

INTRODUCTION

Those of us who have ventured into high power linear circuits with their massive and sometimes liquid cooled heatsinks have a tendency to go ga-ga over the efficiency potentials of the Pulse Width Modulation (PWM) amplifiers. This is OK. But these little switching miracles do bring a new set of challenges to the table.

The PWM amplifier with no filtering is NOT capable of amplitude modulation. It can only change times and maybe polarity. We get lucky once in a while and get a load which will do the filtering for us. Much more often we must design the filter, a job many of us do not place at the top of our list of most cherished activities. In addition to this, the methods required to calculate internal power dissipation and the heatsink size are quite different than those used in the linear world.

The Power Design CAD tool automates Butterworth filter equations found in Apex Applications Note 32 and expands on this base by graphing response with real world components. It then goes on to automate internal power dissipation equations, plus draw a wide variety of graphs on amplifier performance over frequency. As the overall process is usually iterative, the benefit of computer analysis is indispensable.

SOME PWM BASICS

PWM circuits achieve high efficiency compared to their linear counterparts in much the same manner as switching power supplies do versus linear supplies. If the control block is optimized for producing a wide output range rather than a fixed output level, the power supply becomes an amplifier. Figure 1 illustrates a typical PWM amplifier output stage 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.




FIGURE 1. H-BRIDGE OUTPUTS WITH DIFFERENTIAL AND SINGLE-ENDED FILTERING

The H-bridge switches work in pairs to reverse polarity of the drive, even though only one polarity supply is used. Figure 2 shows waveforms of locked anti-phase modulation where S1 and S4 are on during one portion of each cycle, and S2 and S3 are on during the remainder of the cycle.

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

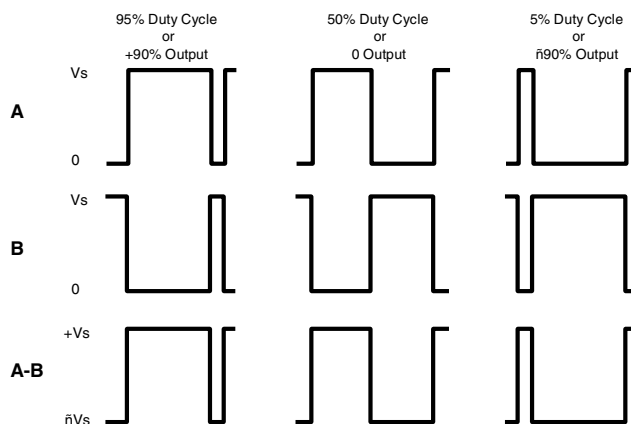


FIGURE 2. H-BRIDGE WAVEFORMS

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, let's 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.

Note that if S1 and S3 were to turn 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 if doing your own design. When you buy the amplifier from Apex, this is all inside the package.

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 C linear stages where zero current crossing brings a discontinuity or dead spot usually referred to as crossover distortion.

National created 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. From the linear or 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 a peak voltage of 50V, this equates to 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.

Refer to Figure 3 (next page) for a pictorial of the filter's job. The relatively flat portion of the curve is the pass band of the filter. The signal frequency of the power drive to the load must fit under this area. Desired attenuation in this area is 0db and the corner frequency is F_c , the cutoff frequency. We then go down the filter slope to the switching frequency, F_{sw} . Allowing one decade between these two frequencies is a good starting point. In this case, the graph tells us the worst case peak ripple voltage at the switching frequency will be a little under 1% of the supply voltage.

Figure 4 (next page) illustrates how a zero output voltage

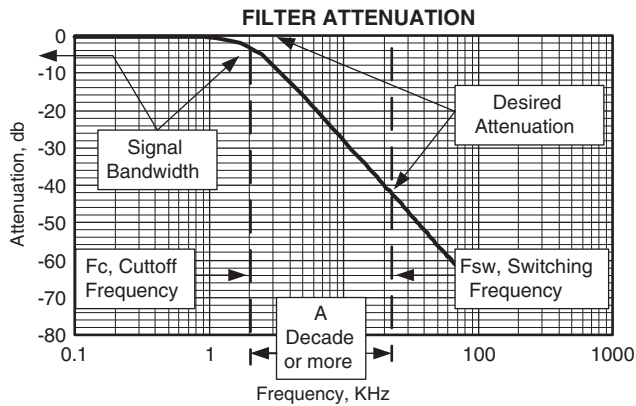


FIGURE 3. PWM FREQUENCY RELATIONSHIPS




FIGURE 4. DUTY CYCLE AND RIPPLE CURRENT VARIATIONS WITHOUT OUTPUT VOLTAGE

corresponds to a 50% duty cycle and produces maximum ripple current. As expected, there is a linear relationship between increased output voltage and increased duty cycle. Not quite as obvious is the curve that indicates the ripple current is reduced all the way to zero if we push modulation all the way to steady state.

The need to squeeze the last ounce of bandwidth from

our designs, along with the physics limitations on switching frequencies, makes it desirable to minimize the distance between signal and switching frequency. Pure theory says adding more poles can increase filter slope. This is true to a point. We would probably question an eight-pole filter in the small signal world. Do you really need that? Can you find high enough quality components to make it work? Can you afford it in terms of size and cost?

In the PWM world these questions are not only valid but are many orders of magnitude more important because power levels have gone from mW to KW! Rule of thumb: Allow at least a decade between switching and signal frequencies.

THE POWER DESIGN APPROACH TO SUCCESSFUL PWM AMPLIFICATION

The usual process is:

1. Select an amplifier model (possibly with the Part Selector sheet).
2. Load circuit data into the PWM Filters sheet.
3. Auto load components into the Filter/load model and sweep the frequency.
4. Tune components and parasitics plus check load variations and fault conditions.
5. Set sweep frequency band from Fmax to at least 10 times Fsw to check high frequency attenuation.
6. Use the graphs to select the heatsink.

The Power Design, PWM Filters sheet data entry screen for step 2 is shown in Figure 5. The pull down Model cell reads the built-in database containing specifications on supply voltage, maximum switching frequency, current levels, and internal resistance. Alternatively, comments in the database area provide instructions to enter data for your own design.

