APPLICATION NOTE 30

PULSE WIDTH MODULATION AMPLIFIER

HTTP://WWW.APEXMICROTECH.COM (800) 546-APEX (800) 546-2739

1.0 INTRODUCTION

This note is divided into three sections. The first section provides general information on Pulse Width Modulation amplifiers and examines a typical block diagram. The second section on designing with PWM amplifiers is NOT intended for optional reading. The family of PWM amplifiers are not equal in protection features and some of the design errors that would cause a linear amplifier to oscillate will destroy some of the less protected PWM amplifiers. The final section examines some ways to use PWM amplifiers.

PWM circuits are taking the same general course of development traveled by op amps and many other electronic functions. Concepts were brought to life using discrete components and were followed by modules, hybrids and then monolithics. Since the introduction of the first hybrid PWM amplifiers from Apex, the product family has broadened to include a wide range of power levels and package types. The family includes models with up to three levels of internal protection circuitry and low cost models that feature no frills PWM power. With digital or analog input capability on most amplifiers, the Apex PWM family can fit the requirements of many applications.

2.0 WHY PWM

As power levels increase the task of designing variable drives increases dramatically. While the array of linear components available with sufficient voltage and current ratings for high power drives is impressive, a project can become unmanageable when calculation of internal power dissipation reveals the extent of cooling hardware required. A 20A output stage often requires multiple 20A semiconductors mounted on massive heatsinks and usually employs noisy fans or liquid cooling in some cases.

Figure 1 illustrates the linear approach to delivering power to the load. When maximum output is commanded, the driver reduces resistance of the pass element to a minimum. At this output level losses in the linear circuit are relatively low. When zero output is commanded the pass element resistance approaches infinity and losses approach zero. The disadvantage of the linear circuit appears at the midrange output levels and is often at its worst when 50% output is delivered. At this level, resistance of the pass element is equal to the load resistance which means heat generated in the amplifier is equal to the power delivered to the load! We have just found the linear circuit to have a maximum efficiency of 50% when driving resistive loads to midrange power levels. When loads appear reactive this efficiency drops even further.

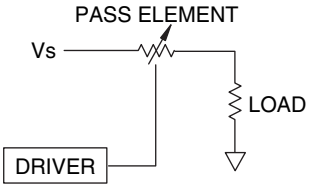

FIGURE 1. LINEAR POWER DELIVERY

Figure 2 illustrates the most basic PWM operation. The PWM control block converts an analog input level into a variable duty cycle switch drive signal. As higher outputs are commanded, the switch is held ON longer portions of the period. Normally, the switch is ON and OFF once during each cycle of the switching frequency, but there are many designs capable of holding a 100% ON duty cycle. In this case, losses are simply a factor of the ON resistance of the switch plus the inductor resistance. As less output is commanded the duty cycle or percent of ON time is reduced. Losses include heat generated in the flyback diode. At most practical supply voltages this diode loss is still small because the diode conducts only a very small portion of the time and this voltage drop is a small fraction of the supply voltage.

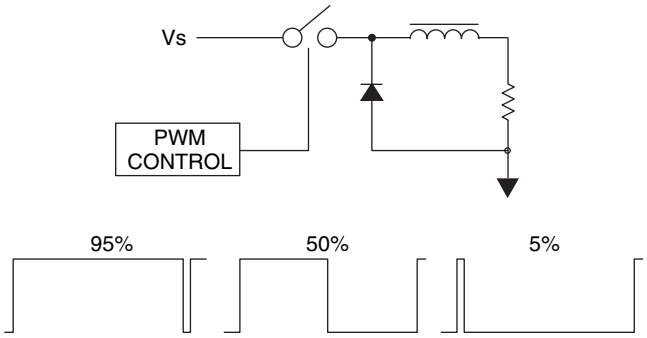

FIGURE 2. PWM POWER DELIVERY

The job of the inductor is storing energy during the ON portion of the cycle for filtering. In this manner the load sees very little of the switching frequency but responds to frequencies significantly below the switching frequency. A rule of thumb is to expect a usable bandwidth to be one decade below the switching frequency. Inductive loads often provide adequate filtering without dedicated filters.

With the PWM circuit, the direct (unfiltered) amplifier output is either near the supply voltage or near zero. Continuously varying filtered output levels are achieved by changing only the duty cycle. This results in efficiency being quite constant as output power varies compared to the linear circuit. Note that efficiency claims on the hybrid PWM amplifier data sheet do not include filter losses. Typical efficiency of filtered PWM circuits range from 80 to 95%.

Almost all power amplifiers (low duty cycle sonar amplifiers are a notable exception) must be designed to withstand worst case internal power dissipation for considerable lengths of time compared to the thermal time constants of the heat sinking hardware. This forces the design to be capable of cooling itself under worst case conditions. Conditions to be reckoned with include highest supply voltage, lowest load impedance, maximum ambient temperature, and lowest efficiency output level. In the case of reactive loads, maximum voltage-to-current phase angle (lowest power factor) must also be addressed.

