DALLAS SEMICONDUCTOR MAXIM

Engineering Journal

Volume Fifty-Four

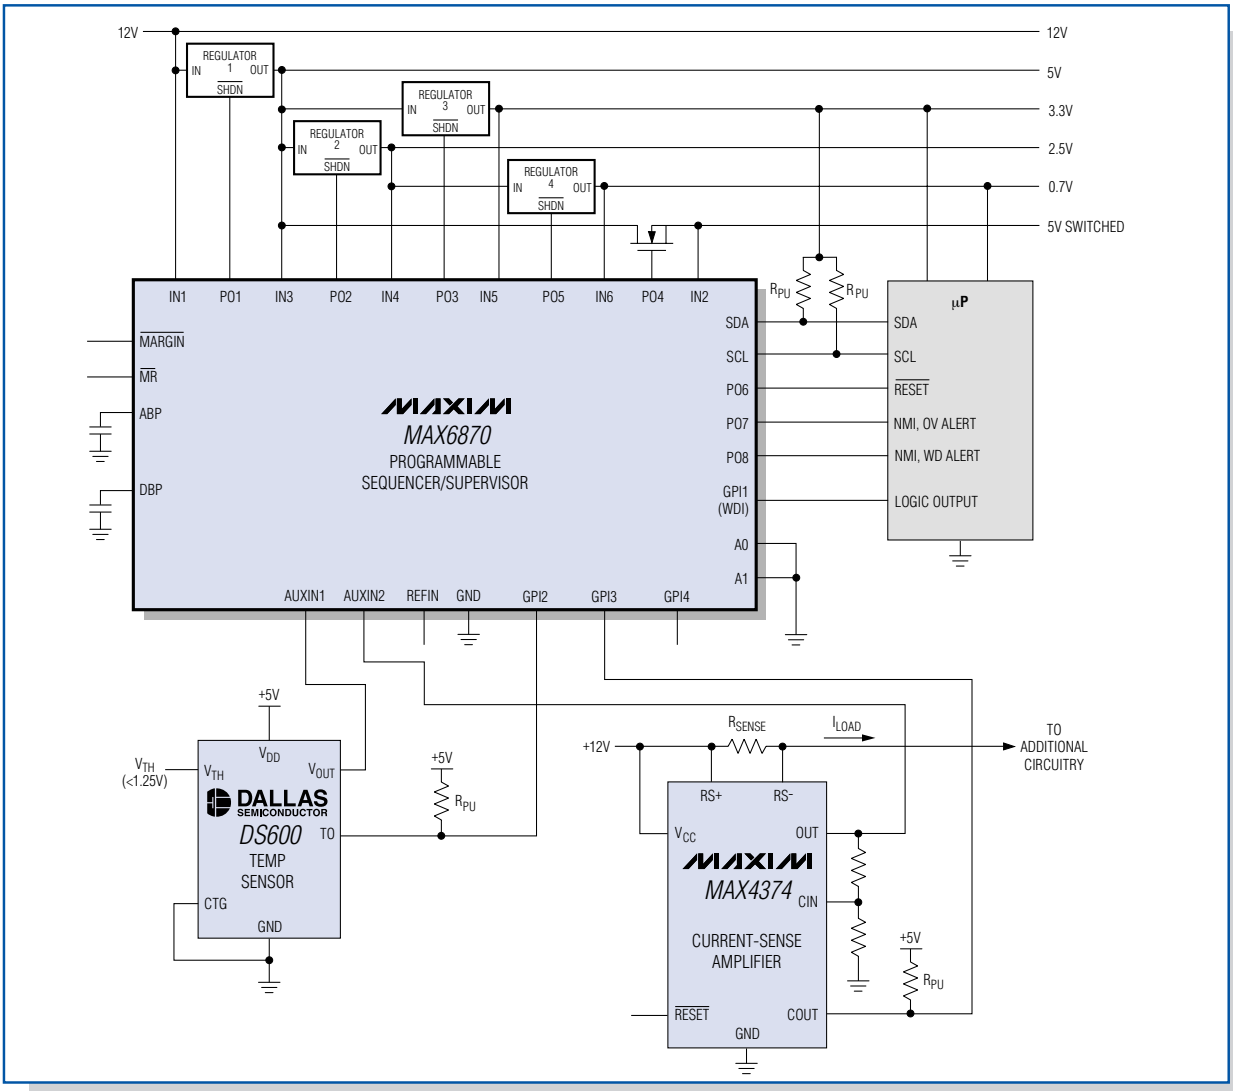

A programmable system-management device provides a flexible means for monitoring and sequencing supply voltages. (For more information on the monitoring/sequencing abilities of the MAX6870, see article inside, page 15.)

News Brief

MAXIM REPORTS REVENUES AND EARNINGS FOR THE FIRST QUARTER OF FISCAL 2005 AND INCREASES QUARTERLY DIVIDEND

Maxim Integrated Products, Inc., (MXIM) reported net revenues of $435.1 million for its fiscal first quarter ending September 25, 2004, a 40.3% increase over the $310.2 million reported for the first quarter of fiscal 2004. Diluted earnings per share were $0.42 for the first quarter, a 68.0% increase over the $0.25 reported for the same period a year ago. Net revenues were 3.4% above the $421.0 million reported for the fourth quarter of fiscal 2004, and diluted earnings per share were 16.7% above the $0.36 reported for the fourth quarter of fiscal 2004. Net income for the first quarter was $144.5 million or 33.2% of net revenues, a 65.4% increase over the $87.4 million reported for the first quarter of last year and a 15.9% increase over the $124.7 million reported for the fourth quarter. Operating income for the Company was $210.7 million or 48.4% of net revenues for the first quarter, compared to the $125.7 million or 40.5% of net revenues reported for the same period a year ago and the $181.2 million or 43.1% of net revenues reported for the fourth quarter of fiscal 2004.

During the quarter, cash and short-term investments increased $88.9 million after the Company repurchased 1.3 million shares of its common stock for $58.5 million, paid dividends of $25.9 million, and acquired $66.3 million in capital equipment. Accounts receivable increased $3.3 million in the first quarter to $200.5 million due to the increase in net revenues, and inventories increased $18.0 million to $135.8 million.

Research and development expense was $79.1 million or 18.2% of net revenues in the first quarter, compared to $87.8 million or 20.8% of net revenues in the fourth quarter of fiscal 2004. The decrease in research and development expense in the first quarter was due to the $9.5 million one-time fourth quarter bonus, which was offset by the hiring of additional engineers to support the Company's new product development efforts. Selling, general and administrative expenses decreased from $26.4 million in the fourth quarter, or 6.3% of net revenues, to $25.1 million in the first quarter, or 5.8% of net revenues, primarily as a result of the $2.1 million one-time fourth quarter bonus.

First quarter net bookings were approximately $377 million, compared to the fourth quarter's level of $535 million. Turns orders received in the quarter were approximately $117 million, compared to the $170 million 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). First quarter ending backlog shippable within the next 12 months was approximately $458 million, including approximately $377 million requested for shipment in the second quarter of fiscal 2005. The Company's fourth quarter ending backlog shippable within the next 12 months was approximately $529 million, including approximately $428 million that was requested for shipment in the first quarter of fiscal 2005.

Jack Gifford, Chairman, President, and Chief Executive Officer, commented: "As expected, first quarter bookings were significantly below the fourth quarter's level, as customers are apparently adjusting their inventory levels. During fiscal 2004, we booked approximately $250 million more than our estimate of consumption of our products by our customers, and we expected booking levels to make an adjustment. We believe that the fiscal 2004 third and fourth quarter booking levels reflected a desire by many customers to avoid supply chain disruptions in a longer-lead-time environment. Now that lead times are shortening, customers appear to be reverting to more of a 'just in time' approach to order placement."

Mr. Gifford continued: "Our San Antonio fab is ramping, as planned, to do 2 million moves at that facility as soon as possible." Mr. Gifford added: "Analysts and investors often view a company's free cash flow per share as one of the best indicators of value, and we agree."

Mr. Gifford concluded: "The Company's Board of Directors has declared an increased quarterly cash dividend of $0.10 per share. Payment will be made on November 30, 2004 to stockholders of record on November 15, 2004."

Testing a power supply for line and load transients

Power-supply specifications for line and load transients describe the response of a power supply to abrupt changes in line voltage and load current. By monitoring the supply as it attempts to maintain regulation in the presence of such transients, you can observe any tendency toward output overshoot or oscillation.

Line and load transients are step functions that inject disturbances into the power supply. A load transient injects disturbance into the output by stepping the load current, and a line transient does so by stepping the line voltage. The supply's output response demonstrates its ability to attenuate the various frequency components of a line or load step. The following discussion describes these tests and the information they provide about a power supply.

How are line and load transients generated?

Besides generating artificial disturbances, the test setup for line or load transients must closely reproduce actual operating conditions for the power supply. The test setup must generate steps in line voltage and load current that are fast with respect to the controller's response time, while simulating the supply's operating conditions or end-user specifications. This task requires special attention to layout and component selection. Otherwise, the parasitic inductance, resistance, and capacitance of PC-board traces and components can prevent a fast step response by limiting the required slew rate.

The rise time of the transient should simulate conditions seen by the power supply during operation. If those conditions are not known, the rise time should be fast enough to pull the output out of regulation, thereby eliciting a maximum response from the controller. To do that, the test transient must be faster than the controller's response time. That is, the step transition for a switching converter should occur in less than half a switching period to ensure that the controller is pulled out of regulation.

Knowing the rise time and magnitude of the voltage or current step defines how critical the parasitic inductance, resistance, and capacitance will be in the test setup. For example, consider the need to apply a 10A step with 200ns rise time at the output (V_{OUT} = 1.8V). If the path between output capacitor and load includes 100nH of parasitic inductance, then the fastest possible rise time is 555ns. Parasitic inductance is clearly critical in that case. For a 10A, 10µs step, on the other hand, the parasitic inductance would account for only 5% of the total rise time.

Line transients

Fast line transients can be generated using two n-channel MOSFETs (with low R_{DS_ON}) switching between two DC sources (**Figure 1**). During interval A, Q1 pulls the power-supply input to 5V, while Q2 disconnects the 3V source. During interval B, Q1 disconnects the 5V source and Q2 connects the input to 3V. Note that the drain of Q2 and the source of Q1 both connect to the supply input, thereby preventing unwanted conduction through the MOSFET body diodes.