Switching frequency, Fsw, is required because some models are programmable, most can be run lower than the maximum and many can be driven with digital signals. Immediately to the right of this data entry cell is the maximum for the model selected. Enter minimum frequency to be amplified as Fmin. Use .001 for DC. Consider using .001 even if the application is a substantially higher fixed frequency, as this may simulate a "lost" input signal condition and some circuits will present their lowest impedance at DC. Enter the maximum frequency

to be amplified in Fmax. Fmin and Fmax set the frequency end points of the sweep that will be run later. Fcutoff is the cutoff frequency of the filter and will be the -3db response point. The next three cells describe three series connected elements forming the load. Vripple is the maximum peak voltage on the load at the switching frequency, your way to specify the attenuation of the filter at Fsw. The bottom two cells specify the magnitude and unit of measure for the output signal.

The Order Calculation section first converts your maximum ripple and power supply ratings into db attenuation. Then by examining the switching and cutoff frequencies, it calculates the order, or number of poles needed. The

Filter Design for PWM Amplifiers				READ ME	Using the Complex Load:		
CAUTION!				Refer to Applications Note 32			
Input Data				Order Calculation			
Model	SA03			Atten. @ Fsw	41.023 db		
Vs	90 Volts	Fsw	22.5 KHz	22.5 N(exact)	1.9513		
Fmin	0.001 KHz	Fmax	2 KHz	N(recommended)	2		
Fcutoff	2 KHz	Matching network				<input type="button" value="60 Load All Data For N=1"/> <input type="button" value="61 Load All Data For N=2"/> <input type="button" value="62 Load All Data For N=3"/> <input type="button" value="63 Load All Data For N=4"/> <input type="button" value="64 Load All Data For N=5"/> <input type="button" value="65 Load All Data For N=6"/>	
Rload	10 Ohms	Cm =	0 uF	Lm =	0 mH	Read Me Yes Auto Sweep on Load?	
Cload	0 uF	Rm =	10 Ohms	Recommended Cleg = 0.3448 uF			
Lload	0 mH						
Vripple	1 Vpk						
Signal	85 Units						
Sig as ?	V peak Note/W						
Notes:							
<input type="button" value="46 Print Filter"/>			<input type="button" value="55 Show Attenuation in db & %"/>				
<input type="button" value="56 Show Attenuation Graph"/>			<input type="button" value="66 Show Filter Components"/>				

FIGURE 5. PWM FILTER DESIGN DATA ENTRY SCREEN

Component Calculations

N	Dual Cap Filter		Single-ended Filter		N	Dual Cap Filter		Single-ended Filter		
	L	C	L	C		L	C	L	C	
N = 1	L =	0.3979 mH			N = 2	L =	0.5627 mH		1.1254 mH	
	P-P Ripple =	2.5133 Amps out of the amplifier				C =	11.254 uF		5.6269 uF	
	Avg. Iout for thermal calculations =	0.6283				P-P Ripple =	1.7772 Amps out of the amplifier			
N = 3	L1 =	0.5968 mH		1.1937 mH	N = 4	L1 =	0.609 mH		1.2181 mH	
	C =	21.22 uF		10.61 uF		C1 =	25.102 uF		12.551 uF	
	L2 =	0.1989 mH		0.3979 mH		L2 =	0.4307 mH		0.8613 mH	
N = 5	P-P Ripple =	1.6755 Amps out of the amplifier				C2 =	6.0909 uF		3.0454 uF	
	Avg. Iout for thermal calculations =	0.4189				P-P Ripple =	1.6419 Amps out of the amplifier			
	L1 =	0.6148 mH		1.2296 mH	N = 6	L1 =	0.6179 mH		1.2358 mH	
C1 =	26.967 uF		13.484 uF	C1 =		28 uF		14 uF		
L2 =	0.5499 mH		1.0998 mH	L2 =		0.6179 mH		1.2358 mH		
N = 6	L3 =	0.1229 mH		0.2459 mH	C2 =	19.124 uF		9.562 uF		
	P-P Ripple =	1.6266 Amps out of the amplifier				L3 =	0.3016 mH		0.6031 mH	
	Avg. Iout for thermal calculations =	0.4067				C3 =	4.1189 uF		2.0595 uF	
					P-P Ripple =	1.6184 Amps out of the amplifier				
					Avg. Iout for thermal calculations =	0.4046				

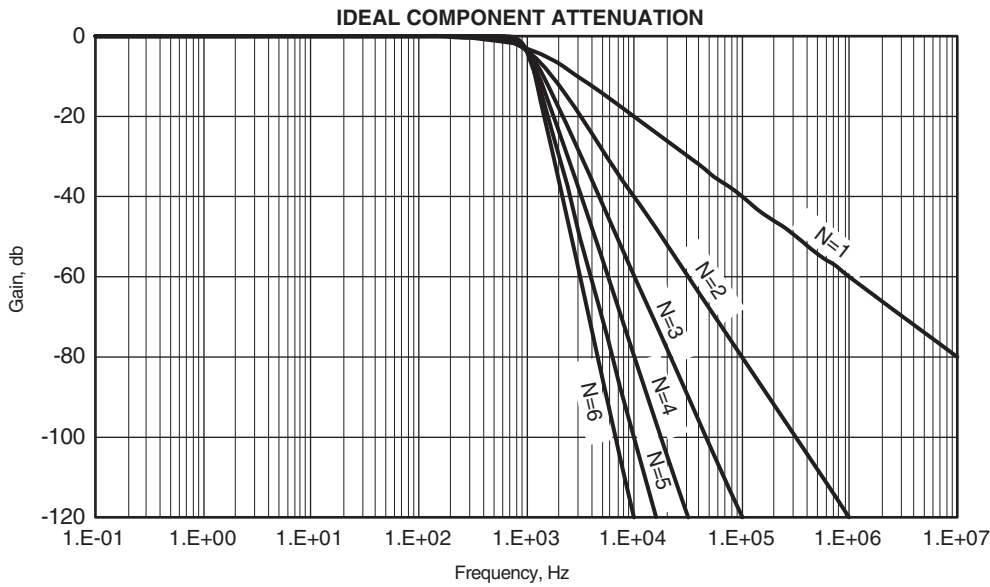
equally well here. To achieve the performance shown in this graph, output impedance of the amplifier must be zero; the filter must contain perfect components; be terminated with the load described; and the specified matching network must be in place. If these conditions are not true, **ALL BETS ARE OFF.**

A quick look at PWM amplifier data sheets will tell us actual output impedance can cause small errors if left unchecked. Use your knowledge of op amp theory and Figure 8 to see how closed loop output impedance of the PWM amplifier is extremely low just as with a closed loop op amp. On the left circuit we know output impedance is reduced by the loop gain. As long as the op amp in the right circuit has no direct DC feedback, and the PWM block with its output impedance (typically ranging from 0.1Ω to 1Ω) is inside the feedback loop, closed loop output impedance will be reduced in the same fashion. With PWM amplifiers being relatively slow compared to op amps, it is easy to obtain high loop gains over the power signal bandwidth to achieve negligible errors in driving the filter. In actual PWM systems, the feedback loop is often much more complex than shown here.

