Tame Your High-Speed Switching Amplifier Circuit Designs

Learn How To Use PWM Switching Amplifiers In Motor-Control And Audio Applications.

DENNIS EDDLEMON
APEX MICROTECHNOLOGY CORP.

In recent years, linear power op amps have basically reached the pinnacle of their capabilities. The output power in these linear hybrid amplifiers has surpassed what was only imagined just 10 years ago. But to achieve these extraordinary power levels, the linear power amplifier may require extraordinary package materials and heat sinks, and may likewise suffer the extraordinary associated costs. And yet, like the television character Tim the Tool Man Taylor, we cry out for "more power!"

Fortunately, hybrid pulse-width- modulation (PWM) switching amplifiers have now appeared on the scene. The first commercial hybrid PWM amplifier, the SA01, was introduced by Apex Microtechnology in early 1996. A number of other models have since followed. These amplifiers have shown the way to deliver far more power in many applications at a lower cost than comparable linear amplifiers. For those unfamiliar with PWM amplifiers, we're talking about amplifiers that can supply hundreds or thousands of watts to the load. After an explanation of basic PWM switching technology, this article will teach readers how to build and tame motor-control and high-quality audio-application circuits using some of these impressive amplifiers.

There are two basic ways of delivering power to a load. The linear approach reduces the resistance of the pass element as more output voltage is needed, approaching zero resistance when maximum output voltage is required *(Fig. 1a)*. When zero output voltage is commanded, the pass element approaches an infinite resistance. At either extreme, the power consumed by the pass element is minimal.

The linear circuit, however, is at its worst when the output voltage is in between these limits. In fact, the efficiency is only 50% at an output voltage of 50% of scale. In other words, half of the total input power to the system is consumed in delivering the remaining half in a controlled manner to the load. The efficiency may even be worse if the load is reactive due to the phase angle between the output voltage and current. It's the power consumed in the controlling element that we want to eliminate to improve efficiency.

Another way to deliver power to the load is with a PWM technique *(Fig. 1b)*. The analog input signal is converted to a variable-duty-cycle switched drive signal by the PWM control block. When a larger output voltage is required, the switch is merely held on for a longer portion of the fixed period. The switch is both on and off during each cycle of the basic switching frequency.

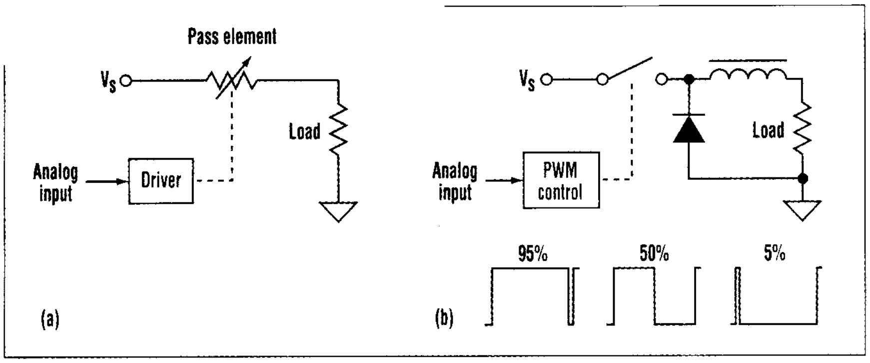

1 There are two basic ways of delivering power to a load. The linear approach reduces the resistance of the pass element as more output voltage is needed, approaching zero resistance when maximum output voltage is required (a). Another way to deliver power to the load is with a PWM technique (b).

Usually, this is all happening at switching frequencies above 20 kHz. The switching frequency can be viewed as a carrier frequency modulated by the duty-cycle information corresponding to the analog input signal. The inductor filters the output switching waveform to produce an average output voltage. The on-resistance of the switch is responsible for most of the power that's lost. The flyback diode and the resistance of the inductive load accounts for the rest.

The function of the inductor is to store energy during the ON portion of each switching cycle and release some of that energy during the OFF portion of each cycle, and therefore acts as a filter. The load sees little of the switching frequency, but responds to signals significantly below the switching frequency. Inductive loads such as motors are often their own filters, and no dedicated external filtering is required to demodulate the PWM carrier frequency. The output of the PWM amplifier is always near zero or near the supply voltage. And because little voltage is dropped across the output, the device's efficiency is always high compared to the linear amplifier. The filtered outputs of PWM amplifiers are typically 80 to 95% efficient.

To compare the linear and PWM approaches on more concrete terms, consider various design choices, each capable of delivering 1 kW to the load. A 90%-efficient PWM circuit generates 50 W of waste heat when delivering half output of 500 W to the load. The theoretically perfect linear circuit will generate 500 W of waste heat while delivering the same 500 W. When both the linear and PWM amplifiers are delivering maximum output, their efficiencies are much more similar, although the PWM amplifier still has the edge. Typically, the PWM amplifier has about a third or less voltage drop than a comparable linear amplifier at similar output currents.

Let's look at three possible approaches to this type of design *(see the Table, p, 52)*. In all three cases, it's assumed ambient temperature is 30°C and maximum case temperature is 85°C. It's also assumed the power rating of the TO-3 devices is 125 W each.