The gate drive for Q1 and Q2 must rise high enough above the drain-to-source voltage (V_{DS}) to fully turn on the MOSFETs. Though posing a problem for high-voltage inputs, that requirement is easily met for systems of 5V and lower using function generators or MOSFET drivers. The MAX4428 MOSFET driver, for instance, has complementary outputs for driving the two FETs out of phase, and produces a maximum 18V gate drive while sourcing and sinking up to 1.5A. The transient-driver source impedance (composed of R_{DS_ON} for Q1 and Q2, the ESR of C_{BP}, and parasitic inductance) must approximate that of the capacitor (C_{IN}) and its power source, as seen during normal operation by the power supply under test.

Parasitics

Parasitic inductance, resistance, and capacitance in the line transient setup (**Figure 2**) limit the ability of the

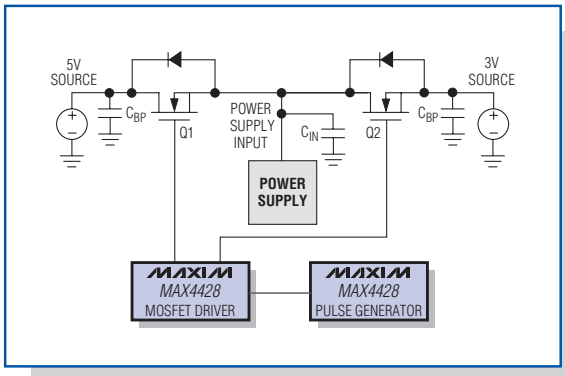

Figure 1. Two n-channel MOSFETs are shown switching between two DC sources, thus generating fast line transients.

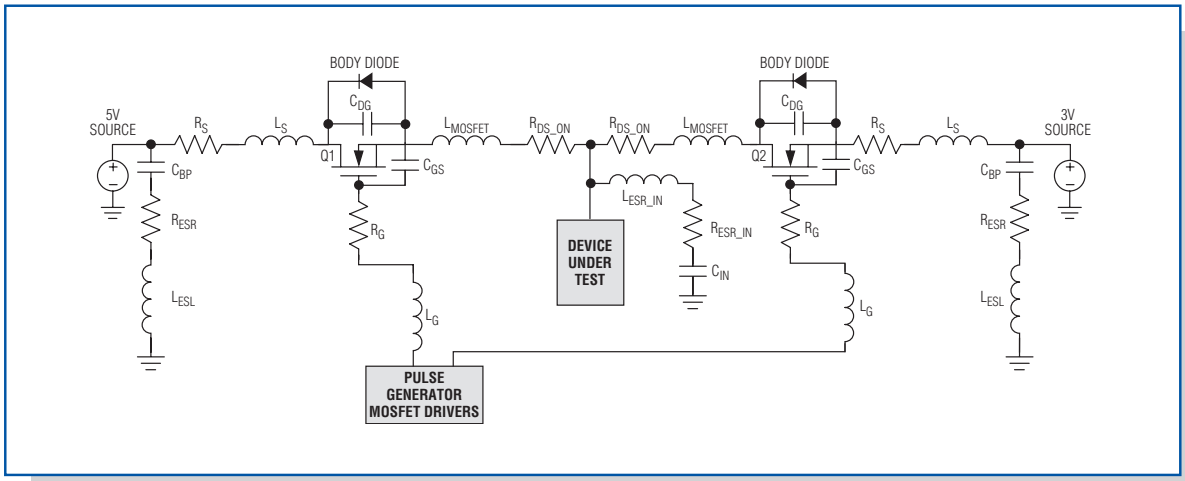

Figure 2. The circuit's ability to deliver a clean step-function waveform is limited by parasitic components in the line-transient setup.

circuit to deliver a clean step-function waveform. To source and sink the large currents necessary, you must minimize series resistance and inductance in the PC board, MOSFETs, and capacitors. Otherwise, the circuit's high capacitance and low series resistance react with the inductance and capacitance at the junction between the MOSFETs and the power-supply input. This produces an underdamped step response (resonance) that causes ringing. The inductance cannot be reduced to zero, but can be reduced to a point where the resonant frequency is high enough to be negligible in comparison with the line transient's rise and fall times.

Source parasitics

Occasionally, the layout requires that C_{IN} must be placed directly at the supply's input. This is because the introduction of a MOSFET between the input bypass capacitor and power-supply input results in unacceptable operation. If this is the case, the line voltage step must be imposed across C_{IN}. To change the voltage by ΔV_{STEP} in time (Δt), the C_{IN} must source or sink a current of

$$I = C_{IN} \times \frac{\Delta V_{STEP}}{\Delta t}$$

The bypass capacitors (C_{BP}) must be low-R_{ESR} ceramics with value much larger than C_{IN}. That condition minimizes the voltage drop across R_{ESR} at the currents necessary to charge and discharge C_{IN} in the time required.

Even with ceramic bypass capacitors, the series inductance of the capacitor (L_{ESL}) and of the connections between C_{IN} and C_{BP} (L_S) can pose a restriction on fast rise times and large currents. A few nanohenries of inductance can limit the current rise time (t) required to produce a reasonable voltage step at C_{IN}. For example, if C_{IN} = 100µF and ΔV = 1V, the circuit must source 100A into C_{IN} to step the voltage in 1µs. If the connection between C_{BP} and C_{IN} includes 100nH of parasitic inductance, the time required to raise the C_{IN} voltage by 1V is 2µs. Too much parasitic inductance can also cause excessive overshoot or ringing, which prevents the line transient from being a clean step function.

Parasitic resistance and inductance can be reduced by connecting smaller-valued ceramic capacitors in parallel. That arrangement reduces the total equivalent impedance by paralleling the R_{ESR} and L_{ESL} of the multiple capacitors. You can also reduce the inductance by using leadless capacitors or multilayer ceramic chip capacitors (MLCCs), or both. The distance from the bypass capacitors to the drains of the MOSFETs can also add impedance. PC-board traces of 2mm-wide, 1oz copper contribute parasitics of about 25mΩ/cm and 5nH/cm. You can minimize this additional impedance by making the connections between the bypass capacitors and MOSFET drains as wide and short as possible.

MOSFET parasitics

The parameters driving the selection of MOSFETs Q1 and Q2 are primarily the on-resistance (R_{DS_ON}), package size, and gate capacitance. R_{DS_ON} is important for the same reasons as the PC-board resistance and R_{ESR} of the bypass capacitors. Higher resistance limits current into the input capacitance (C_{IN}), and also causes excessive voltage ripple due to the pulsed currents of switching power supplies. Low R_{DS_ON} is especially important, because

R_{DS_ON} is the primary source of resistance in the capacitors' charge/discharge path.

The MOSFET series inductance, which includes the drain-to-source inductance and the inductance of the internal bond wires and leads, adds impedance in series with the power sources. Smaller MOSFET packages have less inductance because the bond wires and lead lengths are shorter. Identical MOSFET chips in D^2PAK and 8-pin SO packages, for example, show total series inductances of 10nH and 3.2nH respectively.

MOSFETs with very low on-resistance normally have higher gate capacitance (C_{GS}), which generally requires the use of MOSFET drivers capable of quickly charging and discharging C_{GS}. Drivers like the MAX4428 are useful because they can quickly turn on MOSFETs with several nanofarads of gate capacitance. To reduce the inductance and resistance that slow the rise time of the gate voltage, keep the trace lengths between driver and MOSFET gate short and wide.

When you have minimized the inductance and resistance of the capacitors' charge/discharge paths, the MOSFETs can be connected—either between the power-source bypass capacitors and the power supply's input capacitor or, if possible, directly to the supply input. In the latter case, the power-source bypass capacitors serve as input capacitance for the power supply. In either case, make the connection from the MOSFETs to C_{IN} or from the MOSFETs to the power-supply input as short as possible to minimize parasitic inductance and resistance contributed by the PC board.

Load transients

A good method for generating load-transient steps is to use an n-channel MOSFET to switch between two different load resistances at the power-supply output. For transient tests with large output currents, the MOSFET itself can be the load element (**Figure 3**). The MOSFET drain in that configuration connects to the power-supply output, and its source connects to ground through a current-sense resistor. The load resistance is adjusted by stepping the gate-to-source voltage (V_{GS}). As long as the MOSFET operates outside its saturation region, adjusting V_{GS} will vary the MOSFET R_{DS_ON}, and thus the load current.

To avoid adding extra inductance in the current-measuring loop, you should use a low-inductance current-sense resistor. Too much inductance limits the rise time of the output-current step, and causes ringing between the drain-to-source capacitance (C_{DS}) and the parasitic trace inductance (L_{PARA}), as illustrated in Figure 3. In this configuration, the sense resistor is part of the load.

The MOSFET must be placed directly across the output capacitor (C_{OUT}). Smaller MOSFET packages or MOSFETs in parallel can further reduce the L_{PARA}. To allow fast and clean switching, the connection between MOSFET gate and pulse generator (or MOSFET driver) must be short and wide to minimize trace inductance and resistance (R_G and L_G).

For smaller loads ($R_{LOAD} >> R_{DS\_ON}$), connect the MOSFET between ground and the load resistor R_A (bold lines of Figure 3). Connect another load resistor (R_B) in parallel with the MOSFET. When the MOSFET is on, the load resistance is R_A; when it is off, the load resistance is $R_A + R_B$. For the latter setup, pay close attention to the added inductance of the resistors and leads (L_R) that connect the load to the power supply's output capacitors. You can minimize that inductance by making the connections to R_A and R_B as short as possible. Avoid wire-wound resistors because they add excessive inductance. Power metal-film resistors are preferred.

What do line and load transients tell us about a power supply?

