PWM BASICS

PWM

Pulse Width Modulation

- More Work
- Less Waste

As delivered power levels approach 200W, sometimes before then, heatsinking issues become a royal pain. PWM is a way to ease this pain.

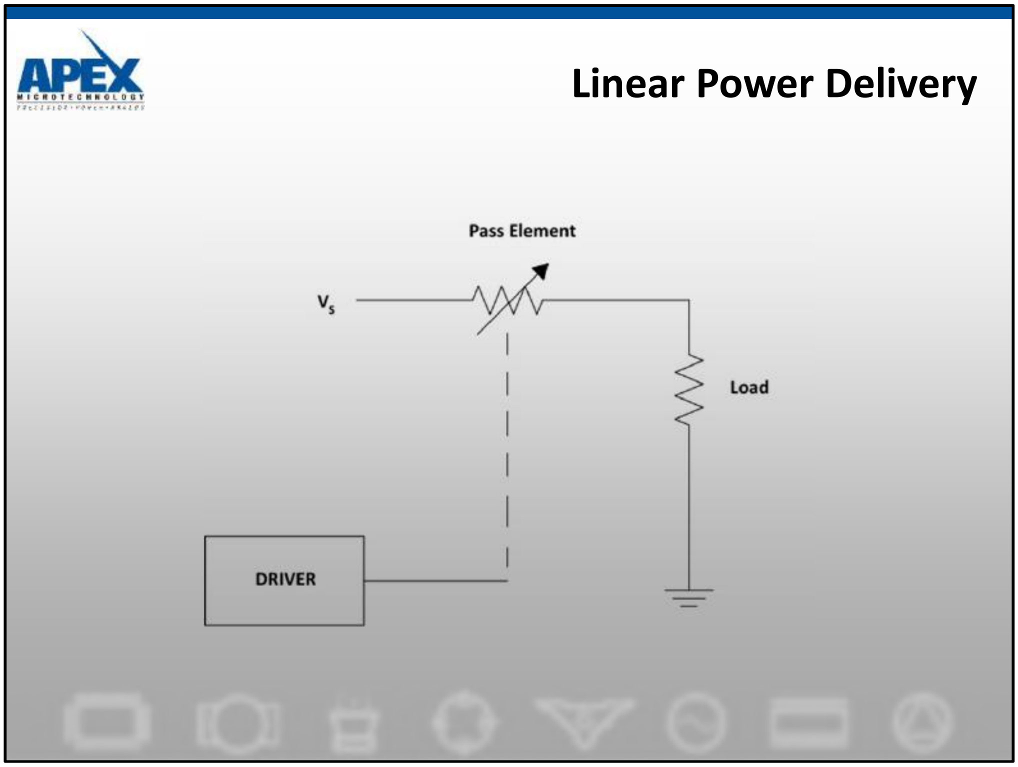

As power levels increase the task of designing variable drives increases dramatically. While the array of linear components available with sufficient voltage and current ratings for high power drives is impressive, a project can become unmanageable when calculation of internal power dissipation reveals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

This slide 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 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 mid-range power levels. When loads appear reactive, this efficiency drops even further.

