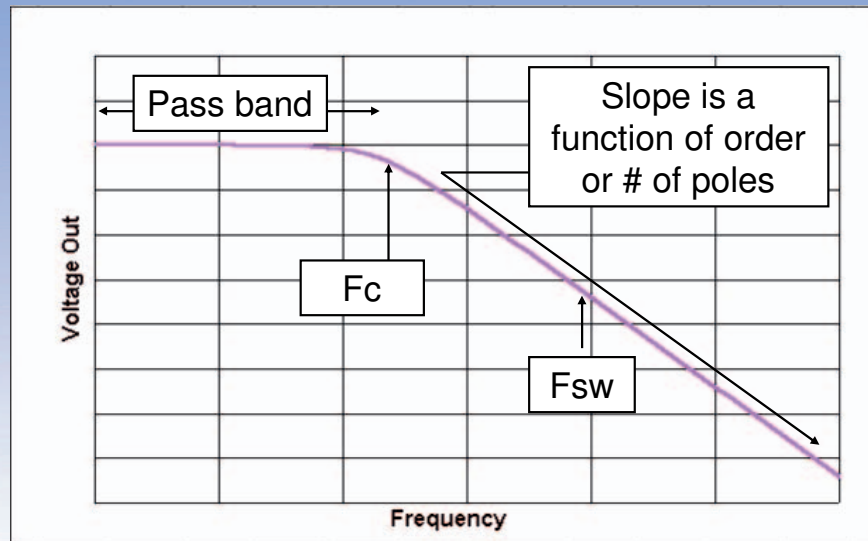


Low Pass Filter Response




PWM filters are normally low pass configuration. These exhibit low attenuation to the frequency spectrum from 0 Hertz to the frequency of cutoff (F_c). This low attenuation region is called the pass band. Beyond the F_c , attenuation increases at a rate determined by the filter type and the numbers of poles (order).

As we speak of the “frequency” of a PWM signal it is very important to realize what area of this response curve is being referred to. In the pass band area, signals are slow and can be thought of as analog. The switching frequency will be well beyond F_c , can be thought of as the carrier frequency containing time modulated digital information. A reasonable analogy for the filter function might be the audio CD technology where high speed DACs translate digital data to analog output, where the D-to-A conversion rate corresponds to switching frequency. Going even further out in frequency, where the high speed transitions of the PWM amplifiers generate spikes, it is best to think of RF energy.

Ref. AN32,AN39

PWM Frequency Relationships



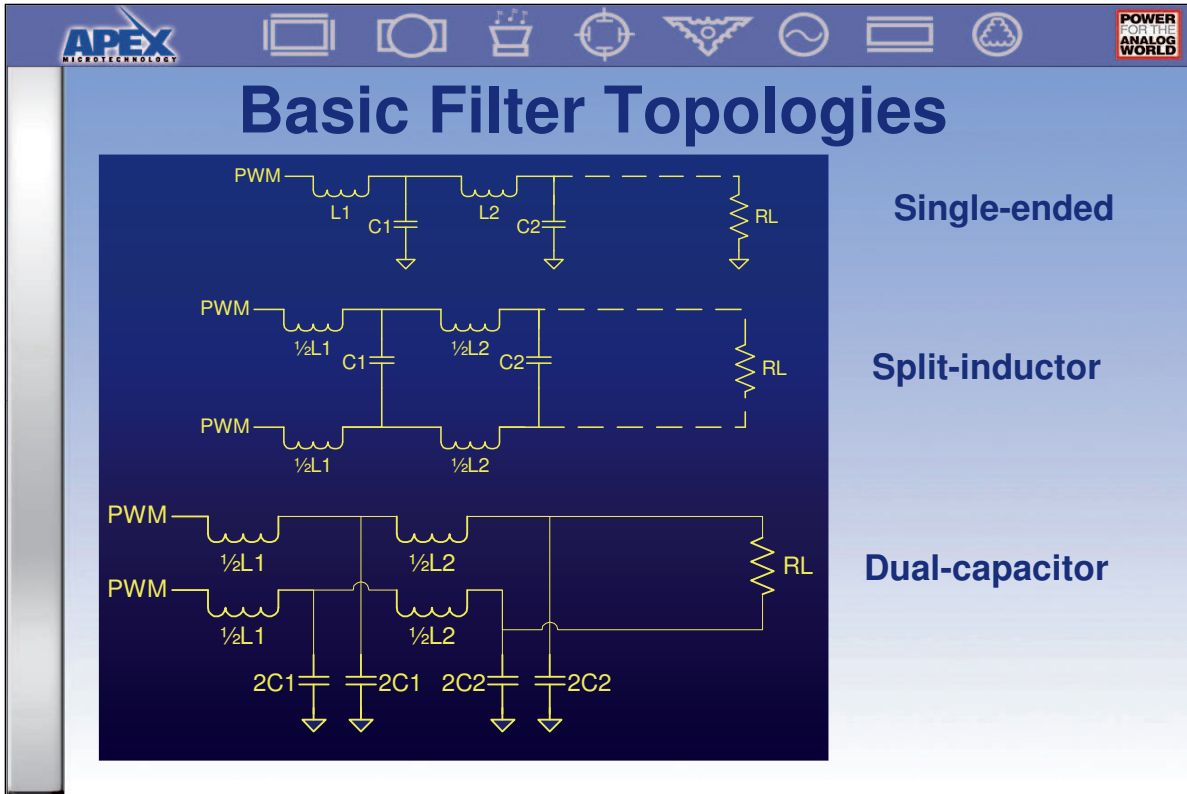
The common ramp generator illustrated the relationship between oscillator and switching frequencies. Some PWM data sheets (such as the SA01) do not mention oscillator frequency because there is no divide by two circuit.

Signal frequency is that of the power drive to the load, power bandwidth. Between the load and the PWM amplifier is the low pass filter (or at least the model of one if the load is also the filter). On the input side of the filter we have the switching frequency. We then go down the slope to a point where the attenuation is adequate. The frequency band we cover while going down the slope is required spacing between the switching and signal frequencies.

Pure theory says filter slope can be increased simply by adding more poles. 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 a decade between switching and signal frequencies.

Ref. AN32,AN39



No matter what topology is used, a first order filter would use only L_1 , a second order adds C_1 , a third order adds L_2 , and so on. Each pole of the filter adds 20db/decade to the slope or roll-off of the filter.

The single-ended filter configuration is the simplest; must be used with half bridge circuits; and can be used with full bridge circuits by substituting the second PWM output for all the ground connections. This substitution is very rarely done because it places the high speed square waves of the PWM output on both load terminals and all the cabling between the amplifier and load. With rise and fall times usually in the tens of nanoseconds, and amplitude nearly equal to supply voltage, this is an extreme RFI problem.


With full bridge circuits, an additional filter requirement is introduced in that common mode voltage applied to both load terminals usually needs to be minimized. The technique to achieve low common mode voltage is to simply split the inductor values in half, applying half to each PWM output as shown in the split-inductor topology.

Capacitors of the split-inductor topology must be capable of bipolar operation and will be very large when the filter is designed for both high current and low signal frequency. While the bipolar capacitors exhibit very low ESL and ESR to provide good roll off in the high frequency spectrum, this leads to very large and expensive banks of capacitors. The dual-capacitor topology can provide a cost savings, at the expense of high frequency performance, by substituting a pair of electrolytic (or possibly tantalum) capacitors of twice the size. To convince yourself this a valid substitution, forget the ground connection and think of two series connected capacitors in place of one. This substitution usually allows the use of unipolar capacitors.