Having explained how line and load transients are generated, we now ask what information do they give about a power supply? As background, we must first understand the action of feedback and the effect of

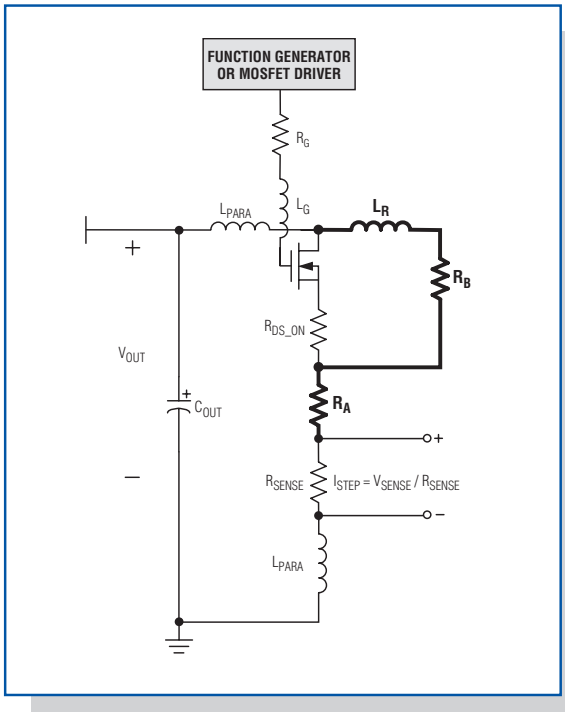

Figure 3. The MOSFET can be the load element in a load-transient test with parasitics labeled.

frequency response on the power-supply control loop and its response to line and load transients.

Loop gain attenuation

First, the effects of feedback on the small-signal gain of the power-supply controller must be understood. A simplified diagram for a buck converter with no feedback is shown in **Figure 4a**. Line and load steps are represented as inputs $I_{LOAD}(s)$ and $V_{IN}(s)$. The effect at the output from line and load disturbances is:

$$V_{OUT}(s) = V_{IN}(s) \times G_{VIN}(s) - I_{LOAD}(s) \times Z_{OUT}(s)$$

where $Z_{OUT}(s)$ is the output impedance.

The controller's power-filter gain ($G_{VIN}(s)$) is the small-signal gain from input to output. The buck converter, for example, has a power-filter gain of:

$$\frac{\Delta V_{OUT}}{\Delta V_{IN}} = \frac{D}{s^2 LC_{OUT} + s\frac{L}{R_{LOAD}} + 1} = G_{VIN}(s)$$

where D is the controller's duty cycle. L, C_{OUT}, and R_{LOAD} are shown in **Figure 4b**.

Output impedance for the buck converter is:

$$\frac{\Delta V_{OUT}}{\Delta I_{OUT}} = R_{LOAD} \left\| \frac{1}{sC_{OUT}} \right\| sL = \frac{sL}{s^2 LC_{OUT} + \frac{sL}{R_{LOAD}} + 1}$$

where R_{LOAD} is the controller's DC operating load.

With no feedback, any disturbance of input voltage or load current propagates through and affects the output voltage directly. For example, the buck converter operating with V_{IN} = 12V and a 50% duty cycle has a 6V output voltage. A 2V step in the input voltage, therefore, produces a 1V change in the output voltage. With the introduction of feedback, the output is made to regulate to a set reference (V_{REF}), shown by the bold lines in Figure 4a. The gain from line and load disturbances is now:

$$V_{OUT}(s) = \frac{V_{REF}(s) \times G_C(s)}{1 + G_{FB} \times G_C(s)} + \frac{V_{IN}(s) \times G_{VIN}(s)}{1 + G_{FB} \times G_C(s)} - \frac{Z_{OUT}(s) \times I_{LOAD}(s)}{1 + G_{FB} \times G_C(s)}$$

The above equation describes the closed-loop gain. The addition of feedback reduces line voltage and load current disturbances by the factor $(1 + G_{FB} \times G_C(s))$, where G_{FB} is the feedback-divider gain and $G_C(s)$ is the controller gain. The controller gain includes the gain for the power filter, the error amplifier, and other gain elements in the control loop. The term $G_{FB} \times G_C(s)$ is the open-loop gain. A Bode plot of open-loop gain shows the effect of feedback on the attenuation of V_{IN} and I_{LOAD} disturbances, with respect to frequency. Of special interest is the crossover frequency (f_C), at which $|G_{FB} \times G_C(s)| = 1$, and its associated phase shift (the phase margin). Phase margin is the difference between 180° and the phase shift at f_C. As the margin approaches 0°, the system exhibits unwanted effects at frequencies near f_C.

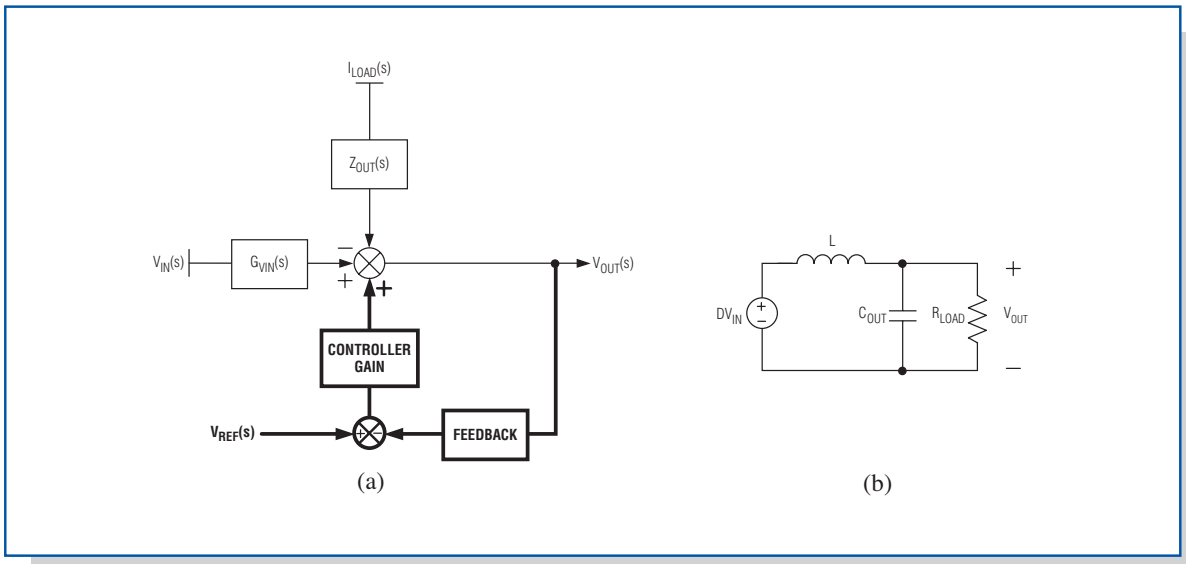

Figure 4. A simplified diagram and schematic for a buck converter is illustrated without feedback (a) and with feedback (b).

The frequency response is related to the transient response, because line and load transients are step functions, composed of an infinite sum of frequencies with amplitudes proportional to $1/f$. A step function is normally written in its frequency domain representation as 1/s, where s is the complex quantity jω.

Time domain to frequency domain

The transient response of a power supply with single-pole open-loop gain is a good example of what happens during a line or load transient. If the open-loop gain has only a single pole, it can be expressed as:

$$G_{FB} \times G_C(s) = \frac{DC_{GAIN}}{\left(\frac{s}{\omega_0} + 1\right)}$$

where DC_{GAIN} is the open-loop DC gain. **Figure 5a** shows a single-pole Bode plot in which the gain rolls off at -20dB/decade and crosses unity with 90° of phase margin. As the open-loop gain decreases with frequency, so does the attenuation of higher-frequency disturbances from V_{IN} and I_{LOAD}. By multiplying the closed-loop gain

$$\left(\frac{1}{(1 + G_{FB} \times G_C(s))}\right)$$

by the frequency-domain step function 1/s, and taking the inverse Laplace transform, you obtain the time-domain response to a step function (**Figure 5b**). A load step (ΔI_{LOAD}) applied to a controller with this open-loop gain exhibits an exponential response in the time domain, with an initial drop in the output equal to $\Delta V = I_{LOAD}(s) \times Z_{OUT}(s)$. At higher frequencies, $Z_{OUT}(s)$ is dominated by the capacitor impedance, which consists of the series resistance (R_{ESR}) and series inductance (L_{ESL}). Recovery takes the form:

$$V(t) = \frac{\Delta V}{1 + DC_{GAIN}} \times \left(1 + DC_{GAIN} \times e^{-(\omega_0 + \omega_0 DC_{GAIN})t}\right)$$

At one time constant ($\tau \approx 1/2\pi f_C$), the output voltage has recovered 63% of the initial drop (ΔV). Similarly, a line step raises the voltage at the output by the $G_{VIN}(s)$ multiplied by $V_{IN}(s)$. After one time constant ($1/2\pi f_C$ again), the output voltage has recovered 63% of the initial excursion.

An open-loop gain expression normally exhibits multiple poles, which results in a phase margin of less than 90°. If the phase margin of the open-loop gain is allowed to approach 0°, the closed-loop system response to a transient begins to show overshoot, and eventually oscillation. These output disturbances, which warn of the possibility of marginal stability or instability in the power-supply control loop, are the indicators to monitor when performing a transient test.