Consider a circuit delivering a peak power of 1KW. A 90% efficient PWM circuit generates 100W of waste heat when running full output and around 50W delivering half power. The theoretically perfect linear circuit will generate 500W of waste heat while delivering 500W. Table 1 shows three possible approaches to this type design. In all three cases it is assumed ambient temperature is 30°C and maximum case

	Discrete Linear	Hybrid Linear	Hybrid PWM
Waste heat	500W	500W	100W
Pkg count	16 x TO-3	2 x PA03	1 x SA01
Heat sink	0.11°C/W	0.11°C/W	0.55°C/W

TABLE 1. CONTRASTING DISCRETE LINEAR, HYBRID LINEAR AND HYBRID PWM 1kW DESIGNS

temperature is 85°C. It is also assumed power ratings of the TO-3 devices is 125W each. Heatsinks for linear designs require multiple sections mounted such that heat removed from one section does not flow to other sections. The linear approaches require five times the heatsink 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 right on the edge of requiring liquid cooling. With its high package count the discrete linear approach will likely have more than five times the heatsink size and weight of the PWM.

3.0 HOW IT WORKS

The simple form of PWM circuit examined thus far is very similar to a number of switching power supply circuits. If the control block is optimized for producing a wide output range rather than a fixed output level, the power supply becomes an amplifier. Carrying this one step further we get the PWM circuit employing four switches configured as an H-bridge providing bipolar load current from a single supply. This does mandate that both load terminals are driven and zero drive results from 50% of supply voltage on both load terminals. See Figures 3 and 4 for the basic bridge operation and typical waveforms.

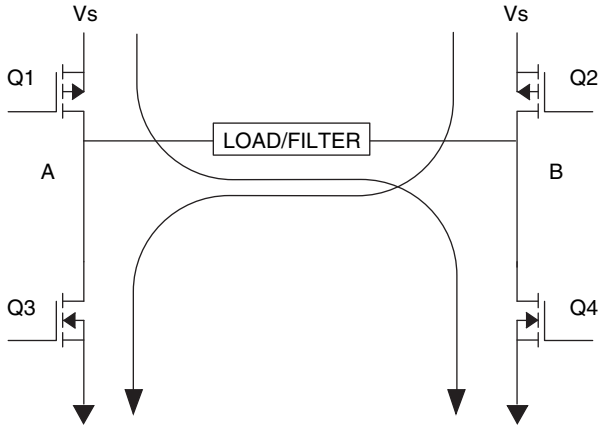

FIGURE 3. BIPOLAR OUTPUT OF THE BRIDGE

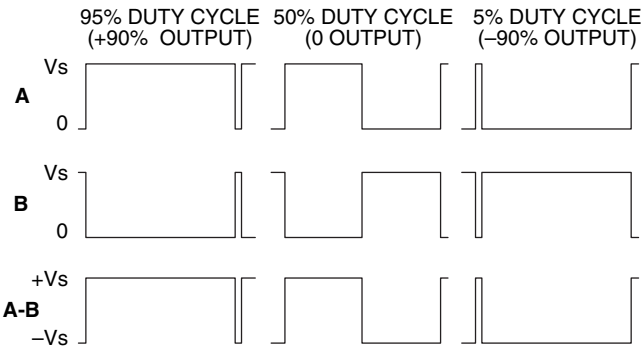

FIGURE 4. H-BRIDGE WAVEFORMS

3.1 HYBRID H-BRIDGE BASICS

The H-bridge switches work in pairs to reverse polarity of the drive even though only one polarity supply is used. Notice how the levels of the A-B waveform are different even though shape is identical to the A waveform. Q1 and Q4 conduct during one portion of each cycle and Q2 and Q3 are on during the remainder of the cycle.

Figure 5 shows a block diagram typical of Apex PWM amplifiers. The hybrid construction of these amplifiers can allow monitoring of temperature directly on the surface of each power die rather than case or heatsink temperature monitoring - the best that a discrete design could implement. Direct die temperature measurement eliminates thermal resistance variables and reduces response time by orders of magnitude. The thermal limit is set at approximately 165°C. Activation of the thermal shutdown circuit will latch all of the H-bridge switches off. Toggle the shutdown pin or cycle the power to reset the amplifier and resume normal operation.

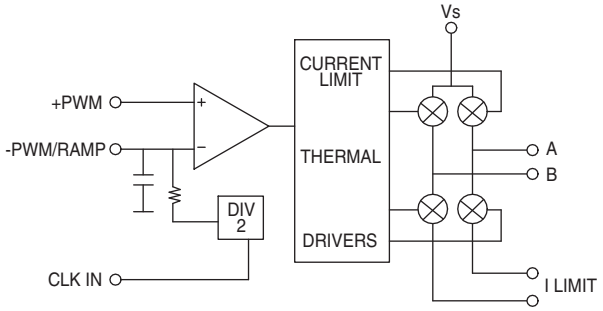

FIGURE 5. BLOCK DIAGRAM

3.2 CURRENT SENSE, CURRENT LIMIT AND CURRENT SHUTDOWN

The first of two current limits in the PWM block is the high side current limit that activates only upon output shorts to ground (assuming the programmable current has been properly configured.) This circuit has a variable response time based on the current magnitude in $+V_S$ line. With a fault current of about 1.5 times the rated output current it will require several cycles of the switching frequency to activate the circuit. As higher currents are sensed the response time decreases. Once a fault has been sensed the amplifier will remain latched off until power is cycled or the shutdown pin is toggled.

The second current limit circuit in the block diagram is programmable and activates upon a load fault or a short to the power supply. An external resistor senses current flowing between ground and the low side of the H-bridge. The sensed voltage is fed to the current limit pin. When this voltage exceeds the threshold (see individual datasheets for this value), all switches in the H-bridge are shut off for the remainder of the switching cycle. Because the sense voltage will have considerable spike content, the hybrid includes an internal filter stage. A second external stage of R-C filtering allows larger peak currents for any given value of current sense resistor. In some models, the current limit function can be used to shut the amplifier down on command of a logic level input voltage.

