

Chapter 3:

Selecting a Power System Architecture

The optimum system power solution is a combination of the best power converters and the most suitable architecture. Selecting the best architecture is non-trivial for both cost and performance - even TTM may be adversely affected by a poorly thought-out power strategy and whole markets may be denied the OEM manufacturer. This chapter details the major architectures available to the system designer and the advantages and pitfalls associated with each.

3.1 Introduction

The first and most basic decision that a power system designer is faced with is what architecture to use. This decision will influence every other aspect of the system design including the types and quantities of power converters that will be needed. More importantly, it will determine the total lifetime reliability and cost of the power system and, in some cases determine the success or failure of the entire project. This chapter will define and describe the most popular power system architectures, present information that will enable you to select the best one for your system, provide some examples of how these architectures are applied in actual systems, and describe how fault-tolerant design can be accommodated by various architectures.

First of all, what is a 'power architecture'? All electronic equipment operates from a power source. These sources are, in order of prevalence, the AC powerline, portable batteries, or a fixed-location battery bank. Each piece of electronic equipment also has certain power needs for its circuitry in term of volts and amps. There are many ways to provide these requirements both in terms of dividing the load up into smaller partitions and in selecting from the wide range of available power converters in several variations of power ratings and package sizes. When all the alternatives and combinations are considered, there are hundreds of possible ways to implement the power system for a given product. Power system architecture is the science and art of selecting one solution from these multiple choices that meets the product's requirements in terms of cost, reliability and functionality.

While most of the possible solutions for a given requirement are truly unacceptable for what will be obvious reasons, there will generally be more than one approach that remain feasible. The final selection from between the best two or three candidate architectures often entails considerable analysis and trade-offs since they may all be acceptable with no one solution.

Selecting a Power System Architecture

appearing to be the best overall. The remainder of this chapter will assist the reader with making the wisest choices when faced with architectural decisions of this type.

3.2 Centralized vs. Distributed Architecture

We will soon see that, when all the variations are taken into account, the number of possible architectures can be overwhelming and confusing. The best way to begin exploring these architectures is to first define the most basic and widely used ones. These architectures clearly demonstrate the types of considerations we will be looking at in this chapter. We will first define a centralized architecture and then the two most widely used types of distributed architecture - shelf-level distributed and card-level distributed.

A **centralized power architecture** has all of the power processing functions packaged at a single physical location within the equipment. The load electronics are packaged separately and connected to the power processing function by DC distribution hardware as shown in Figure 3.1. As a consequence, the number of locations containing power (only one) has no relationship to the number of load packages. At high power levels, the power supply function is sometimes segmented into more than one assembly. This could be done for purposes of simplifying the power supply designs, adding redundancy or to achieve physical packages that are reasonably sized and easy to handle. But if these multiple assemblies are still located at one physical site in the equipment, this is still considered centralized power. If the multiple power assemblies are similar and packaged for easy insertion and removal, we will refer to

this technique as **modular power**. See Figure 3.2 for an example of a centralized power architecture with more than one power converter. Figure 3.3 depicts the packaging of a typical product using centralized power architecture.