If one could acquire a perfect PWM amplifier (equal rise and fall times, no dead time plus an exact out of phase condition) and perfectly matched inductors, current through each of the dual capacitors would be equal and phased such that no current would flow into the ground node. Even with these imperfections, the ground node current will be a small percentage of the capacitor current.

Ref. AN32

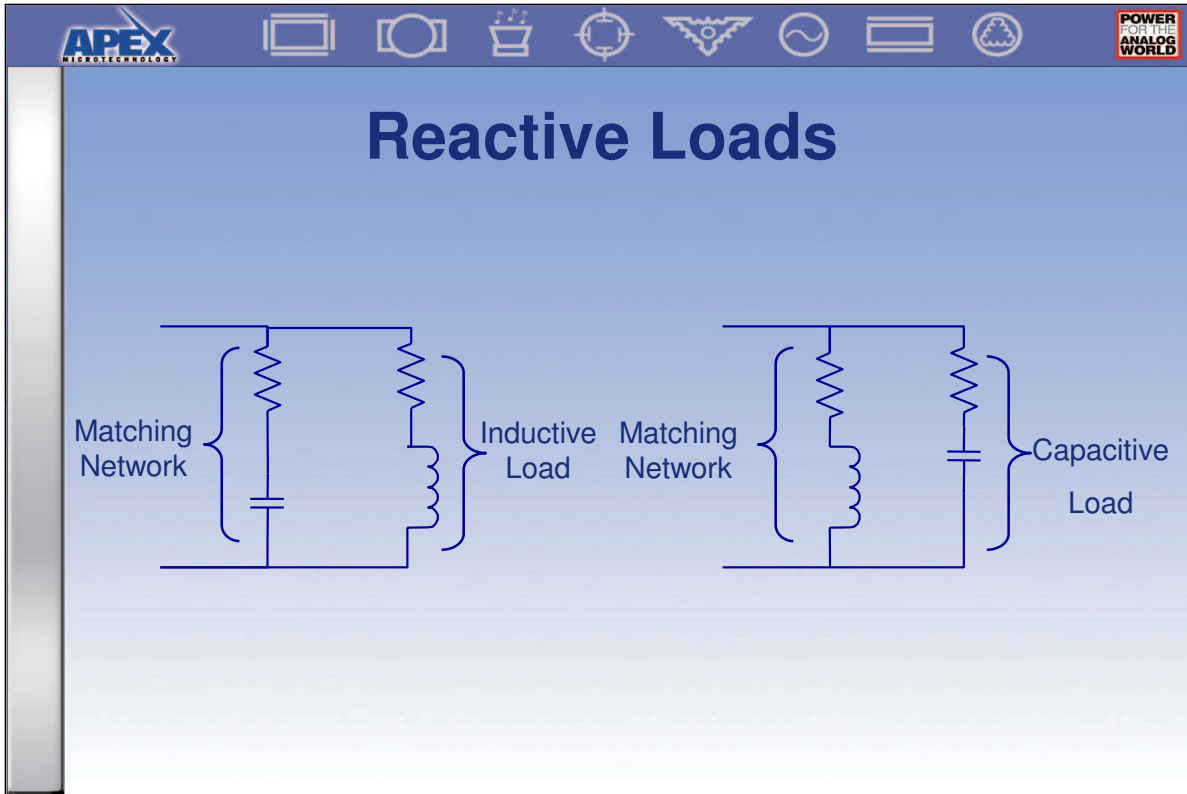
The Most Common Filter



If we apply our previously mentioned fantasy of perfect components to the split inductor filter topology, common mode voltage on the load terminals will be zero. With real PWM amplifiers, the output will contain large amounts of high frequency harmonics. Each application is different, but peak-to-peak noise amplitude may approach the supply voltage. The spectral content of this noise extends well above the switching frequency. A pair of small capacitors added from the output side of each half of L1 to ground will remedy this problem. It is not necessary (and sometimes it is counterproductive) to use more than this one pair of leg capacitors. Placing these small capacitors on the load side of L2 or L3 is not as effective as the placement shown.

Value selection for these ground leg capacitors is less critical than for the main filter capacitors. It has been determined empirically that setting the impedance value of these capacitors at the cutoff frequency, to between 10 to 30 times the value of the load resistance will provide reasonable common mode filtering. The addition of these capacitors will typically produce no more than 0.05db peaking, nor more than 0.2db change at the cutoff frequency in any order filter. From the technical point of view, the two Clegs are in series, and this is in parallel with C1. This means that on all but first order filters, C1 could be reduced by half the value of Cleg to eliminate even these small errors.

Ref. AN32



To achieve even close to these ideal filter responses a constant and purely resistive load termination is required. If a reactive load can be modeled as resistance in series with either capacitance or inductance, a simple conjugate match network can be used to achieve proper termination. The resistor in the network will equal the resistor of the load model. As the network is in parallel with the load, all signals in the pass band will be applied to the network and power dissipation must be checked. Realize that combined impedance of the network plus load is constant and that changing frequency shifts the power between the network and the load. This means a 100W capacitive load drive will require a 100W matching network if DC signals are allowed.

Ref. AN32

APEX
MICROTECHNOLOGY

POWER FOR THE ANALOG WORLD

The Toolkit

- Power Design.exe
 - Very fast, dedicated topology, minimal bandwidth data
- State average spice models
 - Fast, any topology, AC sweeps
- Pulse-by-pulse macro models
 - Slow, no AC sweep, switching data
- The bench
 - Parasitics included
 - Light loads are a danger
 - Low supply voltage is OK, but

Power Design is an Excel97 spreadsheet provided in three zipped flavors; the linear portion, the PWM portion or in total. While it has macros, Apex has taken care to make sure it is virus free. The macros make it easy to work with and fast. The first time running it, you will go to the Tools menu, Addins, Analysis ToolPak. There are many comments covering use and amplifier facts. It is protected, but there are no passwords or hidden sheets.

PWM spice models require Berkeley3 based or PSpice4 based platforms. Linear models usually run on older platforms also. State average PWM models run much faster than pulse-by-pulse models, can run an AC sweep, but provide no switching data.

With frequency components well into the RF range and power into the KW range, capacitive, inductive and resistive parasitics all mandate bench confirmation of a design. If years of linear power design taught you to start with light loads and a low supply voltage, beware:

Improperly terminated filters can generate voltages greater than the supply.

Frequency response and stability change as supply voltage changes.

