

Chapter 2: *Selecting a Power Supply for System Application*

The availability of high quality standard products has greatly simplified the task of the power system designer. However, certain parameters must still be clearly understood before the designer can realise a system optimized with respect to cost, reliability and performance.

2.1 Introduction

With the availability of standard power modules with proven reliability, the task of the power system designer has been dramatically simplified. No longer must he wrestle with the development of custom power supplies of unknown characteristics, but rather can trust that standard products will perform well if applied within their design limits. Nevertheless, the specific standard product and supplier must be selected so that the overall power system design will be optimized in terms of cost, reliability and performance. This chapter will take a general look at that selection process. It will focus on some of the more important criteria in terms of system requirements and how they relate to the types of specifications commonly found in the power converter manufacturer's datasheets.

The Telecom market is one of the most stringent in terms of the electrical and mechanical standards that a power supply must meet. Providing Telecom power on a multi-national basis exposes the power designer to literally hundreds of various national standards in the areas of safety, EMC/EMI, powerline interface and manufacturing quality. Artesyn has been one of the leading suppliers in this market since 1968 and has extensive expertise in all these areas. You, as a customer, deserve peace-of-mind that the resultant power system will meet all its goals in terms of quality, reliability, availability, efficiency and cost while minimizing technical risk. Artesyn is prepared to deliver on this promise.

Much of this chapter will be discussing requirements based on a Telecom point-of-view. This was done because other types of equipment often have less stringent demands, and the understanding of the more severe requirements can be applied to all systems by omitting items that are not relevant. The information here is intended as only an overview. In many cases, a more detailed treatment of these concepts will be found in later chapters of this book.

Selecting a Power Supply for System Application

2.2 System Requirements and Power Supply Specifications

We will now look at system requirements and how they relate to power supply datasheet specifications. There are basically two types of numbers on a datasheet - absolute maximum ratings and performance specifications. The absolute maximum ratings indicate the stresses that may be applied to the device without incurring any permanent damage. For example, there may be a case temperature limitation of +110°C. This means that the case temperature could reach this limit for a brief period and the converter would not catastrophically fail. It does not mean that this would be a wise design point for long-term operation! The numbers used for more normal operation are in the performance specification sections. The supplier may choose to specify a performance parameter in terms of a minimum, maximum or typical specification - and sometimes all three. And, depending on your system requirement, you may be more or less interested in each of these ways of writing the specification. For example, if you are concerned with tripping an overcurrent limit, the minimum value of the trip point would be most useful. If you are calculating the expected heat load of a large number of converters in a system, the typical efficiency would be the most useful specification.

2.2.1 Voltage Regulation

Most power sources for electronic equipment are regulated. That is, there is a feedback loop internal to the design that attempts to hold one of the output parameters constant. The controlled parameter is most often voltage, although some power supplies are designed to maintain output current as a constant (with variable voltage) or to provide constant power output. We will discuss the most widely used constant output voltage power supply.

There are two basic elements to the output voltage regulation specifications - the nominal value, or 'set point', and the variation from this nominal value as a function of other parameters or independent variables. The most commonly used variables for the specification of voltage regulation are load regulation, line regulation and operating temperature*. We will look at each of these along with other methods sometimes used to specify the voltage regulation performance of a power converter.

The set point is actually a range of output voltage within which the converter is guaranteed to operate for a specific set of conditions. For example, assume a 100W DC/DC converter with an input voltage range of 36 to 75V. The manufacturer will specify the set point voltage under what would be considered nominal conditions. In this case, the set point might be specified at an input voltage of 48V, full load of 100W and a temperature of 25°C. A typical set point specification would be $\pm 1\%$. Thus, assuming a 5V converter, the output voltage set point would be between 4.95 and 5.05VDC. The other regulations specifications, to be described below, will be treated as deviations from this nominal set point.

The regulation specification that receives the most attention is the load regulation. In many applications the load current will vary more than the temperature or input voltage, so will consequently have a greater effect on the output voltage than will the other variables. The load regulation is specified either as a percentage of output voltage or as a specific voltage deviation. For our example of a 5V DC/DC converter, typical load regulation specifications could be equivalently expressed as either $\pm 0.1\%$ or $\pm 5\text{mV}$. There should be a load current or power range specified as a condition for the load

* Load regulation is a measure of how the output voltage setpoint is affected by defined changes in the load current. Line regulation describes the effect on the output voltage setpoint of defined changes in the input voltage.

regulation specification - for example, no load to full load or 20% to 100% load. In many system applications, the actual load range is more limited, and the voltage deviation due to load regulation will be significantly less than the datasheet value. There will also be a fixed input voltage defined for this specification to eliminate the effects of line regulation. Modern converters typically have excellent (low deviation) load regulation specifications, the exception being converters designed for slope or droop load sharing, which will be discussed in the section on paralleling. Consequently, static load regulation is not often encountered as a restriction in meeting the system voltage regulation requirements.

Line regulation describes the effect of the input voltage to the converter on its output voltage. The input voltage will be AC in the case of AC/DC converters and DC for DC/DC converters. The specification will be done at a fixed load current to eliminate the effects of load regulation. The input voltage range for the specification is most often the total operational input range for the converter - i.e., 36 to 75VDC for a nominal 48V input converter. The line regulation can be specified either in terms of percentage or an absolute value. The line regulation for a 5V output converter could be expressed as 0.05% or as $\pm 2.5\text{mV}$. In either case, the converter would show almost no change in output voltage over its specified input voltage range.

Another variable which affects the output voltage is the operating temperature of the converter. The temperature dependence of the internal voltage reference circuitry is often the major contributor to this specification. The specification is usually expressed in terms of $\pm\%/^{\circ}\text{C}$ so that the user can size the temperature effect based on the expected temperature range of the actual application.

For our 5V converter example, if the specification was $\pm 0.01\%/^{\circ}\text{C}$, the set point specified at $+25^{\circ}\text{C}$, and the actual expected operating temperature was $+60^{\circ}\text{C}$, then set point could be shifted up to 17.5mV outside of the specified set point range.

Some manufacturers will include a drift and aging specification. The intent is to separate out the long-term effects of component aging on the output voltage from the performance of the converter as it is initially installed in the end use equipment. Other manufacturers do not specify this effect individually, but rather include it within another specification, such as total regulation or total error band.

To simplify the calculations required on the part of the user, some or all of the above specifications are sometimes combined into one number called total regulation. For example, a converter's total regulation specification could include the effects of the operational input voltage range, no load to full load output current, effects of operating temperature over the rated temperature range, and component aging. For a 5V converter, a typical total regulation specification could be $\pm 3\%$ or $\pm 150\text{mV}$. In this case, only the set point specification and the total regulation specification would need to be considered to determine the output voltage operational envelope for static load conditions.