Heat sinks for linear designs require multiple sections mounted in such a manner that heat removed from one section doesn't flow to other sections. The linear approaches requires five times the heat-sink rating of the PWM approach. The bad news with the hybrid linear design is that the heat is concentrated in such a small area that this design is on the edge of requiring liquid cooling. With its high package count, the discrete linear approach will likely have more than five times the heat-sink size and weight of the PWM amplifier. It's clear that the PWM amplifier design approach can save significantly in space, hardware, and operating costs.

How The PWM Amp Works

The Apex SA01 amplifier consists of four basic building blocks: An error amplifier, a ramp generator, the PWM generator, and the H-bridge output drivers *(Fig. 2)*. The difference between the command signal and the feedback signal is integrated by the error amplifier. Its output voltage via feedback will become the precise voltage needed to drive the PWM block to the required duty cycle that will close the loop and reduce the error to near zero.

A triangle wave is produced by the ramp generator at the switching rate of the amplifier (42 kHz in the case of the SA01). This is the carrier frequency referred to earlier.

The PWM generator circuit converts the error amplifier output into a variable-duty-cycle signal *(Fig. 3)*. Another way of looking at it is that the PWM generator modulates the carrier with the analog signal. In Figure 3, the comparator is used to generate a duty cycle by comparing a reference triangle wave with the analog input signal. For each cycle of the ramp, an ON and OFF pulse is generated each time the triangle crosses the analog signal level. In the SA01, two comparators are connected so that complementary duty cycles are generated for each crossing of the input signal and triangle. These signals ultimately drive the H-bridge output MOSFETs.

The H-bridge switches work in pairs to reverse the polarity of the load driver even though only a single-polarity supply is used *(Fig. 4)*. Q1 and Q4 conduct during one portion of each cycle, and Q2 and Q3 conduct during the remainder of the cycle.

To further illustrate the modulation technique and the workings of the H-bridge switching, suppose that the load is a brush-type dc motor. Furthermore, suppose that the PWM generator is producing a drive signal that turns on Q1 and Q4 for 80% of the period of each switching cycle (1/f). Then Q2 and Q3 must be conducting for the remaining 20% of each cycle.

Let's also assume that when Q1 and

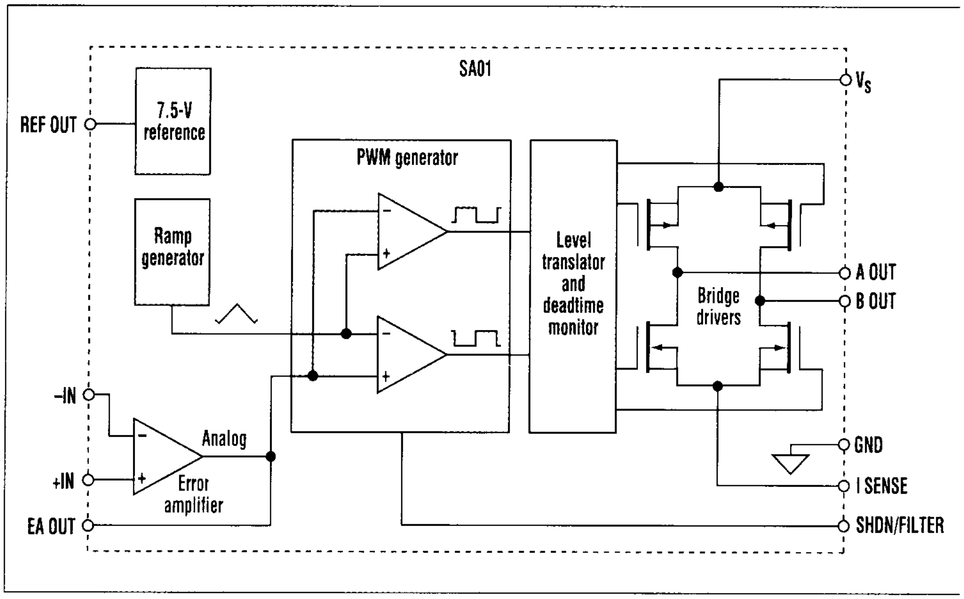

2 **The Apex SA01 amplifier consists of four basic building blocks: An error amplifier, a ramp generator, the PWM generator, and the H-bridge output drivers.**

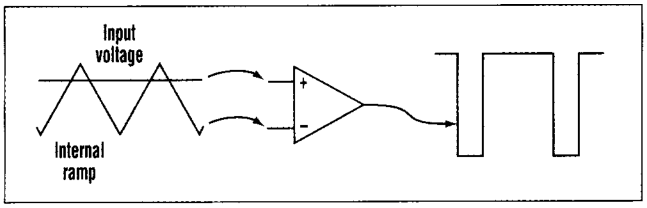

3 In this circuit, the comparator is used to generate a duty cycle by comparing a reference triangle wave with the analog input signal. For each cycle of the ramp, an ON and OFF pulse is generated every time the triangle crosses the analog signal level.