Filter Design for PWM Amplifiers READ ME

CAUTION! Refer to Applications Note 32

Using the Complex Load:

Input Data

Model	SA03
Vs	90 Volts
Fsw	22.5 KHz
Fmin	0.001 KHz
Fmax	2 KHz
Fcutoff	2 KHz
Rload	10 Ohms
Cload	0 uF
Load	0 mH
Vripple	1 Vpk
Signal	85 Units
Sig as	V peak

Order Calculation

Atten. @ Fsw	41.023 db
N(exact)	1.9513
N(recommended)	2

Matching network

Cm =	0 uF
Lm =	0 mH
Rm =	10 Ohms

Using the Complex Load:

- 60 Load All Data For N=1
- 61 Load All Data For N=2
- 62 Load All Data For N=3
- 63 Load All Data For N=4
- 64 Load All Data For N=5
- 65 Load All Data For N=6

Notes:

46 Print Filter 55 Show Attenuation in db & %

56 Show Attenuation Graph 66 Show Filter Components

So maybe filter design is not at the top of your list of most cherished jobs. Application Notes 32, 39 and the Power Design spreadsheet can help. Enter data describing the amplifier circuit, the load and desired attenuation. Placing the cursor in cells with red triangles will display notes of explanation. The order Calculation section converts your maximum ripple spec into db attenuation and by examining the switching and signal frequencies, it calculates the order, or number of poles needed. The matching networks calculated will cause reactive loads to appear resistive to the output of the filter, Finally, a capacitor value is recommended for the leg capacitors for a dual-inductor.

power_design.zip is a free download from www.apexmicrotech.com. When unzipped, Power Design.xls will be extracted, ready to be run with Excel97.

Ref. AN39

Ideal Components

Component Calculations		Shading indicates values for Split Inductor topology					
Dual Cap Filter		Single-ended Filter		Dual Cap Filter		Single-ended Filter	
N = 1	L = 0.3979 mH		0.7958 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. I _{out} for thermal calculations =		0.6283		P-P ripple = 1.7772 Amps out of the amplifier		
					Avg. I _{out} for thermal calculations =		0.4443
N = 3	L1 = 0.5968 mH		1.1937 mH				
	C = 21.22 uF		10.61 uF	N = 4	L1 = 0.609 mH		1.2181 mH
	L2 = 0.1989 mH		0.3979 mH		C1 = 25.102 uF		12.551 uF
	P-P ripple = 1.6755 Amps out of the amplifier				L2 = 0.4307 mH		0.8613 mH
	Avg. I _{out} for thermal calculations =		0.4189		C2 = 6.0909 uF		3.0454 uF
					P-P ripple = 1.6419 Amps out of the amplifier		
N = 5	L1 = 0.6148 mH		1.2296 mH		Avg. I _{out} for thermal calculations =		0.4105
	C1 = 26.967 uF		13.484 uF				
	L2 = 0.5499 mH		1.0998 mH	N = 6	L1 = 0.6179 mH		1.2358 mH
	C2 = 14.235 uF		7.1174 uF		C1 = 28 uF		14 uF
	L3 = 0.1229 mH		0.2459 mH		L2 = 0.6179 mH		1.2358 mH
	P-P ripple = 1.6266 Amps out of the amplifier				C2 = 19.124 uF		9.562 uF
	Avg. I _{out} for thermal calculations =		0.4067		L3 = 0.3016 mH		0.6031 mH
					C3 = 4.1189 uF		2.0595 uF
					P-P ripple = 1.6184 Amps out of the amplifier		
					Avg. I _{out} for thermal calculations =		0.4046


56 Show Attenuation Graph 38 Data Input

Application Note 32 will provide filter coefficient tables and formulas if you insist on calculating component values the hard way.

Values for dual-capacitor and single-ended filters are found under the appropriate columns, for orders up to six. Six is generally higher than is practical because of cost and diminishing returns due to component parasitics and stray coupling. To build a split-inductor filter, use values in the shaded areas from both columns.

P-P ripple calculations refer to current in L1 at the switching frequency when a 50% duty cycle is present. The *Avg. I_{out} for thermal calculations=*, is the average current through one PWM switch and can be used for determining junction temperature.

Ref. AN32,AN39



“Ideal” is a great word. In this case it means most of the work still lies ahead in finding components which work acceptably in the MHz range and whose losses won't kill you at high current levels.

For capacitors, this often means parallel bipolar devices to obtain high value and high frequency performance. You will probably want ceramic for the smallest values and plastic for the higher values. For the largest capacitance values tantalum, or electrolytic types, can often be used in the dual-capacitor topology with some loss of high frequency attenuation.

Finding suitable inductors is also challenging. Air core inductors get away from the magnetic saturation problem and they have less tendency 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 at the switching frequency and can accommodate the flux density of the peak currents to be delivered to the load. Ferrite and powdered iron cores hold the most promise; avoid laminated steel cores.

Ref. AN32,AN39

Modeling the PWM Filter & Load Read Me

Pole 1		Pole 3		Pole 5		Matching Network		Load		Load	
L1	1.125386 mH	L2	0 mH	L3	0 mH	CM	1E+15 uF	LA	0 mH	CC	0 uF
R11	0 ohms	R11	0 ohms	R13	0 ohms	RM	1E+15 ohms	RA	0 ohms	RC	0 ohms
C11	0 pF	C11	0 pF	C13	0 pF	LM	0 mH	CA	1E+15 uF	LC	0 mH
Pole 2		Pole 4		Pole 6		92 Goto Filter Component Work Area		LB	0 mH	CD	0 uF
C1	5.626928 uF	C2	0 uF	C3	0 uF			RB	0 ohms	RD	0 ohms
Rc1	0 ohms	Rc2	0 ohms	Rc3	0 ohms			CB	1E+15 uF	LD	0 mH
Lc1	0 nH	Lc2	0 nH	Lc3	0 nH					CE	1E+15 uF
										RE	10 ohms
										LE	0 mH

Notes:
 Filter: R1= 10, C1= 9999999999999999, L1= 0, Cm= 9999999999999999, Lm= 0, Rm= 9999999999999999
 You may enter some often used parasitics here for repetitive loading on command:
 R1 = 0.2 ohms R4 = 0.2 ohms R7 = 0 ohms R10 = 0 ohms
 C2 = 100 pF L4 = 100 nH C8 = 0 pF L10 = 0 nH

Buttons: 38 Data Input, 44 Print Load, 84 View Filter Component Stress, Frequency Sweep, 65 View Amplifier Out Last Sweep, 71 View Load Last Sweep, Load these Parasitics, Zero All

Pressing one of the “Load All Data” buttons on the PWM Filter sheet transfers your application to the PWM Power sheet. Ideal component values are loaded automatically for all six pole elements, the matching network and on the far right, the load we specified earlier. Extra components in the load modeling area provide more flexibility. As the math (and execution time) would be a significantly larger burden for any other topology, Power Design only analyzes single-ended filters.

Note that the horizontal load model components are “zeroed” with no resistance, no inductance, but an extremely large capacitance. Unused components in the vertical orientation require zero capacitance or extremely large values of resistance or inductance.


The Frequency Sweep button will calculate critical voltages, currents and powers 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

The Goto Filter Component Work Area button will be used to translate component values between the three topologies and for first pass design work, to estimate parasitic values.

