



VENABLE TECHNICAL PAPER # 9

"Oh, By The Way — We Need A Power Supply"

Abstract:

Integrating a power supply into a new system is often left until late in the design cycle. The problems encountered in choosing and mating a power supply with a nearly completed system can seem overwhelming, particularly with the constraints imposed by overall system requirements and development schedules. This paper deals with the problems encountered and how vendors and users can effectively avoid and/or overcome these "last minute" obstacles.

Introduction

"Gee, Steve, the new computer is almost done. That black magnesium cube is totally awesome! There's just one thing left — the power supply. I think we can find some room behind the disk drives and over the motherboard. We just have to wrap the supply around the memory back-up batteries and the serial port connector somehow. I'm sure there's a power supply manufacturer somewhere that can do it!"

Sound familiar? Whether you are a manufacturer or user, you are bound to have run into a similar situation. "And oh yes, the supply has to put out 750 watts, have 5 outputs, occupy no more than 30 cubic inches, be 98% efficient, and pull perfectly sinusoidal current from the line. By the way, the project is two years late, so we need the supply ASAP. And we're already way over budget, so don't bother asking for any NRE money."

Hopefully, by now most people know that a power supply is more than just a transformer, rectifier, and three-terminal regulator. Those of you designing power supplies know that for sure. The question is, what things do users and vendors need to know and do to avoid surprises and problems when the power supply and using equipment meet for the first time?

Size, Heat, Time, And Money

We touched on some typical power supply problems in the first paragraphs. There seems to be no limit to the requests for smaller, cooler, faster, cheaper power supplies. Some government agencies have tried to put limits on the problem. For example, the Navy, through the Willoughby committee, issued guidelines for Navy power supply procurement. The guidelines called for things like 2 watts per cubic inch, reasonable development schedules, and 10% of the program cost devoted to power supply procurement.

We have consulted on several power supply designs destined for use by the Navy, and so far we have not noticed anyone following the guidelines. The Navy seems to want the same high density and low price that everyone else wants when they buy power supplies. Military designs do not have as tight a restriction on schedule and budget as commercial designs, but any military power supply manufacturer will tell you that costs and schedule are important.

It is important though, especially as the size of power supplies gets smaller, to plan ahead as to how the heat is going to get out. I saw one of the new 1 MHz class of power supplies being used at a Boston area computer company, and the heat sink on it was twice as big as the power supply. When vendors calculate watts per cubic inch, they usually do not count the size of the heat sink. They figure that the heat sink is the user's problem, so be aware.

Know The Load Current

Before I founded Venable Industries, I designed high voltage power supplies for traveling wave tubes at Hughes Aircraft Company. We worked closely with the "tube people" during the development cycle, trying to design the power supply and TWT in parallel, so that when the tube was finished the power supply was waiting for it. One of the things that used to amaze me was how far from reality the tube people's guesses were concerning tube operating voltages. After some experience I learned to design for a 30% variation around the stated voltage goal, and most of the time I could adjust the supply to meet the requirements of the tube — but not always and sometimes just barely. I suspect the same is true of computer and other electronic systems, except the voltages are set and the load current is unknown until the system is virtually complete.

Know The Load Capacitance

At Venable, we consult on power supply design problems. We also manufacture frequency response testing equipment. The test equipment is used to stabilize feedback loops, do conducted susceptibility testing, and to measure input and output impedance of power supplies.

Another lesser-known but important application of frequency response analysis equipment is to characterize load impedance by taking data on actual using equipment. This is one way that users can avoid surprises when the vendor delivers the power supply. By measuring and characterizing the actual load impedance and informing the vendor so that his load bank can be modified to mimic the using equipment, the power supply can be designed and tested properly the first time.

Load capacitance is especially important if the supply has remote sense and you plan to use this feature. Phase shift due to inductance of the power cables and capacitance of the load frequently causes power supply feedback loops to oscillate, even though the supplies appeared perfectly stable with the purely resistive load and local sense. The problem is worse with longer cables and higher power levels.

