Introduction to Pulse Width Modulation

An Online Continuing Education Course for Engineers

Course Number: IC-2004

Credit: 2 Hours / 2 PDH / 2 CPD

Introduction to Pulse Width Modulation

Stuart Ball, P.E.

Introduction: This course covers the basic concepts and techniques of Pulse Width Modulation, or PWM. Although PWM-like techniques can be applied to mechanical systems, this course will focus only on electrical applications. The course assumes that the reader is familiar with basic electrical concepts such as resistance, basic transistor concepts, and basic electrical math, such as Ohm's law for resistance, current, voltage, and power.

PWM control is used to control the temperature of heaters, the speed of motors, the brightness of lights, and the volume of sound. With the availability of inexpensive microprocessors, specialized PWM control integrated circuits, and custom ICs, PWM has become a common way to control electrical devices that require significant power.

Why PWM?: Pulse Width Modulation, or PWM, is a method to control analog devices using digital techniques. To explain how PWM works, we will start with an analog circuit. In the simple example below, a resistive load is driven by a driving element such as a transistor:

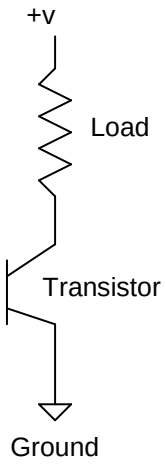

In this example, the load might be a resistive heater controlling the temperature of hot-melt glue, or a crystal heater of some similar application. Say that the load has a value of 10 ohms and the supply voltage (+V) is 24V, and we want the load to be operating at half of its maximum current. In that case, we would want the voltage across the load to be half the supply voltage, or 12V. The power dissipated in the load would then be:

$$\frac{\left(\frac{+V}{2}\right)^2}{RL} = \frac{12^2}{10\ \text{ohms}} = 14.4\ \text{watts}$$

The current through the load is:

$$\frac{12\text{v}}{10\text{ ohms}} = 1.2\text{ amps}$$

The load and the driving transistor are in series, so the same current flows through both of them. Since half the voltage appears across the load, then half the voltage also appears across the driving transistor. Also, since the transistor has the same voltage (half the supply voltage, or 12V) across it, and the same current (1.2A) flowing through it, the transistor dissipates the same power as the load—in other words, half the power supplied to the circuit is wasted as heat in the transistor. If the resistive load were a 1-ohm heater instead of 10 ohms, the amount of power wasted would be ten times as high.

The total power required from the power supply for this approach is the power dissipated in the transistor (14.4 watts) plus the power dissipated in the load (14.4 watts), or a total of 28.8 watts.

The use of analog techniques in driving loads that need significant power requires that the driving transistors be larger. Additionally, a large heatsink or other added cooling to avoid overheating the driving circuitry is often necessary. This is because the driving transistor must dissipate significant power.

PWM: PWM significantly reduces these problems by driving the load with a switched digital signal. A simple PWM circuit might use the same electrical configuration of a resistive heater driven by a transistor. In the PWM implementation, however, the transistor is switched on and off at a high rate. In that case, the voltage across the resistor looks like this:

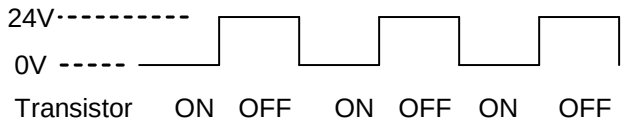

When the transistor is ON (conducting), it pulls the low end of the load to ground, and the full supply voltage of 24V appears across the load. When the transistor is OFF (not conducting), the low end of the load floats, and no current flows through the load.

When the transistor is on, and the full supply voltage is applied across the load, the current in the load is 24V/10 ohms or 2.4 amps. When the transistor is off, no current flows through the load. In this case, the transistor ON time is the same as the transistor OFF time, so the average current is the same as the original analog case is 1.2 amps.

Although the average current in the load, and therefore the average power, is the same as the original analog method, the power dissipated in the transistor is much lower. Like the load, the power in the transistor is the product of the voltage across the transistor times

the current flowing through it. When the transistor is OFF (not conducting), the voltage applied across it is the full supply voltage, but the current is zero, so the power is zero:

$$\text{Power} = \text{Transistor voltage x Transistor current} = 24 \text{ volts x } 0 \text{ amps} = 0 \text{ watts}$$

When the transistor is ON, the full load current of 2.4 amps flows through it, but the voltage is zero, so the power is zero. (This is an ideal case, an actual drive transistor won't quite have zero voltage when in the ON state; we will examine that later).

$$\text{Power} = \text{Transistor voltage x Transistor current} = 0 \text{ volts x } 2.4 \text{ amps} = 0 \text{ watts}$$

Since the power in the transistor is zero when it is OFF and zero when it is ON, the transistor itself wastes no power, yet the same power is delivered to the load as in the original analog example. This is what makes PWM more efficient than analog techniques—less power is wasted. This typically means the driving element can be smaller and requires less cooling.