Figure 2.1 compares the total regulation approach vs. the individual specification of contributors to the output voltage deviation. The figure depicts a worst case scenario in that it directly adds the maximum specified values for each element. In an actual system situation, the calculations should be done for the anticipated conditions in the system. This will result in a smaller and

Selecting a Power Supply for System Application

more realistic total regulation range. As system operating voltages become lower, voltage regulation requirements and specifications become more important. Most modern converters will have no problem meeting the system static voltage regulation requirements for outputs of 3V and above, but at lower voltages the system requirements and the regulation budgets will require closer examination.

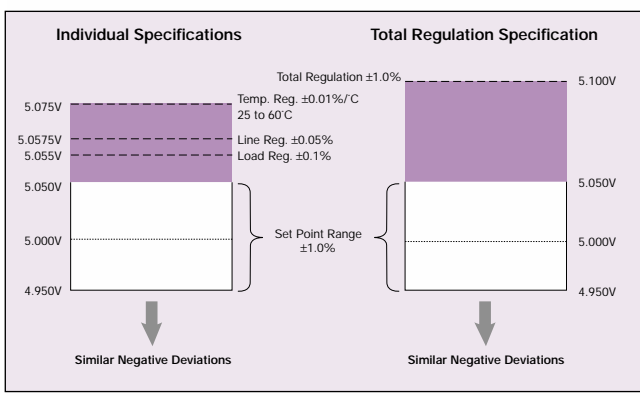


Figure 2.1 - Example of Voltage Regulation Specifications for 5V Converter

2.2.2 Dynamic Responses

The voltage regulation specifications and characteristics described above apply to static output loads. In general, there will be another deviation in output voltage during the time immediately following a change in load current. This deviation will last until the converter has had a chance to adjust to the new value of load current, and is referred to as the dynamic response or transient response. The severity of the deviation in output voltage will be a function of the change in output current along with the rate of the change, or di/dt . Rapid changes will have a larger effect due to the limited response time of the converter and the inductance inherent in the voltage distribution system. A depiction of the converter dynamic response is shown in Figure 2.2.

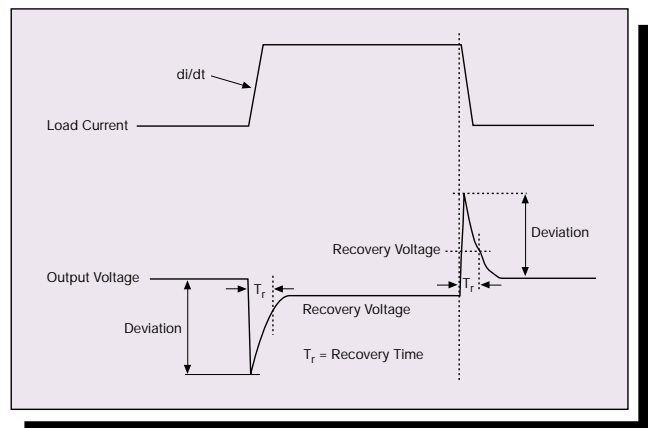


Figure 2.2 - Dynamic Response

When reading a datasheet for a converter, it is important to note the conditions stated for the dynamic response specification. Both the absolute change in current or load should be stated along with the rate of change of load. For example, a typical condition could be 50% to 75% load with a di/dt of $0.1A/\mu s$. The dynamic response to this load change is specified in terms of the maximum voltage deviation and the time required to recover to some defined limit. For example, the specification could be a 2% maximum deviation with recovery within 200ms to 1% of the static value. The voltage deviation will be negative during the increasing current transition and positive during the decreasing current transition.

We will present more detail on how to determine and design for dynamic response requirements in a chapter on distribution and decoupling design. For now, however, it is important to note that, from a system point of view, dynamic performance is determined by the distribution and decoupling network as well as the power converter. In fact, even with a perfect converter (infinite dynamic response) most systems will be non-functional without attention to the distribution and decoupling design. One current extreme example is the dynamic

response requirement for high performance processor chips which need to go from stand-by to fully operational in a short period of time, with values of di/dt up to 30 A/ms. Even with a high performance voltage regulator module located physically close to the processor chip, several decoupling capacitors are required to meet this requirement. For most systems, the dynamic response requirements will be a larger factor in the selection of a converter than will the static voltage regulation.

2.2.3 Output Ripple and Noise

Unlike linear regulators, switching regulators inherently generate some noise during their operation due to the non-linear nature of the voltage and current waveforms, and some of this noise appears on the output voltage terminals. There are actually two distinct components to the output noise, a 'low' frequency component referred to as ripple, and a higher frequency component that is simply called noise. The frequency of the ripple component is tied to the basic operating frequency of the converter. Depending upon the topology of the converter, the ripple frequency will be either the same as the converter operating frequency, or some multiple thereof - most often twice. The ripple exists because, during a portion of the converter's operating cycle, energy is transferred to the secondary from the primary and the output voltage increases slightly. During the time interval when there is no energy transfer to the secondary, the load current is supplied by stored energy in the output capacitance and inductance of the converter, and the output voltage decreases slightly as this energy is depleted. Figure 2.3 shows a typical ripple waveform.




Figure 2.3 - Output Ripple and Noise

Noise is much more variable and harder to predict than ripple. It is caused by ringing in parasitic inductances due to the large values of di/dt that occur internally in a switching converter. The noise is much higher frequency than the ripple - well up into the MHz range with a high harmonic content. The noise occurs in the form of 'bursts' at the time of switching activity in the converter, so therefore appears to be superimposed upon the peaks and valleys of the ripple waveform as shown in Figure 2.3.

The high frequency noise normally has a fairly low energy content, so that its magnitude and waveform will change as a function of exactly where it is measured and the impedance and capacitive loading of the measurement device. Consequently, the manufacturer of the converter will specify a measurement point and technique. The most common convention is to use some sort of low-pass filter network to limit the measurement to the more reproducible and consistent lower frequency components. This is in keeping with the normal system application where the DC distribution system and associated decoupling acts as a low pass filter. 20MHz and 50MHz are the two most common measurement bandwidths specified, and can be easily achieved by using input filters built into many popular oscilloscopes. The chapter on system testing details the noise measurement technique recommended by Artesyn

Selecting a Power Supply for System Application

Technologies. Ripple and noise are most often specified together and measured as a peak to peak value. This is the easiest measurement to make. Some manufacturers also specify the noise and ripple in terms of average and/or rms voltage.

