(DOs, DON'Ts, and Common Pitfalls)

By Lazar Rozenblat

A CompactPCI system integrator may view the power supply as a "black box" that just has to fill the allotted space when the whole system is already designed. This "box" is expected to power the system, fit the budget, and otherwise be "invisible". Unfortunately, it is not.

Here is a possible scenario: *a supply selected from the catalogs seems to satisfy all the system's requirements. One gets the samples just to be sure that they fit the rack. They do fit. But when powered, the system does not always start up, the fault LED blinks, insertion of a second supply resets CPU, and finally the power supply shuts down at noon when room temperature raises a few degrees. For all that, the manufacturer maintains that the supply meets the specification.*

In this article, Lazar provides an analysis of CompactPCI power supplies that would help digitally oriented systems designers avoid last minute problems. It describes how to design a backplane to avoid power supplies' hot swap glitches, cooling of the supplies, how to select a CompactPCI supply for your application, and how to read between the lines of the power supply's specification.

System's start up and minimum load requirement

During the development or debugging of the system, an engineer often does not install all cards into the rack until he or she is confident that the power supply is OK. The most frustrating thing that one may see at this moment is a blinking *Fault* LED of the supply. The problem may be related to a minimum load requirement. A typical CompactPCI system uses a hot-swap controller on CPU boards that measures 5V and 3.3V, and would not connect the board until these voltages are within the required range. Until then, the system presents practically no load for +5V and +3.3V outputs. A supply, on the other hand, may need some minimum load (on one or more outputs) to keep the outputs within the regulation band. This condition appears as a *Catch22* and may cause the system start-up problem and overtime work for a system integrator.

To better understand the mechanism of minimum load requirements, let us consider a simplified block diagram of a typical multiple output supply with forward converter as shown in Figure 1. During conduction time t_{ON} of the inverter transistors (not shown on the drawing), a positive voltage Vsec applied to the secondary winding, diode D1 conducts, and current in the output choke Lout increases by

$\Delta I_+ = (Vsec - 2Vd - Vout) t_{ON} / Lout$, (1)

where **Vd-** voltage drop across the diodes.

When inverter transistors turn off, the choke continues to conduct through free-wheeling diode D2 and its current decreases at the rate:

$\Delta I_- / \Delta t = (Vout + 2Vd) / Lout$. (2)

If this current does not reach zero by end of the cycle, the mode of operation is called continues conduction. The condition for this mode is $\Delta I_- / 2 \le Iout$. Since in steady-stead mode $\Delta I_+ = \Delta I_-$, simple calculations yield:

$Vout = Vsec * D - Vd$, (3)

where $D = t_{ON} / (t_{ON} + t_{OFF})$- duty cycle.

We see from (1) that in continuous conduction mode, the relation between **Vout** and **D** does not directly include load current **Iout.** In steady-stead mode, the supply operates with practically fixed duty cycle, and only small variations of **D** are needed to accommodate load and line variations. The inverter provides specific volt-seconds to all secondary windings, allowing secondary controllers to regulate all other outputs and assuring optimal transient response.

The feedback loop of multiple output supply is normally closed at one output called