Ref. AN32,AN39

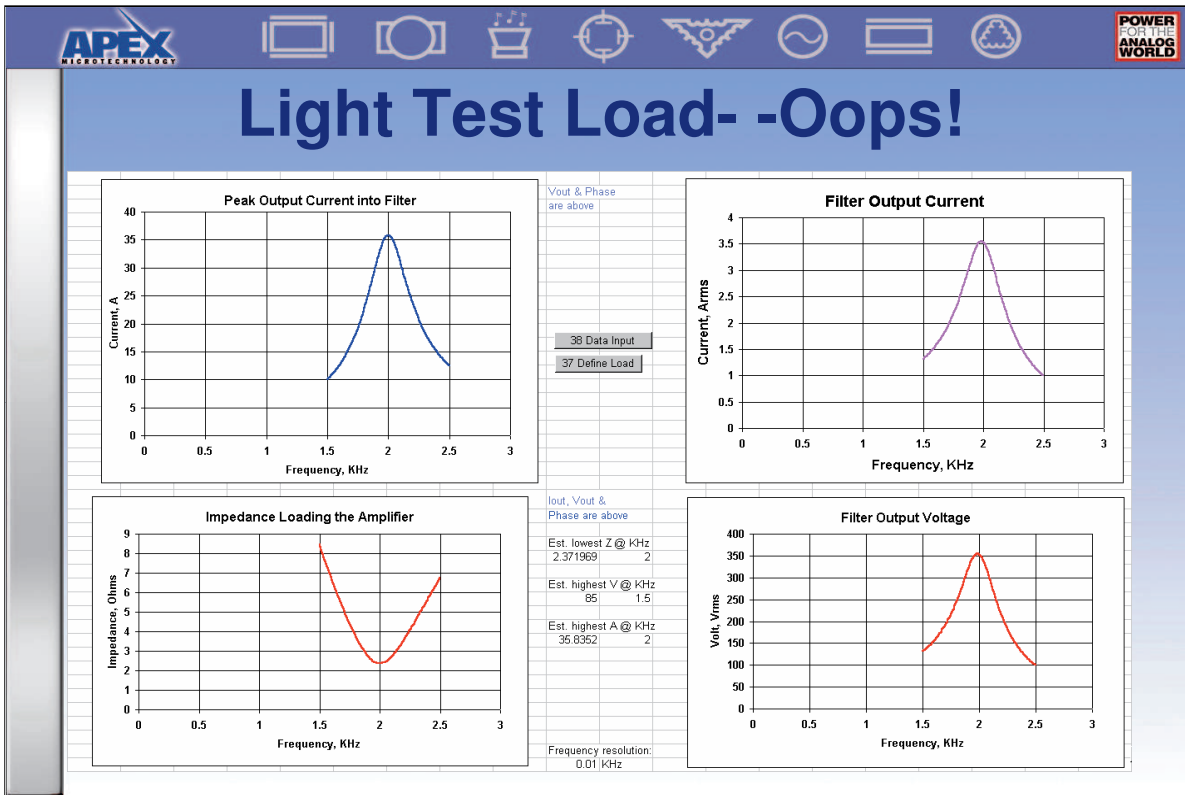


The Penalty of Real Parts



Attenuation is about as expected up to 200KHz, but then the parasitics come into play. We learned earlier that the extremely fast transition times of the PWM amplifiers means high frequency content is powerful well into the megahertz range. This graph is telling us spike content at the filter output is far from ideal. Is this OK? Or should we spend more on better filter components?

Ref. AN32,AN39



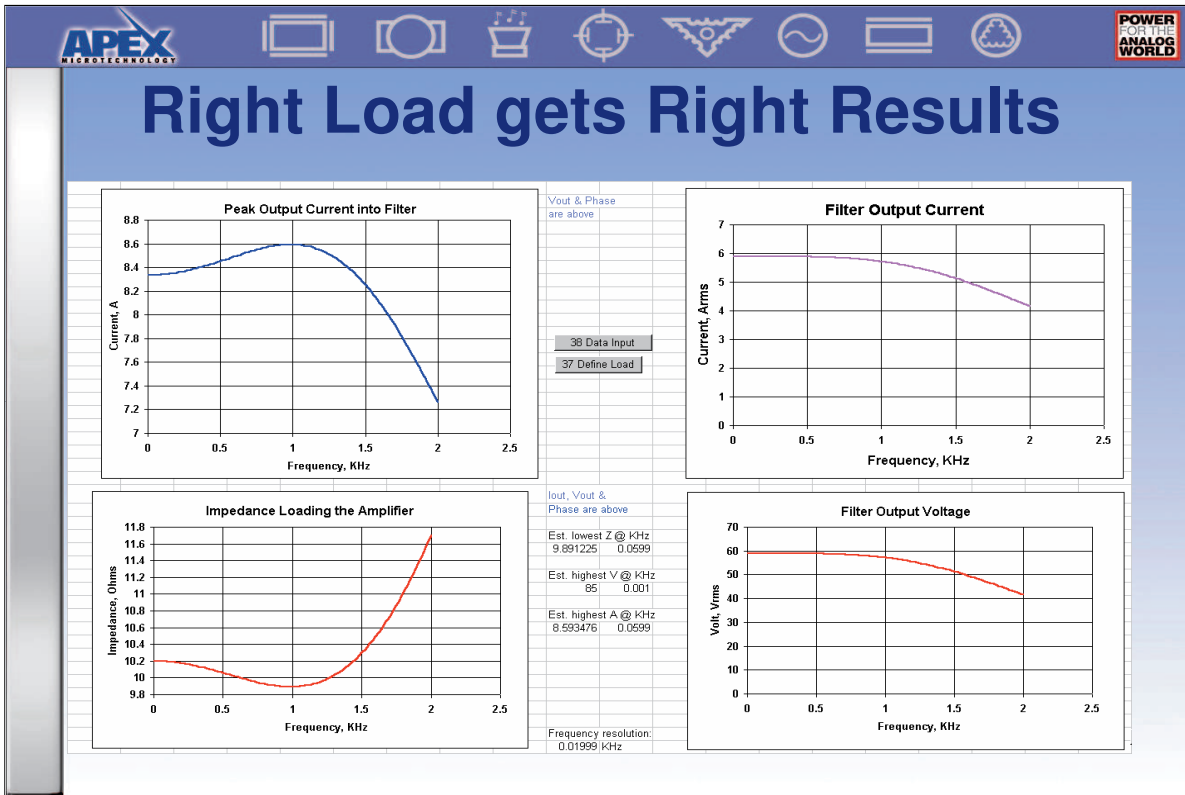
So, you're an old hand with linear power circuits; you fire up the prototype with a light load to make sure everything is working before connecting the real load.

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 the designed 10 ohm load with 100 ohms produces the graphs above. At 2KHz impedance drops to ~2.5 ohms, peak current tops 35A, load voltage is ~355V and load current is 3.5A. 1200W 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.

Ref. AN39



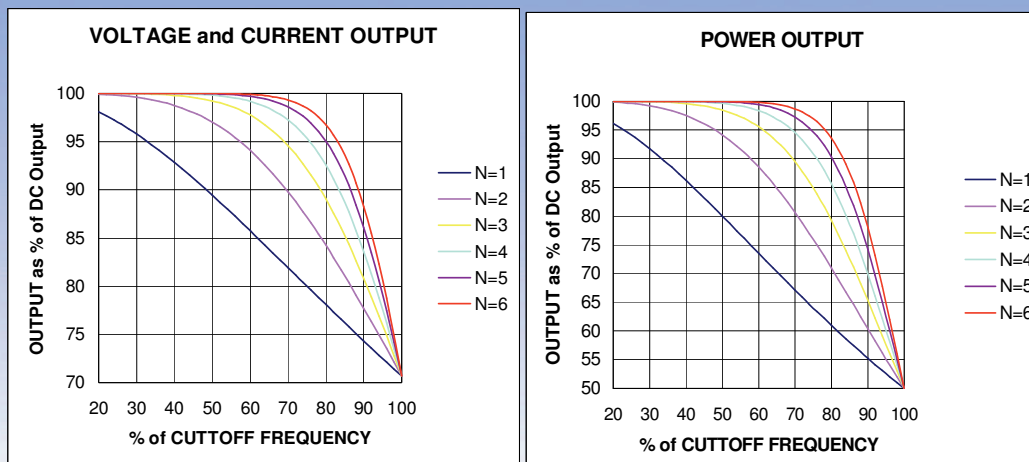
With proper termination of the filter we get a little mid-band peaking amplifier output current but the catastrophic potential of bad filter termination has gone away.

