

# Chapter 8: *Reliability and Availability*

*Reliability is critical to many modern electronic systems. Applications such as routers and wireless base stations generate revenue for the ISP or carrier for only as long as they are available. This chapter details the various methods used to determine the reliability of the power supply and how this can be correlated to the application. Finally, there is a section on how to provide high availability redundant power solutions for particular system needs.*

## **Reliability and Availability**

### *8.1 Introduction*

Reliability has come to be one of the highest priorities in the design of a power system. It ranks along with cost and efficiency as a measure of successful design. This focus on reliability has developed because the custom centralized power supply was formerly one of the least reliable components of the system. The large component count, unique component types, high internal stress levels, complex design, manual assembly and abundant opportunities for design errors led to high overall failure rates. With today's emphasis on DPA power systems and usage of standardized converters, many of these factors are no longer relevant or are much better understood. As a result, today's power converters can be extremely reliable. There still is, however, a need for careful selection and application of power converters in order to arrive at a final power system that meets your reliability requirements.

In this chapter we will first address the basics of reliability, including definitions, units of measurement, infant mortality, wear out mechanisms, and factors that contribute to component failures. We will then discuss various methods of estimating, specifying and measuring power converter reliability. Artesyn's focus on reliability in the design of its converters will be described, including conservative design, component quality, derating practices and reliability demonstration testing. We will then discuss methods that you, as a power system designer, can use to increase the reliability of the total power system. Finally, the concept of system availability and fault tolerant design will be introduced. After a generalized discussion of how "N + M" techniques can be used to differentiate availability from reliability, we will give specific examples of high-availability power system design, including powerline fault immunity, redundant front-end modules and techniques for operating through failures of DC/DC converters.

## Reliability and Availability

### 8.2 Basics of Reliability

**Definitions** - A power converter is composed of a collection of electronic components, materials and manufacturing processes. None of these elements are perfect, and each can "fail" in some fashion that could jeopardize the functionality of the converter and perhaps the entire system. Reliability is a measure of how often these failures will occur. In electronic reliability analysis it is assumed that the failures are **independent** and **random**. This means that a failure in one component, even though it may cause the system to malfunction, will not cause other components to fail and that the failures are distributed in time according to an exponential statistical distribution with a constant failure rate vs. time. The **failure rate**,  $\lambda$ , is the fundamental variable that defines reliability.  $\lambda$  is expressed in terms of failures per unit of time. The time interval used for the definition is arbitrary, and we will shortly describe some of the more common units used in the industry. For now we will assume that the failure rate,  $\lambda$ , is expressed in terms of failures per hour (F/hr).

We will also assume for now that the failure rate strictly meets the above definition and is indeed constant with time, as shown in Figure 8.1. Since we are generally dealing with reliable components, and expect a low probability of failure in any given hour,  $\lambda$  will be a very small number. In the figure, we have shown the failure rate of a single component that is between  $10^{-8}$  and  $10^{-7}$  F/hr. Note that  $\lambda$  is a straight line and that the failure rate is the same over the complete time range shown. This is known as the **Failure Rate**, and the goal in electronic design is to operate components in this fashion so that the failure rate is known and low. The failure rate is a function of the applied stress levels.

Since the failure rate is constant with time, (see Figure 8.1), the probability of a failure does not change with the age of the product. But, of course, the longer the time span over which we expect the product to function, the higher the probability of a failure occurring. For example, the probability of a failure occurring between  $t_1$  and  $t_2$  is greater than between  $t_0$  and  $t_1$  because  $t_2-t_1$  is longer than  $t_1-t_0$ , and not because  $t_2$  is later in time than  $t_0$ . The probability of survival over a period of time is given by the **reliability**,  $R(t)$ .  $R(t)$  is related to  $\lambda$  by the following relationship:

$$R(t) = e^{-\lambda t}$$

Equation 8.1

where  $\lambda$  = Intrinsic Failure Rate (F/t)  
 $t$  = time

That is, at the beginning of the time period,  $R(t)$  will be high indicating a high probability of survival and a high reliability. As the time span increases, and we accumulate more unit-hours,  $R(t)$  decreases and the reliability decreases. This is shown in the plot of  $R(t)$  vs time in Figure 8.2. This is the expected negative exponential curve. The assumed failure rate in Figure 8.2 is  $10^{-6}$  F/hr. During a span of one million hours operation for example, the probability that the component is still operational is about 0.37. That is, we can see that for a small time span (unit hours)  $R(t)$  is very high but decreases with increasing time span.

**Units** - The fundamental unit of failure rate, failures/hour (F/hr), is not often used in practice due to its very small value for most normal components and systems. Instead,  $\lambda$  is more often expressed in terms of one of the following units.

- Failures/Million Hours = F/Mhr = F/hr x 10<sup>6</sup>
- Failures per Billion Hours = Failures In Time = FIT = F/hr x 10<sup>9</sup>
- % Failures per Thousand Hours = %F/khr = F/hr x 10<sup>5</sup>

The preferred unit from the above list for expressing IFR will vary from industry to industry and from company to company. It is fortunate that conversion from one unit to another is done easily by just multiplying or dividing by a power of ten.

It is also possible to express the reliability of a component or system by using the reciprocal of  $\lambda$ . Since  $\lambda$  is in units of failures per unit of time, the reciprocal will have units of time between failures. The technically correct term for the reciprocal is **Mean Time To Fail (MTTF)**. A more frequently used term is **Mean Time Between Failures (MTBF)**. The differentiation between them will be explained in the availability section, but for most practical systems these terms are very close to each other in value and can be used interchangeably. For now we will use the term "m" to describe either. m can be expressed in any unit of time, the most common being hours or years. To summarize:

$$\text{Failure rate} = \lambda = \frac{1}{m}$$

*Equation 8.2*

where  $m = \text{MTTF} \approx \text{MTBF}$

For example, if a power converter has a failure rate of 3000 FIT,

$$\lambda = 3 \times 10^{-6} \text{ F/hr and } m = 333,333 \text{ hr or } 38.05 \text{ yrs}$$

Equations 8.1 and 8.2 can be combined, yielding:

$$R(t) = e^{-\frac{t}{m}}$$

*Equation 8.3*

and at time  $t = m$ ,  $R(t) = e^{-1} = 0.37$

Thus, there is a 37% chance that the device will survive to a time equal to its MTBF - but only if there are no wear out mechanisms, which will be discussed below. The MTBF point is shown in Figures 8.2 and 8.3.

In the power supply industry, it has become traditional to express reliability in terms of MTBF with units of either hours or years. Artesyn follows this tradition in its power converter specifications. The MTBF specification is useful in that the reliability of products from different suppliers can be directly compared. Note that MTBF and failure rate differ in two important ways:

- Failure rates can be directly added to calculate the overall failure rate of multiple component systems. This can not be done when using MTBF.
- A MTBF of, for example, 100 years seems to imply that the component will operate for 100 years. We will see shortly that this is not true. This can create considerable confusion.

It is therefore recommended that any reliability data that is expressed in terms of MTBF first be converted to a failure rate before trying to do any reliability calculations or analysis. This policy will make calculations easier and also eliminate some sources of confusion. As an example, we will calculate the overall failure rate of a circuit that contains the following components:

## Reliability and Availability

Component	Quantity	Reliability Spec
Resistor	10	50 FIT
Capacitor	5	MTBF = 5,000,000
FET	1	0.03%/Khr

Table 1

We can calculate the total circuit reliability by first converting each reliability specification to a common failure rate unit - we will use FIT. These failure rates can then be directly added to obtain the failure rate of the circuit.

Component	Qty	FIT EACH	Total FIT
Resistor	10	50	500
Capacitor	5	200	1000
FET	1	300	300

Total for Circuit 1800 FIT

Table 2

So the failure rate for the overall circuit is 1800 FIT or  $1.8 \times 10^{-6}$  F/hr




Figure 1 - Failure Rate,

**Wear Out** - So far we have assumed that the failure rate is constant with time as shown in Figure 8.1. In actuality this is not the case, as all components and systems will have some type of mechanism that places a limit on their useful life. This phenomenon is referred to as **wear out**. When the component or system starts to approach wear out, the failure rate will increase, signifying the end of the **useful life**, **service life** or **lifetime** of the device.




Figure 2 - Survival Probability - Linear Time Base

Fortunately, most electronic components have extremely long lifetimes (over 20 years), usually well in excess of the effective economic life of the end equipment. If this is the case, then the simplifying assumption of constant failure rate is valid and can be used to model the expected failures for the system.




Figure 3 - Survival Probability - Logarithmic Time Base

There are, however, components that have lifetimes short enough to affect the reliability of the end system. Some examples are:

- Fans (the bearings wear out)
- Air Filters (require service or replacement)
- Electrolytic Capacitors (function of type and application design)
- Contacts / Connectors (number of insertions)
- Batteries (Energy depletion)

Some approximate lifetimes for common electronic components are shown in Figure 8.4. These should be considered only very generalized values, as they will depend greatly upon how the devices or components are rated and stressed. The manufacturing process also

contributes to possible service life considerations. Long-term effects of chemical interactions between materials and thermal cycling stresses affect printed circuit boards and the component interconnection technology and failure may remit. Using power converter suppliers that have well documented manufacturing processes and an emphasis on quality can minimize these risks.