Trying to approach the second "ideal" condition means most of the work still lies ahead in finding components which work as advertised in the MHz range and whose losses won't radically change the pretty graphs. An electrolytic capacitor may perform very well at 60Hz, but rather poorly at 6MHz. If temperatures allow, switching to tantalum should result in a noticeable high frequency improvement. Moving to switching rated plastic capacitors or ceramic capacitors is usually an even better choice.

Not many of us would attempt using laminated steel core inductors here, but please note that not all "high frequency" coils are created equal. Air core inductors get away from the magnetic saturation problem and they have fewer tendencies to become dummy loads at high frequency. The down side will be more turns of wire and more copper losses. When adding a magnetic core, make sure the material can handle the high frequency components of the square wave (manufacturers often rate frequency capability for only sine waves) at the switching frequency and can accommodate

FIGURE 6. COMPONENT VALUES FOR DIFFERENTIAL AND SINGLE ENDED FILTERS



Attenuation in db	Percent attenuation												
500	-9.7E-01	-2.6E-01	-6.7E-02	-1.7E-02	-4.2E-03	-1.1E-03	5.0E+02	10.55728	2.98575	0.772212	0.194742	0.048792	0.012205
750	-1.9E+00	-1.2E+00	-7.1E-01	-4.1E-01	-2.4E-01	-1.4E-01	7.5E+02	20	12.84245	7.863584	4.658634	2.702074	1.547157
22500	-2.7E+01	-5.4E+01	-8.1E+01	-1.1E+02	-1.4E+02	-1.6E+02	2.3E+04	95.55994	99.80247	99.99122	99.99961	99.99998	100
225000	-4.7E+01	-9.4E+01	-1.4E+02	-1.9E+02	-2.4E+02	-2.8E+02	2.3E+05	99.55556	99.99802	99.99999	100	100	100
Hertz	N=1	N=2	N=3	N=4	N=5	N=6	Hertz	N=1	N=2	N=3	N=4	N=5	N=6

FIGURE 7. THE IDEAL ATTENUATION GRAPH AND PRECISE CHECKING OF POINTS OF INTEREST




FIGURE 8. A PURE INTEGRATOR IS THE KEY TO ACCURACY

integer recommendation is rounded up. The matching network that is calculated will cause reactive loads to appear resistive to the output of the filter. Figure 6 shows the actual filter components for filter orders up to N=6, plus expected ripple current at the switching frequency. Figure 7 shows the ideal response graphed and an area where attenuation at specific frequencies can be checked in detail.

"Ideal" is a great word. Just as we use the concept to describe theoretical performance of the linear op amp, it will work

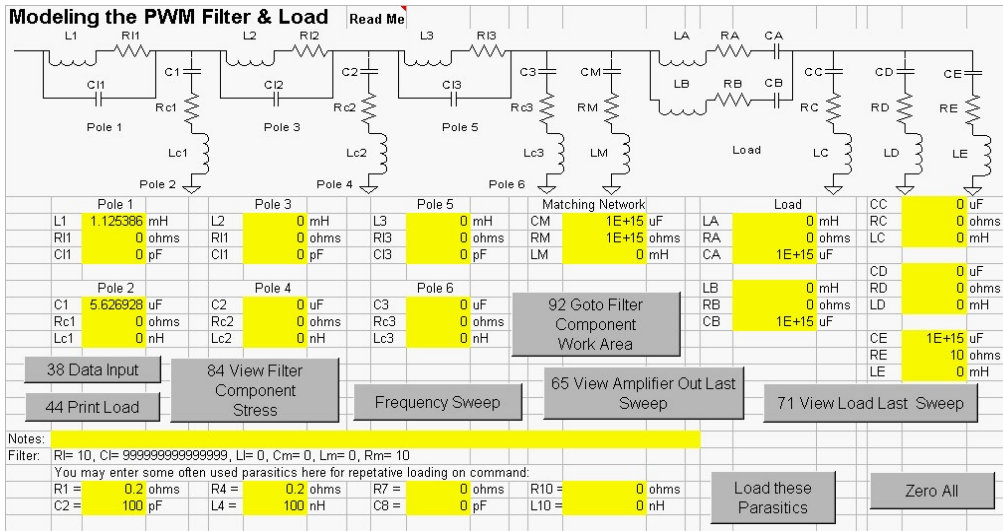


FIGURE 9. LOADING APPLICATION DATA FROM THE FILTER SHEET TO THE POWER SHEET

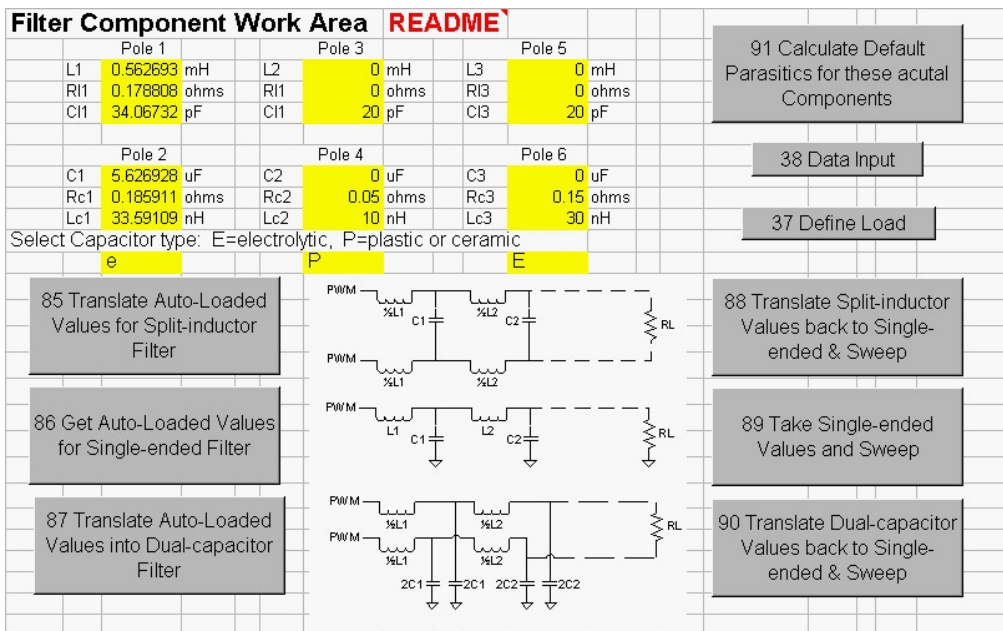


FIGURE 10. TRANSLATION BETWEEN FILTER TOPOLOGIES AND DEFAULT PARASITIC CALCULATION




FIGURE 11. ATTENUATION WITHIN THE PASSBAND

the flux density of the peak currents to be delivered.