In most circumstances (half bridges are a common exception) all of the load current finds its way through a sense resistor to ground. This implies that the choice of resistor should be capable of handling the maximum load current defined by the current limit resistor value.

NOTE: Some amplifiers, such as the SA60 and others, do not have the current limit feature, but do allow the use of current sensing resistors for use in a current feedback loop. Because these amplifiers are unprotected, it can be tempting to use large resistance values for R_{SENSE} to limit the current ($I_{MAX} = V_{SUPPLY} \bullet R_{SENSE}$). However tempted you may be, DO NOT attempt to use high value resistors for R_{SENSE}. When you

do this, there is significant voltage developed across R_{SENSE} - between the source of the lower output FET and ground. This can allow the gate driver circuit to drive the gate of the FETs below the source enough to blow them.

3.3 ACCURACY - CLOSED LOOP

In closed loop applications (the only kind that have any accuracy), an integrating error amplifier is used to eliminate the difference between command signals and feedback signals. Its output voltage will go to the exact voltage required by the PWM block to generate the proper duty cycle corresponding to the desired output. The first job of the error amplifier is responding to input signal changes, but it also compensates other variables inside the feedback loop. Variations in supply voltage will require an adjustment of the input to maintain a output stability. On resistance of the H-bridge, resistance of the filter inductor and sometimes load resistance temperature variations are compensated. Systems such as speed controls may place mechanical factors such as conveyer belt load weight inside the loop where the error amplifier compensates the variations. Closing a voltage or current loop around PWM amplifier is non-trivial. Application Note 41 discusses closed loop PWM in great detail and is recommended reading for designers looking for precision control from their PWM circuit.

The PWM circuit converts the error amplifier output into a variable duty cycle drive signal that includes 0% and 100%. A dead time (all 4 of the H-bridge switches turned OFF) is inserted between each change of polarity at the output. This prevents “shoot through” current spikes caused by both H-bridge switches in the same leg of the H conducting at the same time. If these spikes were allowed to exist they would cause high stress and possibly destruction of both amplifier and power supply components.

3.4 DUTY CYCLE

There are many methods for converting an analog input into a PWM output. Of the methods used in Apex products, the following is most common and provides a good foundation for understanding any of the various techniques used by Apex and others. Accurate analog PWM generation is founded on a triangular reference signal (sometimes simply called RAMP). To provide a good reference, the ramp must have a linear rise and fall. The most economical method for generating this ramp signal is to start with a moderately stable square wave oscillator and use a divide by two circuit to ensure that the duty cycle is exactly 50%. An R-C network can convert this square wave into a reasonably linear triangle wave that is really exponential rise and fall. With carefully chosen peak voltages, the error caused by a non-linear triangle reference can be kept around 1% in an open loop system because most of the error on the rise is compensated on the fall. With a closed loop, the affect of this non-linearity is virtually unnoticeable.

To generate a PWM signal, the analog input voltage is compared to the triangle ramp reference. When the input signal is greater than the ramp voltage the A side of the H-bridge switches to V_S and the B side to GND. Similarly, when the input signal is less than the ramp, the B side switches to V_S. Figure 7 illustrates this principle. If the input is always within the peak to peak ramp voltage then an infinitely variable duty cycle can be achieved. If the input signal is outside the peak to peak ramp voltage there is no switching - the A side or the B side is V_S (0% or 100% duty cycle)

For the safety of the amplifier, the load, and bystanders in the general vicinity, it is important that the two switches on the same side of the H-bridge are never on at the same time. This would essentially create a short from V_S to GND and it