Ref. AN30

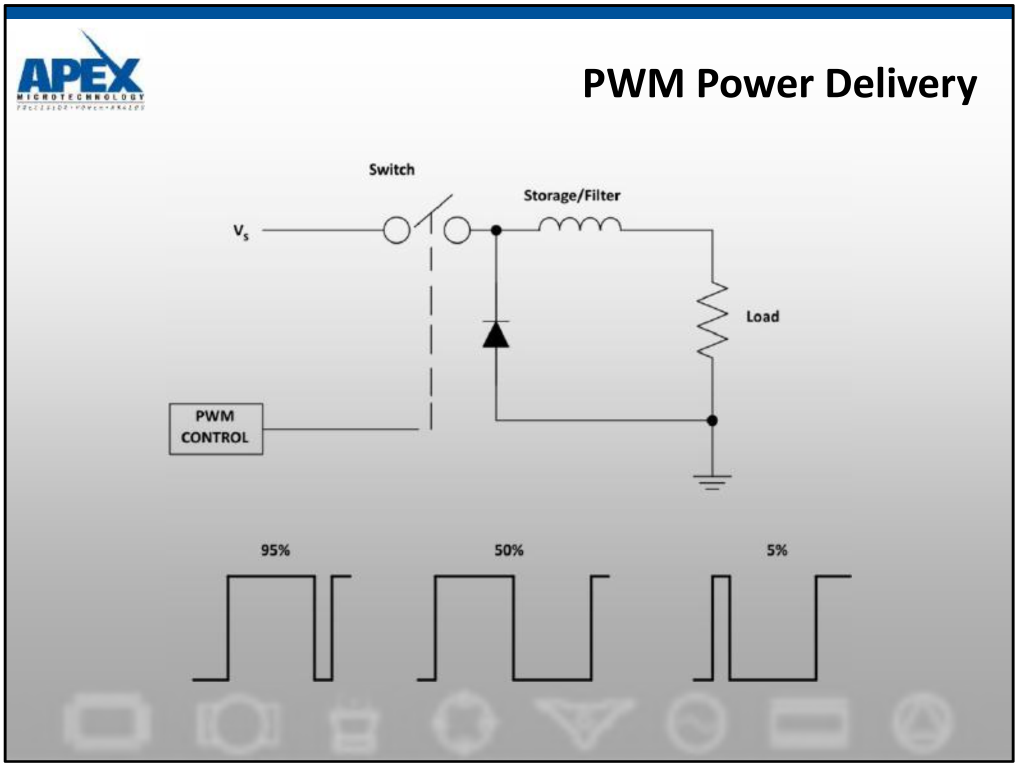

These figures illustrate the most basic PWM operation. The PWM control block converts an analog input level into a variable duty cycle switch drive signal. If high output is commanded, the switch is held on most of the period. The switch is usually both on and off once during each cycle of the switching frequency, but many designs are 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. Note that losses now include heat generated in the flyback diode. At most practical supply voltages this diode loss is still small because the diode conducts only a portion of the time and voltage drop is a small fraction of the supply voltage.

The job of the inductor is both storing energy during the off portion of the cycle and of filtering. Inductors make their living by demanding continuous current flow; they become the energy source during the off time. In this manner the load sees very little of the switching frequency, but responds to frequencies significantly below the switching frequency. When the load itself appears inductive, it is often capable of performing the filtering itself.

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. Typical efficiency of PWM circuits range from 80 to 95%.

Ref. AN30

Contrasting Discrete Linear, Hybrid Linear and Hybrid PWM 1KW Designs

	Discrete Linear	Hybrid Linear	Hybrid PWM
Wasted Heat	500W	500W	100W
$/Year[1]	$438	$438	$88
Package Count[2]	8 x TO-3	2 x PA03	1 x SA01
Heatsink	0.11°C/W	0.11°C/W	.55°C/W

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, and in the case of reactive loads, maximum voltage to current phase angle.

Consider a circuit delivering a peak power of 1KW. A 90% efficient PWM circuit generates 100W of wasted heat when running full output, and around 50W driving half power. The theoretically perfect linear circuit will generate 500W of wasted 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 temperature is +85°C. It also assumes power ratings of the TO-3 devices to be 125W each. Heatsinks for linear designs require multiple sections mounted such that heat removed from one section does not flow to other sections.

Ref. AN30

1 Continuous operation at a cost of $.10/kWH. If equipment is located in a controlled environment total cost will be considerably higher.