Q4 conduct, they tend to drive the motor's shaft clockwise (CW). The motor then is being driven CW for 80% of each cycle and counterclockwise (CCW) for the remaining 20%. As you might expect, the motor can't respond within one period, and therefore filters out the carrier. The motor responds by rotating its shaft at a rate equivalent to a dc signal that's 60% of power supply voltage (80% CW – 20% CCW, neglecting losses). If Q1 and Q4 were ON for 50% of the switching period, Q2 and Q3 would likewise be ON for 50% of the period, and there would be no net motor shaft rotation.

With the basics out of the way, it's now time to consider some actual applications. We'll examine a high-quality audio application and a motor-torque control circuit.

The Apex SA02 can be configured as a high-quality audio amplifier *(Fig. 5)*. Its switching frequency of 250 kHz makes it particularly suitable for this application because that switching frequency results in a usable audio bandwidth of 25 kHz. Because most of the building blocks and other necessary components are built-in, only a few external components can scale and otherwise complete the circuit.

The SA02's maximum bridge supply voltage is 80 V. At full-scale-drive (160 V p-p, 100% modulation), about 56 V rms would be applied to the 8-Ω speaker, producing roughly 400 W of output power. (This is an approximation.) There will be some losses due to the on-resistance in the output transistors, the losses in the filter inductors, and the fact that the SA02 is limited to a linear modulation range of about 3 to 97% because of some internal timing overhead constraints. Still, 400 W of output power would qualify this circuit as a super-power amplifier that can't be matched by many linear amplifiers.

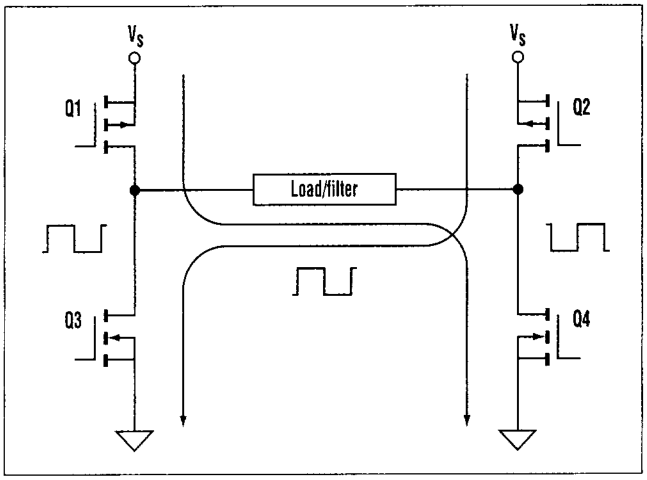

4 H-bridge switches work in pairs to reverse the polarity of the load driver even though only a single-polarity supply is used. Q1 and Q4 conduct during one portion of each cycle, and Q2 and Q3 conduct during the remainder of the cycle.

Differential amplifier A2 is scaled with a gain of 0.015, resulting in an output voltage swing of 2.5 V p-p when the SA02's outputs are modulated 0 to 100% of scale. Capacitors C2 and C3, along with the scaling resistors R3 and R5, form the 20-kHz passband of the circuit. Resistor R5 is returned to a 2.5-V reference instead of ground because A2 is operated as a single ended amplifier with common-mode input-voltage restrictions. In addition, its output always needs to be a positive voltage to match the single ended operation of error amplifier A3.

Note that the feedback point is not across the load, which is typically where the feedback point would be taken if this were a linear-amplifier design. Theoretically this could be done, but not here because the extra phase shift produced by the filter components makes loop stability very difficult. The difference between taking the feedback point from the load as opposed to taking it from the outputs amounts to the losses produced by the filter inductors, a factor that should be very low in a good filter design.

Although the SA02 is capable of sensing the current in each half of the H-bridge, there's no need for that here and the current-sense pins are tied together. R2, the current-sense resistor, was chosen to current-limit the amplifier at about 10 A. This is calculated from the trip voltage of 100 mV and the value of R2, 100 mV/10 mΩ = 10 A. The trip voltage is set internally and compared with the I_{SENSE} voltage at the I_{LIM}/Shdn pin. R1 and C1 filter the switching waveform that appears across R2. Because a good averaging circuit is needed to filter out the switching spikes, it's advantageous to make the time constant of R1 C1 about 10 times the period of the switching frequency, or about 40 µs.

In order to 100% modulate the SA02, the output of integrator A3 must swing approximately 1.25 V to 3.75 V because this is the peak-to-peak (p-p) swing of the ramp voltage. The ramp (which can be monitored at pin 4) is applied to the PWM generating comparator A4. Because A3 is an op amp, it can be operated with any dc gain desired, but here it's used purely as an integrator. Its very high gain ensures loop accuracy. Together, A3 and A4 produce the analog-to-PWM conversion similar to that shown in Figure 3.

Resistors R7 and R8 establish the gain needed between the input signal and the output swing of A2. We've already determined that the output swing of A2 needs to be 1.25 to 3.75 V. Within the operating limitations of A2, this particular range is somewhat arbitrary, but convenient, because it represents the SA02's 0 to 100% modulation range. To scale R7 and R8, suppose that the audio input signal is 1 V p-p. The span of A2's output is 2.5 V. The ratio R7/R8 must therefore be 2.5 because that represents the ratio of A2's input-voltage and output spans.