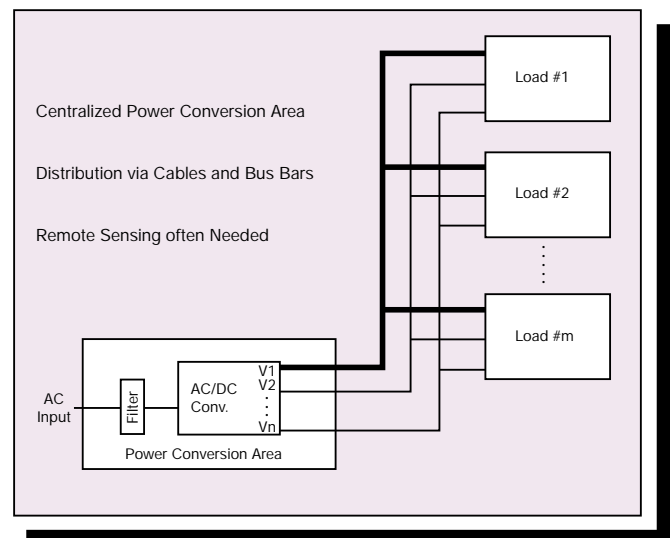


Figure 3.1 - Basic Centralized Architecture

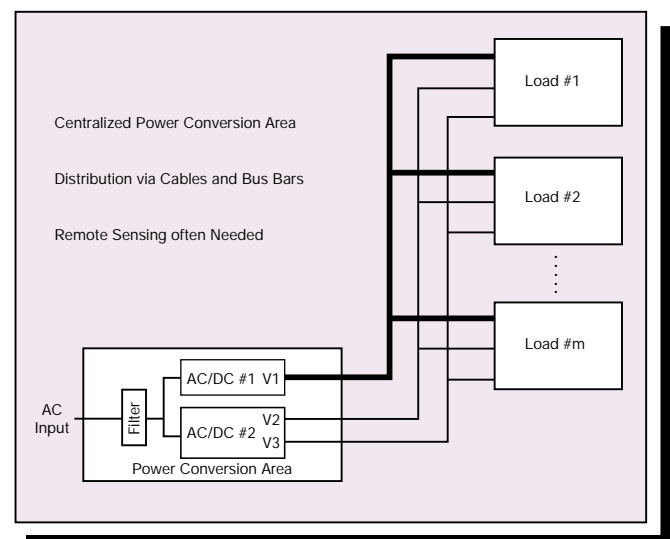


Figure 3.2 - Centralized Architecture with Separate Converters




Figure 3.3 - Typical Packaging for Centralized Architecture

The basic attribute of a **distributed power architecture (DPA)** is that the power conversion hardware is not located at one place within the equipment, but rather placed at more than one location. This placement can be done for a number of reasons, such as reduced distribution distance, better dynamic performance, thermal considerations and ability to use standard hardware. We will explore these considerations and how to use them to optimize the power architecture later in this chapter. For now we will look at the two basic DPA implementations that are in most common use.

The first DPA system, shown in Figures 3.4 and 3.5, is called a **shelf-level** system. This is a popular architectural choice for systems that are packaged in a rack or card cage type of enclosure. Each level of the enclosure (or shelf) contains space for a power converter that provides power for electronics on that particular shelf. These converters obtain their power input from a DC source that is isolated from the AC powerline. This source is most commonly either the telecom central

office battery system or the output of one or more AC/DC converters operating with the AC powerline as their input. The voltage can also be obtained from internal backup batteries for systems utilizing a battery backup powerline fault immunity strategy. This source voltage, sometimes referred to as the '**intermediate voltage bus**', is designed to be below the Safety Extra Low Voltage (SELV) limit in order to simplify meeting international safety standards. The DC/DC converters, in general, provide more than one different DC output voltage depending upon the requirements of the circuits on that shelf. The DC/DC converters are typically in the 50W to 500W power range and could be implemented with either custom designs or a collection of standard power converter modules. Each of the output DC voltages is distributed within a backplane to the circuit cards on that shelf.

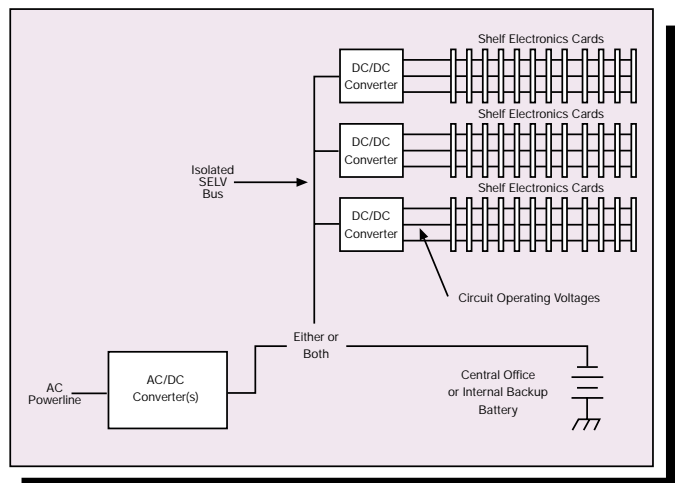


Figure 3.4 - Shelf-Level Distributed Architecture

Selecting a Power System Architecture




Figure 3.5 - Typical Shelf-Level Distributed Architecture Packaging

The second DPA system we will feature is called a card-level system and is shown in Figures 3.6 and 3.7. This architecture is also appropriate for systems with rack or card cage packaging. With this approach, the shelf DC/DC converters are eliminated and the final stage of voltage regulation takes place on the electronics cards themselves. The power conversion on each card is optimized for the card's power requirement and is implemented with low power (mostly less than 50W) DC/DC converters. Two or more low power converters are often used. These power converters are sometimes referred to as **Point of Use Power Supplies (PUPS)**. This approach lends itself to the usage of standard DC/DC converter modules rather than custom designs. This standardization offers several advantages that will be described elsewhere. The goal is to make the on-card converters small enough so that they can be treated as components during the manufacturing process resulting in automated assembly. It is also possible to package nothing but power on certain cards and distribute the final circuit voltages through the backplane to adjacent cards, although this will negate some of the

advantages of a 'pure' card-level architecture. The **dedicated power card** is shown in Figure 3.6 in which the difference between this approach and the pure card-level architecture can be seen. The DC/DC converters in a card-level system operate from the same type of SELV intermediate bus as the shelf converters in the shelf-level system.