Figure 4 - Typical Component Service Life

It is important to note that there is no connection between lifetime and failure rate: they are independent parameters. Failure rate is measured during the normal lifetime, before the wear-out period begins. A component such as a fan, with a limited lifetime of perhaps 10,000 hours, can be extremely reliable during its useful service life, with a MTBF of 100,000 hours or more. This is an example of why using MTBF to express reliability can create confusion. It is less confusing to state that the fan has a failure rate of  $10^{-5}$  F/hr and a service life of 10,000 hours.

**Infant Mortality** - The failure rate is not always constant at the beginning of a component's life. If the component is fresh off of the manufacturing line, it can have a higher failure rate than it will have later in its life. This higher

## Reliability and Availability

initial failure rate is due to manufacturing or material defects that will surface early in the component's life and is referred to as **infant mortality failure**. Systems can also incorporate a higher early life failure rate due to the manufacturing processes that interconnect the components. One example is "cold" solder joints opening after a few heating and cooling cycles during the powering up and down of the product. The infant mortality period will vary with the type of component or system, but is usually in the range of 5 to 500 hours. A proven technique for "eliminating" the infant mortality period is **burn-in**. The component or system is operated in the factory for a period of time so that any infant mortality failures occur before the product is shipped to the customer. The time period to accomplish the burn-in can be reduced by applying various stress levels and thermal cycling profiles to the product.

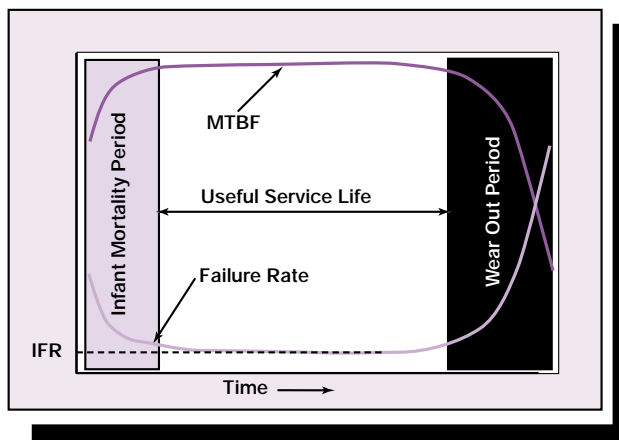


Figure 5 - The Bathtub Curve

If both the infant mortality and service life effects are included along with the failure rate, the overall failure rate vs. time characteristic appears as shown in Figure 8.5. This is sometimes referred to as the "**bathtub curve**". As seen in the figure, the MTBF characteristic will be an inverted curve. The flat portion of the bathtub

curve represents the failure rate during the component or system's useful life. It is this failure rate that should be used for reliability analysis during the product's operating life. As we will see next, this failure rate can be changed by controlling the operating conditions and stresses applied to the component or system.

**Factors that affect Reliability** - The failure rate period is the most consistently reliable portion of the component's life. But even this failure rate can be modified - in either direction - by changing the stresses applied to the component. Some component and converter stresses that are under the control of the power converter and power system designer include:

- Steady-state operating voltage
- Steady-state operating current
- Voltage and current transients
- Shock and vibration
- Humidity and other environments
- Temperature cycling
- Steady-state operating temperature

Of these factors, operating temperature is especially important. It has a major influence on the overall reliability and is also relatively easy to control by the power system designer. Consequently, it is important for all electronics designers to understand the relationship between temperature and failure rate.

The relative change in failure rate, also called the **acceleration factor**, when the operating temperature of a component is changed is given by the Arrhenius relationship:

$$AF = \frac{e^{\frac{\alpha}{kt_1}}}{e^{\frac{\alpha}{kt_2}}}$$

Equation 8.4

where: AF = Acceleration Factor  
 a = Activation Energy in Electron Volts  
 k = Boltzmann's constant =  $8.617 \times 10^{-5}$  eV/K  
 $t_1$  = Initial Temperature in degrees Kelvin  
 $t_2$  = New Temperature in degrees Kelvin

The activation energy is a factor that models the temperature vs. failure rate characteristics of a component. It will vary somewhat from one component type to another. Therefore, to model an overall sub-system such as a power converter, the types and mix of components typically used must be considered and an AF selected that represents an overall characteristic for the converter. 0.55eV is typically used for switching power converters and shows a reasonable correlation with observed measurements. If we use the 0.55eV value for activation energy and solve Equation 8.4 for the relative reliability between 25°C and 35°C (298K to 308K), we see that the acceleration factor is 2.0. That is, the failure rate at 35°C will be twice that at 25°C.

If the relationship in Equation 8.4 is solved over a larger range of temperature and normalized to 25°C, we get the relationship shown in Figure 8.6. This can be a very useful curve for the power system designer, since the reliability of many power converters is specified at a temperature of 25°C. For example, if the converter is specified with a reliability of  $10^6$  F/Mhr (1 million hour

MTBF) at 25°C then the expected failure rate when operating at 50°C would be approximately  $5 \times 10^5$  F/Mhr. Using Equation 8.4, a similar curve can be generated normalized to the base temperature of your choice.

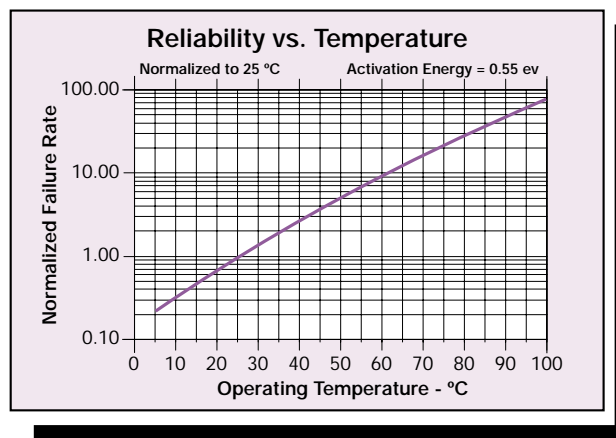


Figure 8.6 - Operating Temperature - °C

**System Impact of Failure Rate** - As we have noted, modern power converters achieve very low failure rates. That being the case, why is there so much emphasis on reliability in power system design? A simple example can answer this question. We will assume that a DC/DC converter with a MTBF of 1,000,000 hours (Failure Rate =  $10^6$  F/Mhr) is used in a system application. This corresponds to a MTBF of 114 years, well beyond the useful service life of the end equipment. The corresponding failure rate will be  $10^{-6}$  F/hr, a very low number. We will also assume that 10 of these converters are used in each system and that 10,000 systems are produced annually over a period of 5 years. Consequently, there will ultimately be a total of 500,000 converters in the field. We will assume that the product is operated continuously. The expected number of field failures per year can then be calculated as:

$$500,000 \text{ units} \times 10^6 \text{ F/hr} \times 8760 \text{ hr/yr} = 4380 \text{ Failures per year}$$

## Reliability and Availability

From a total system perspective, a converter with a 114 year MTBF can result in over 4,000 failures per year in a large population. Thus, even though an individual component or sub-assembly is very reliable, the cumulative effect of operating many of them at the same time can result in a field service situation that is significant. This is why the power system designer should put a high priority on reliability in design, both in selecting the most reliable products and applying them in a way that does not detract from their inherent reliability.

### 8.3 Reliability Prediction and Measurement

Since reliability is so crucial to a successful system design, it is important to be able to quantify it. This is not as straightforward as measuring voltage or current or temperature. There is no method for directly measuring the absolute value of reliability of a circuit. There is no "reliability meter" that can be attached to a circuit and give a real-time readout. Instead, the reliability assessment must be made either by some kind of prediction or by a measurement made over a considerable period of time. In this section we will briefly discuss the most common types of reliability prediction or measurement.

Prediction methods can be accomplished in a short period of time during the design phase of the circuit. Some common prediction methodologies are:

- Parts Stress Methods
- Private Database
- Parts Count Method
- Past Experience

Actual measurement of reliability takes much longer, and must be done after the circuit design has been completed and there are at least prototype production parts available. The two most common types of reliability measurement are:

- Accelerated Reliability Demonstration Testing
- Field History

**Reliability Prediction Methods** - Prediction methods utilize experience with past designs, assumptions and estimates, and knowledge of the design in question to formulate a supportable opinion as to the design's expected failure rate. These techniques are very valuable because they can be done during the design process rather than waiting until the new design is in production. They can also be very economical when compared to time consuming reliability measurement methods.

The most basic prediction technique is **experience with past designs**. If a power supply company is designing a new product that is very similar to an older product, it can be appropriate to assume that the reliability of the new product will be similar to that of the older one. This approach is only valid if the same types of components and manufacturing methodologies are used and if all the stress levels are the same. If this is the case, the reliability history of the older design (obtained by reliability demonstration testing or field history) can be used as an estimate of the new design's reliability.