A3's noninverting input has been biased to 2.5 V to keep the inputs within their common-mode voltage range. When the input signal is at zero volts, the ratio of R7/R8 would not provide 2.5 V at the noninverting input as required to keep A3 operating correctly. This implies that some offset will have to be introduced to balance the circuit. The necessary offset current can be supplied by connecting R9 to a 5-V reference

An equivalent circuit can be used to visualize and calculate the scaling requirements of R7, R8, and R9 *(Fig. 6)*. One method to calculate the values of these components is as follows: Assign values to R7 and R8 such that R7/R8 gives the gain required; in this case, 2.5. Choose these values in accordance with other circuit requirements such as the input impedance of the circuit. In this example, setting R7=10k and R8=4k meets the circuit requirements.

Next, apply the zero-scale input voltage, –5 V in this example, and calculate the required offset to produce zero-scale output (3.75 V). Note that as the input voltage goes from a smaller to a larger value, the output of Figure 6 goes from a higher to a lower value. A convenient way to calculate the offset is to note that the inverting input must stay at 2.5 V if the op amp is to be an op amp. With the zero-scale input voltage and the required output voltage both known, calculate the branch currents and determine the value of R9 that produces the required output offset voltage.

To complete the circuit, the output of A3 is connected through a low-pass filter, formed by C7 and R10, to the +PWM input. A4 compares this signal to the ramp at pin 4 to produce duty-cycle information. The low-pass circuit is optional, but it helps to reduce switching noise from getting into the A4 comparator. Keep the value of R10 low because it becomes part of the hysteresis generating loop around A4. The values of R10 and C7 should provide a passband below 20 kHz if that is the bandwidth required by your application. This circuit provides a small and efficient high-power audio amplifier.

An Apex SA03 can be configured as a motor-torque control driver *(Fig. 7)*. From a block diagram standpoint, the SA03 operates just like the SA02 audio amplifier, except that the SA03 doesn't have an internal error amplifier. A2 fills this function.

The differential amplifier A1 is set up in a gain of ten, and senses each half of the output bridge current via the SA03's I_{SENSE} pins. Resistors R1 and R4, respectively, sense the current in each half of the output bridge. The combined effect of the two sense resistors is to give both amplitude and direction data.