The closed-loop response shows overshoot and oscillation when it goes from attenuating frequency components of the transient to providing gain to them. That action is explained by interactions between the real and complex quantities in the denominator of the closed-loop gain. As the open-loop phase margin approaches 0°, its real portion

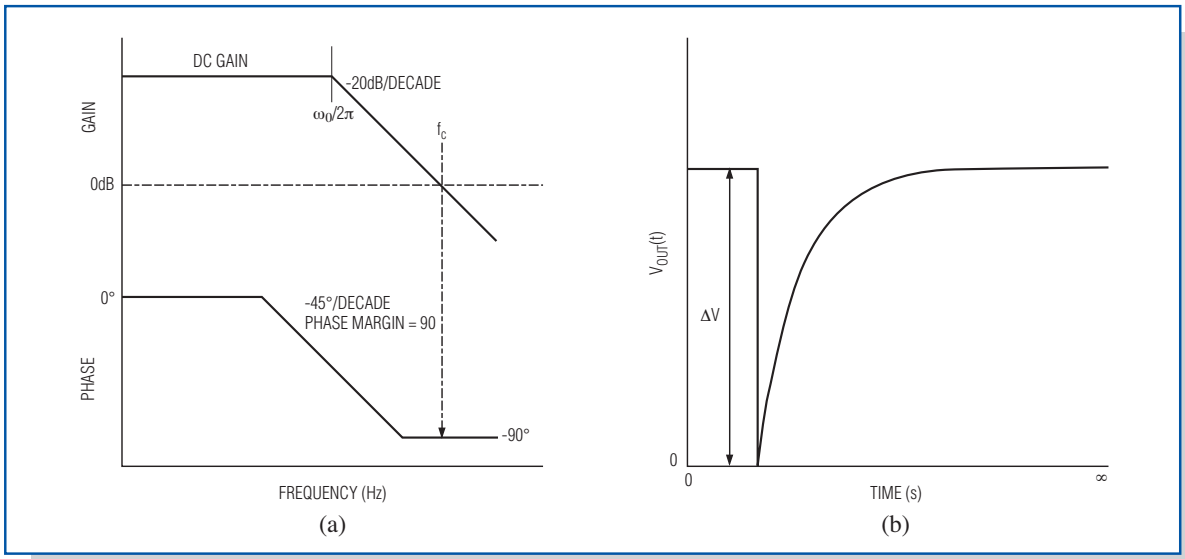

Figure 5. The Bode plot of single-pole loop gain is shown (a), as well as the time-domain response to a step function (b).

approaches -1 and its imaginary portion becomes less than 1. Also, the magnitude of the open-loop gain approaches unity. In the denominator of the closed-loop gain, the real parts sum to zero, leaving only the small imaginary portion. As a result, the closed-loop system that normally attenuates signals now provides gain to them.

A two-pole, open-loop gain expression gives a good example of what happens to a step response as you decrease the phase margin. Consider an open-loop gain with DC_{GAIN} = 60dB, designed with two real poles:

$$G_{FB} \times G_C(s) = \frac{1000}{\left(\frac{s}{\omega_1} + 1\right)\left(\frac{s}{\omega_2} + 1\right)}$$

The closed-loop gain is:

$$\frac{1}{1 + G_{FB} \times G_C(s)}$$

Assume that f_C occurs between ω1 and ω2. You then vary ω1 and ω2 to change the phase margin while maintaining f_C constant. Each separate phase margin has a separate step response in the time domain. By running the step() command in MATLAB®, you can obtain the time-domain response of the closed-loop gain with a two-pole, open-loop gain for various phase margins (**Figure 6**).

Figure 6 shows several transient responses possible at the output of a power supply with different phase margins. The initial step raises the output voltage. As the supply comes back into regulation with decreasing phase margin, the output shows increasing overshoot and eventually full oscillation. The fastest recovery with 0% overshoot occurs for a phase margin near 72°.

Load transient, practical example

Figures 7a/b/c show a 0 to 10A load transient with a buck converter using voltage mode control (see MAX1960EVKIT at www.maxim-ic.com). After adding a high-frequency pole at the COMP pin to reduce the gain above crossover, the open-loop gain and phase are measured at V_{IN} = 5V, V_{OUT} = 1.8V, and I_{LOAD} = 10A. Reducing the frequency of this pole decreases the phase margin. Figure 7a shows the response with an open-loop crossover frequency of 42kHz and 2° of phase margin. A load step from 0 to 10A causes the power supply to go into continuous oscillation. Increasing the phase margin to 11° damps the oscillation (Figure 7b). With 90° of phase margin (Figure 7c), the output response is that of an exponential single pole.

MATLAB is a registered trademark of The MathWorks, Inc.

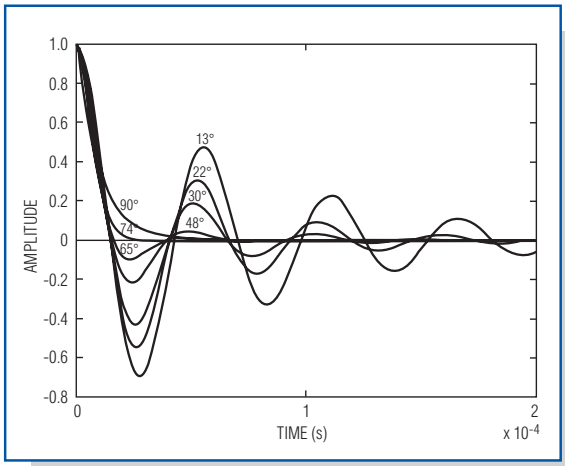

Figure 6. MATLAB step() commands can be run to obtain a time-domain response for closed-loop gain with various phase margins.

The small-signal frequency response for this circuit predicted that, with 2° of phase margin, the power-supply output would oscillate. Why does it not oscillate when the load steps back to 0A? The conditions at 0A do not support oscillation, because the small-signal response changes with the operating condition (bias point). The initial phase margin was measured at I_{LOAD} = 10A, V_{IN} = 5V, and V_{OUT} = 1.8V. When any of those parameters change, the small-signal gain and phase margin also change. This effect is seen in the 10A to 0 load step. When I_{LOAD} steps to 0A, the oscillation becomes damped. This is because the phase margin improves when operating with a 0A load (as opposed to a 10A load).

Be careful when attempting to equate the small-signal response predicted by open- and closed-loop transfer functions to large-signal changes like line and load transients. The small-signal responses are valid only for small changes around a specific operating point. Large-signal changes alter the bias point, which in turn changes the open-loop gain and phase. Only transients that occur at the specific operating point can be predicted by the small-signal response.

The load transient was generated using a single n-channel MOSFET (IRLR024N) in a DPAK package. The MOSFET was placed directly on top of one of the output capacitors, with a 37.5mΩ low-inductance sense resistor between the source and ground. A pulse generator (HP8112) directly stepped the gate from 0 to 4V, generating step responses from 0 to 10A in 200ns with virtually no overshoot or ringing.

Line transient, practical example

Figure 7d shows the line-transient response, with the same phase margin as in Figure 7c. The input voltage is stepped from 3.3V to 5V. Using the setup from Figure 1, two 9mΩ, n-channel MOSFETs (IRF3704, TO-220 package) switch the input of the supply under test between a 3.3V source and a 5V source. Each switch was placed between the MAX1960 input and two 470µF POSCAPs (6TPB470M) in parallel. Rise times of 400ns with 250mV overshoot were developed to simulate the line step.

Small-signal Bode plots predicted that the response should be exponential, which was verified. Stepping V_{IN} from 3.3V to 5V is a large-signal swing, and thus a different operating point is obtained when the power supply operates at V_{IN} = 3.3V. This transient test indicates that, at either input voltage, the phase margin is sufficient to prevent overshoot or ringing.

A similar article appeared in the July, 2004 issue of Power Electronics Technology.

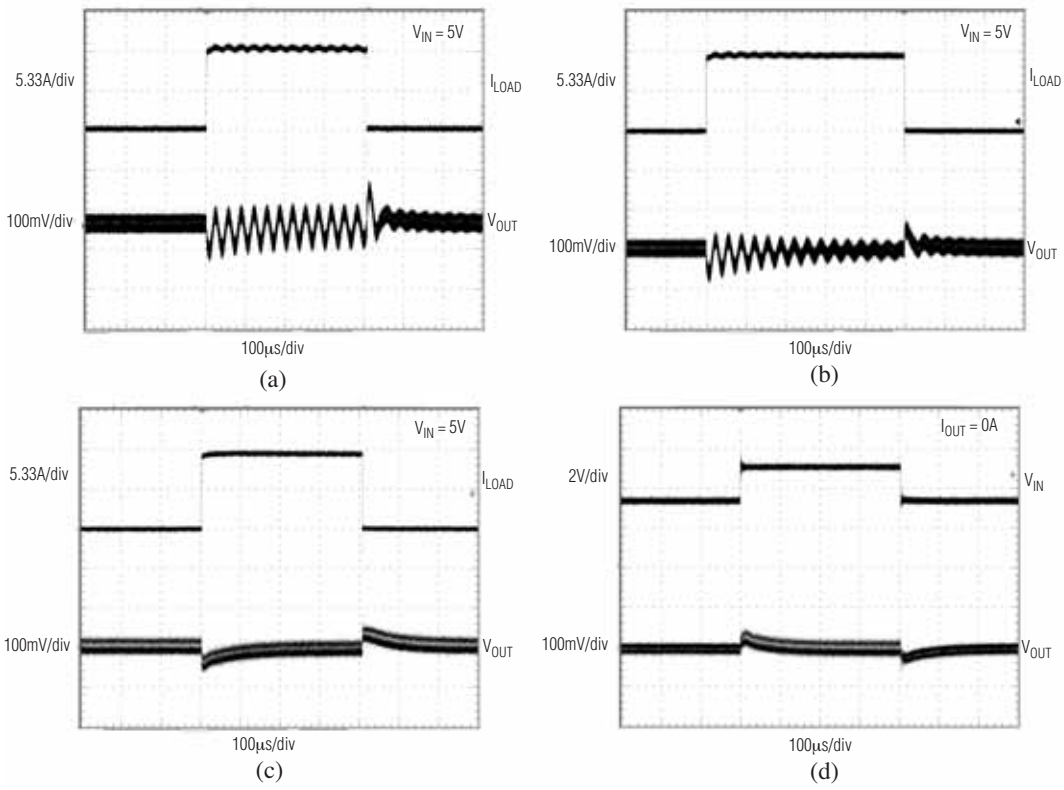

Figure 7. EV kit transient responses are shown for: a load step with 2° phase margin (a); a load step with 11° phase margin (b); a load step with 90° phase margin (c); and a line step with 90° phase margin (d).

Spectral content of NRZ test patterns

Non-return-to-zero (NRZ) signaling is widely used for data transmission in digital communication systems. Many NRZ test patterns have been created for system test and verification. These patterns are usually designed either to simulate actual data or to stress certain aspects of the system. To understand the effects of the various test patterns on a particular system, it is important to understand the frequency characteristics of both the test pattern and the system under test.

This article shows straightforward relationships between the time-domain characteristics of NRZ test patterns, such as data rate and pattern length, and their frequency-domain spectral components. Topics include an overview of NRZ test patterns, computation of the power spectrum, lab measurements of the power spectrum, and application of these concepts to system understanding.

Overview of NRZ test patterns

In NRZ signaling, each binary bit is assigned a unique time slot of duration, called the bit period (T_b). The signal is either high (representing a one) or low (representing a zero) during the entire bit period. NRZ waveforms are defined and measured as functions of time. Specifically, they are time-domain signals.