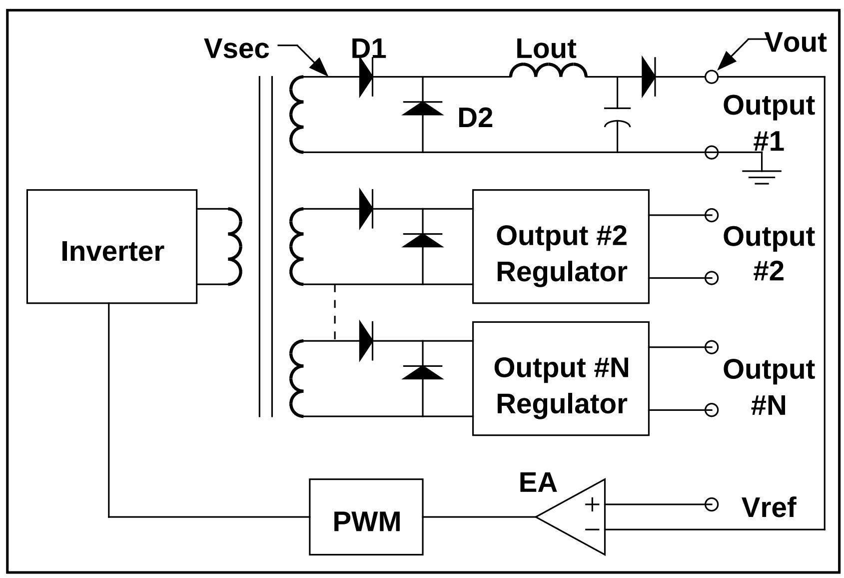

Figure 1

main (usually 5V). An error amplifier (EA on Figure 1) compares output voltage **V_{out}** with a reference **V_{REF}** and controls **D** via a pulse-width modulator (PWM). If **Iout < ΔI- /2,** the choke *dries out* before the beginning of the next cycle which causes discontinued mode of operation. In this mode, duty cycle becomes proportional to **Iout** and diminishes at no load. While PWM may still regulate main output, the inverter may not provide enough energy to regulate all other outputs and its transient response degrades.

Equations (2) and (3) yield minimum load requirement as a function of output inductance Lout:

Iout(min) > (Vout+2Vd) / 2FLout,

where **$F=1/(t_{ON}+t_{OFF})$** – operating frequency.

Placing high wattage *bleeding* resistors inside the supply reduces efficiency and is not always feasible (or even possible) in high-density supplies such as CompactPCI. Manufacturers may address the start up problem differently. For example, Todd's new 3Ux8HP supplies use *swinging* output chokes (with inductance increasing at light loads) and dynamic preload that allows the supply to start up at no load.

If a supply does require minimum start up load, the preload resistors maybe placed on the backplane. To determine the required preload, measure actual currents that your system draws until hot swap controller connects CPU. Ask the vendor to determine minimum load current required to power the system at such condition: this value may be less than minimum current specified in the supply's catalog (the latter is what is normally required to achieve full power on all outputs). In redundant systems, the pre-load will be needed for each supply.

Output power and proper cooling

The rated power and currents are always given within a specific ambient temperature range and usually at specific forced air-cooling. At that, the power supply's ambient temperature is not the temperature in the room where the system operates — it is the temperature INSIDE your rack.

Except for encapsulated supplies, the airflow should provide cooling of not only chassis and heatsinks, but also of supply's internal components including magnetics. International safety agencies check temperature rise on the supplies' parts (such as windings' insulation) at a specified airflow. The required airflow is indicated as one of the conditions of acceptance. Inadequate airflow may cause excessive temperature rise on electrical insulation—the failure of which can cause a hazard and is a violation of the condition of acceptance.

It is not easy to determine whether your system actually provides required airflow. In most cases, not all volume of the air driven by a fan passes through the supply because of different geometry of the CompactPCI supplies and conventional fans and because of turbulence and obstructions to air created by all the components in the enclosure. Therefore, the rating of your fan would not give you a direct answer.

Catalogs of power supplies usually specify airflow in terms of cubic feet per minute (CFM) or linear feet per minute (LFM). If the cross-sectional area of the supply that faces the airflow is Ac (in^2), then: **CFM=LFMxAc/144.** The first thing to do is to find out what the supply's specification really means and how the supply was tested for safety agenices' compliance. If a supply needs (say) 15 CFM, is it a rate of:

- air that just has to reach the supply
- air that passes through the area allotted for the supply when supply is removed
- air volume actually forced through the supply

If the spec calls for specific LFM, is it the rate of airflow in your enclosure outside the supply (that is between its chassis and adjacent cards), or is it an airflow through the inside of the supply (measured at its exhaust)? Design of your cooling system will depend on the answer to this question.

Linear airflow rate (LFM) in your enclosure can be measured with an anemometer. When you move an anemometer's probe, you will likely get different readings. Take an average value; it should not be less than that required by the supply's specification.

If performing such measurement is non-practical, you may try estimating the portion of the airflow that reaches the supply. For rough estimation, disregard the fact that airflow of a fan is distributed unevenly along its frontal area. Make geometrical projection of the power supply onto the fan (see Figure 2) and, by comparing cross sectional areas, estimate the volumetric rate **V** of airflow that reaches the supply as:

V≈CFM*As / Af, (2)

where CFM- rated airflow of your fan, As- area of the supply's projection onto the fan, **$Af=\pi D^2/4$** — area of the fan, D-diameter of the fan.