The sense voltages are filtered by the R2-R3-C1 network. The common

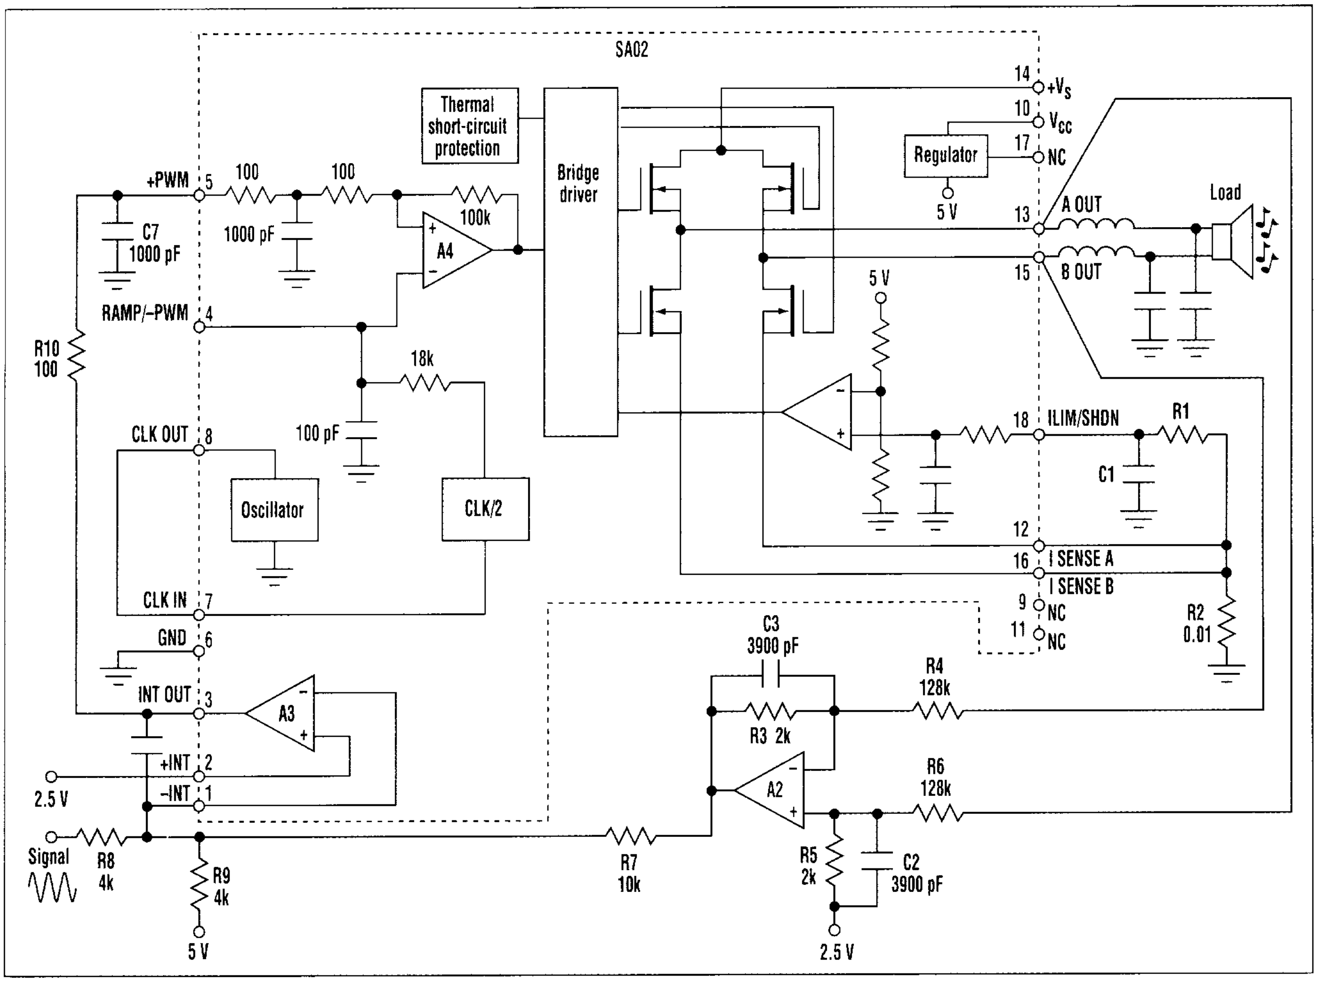

5 The SA02 an be configured as a high-quality audio amplifier. Its 250-kHz switching frequency results in a usable audio bandwidth of 25 kHz.

node connection at R2 and R3 sums the signals. The combined current-sense signal is applied to the I_{LIM}/Shdn pin, which has a trip point that's internally set at 100 mV. Because the sense voltages at the I_{SENSE} pins are divided down, the I_{LIM}/Shdn pin sees only half of the voltage at the I_{SENSE} pins. Full-scale current limit now occurs when either I_{SENSE} pin develops 200 mV.

Each sense voltage is also connected to the differential amplifier A1. Similar to the SA02 application, the feedback components determine the passband of the amplifier. Motor drive circuits normally have passbands under 100 Hz, so the SA03's relatively low 22.5-kHz switching frequency (compared to the SA02) is not detrimental. A1's output swings 3 to 7 V, depending on the magnitude and direction of current in the motor that's due to the noninverting input being returned to a 5-V reference and the gain of 10. The range of 3 to 7 V represents the SA03's 0 to 100% modulation range in a manner similar to that discussed in the SA02 application.

The scaling from input voltage to output current is set by R9, R10, and R11. This relationship is similar to the gain arrangement in the SA02 audio amplifier. The SA03 in this application circuit is capable of delivering 100 V and 30 A to the motor.

Now you understand the PWM technology and have seen the theoretical side of putting two application circuits together. There is, of course, the practical side of actually building these kinds of circuits. For those inexperienced with amplifiers offering these power levels, there's plenty to know. The amps, watts, and switching speed that these amplifiers produce are awesome. Several design areas need particular care. We will extend techniques of the internal design to the external application circuit as well.

PCB LAYOUT RULES

The application designer needs to appreciate that PWM amplifiers combine both high-speed, high-power switching and low-level analog signals in one circuit. There is little leeway for layout compromises. Still, as demanding as the engineering requirements are, successful PCB designs are well within reach with reasonable care and a little elbow grease. Ordinary layout design techniques that are acceptable for most circuits will almost certainly cause major problems in power-switching circuits. The design rules of autorouting layout programs may have to be changed to avoid creating some serious problems. The following layout "rules of thumb" will help you jump these hurdles:

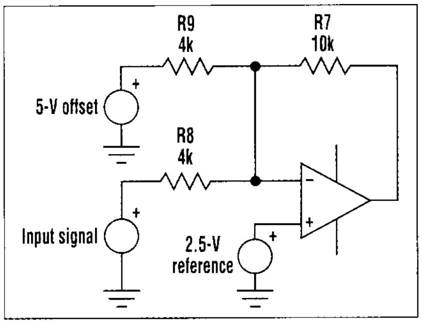

6 This is an equivalent circuit used to visualize and calculate the scaling requirements of R7, R8, and R9 in the audio amplifier.

1. Bypassing of the power supplies is critical. Connect capacitors directly to the power-supply pins with very short lead lengths (well under one inch). Ceramic chip capacitors are best. Bypass capacitors with long lead lengths simply will not work.

2. Make all ground connections with a star pattern at the ground pin of the amplifier. Single-point grounding is imperative. The dI/dt produced by switching circuits causes even short lengths of conductors to produce substantial voltages that interfere with low-level analog circuits.

3. Beware of capacitive coupling between output connections and signal inputs. This occurs through the parasitic capacitance between layers of multilayer PCB designs.