Pressing one of the “Load All Data” buttons on the PWM Filter sheet transfers your application to the PWM Power sheet. Starting on the left of Figure 9 we find single-ended components for up to a sixth order filter have been entered. Next we find the matching network. On the far right you will find the simple three-element load specified on the PWM Filters sheet plus space to model a more complex load. Parasitics for the filter components have all been zeroed.

The sweep function handles only single ended filters, but Figure 10 shows the area where these component values can be translated into values for either split-inductor or dual-capacitor designs. While there is absolutely no substitute for finding real parasitic values for filter components, button 91 provides a default parasitic calculator for first pass design efforts. Notice the cells where capacitor type can be selected individually for all three capacitors. Parasitics vary WILDLY from part to part. The default calculator is ONLY intended to get somewhere in the ballpark. These defaults are reasonable for parts suitable for switching applications. Your real parts could be better, but could easily be much worse.

Consult manufacturer’s data sheets or measure the parts to get accurate data for subsequent analysis. Values of purchased components and their real parasitics should be entered directly into the yellow cells and then be translated with button 88, 89, or 90.

This picture is part of the result of loading our sample application from the PWM Filter sheet (no auto sweep); going to the Filter Component Work Area with button 92; translating component values for a split-inductor design with button 85; and calculating default parasitics with button 91. Button 88 will translate prime component values and the parasitics back to single-ended equivalents and then run the frequency sweep to calculate critical voltages, currents,




FIGURE 12. SMALL PARASITICS CAUSE A LARGE DEPARTURE FROM THE "IDEAL" PICTURE AT HIGHER FREQUENCIES

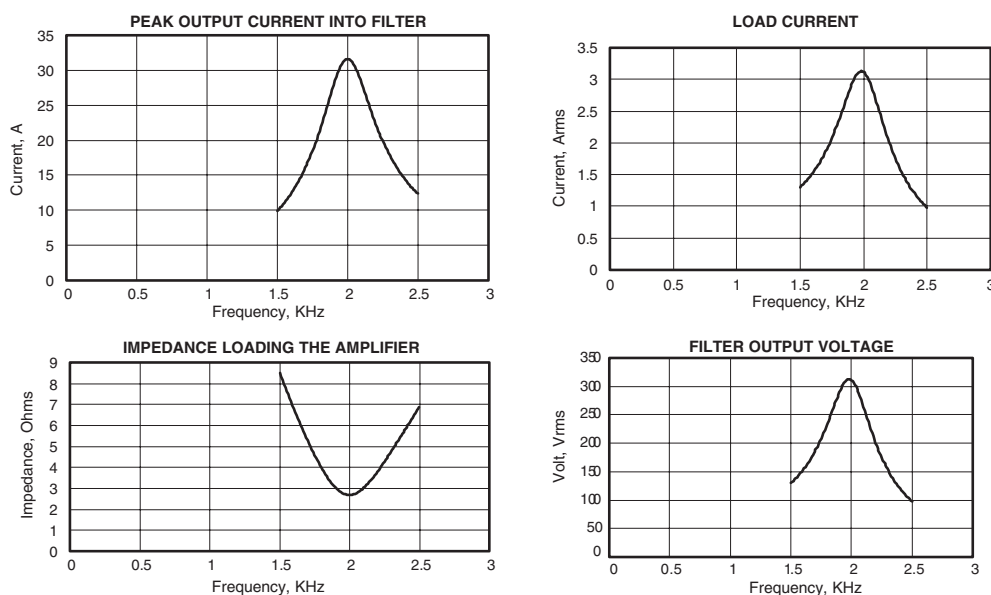


FIGURE 13. A GOOD PROCEDURE FROM THE LINEAR WORLD MAY BE DANGEROUS IN THE PWM WORLD

powers and phase angles over the frequency range we specified. 100 frequency points will be examined. If this takes less than 10 seconds, you should be proud of your computer. If it takes more than a minute -----, Frequency sweep requires Analysis ToolPak. If you see cells with #NAME? or a runtime error, try TOOLS, ADD-INS, Analysis ToolPak and then do the sweep.

Figure 11 shows attenuation of signal frequencies is close to the ideal except the entire curve is about 0.3db lower than expected. This drop is due to inductor resistance. We learned earlier that the extremely fast transition times of the PWM amplifiers means high frequency content is powerful well into the megahertz range. To check performance in this range; go to the data input screen; enter the cutoff frequency as Fmin; and at least the tenth harmonic of the switching frequency as Fmax; and rerun the sweep. The graph in Figure 12 tells us spike content at the filter output is far from ideal. An ideal second order filter for this example has about 82db attenuation at 225KHz, but parasitics reduce this figure to about 73db, or roughly a factor of 3 less attenuation with electrolytic capacitors in this dual-capacitor design. Is this OK? Or should we spend more time looking for better filter components? Or should we

consider one of the other two topologies which will perform better at high frequencies?

So, you're an old hand with linear power circuits. You fire up the prototype with a light load to make sure everything is working before connecting the real load.

While this procedure is commendable for linear drives and may work fine for a PWM drive, watch out for tuned circuits in the filter/match network/load. Replacing a designed 10 ohm/1mH load with a 100 ohm purely resistive load (matching network removed), produces the graphs of Figure 13. At the 2KHz cutoff frequency, impedance presented to the amplifier drops to ~2.7 ohms, peak current tops 30A, load voltage is ~313V and load current is 3.1A. 970W delivered to the light 100 ohm load!

**Be careful!
Deadly voltages
easily generated.**

The second order filter driven at the designed cutoff frequency, with no load, is a series resonant circuit which presents a theoretical zero impedance to the amplifier and develops a theoretical infinite voltage at its center. Higher order filters generally produce lower amplitude peaks at lower frequencies relative to the cutoff frequency.