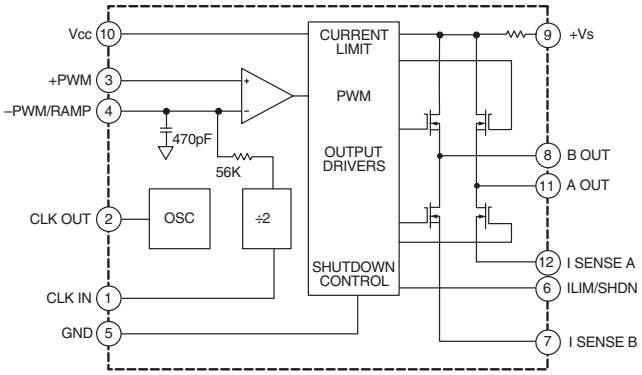

FIGURE 6. PWM CONTROL BLOCK

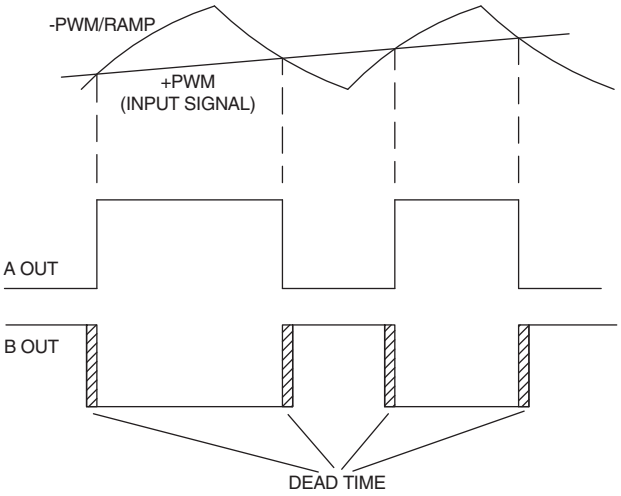

FIGURE 7. PWM WAVEFORMS

would not be pretty. Dead time is the solution to this potential problem. During the transition from the A side to the B side of the H-bridge, there is a period of time in which neither side is high - this is the dead time. Each amplifier has a fixed dead time based on the properties of the switches used inside. Dead time causes a small non-linear region in the variable duty cycle near 0% and 100%. Notice in each product datasheet that the output duty cycle will jump from approximately 97% to 100% (or 3% to 0%). Dead time also contributes a very small amount of inefficiency to the system. It is important to note that some models have variable switching frequencies. Since the dead time is not variable, efficiencies will drop as switching frequency increases - the dead time accounts for a larger percentage of the switching cycle.

In figure 7, note that the input voltage depicted is a straight line. The input voltage will obviously vary, but keep in mind that the useful bandwidth of the PWM amplifier is significantly less than the switching frequency. If the slew rate of the input voltage were allowed to approach that of the triangle wave, dead time circuitry could be ineffective at preventing shoot through.

The output waveform during dead time is primarily a function of load impedance. Current flow is interrupted by the dead time and the load or filter inductance will discharge its flyback energy at this time. While generally not shown in the block diagrams, each power switch has a diode to conduct the flyback current.

The outputs of this block labeled A OUT and B OUT do not directly represent high and low states of the two amplifier output pins. When the A Drive line is high it turns on the switch between the A output and the power supply and also

the switch between the B output and ground. When the A Drive is low, both these switches are off. B Drive controls the other two switches in the bridge.

4.0 FREQUENCY, FREQUENCY, FREQUENCY

In designing with and using PWM technology, the switching frequency is one of the critical system parameters. Note however, that most Apex PWM datasheets specify a clock frequency. Circuitry in the amplifier divides this clock frequency by two to provide the basis for the triangle or ramp signal. For example, the SA08 has an internal clock frequency of 45kHz and therefore has a switching frequency of 22.5kHz. The switching frequency is the basis for any filter design that might be required.

Some PWM amplifier models include CLK OUT and CLK IN pins. Generally these pins are simply connected together so the amplifier generates its own clock source. You can, however, use the CLK IN pin to synchronize the clock frequencies of multiple amplifiers in a system or to some other external source. In some cases this can be helpful for managing EMI. However, the divide-by-two function does not guarantee that the ramp signals of multiple devices are in phase, so simply sharing the same clock source does not allow multiple amplifiers to be used in parallel. A parallel PWM design is not trivial and is far beyond the scope of this text.

The third frequency important to consider in any PWM system is the signal frequency. As mentioned previously, the PWM signal must be averaged over time to achieve an analog type of output. The precision and cleanliness of the signal depends on how many output pulses are used to achieve the signal level. As a rule of thumb, the analog signal frequency should be at most 10% of the switching frequency. This relationship of signal and switching frequency is analogous to analog to digital conversion and the necessity to have many digital samples to accurately represent the changes in the analog signal.

4.1 CHANGING THE CLOCK FREQUENCY

At times it may be useful to reduce the switching frequency of an Apex PWM amplifier. In some amplifiers, this is possible with the addition of an external capacitor and a clock signal of the desired frequency. The Operating Considerations discussion in the datasheet will indicate if this is possible. The clock signal should be connected to the CLK IN input pin and should swing from GND to V_{CC}. The properly sized external capacitor connected between the -PWM/RAMP pin and ground will return the peak to peak ramp voltage to the datasheet specifications.

The exponential charge and discharge of the capacitor in the R-C network determine the clock frequency. However, only a simple ratio is required to calculate the size of the external capacitor for the new switching frequency. Note that the tolerance of the natural switching frequency is ±2%.

$F_{SW(natural)} \bullet C_{(internal)} = F_{SW(desired)} \bullet C_{(total)}$

So,

$C_{(external)} = ((F_{SW(natural)} \bullet C_{(internal)}) / F_{SW(desired)}) - C_{(internal)}$

SA14 Example:

Desired switching frequency = 10kHz

Natural switching frequency = 22.5kHz

Internal Capacitor = 470pF

Solving for C:

Total C = 1058pF

588pF required externally.

560pF is the closest standard value and corresponds to a switching frequency of 10.27kHz. If you insert a 10kHz clock signal, the peak to peak ramp voltage will vary slightly from the datasheet, but the center voltage will not change. Non-linearity will result if the ramp voltages are extended far beyond the datasheet specifications.

5.0 DESIGNING WITH PWMS

PWM amplifiers are high power switching devices whose voltage and current slew rates often surpass those found in either digital or analog circuits. Even though signal bandwidth may not top 1KHz, adopting the viewpoint of an RF designer can be very wise. Here are a few useful things to keep in mind:

** Wire inductance @ 20nH per inch

** Inductor voltage = dI/dt • L

** Capacitor current = dV/dt • C

** A good square wave = very large harmonic content

5.1 POWER SUPPLY BYPASS

It is difficult to over emphasize this aspect of the PWM design. Most of us are familiar with the good design practice of including a supply bypass capacitor at every IC in a low level logic design. If this is not done, the high switching rates cause problems on the power bus. As the most common fault in switching circuit design, inadequate bypass causes ripple and spikes on the supply line which make circuits inoperative and can even destroy components. Careful attention to location, size, ESR and ripple current capacity can result in a good design.