Capacitive coupling also might include trace-to-trace coupling with a solder mask acting as a capacitor dielectric. PWM-amplifier outputs can switch several hundred volts in 50 ns. It's not hard to imagine (or calculate) that a few pF of coupling between outputs and analog inputs will wreak havoc with your circuit.

4. Do not run small signal traces between output pins.

5. Do not allow high currents to flow into the ground plane.

6. Separate switching and analog grounds, and connect the two only at the ground pin of the amplifier as part of the star pattern.

Grounding is particularly picky in these circuits due to all the RFI and possible ground loops caused by the high-speed, high-current switching. In practice, the analog ground return should have very little or no current in it. If your circuit has low-level logic, separate the logic ground return from the analog and power grounds. Ideally, you want no step-change signals in the analog ground return. The power ground connecting the bypass capacitor should be as short as possible and as wide as possible to lower IR drop and keep inductance to a minimum.

POWER-SUPPLY BYPASS

The most common fault in application circuits is poor bypassing of the power supplies. It's difficult to overemphasize this aspect of the PWM application design. Without proper bypassing, your circuit is headed for trouble.

Consider both the high-frequency and moderate- to low-frequency aspects of the bypassing. The PWM amplifier's outputs can switch in 10 to 50 ns. The lead inductance connecting the power supply to the amplifier thus prevents the supply voltage from becoming stabile without proper bypassing. This is a job for a ceramic capacitor because of its low ESR and low self inductance. A type-X7R ceramic capacitor is recommended. Its job is to supply the load current and stabilize the supply voltage during the switching interval. The best choice is a surface mount ceramic chip capacitor because there is no lead length and therefore no lead inductance (Novacap makes 1-µF models ranging up to 500 V).

Low-ESR aluminum electrolytic or tantalum capacitors can be used to bypass the lower-frequency components of ripple. Use a ceramic bypass capacitor with a value of at least 1 µF. Back that up with an electrolytic capacitor that's 1 to 10 µF in value per amp of load current. Use electrolytic capacitors rated for switching applications. The low-ESR aspect of the capacitors is extremely important. The high-frequency ripple current in a capacitor with a high ESR value may get so hot it explodes.

Connect the bypass capacitors directly between the amplifier's supply and ground pins. To help reduce the lead length, mount the bypass capacitors on the bottom side of the circuit board. Keep the lead lengths well under one inch to prevent excessive ringing at the outputs.

OUTPUT FILTERING

A PWM amplifier driving a resistive load with no filtering is unable to modulate the output voltage—it can only switch polarity. Loads or filters with low inductance may overheat with high ripple current (I p-p). Set a design limit for I p-p, and calculate minimum total inductance:

$$L=V_S/(2\times f\times I\ p\text{-}p)$$

where L is the minimum inductance, V_S is the supply voltage, and f is the switching frequency. This is the maximum peak current that the inductor will see. After the filter is charged up, the ripple current will shrink to about one half this value.

As an example, the SA01 (switching at 42 kHz) on 100 V needs 300 µH to keep I p-p down to 4 A p-p. This inductance could be a motor without specific filter components, or it could be an L-C network on each output using 150-µH inductors.

TROUBLESHOOTING

When an engineer is unfamiliar with troubleshooting power-switching circuits, he is often appalled by all the raucous and ragged wave forms on his oscilloscope. The ringing and ripple seem so unsanitary. He might be more accustomed to the pristine sine waves and damped square waves of low-bandwidth analog circuits. Are all those ugly waveforms really there? If you haven't followed the PCB layout guidelines, ugly waveforms will be the least of your problems. However, touching the probe tip to the ground clip may reveal another common problem. If the scope shows a waveform with a "grounded" input (all high-impedance nodes appear to have spikes), there are several possible causes.

The most common problem is that ubiquitous 3- to 6-in. oscilloscope ground lead. It's forming an inductive pickup loop for all that high-frequency current the PWM amplifier is moving around. Many scope accessory kits offer an RF adapter capable of providing a ground lead less than 1/4-in. long. If the adapter is not in the kit, consider buying one or making your own from a length of spring wire. Shortening the scope grounding lead will dramatically improve the waveforms you observe.

Another possibility is that the local ground of the PWM amplifier may be quite different from the local ground of the oscilloscope. Try disconnecting all other signal cables to the scope to remove any interaction with other local grounds. Try a battery-operated scope if one is available. You might