The very nature of PWM amplifiers demands reactive elements

be driven. Inductance is mandatory and capacitance is very common, meaning resonance will exist. A properly designed and terminated filter will yield a response close to the text book curve. The trick is to design the circuit to accommodate load variations and possibly certain fault conditions such that these conditions will not place undue stress on components or produce extreme high voltage hazards.

Figure 14 (next page) shows the well behaved performance of this example modified for a 10Ω purely resistive load. At DC, impedance loading the amplifier is the sum of the load and the parasitic 0.36 ohms of the filter inductor(s). The amplifier sees about a 3% impedance dip at mid-band and drives a corresponding peak output current. This is normal and Power Design will search for these peaks and dips when calculating heatsink requirements.

While this operation is proper, is it what you wanted? The cutoff frequency of the filter is where the load voltage is down 3db. Does -3db equal .707 or .5? Both, .707 applies to the voltage and the current ratio but .5 is correct for the power ratio. In this example output power drops from 337W at DC to 172W at 2KHz. Many times half power at maximum fre-

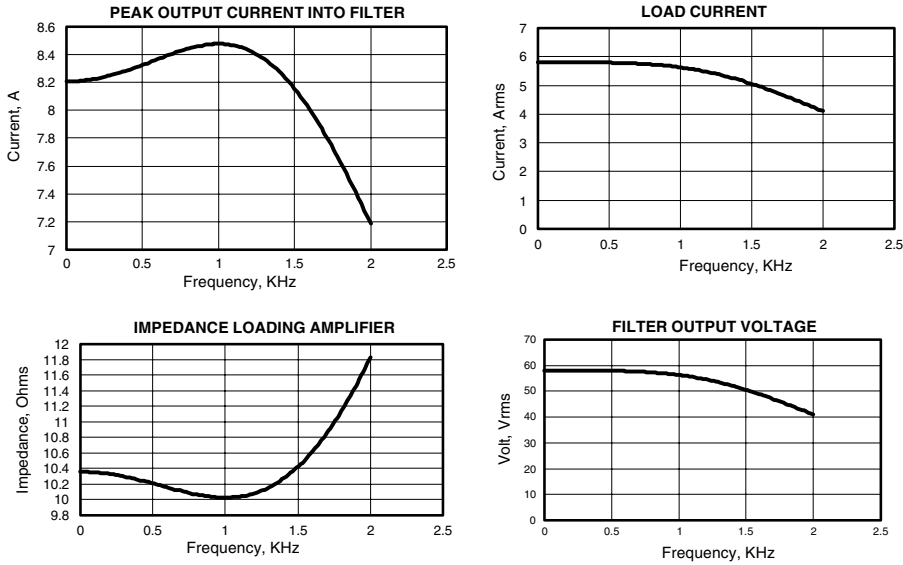


FIGURE 14. PERFORMANCE OF THE PROPERLY TERMINATED FILTER

quency is not acceptable. This is why some designers routinely start their filter calculations using a cutoff frequency twice the required maximum signal frequency of the application.

Doubling the design cutoff frequency of the filter enables the circuit to deliver a lot more power at the desired 2KHz as shown in Figure 15. This is still a 2 pole filter and power loss is only about 6% at 2KHz. As an added benefit, inductors for a 4KHz filter are half the value of those for a 2KHz filter and likely will have substantially lower parasitic resistance.

Yes, you could double again to achieve an even flatter pass band. No, there is no free lunch. Every time you move cutoff frequency up, you allow more switching frequency power in the load. Yes, you can add more poles to the filter. Analyze as many combinations as you wish, it won't take long. The question becomes one of cost in terms of money, extra loss in the




FIGURE 15. DOUBLING F_c YIELDS INCREASED OUTPUT POWER

filter, size and weight.

Speaking of properly terminated filters, we need to look closely at the matching network. While the conjugate matching network performs almost like magic in terms of forcing the attenuation graph to approach text book shape, there is a cost involved. This cost is slight when the load is mostly resistive, but the power dissipated in this network approaches power delivered to the load as the load approaches pure reactance.

The graphs of Figure 16 are for an application driving a 1uF piezo stack with 12 ohms series resistance, to 75V peak from 1KHz to 20KHz. The second order filter cutoff frequency was designed for 40KHz providing a quite flat response. The V-A output falls at low frequency because the load impedance is increasing. To keep filter termination impedance flat, the matching network impedance moves in the opposite direction, giving rise to large power levels in the matching network resistor, especially at low frequencies. As this power is not delivered to the load, efficiency is far from the desired level.




FIGURE 16. PERFORMANCE OF THE PROPERLY TERMINATED FILTER

Upon seeing this power loss, some designers immediately want to see what happens if they simply remove the matching network. With no matching network we cannot lose this power, but this leaves the filter with an improper termination. This would result in a unwanted resonant circuit causing almost 4db peaking as shown in Figure 17 (next page). In terms of V-A in the load near the upper end of the band, power goes from ~180 to over 450V-A. If you wish to see more on this subject, use the Filter Hazard on the Power Design Examples sheet. Comments explain each situation and a macro sets up and runs the analysis.