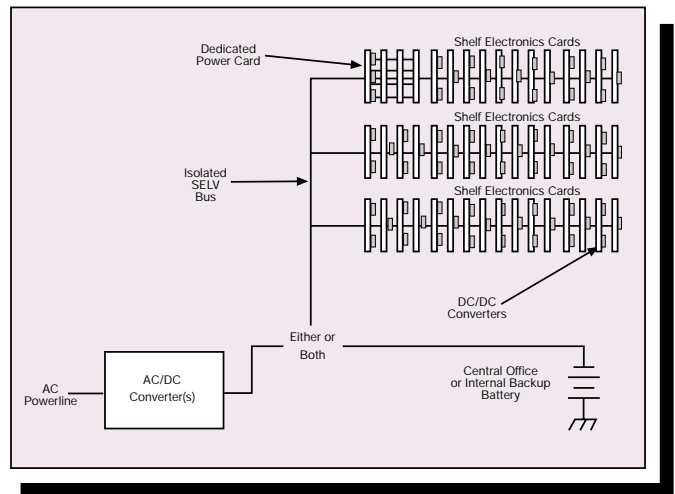


Figure 3.6 - Card-Level Distributed Architecture




Figure 3.7 - Typical Card-Level Distributed Architecture Packaging

3.3 Distributed Power Variations

We have defined and briefly described the two major types of distributed power architectures. Before addressing their advantages and disadvantages, we will now look at some common variations of the DPA approach.

The first variation is a combination of centralized and distributed architectures, and is referred to as a **hybrid** architecture. This could actually take many forms, as shown in Figure 3.8. Some systems have a basic centralized architecture, with most power derived from a multi-output AC/DC converter, but also have a portion of the system implemented with a distributed converter approach. One form of this approach is shown in Figure 3.8a. In this system, a DC/DC converter is located outside of the main power conversion area, perhaps for the purpose of providing low distribution drops to some critical circuitry. Figure 3.8b shows a similar system where there is a remote area of card-level distributed

power. This could be done in order to provide modularity and the ability to plug in different circuit cards that could change over the life of the equipment. With the card-level distributed power in this area, the power for the pluggable cards will automatically be configured without the need to change any of the basic AC/DC area power as long as the power demand on the front-end does not increase significantly.

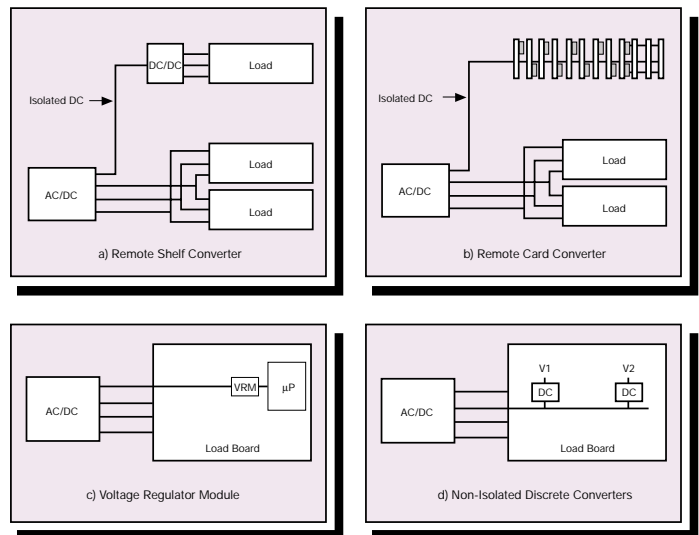


Figure 3.8 - Examples of Hybrid Architectures

Perhaps the most ubiquitous example of a hybrid power architecture is the personal computer. PC power supplies are classical centralized multi-output AC/DC converters, which served the industry well for many years. The latest high-speed processor (μP) chips, however, require a dedicated local **voltage regulator module (VRM)** in order to provide the low processor voltage at high currents and large dynamic current swings. These VRMs are small high-density non-isolated converters that use an isolated DC voltage such as 5V or 12V as their input. They are mounted in close physical proximity to the processor chip. This technique is shown in Figure 3.8c.