The average power required in the PWM circuit is the power delivered to the load—14.4 watts. So the average power is less than that required by the analog circuit. However, the peak power is the power used by the load when the transistor is ON. In this case, that is 24 volts x 2.4 amps, or 57.6 watts. The *average* power required from the supply is lower than for the analog circuit, but the *peak* power is the same as when the analog circuit turns the heater fully on. Peak and average power are the same at 100% duty cycle.

In this example, PWM implementation, or the average power in the load, is controlled by adjusting the *duty cycle* (the ratio of ON time to OFF time) of the transistor. When the ON time is higher, the average current is higher, so the average power is higher. The duty cycle may be expressed as a percent (50% in our example) or as a decimal (0.5 in the example). The figure below shows what the voltage across the load would look like at 10%, 50%, and 90% duty cycles. In each case, two full PWM cycles are shown.

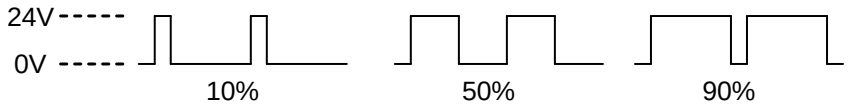

PWM techniques do come at a price—the driving circuitry is more complex than in the analog circuit. On a schematic, the driving transistor may look the same, but the transistor used for PWM typically must have faster-switching characteristics and lower voltage drop when in the fully ON state than the transistor used for analog control. In addition, the circuitry to drive the transistor is usually more complex. That complexity may be in hardware, such as digital logic in an integrated circuit or in software such as a program in a microcontroller.

Actual (non-ideal) Transistors: The ideal transistor, as previously described, has zero volts across it when in the ON state, and it turns on and off instantly. As a consequence,

the power dissipation of the ideal transistor is zero. Real transistors do not have quite zero volts across them when ON, and they do not switch instantly. Therefore, a real transistor does dissipate some power when ON, and it dissipates power when switching between the ON and OFF states. The figure below illustrates this.

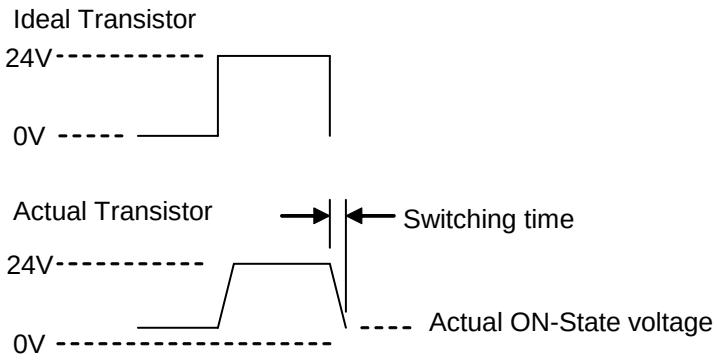

The ideal transistor switches between 0V and 24V and has zero switching time. The actual transistor does not go quite to zero volts when ON, and it has a finite switching time. Most transistors used as PWM switches to have very close to zero current when in the OFF state, so they dissipate negligible power in that state; in that respect, they are very similar to the ideal transistor.

In our example, say that the transistor has 1V across it when in the ON state. In that case, the ON-state power of the transistor is:

$$\text{Power} = \text{Transistor voltage x Transistor current} = 1 \text{ volt x } 2.4 \text{ amps} = 2.4 \text{ watts}$$

At 50% duty cycle, the average dissipation of the transistor would be half of this, or about 1.2 watts. This is significantly less than the power dissipated by the transistor in the original analog circuit. The ON voltage of the transistor will vary with the type of transistor and the amount of current in the ON state. Typical values for common transistors are less than one volt, although high-current loads may result in higher ON voltages.

The switching time or transition time also dissipates power. When the transistor is transitioning between the ON and OFF states, it has a voltage across it, and a significant current flowing through it. As a consequence, during the transition, the transistor can dissipate significant power.

The switching time of the transistor also limits the maximum switching frequency of the PWM control. When the period of the PWM waveform is much greater than the switching time, the average dissipation of the transistor is low because most of the time the transistor is in the ON or OFF state; the transition time is a small percentage of the total period. As the frequency is increased, the edges become closer together and the transition time is a greater percentage of the total period. If the period were to become shorter than the transition time, the transistor would never completely turn on or off and

the power dissipation would potentially be very large. Of course, the PWM control would stop working effectively well before that point was reached. Therefore, one of the factors that govern the upper limit of the PWM frequency is the switching speed—the transition time—of the driving transistor.

The power dissipated in the PWM drive transistor consists of two components; the power dissipated during the ON state and the power dissipated during transitions between ON and OFF states. The ON state power will vary with the duty cycle because the transistor only dissipates power when ON. The power dissipation during transitions will be essentially constant as long as the PWM duty cycle doesn't change.