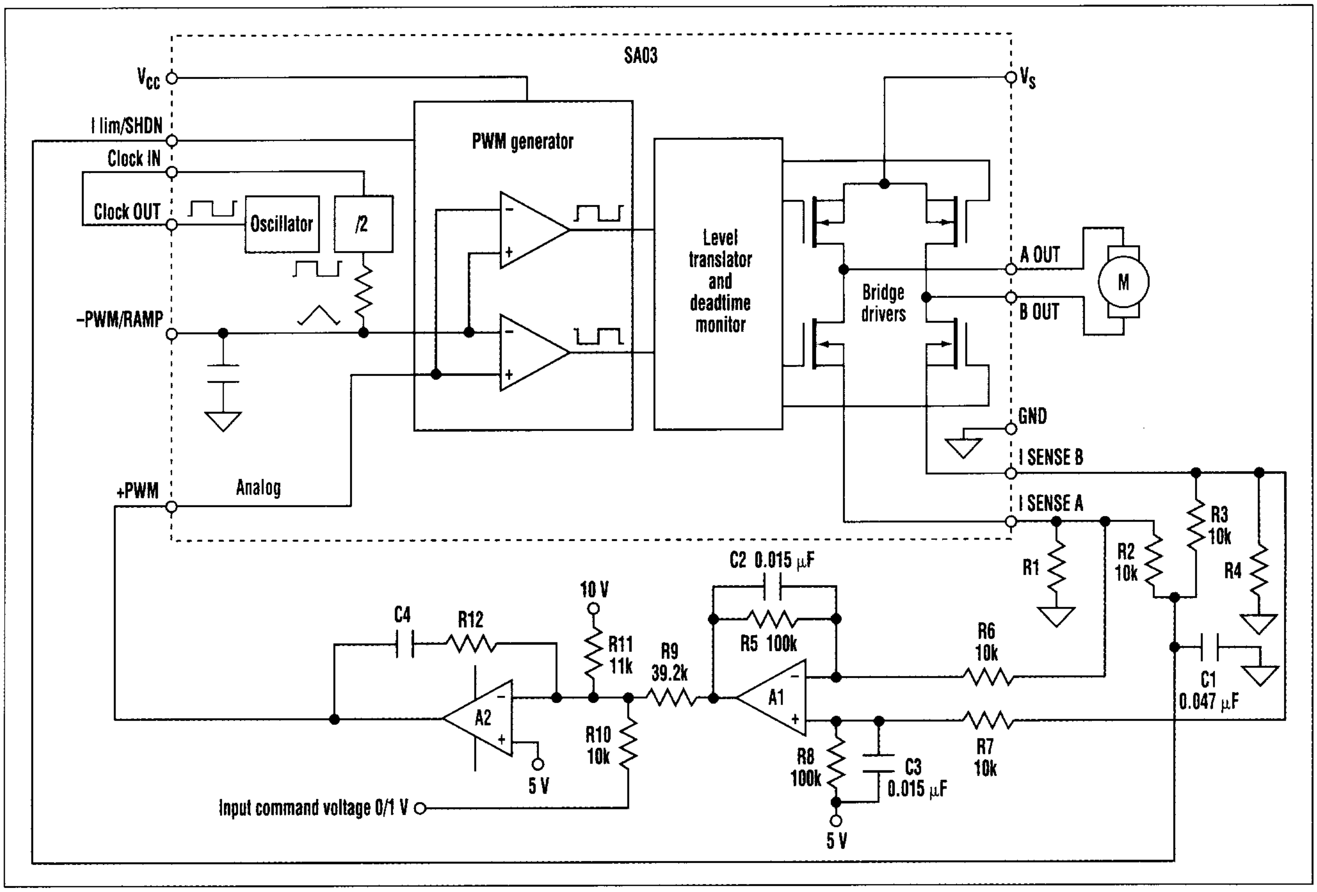

7 An SA03 can be configured as a motor-torque control driver. From a block diagram standpoint, the SA03 operates just like the SA02 audio amplifier, except that the SA03 doesn't have an internal error amplifier. Instead, A2 fills this function.

also try a groundbreaker installed on the oscilloscope's power cord. Make sure that the oscilloscope is grounded at the same point as the other circuitry (at that star pattern on the ground pin as mentioned earlier).

Minimize the use of extenders, grabbers, or clips. Try using only shielded probes. High-impedance nodes are the most susceptible to capacitive coupling from high-slew-rate signals, and switching amplifiers have plenty of these signals that will get you in trouble.

Also highly recommended is a current probe: Either (or both) the Hewlett Packard 1146A and Tektronix AM503A. These probes are easily connected to the oscilloscope and are particularly useful in examining what's happening in the load in terms of average current and peak ripple current. A digital storage oscilloscope (DSO) is not recommended unless the sampling rate is 1 GHz or better.

INTERNAL DISSIPATION

Although pulse-width-modulation amplifiers are considerably more efficient than their linear counterparts, thermal management issues remain the same. However, there are two major differences in the thermal aspects of linear power amplifiers and PWM amplifiers. First, power in the PWM amplifier due to loading can be calculated without knowing the output voltage or the supply voltage. This is simply because no matter what the supply voltage, the outputs always switch to the supply voltage and the power dissipated is the product of the switch's on-resistance and the load current.

The second difference is a little more subtle, but affects the very reason that a PWM amplifier is used: Efficiency drops rapidly as junction temperature increases. Heat sinking, then, is more than a reliability issue. Thermal design of the PWM amplifier has a first- order effect on circuit performance.

Like linear power amplifiers, nearly all of the PWM amplifier's power dissipation occurs in the output transistors. Quiescent power-dissipation overhead is only a small fraction of the total. A first-order calculation of power dissipation due to loading involves only the output current and the amplifier's total on-resistance.