Power supply bypassing is a wideband job requiring at least two components for satisfactory operation of the amplifier. Use at least 10µF per ampere of load current to bypass the lower frequencies. Some applications appear to require many times this amount of capacitance. Capacitors with lower ESR ratings may ease the burden of finding space for such large devices. Locate this capacitor within a few inches of the amplifier. The high frequency bypass is absolutely critical! Think of frequencies in the 1 to 10Mhz range. Remember that many capacitors appear inductive in this range. Use ceramic capacitor(s) totaling 1µF to 10µF. Type X7R is recommended to ensure low ESR and ESL. Connect these capacitors directly between the supply and ground pins of the amplifier minimizing trace length, i.e. trace inductance. To illustrate the importance of this, consider a design having 3" between the supply pin and the ground point that terminates the capacitor: At 20nH per inch of parasitic trace inductance, the supply pin can have spikes equal to the supply voltage! When this happens, signal integrity is in question and peak voltages applied to components may be twice expected values. Connect the capacitor right at the amplifier pin.

The function of bypass capacitors is to satisfy AC current demands of the amplifier, which is isolated from the power supply by the inductance of the very same line that connects them. The degree of isolation increases with current magnitude, frequency and distance. When this isolation prevents current flow to the power supply, it must come from the bypass capacitors. Attempting to calculate capacitor currents is a questionable investment but ignoring them is no solution. Keep the requirement in mind when selecting components and follow up with temperature measurements on the prototype. Run the system at maximum frequency and power until temperatures stabilize. During this process, keep in mind that under-rated capacitors can explode.

5.2 HOW MUCH INDUCTANCE?

PWM amplifiers driving resistive loads with no filtering are unable to modulate the output voltage, they can only switch polarity. Loads with small amounts of inductance may over heat with high ripple current even with a 50% duty cycle (zero

output) drive. Other types of loads may suffer performance degradation if ripple currents exceed 1% or even 0.1% of the full-scale current. Once a design limit on peak-to-peak ripple current has been set, calculate minimum total inductance. It is proportional to supply voltage and inversely proportional to I_{P-P} and switching frequency:

$$L = V_S/(2 \bullet F \bullet I)$$

where V_S is the supply voltage and F is the switching frequency. As an example, this means the SA01 (switching at 42Khz) on 100V needs 300mH to keep ripple current down to $4A_{P-P}$.

5.3 GROUND CONCERNS AND LAYOUT CONSIDERATIONS

The circuit layout allows little room for compromise because the PWM amplifier combines both high-speed switching of large current and small-signal analog levels in one circuit. Though challenging, successful PBC layouts are not unreachable with some effort and care. However, some ordinary layout techniques and practices (including those used by auto-routing software) often cause major problems in switching applications, so be sure not to violate the following rules.

1) Power supply bypass. Connect all capacitors directly to the power-supply pins keeping trace/lead length as short as possible. Use ceramic capacitors for the high frequency content. Review the previous bypass discussion for more details.
2) Star ground. Make all ground connections in a star pattern with the ground pin of the amplifier at the center of the star. The dI/dt produced by switching circuits can produce substantial voltages over even short conductors that may interfere with low-level analog signals.
3) Separate small signal and power grounds. Connect your small signal ground the power ground star at only one point - the center of the star. If your circuit uses low level logic as well as small signal analog and high current PWM devices, keep these three grounds separate, connecting only at the center of the star.

Improper grounding can allow switching noise to enter the small signal analog channels disrupting the control of the signal. Using the star ground technique will ensure that high currents in one ground path do not induce voltages at other ground points. This is the only way to guarantee that the ground reference is clean and stable at all important points.

4) Avoid capacitively coupled feedback. Parasitic capacitance between traces or layers can couple unintended feedback from the output to the input section of the amplifier. With PWM outputs switching hundreds of volts in a few nanoseconds, it is not difficult to calculate what a few pF of parasitic coupling might do in the circuit.
5) Keep small signal traces away from outputs. High current outputs can magnetically couple to the small signal sections of your circuit. Never, ever, consider running small signal traces between the output pins.
6) Keep ground plane current to a minimum. If you are convinced that you must have a ground plane, do not under any circumstances connect high current return lines to it.
7) Keep traces to and from current limit resistors as short as possible. The current limit threshold voltage on most Apex PWM amplifiers is 100mV. Fast rise times in the current limit resistors and parasitic inductance can cause false current limiting or confuse the control logic in the device. This particular bit of the layout can cause high levels of hair loss if the proper attention is not given.

5.4 IS THE SCOPE TELLING THE TRUTH?

Could be, but touching the probe tip to the ground clip may reveal otherwise. If the scope shows a waveform with this "grounded" input, or all high impedance nodes appear to have spikes that they should not, there are at least three possible sources of error. The amplifier local ground may be quite different from the local ground seen by the scope's input amplifiers and their common mode rejection is less than perfect. First, disconnect all other signal cables from the scope to remove interaction with any other local grounds. If a battery-operated scope is available give it a try. If not, install a ground breaker on the scope power cord.