The **parts count method** provides for additional sophistication in the reliability prediction analysis. All components used in the design are classified into groups. For example, some typical groups could be 0.1

W resistors, ceramic capacitors, small-signal transistors and diodes, power FETS, and connector pins. Each of these categories is assigned a failure rate that is determined by the type of component and the typical stresses applied to components in that group by the company's design guidelines. The reliability estimate for the overall product design is then determined by summing the assumed failure rates for each component used. This can be done in a matter of minutes concurrently with the design process as the design progresses and changes. As such, it can be a "real-time" indicator of the effect of proposed design changes on the product's reliability. As with all prediction methods, the parts count methodology is very well suited to making comparisons of the expected reliability between competing designs.

An additional level of detail can be added by using one of the **parts stress** prediction methods. These methods assign a specific "baseline" failure rate for each of the dozens of types of components used in electronic design. This baseline failure rate is then modified by the application of multiplier factors for each of several considerations involving the specific component and the stresses applied in the actual application. These factors can include:

- Quality of the component
- Voltage/Current stresses applied
- Operating Temperature
- Mechanical Stress Environment (Shock, Vibration, etc.)
- Complexity of design
- Construction Techniques

The three most commonly used parts stress methods are those developed by the US Military Agencies (MIL-HDBK-217), the US Telecom Industry (Bellcore 332) and the British Telecom Industry (HRD5). While the general methodology is consistent between these three agencies, the component databases and stress factors differ so that they will give three different predictions for the same design. A slight variant of this approach, used by some manufacturers of power converters, is to use some kind of **private database** for the baseline component failure rates. Artesyn Technologies does not recommend the use of private databases for accurate, comparable reliability measurements. It is extremely costly and time-consuming to develop a proper statistical database. Also, the database must be consistently and exclusively employed by the owner or the results are meaningless. The only true measure of reliability is derived from demonstrated reliability testing and this is the method employed by Artesyn Technologies in its latest power converter designs.

There are two inherent problems with using parts stress methods for reliability prediction. First of all, they are very time-consuming to accomplish. Each of the sometimes hundreds of components in a design must be assessed in terms of its specific operating conditions - i.e. voltage, current, temperature, interconnections, etc. Secondly, the databases used in the parts stress methodologies tend to be outdated for modern power converter design. Every month new integrated circuits, capacitors, FETS, thermal materials, and magnetic packages are released. It takes years for these design options to show up in the parts stress databases. In the meantime, unsupported estimates must be used for the newer parts. As a consequence, the designer has to choose between using outdated components and processes and having a less efficient design or doing a

## Reliability and Availability

state-of-the-art design and not being able to do a detailed reliability prediction. The same is true when using the parts count method.

The confidence level in the test is another important factor when comparing failure rates of different power converters - the higher the confidence level the more trustworthy the figure. For example, it's easy to predict a zero failure rate (infinite MTBF) if we use a 0.0 confidence level!

**Reliability Measurement Methods** - We now turn our attention to actual measurement of reliability as opposed to prediction. These methods must wait until after production hardware is available and tend to be time consuming, but are more accurate than the prediction methods discussed above. They are valuable as a source of historical data for application to future designs that are similar in nature.

The most basic method of measuring reliability is to record each failure as it occurs in the end product. If the total operating hours for the product are also known, the failure rate can then be calculated by dividing the number of failures by the total operating hours. This method is referred to as the field history method. There are, however, a number of practical barriers to using this method.

- Return of field failure reports tends to be incomplete. Repair personnel are much more interested in fixing the problem and getting the customer back on-line than they are in filling out a failure report form.
- For many types of equipment, it is difficult to assess the actual operating hours.

- The temperature stresses on the device are difficult to assess, since each customer may have a unique operating environment in terms of external temperature.
- Failures are often not analyzed correctly by field personnel. Repair strategies sometimes call for replacement of multiple parts to fix a problem rather than taking the time for analysis to a specific part. Sometimes parts are replaced unnecessarily.
- Some parts fail because of overstress from a different failure. These "chain reaction" type of faults should not legitimately be counted in the failure rate of the part with the secondary failure. Unfortunately, the failure reporting mechanisms are not usually sufficiently sophisticated to make this kind of differentiation.
- It takes a long time - at least a year - to build a valid field history.

If the above difficulties can be minimized, the field history method can provide very useful data. Large vertically integrated companies that provide field service on equipment designed by their own engineering departments have achieved the most success with this method. Some examples are large computer and telecommunications equipment manufacturers. The data obtained from studies by such companies provides a valuable source of feedback that can be used to "calibrate" the prediction methods to the actual failure rates experienced in the field.

A second reliability measurement technique is **reliability demonstration testing**. In this approach, the manufacturer of the device being tested puts a fixed number of the same device on test for a controlled period of time. This can be done under much more

controlled conditions than the field history method. With power converters, for example, such variables as load current, input voltage and operating temperature can be precisely defined along with the actual number of operating hours. With the low failure rates of today's converters, the test plan would need to either use a very large number of converters (very expensive) or operate them for a large number of hours (very time-consuming).

The real value of a reliability demonstration test is to get results in a timely fashion so that the reliability estimate is available soon after the product's introduction into the marketplace. This has led to the introduction of a method called **accelerated reliability demonstration testing**. Accelerated testing takes advantage of the failure rate vs. operating temperature characteristic defined in Equation 8.4. By operating the converters at elevated temperatures (but still within their maximum operating temperature rating), the time required to provide a statistically valid failure rate measurement can be reduced.

As an example of the time reduction possible with acceleration, consider the case of verifying the reliability of a converter with an estimated MTBF of 1,000,000 hours (Failure Rate =  $10^{-6}$ F/hr) at 25°C. We will assume that the manufacturer wants to limit the quantity under test to 100 and to test for a period that will result in 3 expected failures. The test time at 25°C can then be calculated as:

$$\text{Test Time} = 3 \text{ Failures} / (100) (10^{-6} \text{F/hr}) = 30,000 \text{hr} = 3.4 \text{ years}$$

This is clearly too long to provide any assurance to

customers who want to purchase this product soon after its introduction!

Now consider the test time if the test is conducted at a temperature of 80°C. From Equation 8.4 and Figure 8.6, we see that the normalized failure rate at 80°C is about 29. This can be thought of as the acceleration factor for the reliability demonstration test. The test time now becomes:

$$\text{Test Time} = 3 \text{ Failures} / (100) (29 \times 10^{-6} \text{F/hr}) = 1,034 \text{ hr} = 43 \text{ days}$$

43 days is a time frame that is feasible, both for the manufacturer and the customer. The use of acceleration makes the reliability demonstration test a practical alternative.

The above example was somewhat extreme for purposes of illustration, and smaller acceleration factors (between 2 and 5) are used in most power converter testing. The acceleration factor selection is a trade-off between the time and expense of the test and the risk of inaccuracy of the failure rate vs. temperature projection (activation energy). The test plans are formulated based on industry accepted statistical standards to provide results accurate to a predetermined confidence level. Artesyn, for example, uses a **Probability Ratio Sequential Test Plan (PRST)** as defined in **MIL-HDBK-781**. The test plan has predetermined acceptance and rejection criteria based on the hypothesis that the converter meets its reliability claims. Figure 8.7 depicts a typical test plan and shows three possible scenarios. In the plan shown, the claimed MTBF will be accepted if either 1.8 Mhr is reached with no failures, 3.1 Mhr is reached with one failure or 4.6 Mhr is reached with two failures. If three failures occur before 4.6 Mhr, the reliability claim is rejected. In the case of acceptance, it is possible to continue the test for

## Reliability and Availability

additional time and verify a claim for a lower failure rate than the original hypothesis.

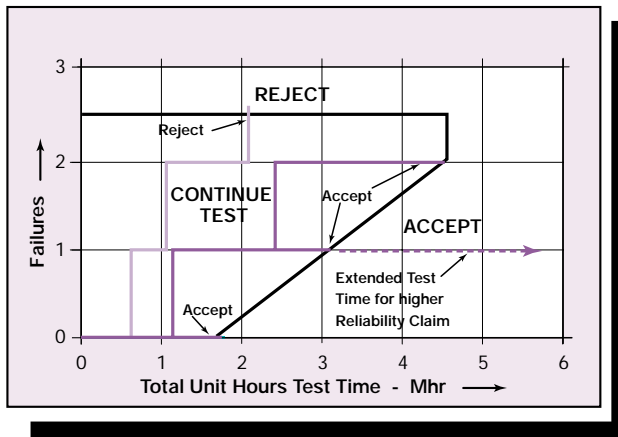


Figure 8.7 - Typical Reliability Demonstration Test Criteria

Even with accelerated techniques, the commitment of expense and time for qualifying a power converter is non-trivial. Consequently, most suppliers will not undertake this process. Artesyn has begun a program of accelerated reliability demonstration for most of its DC/DC converter products because this provides an additional source of information and confidence for its customers.

**Practical Approaches** - As we have seen, obtaining reliability data is not easy, fast or inexpensive. As a consequence, some practical approaches have been developed over the years that allow power converter and power system designers to proceed with their designs with a fair amount of confidence as to the reliability of the end product. During the design phase, the parts count method is a valuable method for accomplishing "real-time" assessments. This is especially true for established companies that have extensive prior experience with their component sets and know the typical stress levels applied and the areas needing

special attention.