COMPARISON OF AMPLIFIER DESIGNS

	Discrete linear	Hybrid linear	Hybrid PWM
Waste heat	500 W	500 W	100 W
Package count	16 × TO-3	2 × Apex PA03	1 × Apex SA01
Heat sink	0.11°C/W	0.11°C/W	0.55°C/W
Cost		$700	$250

HEAT-SINK REQUIREMENTS

For each amplifier model, the maximum total on-resistance at 25°C (R_{ON}) will be specified in its datasheet performance table. Typically in these amplifiers, the output transistors are MOSFETs. MOSFET on-resistance usually doubles between junction temperatures of 25°C to 150°C. Choose a maximum junction temperature consistent with your design standards, and estimate the R_{ON} at the junction temperature chosen. $I^2 \times R_{ON}$ now yields power due to loading.

With the total power dissipation (P_{TOTAL}) now known, determine the heat sink requirement. Again, consistent with your design standards, choose a maximum case temperature (T_{CMAX}). Determine the maximum ambient temperature (T_{AMAX}). R_{CS} is the thermal resistance of the package to the heat-sink interface. The lowest R_{CS} can be obtained from thermal grease or an aluminum thermal washer available from Apex.

$$R_{SA}=[(T_{CMAX} - T_{AMAX})/P_{TOTAL}] - R_{CS}$$

where R_{SA} is the required minimum heat-sink thermal resistance.

The last item to check is the junction temperature (T_J). Find the thermal resistance of the amplifier (θ_{JC}) in the performance table of the data sheet.

$$T_J = P_{TOTAl} \times \theta_{JC} + T_{CMAX}$$

As an example, consider an SA03 with a R_{ON} maximum of 16 mΩ delivering up to 15 A at maximum ambient of 35°C. The thermal resistance θ_{JC} of the SA03 is 1°C/W. Design rules allow case and junction temperatures up to datasheet maximums.

$P_{TOTAL} = 15^2 \times 0.16 = 36\text{ W}$
$T_{CMAX} = 85°\text{C}$
Allow 0.1°C/W for R_{CS}

$$R_{SA} = [(85°C - 35°C)/36\text{ W}] - 0.1°C/W = 1.29°C/W$$

$$T_J = 85°C + 36\text{ W} \times 1°C/W = 121°C$$

FEATURES

The features offered in the new PWM amplifier families differ, thus some models are more suited for certain applications than others. For example, some of the models do not have an error amp, discussed previously as one of the building blocks of PWM amplifiers, and so must be supplied externally. Here are the main features offered and what they do for the user:

Dual current sensing—The low side of each half of the H-bridge is accessible. Current-sense resistors tied from these points to ground provide both amplitude and direction information. This is a very useful feature for motor-torque-control circuits.

Synchronizable clock—Clock input and output pins are provided. Normally, these pins are tied together. But it's possible to assign one amplifier as the master and have its clock drive the clock input pins of other similar amplifiers so that they all switch at the same frequency. It's also possible to use an external clock of a lower frequency as needed by the application.

Logic inputs—Several models have inputs that are normally used for the analog aspects of the amplifier, but can be overdriven by logic inputs to control the duty cycle. For example, these inputs can be connected to motion-control processors.

External shutdown—An external logic signal applied to this pin turns off all of the output transistors. It can be used with external protection circuits or as an aid in lowering quiescent power consumption when the amplifier is idling on battery power, for example.

Protection circuits—Several models protect against short circuits between output and ground, or output to output, by sensing the current in the high-side switches and latching the outputs off until the fault is cleared and power is recycled. Thermal overloads produce similar results by directly measuring the silicon temperature of the output transistors. These features make the protected amplifiers virtually bulletproof.

Linear versus nonlinear ramp—Some models offer a highly linear ramp. In others, the ramp is merely the voltage across the capacitor in a RC circuit. At first glance, you may think that the linear ramp would be far superior for linear modulation of the output. It turns out, however, that little is sacrificed with the nonlinear ramp: About 0.7% worst case with no feedback and much less, of course, with feedback. Most of the nonlinearity produced when the ramp is going positive is compensated for when the ramp swings back negative. The nonlinear-ramp approach is lower in cost and has the advantage that it's easy to overdrive with external signals as mentioned above.

Hybrid PWM amplifiers offer highly reliable solutions in high-power motor drive, motion-control, or audio applications where linear designs would otherwise require larger sizes or an inordinate amount of heat-sinking hardware. The hybrid approach coupled with rugged hermetic packaging best suits the demanding requirements of military or high-reliability commercial applications, such as aircraft systems. By applying a few rules of thumb for bypassing and PCB layout, the PWM amplifiers are easy to use. The arrival of the new hybrid PWM amplifiers speeds the design process and greatly enhances fault tolerance by offering protection circuits simply not possible in discrete or module implementations. PWM amplifiers meet the Tool Man's cry for more power.

DENNIS EDDLEMON, a senior design engineer at Apex Microtechnology for 12 years, holds a BA in sociology from Valparaiso University, Indiana, as well as four patents in hybrid and monolithic power amplifier design. He can be reached at Apex Microtechnology Corp., 5980 N. Shannon Rd., Tucson, AZ 85741-5230; (520) 690-8600; fax (520) 888-3329.