FIGURE 17. THE BAD NEWS IF THE MATCHING NETWORK IS OMITTED

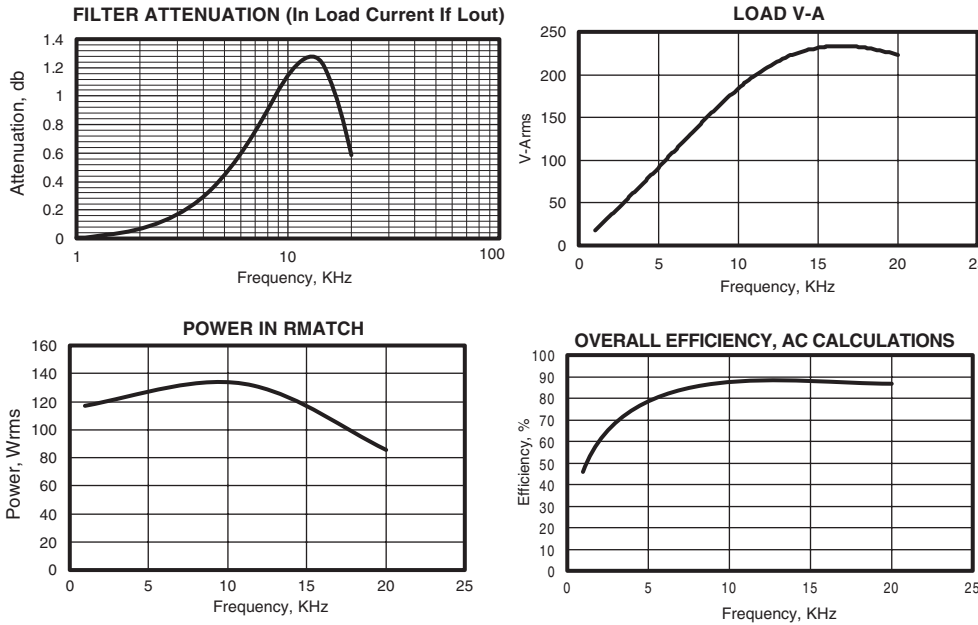


FIGURE 18. RESULTS OF A MODIFIED MATCHING NETWORK

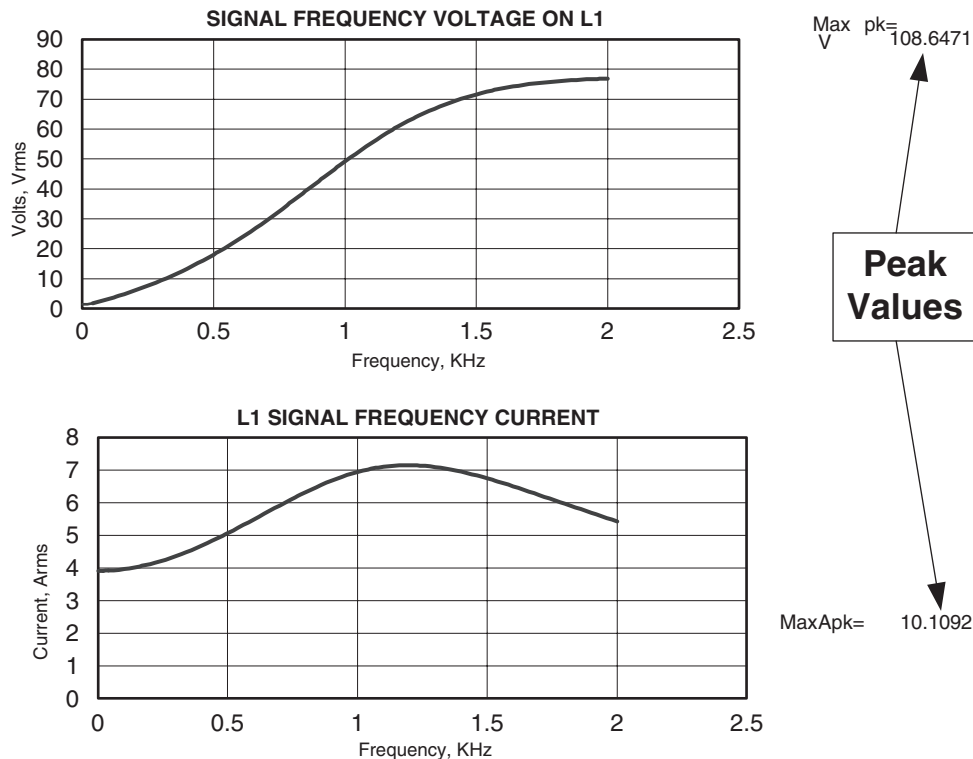


FIGURE 19. STRESS LEVELS ON L1

Here lies part of the beauty of the Power Design tool; investigating possible compromise circuits is a snap. See Figure 18 for results of doubling the resistor value in the matching network which may provide a workable compromise. Peaking at the load is down substantially from not using any network and wasted power is down substantially from using the ideal network.

DETERMINING FILTER COMPONENT STRESS LEVELS

This section uses the original coil driver example, with the second order filter designed for cutoff frequency=2KHz; load resistance=10Ω; and load inductance=1mH. However, we are assuming the load has gotten hot and the resistance has gone up to 15Ω. This change affects all the performance graphs covered so far and they should be checked. Power Design calculates voltage and current stress levels on L1 and L2, plus C1 and C2 for all designs. Resonance of these filters can produce voltages and currents larger than the load levels. Button 84 will place the first graph on the screen, then scroll up and to the right to view other graphs. The currents shown in Figure 19 can be used directly for all filter topologies. If L1 is actually two inductors, half the voltage shown will be across each individual inductor. Note that our circuit example only has a 90V supply; the drive signal is only 85Vpk; the load resistance is 15Ω; but L1 has current peaks of 10.1A and voltage peaks of 108V.

In Figure 20 (next page), we find that in addition to the switching frequency current, C1 has 3A flowing at 1.6KHz. Three peak voltages are given; 96Vpk for a differential capacitor; a positive peak of 93V for a grounded capacitor; and a negative peak of -3V. Anytime the negative peak is below ground, the warning to use