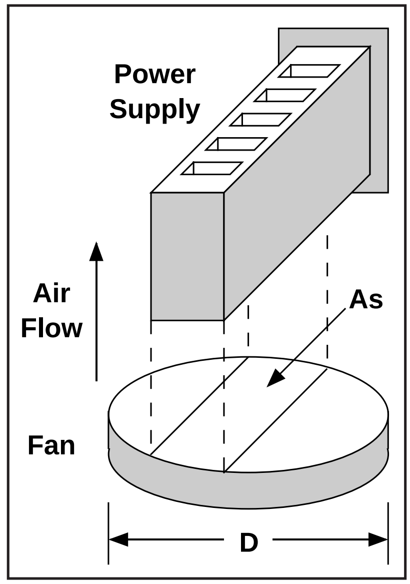

Figure 2

For example, if you use 120 mm fan for 8HP-wide supply, V obtained from (2) will be only about 40% of fan's rated CFM. If you use two parallel 60 mm fans, this portion will be about 80%. However, with two parallel fans, airflow does not double up because of mutual interaction. Of course, since power supply components block some airflow and a portion of air always reflects back, the actual volumetric airflow that passes through the supply will be several times less than the one that reaches the supply.

As a rule for the proper cooling of a CompactPCI supply, a fan with rated CFM several times larger than required by the supply's specification should be selected.

Suppose you cannot provide the needed airflow. Can you derate the power to live with less cooling? Yes, but usually only up to some point (ask the vendor for power derating curve). There are components (such as aluminum electrolytic capacitors, pre-load resistors, ICs, magnetic cores, bias supplies) whose power dissipation may not reduce with decreased power and, in some cases, even increases (i.e. components of magnetic amplifiers circuits). Operating at higher temperature on these components may also reduce MTBF of the supply.

Live insertion and removal

The CompactPCI specification requires live insertion and removal of the supplies without causing any systems glitches. The potential problems at insertion and removal may be of a different nature.

During insertion of the supply, high, inrush current flows through input power line. Inrush current is usually specified in catalogs as *cold turn on* which has little practical meaning. Most supplies use NTC thermistors to limit inrush current. After warm up within a few seconds, their resistance drops to less than an ohm. If you then retract the hot supply and quickly insert it again the inrush current may be many times more than specified in catalog. This condition (although unusual) may occur if the system looses the power for a short period of time. If you use a system fuse or a circuit breaker, they should not trip under such condition. You may ask the manufacturer about the maximum value of $\mathbf{I^2t}$ of the input current pulse. The correct answer will depend on the impedance of your input voltage source and its cables. If you want to estimate it yourself, run the system with all supplies installed under maximum load condition for a few minutes and then record *hot* inrush current with a current probe and oscilloscope. The value of $\mathbf{I^2t}$ will be approximately $\mathbf{0.5 * I_{PK}^2 * t_P}$, where $\mathbf{I_{PK}}$ – pick inrush current and $\mathbf{t_P}$ is the duration of the inrush current pulse. The obtained value should be at least twice less than the rated $\mathbf{I^2t}$ of your system's fuse or circuit breaker.

Also, be sure that input voltage sag due to the inrush current will not bring the input voltage below brownout threshold. It is especially important for DC-DC supplies, which may not have input rectifiers that isolate their input from the power line. For such supplies the input voltage on the first one may rapidly sag to the half of its original value upon insertion of a second supply (say, from 48V to 24V) which may cause output dropout. Don't use long input cables. Do place an electrolytic capacitor on backplane in close proximity to the inputs of your DC-DC supplies. The value of this capacitor should be several times larger than the net input capacitances of individual supply.

A potential problem that the system may see upon failure of one of the supplies is of a different nature. Let us examine Figure 3, the simplified block-diagram of redundant connection of two supplies, PS1 and PS2. On this diagram: ESR and ESL-equivalent series resistance and inductance of distributed backplane capacitance Cbp, L1 and L2: net inductance between the supplies' outputs and the load (such as CPU card) that includes distributed inductance of PC board traces and inductance of the connector's pins.