Use only shielded probes, and do not use any extenders, grabbers, or clips which do not have nearly complete shielding. Capacitive coupling into high impedance nodes works best when voltage slew rate is high and these switching amplifiers have plenty to get in trouble.

That 3" to 6" ground lead may have to go. It is forming an inductive pickup loop and the PWM is moving lots of high frequency current. If luck holds, the scope accessory kit will yield an RF adapter capable of providing a ground lead less than 1/4" long. If not, consider buying one or making your own from a length of spring wire.

5.5 INTERNAL POWER DISSIPATION

PWM amplifiers share the following thermal principles with their linear counterparts.

- Quiescent current and supply voltage determines standby power.
- Driving the load generates additional heat.
- The heatsink must dissipate both the above.
- The case temperature range must not be exceeded.
- Load related power elevates power transistor junctions above case temperature.
- Maximum junction temperatures must be observed.
- Lower temperatures (case and junction) increase reliability.

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. The second difference is subtle but affects the very reason a PWM amplifier is used: Efficiency drops rapidly as junction temperature increases. This means heatsinking the PWM is more than a reliability issue. Thermal design of the PWM amplifier has a first order affect on circuit performance.

First order calculation of power due to loading involves the output current and the total ON resistance of the amplifier. The high-speed waveforms present at the output pins do indicate second order calculations could be made but this document will concentrate only on the basic elements of power dissipation.

Total On resistance includes impedance of the H-bridge power switches (most often FETs) plus resistance of the metal interconnects. Consult the amplifier data sheet to find the contribution of each element. If interconnect resistance is not specified, consider it to be insignificant. Consider interconnect resistance to be constant over temperature. Because FET ON resistance is a function of temperature, choose a maximum junction temperature consistent with your design standards (not to exceed the data sheet absolute maximum). Find FET ON resistance(s) at your maximum junction temperature. $I^2 \bullet R$ now yields power due to loading. This is a single calculation on lower current amplifiers using all N-channel FETs, but requires another calculation if P-channel FETs are used and a third if interconnect resistance is broken out separately. Sum

the above calculations with standby power to obtain total heat loading on the heatsink. If the amplifier has a separate low voltage supply pin, don't forget to include it in the total power calculation.

With total internal power dissipation now known, it is time to determine the heatsink requirement. Again, consistent with your design standards, choose a maximum case temperature. Do not exceed the product operating temperature range listed on the last line of the specifications table. $R_{\Theta CS}$ is the thermal resistance of the package to heatsink interface.

The last item to check is the junction temperature. Multiply

$$R_{\Theta SA} \leq \frac{T_C \text{ max} - T_A \text{ max}}{\text{Total Power}} - R_{\Theta CS}$$

power in a single FET by the thermal resistance of the amplifier and add to the maximum case temperature. In models that have higher P-channel on resistance use the P-channel power level and realize the N-channel devices will run cooler. An alternative to finding the specific junction temperature is to find the appropriate fraction of total power and then use the power derating graph to make sure junctions do not exceed 150°C.

As an example consider an SA01 delivering up to 10A from a supply of 70V in a maximum ambient of 35°C. Design rules allow case and junction temperatures up to data sheet maximums.

- Standby power = 70V • 90mA = 6.3W
- N-channel power = $10A^2$ • .145Ω = 14.5W
- P-channel power = $10A^2$ • .26Ω = 26W
- Interconnect power = $10A^2$ • .05Ω = 5W
- Total power = 51.8W
- Maximum case rise = 85°C - 35°C = 50°C
- Allow .02°C/W for case to heatsink thermal resistance
- Heatsink maximum rating = 50°C/51.8W - .02°C/W = .95°C/W
- Junction temperature = 85°C + 26W • 1°C/W = 111°C

This example would actually run cooler than the above calculations would seem to indicate because junction temperatures are lower than the assumed starting point and FET ON resistance is lower. An iteration of the above based on an assumed maximum junction of 110°C would yield a heatsink rating of 1.1°C/W and result in maximum junction temperatures of 106°C. This will still have a small safety margin because the N-channel junctions run cooler than the P-channel junctions.

It is interesting to note that heatsinking will not only protect the amplifier from thermal problems, but it will also increase the efficiency of the amplifier by maintaining low on resistance. At lower temperatures, the FET ON resistance is lower and there is less dissipation. Essentially, using a larger heatsink will increase efficiency as well as protect the devices.

6.0 PWM TRANSFER FUNCTION

The transfer function of an op amp circuit is simple: $V_{OUT} = V_{IN}$ • gain. With the high PSRR of most op amps, supply voltage variations are largely ignored. Thermal effects are mostly second order, and changes in the load have little affect on the output voltage. Of course, op amps are almost always run closed loop.

PWMs can offer similar performance in a closed loop system. How to close the loop is a topic for another discussion, but it is important to understand how a PWM system behaves open loop. PWM amplifiers are affected by a variety of influences that rarely come into the picture in a linear design, or they are present in a much more subtle way than with PWMs. Hopefully, this discussion impresses the importance of closing the loop in most PWM systems.

The transfer function of a PWM amplifier is stated as follows:

$$V_O=(((V_{MID}-V_{IN})/V_{PK})\bullet V_S)-(I_{OUT}\bullet R_{ON})$$

Where:
V_O = output voltage (averaged over time, i.e., filtered)
V_{MID} = midpoint of the ramp signal
V_{PK} = + of the ramp p-p voltage
V_{IN} = input voltage
V_S = supply voltage
I_O = output current
R_{ON} = total on resistance (1 switch if 1/2 bridge, 2 switches if full bridge)

The equation above relates the output voltage (averaged over time) to the input voltage, the ramp signal, and the losses of the switches. There are several things to notice from the equation above.