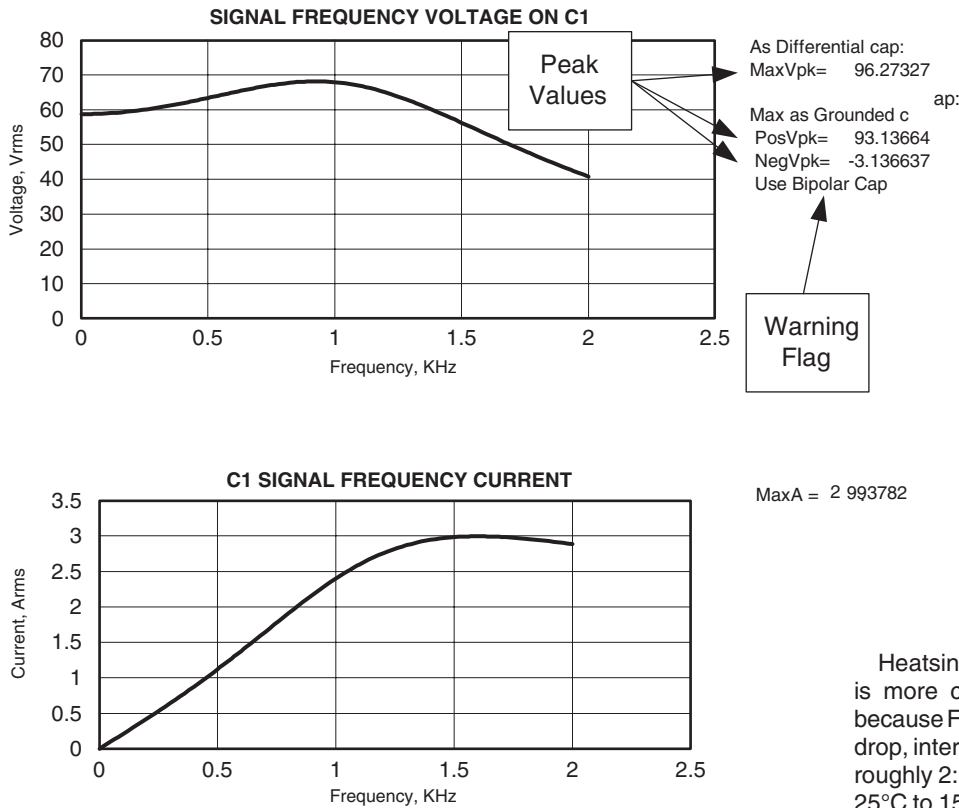


FIGURE 20. STRESS LEVELS ON C1




FIGURE 21. VOLTAGE AND CURRENT PHASE IN THE LOAD

bipolar capacitors also appears.

PHASE ANGLES AT THE LOAD

These filters are notorious for introducing large phase shifts. This is usually not a problem when feedback is taken directly at the output of the PWM amplifier. In applications such as servo loops, feedback is taken after the filter and any phase shift introduced here affects system phase margin. Figure 21 shows both voltage and current phase in the load for this example. This phase shift is reduced as the ratio between Fmax and Fcutoff frequencies widens, and is lower with lower order filters.

CALCULATING INTERNAL POWER DISSIPATION FOR PWM AMPLIFIERS

Heatsink selection for most PWM amplifiers is more complex than for a linear amplifier because FET ON resistance (and hence voltage drop, internal power and dissipation) increases roughly 2:1 as junction temperature goes from 25°C to 150°C. PWMs have the same concern over temperature vs. life expectancy as linears, but changes of circuit performance over temperature are much more pronounced than with linear amplifiers.

For a first order estimation of power dissipation in the PWM, simply multiply the output current (a given application requirement) and the voltage drop at that current (read from a graph on the product data sheet). This points out the PWM advantage over linear power delivery; supply voltage is not part of the equation. With a first order approximation, the voltage drop divided by supply voltage yields efficiency (quiescent current of both Vcc and Vs will reduce this a little). Unfortunately, first order approximation is not good enough unless you have the luxury of using overkill amplifiers and massive heatsinks.

Looking a little deeper, there are two points of confusion. First, the voltage graph offers multiple curves based on various case temperatures. We know cooler is better for life expectancy and efficiency, but there are no rules regarding which one to choose. Something not presented in any direct way is the second problem: methods to calculate junction temperature are not given. This is a parameter every power designer should know and it is often specified by contract.

While the linear Power sheet simply provides you with a minimum heatsink rating, the PWM Power sheet gives you graphs of junction temperature, internal dissipation and efficiency. With this data, you can make intelligent tradeoffs concerning circuit operation vs. investment in the heat removal system.

Our exercises on filter design have already taken us to the PWM Power sheet. From numer-

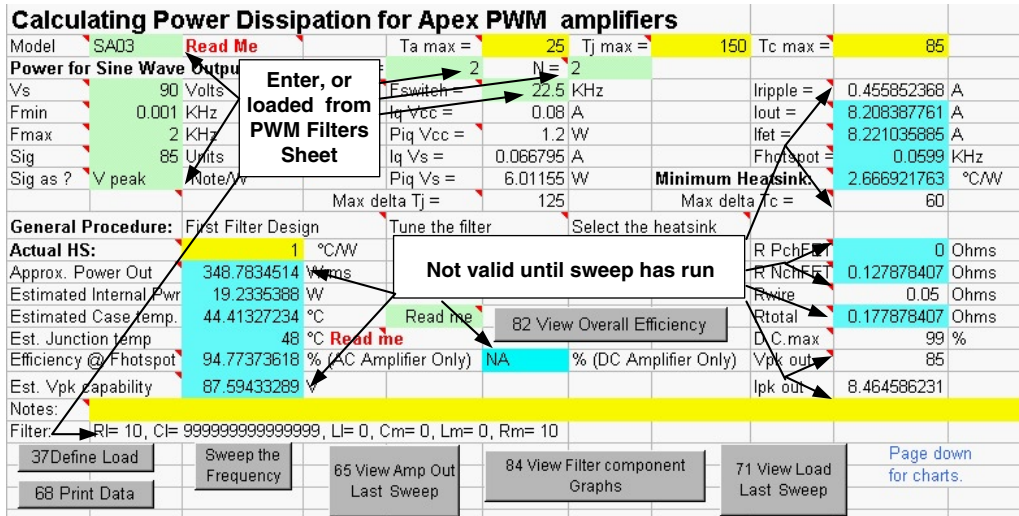


FIGURE 22. THE PWM POWER SHEET DATA INPUT SCREEN

ous locations you can use command button 38 to see the Data Input screen as shown in Figure 22. The green input cells are normally filled in with one of the Load Data command buttons on the PWM Filter sheet. These values may be changed at will but neither the graphs nor circuit parameters in the blue cells will necessarily be valid until a Frequency Sweep has been run.

In the center area, find data on quiescent current and resulting power dissipation. Data for these calculations comes from application parameters and the built-in database containing information on quiescent current variations with supply and switching frequency. In the upper right, applying a sine wave at the switching frequency and using a correction factor have approximated ripple current. If the filter and load are close to those entered in the Filter sheet, this ripple approximation will be close to ripple predicted by the Filter sheet. Hotspot frequency is the frequency where load current is producing the highest junction temperature. This is not necessarily coincident with the frequency producing the highest peak current. Consider the case of a DC current of 10A rising to 11A peak at 1KHz. At 1KHz heat generation is alternating between pairs of transistors fast enough to find them running cooler than at a steady state level of 10A.

The Iout cell will report amplifier output current at signal frequencies using peak values for hotspot frequencies below 60Hz or RMS values at higher frequencies. The Ifet cell is the RMS addition of the ripple current (at the switching frequency) and the output current. If hotspot frequency is 60Hz or more, ripple current is reduced 30% because ripple decreases as duty cycle increases (there is no ripple at 100%). The Minimum Heatsink is the thermal rating which will keep both the case temperature and the junction temperature within the boundaries entered at the top of the screen.