For a random NRZ data stream, each bit in the sequence has an equal probability (50%) of being a one or a zero, regardless of the state of the preceding bit(s). It is therefore possible to have large sequences of consecutive identical digits (CIDs). Because of the very low frequency content produced by long sequences of CIDs in the data signal, designing high-speed systems that can work with random data can be difficult.

Data encoding, or scrambling, is often used to format the random data into a more manageable form. One of the most widely used encoding methods in high-speed systems is known as 8b10b, which is used in Ethernet, Fibre Channel, and high-speed video applications. 8b10b encoding takes 8 bits of data and replaces it with a 10-bit symbol. The extra bits are added to balance the pattern (make the number of ones equal the number of zeros for a given bit interval) and limit the maximum number of CIDs. The encoding algorithm is also used to improve bit-error ratio (BER) by mapping the 8 bits to specific symbols in the 10-bit signal space that can be easily distinguished from other 10-bit symbols. Other methods, such as scrambling or 64b66b encoding, are common to SONET and SDH telecommunication systems. Scrambling and 64b66b encoding also work to balance the pattern and improve the BER, though much larger runs of CIDs are possible with these methods.

For a given application, there may be several types of test patterns that stress various performance aspects or system components. For example, a K28.5± pattern (11000001010011111010) is often used to test the deterministic jitter performance of systems that use 8b10b encoding. Likewise, a pseudorandom bit stream (PRBS) is used as a general-purpose test pattern in encoded, random, and scrambled NRZ applications.

The PRBS is typically denoted as a 2^X - 1 PRBS. The power (X) indicates the shift register length used to create the pattern. Each 2^X - 1 PRBS contains every possible combination of X number of bits (except one). A short PRBS, such as the 2^7 - 1 PRBS (127 bits), is often used in Ethernet, Fibre Channel, and high-speed video applications, because it provides a good approximation to an 8b10b-encoded NRZ data stream. A 2^{23} - 1 (≈ 8.4 million bits) PRBS is commonly used in both SONET and SDH telecommunication systems, which require a test pattern with lower frequency content and provide a better representation of scrambled or random NRZ data.

Computing the power spectrum of an NRZ test pattern

Each NRZ test pattern has an associated power spectral density (PSD) that indicates the frequency distribution of the power in the pattern. The two primary methods of computing PSD are to: (a) square the magnitude of the Fourier transform of the pattern; or (b) compute the Fourier transform of the pattern's autocorrelation function[1]. The first method, (a), is generally simpler for signals that can be mathematically written in a finite, closed form (e.g., $s[t] = A\cos[2\pi f_0 t]$). The second method, (b), is used for more complicated signals, such as long sequences of NRZ data (like test patterns) or random bit streams. To apply these methods, a review of some of the basics of Fourier analysis[2] is useful.

- The delta function, $A\delta(t)$, can be thought of as an infinitely narrow rectangular pulse with area (A). It has a nonzero value only when the argument of the function is equal to zero and is represented graphically by a vertical arrow.
- The comb function, $A\Sigma_n\delta(t - nT)$, is composed of an infinite number of equal-area delta functions spaced at uniform intervals (T).

- The Fourier transform of a comb function is also a comb function, where the interval is inverted (e.g., n/T) and the areas of the delta functions are modified by the inverted interval (e.g., A/T).
- Convolution in the time domain (represented symbolically by *) is equivalent to multiplication in the frequency domain, and vice versa.
- Convolution of a signal with a delta function results in a copy of the signal that is shifted to the location of the delta function.
- Multiplication of a signal with a delta function, or "sampling", results in a delta function with an area modified by the signal's magnitude at the location of the delta function.

As an example of an application of the above rules, the PSD of an NRZ test pattern is computed (**Figure 1**). The test pattern can be represented by a sequence of high and low levels (representing ones and zeros) with a defined T_b and total pattern length, $L = nT_b$. An infinite repetition of the pattern results from the convolution of the finite-length test pattern with a comb function that has a spacing interval equal to the pattern length (Figure 1a). Next, the autocorrelation functions for each component of the test pattern are separately computed (Figure 1b). Note that the autocorrelation of the test pattern approximates a triangle (the accuracy of this approximation improves as the length and "randomness" of the pattern increase). Finally, the Fourier transform of the autocorrelation functions are used to compute the power spectrum (Figure 1c).

The power spectrum resulting from the example in Figure 1 shows an infinite sequence of discrete spectral lines (delta functions) scaled by a "$\text{sinc}^2(f)$" envelope, where sinc(f) is defined as $\sin(\pi f)/(\pi f)$. Important observations that apply to test patterns in general include: (a) the nulls in the $\text{sinc}^2(f)$ envelope occur at integer multiples of the data rate; (b) spectral lines are evenly spaced at an interval that is the inverse of the pattern length; and (c) the magnitude of the $\text{sinc}^2(f)$ envelope decreases (i.e., "flattens out") as the data rate and/or pattern length increase. In the limit, as the pattern length approaches infinity, the spacing between the spectral lines becomes infinitesimally small, and the spectrum shape approaches a continuous $\text{sinc}^2(f)$ function.

As an example, if the 6-bit pattern shown in Figure 1a is transmitted at a data rate of 1.25Gbps, the spectral-line spacing, amplitude, and spectral nulls can then be calculated as shown in **Figure 2**. Note that $\text{sinc}^2(f)$ envelope shown in Figure 2 is an approximation of the 6-bit pattern. The accuracy of this approximation improves as the pattern length or randomness increases.

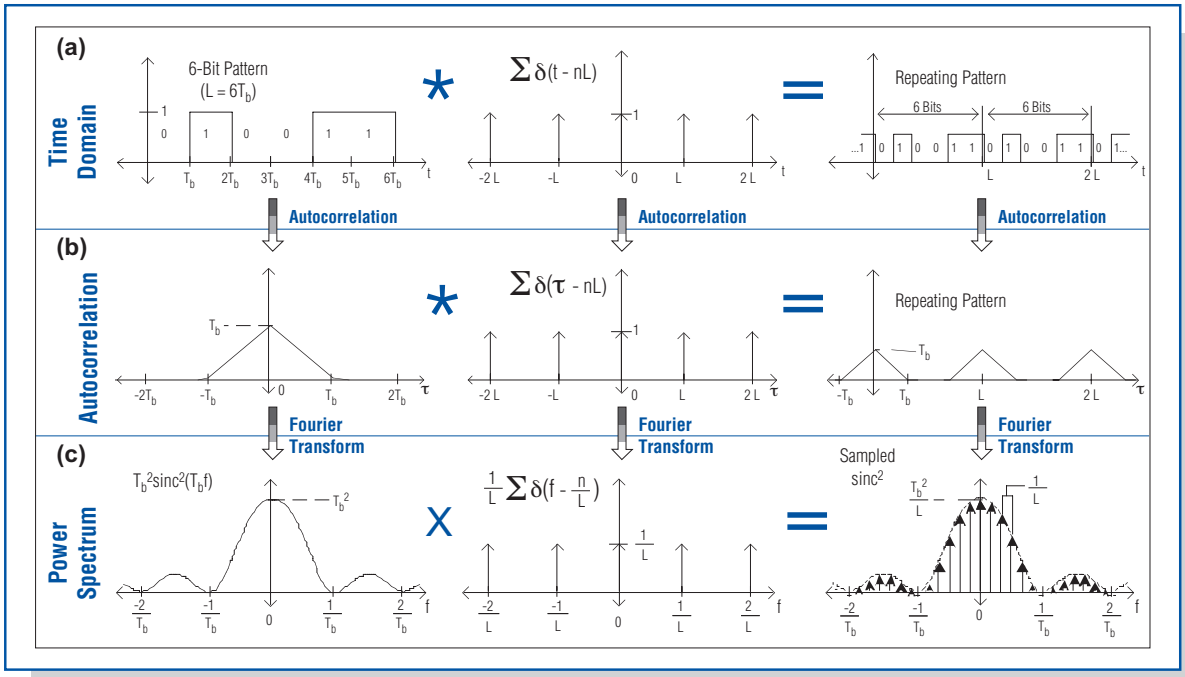

Figure 1. The test patterns illustrate time domain (a), autocorrelation (b), and power spectrum (c) of NRZ signaling.

Spectrum analyzer measurements

The equations and principles described previously can also be demonstrated by lab measurement. This involves using a high-speed pattern generator to create the test patterns, and a spectrum analyzer to measure the PDS of the signal. Starting with a simple example, the measured spectrum of a 4-bit pattern (1110 test pattern) transmitted at 1.25Gbps can be seen in **Figure 3**. The spectral nulls are measured at 1.25GHz ($1/T_b$) and 2.5GHz ($2/T_b$), and the line spacing is 312.5GHz (1/L). The power spectrum envelope is also seen to be approximately $sinc^2(f)$. The slight deviations in magnitude result from the short pattern used in this example.

Increasing the pattern length to 20 bits (K28.5± test pattern) and keeping the transmission rate at 1.25Gbps (**Figure 4**), the spectral nulls are measured to be in the same locations (1.25GHz and 2.5GHz). Meanwhile, the spectral-line spacing is reduced to 125MHz due to the longer pattern length. Additionally, the envelope of the spectral lines more closely matches the $sinc^2(f)$ function than the 4-bit pattern example.

The K28.5± test pattern presents an interesting aspect to this topic. It is a 20-bit pattern. However, the spectral-line spacing is measured to be 125MHz when transmitted at 1.25Gbps, which would correspond to a 10-bit test pattern. This discrepancy is because the K28.5± pattern is composed of a K28.5+ sequence (1100000101) and its inverse, the K28.5- sequence (0011111010). In the frequency domain, the K28.5- sequence contains the same spectral information as the K28.5+ sequence. The pattern therefore repeats spectrally every 10 bits, which is the reason for the 125MHz spacing in Figure 4.

The $sinc^2(f)$ envelope becomes very apparent as the pattern length is increased further. **Figure 5** illustrates this point by using a 2^7 - 1 PRBS pattern (127 bits) transmitted at 2.5Gbps. At this longer pattern length, delta spacing is reduced to approximately 19.7MHz. And, corresponding to the higher data rate, spectral nulls are at 2.5GHz and 5GHz. Given the small spectral-line spacing in respect to the data rate, the $sinc^2(f)$ envelope and spectral nulls are clearly seen in the power spectrum (Figure 5).