2 Package count must be doubled for the discrete design if bipolar output is required.

Benefits Resulting from PWM Efficiency

- Operating cost savings
- Capital cost savings
- Reduced heatsinking 5:1
- Smaller, lighter finished product

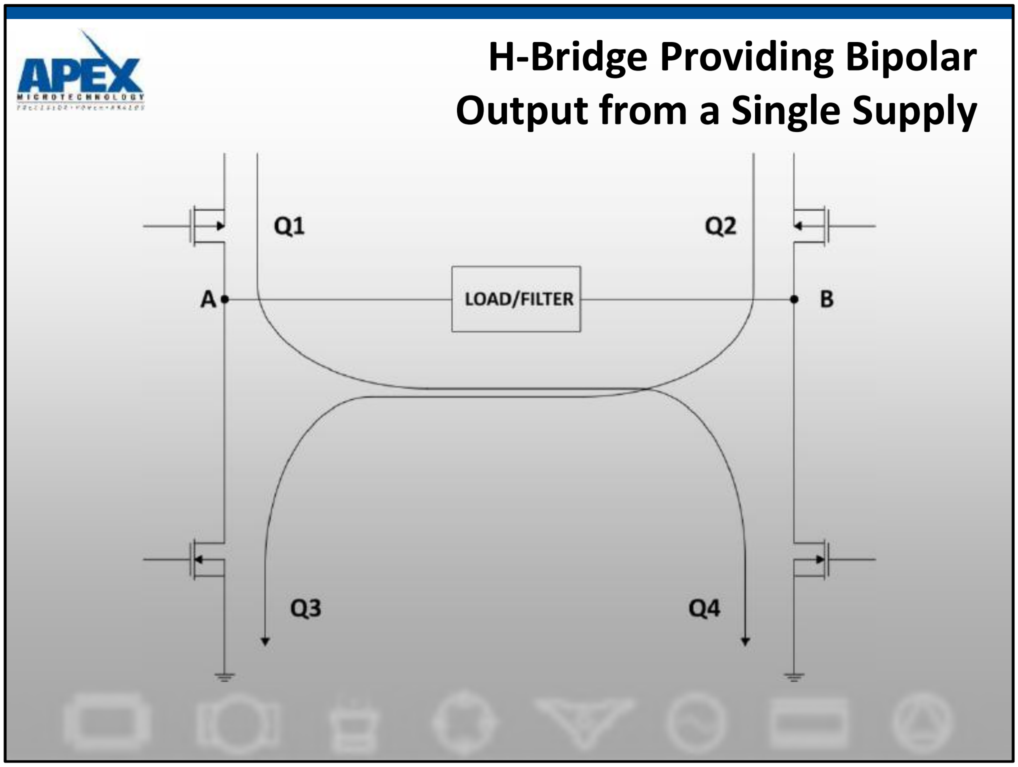

The simple form of a PWM circuit examined thus far is very similar to a number of switching power supply circuits. If the control block is optimized for producing a wide output range rather than a fixed output level, the power supply becomes an amplifier. Carrying this one step further results in the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

The H-bridge switches work in pairs to reverse polarity of the drive even though only one polarity supply is used. Q1 and Q4 conduct during one portion of each cycle and Q2 and Q3 are on during the remainder of the cycle.

Note that if Q1 and Q3 turned on simultaneously, there is nothing to limit current. Self-destruction would be only microseconds away. The fact that these transistors turn on faster than they turn off means a “dead time” needs to be programmed into the controller.