** There is no supply line regulation. As the power supply varies, so does the output voltage.
** There is poor load regulation. As the load changes, so will the output current, and the output voltage will vary by $I_O \bullet R_{ON}$
** The PWM amplifier is temperature sensitive. As discussed in the previous section, the ON resistance changes with temperature. Since R_{ON} is a factor in the output voltage, this must be dealt with.

So, it is easy to see that open loop operation of a PWM amplifier is far from a precise endeavor. Closing the loop locally around the PWM amplifier or using system feedback to close the whole system are equally valid. It is sometimes easier, but far from trivial to close the loop locally around the PWM amplifier. This can eliminate all of the effects of the sources mentioned above - and others that have not been mentioned. Application note 41 discusses several methods for analog closed loop solutions and how to use the PowerDesign spreadsheet to aid in the design process.

7.0 DIGITAL INTERFACE TECHNIQUES

Apex PWM Amplifiers are analog products by design. Most of the previous discussion has assumed that the input signal is analog and a comparison is made to a triangle ramp reference to generate the PWM signal. But, with a few exceptions, Apex PWM amplifiers can also be driven with digital input signals. The following discussion explains how to interface an Apex PWM amplifier with a digital signal from a DSP or microcontroller.

Most Apex PWM amplifiers have 2 inputs, PWM+ and PWM-/RAMP. To control the amplifier digitally, apply the logic signal to PWM+. For the input signal to be processed properly the PWM-/RAMP pin must be biased to 1/2 of the logic level. For example, for 5V logic signal, bias PWM-/RAMP to 2.5V. The exact voltage of the bias is not critical as the two inputs go directly into a comparator, however it is very important to use a low impedance reference. The PWM-/RAMP pin has an internally generated triangle ramp signal that must be overcome sufficiently that the comparator can operate properly.

There may be a temptation to disable the internally generated clock signal when using digital control. After all, if the DSP or microcontroller is generating the PWM frequency, there is no obvious reason for the internal clock. There is however, a not-so-obvious reason to keep the CLK IN pin connected. For amplifiers with current limit circuitry, a low-side current limit activation is reset once every clock cycle. If there is no signal at CLK IN then any signal on the current limit pin that activates the low-side current limit will shut down the amplifier

outputs until power is cycled.

7.1 SPEED KILLS

In most cases, the natural internal clock frequency of an Apex PWM is balanced to provide reasonable efficiency while maintaining reasonable bandwidth. There are cases that require more bandwidth than the natural internal clock frequency will allow. By driving the PWM digitally, you can extend the usable bandwidth. If digital control is approached without caution, however, the consequences can range from undesirable to devastating.

Really fast switching speeds can be dangerous. As input cycle time approaches the rise time of the output drive signal there is a real potential for destruction of the amplifier as well as surrounding components and the load. The controlling circuit inside the amplifier simply becomes overwhelmed and confused when the input changes more quickly than the output can follow.

Always observe the following rules when selecting a PWM frequency:

1) Never increase the switching frequency more than twice the internal PWM frequency of the amplifier given in the datasheet.
2) Never apply any pulse shorter than 3% of the natural PWM period of the amplifier.

8.0 TYPICAL APPLICATIONS

The following are a few example design discussions. In most cases, the actual amplifier model is not important and almost any model could be used in the same circuit topology. However, there may be design considerations mentioned that are specific to a supply voltage issue or PWM frequency issue.

8.1 MOTOR SPEED CONTROL

The design steps of the PWM speed control employing a tachometer feedback shown in Figure 8 are as follows: The 7.5V reference output is used to bias the non-inverting input of the error amplifier to the middle of its 2V to 8V common mode voltage range. The gain adjust potentiometer corrects initial inaccuracies stemming from error amplifier voltage offset, tolerance of the 3.83KΩ bias resistor in the inverting input and possibly even for offsets in the input signal. The 470Ω resistor and the two associated capacitors form a low pass filter to attenuate components of the switching frequency which may be coupled to the tachometer through the motor. The gain adjust potentiometer compensates tachometer variables of accuracy and internal resistance plus tolerances of other resistors in the feed back path. The 10KΩ input resistor sets overall gain to 3.4.

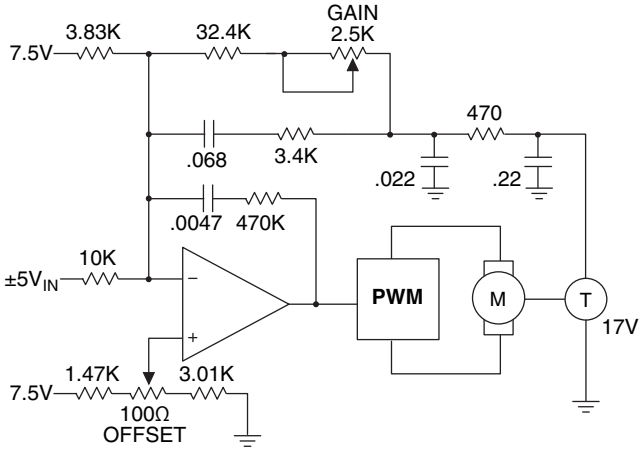

FIGURE 8. PWM SPEED CONTROL

The 3.83KΩ resistor was selected to pull the inverting input of the error amplifier up to 5V when both the input voltage and tachometer output voltage are zero. The two R-C networks were selected to provide circuit stability while maximizing system response time. Specific values will depend on both motor parameters and mechanical load characteristics.