Selecting a Power System Architecture

Another example of a hybrid architecture is shown in Figure 3.8d. This system is similar to the VRM approach except that the final DC/DC conversion stage is accomplished with discrete circuitry (also known as on-board power) rather than by a VRM. The discrete converters are non-isolated and are designed around control chips from semiconductor companies. This technique is most often used when one or more unique voltages are required at low current and is most practical at power levels of 10W or less.

The next architecture we will look at, **shelf & card**, is a DPA variant that includes an additional stage of power conversion. It is essentially a combination of shelf-level and card-level architectures with the on-card converters implemented without isolation. The shelf DC/DC isolated converter is often configured to provide an output voltage of around 12V, which is used as the input to the on-card non-isolated down converters. This architecture is illustrated in Figures 3.9 and 3.10. It is used when the cost saving of using non-isolated card converters outweighs the disadvantages of the loss of efficiency and additional volume of the extra conversion stage.

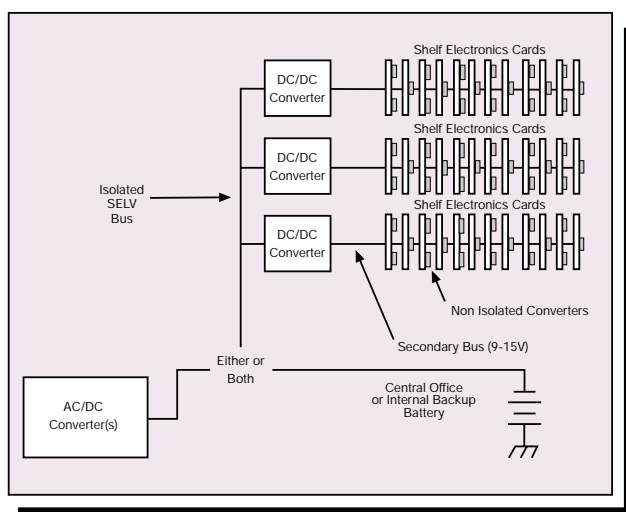


Figure 3.9 - Shelf & Card Distributed Architecture




Figure 3.10 - Typical Packaging for Shelf & Card Distributed Architecture

All of the DPA systems we have discussed use an isolated DC intermediate bus voltage. There are other architectures that are sometimes referred to as distributed that do not use the SELV bus. One of these distributes the AC powerline voltage around the system to locally placed AC/DC converters. This system is indeed physically distributed, but it lacks most of the positive attributes of DPA systems that we will soon identify. Consequently, we will not include such systems within our definition of DPA. A variant of the distributed AC approach is to distribute the rectified but non-isolated powerline voltage or the output of a PFC converter throughout the system. These high voltage (200V to 500V) approaches result in a non-SELV distributed bus, and we will not include them within our DPA definition. Power architectures of this type are primarily used for large data processing and communications systems that have no user access and are serviced by trained personnel.

3.4 Bus Voltage Selection

The most significant common feature of the DPA systems we have defined is the usage of an isolated intermediate bus as the input power source for DC/DC converters. Conceptually, the voltage on this bus could be any level. We will now see that the selection of this voltage, which used to be a controversial topic, is actually fairly straightforward. Since the load for the intermediate voltage bus is a constant power (a function of the load circuits), the current on the bus will be inversely proportional to its voltage. Minimizing the bus current is advantageous, as it will minimize the amount of copper in the bus, the cost of the bus, and the number of connector pins required to move the current in and out of the bus. As a result of this basic consideration, it would seem that the bus voltage should be as high as possible.

There are other considerations, however. Obtaining safety agency approvals without adding costly special packaging and safety shielding to the system requires keeping the bus voltage below the SELV limit, which is 60V. The advantages of being below this limit have resulted in the vast majority of DPA systems selecting a voltage that is less than 60V. With allowance for tolerances, the existing nominal 48V telecom battery voltage nicely fits this requirement. Even at power levels of 4 KW, the bus current will be less than 100A. Since the telecom standard has been in existence for many years, standard DC/DC converters were first developed to utilize this input voltage. The large base of available standard converters is, in itself, a powerful argument for using a 48V intermediate bus. The advantages of using standard converters as opposed to custom designs will be covered in detail later, but they center around design flexibility, fast time-to-market, minimal technical risk and proven reliability. Another similar 'standard' intermediate

voltage has evolved at a nominal level of 24V. The 24V nominal level is used mostly for radio communication systems and industrial electronics, and is also supported by a growing array of standard DC/DC converter products. Since, for a given power level, the bus current will be twice as high as in a 48V system, 24V systems tend to be used more often at lower power levels. Internal battery backup can be less expensive for low power equipment at 24V since the number of battery cells needed will be half that of a 48V system.