This filter design technique assumes amplifier output impedance is low compared to the load impedance and that the combined impedance of the load plus matching network is constant over frequency. The demand for circuit efficiency will insure the impedance relationship requirement is met. Beware that changing load element values, without corresponding matching network value changes, will alter the filter response curve. With some loads, such as solenoids or valves that tend to change inductance with position, the textbook response curve is nearly impossible to achieve. In these cases, try designing for the highest impedance, and then check performance driving the lower impedance.

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 is the voltage or current ratio and .5 is the power ratio. Many times the half power at maximum frequency is not acceptable.

Ref. AN32,AN39


Passband Attenuation



Designing the cutoff frequency at twice the actual maximum signal frequency is a very common technique to obtain a flatter response in the portion of the pass band actually used. You can see that in cases where amplitude flatness is critical, higher order filters and a wider ratio between actual signal frequency and F_c both help.

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. The question becomes one of cost in terms of money, extra loss in the filter, size and weight.

Ref. AN32

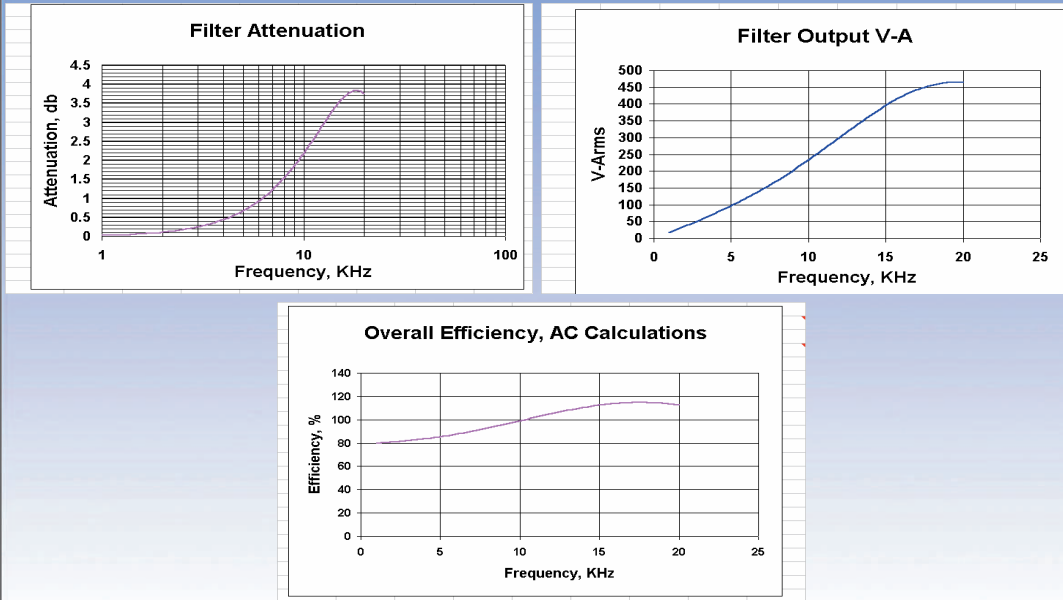


While the conjugate matching network performs almost like magic in terms of forcing the attenuation graph to near text book shape, there is a cost involved. This cost is slight when the load is mostly resistive but power dissipated in this network approaches power delivered to the load as the load approaches pure reactance.

These graphs are for an application driving a 1uF piezo stack with 12 ohms series resistance, to 75V peak from 1KHz to 20KHz. The filter cutoff frequency was designed for 40KHz providing 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. As this power is not delivered to the load, efficiency is far from the desired level.


Ref. AN32,AN39

Without a Matching Network



With no matching network we cannot lose any power there, but this leaves the filter with an improper termination. The result is a resonant circuit causing almost 4db peaking. In terms of V-A in the load near the upper end of the band, power goes from ~180 to over 450V-A. The efficiency graph looks like a patent should be applied for. The reason for this is recirculating currents in our newly formed resonant circuit.

Ref. AN32,AN39

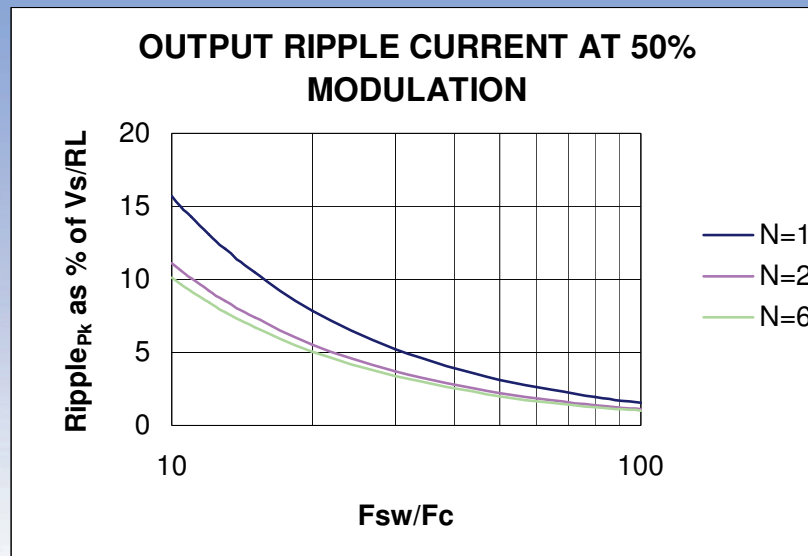


Here lies part of the beauty of the Power Design spreadsheet; it took more time to prepare this slide than it did to discover that doubling the resistor value in the matching network may provide a workable compromise.

Peaking at the load is down substantially from not using any network and wasted power is down substantially from using the ideal network.

Ref. AN32,AN39

Minimum Inductance




Application Note 30 admonishes us to make sure the PWM amplifier is driving enough inductance to keep ripple current at the switching frequency to a reasonable level. When designing filters according to Application Note 32, this concern becomes part of the filter design.

A full bridge PWM amplifier driving a first order (single pole) filter with F_c set at $1/10$, the switching frequency will be required to deliver approximately 15% of the peak output current as peak ripple current. The ripple is at the switching frequency; measured when the modulation level is 50%; and assumes peak output current equals V_s/RL . Changing to a second or higher order filter will drop this to almost 10%. A second and even more effective way to reduce this ripple current is to widen the ratio between signal and switching frequencies. As switching frequencies of Apex PWM amplifiers range from 22.5KHz to 500KHz, this technique has obvious limits.

This ripple current flows through the first inductor of the filter, meaning high frequency core loss is of concern. With first order filters driving resistive loads, it also flows through the load. With higher order filters, most of the ripple current flows in the first filter capacitor, affecting the ripple capacity rating of these components.