For a 5V output converter, a typical ripple and noise specification could be 75mV peak to peak with a 20MHz bandwidth. Performance in this range, along with standard distribution and decoupling methods, should be more than adequate for most systems. Exceptions might be extremely sensitive analog circuitry, in which case additional attention to filtering and layout may be required.

2.2.4 Static Input Specifications

The most basic input specification is the input voltage operating range for the converter. This could be either an AC voltage range or a DC voltage range depending upon the type of converter under consideration. For example, an offline converter might be specified with an input range of 170 to 264VAC for a 200/240V nominal system and a DC/DC converter with a nominal 48V input might be specified from 36 to 75VDC. This is an important specification for the power system designer, because it must be guaranteed that the converter will operate over the minimum and maximum values of the input voltage actually seen during the operation of the end use equipment.

For AC systems, the input voltage range should encompass the normal range of static AC powerline voltages, as defined by the appropriate line voltage standard. For DC input systems, such as distributed power systems, the broad input range is needed to allow

for such things as hot-plugging, bus switching transients, switchover to battery back-up, and battery depletion.

When selecting a DC/DC converter the input voltage levels at which the converter turns on and off should be considered as well as the overall input voltage operating range. The turn-on voltage should be higher than the turn-off voltage when the converter is used in DC input telecom systems or systems that utilize battery backup. As the batteries discharge and the input voltage decreases, the converters will begin turning off at some point. When they turn off, the load on the batteries will decrease and the battery voltage will rise, turning the converters back on again. This results in a type of oscillation, with converters turning on and off repeatedly, an undesirable situation.

This situation is remedied in the converter design by including hysteresis in the input voltage detector as shown in Fig. 2.4. With the converter in the off condition the converter will start up at some point as the input voltage increases. This 'point' is actually a small range as indicated with a '?' in the figure. If the input voltage decreases after the converter is running, the turn-off range will be at a lower voltage than the turn-on range. This will allow the converters to remain off even if the battery voltage raises slightly. This feature protects the batteries from further discharge as well as preventing the turn-on, turn-off oscillation in the system. For a typical nominal 48V input converter, the turn-on range might be in the area of 35V with the turn-off range around 30V.

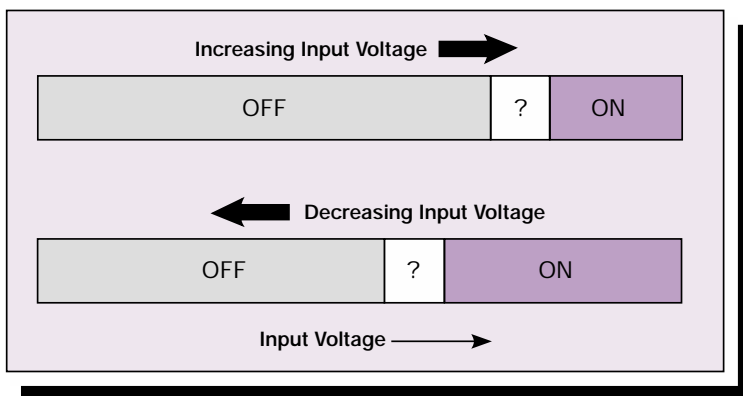


Figure 2.4 - DC/DC Converter Input Hysteresis

More and more of today's systems that interface with the AC powerline are expected to meet some form of line harmonic or **power factor correction** (PFC) standard, the most ubiquitous being EN61000-3-2. In the past only higher power converters were subject to these requirements, but now even converters in the 200W range often have a PFC requirement. Fortunately, the technologies and topologies for implementing PFC are now readily available and cost effective for most applications. These circuits can achieve power factors of >0.99, and the better AC/DC converters are often specified accordingly. PFC is considered in more detail in a forthcoming chapter.

Some AC/DC converters will have a **supply voltage ripple rejection** specification, although this is becoming less and less common. Ripple rejection can be thought of as line regulation, but at the line voltage frequency rather than at DC. Because of the high effective gain of the conversion circuits used in today's switching AC/DC converters and the high conversion frequency relative to the line frequency, the ripple rejection at the line frequency is essentially as good as the static line voltage regulation, and is not specified separately. In cases

where it is specified, the performance is usually stated as the peak to peak voltage component on the output at the frequency of the rectified line voltage. Alternatively, it is sometimes specified in terms of attenuation of the input AC voltage in terms of dBV. For example, a converter with a 100VAC input and specified at 80dBV ripple rejection would have a 10mV rms voltage content at the line frequency on its output.

As a consequence of the internal switching activity, switching power converters will not have a constant DC input current, but rather an AC current component superimposed on the average DC current. This AC component of the input current is called **reflected ripple**. The magnitude and waveform of the reflected ripple is a function of the circuit topology used and the input filtering internal to the converter. The reflected ripple is useful to the power system designer when designing the input distribution system and EMC/EMI filtering.

AC/DC converters also typically specify the input **frequency range**. This is to a large extent a holdover from the time when there were line frequency transformers used within the power supply, and the frequency of the input voltage could have severe effects on the efficiency and operation of the power supply. Today's switching converters can inherently operate over a much wider range of input frequencies, but the input frequency can still play a role in the specification in the case where there are AC powered fans or auxiliary power supplies within the converter. A typical frequency range specification for a 50/60Hz nominal system would be 47 to 63Hz.

Selecting a Power Supply for System Application

2.2.5 Dynamic Input Specifications

It is desirable for a converter to operate through a temporary reduction in its input voltage below the static input voltage range. This results in a much more robust system that is not affected by such things as powerline dips and sags due to load switching activity or short-term outages. This converter characteristic is often specified and is referred to as **hold up time, holdover, ride-through or dropout immunity**. For AC input systems, the specification is often in the range of 20ms for a complete dropout of the AC input voltage and can be much longer for less severe deviations below the lower static input voltage limit. Hold up time is less often specified for DC/DC converters intended for use in DPA systems. Hold up time for DC/DC converters is provided by additional capacitance on the input. The required capacitance is best packaged as part of the board level DC input circuitry as described in a later chapter.

Both AC/DC and DC/DC converters typically have specifications for **inrush current** or **surge current** such as ETSI 300-132-2. All of Artesyn Technologies' low power DC/DC products are designed to meet this specification. This specification is important in terms of understanding the interaction of the converter with the input power source when the connection with the converter is established. Most converters have some form of capacitance on the input which can result in initial peak input currents in excess of the maximum steady-state currents. Both the magnitude and duration of this current will affect the selection of appropriate fuses and circuit breakers as well as the distribution conductors. For AC/DC converters, it is specified in terms of the peak amps that would be seen on the powerline if the connection was made at the peak of the AC input voltage. For DC/DC converters, the

specification is sometimes defined in terms of peak amps as above. Alternatively, the DC/DC converter manufacturer can specify the actual value of equivalent capacitance on the input of the converter. This allows an actual calculation of the inrush current waveform, which will be strongly influenced by the source impedance of the power distribution network ahead of the converter, the actual input DC voltage and the DC input voltage filtering/protection network on the system board.