Under normal condition, each supply delivers 50% of the load, so $\mathbf{I_1 = I_2 = I_0/2}$. Suppose output capacitor Cout fails into short circuit in the supply PS1. We know that according to Faraday's Law, a current across inductance cannot be changed instantaneously. It will take certain amount of time until PS2 will raise the current in **Lout2** and **L2** to the level of the load current $\mathbf{I_O}$**.** Until **Lout2** and **L2** slew to the full load current, what supplies the remaining 50% of the load?

Initially, a portion of the load current still flows through L1 and will decay with the rate $\mathbf{dI_1/dt\_Vout/L1}$**.** The remaining portion will flow through capacitance Cbp and will start ramping up with the same rate causing additional voltage drop across ESR. The initial deviation of the voltage at the load will be:

$$\Delta Vout \approx ESL * dI_1/dt = Vout * ESL / L1 \quad (3)$$

We see from (3) that in our model initial deviation of the voltage does not depend on power supply, but is a function of backplane equivalent impedance: it depends on the ratio between ESL of Cbp and L1. The former, in this case, acts as *bad* inductance and the latter as *good* inductance. To calculate the exact relation, we would have to include stray capacitance and inductance between adjacent copper planes in our model, which makes analysis too complex and is beyond the scope of this article. However, as a rule of thumb, in order to keep the initial voltage deviation below say 1%, the ESL of distributed backplane capacitance should be many times less than inductance between each supply and CPU. We can see that, contrary to our intuition, some backplane inductance in this case is beneficial. For practical estimates, consider that surface-mount ceramic capacitors may have ESL less than 2 nH each and paralleling the capacitors respectively reduce the overall ESL of Cbp.

During subsequent process of the recovery, current through Cout2 increases that causes more output voltage excursion. Depending on the supply's response time, additional electrolytic or tantalum capacitors may be needed on the backplane to support the load during this time. These capacitors along with distributed ceramic capacitors will also supply fast load transients. A few hundred microfarads would be a typical value of the bulk backplane capacitance.

Backplane designers should be aware however that excessive capacitance may cause:

- Supply's overcurrent condition at start up and longer rise time
- Lower bandwidth of the feedback loop
- Instability

Work with the vendor's application team to determine the optimum capacitance.

Current sharing

For redundant operation, the supplies have to share the load currents. Methods of providing current sharing may be generally divided into two classes: *single-wire* methods and droop (or zero-wire) methods. The former may use a separate wire (in addition to power and return) that lets power supplies share information about their individual currents. That is why sometimes it is

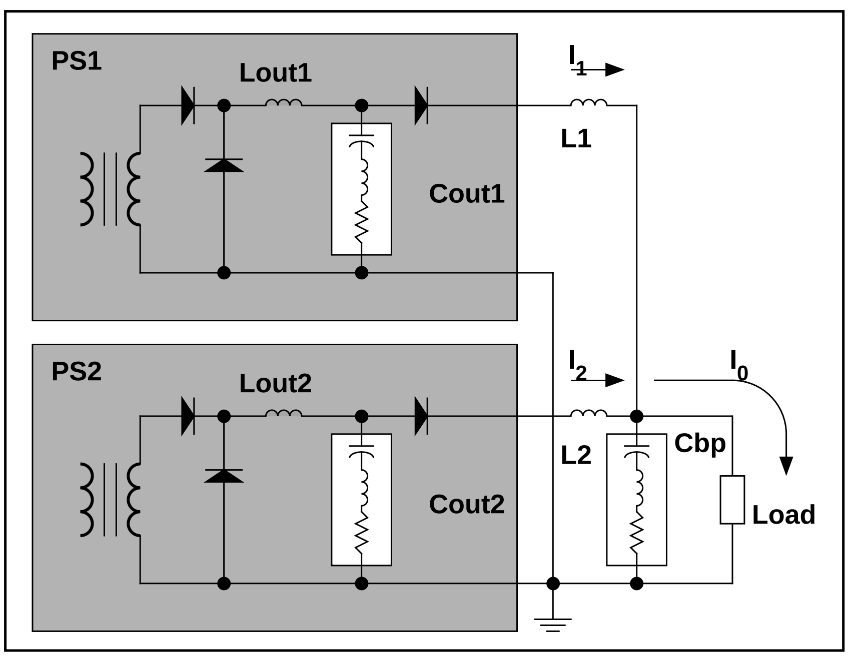

Figure 3

called *third-wire* current sharing. These methods assure the best voltage regulation but have some drawbacks. The supply, however, may be more expensive because of additional control circuit. Fault in the current sharing bus may introduce a single point system failure or drive the bus voltage out of tolerance. In addition, since CompactPCI Power interface specification PICMG 2.11 does not specify any particular scheme of load sharing, supplies from different manufacturers may not necessarily share the load.