Many suppliers of power conversion products, including Artesyn, use the parts stress methods. If they are properly applied, these methods will result in viable reliability estimates. The answers obtained, however, will be strongly influenced by the stress factors applied during the estimation process. Therefore, you need to know the assumptions used when doing the estimate before you can directly compare the result with reliability claims of other products.

Some "rules of thumb" have developed over the years that can be useful guidelines for rough reliability assessment and comparison. If MIL-HDBK and Bellcore reliability estimates are done for the same design using the same assumptions and stress levels, the Bellcore estimate will generally give a lower failure rate - perhaps 0.25 to 0.33 of the MIL-HDBK value. Similarly the HRD4 analysis will also show a lower failure rate than the MIL-HDBK, and should be somewhat close to the Bellcore estimate. All of these methods tend to be conservative (when properly applied) compared to actual field reliability. It is not uncommon for field history to give results up to an order of magnitude higher in reliability (lower failure rate) than a MIL-HDBK estimate. Knowing this, some companies will use less conservative factors when doing the parts stress estimate. Thus, it is even more important to inquire about the assumptions and stress factors used in any parts stress method analyses.

The trade-offs between time expended and expense are summarized in Figure 8.8. The accelerated reliability demonstration test is especially valuable, as it gives a more accurate "answer" than the design-phase estimation methods and yet has this information available

during the early part of the product's manufacturing cycle.

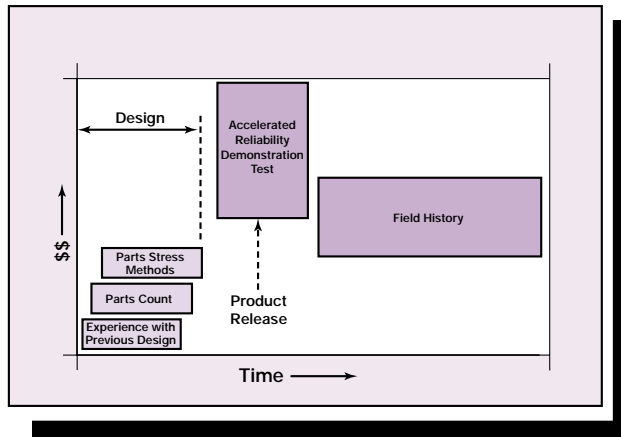


Figure 8.8 - Reliability Prediction and Measurement Techniques

#### 8.4 Artesyn's Reliability Focus

In this section, rather than discussing reliability in general, we will address the more specific actions that Artesyn takes to insure that its products are highly reliable and its customers have access to the information needed to properly analyze the reliability of their overall power system. This process begins with initial design and extends through the manufacturing process, reliability testing and customer field support.

We have seen how important conservative design and component derating is to achieving a reliable design. Artesyn engineers put significant emphasis on this. The electrical and thermal stresses on each component are evaluated and kept in compliance with conservative design standards. Voltage and current stresses on components are derated at least 20%. That is, under

worst case conditions, the electrical stresses applied to components will be at most 80% of the component rating. Tantalum capacitors are derated 45% so that they will not become a limitation on the overall reliability of the converter. Thermal stresses receive special focus due to the strong dependence of reliability on operating temperature as seen in Figure 8.6. Artesyn's design guidelines keep components operating well below their rated temperature limits. The chart in Figure 8.9 shows typical Artesyn thermal derating practices.

Component	Temperature Rating - °C	Temperature Operating - °C	Derating %
MOSFET	150	74.5	50.3
Diode	150	66.7	55.0
Planar PCB	130	44.0	66.0
Inductor	110	57.5	47.7
Film Capacitor	110	51.8	52.9
OptoCoupler	105	52.5	50.0

Ambient Temperature = +35°C

Figure 8.9 - Typical Artesyn Thermal Derating Practices

Artesyn selects high quality components from well-known established suppliers so that the risk of unexpected component quality problems is minimized. The detailed knowledge of the component characteristics and ratings along with the conservative design guidelines results in rugged and robust converter designs with an excellent reliability record. The parts count method is used during the design process to monitor and compare the expected reliability. This is made possible due to Artesyn's prior experience with and knowledge of the reliability history of each component type when applied in accordance with their design practices. As the design nears completion, a detailed parts stress analysis is done using one or more of the MIL-HDBK, Bellcore 332 or HRD4 methodologies. The results of these analyses are

## Reliability and Availability

used to determine the claimed reliability in the datasheet. The assumed environment (usually Ground Benign), operating conditions and the operating temperature will be stated.

Note that Artesyn specifies reliability for some of its products at temperatures higher than 25°C. This is done so that the specified temperature is similar to the application temperature in its intended market. To compare these products with competing converters specified at 25°C, the reliability specifications must be converted to the same operating temperature. For example, the Artesyn EXB50 series DC/DC converter has a specified MTBF of 776,000 hours ( $1.29 \times 10^6$ ) at 40°C. Using Figure 8.6, we see that the relative failure rate factor between 25 and 40°C is about 2.9. Thus if the EXB50 is being compared with competitive converters specified at 25°C, the EXB50 MTBF at that temperature would be  $776,000 \times 2.9 = 2,250,000$  hours ( $4.4 \times 10^6$ ).

The manufacturing process is an important element in assuring that the predicted reliability is achieved in practice. A reliable design is no good if defects or opportunities for failure are introduced during assembly and test. Artesyn uses advanced manufacturing processes and the latest automated equipment to insure that all aspects of the product, including interconnections and packaging can achieve the highest reliability possible. Cleaning, soldering and washing processes receive extra focus and attention to eliminate long-term problems with interactions between materials that could prematurely limit the useful life of the product. The manufacturing process for most of Artesyn's products includes a burn-in cycle. This allows for extended electrical testing of the converter and also prevents any infant mortality type component failures from showing up

in our customer's product. The converter as received from Artesyn should meet its claimed reliability specifications.

Extensive stress testing of early prototypes is an important part of the Artesyn design and production cycle, and something not done by all converter suppliers. A sample size of at least 25 units, from two or more production runs, is typically used so that any component tolerancing problems can be identified. These units are subjected to dynamic testing so that the major usage parameters such as input voltage, output load and temperature can be independently cycled. This applies thousands of different combinations of operating conditions and can quickly identify any unexpected areas of instability or performance degradation. At least one unit is exposed to HALT testing, which imposes high levels of environmental stress. Several units are then electrically and thermally stressed to failure to identify the design and specification margins. The end result is that the final design has a proven ability to withstand stresses well in excess of its specified ratings. This will allow for excellent long-term reliability when the product is used within its specified limits.

In 1995, Artesyn Technologies began an extensive reliability demonstration testing program on most of its new DC/DC converter designs. The testing program will benefit both Artesyn and its customers. It will verify in a timely fashion, via actual test hours, the reliability specifications for each converter. This should increase customer confidence in using new converter designs in their equipment. Where practical, the demonstration testing will be continued beyond the "accept" criteria to allow for enhanced reliability ratings for some products. Upon inquiry, Artesyn will keep its customers informed as

to the test plans and the progress of the reliability demonstration testing activity. Artesyn will also publish demonstrated reliability data, as it becomes available, on its updated product datasheets. For example, the BXA30 series DC/DC converter now has demonstrated reliability of over 7,000,000 hours!

Artesyn is continuously developing new products and continuing reliability demonstration testing on recent products. We cannot, therefore, include up-to-date detailed reliability information on specific products in this document. The Artesyn website is a valuable source for the latest reliability information. The most recent datasheets can be downloaded, and there is detailed technical information available that focus on reliability analysis.

### *8.5 Achieving High Reliability in System Applications*

The power system designer can have a substantial influence on the reliability of the product. Both the selection of power converters and their operating environment will determine the overall failure rate that can be expected in the system application. The key factor is the individual component temperatures. These can be controlled by reducing their heat rise or by external cooling. In this section we will look at methods to control the component temperatures in the power converter and give an example of how each can affect the field failure rate and associated costs.

**Converter Selection** - An essential step in achieving high power system reliability is to select **high reliability**

**converters** for the main power conversion functions. You will need to rely upon the credibility of the converter suppliers and the datasheets and other information they supply to make this determination. It is hoped that the information supplied in the preceding section will be helpful in establishing Artesyn as one excellent choice for converters with predictable and low failure rates.

In making your converter selection, the importance of operating **efficiency** cannot be over-emphasized. Efficiency determines the power dissipated in the converter and consequently the temperature rise between ambient and the internal component and junction temperatures. As we saw in Equation 8.4 and Figure 8.6, the converter operating temperature will have a pronounced effect on the reliability of the converter in the system application. Consequently, a more efficient converter - even at a price premium - is almost always an excellent choice, with an overall system cost saving when the cost of failures is considered.