Keep in mind that the 24V and 48V design points are only the nominal voltages. Each bus, as described in the telecom standards, can actually have a wide range of voltage due to transients, switchover to battery power and battery charge and discharge. The better standard DC/DC converters are designed to operate over this range, as noted in Chapter 2. The ability to operate over a wide input range also provides benefits to the intermediate bus design. Unlike distributing centralized power, where the distribution bus needs to hold tolerances of millivolts, the intermediate bus voltage can vary by several volts with load, reducing the size and cost of the bus.

With the recent trend toward very low circuit voltages (less than 2V), some conversion topologies are more effective and efficient with input voltages lower than 48V. This has resulted in the increased usage of input voltages of 5 to 12V for converting to low voltages at low power levels. Examples are the usage of VRMs, discrete low power converters and shelf & card architecture as described earlier. We will refer to intermediate voltages of less than 24V as 'low voltage'. The following table describes, in general terms, the recommended intermediate bus voltage for specific equipment power ranges.

Selecting a Power System Architecture

Equipment Power Range (W)	Recommended Bus Voltage
<100	Low Voltage, 24V
100 to 500	24V, 48V
>500	48V

3.5 Architectural Trade-offs

There is no single criterion that can be used to make a wise decision between the available architectures.

Rather, several factors must be considered and these will have a different importance or weighting for each system and its designer. In this section we will provide a general discussion of the criteria that are most often used to make the decision between alternative architectures. Some of them will be important to your system while others may not be relevant.

Hardware Cost - The cost of the power hardware for a distributed power approach is rapidly decreasing, but it is still normally higher than that of a custom centralized approach for a system with large manufacturing volumes. Large volumes (over perhaps 50,000 per year) allow the substantial cost of custom power converter development to be amortized to a reasonable amount per system. As long as the design remains stable over the product life, this advantage in the power hardware cost can be realized. Distributed approaches use a larger number of high-technology converters, which tends to increase the total power converter cost. It is Artesyn's belief that the Total Cost of Ownership (TCO) is a more meaningful measurement of the overall affordability of the power system, but for those companies that use only power converter hardware cost as the cost criterion the centralized approach can be attractive.

Total Cost of Ownership (TCO) - TCO includes not just the cost of the power converters, but also costs associated with distribution hardware, connectors, power thermal management, power system design, development and qualification, sourcing costs, change activity, reliability costs and, sometimes, energy costs. When these factors are included, the cost differential between DPA systems and centralized systems are reduced or reversed. Hybrid systems tend to be intermediate in terms of TCO.

The largest factor in the reduction in TCO for DPA systems is usually the huge reduction in development cost. Designing, sourcing, qualifying and maintaining engineering change activity on a custom power supply can cost millions of dollars. Usage of standard DC/DC converters and standard AC/DC front-end converters eliminates most of this cost. Also, the high reliability and quality levels of the latest standard converters often result in savings in system repair and maintenance activity. Companies that utilize a TCO costing approach will each have their own methodology for doing the detailed costing, but a generalized approach will be discussed in a later section on cost.

System Manufacturing Volume - The system annual production quantity is sometimes a factor in the architectural decision. For high volume systems with stable designs, the centralized approach is often less expensive - at least in terms of power converter hardware cost. As the production quantity decreases, DPA approaches - especially those utilizing standard converters - become more attractive.

System Operating Life - The equipment's expected service life, along with its reliability, will determine the average number of repairs needed. Since the repair strategy is usually easier with DPA, equipment that is expected to be maintained in the field for several years often uses a DPA design. One good example of this is telecom equipment, which is often designed for a service life of over 20 years. With DPA, as new circuit cards are designed for such systems, the DC/DC power converters can be upgraded to the latest technology and output voltage levels over the system's life without the cost or disruption of an overall power system re-design. This would not be possible with centralized power. If the system has a short lifetime, a PC for example, then the centralized approach has merit since most systems would not experience a single power failure during their useful life.

Reliability - High levels of reliability are achievable with any architecture. But to accomplish it with a custom centralized approach requires that the custom design be perfect and that none of the components in the converter experience unexpected failures. This usually requires more than one design iteration and extensive testing and qualification, which add cost and time. The better current standard DC/DC converters are highly integrated and automatically assembled. They are proven designs and use high quality components. This results in very high reliability levels, with MTBF values between 1 and 5 million hours. Even when the larger number of converters in a DPA system is taken into account, the DPA system will be extremely reliable - and without the need for extensive specialized analysis or testing on the part of the system designer.