Ref. AN32

Ripple Varies with Modulation




In applications where full modulation is expected (output current is expected to approach V_s/RI), the workload imposed on the amplifier by delivering the ripple current is of minor concern. While 15% (or less as) of maximum output may seem more than minor, this ripple current decreases as modulation percentage moves away from 50% (a graph of zero to 50% would produce a mirror image curve). In other words, heatsink size is not increased 15% because maximum DC output and maximum ripple output never occur at the same time. The heatsink will be sized to handle the much larger output current. The ripple current curve is also valid for half bridge circuits, but the V_{out} curve would need to be re-scaled from 0.5 at 50% modulation to 1 at 100%.

For applications spending a major portion of the time near the 50% modulation level, the ripple current will be quite noticeable in terms of lowered efficiency (power supply loading and heatsink temperature). These circuits include full bridges spending most of their time delivering small signals compared to peak output capability; full bridges whose peak output voltage is considerably less than supply voltage; and half bridges spending most of their time delivering half the supply voltage.

Ref. AN32

Load current ≠ Amplifier Current




When using second and higher order filters, impedance presented to the PWM amplifier will dip below the load impedance as signal frequency approaches F_c . The graph above shows this in reciprocal form. Putting some numbers to go with the worst point: $N=6$, $F_c=1\text{KHz}$, $F_{\text{signal}}=900\text{Hz}$, $I_{\text{load}}=10\text{A}$, amplifier output=12.3A. This “extra” current flows in the output devices of the PWM amplifier increasing internal power, increasing ON resistance, increasing junction temperatures and reducing efficiency. This effect should be considered also with regard to amplifier and power supply current ratings and design of current limit circuits. We will see what looks almost like a duplicate of this graph when discussing filter component stress levels.

Again, this graph shows the advantages of lower order filters and wider ratios between actual signal frequency and F_c .

Ref. AN32,AN39


Speed & Order vs Efficiency



This is efficiency data for a perfect component filter (no parasitics) designed for an SA03 running maximum output voltage into a 10Ω load while mounted on a 0.1°C/W heatsink. At 10% of F_c , about 3.3% is lost in the amplifier and the filter is having very little affect on efficiency. As signal frequency increases, three effects combine to bring high frequency efficiency down further. First, quiescent power remains constant even though the output signal is rolled off. Secondly, the peaking output current demanded by second and higher order filters increases internal PWM losses. The last item is the positive non-linear temperature coefficient of the ON resistance of the PWM, which increased about 1% in this example.

The point here is that filter choices can double efficiency loss even before allowing for filter component loss. Importance of this data varies with the spectral content of the signal being amplified. Consider an audio application versus a fixed 400Hz inverter application.

Ref. AN32



Multi-pole filters are a combination of one or more series resonant circuits and they do develop currents and voltages above the input and output levels as the signal frequency approaches the cutoff frequency. The highest stress levels will be born by L1 and C1. Higher order filters produce higher amplification levels. The last two components of the filter do not see stress levels above the signal level. In these graphs, voltages and currents are normalized to the DC or very low frequency output signal amplitude and are based on ideal components.

Data on current can be used directly for any filter topology for both inductors and capacitors. If a split inductor topology is used, the inductor voltage data must be divided by two. Voltage data can be used directly for capacitors not connected to ground. Ground terminated capacitors have a DC bias equal to 1/2 the supply voltage which must be added to half the peak voltage calculated from the graphs. Do this calculation for BOTH the positive and negative peak output voltages. Note that if output voltage is nearly equal to supply voltage, and the filter order is three or more, the most negative going peak for C1 will be negative with respect to ground. The same is true for C2 with fifth and sixth order filters. This means even a ground-terminated capacitor can have BIPOLAR voltages applied. From a practical point of view, this situation implies the use of unipolar capacitors limits filter order to two.

As an example, consider filter options for an SA06, which is to deliver ±470V to a 332Ω resistive load at 1KHz. Current will be 1.414A peak or 1A RMS. Power will be 665W peak or 332Wrms. A supply of 480V will provide plenty of headroom for internal losses and maximum linear duty cycle limitations. The worst case for voltage and current extremes will be a sixth order filter.

$$L1 \text{ peak current} = 1.414A * \sim 1.23 = 1.75A$$

$$L1 \text{ peak voltage} = 470V * \sim 1.82 = 850V \quad 425V \text{ each if dual}$$

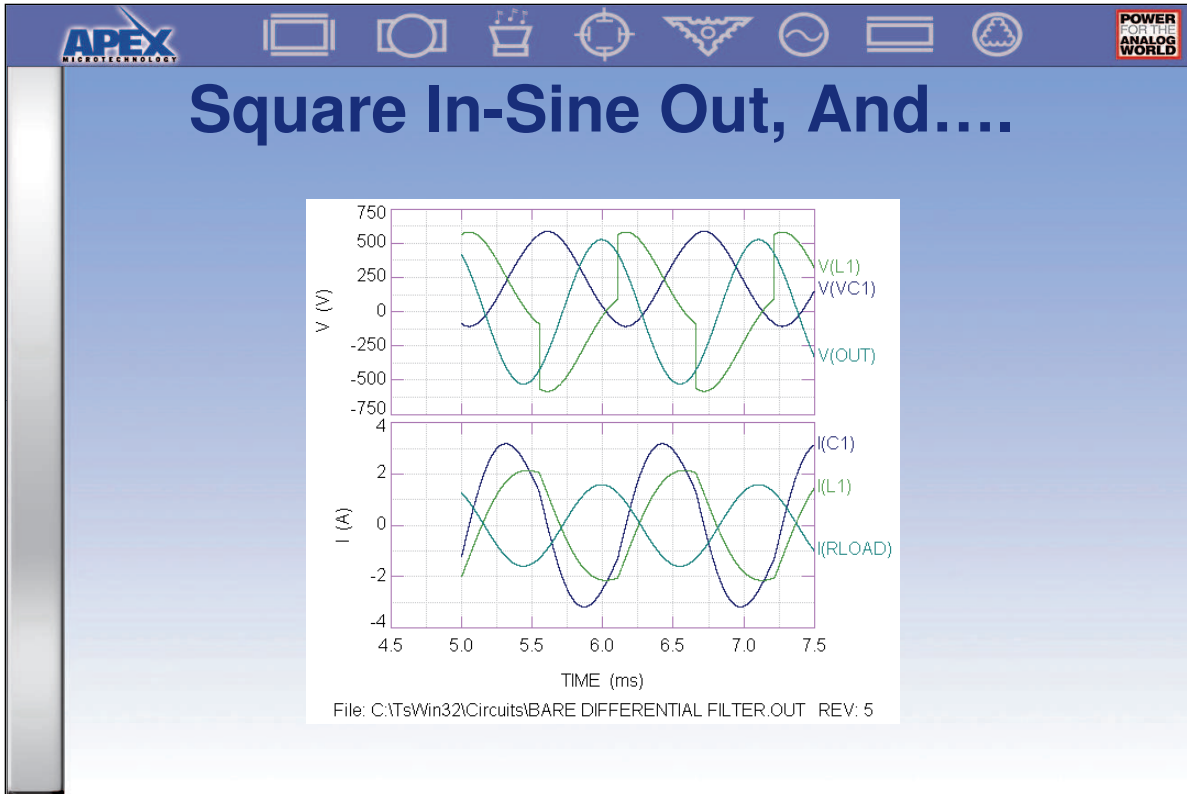
$$C1 \text{ RMS ripple current} = 1A * \sim 1.82 = 1.82A$$

$$C1 \text{ peak voltage (differential)} = 470V * \sim 1.17 = 548V$$

$$C1 + \text{ peak voltage (grounded)} = 240V + 274V = 514V$$

$$C1 - \text{ peak voltage (grounded)} = 240V - 274V = -34V \quad \text{Must be bipolar}$$

These stress levels are normal, even though the output ratings of the circuit are only 470V peak and 1Arms and the filter is properly designed and terminated. Before we go to the next slide, note that the input signal for this circuit is a sine wave. Ref. AN32




This is a Spice simulation of the previous example showing L1 and C1 stresses when the input signal is a 900Hz, 470V square wave instead of a sine. The modeled filter topology was a dual capacitor design.