Dynamic input specifications, for both AC/DC and DC/DC converters, are also tied to several international standards that address issues in the areas of safety, EMC and transient input energy. These issues and standards will be explored in more detail in the later chapters that address input voltage network design for DPA systems and telecom standards.

2.2.6 Efficiency

Converter efficiency is perhaps the most important single specification as it will influence many aspects of system operation and performance. The power system designer is encouraged to understand thoroughly the impacts of efficiency on the overall system design and cost and to consider it carefully when selecting a converter. To further understand why efficiency is so important, we will look at a typical DC/DC converter application. Assume that the output power of the converter is 50W. Let's now look at the power dissipated internally. This will be the difference between the input power and the output power, where the input power is $P_{out}/\text{efficiency}$. Therefore, for our 50W converter, the dissipated power can be expressed as:

$$P_{diss} = 50 / \eta - 50 \quad \text{where } \eta = \text{efficiency}$$

The above expression is evaluated for various values of η and the results displayed in Figure 2.5. We can now graphically see the huge influence of efficiency on the power dissipated. DC/DC converters are available with efficiencies between 75% and 90%. At 90%, the dissipation would be 5.6 watts. At 75%, it is 16.7 watts - 3 times larger. **A 15% change in efficiency causes a 300% increase in power dissipated!**

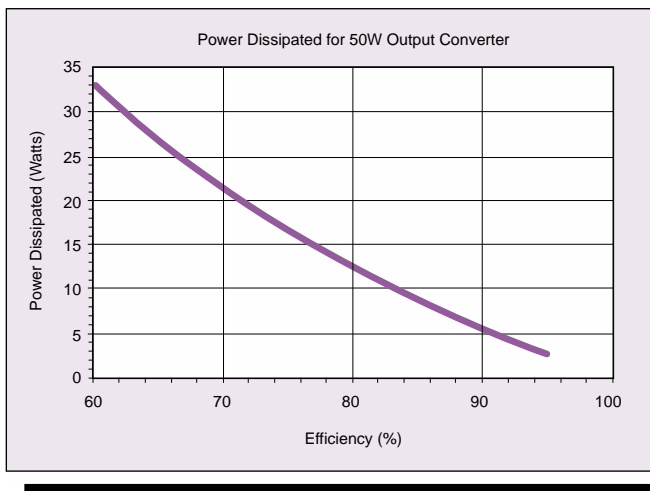


Figure 2.5 - Effect of Efficiency on Power Dissipated

Why is the dissipated power so important? For several reasons. Extra dissipated power will raise the operating temperature of the components in the converter and this reduces the reliability. It may require additional physical space and cost in the system for heatsinks and/or increased airflow. This will increase the size of the converter. It will increase the cost of energy for operating the system over its lifetime. It will increase the power requirements for all the conversion and distribution functions ahead of the converter (i.e. at the input side). Clearly, the dissipated power should be minimized if an optimal system is the objective. See Figure 2.6 for other system parameters influenced by efficiency.




Figure 2.6 - Effects of Converter Efficiency on the System

All converter manufacturers specify efficiency, sometimes as a typical value, sometimes as a guaranteed minimum and sometimes as both. If it is a reputable supplier, the typical values can usually be relied upon to reflect an average value for the conditions specified. Some converters have a fairly flat curve of efficiency vs. power output, while others have an efficiency peak at over a specific load range. For applications where the exact value of efficiency is critical, it is important to understand these curves. The efficiency will also vary with the input voltage. Most suppliers will specify efficiency at the nominal input voltage for the device. If your application is in this nominal range, as most will be, then the datasheet values will be adequate. If you are operating the converter at a steady-state input voltage other than nominal, then you should investigate the efficiency under this condition. Some suppliers will include such data on the datasheet, while others will provide it upon request.

Selecting a Power Supply for System Application

There has been significant progress in improving efficiency since the introduction of the first DC/DC converters. A few years ago, commonly available 5V converters had an efficiency of only 75%. The later generation devices have improved to the point where 85% is realizable. As this manual was being written in early 2000, it was a time of rapid improvement in efficiency due to the more general availability of synchronous rectification in DC/DC converters. The very latest devices using synchronous rectification provide efficiencies of up to 93% at 5V. Figure 2.7 summarizes some efficiency benchmarks for both AC/DC and DC/DC converters.




Figure 2.7 - Year 2000 Converter Efficiency Benchmarks

2.2.6 Reliability

Power converter reliability has become increasingly important as the expectations for availability of end-use systems become more severe. With the increased use of distributed power architectures (DPA), systems now contain dozens of power converters instead of one or two. As a consequence, each of them must be much more reliable than their previous generation products. Another chapter addresses reliability issues in much more detail, but here we will take a brief look at some of the most important reliability related criteria to keep in mind when selecting and applying power converters in your system.

Reliability will most often be specified in terms of Mean Time Between Failure (MTBF), with values extending up into the millions of hours. We will discuss how this relates to failure rate in the chapter on reliability. The MTBF value is only half of the specification, however. Equally important is the converter ambient or case temperature at which the specification applies. This number will vary from supplier to supplier. The actual converter reliability you will experience has a **very strong** dependence on

the operating temperature in your system, as shown in Figure 2.8. Consequently, you must adjust the published reliability data to reflect the operating conditions in your system before you have any meaningful projection of the actual expected reliability. The chapters on thermal design and reliability will assist with this process.




Figure 2.8 - Effect of Operating Temperature on Reliability

When selecting a converter and its supplier, it is important to understand the basis for their published reliability data. Some suppliers use test data such as field history or accelerated life tests. Others will use one of the generally accepted prediction methodologies such as MIL-HDBK, Bellcore, internal component databases or comparison with similar models. All of these approaches can result in valid data. The most important thing is to have confidence in your supplier and in their prediction or testing methodology. This confidence is established based upon the supplier's reputation and on the nature of your interaction with their engineering and marketing people.

DC/DC converters tend to be more reliable than AC/DC converters. This is because they are functionally less complex, with only one power conversion stage. Even more importantly, they are implemented with extreme levels of integration, with low component counts, high quality materials and components and automated assembly processes. Reliability in the range of 1 to 5 million hours MTBF is achievable with the latest designs. AC/DC converters will have a smaller MTBF number due to the factors mentioned above and the inclusion of higher failure rate components such as fuses and fans.