Availability - Reliability is a measure of the number of failures. Availability is a measure of the ability to remain operational in spite of failures by means of fault-tolerant design. A following section will discuss fault-tolerant architectures in more detail, so we will only provide a summary here. Basically, it is much more difficult to do a fault-tolerant design with a centralized architecture than with a DPA approach. This is due both to the modular nature of the DPA system and to the ease of obtaining immunity from powerline faults with the DPA system. As a consequence, systems with a requirement for high levels of availability will tend to use some form of distributed approach.

Flexibility - In the real world product development does not proceed entirely smoothly with all the requirements defined at the beginning of the process and remaining unchanged over the duration of the project. The voltage and current requirements for some circuit functions are usually late to be defined. This will impede the progress of centralized architectures with custom power supply designs. The design process will either be late to get started or there will be engineering change activity along the way, and either one will delay the completion of the project. Hybrid and DPA approaches will reduce or eliminate this problem. A hybrid approach where low power non-isolated voltages are developed locally on-board from 5V or 12V with discrete or modular converters is one way to minimize the effects of change activity when there is a commitment to a centralized custom power converter. Shelf-level DPA systems will contain the change activity to the shelf DC/DC converter assemblies so that the front-end AC/DC converters and intermediate voltage distribution can remain stable. The ultimate protection against schedule delays from changes in the product requirements is provided by card-level DPA. With this approach, changes in circuit

Selecting a Power System Architecture

requirements are contained to the circuit card itself.

The above arguments also apply when considering how to design power systems that are flexible with regard to supporting several models or variations of a product. It is common in the telecom and datacom world to offer the buyer a choice of features and functions, such as multiple processors, I/O selection and storage options. Ideally, the power system should scale with the selected features so that basic entry-level systems are not penalized with the total cost of power for the most elaborate system. Furthermore, it is often desirable for these configuration changes to be made in the field, sometimes by the end user. These objectives are best met with the card-level DPA approach.

Time-to-Market (TTM) - There is often a significant competitive advantage to getting your product to market earlier. The additional sales and profits that are captured by an earlier introduction can mean millions of extra dollars on your bottom line. Centralized custom power converters, while sometimes less expensive in terms of hardware cost, can take over a year to design and qualify. By using DPA architectures along with standard converters, the time for development of power can be drastically reduced. Shelf-level and card-level DPA will allow the main front-end power to be designed before the detailed power requirements of the load circuitry have been defined. Prototype power converters are available within days or weeks and, since they are standard, they will perform the same as the final production hardware. It is not unusual for a DPA approach to reduce the TTM for a complex product by at least 6 months. Weighing the additional profit that accrues from the earlier introduction along with the other costs is sometimes the deciding factor in which architectural choice is the best.

Technical Risk - Technical risk is the probability that an unexpected event will occur during the power system development that creates extra cost or schedule delays. Some examples are:

- Design problem with a power converter
- Poor quality level in a component
- Difficulty with procurement of a component
- EMC problems
- Manufacturing or tooling problem with power supply
- Failure during internal qualification testing
- Failure during agency approval testing

These risks are significantly higher with the development of a custom power supply, even with a high quality supplier. Designing multiple outputs into one assembly usually entails custom packaging, new circuit board layouts, new components and magnetics to design and source and a different EMC environment. These risks are best tolerated when the initial design point is stable and there is enough room in the schedule for at least one re-design cycle. All of these concerns are minimized when using standard converters, since the converter manufacturer has established a proven design prior to releasing the product. With DPA systems, if there is any problem it will tend to be contained to one section of the system, allowing the development of other areas to proceed.

Thermal Considerations - The thermal environment is different in a centralized system than in DPA systems. Since all of the power conversion takes place at one area of the equipment enclosure, all of the power dissipation associated with the inefficiency of the power converter(s) will be located here. Assuming an overall conversion efficiency of 80%, this power dissipation will be 25% of the load power and will normally require some sort of

forced convection cooling. This approach is certainly feasible and widely used, but will require a dedicated fan or blower and/or customized air plenums and airflow analyses.

With DPA systems, the heat load associated with power conversion tends to be much more distributed within the equipment enclosure, eliminating some of the dedicated thermal solutions associated with centralized architectures. The specific DPA variation is often dictated by the overall system thermal approach. For example, to accomplish shelf-level DPA within a reasonable volume will require the availability of forced convection, at least within the section containing the shelf converters. Card-level DPA systems can be either forced convection or free convection. Forced convection allows for enhanced packaging density but with the addition of acoustic noise and maintenance requirements for the fans and filters. Long lifetime telecom equipment often uses free convection cooling for its simplicity and reliability. Such systems are admirably suited to a card-level DPA approach because of the distributed heat load. With either forced or free convection cooling, on-card DC/DC converters will require some of the system cooling capacity for removal of heat from the converters. However today's better converters, such as Artesyn's EXB50 series, are approaching and beating 90% efficiency figures, minimizing this heat load and imposing a very minimal thermal penalty for inclusion of power on the circuit card.