**Derating** - The converter manufacturer's estimate of failure rate will be predicated on an operating temperature and output power. While the values selected for these criteria have achieved some industry standardization, there is still wide variability between suppliers, and the final choice is somewhat arbitrary. If the converter is used at an output power or current different from that used for the reliability specification, the actual system reliability will be different to the prediction. This effect can be used to your advantage by operating the converter at an output power less than the maximum permissible value - a practice referred to as derating. The power dissipated in many internal components (including most of the high power devices in the power conversion chain) is at least proportional to the output

## Reliability and Availability

power or current. This means that derating the output current by  $x\%$  will reduce the component temperature rise by at least  $x\%$  and by more for components such as MOSFETS and magnetics where the power dissipation is proportional to the current to be squared. The lower power dissipation leads to a corresponding lower failure rate. Consequently, the power system designer can decrease the converter's failure rate to a desired value by reducing its output power.

The exact correlation between derating and failure rate cannot be calculated from datasheet information, since the operating temperature stress is not reduced uniformly on all internal components. The power dissipation of some of the converter circuit components, such as control and support circuitry, will not change with output power. The overall converter physical package will influence how dissipation from components affect others. Internal operating temperatures would also need to be known to accurately calculate the improvement in failure rate, and these temperatures would vary between component types. Suffice it to say that the reliability enhancement effects of derating can be substantial. The reduced power dissipation and lower internal operating temperatures will reduce component stresses and achieve a reduced failure rate and perhaps increased operational life.

**Operating Temperature** - Equally effective is decreasing the ambient operating temperature of the converter. With a fixed temperature rise from ambient to the converter's internal circuitry, this will reduce the average operating temperature of the converter's components and reduce the failure rate. The effect of reduced operating temperature on reliability is easy to calculate, and an example will be given shortly. There are several methods

for reducing the operating temperature. They include, increasing the airflow rate in forced convection systems and using less dense system packaging in free convection applications.

**Electrical Stresses** - So far we have addressed thermal stresses on the converter, but there are electrical stresses as well. The most common electrical stresses are in the form of transients on the input of the converter. Input transients were addressed in Chapter 5, along with techniques for minimizing them. It is important to use the appropriate diode, capacitor and filter networks on the input to prevent voltage excursions outside the maximum rated input voltage range of the converter. This will prevent damage and electrical stresses that could reduce the converter reliability. A summary of how the power system designer can enhance the system reliability is shown in the design checklist of Figure 8.10.

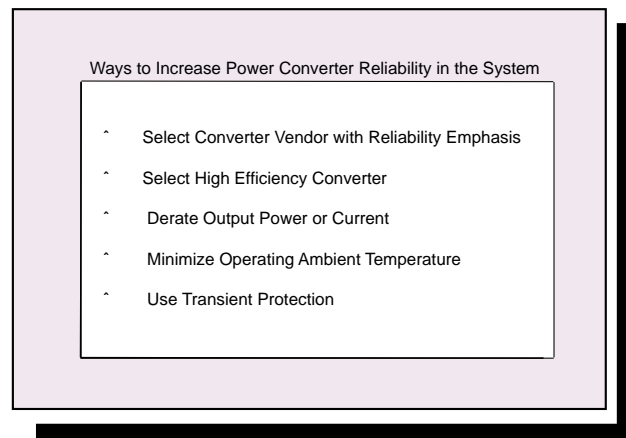


Figure 8.10 - High Reliability Design Checklist

**Design Example** - It is instructive to consider an example of how seemingly minor changes in the system thermal design can drastically influence the reliability and resulting costs of the overall system. We will use the Artesyn BXB100 DC/DC converter in this example, which builds on the results obtained in the example shown in Section 7.6 of Chapter 7. The converter used is a 48V

input 5V output device with a maximum output power of 100W (20A). The output power is derated to 75W in the example. The converter's efficiency is 83%, and its maximum baseplate temperature rating is 100°C. In the system application, it is cooled via forced convection with an ambient air temperature of 50°C at a velocity of 2 m/s (400 lfm).

There are alternative heatsinks available for this converter. These heatsinks differ in their height, with the higher heatsinks offering better thermal resistance performance. The baseline design for this application used the 0.45 inch high heatsink which has a thermal resistance specification of 2.4°C/W. From the 75W output power and the 83% efficiency, we can calculate the power dissipated in the converter to be 15.36W. The converter baseplate temperature in the application can be calculated as:

$$T_{\text{baseplate}} = 50 \text{ }^{\circ}\text{C} + (2.4^{\circ}\text{C/W}) (15.3 \text{ W}) = 86.9^{\circ}\text{C}$$

This is within the 100°C maximum rating and therefore is a valid design from a thermal point of view. But let's now consider the reliability implications. The BXB100 datasheet specifies the reliability at 1,400,000 hours MTBF ( $7.14 \times 10^{-7}$ ) at a baseplate temperature of 40°C at full load (calculated with Bellcore 332). As we pointed out earlier, the derating to 75% output power will result in improved reliability, but in this example we will be conservative and use the full load specified reliability value of  $7.14 \times 10^{-7}$  F/hr.

To determine the expected reliability in this system, we must convert the 40°C specification to an estimated value at 86.9°C. This is done using Equation 8.4, with an

activation energy of 0.55 eV,  $t_1 = 313^{\circ}\text{K}$  and  $t_2 = 359.9^{\circ}\text{K}$ . Solving Equation 8.4 with these values gives an acceleration factor of 14.26 and a resulting failure rate of  $1.02 \times 10^{-5}$  F/hr.

We will now calculate the effect on the converter's reliability of using a larger heatsink, with all other system parameters remaining the same. The next larger heatsink has a height of 0.95 inches and a thermal resistance of 1.65°C/W. The resulting baseplate temperature is calculated as 75.4°C, a reduction of 11.5°C. Again using Equation 8.4 with a  $t_2$  value of 348.4°K gives an acceleration factor of 7.94 and a failure rate of  $5.67 \times 10^{-6}$  F/hr - a 44% reduction.

A 44% reliability increase is certainly a good thing, but what does it mean in practical terms? One measure is the expected number of failures experienced during the system's operating life. If we assume a field population of 20,000 converters each operating 24 hours per day and a five year system service life the expected number of failures can be calculated as:

$$0.45\text{in. HS} = (20,000 \text{ units}) (5 \text{ yr}) (8760 \text{ hr/yr}) (1.02 \times 10^{-5} \text{ F/hr}) = 8935 \text{ Failures}$$

$$0.95\text{in. HS} = (20,000 \text{ units}) (5 \text{ yr}) (8760 \text{ hr/yr}) (5.67 \times 10^{-6} \text{ F/hr}) = 4967 \text{ Failures}$$

Thus merely using a larger heatsink results in 3968 fewer expected failures during the life of the system! This can have a profound economic impact. In Chapter 12 on value analysis, we will discuss the assessment of the cost impacts of reliability in more detail. Here we will make a simple assumption that each converter failure

## Reliability and Availability

has a "cost" of \$500 due to system downtime, repair actions, etc. With this assumption, the increased reliability of using the larger heatsink results in a cost saving of:

$$(3968 \text{ Failures}) (\$500/\text{Failure}) = 1.98 \text{ Million Dollars}$$

If this saving is allocated over the 20,000 converters in the field population it results in a value of \$99 per converter. Another way of looking at this is that the system designer should use the larger heatsink if it has a cost premium of less than \$99 over the smaller one. Since the cost premium is actually well less than \$1, the economic aspects of this decision should be a "no brainer"!

### 8.6 Basics of Availability and Fault Tolerance

High reliability, while essential, is no longer sufficient for many of today's electronics systems. Banking systems, reservation systems, automated manufacturing control systems and communication networks are now specified in terms of availability as well as reliability. Terms such as "five nines availability" and "24/7 operation" have become common in our vocabulary as the need for such high performance systems has become more completely understood and accepted. The power systems of these high-availability products must support their availability goals by features such as system operational survivability through power converter failures as well as immunity from outages in the AC input power source.

In this section we will cover the basic principles of availability and fault tolerance, including the definition of terms and the general operation of fault tolerant

architectures. The differentiation between redundancy and fault tolerant design will be explained. The explanations and examples will primarily be at a "block diagram" level and generic in nature rather than attempting to capture the intricacies of the detailed interface designs for power converters, which are addressed in the next section. The mathematics will be kept to a minimum, with intuitive understanding being a more important goal than mathematical preciseness.

Definitions - The most basic definition of availability is the ratio of actual service to required service:

$$\text{Availability} = \frac{\text{Actual Service}}{\text{Required Service}}$$

*Equation 8.5*

For example, if a system is required to operate continuously and it is out of commission due to repair of failures for 12 hours per year, its actual availability is:

$$\text{Availability} = (8760 - 12) \text{ hours} / 8760 \text{ hours} = 0.9986 = 99.86 \%$$

which would be approaching "3 nines" availability.

Availability can also be defined in a form that better lends itself to analysis using reliability terms. Consider that system unavailability is the result of a failure, and that the expected time to failure is the MTTF of the system. The system must then be repaired before it can become operational again. The average expected time interval for executing the repair is called the Mean Time To Repair (MTTR). The total expected outage time can then be considered to be the MTTR. The expected interval between outages then becomes the MTTF plus the MTTR, which is also referred to as the MTBF for the system. System availability can then be defined as:

$$\text{Availability} = \frac{\text{MTTF}}{\text{MTTF} + \text{MTTR}} = \frac{\text{MTTF}}{\text{MTBF}}$$

Equation 8.6

Thus even if the system has a finite failure rate, if the repair is instantaneous the system can have an availability of 100%. As we will see later, the intent of fault tolerant design is to achieve high availability by making the repair time appear to be essentially instantaneous.