Ref. AN30,AN39

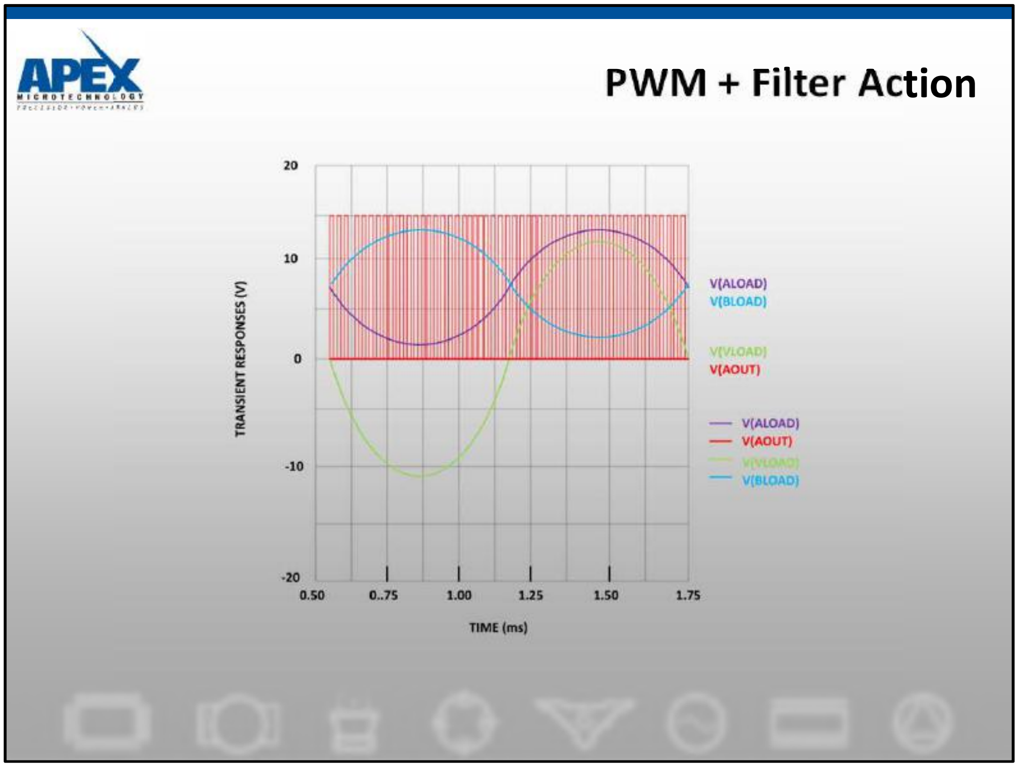

This picture shows the B output, switching at 42KHz, modulated at a 1KHz rate, along with the two filtered outputs and voltage as seen by the load.

As the A output spends most of its time in the low state, its filtered counterpart is swinging low. At the same time the A output (not shown, but out of phase with B) is mostly high and results in the filter A voltage swinging high.

With the load looking at the two filtered outputs differentially, it swings plus and minus. If you would zoom in on a ripple, it would be visible.