High quality AC/DC converters will have MTBF specifications in the range of 100 to 500 thousand hours.

2.2.8 Remote Sensing

The remote sensing decision affects both the power system design and the power converter selection. When most power systems used centralized power converters, remote sensing was almost always used on low voltage outputs in order to compensate for the voltage drop due to the distribution distance between the converter and the load. This is still the case today for low voltage centralized systems. With DPA, the situation is different. The intermediate bus voltage does not need precise regulation due to the input voltage range of attached DC/DC converters. Individual DC/DC converters are often assembled on circuit board along with their load with a short and low impedance distribution path between them. As a result, remote sensing is often not a requirement.

The advantage of remote sensing, of course, is that it allows the system designer to tightly control the absolute value of the voltage at one point in the power distribution network. This is sometimes useful for a very critical load circuit that is located some distance from the power converter. The most common applications involve low values of voltage, high values current and situations where the DC current level has a large variation. Remote sensing is not without disadvantages. Some are listed below:

- It adds additional complexity. At least two more conductors must be used for the sense lines along with resistors between the sense lines and outputs to prevent excessive voltage in the event of an open sense line.

Selecting a Power Supply for System Application

- It adds additional cost. Both system costs and the cost of the converter itself are increased.
- When using remote sensing, the inductance and capacitance of the distribution network between the converter and the sense point is inside the converter's control loop. This complicates the stability analysis for the converter. The system designer must understand all of these circuit elements and, in conjunction with the converter manufacturer, conduct stability analyses that integrate both the converter and the distribution system.
- If there is a need to parallel power converters, remote sensing complicates the process. Some converters use a slope compensation technique for paralleling, which is not compatible with remote sensing.
- Remote sensing adds another failure mode for the system. An open sense line must be diagnosed and repaired, neither of these tasks is particularly straightforward.

Because of these disadvantages, it is recommended that remote sensing only be used when absolutely necessary. It is a common feature in low voltage AC/DC converters and in high power (> 50W) DC/DC converters. The power converter will specify the maximum DC voltage that can be compensated for by the remote sensing circuit - typically 0.5V or less.

2.2.9 Paralleling

Paralleling is the connection of two or more power converters to simulate a single converter, and is used in two situations. In cases where the current demand is greater than the output capability of a single power converter, converters are sometimes operated in parallel for increased current and power capability. In high

availability system designs, two or more power converters are connected together in a redundant 'N+m' configuration so that the failure of a single converter will not affect the system operation. Both of these situations require that converters be operated in parallel. The design issues of paralleling will be discussed in more detail later, but we will provide a short summary here as it applies to the selection of a converter.

There are three commonly used techniques for implementing the paralleling. The first is a direct connection of the converter outputs without provision for current sharing between the converters. This method, while feasible, has the disadvantage that one converter will provide the majority of the load current and consequently run hotter and have a higher failure rate than the others. The degree of this imbalance can be reduced by increasing the resistance of the output voltage distribution network, but at the cost of additional power losses. The imbalance can also be reduced by careful trimming of the converter output voltages so that they are very close to each other, but this operation adds manufacturing and field repair complexity and cost.

The second technique, slope compensation, also uses a direct connection of the converter outputs as described above, but with power converters specifically designed for slope compensation. A converter with slope compensation has an intentionally 'soft' voltage regulation characteristic so that the output voltage will drop slightly with increased load. This characteristic will automatically force current sharing between the converters without the need for specialized circuitry or interconnections. The downside is reduced voltage regulation, which may not be acceptable in some systems.

The most sophisticated technique is active current sharing. This approach requires converters with internal circuitry to achieve balanced current sharing between the converters and a separate current sharing connection between the converters. This is the highest performance approach but also the most complex and costly.

The complexities introduced by paralleling should be carefully considered and paralleling avoided if possible. The most common application in DPA systems will be the paralleling of front-end AC/DC converters for the intermediate bus voltage generation. With the wide range of power levels now available in standard DC/DC converters, and the wise usage of circuit partitioning techniques, there should be very little need for paralleling DC/DC converters for increased power capability. For systems requiring redundancy, it is often possible to provide redundancy at a higher functional level by using multiple functional cards, each with its own power converter. It should also be noted that some converter topologies using the latest high efficiency synchronous rectification techniques are not easily paralleled.

2.2.10 Remote On/Off Control

Many system applications will require that the power converter be turned on or off remotely by means of a signal from elsewhere in the system, and it is now common for power converters to have a remote on/off feature to accommodate this need. This feature is useful for enabling individual voltages in a desired sequence or for delaying enabling of a voltage in the case of a hot-plugging situation. The remote on/off pin on the power converter is typically able to be driven by a low-level logic signal. It is important that this interface be completely specified by the power converter

manufacturer so that you will know how much current the system signal will need to source or sink, the maximum open-circuit voltage stress that will be seen and the maximum down-level voltage. This information will allow for a robust interface design. Some manufacturers offer two versions of their converters - one that enables the converter with a down level and one with an up level. Many systems will not require usage of the remote on/off feature. It can be bypassed by connecting two of the converter pins together so that the device will automatically power up when the input power is applied.

Most DC/DC converters have the majority of the control and conversion circuitry on the primary side of the isolation transformer, and the remote on/off signal interface is referenced to one of the primary pins rather than to the secondary. The system designer must be aware of this and provide an appropriately referenced control signal in cases where isolation is required between primary and secondary ground systems. If no primary-referenced signal is available, one can be created easily by using an optical isolator in conjunction with a secondary-side signal.

2.2.11 Thermal

Thermal design is one of the most important aspects of power system development and one of the most overlooked. The trend toward DPA has changed the physical environment of power converters and consequently the type of thermal analysis required. More first-time DPA designs are unsuccessful due to inadequate thermal design than for any other reason. It is imperative that today's power system designer be as comfortable with thermal analysis as with electrical. For this reason, there is a complete chapter devoted to

Selecting a Power Supply for System Application

thermal analysis and design. In this section, we will only cover some of the highlights as they relate to the selection of a power converter.

There are several places in the system and in the power converter where temperatures are specified and measured, and this is often a source of confusion for the power system designer. Confusing ambient temperature with power converter surface temperature, for example, can result in an error of 20°C or more with dire consequences on the reliability or even the operability of the system. There are four major categories of temperatures to deal with. First is the air temperature external to the system - the temperature in the room itself. Second is the air temperature inside the equipment in proximity to the power converter, referred to as ambient temperature. Note that the ambient temperature inside the equipment is typically significantly elevated above the external temperature. When power supply manufacturers refer to ambient temperature, this is what they mean - not the room temperature.