L1 voltage = $\pm 582\text{V}$ and is for $\frac{1}{2}$ the total inductance (a single-ended design would place $\pm 1164\text{V}$ across the inductor). L1 current peaks at $\pm 2.14\text{A}$. C1 current peaks at $\pm 3.18\text{A}$. C1 is grounded and has voltage peaks of 587V and -107V . Watch out with that electrolytic capacitor! The output is a very good looking sine wave instead of a square, and peak output amplitudes have risen from 470V to 527V , from 1.414A to 1.59A and from 665W to 838W .

Points to consider:

1. Other than this slide, all input signals have been sine waves.
2. Input waveforms other than sine, can produce stress levels even higher than Power Design predicts.
3. As signals approach F_c , filters REALLY like to output sine waves.
4. If you really do need constant frequency sine waves with peak amplitude higher than the supply voltage, this is a possible circuit.

Ref. AN32



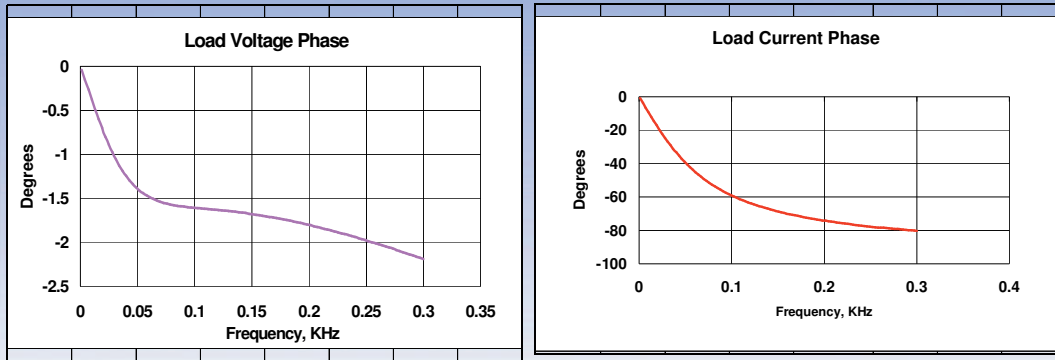
Power Design calculates voltage and current stress levels on L1 and L2, plus C1 and C2 for all designs. Resonance of these filters can produce voltages and currents larger than the load levels. Button 84 will place the first graph on the screen, then scroll up and to the right to view other graphs. The currents shown here can be used directly for all filter topologies. If L1 is actually two inductors, half the voltage shown will be across each individual inductor.

This circuit example only has a 90V supply; the drive signal is only 85Vpk; the load resistance has risen to 15Ω even though the filter design was for a 10Ω load. We might initially expect the 85Vpk signal and the 15Ω to limit inductor current to about 5.7A, but L1 has current peaks of 10.1A and voltage peaks of 108V. These peak values are pointed out on the right.

For capacitors, peak voltage is calculated for both differential and grounded capacitors. If a ground capacitor would experience a negative voltage, The red flag pops up.

Ref. AN32

Filters in Servo Loops



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 a servo loops, feedback is taken after the filter and any phase shift introduced here affects system phase margin. Power Design calculates both voltage and current phase in the load.


Voltage phase shift through a properly designed and terminated filter will be 45° per pole at F_c . This phase shift is reduced as the ratio between F_{max} and F_{cutoff} frequencies widens.

These graphs are from a 5mH, 2Ω magnetic bearing application featuring current output, third order filter with a cutoff frequency of 3KHz, and a modified matching network.

Ref. AN32

Keep it Cool!

$$R_{on} = f(t_j)$$



The “on” resistance of a power MOSFET increases about two times as junction temperature rises from +25 °C to +150 °C. This means a larger heatsink increases both output capability and efficiency. If there’s good news to this story it’s the non-linearity of the curve: The first few degrees we lower temperatures buys the most. Here’s a way to approach the problem.

First order power dissipation in the PWM is a function of the output current and the voltage drop at that current. This is the PWM advantage over linear power delivery; supply voltage is not part of the equation. Start with the 60 °C curve (interpolate if required). Find your current (PEAK if below 60Hz, otherwise RMS) and read the voltage drop. The product is power dissipation. The voltage drop divided by supply voltage approximates efficiency (quiescent current of both Vcc and Vs will reduce this a little). The heatsink rating is 60 °C minus ambient temperature, divided by power.

Are these numbers all affordable? Remember that a bigger heatsink actually reduces the watts to be dissipated (unlike linear systems).

Ref. AN32,AN39

PWM Internal Power Dissipation

- $P = I^2 * R$
 - I = switch current, R = switch on resistance
 - Vs and Vo are NOT in the equation!
- Both terms are moving targets
 - $\Delta R \cong 2:1$ from $T_j = 25$ to 150°C
 - Reactive elements cause I to move vs. frequency
- Heatsink rating affects efficiency

First order approximation of internal power dissipation looks deceptively simple. The advantage of PWM power delivery is that supply and output voltages are NOT part of the equation. This means that the power supply can be substantially higher than the output voltage without a significant penalty in efficiency.

The fact that both terms of the prime equation move around makes the calculation task more complex than with linear amplifiers. Again, this becomes a good job for a spreadsheet.

Just as with linear amplifiers, the heatsink must keep the semiconductors cool enough to provide a reliable circuit. In addition, larger heatsinks on PWM amplifiers actually increase efficiency.