Ref. AN30.AN39

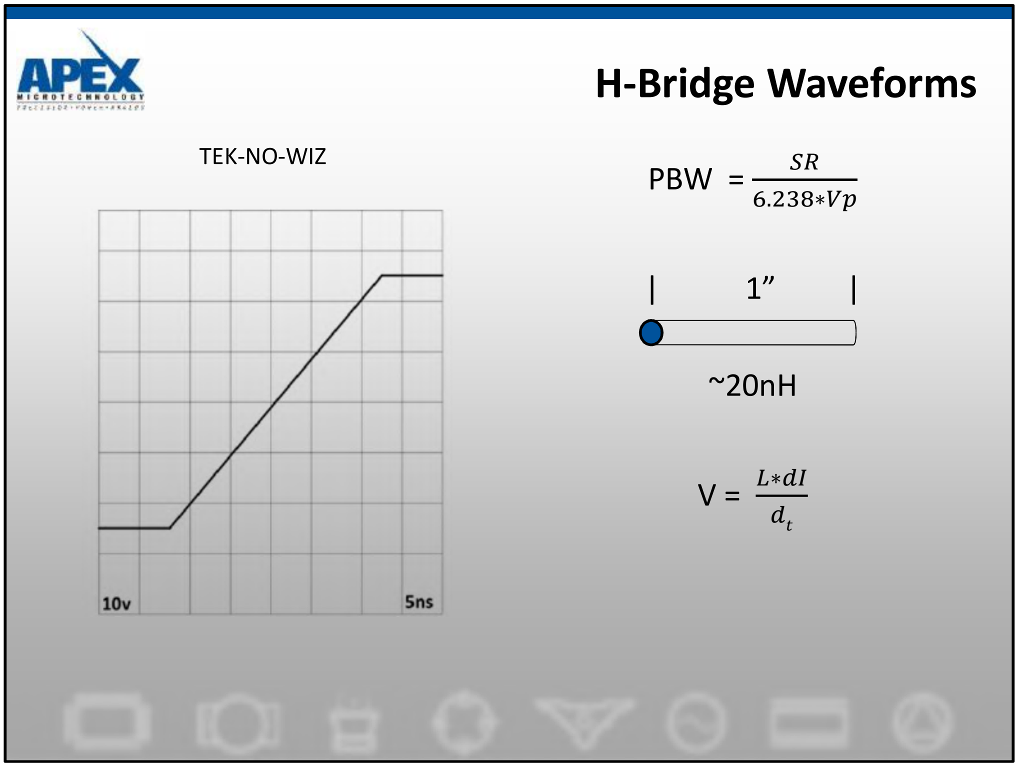

National had their FAST and DAMN FAST buffers, but they can’t hold a candle to these guys. In fact, that’s the problem with switchers- -they move voltages and currents around so fast it’s difficult to keep the noise down. Here are a few items you may not have had a chance to use lately.

From the analog world we borrow the equation relating slew rate to power bandwidth. If your PWM amplifier switches 50V in 25ns, the slew rate is 2000V/us. With peak voltage of 50V, this is over 6MHz. With 5 or 10 amps flowing, those transitions contain RF energy similar to a moderate radio transmitter. Spending a few minutes thinking like an RF designer may be worthwhile.

Currents are also changing very rapidly in these circuits. The picture above is of voltage, but keep in mind this voltage is on one end of an inductor where a power MOSFET just interrupted current flow. Look at the positive going transition: the lower MOSFET was conducting and the inductor is driving the voltage positive, above the positive supply, to maintain the previous current flow. The path will be through the body diode of the upper MOSFET, into the supply bypass capacitor. If current changes 5A in the same 25ns, two 1 inch capacitor leads will develop an 8V spike. On high speed PWMs this spike will cause the controller to freak out, rendering the circuit useless.

Ref. AN30, POWER SUPPLY BYPASS

Layout Sensitive? You Bet Your…

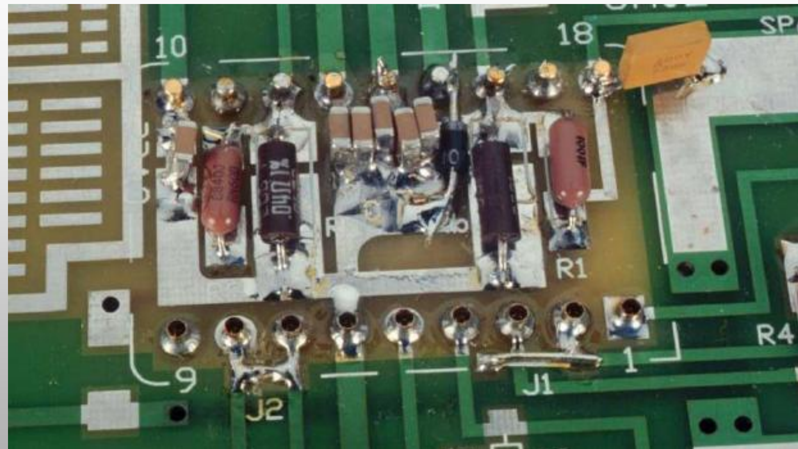

Apex has an Evaluation Kit for every PWM model

Evaluation Kits for PWM amplifier prototyping are a must. A bad layout will produce ample frustration and can cause dead amplifiers!

At a minimum, each kit provides a PC board, a way to get the amplifier plugged in, a moderate sized heatsink, and enough hardware to get it all put together. Several models also provide chip capacitors for low inductance bypass of the supplies.

In this example, the amplifier is on the opposite side of the board. Note the chip capacitors DIRECTLY between supply and ground pins of the amplifier. The two large black resistors are the current sense resistors which need to be a noninductive type.

Separate high current traces from low level traces as much as possible. Include ground plane under low level traces, but NOT under high current traces. Do NOT run high currents through the ground plane. Specify at least 2 ounce copper for the PC board. Make the ground pin of the amplifier be the center of the star ground system.

Ref. AN30

Does every low level scope observation yield the same spike-laden waveform? Here are a few causes and helpful hints.

The typical 3” to 6” ground clip on the probe has to go because it is forming an inductive pickup loop. If luck holds, the scope accessory kit will yield an RF adaptor capable of providing a ground lead about ¼” long. If not, consider buying one or making your own by winding a length of spring wire (check for piano wire at your local hardware store) on a shaft slightly smaller than the probe tip (a set of drill bits would be handy).