As an example, consider a simple power system with one converter. The MTTF of the converter in the application is 500,000 hours. The average time to diagnose the failure, replace the converter and get the system operational again is 4 hours. The power system availability is then:

$$\text{Availability} = 500,000 / 500,004 = 99.9992 \%$$

Since the availability number for high availability systems will be a number typically beginning with several nines, it is often preferable to use **unavailability** as a criterion instead. Unavailability is defined as:

$$\text{Unavailability} = 1 - \text{Availability}$$

Equation 8.7

and is usually expressed in terms of time per year. In our previous example, the unavailability would be:

$$\text{Unavailability} = 0.0008 \% \times 8760 \text{ hours} = 252.3 \text{ seconds per year}$$

While this may seem like a small amount of outage time, some of today's systems have unavailability requirements

of milliseconds or even microseconds per year! To meet an unavailability requirement of 10 milliseconds per year with a repair time of 4 hours would require a MTTF of approximately 13 billion hours. Clearly this is not achievable even with high reliability design techniques. This situation leads to the use of redundancy and fault tolerant design techniques.

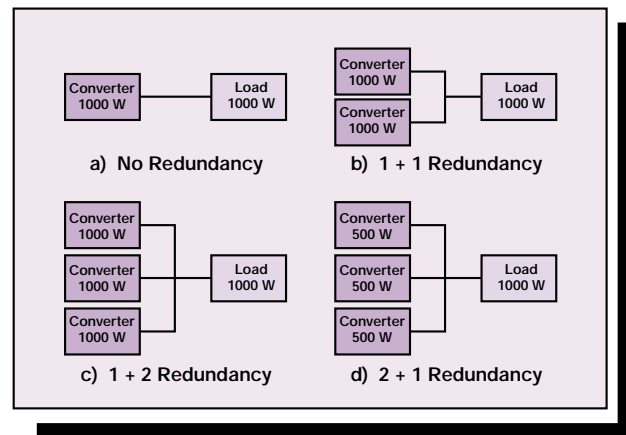


Figure 8.11 - Redundancy Techniques

**Redundancy** - This is the use of additional hardware assemblies in parallel so that one or more of them may fail without bringing down the system. For our discussion here, we will consider the hardware assembly to be a power converter. Figure 8.11 shows the basic principles of redundant architectures. A system load of 1000W is assumed in all the examples. A non-redundant system is shown in Figure 8.11(a), where a single 1000W converter is used to supply this load. Note that if there is a failure in the power converter, the system will go down.

In redundant architectures, the number of assemblies required for system operation is referred to as "N". The number of additional redundant or "spare" assemblies is referred to as "M". In Figure 8.11(b), for example there is an additional 1000W converter available. Since the

## Reliability and Availability

system can operate with one converter and there is one spare converter, this would be referred to as a "1 + 1" redundant system. A "1 + 2" redundant system is shown in Figure 8.11 (c). Here only one converter is required for operation but there are two spares available. Figure 8.11 (d) shows an architecture using 500W converters. In this case, a minimum of two converters will be required for system operation and there is one spare. Therefore, this is referred to as a "2 + 1" redundant system.

There are two possible ways to implement each of these redundant architectures. If all of the converters, including the spares, are operational at the same time it is referred to as an **active redundant system**. In this case, it is assumed that the load current is approximately equally shared by all of the operational converters. It is possible to design a system in which the "spare" converters are kept in a powered-down state until there is a failure of one of the "main" converters and then switched on. This approach is referred to as a **standby redundant system**. The standby approach, although having some theoretical advantage in terms of total system failures, is much more complex and intricate in terms of its implementation and is therefore not often used in the types of systems shown in Figure 8.11. The standby technique is commonly used, however, in some UPS systems. Unless stated otherwise, it is assumed that the active approach is used in the remainder of the systems and discussions in this chapter.

Note that some systems can meet their availability requirements even if they are down for service for substantial periods of time. For example, an automated retail checkout system may need very high reliability and availability for 12 hours per day when the store is open, but could be taken down for maintenance for several

hours each evening without compromising the system's availability objectives. In situations such as this, adding redundant power converters can increase the system availability, even though there is no provision for on-line diagnosis and repair, which can be accomplished during "off-service" hours. The positive benefits of this approach are shown in Figure 8.12. In this example, it is assumed that each power converter has a failure rate of  $2 \times 10^{-5}$  F/hr. The figure shows the reliability function for the system,  $R(t)$ , which can be thought of as the probability of system operability. The curve for the non-redundant approach is the same exponential result as was shown in Figure 8.2. The usage of one and two additional converters in "1 + 1" and "1 + 2" redundant architectures are depicted in the two blue curves. Note how the redundancy increases the system reliability and availability, even with no provision for on-line diagnosis and repair. This is an example of **high-availability system design**.

Note that even the high-availability design techniques shown in Figure 8.12, while improving the system reliability, do not result in a completely reliable and available system. Using **fault tolerant** approaches, however, can result in a system that appears to have a high reliability that is constant over the system's operational life, as shown in the green curve in Figure 8.12. This is accomplished by means of on-line fault diagnosis and repair which does not interrupt the operation of the system. We will explore fault tolerant design in more detail below. Note that for the fault tolerant approach shown in the figure, there will be power converter failures during the equipment's operational life but that they will not affect the system operability, which will appear to be constant.

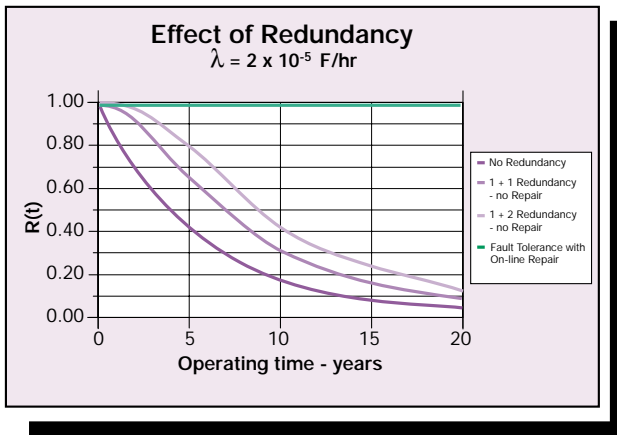


Figure 8.12 - Effective System Reliability vs. Redundancy Strategy

**Fault Tolerance** - Fault tolerant design is an extremely powerful tool, and perhaps the only way to achieve the high levels of availability needed for some continuously operating systems. The principles of fault tolerant design are relatively simple, but attention must be paid to many details. Also, fault tolerance is far from free or inexpensive. The designer of a fault tolerant system must be prepared for additional hardware expenditures and for a structured maintenance and repair program to implement the operational repair strategy. Each fault tolerant system will have different architectures and hardware partitioning as well as a unique repair strategy. There are, however, four capabilities that all fault tolerant designs must contain.

- **Redundancy** - In the event of a hardware failure, there must be at least one redundant unit to take its place until the repair can be affected.
- **Fault Isolation** - The failed unit must isolate itself from the system so that it does not create failures in other units or impair operation of the system. A common example in paralleled power converters is the use of "ORing" diodes.

- **Fault Detection and Annunciation** - After the occurrence of a failure, and before the repair is completed, the system will be operational (but at a reduced level of reliability due to fewer operational redundant units). Consequently, it is important to detect failures as early as possible so that the repair time is minimized. This involves two steps. First, the system must somehow sense that there has been a failure. Secondly, the sensed failure must be made known to the outside world. This could be in the form of a signal to a control computer or an alarm or light for observation by a human observer.
- **On-line Repair** - To maintain the system availability, the fault repair must be made without disrupting the system operation. This typically takes the form of having easy repair access to failing units and being able to do "hot plugging" of power converters without shutting down power to the system.

More detail on how to accomplish the above requirements in power systems will be given in the final section of this chapter.

**Characteristics of Fault Tolerant Systems** - The achievement of high levels of availability is not without disadvantages. Addition of redundancy will involve adding additional hardware in the form of excess converter capability. Control and diagnostic hardware will also be needed along with provisions for fault isolation. All these items add cost to the design, both in terms of hardware cost and in product development time. Fault tolerant systems are therefore almost always **more expensive**.

While the cost impact of fault tolerance is generally understood and intuitive, there is another characteristic

## Reliability and Availability

that does not have the same level of awareness, especially with the end-user of the equipment. Because a larger number of converters are used in a fault-tolerant system, and additional hardware is added for diagnostics and isolation. The fault tolerant system will have a **higher repair rate** than a similar non-redundant system. Consequently, there must be a very structured and deliberate repair and maintenance strategy in place. The beauty of fault tolerance is that these failures and resultant repair actions will not interfere with the operability of the system. In many cases, the repair can be accomplished by customer operating personnel rather than trained service technicians. But the additional failures do impose additional cost in terms of replacement units and field labor. This can also create a negative customer perception due to the seemingly frequent repair activity on a "high availability" system.