8.1 MOTOR POSITION CONTROL

While one of the simplest forms of position sensing is shown in Figure 9, options such as optical encoders, LVDT sensors and variable capacitance transducers are also viable. Again, error amplifier inputs are biased to 5V. While 20KΩ input and feedback resistors would have set proper gain and biasing for the inverting input, they would have allowed common mode violations at the error amplifier. This could happen if the system was at one position extreme while a very quick command came in to travel to the opposite extreme. The three 30KΩ resistors prevent common mode problems by increasing impedance from summing junction to the two 10V signal levels at the output and at the input while adding an impedance to ground.

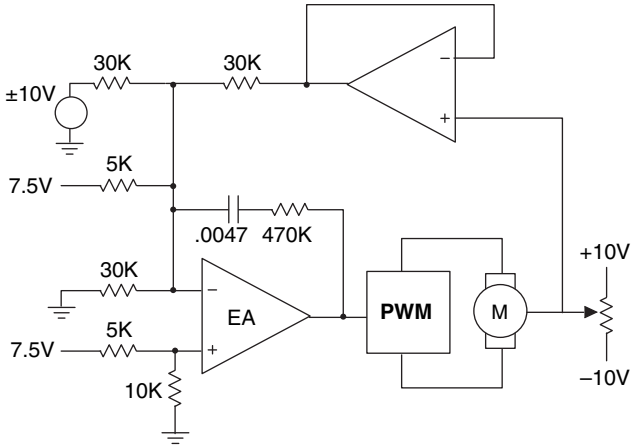

FIGURE 9. PWM POSITION CONTROL

8.2 COMMON VOLTAGE FEEDBACK CIRCUIT

Figure 10 shows a differential input, voltage controlled output circuit resembling the familiar differential op amp configuration. The SA01 includes the error amplifier within the package so the circuit can really be quite minimal. Signal gain is 2•RF/RI. Again, two pull-up resistors are used to bias error amplifier

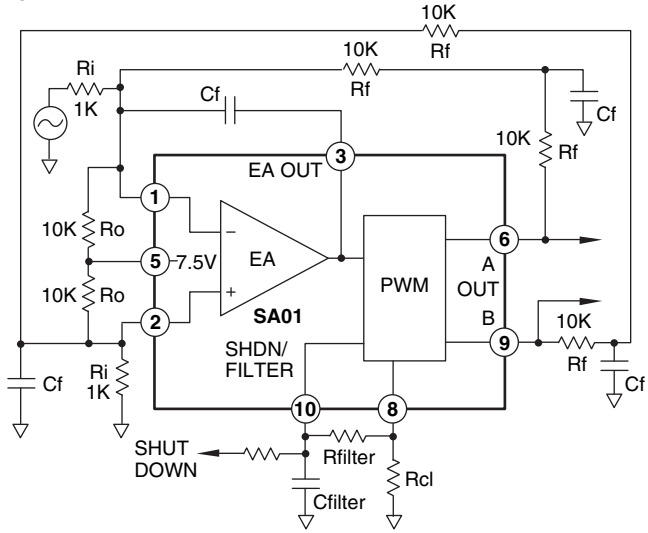

FIGURE 10. VOLTAGE FEEDBACK - SA01

inputs within the common mode range. Select this value to get 5V bias when both inputs are zero and both outputs are 1/2 the supply voltage (50/50 duty cycle.) At zero drive to the load, this differential stage is rejecting 1/2 the supply voltage present on both outputs. This means resistor ratio matching becomes critical. It should also be noted that even though the signal gain is 20, the gain of offset errors is 50 because the effective input resistance is the parallel combination of the signal input resistor and the pull-up resistor.

8.3 PROGRAMMABLE CURRENT SOURCE - DIGITAL

Figure 11 shows a digital implementation of a programmable current source. Only a half bridge is required for this single sided current source. In a half bridge, the current sense resistor (R_{SENSE}) must be in series with the load to sense source and sink currents. An instrumentation amplifier with a high common-mode voltage range senses the voltage across the sense resistor and feeds the information back to the controller.

9.0 CONCLUSION

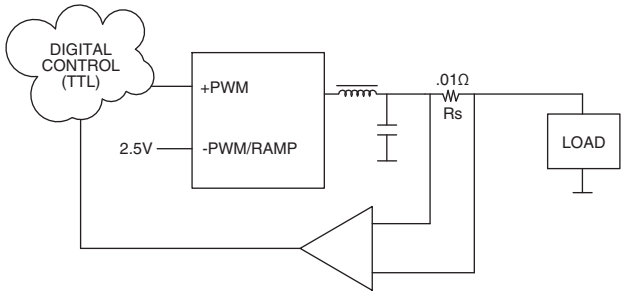

FIGURE 11. DIGITALLY CONTROLLED PROGRAMMABLE CURRENT SOURCE

The switching amplifier provides a solution to high power drives that could otherwise require an inordinate amount of heat sinking hardware. The arrival of the hybrid PWM speeds the design process and in many cases greatly enhances fault tolerance by offering protection circuits simply not possible in a discrete implementation.