Continuing down on the right side are On resistances of the H-bridge switches at the hotspot frequency. The P channel number will be for one FET if P channel devices are used in the amplifier. If this is the case, the N channel number will also be for one FET. When only N channel FETs are used, this number will be for two FETs. If you specify a heatsink that will not keep junction temperatures below the specified maximum, both FET resistances will be forced to 10 ohms. Rwire represents internal conductor losses in the amplifier and Rtotal adds it all up. For amplifiers using IGBTs, all resistance cells will be blank.

Most PWM amplifiers can hold their output switches in one state. To rephrase, PWM amplifiers can be driven to zero or

100% duty cycle; however, propagation delays and dead time requirements limit the linear modulation range to less than these levels. D.C. max is the maximum percentage of the power supply voltage delivered before encountering the non-linear jump to being latched in one state. Vpk out and Ipk out are from anywhere between Fmin and Fmax.

Back on the left side under Actual HS, Approx. Power Out is the power factor corrected VA output directly at the amplifier at the hotspot frequency. If

you really intend to look at DC or want a peak value, multiply by two. Estimated Internal Pwr includes losses due to driving the load and the quiescent power.

Efficiency of a DC power supply would compare DC, or peak power out, to input power. Efficiency of an audio amplifier would likely compare RMS power out to input power. These two approaches to efficiency will produce different answers for the same amplifier, driving the same peak current from the same power supply. Power Design always calculates an efficiency based on the AC thought process but will give you a DC based answer only if both Fmin and Fmax are less than 0.003 (remember to Sweep before reading the answer). In both cases, the numbers appearing here do not include filter loss. Est. Vpk capability subtracts internal losses and duty cycle limitations from the supply voltage.

Not shown in Figure 22 are four data dependent red warning flags. The first warning appears if the Actual HS is too small to maintain either specified case temperature or junction temperature. If Est. Vpk capability is less than the signal voltage, the next flag will become visible. The third flag warns of output current beyond the peak rating of the amplifier. The last flag concerns application supply voltage compared to amplifier ratings.

We know many of these numbers are a moving target relative to selection of the actual heatsink rating. Refer to Figure 23 (next page) for this important step. All heatsink references assume fresh thermal grease has been properly applied or an Apex thermal washer has been used.

In the upper left graph it looks like a quite small heatsink will keep junction temperatures below the maximum of 150°C specified by almost all PWM amplifiers. However the graph below says there is little difference between junction and case temperature and we surely want to keep case temperature much lower than 150°C.

On the top right we see that internal power dissipation of the amplifier changes with junction temperature. Settling for the minimum heatsink size instead of investing in a 1°C/W heatsink (easy to do without a fan) will increase internal power almost 8W or nearly 40%. Below we see this same effect expressed in terms of efficiency. At first, moving only a few percentage points may not seem like much, but remember these points are relative to a quite large power level. Enter 100 as the Actual HS if you plan to not use a heatsink. More likely, a rating will come from the Heatsink sheet of Power Design, a manufacturer's data sheet, or your own design efforts.

The efficiency graphs above refer to performance only at

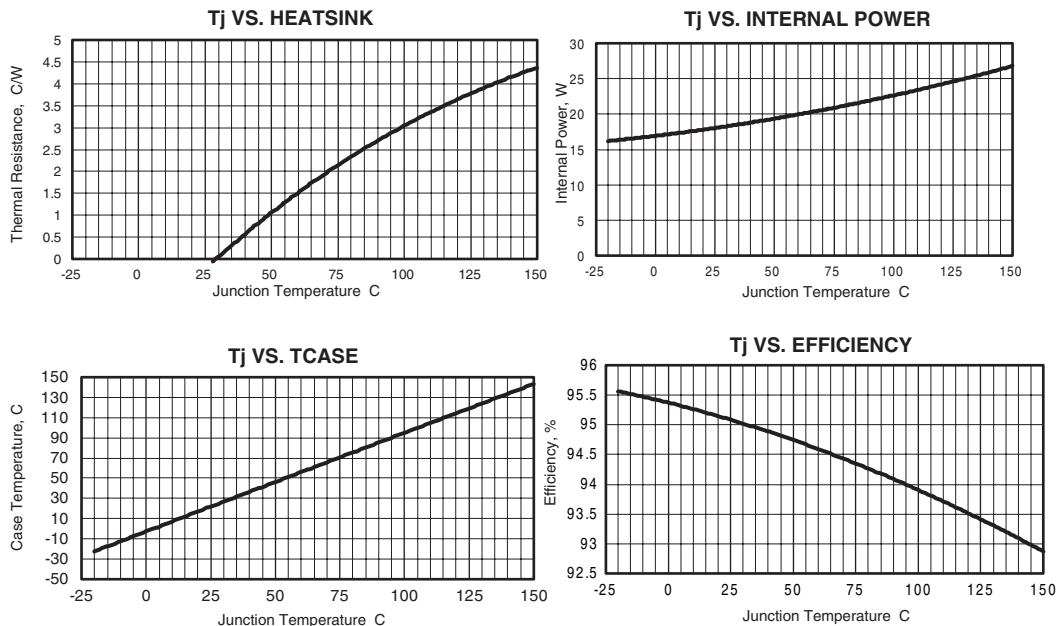


FIGURE 23. SELECTING A MAXIMUM JUNCTION TEMPERATURE IS THE KEY STEP

the hotspot frequency and do not include filter losses. Figure 24 however, includes losses in the filter and matching network and provides frequency data. The curve is based on the AC thought pattern discussed earlier, input power compared to RMS power delivered. If there is a glitch at 60Hz, it is due to the instant change (mathematically) from peak power heating effect to RMS power heating effect. Fig. 24. Including filter losses in the big picture.




FIGURE 24. INCLUDING FILLER LOSSES IN THE BIG PICTURE

CONCLUSION

The Power Design tool is even more important to the PWM designer than the linear designer. PWMs are not as widely understood and worse yet, literature is not as widely available. The comments and automated examples built into this spreadsheet serve well as a text on the subject. PWMs also tend to require more iteration to approach an optimum design and are more frequency sensitive than linears. Again, tackling all this with computer aided design is the only way to go and the tasks of filtering and heatsinking are better handled by Power Design than by many Spice machines.

Not to harp on it, but do not let all this quick and easy data lead you into the trap of accepting theory with a smile while sticking your head in the sand when it comes to the hardware world and parasitics. And one more time: Without case temperature measurements, your design effort is NOT complete!