Figure 6 illustrates the difference in the spectral-line magnitude and spacing for a 2^7 - 1 PRBS pattern at 1.25Gbps and 2.5Gbps. As seen in this figure, when measured at the same frequency, the magnitude of the

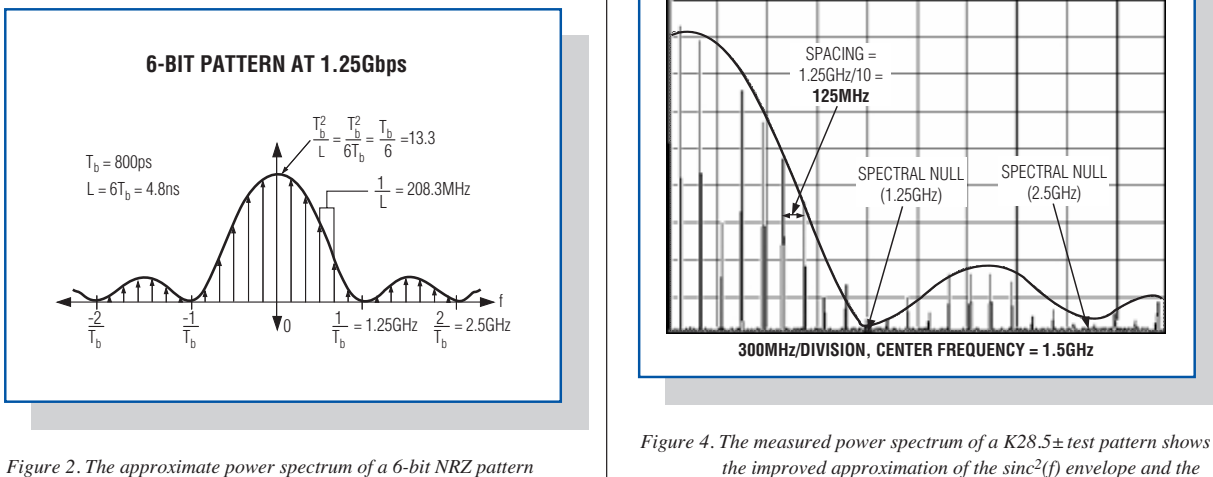

Figure 2. The approximate power spectrum of a 6-bit NRZ pattern shows spectral-line spacing and the $sinc^2(f)$ envelope.

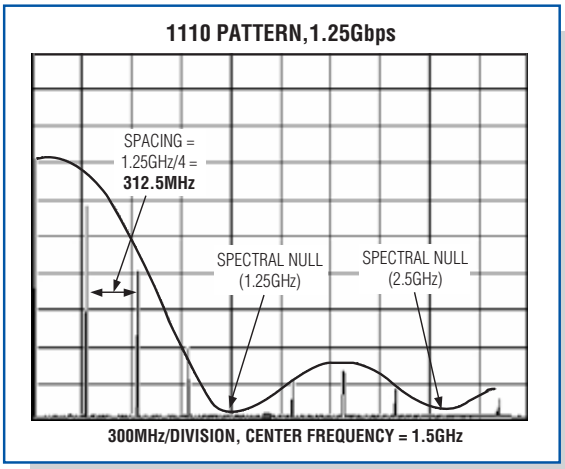

Figure 3. The power spectrum of a 4-bit pattern shows slight deviations of the spectral-line magnitude from the $sinc^2(f)$ envelope. As the pattern length increases, the deviation is reduced.

Figure 4. The measured power spectrum of a K28.5± test pattern shows the improved approximation of the $sinc^2(f)$ envelope and the reduced spectral-line spacing due to the longer pattern.

spectral lines and spacing is larger at 2.5Gbps data transmission rates than at 1.25Gbps.

Application examples

Knowledge of the power spectrum of NRZ test patterns can lead to significant improvements in digital communication-system design. This is illustrated through examples of three different applications: receiver bandwidth, adaptive equalizers, and electromagnetic interference (EMI).

Receiver bandwidth

The design process for a receiver inevitably includes questions about the necessary bandwidth. If the bandwidth is too low, the high-frequency components of the received signal are attenuated, and the signal is distorted. If the bandwidth is too high, excess noise is admitted to the receiver, causing a reduction in signal-to-noise ratio (SNR). Also, to achieve the higher bandwidth[3], an increase in complexity and cost is necessary. Knowing the spectral content of the signals that will be received, the bandwidth decision can be made in a manner that includes only the critical spectral components.

Adaptive equalizer

Adaptive equalizers are designed to reverse distortion effects caused by nonideal transmission media. The MAX3800 adaptive cable equalizer, for example, reverses the distortion caused by skin-effect losses in copper cables at data rates as high as 3.2Gbps[4]. It accomplishes this by comparing the power of the input signal at two discrete frequencies (f_1 = 200MHz and f_2 = 600MHz). Based on the $sinc^2(T_bf)$ envelope of the power spectrum with the first null at 3.2GHz, the power ratio at these two frequencies should be $sinc^2(T_bf_1)$ / $sinc2(T_bf_2)$ = 0.987/0.890 = 1.11. If the measured ratio is different than expected, the equalizer changes the amount of skin-effect compensation in order to restore the correct ratio. This works well for high data rates and long data patterns. However, using our knowledge of spectral content of NRZ test patterns, we can predict that some patterns may cause problems.

If, for example, the data rate is reduced to 622Mbps, the $sinc^2(f)$ envelope with first null at 622MHz results in a 200MHz to 600MHz power detector ratio of 0.703/0.00134 = 525, instead of the expected 1.11. As the equalizer tries to restore the expected 1.11 power ratio, the output may be distorted. As another example, consider a short test pattern with a pattern length of 10 bits. For shorter patterns, the spectral lines are spaced at larger intervals. In the specific case of the 10-bit pattern at a data rate of 3.2Gbps, the spectral lines are separated by 320MHz, with the first few at 0, 320, and 640MHz. For this type of pattern and data rate, there may be little or no power to detect at 200MHz or 600MHz. This, in turn, can cause signal distortion, because the equalizer is not able to adapt correctly.

Electromagnetic interference (EMI)

The effects of EMI in a system can be reduced or eliminated by altering the magnitude and/or frequencies of the power spectrum. The alteration(s) can be done by changing the data rate or pattern length.

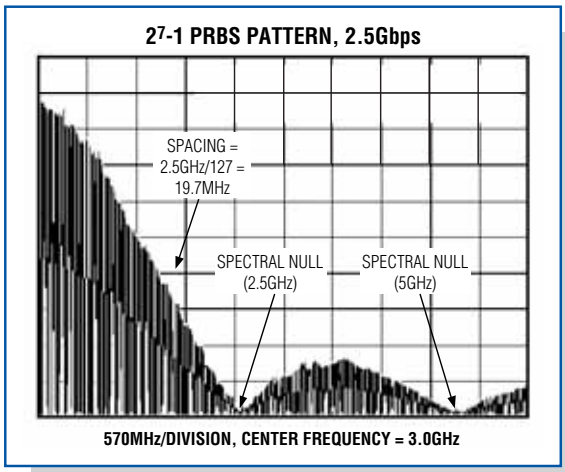

Figure 5. The power spectrum of the 2^7 - 1 PRBS (127 bits) clearly shows the spectral nulls and $sinc^2(f)$ envelope.

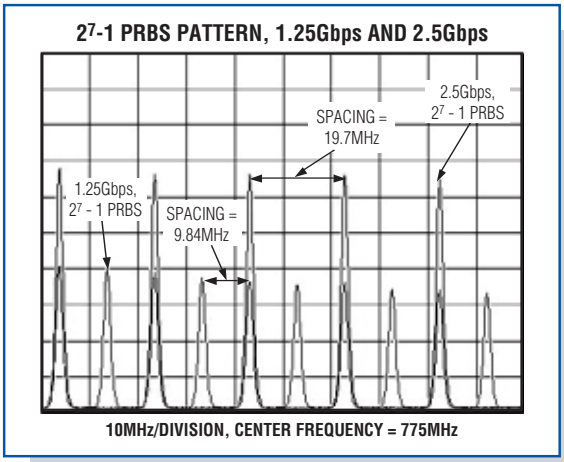

Figure 6. The measured power spectrum of a 2^7 - 1 PRBS pattern transmitted at 1.25Gbps and 2.5Gbps (as viewed at 725MHz to 825MHz) shows the spectral-line magnitude and spacing difference as the data rate changes.

As the data rate increases, the spectrum nulls are spread farther apart. Also, the magnitude of each spectral line is reduced by pushing some power to higher frequencies. Spreading the power over a larger frequency range leaves less at the frequencies of interest. One way to achieve this effect is by adding extra bits to the original data stream to effectively increase the data rate.

Pattern length also plays a role in EMI, because spectral-line magnitude and spacing vary as the pattern length changes. A longer pattern reduces the magnitude and spacing, while a shorter pattern increases the magnitude and spacing. To reduce EMI at a specific frequency, the pattern length can be changed to shift the spectral line away from a particularly sensitive frequency range. Alternatively, a longer pattern can be used to reduce the magnitude of the EMI.

Conclusion

A clear understanding of the frequency-domain spectral content of NRZ data is critical to success in high-speed, digital-communication system design. The principles presented in this article establish basic relationships between NRZ data-time domain characteristics (pattern length, data rate, etc.) and their corresponding frequency domain characteristics (spectral magnitude, envelope, and line spacing). These principles can be applied to a variety of circuit design issues, including filtering, signal equalization, and EMI.