**Selection of N and M** - The mathematics determining the effects of N and M on the overall system availability are well understood and documented and will be summarized here. The prudent selection of these values, however, involves art and experience as well as science. Some of the design trade-offs involved in making this important decision will also be discussed.

The calculation of system availability as a function of N and M is done by using the discrete binomial distribution function:

$$A_{\text{Sys}} = \sum_{i=N}^{N+M} \binom{N+M}{i} A_{\text{Conv}}^i (1 - A_{\text{Conv}})^{N+M-i}$$

Equation 8.8

where:

$A_{\text{Sys}}$  = Availability of System

$A_{\text{Conv}}$  = Availability of Single Converter

A similar relationship is used when doing the calculations in terms of unavailability:

$$U_{\text{Sys}} = \sum_{i=M+1}^{N+M} \binom{N+M}{i} U_{\text{Conv}}^i (1 - U_{\text{Conv}})^{N+M-i}$$

Equation 8.9

where:

$U_{\text{Sys}}$  = Unavailability of System

$U_{\text{Conv}}$  = Unavailability of Single Converter

As an example, we will use the system discussed earlier with a converter MTTF of 500,000 hours and a repair time of 4 hours. We saw that  $A_{\text{Conv}}$  for this situation was 0.999992. If a redundant array of two such converters is used in a 1+1 configuration, the resulting system availability can be calculated from Equation 8.8 using a value of 1 for both N and M. The resulting system availability is 0.999999999936, 5 orders of magnitude higher than the availability of the non-redundant system, demonstrating the astonishing power of redundancy. Adding another redundant unit (1 + 2 redundancy) would increase the system availability by another 5 orders of magnitude, but this would be "overkill" in all but the most extreme applications.

Configuration	System Availability	Unavailability (ms/year)	Mean Time Between Unavailability (Hours)	Mean Time Between Repairs (Hours)
Non-Redundant	0.999992000000	252.5	500,000	500,000
1 + 1	0.999999999936	2.02	$6.25 \times 10^{10}$	250,000
2 + 1	0.999999999808	6.06	$2.08 \times 10^{10}$	166,667

Converter MTBF = 500,000 Hrs ( $5 \times 10^5$ )

Repair Time = 4 Hours

Figure 8.13 - Effects of Redundancy Partitioning on System Performance

Using three lower power converters in a "2+1" configuration results in a system availability of 0.999999999808, just slightly less than the "1 + 1" partitioning (assuming no change in the converter failure rates and repair times). Note that this partitioning will also result in lower mean time between repairs for the system. A comparison of results is shown in Figure 8.13. Note how the fault tolerant design increases system availability but at the cost of more system repair activity. Using the data from Figure 8.13, it would seem that the best choice for N would always be one, since this always gives the best system availability and a lower repair rate than using higher values of N. Using N = 1 also gives the highest minimum power derating for the converter (0.5), increasing its field reliability. But in practice, larger values of N are commonly used (up to 6 is not unusual). Why is this so? There are some practical benefits to using two or more power converters to provide the system load. These include:

- Smaller power units (less size penalty for the redundant unit or units)
- Reduction in cost of redundant power capacity
- Scalable power for entry-level systems

For example, with 1 + 1 redundancy, there is a 100 % size and cost penalty for the redundant converter. For a 3 + 1 redundant system, the penalty is 33%. Thus the basic trade-off comes down to cost and size and packaging flexibility vs. system availability and repair performance. The most common resolution is to use values of N between 1 and 6 and a value of 1 for M. Higher values of M are not normally needed to achieve the desired system availability performance.

## 8.7 High Availability Power System Design

In this section we will address some of the more pragmatic aspects of high availability and fault tolerant power system design. While the concepts discussed can apply to any type of power system architecture, they are most commonly used with distributed power architectures. Designing a fault tolerant power system using a centralized approach is an extremely complex and expensive undertaking, due mainly to load partitioning and multiple voltage DC distribution issues. Consequently, we will focus here on distributed power architectures, with emphasis on board-level distributed systems. This approach is used in the vast majority of fault tolerant system designs.

**Overall Architecture** - A typical fault tolerant board-level distributed power system uses three different techniques

## Reliability and Availability

to achieve its fault tolerant behavior. First, the board mounted DC/DC converters are partitioned and packaged with the loads themselves. This eliminates the need to provide redundant converters and redundant loads and to implement an interconnection scheme between them that allows for fault isolation. The load boards themselves (including their on-board DC/DC conversion functions) are redundant. A fault on one board (either load circuit or DC/DC converter) will cause that board to lose functionality, but not affect the remainder of the system. Multiple such boards are used so that the loss of one will not significantly affect the system performance or throughput during the time needed to replace the failed board. With this approach, there is no need for specialized detection, annunciation and isolation circuitry for the outputs of the DC/DC converters - the board-level functional diagnostics are sufficient to signal the need for replacement of the board and its self-contained power. Of course, there must be provision for isolating a DC/DC converter input failure, such as a short circuit, and preventing it pulling down the intermediate voltage bus. Techniques for doing this are discussed below.

Secondly, a N + 1 approach is normally used for the AC/DC converter section. This can easily be done by using one of the new standardized power shelf and modular AC/DC converter systems such as the Artesyn AFE series. These sub-systems include power factor correction, redundancy in your choice of several partitions, fault isolation by means of output ORing diodes, built-in fault detectors and indicators, and external access hot plugging of converter modules. They include all of the required capabilities for building a fault tolerant front-end system as part of their standard hardware and very little design effort is required by the power system developer.

The third aspect is providing for immunity from faults in the AC powerline. This is an extremely important consideration, since the failure rate for the powerline is much higher than that of any of the power conversion hardware. Published values of MTBF for the powerline are as low as 200 hours. In central office Telecom applications, this immunity is achieved by using the centralized -48V battery bank as an input source to the system. In other systems, the powerline immunity is more complex to implement, but still very feasible when using a distributed power architecture. The immunity is achieved by different methods, depending upon the duration of the powerline outage. Fortunately, most outages are short in duration (less than 10 minutes) so that all of these techniques are not needed in all systems.

Short duration outages (less than one cycle) are handled by means of using sufficient hold-up capacitance on the AC/DC converters and intermediate voltage bus. Longer duration outages (up to an hour or so) are addressed with either a battery backup system on the intermediate voltage bus internal to the system or an external battery powered UPS system. The battery backup alternative is the preferred approach for most applications, as it is a lower cost and higher reliability solution. For powerline outages of extreme duration (hours or days), localized generation techniques such as petroleum powered generators, microturbines or fuel cells are sometimes needed. Figure 8.14 summarizes the fault tolerance aspects of a typical board-level distributed power system.

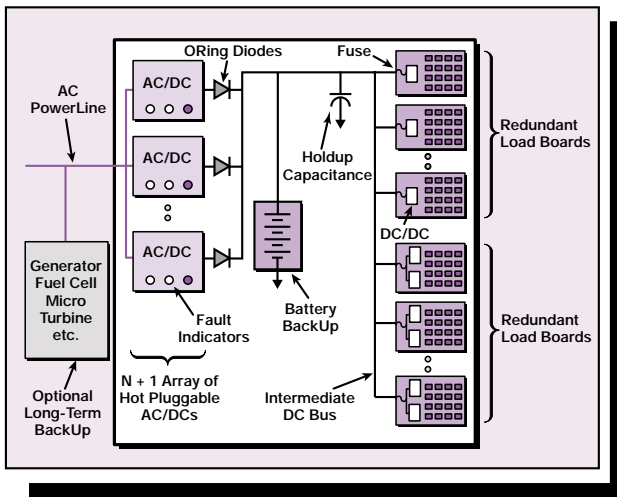


Figure 8.14 - Typical Fault Tolerant Power Architecture

**Isolation of Failed Unit** - We have already seen the power of redundancy in reducing system down time. All of this is for naught if a failure in one unit causes additional failures in units of the same or different type. If this happens, the original, single failure really becomes multiple failures that can easily cause a system malfunction. This shows that redundancy is necessary, but not sufficient to create a high availability or fault tolerant system. Also necessary is that a fault in any one unit be contained, or isolated, to within the failing unit. A failing power unit can cause the host system to malfunction in a number of ways. The failing unit can:

- Force the output voltage rail to go out of specification or fail completely
- Disrupt the input power source so that other converters are deprived of input power
- Cause malfunction of, or damage to, the load,
- Cause malfunction of, or damage to, other power converters

A properly designed power system must not allow any of

these to occur when there is a fault of any kind in any one (or more, if necessary) converters. Each of the following interfaces must be examined in detail to assure that a failure in one unit will not affect the system in any way:

- Power (AC or DC) output stage of any converter
- Power (AC or DC) input stage of any converter
- Power input to, and distribution within, any system load
- Signal connections between the power converters
- Signal connections between power converters and a Power System Controller
- Signal connections between the power system and the system being powered

Examples of concerns that must be addressed in each of these areas include:

- Shorted output capacitor or diode in a power converter that can disrupt the power distribution to the load
- A catastrophic short in the input stage of one of several paralleled power converters that causes the circuit breaker or fuse in the input power distribution circuit to open, causing all units to shut down
- A shorted component on a circuit board in the load that disrupts the power distribution to other circuit boards
- A shorted or open current share signal line
- A Power System Controller fault that fails to identify a failed power converter
- A fault in the host system that erroneously sends a command to turn all power converters off

## Reliability and Availability

These concerns will be briefly examined at each of the major power and signal interfaces in a power system. In the **power converter output stage**, the main problem is that a shorted component, typically a capacitor or diode, can cause the output to fail short. Without an ORing diode, this would cause the output voltage bus to fail - an unacceptable consequence. Another problem that can disrupt a system is a runaway converter. If the system load is less than the output capacity of a given converter, a runaway condition could drive the output voltage up to levels that would damage or destroy the load. Overvoltage detector and shutdown circuits are used to prevent this. However, when converters are paralleled, it is important that only the faulty unit shut down. Controllers can compare the voltage on each side of the ORing diode. If the diode is reverse biased, that converter is not at fault and should not shut down for an overvoltage fault.