From a power system designer's point of view, perhaps the most important temperature to get a handle on is the power converter surface temperature. The converter manufacturer will define the location of this point. Sometimes it will be on the surface of the converter case or baseplate. Sometimes it will be a temperature on a pin of the converter. In the case of an open-frame converter the manufacturer will sometimes define the surface of a specific component for this measurement point. The surface temperature, no matter where it is defined to be, is the most useful because it can be easily measured and, with the help of the converter's datasheet, will give the most meaningful data on the reliability of the converter in its actual operational environment.

The final temperature of interest is the junction temperature of the power semiconductor devices internal to the converter. The power system designer, of course, cannot directly measure this temperature. But, if the converter's surface temperature is known, the junction temperatures can be estimated by using thermal impedance data obtained from the supplier. It is the junction temperature, ultimately, that will determine the converter's reliability. Figure 2.9 shows how these temperatures relate to each other.




Figure 2.9 - Important Power System Temperatures

The other significant parameter that the system designer must understand and quantify is the airflow environment in the vicinity of the power converter. For forced convection systems, the magnitude of the airflow in linear feet per minute (lfm) or m/s is equally important as ambient temperature in determining the thermal performance.

With estimates of case or ambient temperature and the airflow, the power system designer is prepared to select an appropriate converter with the help of the manufacturer's datasheets. Each supplier has a somewhat unique way of presenting this data, but each

will have one very important data point - the maximum surface or case temperature for full rated operation. This will be the surface temperature the manufacturer has determined corresponds to the maximum desirable junction temperature of the components internal to the converter. Given that the converter has a good internal thermal design, the higher the maximum surface temperature specification the better. A high case or surface temperature rating will allow maximum power to be extracted from the converter at elevated ambient temperatures.

Another important set of data is the derating curve. These will enable the power system designer to determine the power output capability of the converter as a function of the ambient temperature, airflow and auxiliary heatsinks used. An example of a typical curve is shown in Figure 2.10. Some power converters will have internal circuitry that monitors the operating temperature and initiates a warning signal and/or an automatic shutdown at specific temperature points. These functions can be helpful in some applications. If they are implemented, the datasheet should define the temperature sensing thresholds, the type of shutdown used, and the temperature hysteresis between shutdown and turn-on.

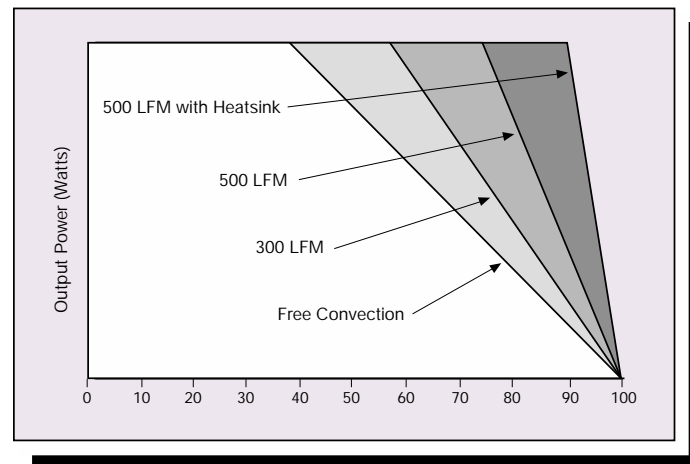


Figure 2.10 - Typical Derating Curve

2.2.12 Safety

From a system design point of view, the most basic safety-related issue for AC input systems is isolation from the powerline, which is provided as one of the functions within the AC/DC converter. The isolation test voltage should be listed in the AC/DC converter datasheet, as a function of the standard it was designed and tested to. Once this basic powerline isolation has been accomplished at one point in the system, the system designer's options relative to additional isolation become much more flexible. In DPA systems, a second level of isolation is often used in the DC/DC converters, but in most cases this design decision relates more to grounding configurations and noise control than to meeting required safety standards.

Another safety-related design decision that will need to be made in DPA systems is the voltage level of the intermediate voltage bus. The total system design including packaging, shielding and field maintenance becomes greatly simplified if the intermediate bus voltage

Selecting a Power Supply for System Application

is below the Safety Extra Low Voltage (SELV) level as defined in the appropriate safety standards. The SELV level is 60 Vdc, so almost all recently designed DPA systems use an intermediate bus voltage below this level. The most common 48V nominal bus complies with this standard.

It is also important to understand early in the design process where the product will be marketed, as this will have some effect on the standards and agency approvals required. Your own standards compliance people will need to be involved in the details of the standards work and determine the standards and approvals needed for the power converters. In general, however, you will probably be looking to meet the requirements of EN60950, which is the most universal international safety standard. The three approval agencies most often used to verify compliance to this standard are UL in the USA, CSA in Canada and VDE for international markets. Your time-to-market and cost of compliance for the entire system will both be reduced by selecting power converters that have pre-approval to the EN60950 standard by the agencies listed above. This subject is discussed in greater detail in the chapter on safety standards for power supplies.

2.2.13 EMC/EMI

As with safety, the general category of EMC and EMI performance is a system level requirement that has implications for the power system design and the selection of power converters. This topic will be discussed in much more detail later, but for now we will highlight some of the more important considerations. In general, the situation is a very positive one. Both the requirements for EMC/EMI and the methods of

successfully designing to them are better understood now than in the past. DC/DC converters, in particular, have made significant progress in this area, making selection and application easier with respect to the EMC performance.

The EMC/EMI design concepts for centralized power systems using AC/DC converters have remained fairly constant over the years. Centralized supplies tend to be well shielded and located near the input powerline filter and some distance away from the load circuitry. As a consequence, most of the recent questions about EMC/EMI performance and design tend to revolve around the use of DC/DC converters located in proximity to the load circuits in systems using a distributed architecture. This can, at first, seem somewhat overwhelming since the designer must consider differential and common-mode noise on both the converter's input and output, both radiated and conducted noise and both converter emissions and converter susceptibility. Each situation is unique, but some general comments about the EMC/EMI performance of distributed DC/DC converters in telecom systems may prove comforting.

With good basic power system design practices and the use of converters from reliable suppliers, telecom systems usually have no problems meeting their EMC/EMI objectives. Radiated emissions and susceptibility rarely present a problem. Most of the focus tends to be on the input side of the converter, both in terms of conducted emissions from the converter back onto the distributed bus and the converter's susceptibility to transients on its input in a telecom environment. Both of these concerns can be easily resolved by the addition of a few inexpensive components on each assembly

containing a DC/DC converter. As will be seen in a later chapter, normally a few components are already needed here for fusing, input energy storage and hot-swap capability. Adding a small filter network in this area tends to be inexpensive and allows for meeting even class B EMC/EMI requirements at the system level.