Calculating Power Dissipation for Apex PWM amplifiers

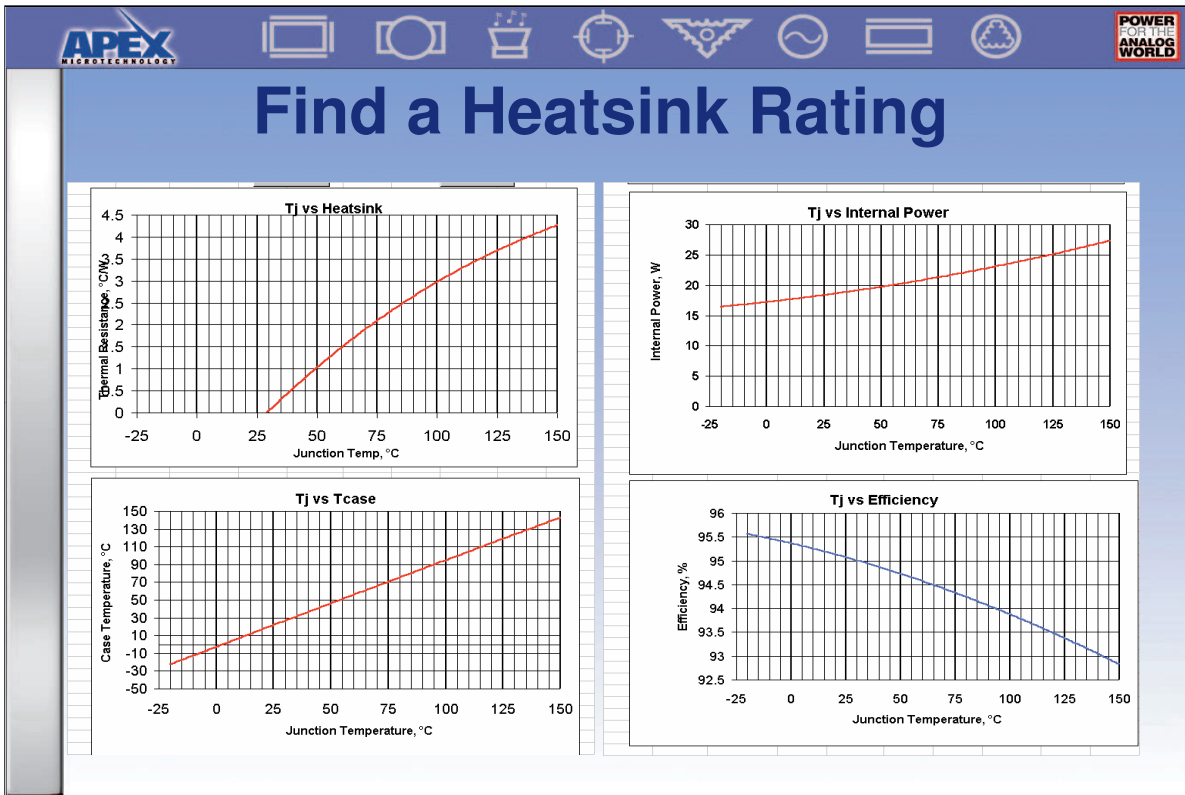
Model	SA03	Read Me	Ta max =	25	Tj max =	150	Tc max =	85
Power for Sine Wave Outputs			Fc(KHz) =	4	N =	2		
Vs	90	Volts	Fswitch =	22.5	KHz		ripple =	0.932806782
Fmin	0.001	KHz	Iq Vcc =	0.08	A		Iout =	8.500476297
Fmax	2	KHz	Piq Vcc =	1.2	W		Ifet =	8.551504299
Sig	85	Units	Iq Vs =	0.066795	A		Fhotspot =	0.0599
Sig as ?	V peak	Note/W	Piq Vs =	6.01155	W		Minimum Heatsink:	2.54225414
			Max delta Tj =	125			Max delta Tc =	60
General Procedure:	First Filter Design		Tune the filter		Select the heatsink			
Actual HS:	1	°C/W					R PchFET	0
Approx. Power Out	361.2499818	Wrms					R NchFET	0.12956009
Estimated Internal Pwr	20.3424608	W					Rwire	0.05
Estimated Case temp.	46.0680827	°C	Read me	82	View Overall Efficiency		Rtotal	0.17956009
Est. Junction temp	50	°C	Read me				D.C.max	99
Efficiency @ Fhotspot	94.66906088	% (AC Amplifier Only)	NA		% (DC Amplifier Only)		Vpk out	85
Est. Vpk capability	87.52934768	V					Ipk out	8.747223943
Notes:								
Filter:	RI= 10, CI= 9999999999999999, LI= 0, CM= 9999999999999999, Lm= 0, Rm= 9999999999999999							
37 Define Load	Sweep the Frequency	65 View Amp Out Last Sweep	84 View Filter component Graphs	71 View Load Last Sweep	Page down for charts.			
68 Print Data								

Did someone complain about lack of details on the previous page? Here they are, and the inputs were transferred from the PWM Filter sheet. If you change a green cell value, blue cell answers will not be valid until you run a frequency sweep.

If you get errors when you do this at home, check the READ MEs. You need the Analysis Toolpak add-in. Now you can see in the upper half, quiescent powers calculated, plus output current, FET current, hotspot frequency and best of all, minimum heatsink.

A little lower, notice I have already input an acceptable heatsink value and operating points have been calculated. Please read the comments. The Power Output assumes a properly terminated, zero loss filter, and a power factor of 1 in the load. Use button 82 to see efficiency including filter losses. If you enter too small a heatsink, most of these answers will be forced to ridiculously large numbers and a red TOO HOT warning will appear.

Ref. AN39



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

On the top right we see that internal power dissipation of the amplifier changes with junction temperature - - or with the size of the heatsink. Below we see this same effect expressed in terms of efficiency. This is a relatively low power PWM application. With higher power applications the percentage point change shown on this graph will increase.

Ref. AN39

Selecting an Apex Heatsink See ACCESSORIES INFORMATION data sheet for specifications

Thermal Resistance	Package	Velocity Calculator:				Units of Measure:
2.3 °C/W	MO127	100 CFM	9.5	Inch Width	English	
	Update heatsink List	10 Inch Dia	12	Inch Length		
READ ME		183.3465 Ft/min	126.32	Ft/min		
		0.931412 M/sec	0.6417	M/sec		

Notes: **Beware: Flow rates change as you enter Thermal Resistance, but may be wrong until the Command Button is used!**

Model	Fluid	Thermal resistance, free air, °C/W	Your rating requires FPM or GPM flow	Package(s) accepted	Style	Length, inches or cm	Width, inches or cm	Height, inches or cm	Weight, ounces or grams	Singles Price USD Domestic
HS06	Air	0.96	0	MO127		4.5	6.25	2	19.8	\$42.25
HS11	Air	0.68	0	TO-3, MO127		6	8	2	44.8	\$214.80
HS11	H2O	0.68	0	TO-3, MO127		6	8	2	44.8	\$214.80
HS18	Air	1	0	MO127		5.5	4.612	1.5	14.1	\$80.95

Believe me, heatsinking is NOT the easiest science in our universe.

Let's start with "the" heatsink rating. The HS03 is rated at 1.7°C/W in free air. True, when power dissipation is about 45W, but check the actual curve at 10W and you'll find a rating more like 2.3°C/W. On top of that, "free air" means no obstructions to air flow and the flat mounting surface must be in the vertical plane. Demands for higher performance in smaller packages can be at odds with optimum heatsinking. Poor installation choices can easily reduce effectiveness 50%.

Moving on to this selector software. Air velocity curves from the heatsink data sheet (when available) have been approximated with polynomial expressions. While these errors are minor compared to the previous paragraph, it would be good to allow 10% for velocity ratings over 150 feet per minute and 20% below that.

Adding a fan to your design enables you to use smaller heatsinks. Please remember: Most fans are rated in cubic delivery and this rating varies with working pressure. A 5 inch diameter fan delivering 100 CFM produces over 700 FPM right at the fan. If this air is flowing through a 19 x 24 inch rack, theoretical velocity is down to 32 FPM, will vary with location and goes lower as the rack is sealed tighter.

The bottom line: Without case temperature measurements, your design effort is NOT complete!