Specify Predictable Designs (If Possible)

If you're a power supply user, you may not have much say in the inner workings of the design. If you do, though, ask for as much predictability as you can get. There are three main culprits when it comes to unpredictability:

Transconductance Amplifiers

The first is transconductance amplifiers used with the compensation connected from the compensation pin to ground. Many people say, "Dean Venable doesn't like transconductance amplifiers." That isn't true. It's not a matter of liking or disliking; it's a matter of predictability in the design. With a real op-amp, or even with a transconductance op-amp in many cases, connecting the compensation around the op-amp (from output to the inverting input) makes the gain a function of the ratio of resistors and capacitors. In a transconductance amplifier, connecting the compensation from the compensation pin to ground makes the gain a function of the transconductance of the error amplifier. Transconductance varies from chip to chip, and as much as 3 to 1 with different manufacturers. After a design enters production, the purchasing department will make the decisions about PWM chip manufacturers — usually based on price. It is not uncommon to have a previously stable loop oscillate because a different vendor was selected for the PWM chips at some point in the manufacturing cycle.

Opto-Couplers

Opto-couplers are a second common source of unpredictability. Because the outputs of a power supply are almost always electrically isolated from the input, some way has to be found to couple a signal from the output side (where the error amplifier is) to the input side (where the PWM chip is). Opto-couplers are frequently used to do this. Since the whole idea is to have the input and output isolated, there is no way to use feedback to stabilize the gain of the opto-coupler. The gain of this device also varies from chip to chip, and varies significantly from manufacturer to manufacturer. Better designs use transformers to couple signals across the isolation barrier.

Current reset in mag-amps

Magnetic amplifier post-regulators using current reset are a third source of unpredictability. Since there is no direct relationship between current and flux in a mag-amp, there is no sure way to predict the gain. Flux is proportional to volt-seconds. If voltage reset were used, the results would be predictable and would not vary with the particular core being used. With current reset, voltage is generated by driving current through the parasitic core loss of the magnetic core. Since core loss is related to many unpredictable things (such as whether or not the core was ever dropped), accurate analysis of the performance is virtually impossible. There's not much you can do to eliminate this unpredictability, though. Almost every vendor uses current reset because it's simpler and easier.

Input Voltage Range

A recent advance in power supply design is automatic switching of the input voltage, either by automatically emulating the jumper move required on most supplies now or by making the supply able to tolerate a very wide input voltage swing. This is an encouraging trend, and if you plan to market or use equipment in both Europe and North America, automatic switching is a feature to ask for. As an example, we make welding equipment, which is able to work on 115, 230, or 460 volts, with automatic switching. This feature has prevented an accidental "Oops!" on several occasions.

Efficiency And Power Factor

Efficiency has always been of some concern because of the problem of getting the heat out of the power supply and/or the unit it is mounted in. As customers demand more and more features in their equipment, the using equipment needs to pull more and more power from the line to deliver these features.

A limitation, at least in the United States, is that wall sockets are generally designed to supply only 15 amps at 120 volts. At first, as equipment power input climbed near the kilowatt level, improvements in efficiency solved the problem of blown circuit breakers. Now, with efficiencies about as high as practically possible, improvement can come only with improved power factor. Some new power supply designs have a "front-end" regulator, whose main purpose is to draw current from the line in proportion to the instantaneous line voltage, thereby simulating a resistor and a perfect power factor. Older designs pull power from the line in current "spikes" or "peaks," which occur at the peak of the input sine wave. Essentially no current is drawn during the remainder of the input voltage cycle. In some office buildings, especially those with a lot of electronic equipment, this has resulted in a significant "flattening" of the peak of the input voltage sine wave. It is quite probable that electric utilities will soon monitor this phenomenon and either legislate against it or charge extra for it. If your power requirements are in the 1000 to 1500 watt range, power factor is something you should consider specifying.

Hold-Up Time

If momentary power outages are a serious problem, you should consider an uninterruptible power supply (UPS). This can be built-in to a new custom power supply design, or a stand-alone UPS can be used in selected locations to power a normal supply. There are two types of UPS systems, standby and on-line (continuous). Standby systems sense dropout of the power line and switch a relay to change the source of power from the ac lines to the battery-powered inverter. Since this sensing and switching requires a finite amount of time, hold-up in the power supply is important. "Hold-up" has nothing to do with Bonnie and Clyde, it simply is a measure of how long the input filter capacitor will continue to keep the power supply outputs in regulation after the input line voltage has gone away. This has to be longer than the UPS switching time.

