical wiper positions of U1 and U2, applying a logic-low to the LOCK input allows the output resistance to change with V_{IN}; applying a logic-high holds the resistance value indefinitely.

LOCK may be connected permanently to ground but, in that case, the output resistance toggles continually between two consecutive states, even if V_{IN} is constant. For example, if the potentiometer is 10kΩ and the wiper is set to 5kΩ, the output resistance will toggle between 5kΩ and 5.3125kΩ on every clock cycle. If necessary, that effect can be filtered by connecting a capacitor to the output wiper. A clock frequency between 100Hz and 10kHz is acceptable.

Output resistance does not change instantly with V_{IN}, but takes a number of clock cycles to reach its final value. The number of cycles (32, max) depends on the initial wiper position and the input voltage.

If higher resolution is needed, substitute a 6- or 8-bit digital potentiometer for the 5-bit model in Figure 1. Note that the MAX5160 has a power-up reset that sets the wiper position to midscale, thereby allowing the two digital pots to synchronize to the same resistance. Choose a digital pot that has a known output-resistance value at power-up.

A similar design idea appeared in the June 7, 2004 issue of Electronic Design.
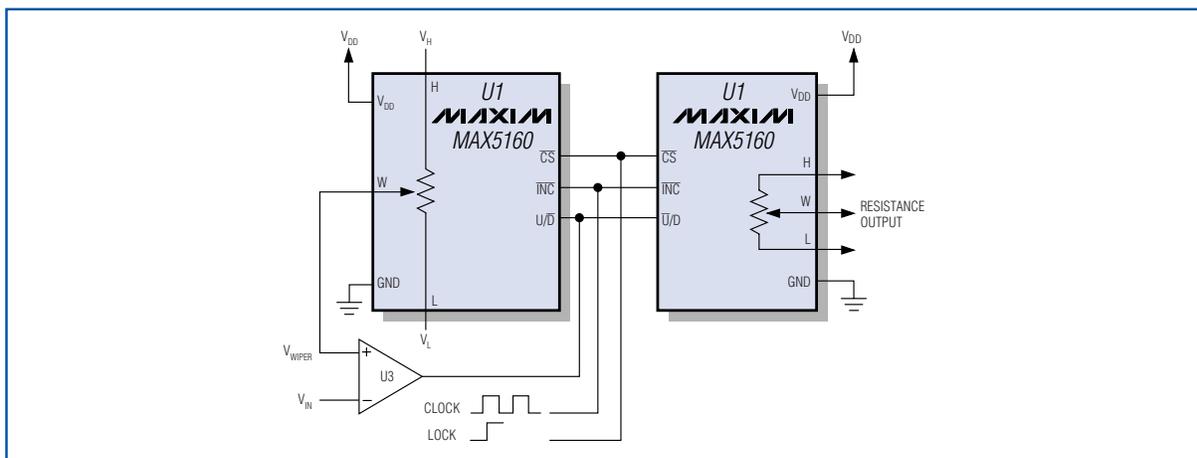
Figure 1. Two identical digital potentiometers (U1 and U2) enable this circuit to implement voltage-to-resistance conversion.

GPIO Expander Provides Charge Pump for Blue LEDs

Blue LEDs are becoming prevalent in the industry, and sometimes make their way onto control panels as colorful indicators. Their high forward voltage, however, runs counter to the downward trend in supply voltages. A panel that includes a 3.3V supply voltage may have a difficult time driving LEDs that sometimes require upwards of 3.5V for illumination. To drive multiple blue LEDs, you may require a charge pump or boost regulator. However, for a single LED you can use a few discrete components and an IC that may already reside on the PC board.

$I^2C^\ddagger$ GPIO expanders are often designed onto control panels to accommodate their remoteness. (Routing a serial bus in place of 8 or 16 parallel wires is much more convenient; the smaller cable and connector save cost.) You can drive the blue LED by using two GPIO lines. Thus, a GPIO expander with built-in pulse-width modulation (PWM) and 50mA sink capability on each I/O forms an inexpensive discrete-component charge pump (**Figure 1**). While the charge pump is limited by the oscillator's low switching frequency (2kHz), this approach works because the LED needs only about 10mA.

Adding a capacitor and Schottky diode to one of the eight outputs (P4) and placing that output in PWM mode (at about 50% duty cycle) generates enough voltage to drive the LED. A second output (P6) turns the LED on and off. On-off capability is needed if the LED forward voltage is near or below the power-supply voltage, because current can flow though CR1 even when the oscillator is turned off.

Capacitor C1 charges to approximately 3V when P4 goes low. When P4 goes high, the capacitor voltage adds to the power-supply voltage and delivers current to the LED through R1. The LED current discharges C1 during P4's off-time, so the capacitor must fully recharge on the following cycle.

Using a 3.3V supply voltage, the circuit delivers about 15mA (typ), which produces a brightly glowing blue LED. At 3.0V, the current drops to about 10mA, and at 2.0V it drops to about 3mA. By altering the values of R1 and C1, you can modify the current to suit a particular LED. Note that you must stop the PWM before P6 is turned off. The charge pump places up to twice the supply voltage on P6, which comes precariously close to that output's maximum application voltage rating.

‡*Purchase of I^2C components from Maxim Integrated Products, Inc., or one of its sublicensed Associated Companies, conveys a license under the Philips I^2C Patent Rights to use these components in an I^2C system, provided that the system conforms to the I^2C Standard Specification defined by Philips.*

A similar design idea appeared in the January 17, 2005 issue of EE Times.
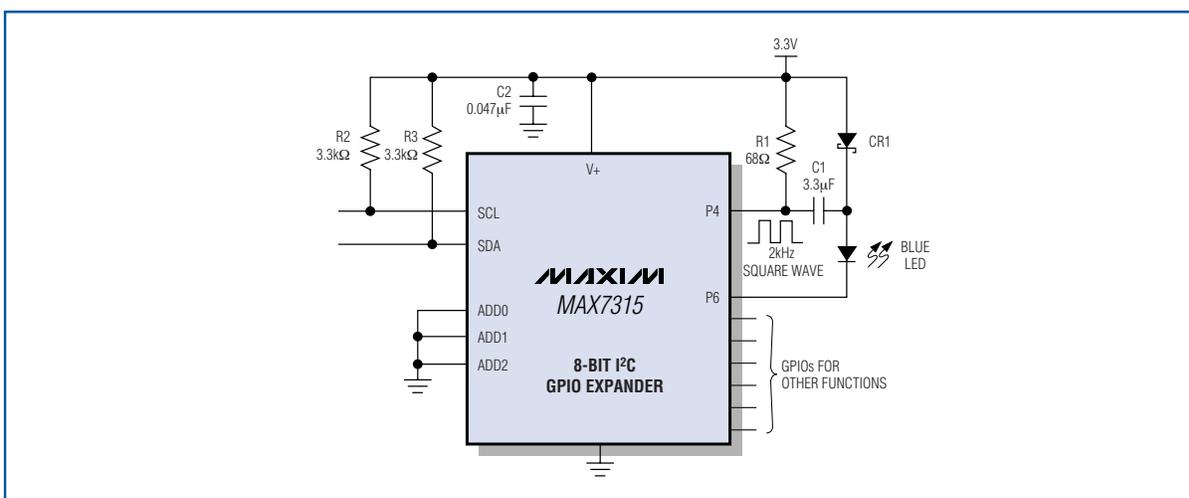
Figure 1. While performing other functions, this GPIO expander can easily drive a blue LED as well.