In terms of selecting a converter for such an application, the key is again the credibility of the supplier. Manufacturers who have a successful history of integrating their converters into telecom products will have the data, design suggestions and answers you need in order to meet your system EMC/EMI requirements. In terms of what to look for on a datasheet, you want to see that EMC/EMI has been considered in the converter design. This will be evidenced by, at a minimum, a listing the EMC/EMI standards that the converter will meet and the external components, if any, required to do so. But, in the final analysis, EMC/EMI is a system issue and your confidence in the expertise of the supplier and your ability to interact with them on system design issues will be the major factor in a successful EMC/EMI design.

2.2.14 Short Circuit Protection

Short circuit protection is one of the main goals of system safety design. The objective is to prevent possible overheating and fire hazards in the event of a fault in either a power converter, the load circuits or the power distribution network. In a typical telecom system implemented with a distributed architecture, short circuit protection for the AC line itself is provided by circuit breakers at the power input to the equipment. A fuse at the input to each assembly provides protection to the card and board distribution from possible internal faults in distributed DC/DC converters. The remainder of the

system circuit protection is usually achieved by some form of electronic overcurrent protection in a power converter. The front-end AC/DC converter provides protection for the intermediate bus distribution system and the DC/DC converters provide protection for the low-voltage output distribution and the load circuits. In fact, unlike other control and fault-protection features, some form of overcurrent limiting is standard on all modern power converters. In the remainder of this section we will discuss briefly the types of electronic overcurrent protection as they relate to power converter selection.

There are four types of current limiting in widespread use in power converters. The first is called continuous or brickwall, and is illustrated in Figure 2.11. When the overcurrent trip point is reached, the power supply becomes a constant current supply at this current level, and the output voltage goes to near zero. The absolute value of the output voltage will depend upon the system's distribution resistance on the output of the converter. This value of current will continue to flow from the power converter until the output overload is removed, at which point the converter will automatically resume normal operation. This approach is simple to implement both from a converter design and system design point of view.




Figure 2.11 - Types of Electronic Current Limiting

Selecting a Power Supply for System Application

Foldback current limiting is a variation of the continuous approach. Once the current limit point is reached, the converter becomes a constant current supply, but the value of this current is a function of the degree of the overload. This characteristic results in a sloped output voltage vs. output current line as shown in Figure 2.11. The advantage of using foldback is that the continuous current and power in a short circuit is lower than for the continuous approach. The disadvantage is that, depending upon the load impedance, these circuits sometimes 'latch up' and will not recover automatically. For this reason, more care must be taken with the system distribution design. Foldback approaches were used extensively in linear regulators, but are less commonly used in switching converters.

Another approach is the auto retry or 'hiccup' type of current limiting. When the overcurrent trip point is reached, the internal control circuitry turns off the converter so that both the output current and the output voltage is zero. Then, after a pre-set time delay, the converter is turned on again and the output voltage ramps up. If the fault is still present, the output voltage will rise until the overcurrent limit is reached at which point the converter will again shut down. This process will be repeated, as shown in Figure 2.11, until the fault on the output is cleared. This approach minimizes current flowing through the fault and offers the opportunity for 'self-healing' systems in the event of a temporary overload.

The latching overcurrent protection approach is similar to the auto-retry in that the power converter is completely turned off in the event of an overcurrent situation. But there is no automatic retry, and the converter stays in the off condition until it is reset by either manual or system

intervention. Often this reset is in the form of removing and reapplying input power to the converter. This approach lacks the 'self-healing' feature of the auto retry method, and is used when it is desirable to either shut down the converter completely or notify an operator in the event of a fault.

Some more sophisticated converters may incorporate some elements of two or more of these techniques. For example, the Artesyn AFE series of AC/DC front-end converters operates in the continuous overcurrent mode for 30 seconds to allow any temporary fault to clear and then switches to the latching approach so that the converter is disabled. From a safety point of view, none of these methods is inherently better than any other, and you will find examples of each in the product selection of converter suppliers. However, one approach may be superior for your specific application, depending upon the diagnostics and repair philosophy for your system. In any event, it is good to be aware of all of the techniques so that you can make the best decision when selecting a converter.

2.2.15 Overvoltage

Like overcurrent protection, overvoltage protection is essentially a standard feature on all modern power converters. It serves to protect the load circuitry from excessive voltage levels in the event of a catastrophic failure in the power converter. In transformer-isolated converter topologies there are actually very few failure modes that will result in overvoltage conditions on the output, but overvoltage protection is still used as an extra margin of safety. Its use is more critical when using topologies such as the non-isolated buck where the failure of a semiconductor device can directly apply the

converter input voltage to the load circuitry on the output.

On low power converters, the simplest overvoltage implementation is a zener diode clamp on the output which, when its voltage level is exceeded, will conduct enough current to activate the converter's overcurrent protection circuitry. A more common and sophisticated approach is to have a separate voltage detector monitoring the output voltage and set to activate at the desired overvoltage value. This detector then shuts down the converter through the control circuitry. This shutdown can be either latching or non-latching. Artesyn Technologies' EXA40 allows the system designer choose between a device with a latching OVP and one without this feature.

More important to the system designer than the type of overvoltage circuitry used is the trip point. The trip point, of course, must be low enough to protect the load circuitry. However there is a danger in setting it too low. It must be higher than the worst case output voltage, including effects of line/load regulation, temperature effects, ripple and long-term drift and aging. There must also be allowance made for the output voltage deviations due to dynamic loading, as described in section 2.2.2. Neglecting to account for this is the most common reason for nuisance tripping of overvoltage circuitry.

2.2.16 Diagnostic and Control Features

The manufacturers of power converters offer a wide range of diagnostic and control features that we have not yet discussed. Unlike overvoltage and overcurrent, they are not essentially standard, and the mix of offerings will vary significantly from manufacturer to manufacturer and

from product to product. We will here make some brief comments on the types of features available and how they might be beneficial in some system environments.

Voltage Adjust, Voltage Trim and Margining are all terms used to describe user or system adjustability of the nominal output voltage set point of a power converter.

The three most common applications are:

- Generating a non-standard output voltage level for circuitry with unique requirements while retaining the benefits of using a standard converter.
- Compensating for output distribution drops for systems without using remote sensing.
- Verifying operation over a range of output voltage during system testing (margining).