References

1. J.W. Goodman, Statistical Optics, John Wiley & Sons, New York, NY, 1985.

2. J.D. Gaskill; Linear Systems, Fourier Transforms, and Optics; John Wiley & Sons, New York, NY, 1978.

3. HFAN-09.0.1 NRZ Bandwidth-HF Cutoff vs. SNR, Maxim Integrated Products, Inc., 2002.

4. MAX3800 data sheet, Maxim Integrated Products, Inc., 2001.

System-management IC meets monitoring and sequencing requirements in multivoltage systems

In most electronic systems, it is important to monitor system voltages both to ensure that processors and other ICs remain reset during power-up, and to detect when brownout conditions occur. This monitoring minimizes code-execution problems, which corrupt memory or cause systems to execute improperly. In high-end systems, it is also critical to ensure proper sequencing of the many power supplies within these systems. Proper sequencing prevents latch-up conditions, which create system problems or damage important components such as microcontrollers (µCs), DSPs, ASICs, or microprocessors (µPs). Commonly, one or more supervisory products are needed to implement the proper sequencing and monitoring functions noted here.

Traditionally, many of these functions have been performed with power-on resets and other µP supervisory circuits. In recent years, as the number of supply voltages increased, the number of devices required to perform this duty also grew, thus compounding complexity, cost, and the board space consumed.

Monitoring and sequencing complex systems

The easiest way to monitor a supply voltage is with either a power-on reset (POR) or a voltage detector circuit. These devices can monitor a single voltage or multiple voltages. After the monitored supply voltage has powered up and exceeded the POR's voltage threshold, the POR's output does not de-assert until after a specified time period. This allows both the system clocks to stabilize, and the system boot routine to initialize before permitting the µC to operate. These PORs and voltage detectors can also be used to sequence power supplies. Connecting the output of a POR that is monitoring one regulator to the shutdown pin of the next regulator (i.e., daisy-chaining them), one regulator will then come up after the other, once the POR's time delay has elapsed.

As the number of system supply voltages increases, voltage monitors and supervisors that monitor multiple voltages become necessary. However, because it is common for ten to fifteen voltages to power a complex system, several such devices are often needed.

Challenges when using multiple supervisors

Using this multiple supervisor approach has its own problems. One problem is finding devices with the correct thresholds. Although there are a number of standard voltages, such as 3.3, 2.5, 1.8, 1.5, and 1.2V, many non-standard voltages need to be monitored. This requires external resistor-dividers to set the monitored thresholds. If system-supply voltages change (e.g., you lower an ASIC's core voltage to reduce power consumption, or increase it to enhance the ASIC's performance), you would have to change the resistor values to accommodate these new voltages. Obtaining this flexibility requires these additional external resistors, and thus more board space and cost. The same problems occur when selecting the correct reset timeout periods.

Another problem with multiple supervisors occurs when a system must provide a specific power-up sequence. When larger numbers of supply voltages power a system, the daisy-chaining technique outlined above may not be capable of handling the timing for when various supplies are coming up. Also, as the sequencing requirements change during development, altering circuitry to accommodate those later changes becomes problematic.

An additional sequencing problem can occur when these large systems use "silver box" or "brick" power supplies. These supplies simplify power supply design, but pose a problem when a particular power-up sequence is required. For example, a brick supply that provides multiple output voltages might only have a single enable pin. Therefore, all its supply voltages turn on and off at the same time under the control of that one pin. A brick supply with multiple enable (or shutdown) inputs can resolve this issue. However, if multiple ICs share the same supplies (for example, a 3.3V I/O logic supply and a 1.8V core supply), the two ICs may have conflicting requirements. One device might require its core supply to come up before its I/O supply, while the second device might require its supplies to sequence in the opposite order.

This issue can be resolved with an external switch, such as a MOSFET. For low-power applications, you can use a p-channel MOSFET, which is generally more expensive than an n-channel MOSFET, but simpler to use. An

n-channel MOSFET is optimal for higher current applications, because its lower on-resistance reduces the voltage drop across the switch. An n-channel can be used for very-low-voltage cores as well. However, to fully enhance an n-channel MOSFET, a high enough supply voltage must be available to provide a suitable gate-to-source voltage. In systems without this higher voltage, you can use ICs such as the MAX6819/MAX6820 power supply sequencers to control the sequencing process; the internal charge pump for these devices guarantees a 5V gate-to-source voltage. This voltage drop is too high for some systems. Consequently, board designers sometimes double the number of regulators so this sequencing problem is avoided.

As the number of supply voltages increases, using multiple MAX6819/MAX6820 circuits can work. As when using multiple PORs, these power-sequencing ICs can be configured to daisy-chain the power supplies. However, for a large number of voltages, this solution requires too many discrete ICs, thus increasing the overall system cost and consuming too much board space.

The margining function

The ability to monitor and sequence supply voltages is critical to providing high levels of reliability. In many large, complex systems such as those found within telecom, networking, server, and storage equipment, additional testing of key components is required. One example is a margin test, which checks a system's performance as its voltages are temporarily varied up and down. Margining is often performed when a system is under development, but is also commonly implemented during the manufacturing process. The margining process is used to improve a system's long-term reliability.

The supply voltages can be adjusted by trimming the regulator's reference input (for voltage-regulator modules), altering the voltage regulator's feedback loop, adjusting a "brick" power supply's trim input, or by programming the regulator through an interface. There are various degrees of margining control. One method is an "all or nothing" approach in which the supplies are increased/decreased by some fixed amount (e.g., ±5% or ±10%). Another, more exacting method increases or decreases the supply voltages in smaller steps (e.g., 10mV or 100mV) to allow the evaluation of the system performance in greater detail. To gain more detailed information about the system voltages during normal operation and during the margining process, an analog-to-digital converter (ADC) can be used to accurately measure these values. Note that the POR controlling the μC needs to be disabled during the margining process to prevent the system from being reset.

Performing these margining functions can be quite tedious when working with large systems. You can use multiple supervisor devices to manage the margining process, along with the monitoring and sequencing tasks. However, this approach can be problematic. Aside from the cost of the ICs and the additional board space consumed, changes to the supply voltage levels or the sequencing order of those supplies are difficult to accommodate. This is because the requisite design changes are not trivial.

An integrated system-management device

One way to minimize these monitoring and sequencing issues is through the use of a fully integrated, EEPROM-configurable system-management device such as the MAX6870. This type of IC integrates the functions required to monitor and sequence the system power supplies and to simplify the margining process. The MAX6870 provides the flexibility to: change the voltage thresholds at multiple inputs easily; sequence the outputs in any order; and configure the outputs' structures to be push-pull, open-drain, or charge-pump enhanced. In addition, the digital inputs and outputs can be configured for active-high or active-low logic. Furthermore, the outputs can be either disabled or set to a predetermined state during the margining process.

Figure 1 shows a block diagram depicting the MAX6870's features. This device's six inputs can monitor a system's various supply voltages, among other tasks. The two threshold levels that you program for each input can be set to detect two undervoltage conditions or an undervoltage and an overvoltage condition (i.e., a window detector). These threshold levels are programmed through an I^2C* interface and stored in the configuration EEPROM. You can designate thresholds that range from 0.5V to 5.5V in 10mV and 20mV increments, depending on the threshold voltage range you select. One input, IN1, can monitor a voltage as high as 13.2V, and therefore can allow a 12V system bus voltage (or lower) to be monitored directly. A second input, IN2, allows either a second high voltage or a negative voltage to be monitored. The remaining inputs, IN3–IN6, monitor supply voltages with 0.5V to 5.5V levels.

An internal multiplexer routes the six detector inputs, along with the two auxiliary inputs, to a 10-bit,

**Purchase of I^2C components from Maxim Integrated Products, Inc., or one of its sublicensed Associate 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.*