Standard vs. Custom - The majority of DPA systems are configured with standard power converters rather than custom designs. This makes it difficult, in a pure sense, to separate the effects of the architectural choice from the choice to use standard products vs. custom

designs. In fact, many of the benefits of DPA we were discussing are due to the hardware standardization rather than to the architecture itself. For practical purposes, both must be considered together, and that is the approach we will take here. DPA designs and standard hardware tend to support each other as do centralized designs and custom hardware.

Distribution Hardware and Connectors - As the average circuit operating voltage has moved steadily downward with no reduction in total power, the current levels to supply the end circuitry has been steadily increasing. With centralized approaches, this current must travel long distances over customized bus bars or cables, while remote sensing is needed to compensate for the associated voltage drops. The cost of providing this distribution has been one of the major factors that has influenced the trend toward distributed architectures. So, if your system has requirements for low voltages (3V or less) at high current levels, it could be a candidate for some form of distributed or, perhaps, hybrid architecture. With DPA, most of the inter-board power distribution occurs at reasonably high voltage (>24V) levels and this minimizes the current levels and the number of connector pins required.

Manufacturing - The architectural decision will affect the manufacturing process. Customized centralized power will require some customization of the manufacturing process, often involving large and heavy assemblies that need manual assembly and interconnection. Labor hours will be spent on fabricating and interconnecting complex bus bars and wiring harnesses. With DPA approaches, more automation can be used in the manufacturing process since many DC/DC converters use standard solder profiles and are small enough and light enough to

Selecting a Power System Architecture

be handled with standard pick and place equipment. If standard AC/DC front-end converters are used for the system, they can be quickly inserted into the associated standard racks with all connections automatically established by means of the backside floating connectors. Because of factors such as these, the manufacturing hours for a DPA system are significantly less than for a centralized one.

Fault Diagnostics and Repair - The system repair strategy sometimes influences the architectural choice. Non-repairable systems such as consumer electronics often use a centralized approach. Diagnosis and repair of a large centralized power system tends to be a complex process. Power converters, distribution hardware, load boards, sensors and remote sensing are all interconnected into one large interactive network. A fault at any point in this system can bring down the entire network. For example, a short on a circuit board will shut down the power supply. Finding the location of the short can take an excessive amount of time, even for specially trained service personnel. For systems that need to be diagnosed and repaired quickly utilizing minimal skill levels, DPA has some definite advantages. Faults tend to be contained to either the card involved or to a well-defined area of the overall system. A fault in a shelf DC/DC converter will only affect that shelf. DC/DC faults in a card-level DPA system will only affect the card involved and not be propagated to other cards or to the front-end power system. The faulty card or AC/DC converter can easily be replaced by unskilled personnel, especially if hot-plug technology is used.

The above discussion is summarized in Figure 3.11. This table lists the main advantages of the three most widely used architectures - centralized, hybrid and card-level

distributed - along with the types of equipment and applications in which they are most commonly employed. While card-level DPA is used in the chart for the purposes of delineating the advantages of DPA, many of the same advantages will be present (but to perhaps a lesser degree) with shelf-level and shelf & card approaches. Your particular application may fall neatly into a category that makes the architectural selection straightforward. or you may need to spend a fair amount of time with the considerations outlined above before being able to finally select an appropriate architecture. This may require a lot of soul-searching, especially for a first-time DPA user. The overview presented here along with the more detailed information contained in other chapters should make the process an easier one.

Architecture	Advantages	Where Used
Centralized	Less hardware \$ in large volume	Consumer electronics High volume applications Non-repairable products Products with short service life Equipment wth skilled service
Hybrid	Quick mods to Centralized Systems Moving power supply closer to load	High performance PCs Datacom with high DC currents
Card-Level DPA	Reduced TCO Enhanced TTM Reduced technical risk Ability to use standard converters Easy to design fault tolerance Ease of product changes over life Upgrade power technology over life More automated assembly Fewer power connector pins needed Simpler DC distribution Distributed thermal load Easier fault diagnosis	Low/mid volume systems Equipment with long service life High availability systems User maintained equipment

Figure 3.11 - Summary of Common Architectures

3.6 Fault Tolerant Architectures

The architectures we have considered to this point do not allow for system operation in the event of a power converter failure. Many of today's systems require guaranteed up-time and some form of fault-tolerant design is used in such cases. The ease of accomplishing this fault-tolerance is strongly dependent upon the power architecture selected, with DPA approaches having a distinct advantage over centralized systems. In this section we will look briefly at fault-tolerance and redundant design as they relate to architectural choices. Calculation of system availability as a function of the number of redundant converters will be discussed in the chapter on reliability design. Details of how to implement the paralleling of converters and designing for hot-plugging will also be covered in later chapters.