The change in output voltage is achieved either by means of resistors external to the converter for user-defined adjustments or by connecting converter pins together to activate pre-set values of voltage adjustment. Voltage regulator modules intended for high-performance processors often use a four bit voltage programming bus. When using any trimming arrangement, the overvoltage (and possibly undervoltage) specifications of the converter must be considered so that they are not triggered by the voltage adjustment processes. Some converters, such as Artesyn Technologies' EXB30D series 30W dual positive output DC/DC converters, automatically adjust the overvoltage trip point to reflect the programmed value of output voltage, while others provide a trip point that allows a reasonable range for output voltage adjustment.

Selecting a Power Supply for System Application

Undervoltage or, conversely, **power good** is an output signal sometimes provided by a power converter to inform the user or system whether the converter's output voltage is in its normal operating range. This signal is more commonly used with AC/DC front-end converters than with distributed DC/DC supplies. The signal is also often used to drive a front panel LED display on front-end AC/DC converters. Most modern systems contain some equivalent output voltage monitoring at the system level, so this feature is often not needed or not used if provided.

2.2.17 Stability

Stability analysis determines whether the power converter will operate without oscillation or other abnormal characteristics over the complete range of line voltage, load currents and dynamic loading. It is especially important to do stability analysis when using remote sensed configurations with reactive distribution components inside the power converter's feedback loop. The most common evidence of stability considerations in a converter datasheet will be limitations on the amount of capacitance that can be placed on the output and the specified need for some value of capacitance on the input. It is essential that these recommendations be followed in order to achieve a robust power system that meets your performance expectations.

Historically, it has been expensive and difficult to do the proper system-level stability analysis because of the need for specialized software and analysis skills. Artesyn has addressed this problem by developing and making available a sophisticated stability analysis tool that power system designers can use to evaluate their systems. It is included in the SimScope™ tools that can be accessed

in the Virtual Power Lab™ on the Artesyn website (www.artesyn.com). The user can easily enter his design parameters for the supported Artesyn power converters, including the system output distribution system parameters, and SimScope will generate a stability analysis that includes phase and gain margin information. A complete operating guide and tutorial is provided for beginners. This service is provided free of charge.

2.2.18 Startup

Most converters will have a specification that defines their output voltage waveform during startup. Rather than turning on as a step function, it is desirable that the output voltage 'ramps up' over a controlled amount of time. This will allow any capacitance on the output to charge up gradually, reduce peak output currents, and avoid the possibility of tripping the converter's overcurrent limit during startup. This ramp-up time, or risetime, is usually on the order of 5 to 50ms for DC/DC converters and up to a second for high power AC/DC front-end converters. Using reasonable values of ramp-up time will usually make any voltage overshoot at turn-on nonexistent or very small.

In addition to the ramp-up time, there can be some sort of fixed delay internal to the converter before the output voltage starts to rise. This delay is defined by an additional specification usually called either 'turn-on delay' or 'start-up delay'. The exact definition will vary from supplier to supplier, but this value most often will include the ramp-up time in addition to the delay time. The generalized definition of these terms is shown in Figure 2.12.




Figure 2.12 - Converter Startup Definitions

2.2.19 Mechanical

There is a natural tendency to focus on the electrical characteristics of a power converter, but sometimes the mechanical parameters are equally important. The basic mechanical information that should be specified for all power converters includes the following:

- Footprint dimensions
- Height
- Weight
- Connector and pin-out information
- Mating connector part numbers (if applicable)
- Mounting provisions and/or restrictions
- Heatsink interface (if applicable)

Several of these parameters will have major implications for the manufacturing engineer as well as the system designer. Size and weight will largely determine whether the converter will be capable of automatic placement in an SMT environment. The converter's height will often

determine the ultimate card-to-card pitch (spacing) that can be achieved in the system.

2.2.20 Process Interface

In addition to the mechanical attributes of the converter as mentioned above, the better suppliers realize that the manufacturing process interface extends into areas such as **soldering profiles**, **solvent compatibility**, and **packaging** for efficient handling. As the market moves more towards increased use of small (<50W) surface-mountable DC/DC converters, there is increased attention toward specification of additional process interface information as a part of the datasheet.

This information will include soldering temperature profiles that are approved for use with the converter including, in some cases, the materials and temperatures used when the converter itself was manufactured. Also helpful will be specifications on compatibility with cleaning solvents, solderability specifications, and the specific JEDEC formats used when the converters are packaged for automatic placement. The leading suppliers of SMT DC/DC converters are now providing this data.

2.2.21 Environmental Specifications

Power converter datasheets now more commonly contain extensive information about the environmental conditions suitable for the product. Most of these specifications will actually apply more often to possible conditions experienced by the converter during shipping and handling of the converters and the final product than operational conditions inside the end-use equipment. One exception is **altitude**. The output power of some converters will need to be derated when operated at high

Selecting a Power Supply for System Application

altitude (> 3000 meters). Examples of environmental specifications that are most often encountered in a non-operational situation include **storage temperature, thermal shock, vibration, shock and humidity**. The leading manufacturers of power converters will list the environmental standards in the above areas that the converter has been designed to withstand.

2.2.21 Acoustics

Almost all DC/DC converters contain no moving parts and therefore are acoustically silent. Some AC/DC converters, especially those intended for usage as front-end supplies for DPA systems, will contain internal fans for forced-air cooling. These fans will have an acoustic noise output, which should be specified in the datasheet. The levels involved are normally low enough to not present any problem at the end product level.

2.3 System Design Checklist

We have covered a lot of diverse material in this chapter, and there will be more to come on much of it in later chapters. As an aid to the reader in focusing on the most important system design issues we present Fig. 2.13 - a system design checklist. Here, on one page, are the most important considerations to keep in mind when designing a distributed power system for a telecom environment. By checking off each item as you obtain the data, you will avoid overlooking what could be critical information for your system.

<input type="checkbox"/> Use Standard Products when Possible - Cost - Time to Market - Technical Risk	<input type="checkbox"/> Focus on Efficiency - Power Dissipation - Reliability - Cooling Requirements/Cost
<input type="checkbox"/> Do a Reliability Prediction - At operating temperature!	<input type="checkbox"/> Understand the Thermal Design - Ambient Temperatures - Airflow
<input type="checkbox"/> Consider your design Margins - Reliability vs. Density	<input type="checkbox"/> Do a complete DC Distribution Design - Decoupling - Sensing Strategy - Contact and Distribution Resistance
<input type="checkbox"/> Don't forget the Input Circuit - Hold-up Time - Filtering - Surge Current - Hot Plugging	<input type="checkbox"/> Get Mechanical and Process Data - Height / Weight - Solder Profiles
<input type="checkbox"/> Consider Second Source Strategy - Technical Risk	<input type="checkbox"/> Minimize Number of Output Voltages - Reduced interaction / Technical Risk in Multi-output converters

Figure 2.13 - System Design Check List