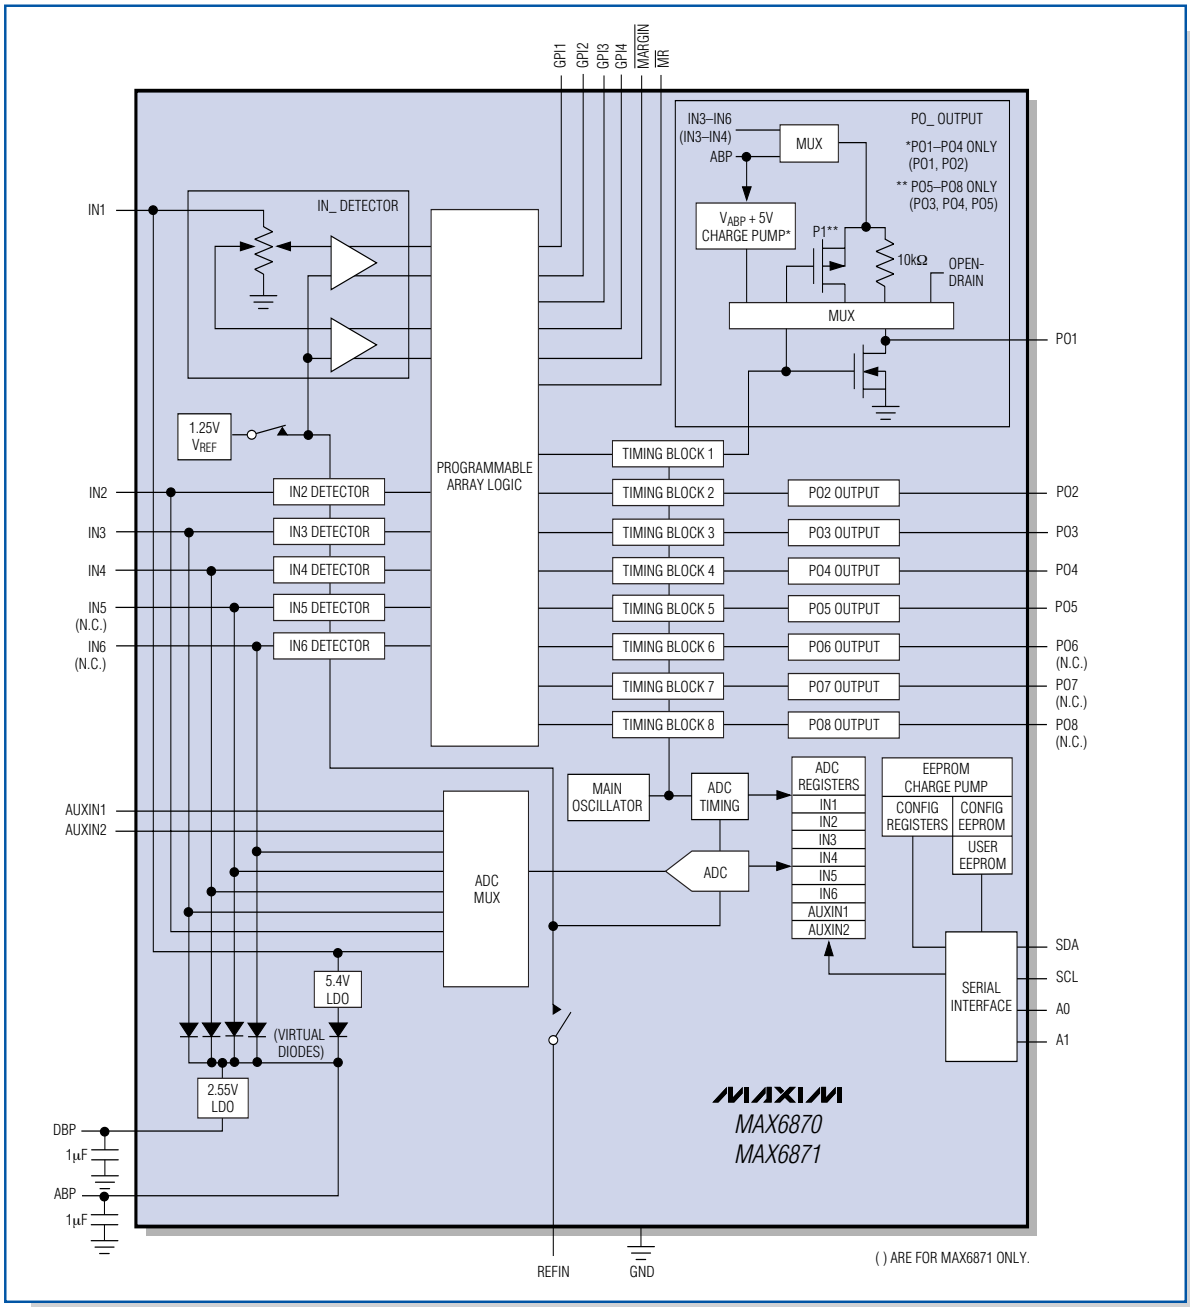

Figure 1. This IC can monitor and sequence multiple supplies, while providing ADC read capability. Internal EEPROM allows key parameters, such as thresholds, timing, logic dependencies, and output structures to easily be adjusted.

1%-accurate ADC. The ADC writes the digitized version of the voltage at each of these eight inputs to internal registers. Available through the I²C interface, these stored values are commonly used during the margining process, when trimming a power-supply output, or when checking the long-term stability of system voltages. Also, by using the two auxiliary inputs, you can read two additional voltages, such as the analog outputs from a current-sense amplifier or a temperature sensor.

The device operates when any of the voltages on IN3–IN6 exceeds the minimum operating voltage of 2.7V, or when IN1 exceeds 4V. Any of these inputs can power the device through the diodes pictured in Figure 1.

The six detector inputs and the four general-purpose inputs (GPIs) can assert any of the eight outputs, depending on the connections set up within the programmable array logic. Also, the outputs can be controlled by the device's other outputs, or by a mix of input and output signals. Each output's timing delay is independently programmable and can be stored in the IC's EEPROM.

The outputs are also configurable. You can set them up as open-drain outputs with internal or external pull-ups, or as push-pull outputs that can be connected inside the IC to any of the monitored supply voltages. All outputs can be active low or active high. Also, as mentioned above, different combinations of inputs and outputs can drive specific outputs; the MAX6870's programmable-array logic allows a wide variety of connections. For instance, OUT2 can be controlled by IN2 and also by OUT1. This type of connection is useful when the supply powered up by the OUT1 signal must come up before the supply powered up by the OUT2 signal.

The MAX6870 also includes an internal charge pump that allows OUT1–OUT4 to fully enhance an external n-channel pass element without the need for an additional supply voltage.

Other features in this device include two watchdog timers with configurable watchdog timeouts and startup delays. The startup-delay feature of the watchdog timer provides a long delay after a reset condition. This delay allows extra time for the system initialization process, allowing memories to be uploaded and software routines to fully load.

A manual-reset input permits a test technician to manually assert all of the IC's outputs. The device's margin input can be used to latch the outputs in their current state to prevent the system from resetting during the margin process. The margin input can also be used to set the outputs into a predetermined state by programming the associated EEPROM registers. The MAX6870 also provides 4kb of user EEPROM to store items such as serial board identification, board revision history, and other programming information.

The MAX6870, moreover, includes configuration registers and configuration EEPROM. During the prototyping phase of a project, you can write updates to the configuration registers to make immediate changes. If these changes are acceptable, you can then write them into the configuration EEPROM. To reload the saved EEPROM configuration, the system can be rebooted through a software reboot or through a complete hardware reboot. The configuration EEPROM downloads its data to the configuration registers during the reboot phase.

The MAX6870's EV kit

To simplify the MAX6870's configuration process, an EV kit is available that allows you to point and click on your computer screen to load the correct configuration information. Each screen allows you to configure a part of the device without referring to the register tables. The screens allow you to set thresholds levels, timing delays, logic operating states (active low/active high), and logic input configurations, as well as to configure the outputs.

Figure 2 shows the main configuration screen from the *MAX6870 Evaluation Software*. You can click on the blocks within the block diagram or on the tabs. Clicking on one of the blocks allows you to configure that specific block. Clicking on a tab brings up a screen specific to a particular function. For instance, when you select the Voltage Monitor tab (**Figure 3**), the screen that appears allows you to easily select the thresholds and configure the inputs. Or when you click on the Output tab (**Figure 4**), you can set up the output as open-drain, push-pull, or charge-pump enhanced, as well as set up the logic terms that determine the output's state.

Once you have completed the device's configuration, you can save the configuration data in its EEPROM. In addition, this data can be saved to a file that can be loaded to another device. Of course, you can write directly to all the configuration registers and the configuration EEPROM through the I^2C interface. The data sheet provides the values needed to do this. Using this approach, however, takes more time and increases the chances of error.

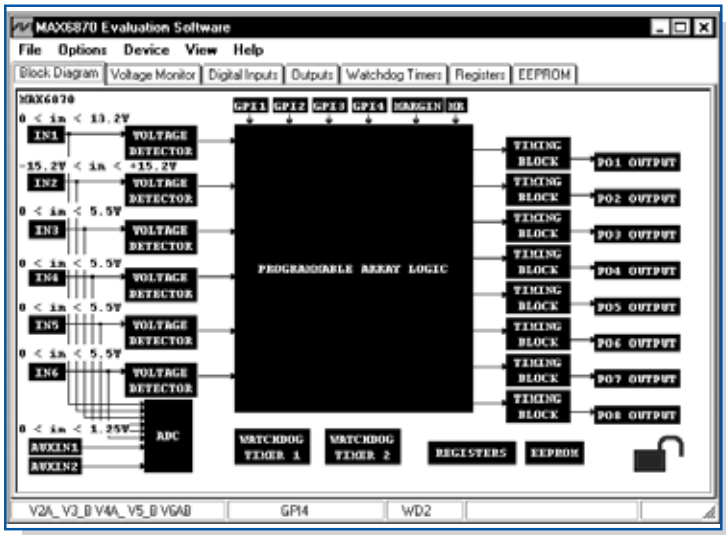

Figure 2. Click on the appropriate box or tab to set up thresholds, delays, output configurations, and logic.

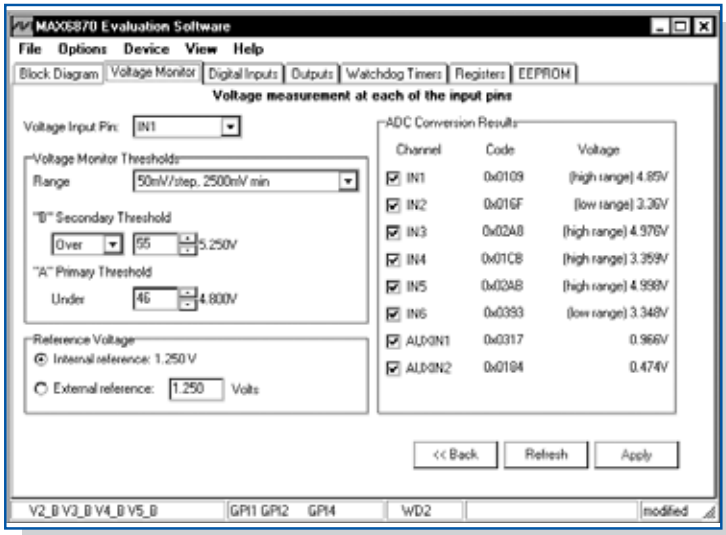

Figure 3. By clicking on the Voltage Monitor tab, you can determine for each input whether it monitors two undervoltage levels, or an undervoltage and an overvoltage level. You can also set the thresholds and select which digitized inputs can be viewed.

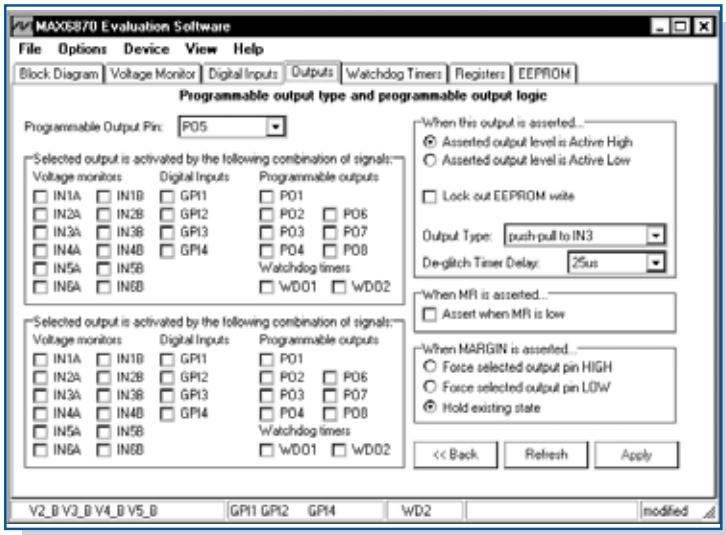

Figure 4. By clicking on the Output tab, you can configure each output to be open-drain, push-pull, or charge-pump enhanced. Each output is connected to the IC's programmable logic array and can be used to control other outputs.