At the other end of the frequency spectrum, the load must not respond to the individual ON and OFF states of the transistor, it must respond to the average. If the ON and OFF states in the previous heater example were each several seconds long, it is possible that the temperature of the heater would cycle with the PWM signal, as shown below:

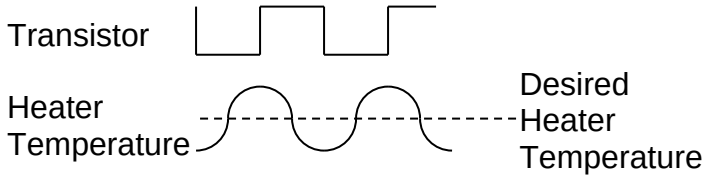

The amount of fluctuation would depend on the mass of the heater and whatever load it was heating and the actual frequency of the drive signal. But the important point is that in most PWM systems, it is undesirable for the PWM signal to be so slow that the load responds to the individual PWM transitions—the load should instead respond to the *average* PWM signal.

If the hypothetical heater were heating glue at the point of application, this fluctuation might cause the glue temperature to vary too much, causing improper application of the glue. However, if the heater is being used to heat a pot of glue, then the fluctuation would have almost no effect if the glue pot was large enough—it would then have a sufficiently long thermal time constant that it would not respond to the fluctuations in the heater temperature, only to the average. So there are some situations where it would be acceptable to allow the kind of fluctuation described above and to operate the PWM signal at a lower frequency.

Efficiency: The efficiency of the PWM circuit is the power delivered to the load (useful power) divided by the total power delivered to the circuit. For the analog circuit originally described, with the load operating at half the supply voltage, the transistor and the load dissipate the same power, so only half the power supplied to the circuit goes into the load. The efficiency of the analog circuit with half the power across the load is 50% (14.4 watts dissipated in the load divided by 28.8 watts total power). An ideal PWM circuit always has 100% efficiency because no power is dissipated in the ON state, and no power is dissipated during the switching transitions. Therefore, all the power goes into the load. A PWM circuit with actual transistors will have efficiency lower than 100% due

to the power dissipated by the drive transistor when it is ON and due to the power dissipated by the drive transistor when it is transitioning between ON and OFF.

The implication of this is that the efficiency of the PWM circuit is affected by the PWM frequency. As the PWM frequency goes up, the transition time becomes a higher percentage of the total PWM period, lowering efficiency.

The analog and PWM examples have so far been based on operating the load at half the supply voltage. If the load needs to operate at maximum power (the full supply voltage applied across the load), then the PWM and analog circuits do the same thing—both will turn the transistor fully on and the power is all dissipated in the load (except for the ON state power in the transistor). Assuming the transistors have identical characteristics, the analog and PWM methods both have the same efficiency at 100% power. At other power levels, the PWM method is more efficient because the drive transistor dissipates less power.

As an example, assume that the PWM transistor dissipates 1 watt when ON, and 1 watt of additional power during the transitions between the ON and OFF states. The average power dissipated in the PWM transistor would then be the 1W transition power (which doesn't change with duty cycle) and the average ON-state power (which does vary with duty cycle). The dissipation of the PWM transistor is calculated as:

1W + Duty cycle x 1W

So at varying duty cycles, the PWM transist s:

Dut
25%
50%
75%

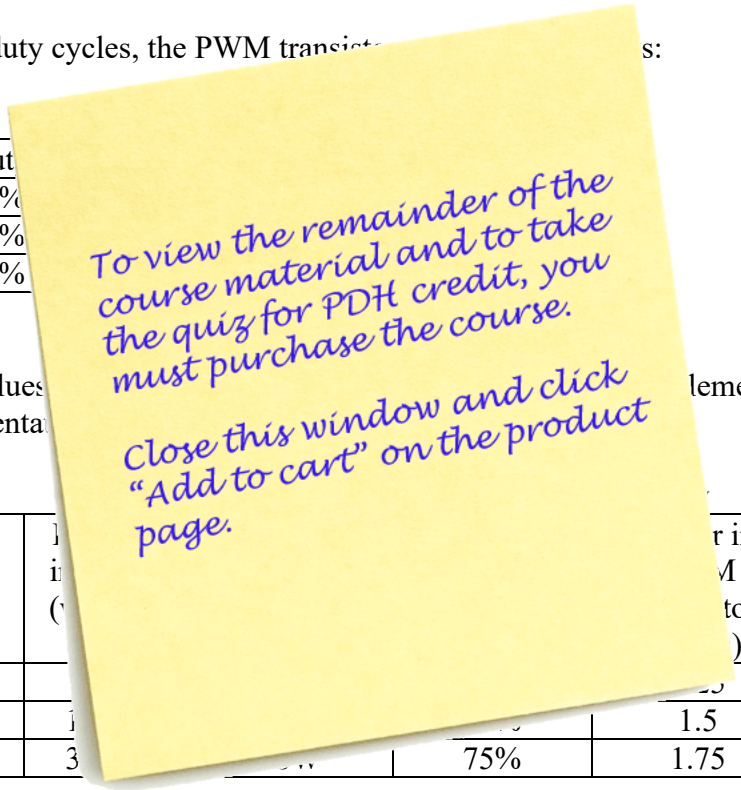

Using these values lementation to the PWM implementa

Current (percent of maximum current)				r in M tor)	PWM efficiency
25%					74%
50%				1.5	90%
75%	3		75%	1.75	95%