The system must also be protected from disruption by a short circuit or other high current fault on the **power converter input stage**. Such a fault must not damage or disrupt the power source, be it the AC mains or an intermediate distribution bus. A current limiting device such as a fuse or circuit breaker generally accomplishes this protection. In the case of a DC intermediate bus distribution system, the choice of fuse or circuit breaker can have unexpected consequences. Clearing the fault quickly is important in terms of minimizing the amount of time that the intermediate bus is pulled down.

Experience has shown that typical glass cartridge fuses may not clear for tens of milliseconds-long enough to disrupt system operation. The use of small, surface mount, very fast blow pico-fuses is encouraged wherever possible. The system can be made more immune to these effects by incorporating a diode and holdup capacitor network at the input to each DC/DC assembly

as described in Chapter 5.

In any power system, it is not only the power converters that have power inputs. The **system load power input** must also be protected against faults in the loads. If the power is fully distributed and partitioned so that each load element has its own power source, this is easily managed. For example, the output overcurrent protection of the converter powering the load prevents a short circuit in the load from disrupting the intermediate distribution bus. Systems that use single power distribution bus to power multiple loads require protection circuits or devices in the loads. For example, consider a system with multiple circuit cards powered from a common +5V source. A short circuit on any one of these cards would interrupt the +5V power. In the same way that DC-DC converters powered from an intermediate distribution bus need input current limiting, so do loads powered from a common bus.

The **inter-converter and intra-power system signal connections** must also be considered. There are typically very few signal connections between power converters in a distributed power system. The most common is an output current or load-sharing signal. The current sharing circuit must be designed such that no disruption of the output power occurs if the current share signal line is:

- Shorted to output return
- Shorted to any output voltage
- Shorted to a voltage internal to the converter
- Open circuited at any point

Current sharing among the paralleled converters will

typically be lost under these conditions. It is important that the loss of sharing be detected and signaled (see fault detection and annunciation below). An example of a common practice that cannot be used in fault tolerant systems is a daisy-chained clock signal that drives auxiliary power converters. In most implementations, a fault on the clock signal would cause one more additional power units to stop operation, which would, in turn, cause the output voltage to fail.

Problems can also occur when other types of signals are daisy-chained among paralleled converters. For example, consider a remote on/off signal that causes the converters to turn off when pulled low. If all of the on/off signals were connected in parallel, then an on/off input shorted to ground in one converter would cause all of the converters to turn off. However, driving each converter independently may not hold up under a cost-benefit analysis. This is another example of the cost and complexity versus availability tradeoffs that must be made when implementing a highly available power system.

It is also important to consider the fault isolation implications of the **power system to host system signal connections**. Customers specifying a highly available or fault tolerant power system often ask for Power System Controllers (PSCs) with advanced functionality. This often requires the use of microcontrollers. Functions desired include the ability to turn individual power converters on and off, margin output voltages up and down, monitor voltages and currents, perform periodic diagnostic checks, control battery connect and disconnect devices, and similar functions.

These PSCs may fail and bring down the entire system.

Consider, for example, a system with a PSC that has a single signal line that goes to each Front End Power Supply for remote on/off control. If the controller were to fail such that all the Front End Power Supplies were turned off, then the system fails. This violates the requirement that no single failure disrupts the system. This is a difficult situation. Implementing redundant controllers is a complicated and expensive proposition. Often the choice is to do what can be done to minimize the opportunities for, and likelihood of, such fails. One tactic is to make the controller more of an observer - reporting power system status to a system maintenance function but without direct power system control.

**Fault Detection and Annunciation** - At this point the system has redundancy and no single failure will propagate and disrupt the system. These are necessary, but not yet sufficient, conditions for the system to be considered highly available or fault tolerant. A key part of the availability calculations was the Mean Time To Repair (MTTR). A system that has a failed component or sub-system must first detect that failure. Once detected, it must announce that malfunction in some manner so that a repair can be made. If the failure is not detected and repaired, then the system is no longer redundant. The increased availability that redundancy brings is now lost, and the next failure may very well cause the system to malfunction. Such a system cannot be said to be fault tolerant or highly available.

The detection of faults is not limited to those that cause an obvious loss of function. For example, if an output ORing diode were to fail short, the power converters and the system being powered would continue to operate with no indication that there had been a failure. However, if that unit were then to experience a shorted output

## Reliability and Availability

capacitor, the output voltage would be disrupted. This is an example of a **hidden or latent failure**.

There are two basic ways of determining if a unit has a latent failure:

- Exercise the fault detecting functions through an external stimulus and check for the proper response
- Separate and independent means of verifying the converter's operating status

Note that both of these options generally require an external agent. Extensive self test can be built into converters, but this often adds excessive complexity and cost and an unacceptable reliability penalty.

An example of finding a latent failure through an external stimulus and response check is the problem of finding a shorted output ORing diode. Finding a shorted diode by directly measuring the diode's forward voltage drop is nearly impossible. When operating normally, the diode forward voltage is as small as possible. When a diode shorts, it still has a finite resistance that causes a small voltage drop. Reliably distinguishing between a normal forward drop and shorted voltage drop is essentially impossible.

One method for detecting an ORing diode that has failed short is to margin down the output voltage of just one of the units operating in parallel. The voltage is reduced to less than the threshold of the output low voltage detector. The detector senses the voltage on the converter side of the ORing diode. If the converter does not signal an output low voltage alarm under this condition, then one of the following has happened:

- The ORing diode has failed short
- The margin circuit has failed
- The low voltage detector has failed
- In any case, the unit has a fault and needs to be removed from operation

An example of a latent failure that can be found by an independent measurement of a converter's operating condition is a current sharing fault. Suppose each of the paralleled converters had a current monitor signal with an output that was proportional to its output current. Then an external diagnostic agent can check the output current of each converter and compare it against the other units. A unit that is operating with an output current much higher, or lower, than that of the other units has probably failed.

Once a failure has been detected, it must be **annunciated**. The nearly universal way to indicate a failed unit is with a visual indicator, usually an LED or other type of lamp that either illuminates or extinguishes when the unit has detected a fault. If service personnel continuously monitor the equipment, a visual indicator may be sufficient. This is rarely the case and the equipment must generally provide a signaling function. Relay contact closure or optically isolated open collector signals are the simplest way to signal a failure. More complicated, and thus perhaps less reliable, schemes may involve some form of serial digital communication between a power system controller/observer and the host system's maintenance processor.

**On-line Repair** - There is still one more criteria to meet before the system can be considered fault tolerant. So far the system has redundancy, failure isolation, and as

many faults as economically possible are detected and signaled to a maintenance function. All of these are necessary but not yet sufficient for true fault tolerance. The calculations that led to high availability did not allow for system down time except in the rare instance that  $M + 1$  units had failed. This means that a system with a failed unit cannot be made unavailable while the system is being serviced. When this requirement is met along with the other three, all of the conditions for a highly available or fault tolerant system have been established. There are some exceptions to this final requirement. We saw an example of the sales processing system where it may be possible to schedule a system down time during the evening in order to make a repair. Another case where on-line repair may not be required is systems that are themselves  $1 + 1$  redundant.

In general, making an on-line repair involves hot-plugging either a power converter or a load module. For power inputs and outputs, this is often managed with inrush limiting circuits and ORing diodes. There are several alternatives available for inrush current limiting, ranging from simple mechanical solutions with long and short connector pins to utilization of hot-plug control integrated circuits. The range of solutions is discussed in Chapter 5. Connectors must be rated to maintain their characteristics after repeated insertions - and one must not forget to allow for the number of contact cycles that will occur in the manufacturing and test process. It is also important to use appropriate keying or another mechanical device so that only the correct converter or load board can be inserted into any given connector.

Attention must also be paid to the signal lines. For example, the current sharing circuit must be well behaved during the removal and insertion of a new unit.

If an intelligent digital controller is used, care must be taken in the fault detection and monitoring functions. Generally, if a particular module stops responding it is presumed bad and an appropriate fault indication is sent to the host system. But what if the unit comes back on line? Does this mean that the unit is intermittent or that it was removed and replaced? The monitoring program must be able to determine the difference and respond accordingly.