The ground at the amplifier contains high levels of high frequency signal relative to the ground at the scope and common mode rejection of the scope is limited. Disconnect all other signal cables from the scope. Use a battery operated scope or a ground breaker on the power cord. Use a high frequency toroid to construct a low pass common mode filter for the probe as shown.

Fast slewing signals can easily be coupled to high impedance or unshielded probes. Use only probes with nearly complete shielding. Forget the grabber clips, extenders or any single conductor connections to the scope.

Ref. AN30

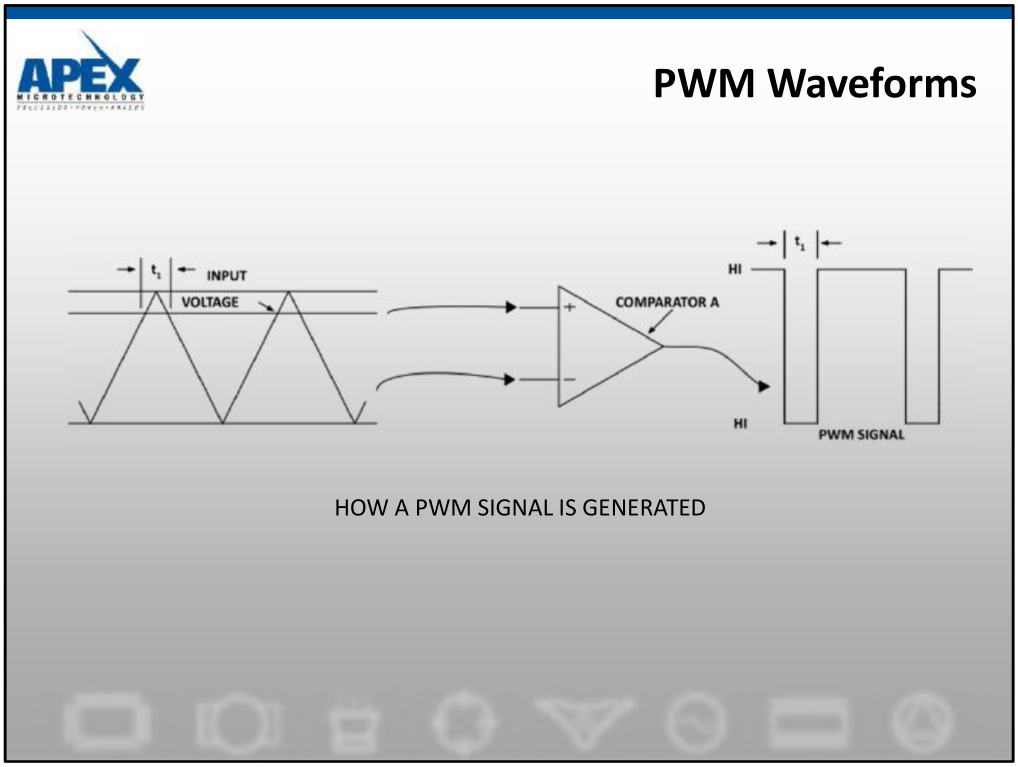

This picture illustrates basic control of one diagonal pair of switches. As the input voltage rises toward the upper peak of the triangle wave, the high portion of the waveform increases. Pure theory dictates that any duty cycle can be programmed, but propagation delay, rise and fall times and hysterias pose practical limits to how narrow a pulse can be. Typical minimums range from about 100ns to 1us. This typically translates into a 1 to 5% band on each end of the duty cycle range which cannot be sustained. Let us assume the signal is 90% and increasing. When the band is encountered, modulation will jump to 100%. On the way back down, there will be a jump from 100% to just under the band (down to 99 to 95% typically).