Continuous UPS systems process all the power all the time. They do not have to switch, so hold-up in the power supply is not important. Continuous UPS systems also remove line transients, brownouts, and similar phenomenon, which are harmful to the health of power supplies, and/or the equipment they are intended to drive. As you might have guessed, continuous UPS systems not only do more they cost more, so economics will have to be factored into your final decision.

Electromagnetic Interference (EMI)

Agencies in many countries are placing more restrictions on the amount of conducted and radiated emissions a power supply can put onto the power lines and into the air. The advent of high-speed power transistors, particularly FET's, have enabled power supplies to put out radiated energy well into the radio frequency bands. If you have control of the design, try to have the manufacturer limit the switching speed of the transistors somewhat, for example with a 50 ohm resistor in series with the gate lead of the power FET. Often, with a very minor effect on efficiency, a great deal of EMI can be eliminated. If you have no control over the design, make sure that the various applicable agency approvals have been granted, in particular the German agency VDE that is the most stringent.

Type of Current Limiting

Something you might not think of in the beginning, until it is too late, is what you want to happen in the event of a short circuit on one of the power supply outputs. Of course you don't want the supply to fail, but there are other things to consider. Some people want the outputs to limit at some amount slightly over the rated output current of the supply, and to become constant-current sources after that. The theory here is that if a problem occurs, it will be easy to find — just look where the smoke is coming from.

Another approach, which is probably the most popular, is for the supply to "motorboat" or shut down momentarily and then try to restart. This approach normally will not generate an excessive amount of smoke during a fault, and will not shut the system down permanently should some random fault occur. A third type of current limiting causes the supply to shut down and remain off until some action is taken by the user, typically turning the ac power off and then on again. This is the best approach if the periodic attempts to re-start may do more damage, or if re-applying power soon after a shut-down would cause system problems, such a loss of memory in a disk drive. Personal computer power supplies typically use this type of current limiting.

Excessive Specifications

If money is a consideration, you should avoid over-specifying the power supply. This typically takes the form of specifications that are too tight for regulation and/or ripple, far beyond what the using equipment needs. I remember an actual experience once where a supply was specified to have essentially zero output variation. When the question "Why?" was asked, it turned out that a 5% variation was OK for the using equipment, but that the voltage drop in the wiring used up the entire 5%, so there was none left for the power supply. Modification of the cable specs solved the problem in a more economical manner.

Low ripple specifications can actually be a detriment, since they may require the use of low ESR (equivalent series resistance) output filter capacitors. Not only are low ESR capacitors more expensive, in many cases the stability of the feedback loop is dependent on that resistance. Substituting low ESR capacitors could reduce the phase margin to the point where the load transient response is severely degraded, even to the point of loop oscillation.

Feedback Loop Stability

Last but not least, we come to my favorite subject: feedback loop stability. I recommend measuring the phase margin and, if possible, the gain margin of all feedback loops within a power supply. One reason for doing this is that almost anything that affects the normal operation of a power supply affects the Bode plot. (For those of you who may not remember, a Bode plot is a plot of open loop gain and phase shift versus frequency for the control loop of the power supply.) This makes a Bode plot an excellent quality assurance tool, often detecting defective parts that slip through any other screening or test.

Step load testing shows up some loop problems, but conditionally stable systems have good step load response. Conditional stability only shows up in a Bode plot.

Bode plots give a quantitative measure of the margin of stability of a feedback loop. A series of plots at various line, load, and temperature conditions often points out weaknesses in a design, which can then be corrected before a field problem occurs at a customer location.

Summary

Several of the things I mentioned in this paper are, of course, out of the hands of the power supply user. They were included in part to help the user gain a better understanding of the problems faced by the power supply designer, and the need to give vendors the best possible information of system power requirements and the most realistic possible specifications. Let me re-emphasize a couple of points made earlier:

Complete information on load capacitance, type of sensing scheme, hold-up time requirements, etc. will greatly help the power supply vendor meet your needs. This type of information is often overlooked until the supply is designed and found not to work. Feedback loop stability analysis is an excellent - and under used - method of determining and verifying power supply performance. We are involved with a number of customers who use Bode plots not only for design, but also to assure product performance during production, or as a product acceptance tool during incoming inspection (we, of course, would be delighted to talk with you individually about what types of equipment and services you need to implement such programs).

In summary, I have tried to point out some of the power supply integration problems that I have run into over the years, and how to avoid or at least watch out for them. I hope you found this helpful.