Droop methods provide automatic current sharing without additional current share pins. When output current of one supply increases, the circuit will slightly reduce its output voltage to force the other supply to provide more current. These methods cannot introduce a single point failure (except for a short after an OR-ing diode) and use fewer components, thereby increasing system reliability and reducing its cost. Some droop methods provide poor regulation, but others (such as those that have predetermined linear slope) can feature fair regulation of ±1-2%. Although it's slightly worse than that of single-wire methods, it meets PICMG 2.11 requirements and is sufficient for most CPUs. Also, a supply with programmed slope current sharing will share the load with supplies that use *single-wire* methods (provided it's adjusted to the same nominal voltage at center load), which gives the system's integrator more flexibility.

Regardless of current sharing method, the supplies never share current precisely. This has to be accounted for when selecting the number of supplies needed to power the system. For example, if the system draws 50A and each supply can provide 25A, two supplies will not be enough. The minimum amount of supplies should be determined as:

$$N > (1+ \delta/100)*I_{LOAD} / I_{MAX},$$

where: I_{LOAD} – the system's current, I_{MAX} – maximum current of individual supply, δ — maximum current imbalance between two supplies in percentage. Of course, add one more supply for N+1 operation.

Compliance issues

A power supply is generally submitted to and approved by the safety agencies as a component. This does not guarantee compliance of the entire CompactPCI system especially if it uses multiple supplies. For example, an individual supply may meet UL1950 requirement on earth leakage current (<3.5 mA), but the N+1 system may not because it uses more than one supply, and also because its components (such as system EMI filter) or connections may provide additional leakage current. In this case, if the system's leakage current exceeds the 3.5 mA, but is less than 5% of the input current, a label stating: "High leakage current. Earth connection is essential before connecting supply" should be affixed. Additionally, protective system grounding should be provided per UL1950.

Since power supplies are not synchronized with each other, their operating frequencies will never be exactly the same. The EMI level, generally, will not directly add. Nevertheless, the system with multiple supplies may have higher EMI levels. In addition, supplies are tested with resistive loads that do not generate high frequency noise as digital circuitry does. This noise, through stray capacitances and via radiation, may affect the system's EMI level. Also, a new standard imposed by EN 55022:1998 requires additional measurement of conducted emissions on data cables. The latter is a function of the system and has to be accounted for by the CompactPCI system designer. The new standard will be in effect within the next two years and CompactPCI system designers may want to get an early start. EMI testing of a CompactPCI system under its actual working condition is crucial and, depending on the results additional system's EMI filter, may be required.

Conclusion

A successful CompactPCI system design, from its early stage, needs to take into account properties of the power supplies. Choosing the supply that best fits your system is the task that requires a basic understanding of the performances of the supplies, interfacing with a vendor's application team, careful testing at the system level, and is usually a trade-off between the features and the costs.

Lazar Rozenblat is a manager of R&D Group at Condor DC Power Supplies Inc. (Todd Products group). He has worked for more then nine years for Todd Products in Brentwood, New York as senior design engineer and R&D group manager. Prior to joining Todd, Lazar worked for Bertan Associates as an electrical engineer and for the Research Institute for Electronic Instrumentation (KNIIEP) in Kishinev, Moldova as a member of the research staff. Lazar holds The Engineer Degree (equivalent to MSEE) from the Leningrad Electrotechnical Institute for Electronic Communications (LEIS) and has 15 years experience in the design of AC-DC and DC-DC power supplies.

Acknowledgement. *The author greatly appreciates the valuable input of Craig VonIlten at Ziatech Corp. regarding the interactions between CompactPCI systems and the power supplies.*

If you have questions about this article, or if you would like more information about Condor DC power supplies, you may contact the author at:

Condor DC Power Supplies
(Todd Products Group)
50 Emjay Blvd.
Brentwood, NY 11717
Tel: 516-231-3366 ext. 256
Fax: 718-382-1303
or 516-952-3048.
E-mail:
Lazar_Rozenblat@condordc.com
Website: www.condordc.com