Notice that there is no discontinuity in the middle of the range (50/50 duty cycle and zero volts out for a full bridge. This is an advantage over another popular modulation scheme, sign magnitude modulation where different switches are turned on and off for positive and negative outputs. Borrowing from the linear world, this is similar to the difference between class A-B output stages and Class B output stages where Class B designs have a dead band (crossover distortion) where the circuitry changes output drive transistors. On the down side for locked anti-phase modulation, is the fact that near 50/50 duty cycles, the filter must do more work (efficiency at midrange is lower than sign magnitude modulation). This is generally only a concern when the system spends a large portion of the time driving very low currents.

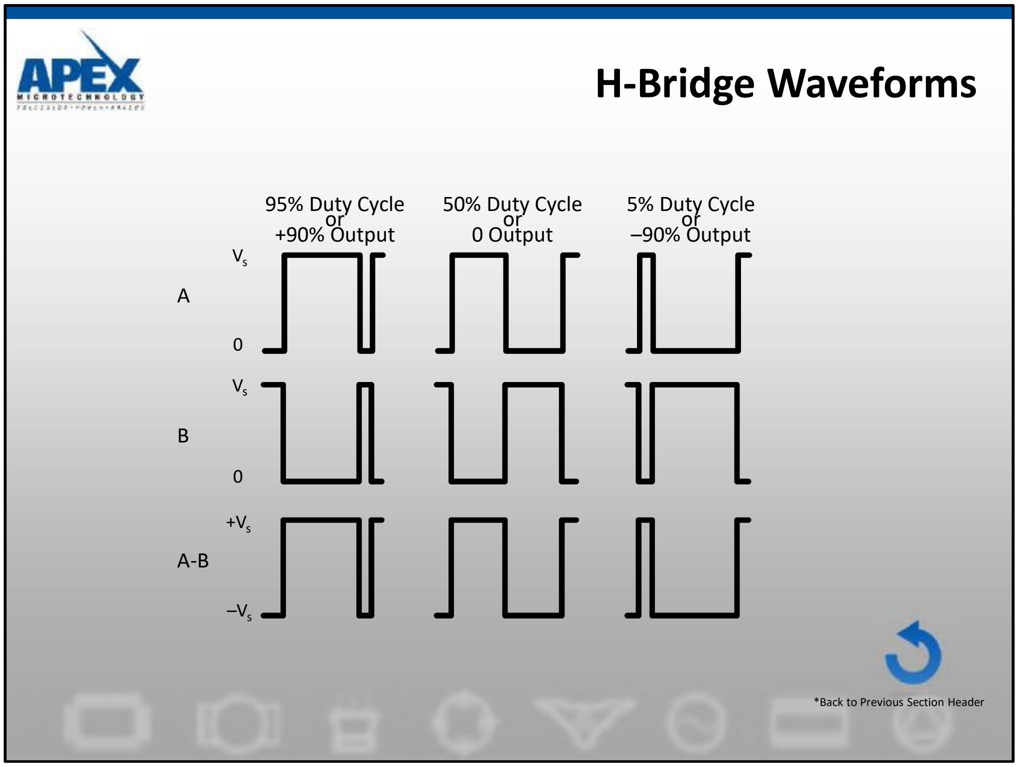

To help understand the conversion of the time modulated data to analog levels, visualize each waveform segment of Figure 2 run through a low pass filter whose cutoff frequency is at least 10 times lower than the switching frequency. The A and B voltages of the 50% duty cycle waveforms will both be equal to 50% of the supply voltage. With both terminals of the load connected to the same voltage, the load sees 0V across itself. The A-B waveform represents this differential connection of the load, and the filtered voltage of this waveform equals zero.

To examine the 95% duty cycle waveforms, lets assume a supply voltage of 100V. The filtered A value will be 95V, B will be 5V, and the load will see 90V; the same as the filtered value of the A-B waveform. When the duty cycle shifts to 5%, the filtered A value will be 5V, B will be 95V, and the load will see –90V, again matching the filtered value of the A-B waveform.

Changing duty cycle through 50% is a continuous function, meaning there is no inherent discontinuity as exists in sign magnitude modulation. This is analogous to the much improved distortion levels of class AB linear stages versus class B linear stages where zero current crossing brings a discontinuity or dead spot usually referred to as crossover distortion.

Ref. AN30,AN39