Well designed modern power converters are extremely reliable - much more so than in the past. As a consequence, a typical system will experience loss of power from a powerline outage orders of magnitude more often than from power converter failures. Therefore it very seldom makes sense for an AC powered system to use redundancy techniques in its power system without some provision for buffering from powerline faults. With this in mind we will first look at the most common techniques for providing this immunity from powerline faults.

The most common technique for centralized systems, as shown in Figure 3.12a, is the **Uninterruptible Power Supply (UPS)**. This device is a battery-powered inverter that creates an AC source that is substituted for the powerline voltage in the event of a powerline failure. The battery is charged during normal operation and is sized as appropriate for the desired duration of powerline

immunity. For DPA systems, the most often used technique is the internal **battery back-up (BBU)** on the intermediate voltage bus as shown in Figure 3.12b. The front-end AC/DC converters charge this battery during normal system operation, and it automatically picks up the system load in the event of a powerline outage. Since the UPS needs a high power inverter circuit in addition to the battery, it tends to be less reliable, larger and more costly than the corresponding BBU approach. Also, it is normally packaged external to the equipment - a disadvantage for most users.

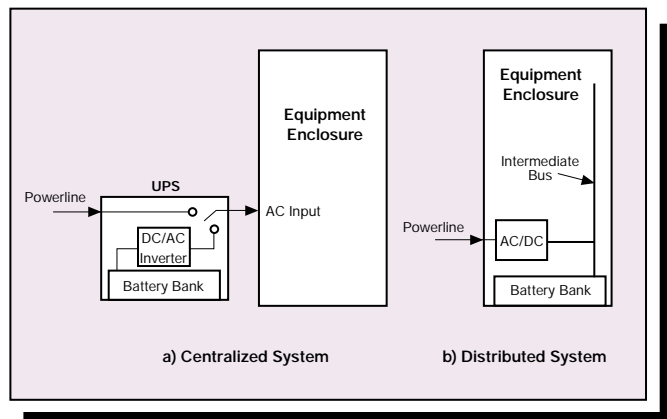


Figure 3.12 - Powerline Immunity Techniques

Once provision has been made for powerline immunity, the effects of power converter failure can be considered. The type of design required to achieve fault-tolerance from power converter failure depends on the power architecture. Fault tolerant design for centralized power systems is possible, but complex. The entire power supply or supplies must be replicated so that twice the amount of packaging volume is required. For complete redundancy, the distribution bus and remote sensing system must also be duplicated and all output levels diode isolated. Such designs have been done, but they tend to be difficult and expensive.

Selecting a Power System Architecture

With DPA, the problem becomes much more approachable. The front-end AC/DC converters can easily be configured in an 'N+1' arrangement. If the load power is, for example, 1000W and each AC/DC converter is rated at 500W, two converters would be used for a non-redundant design. By adding a third converter, each one will operate at 333W during normal operation, increasing their reliability compared with operation at 500W. In the event of one of them failing, the remaining two will automatically pick up the entire load so that the system will operate through the failure. Since there is only one output voltage at a reasonably high output voltage level and usually no remote sensing requirement, the addition of the ORing diode adds very little in the way of reduced efficiency or increased cost. And this is accomplished with only a 50% increase in the volume required for the AC/DC conversion function. Standard 'power shelf systems' are available that will accomplish all of these functions with no need for detailed design by the OEM.

Redundancy techniques can also be used with the DC/DC converter implementation. In a shelf-level DPA system, it is possible to use a redundant shelf converter. This is easier than duplicating the AC/DC converter in a centralized system, but still can be challenging in terms of the packaging volume required and the issue of providing redundancy in the shelf DC distribution. For these reasons, most DPA systems that have a firm requirement for complete redundancy usually would either have completely redundant shelves (DC/DC converter and load cards) or use some form of card-level architecture.

In a card-level system, designing for redundancy becomes much easier. If it is desired to only provide

fault-tolerance for the power system and not the load circuit, two DC/DC converters can be paralleled. Isolation diodes are required on the outputs to isolate the failed converter. But since the on-card DC/DC converter can achieve reliability levels in the same range as the load circuitry, it usually makes more sense to duplicate both of them rather than just the DC/DC converter. This approach often takes the form of redundant cards, so that both the DC/DC converter and the load circuit are replaced in the event of a failure. This also eliminates the need for the isolation diodes and for the detection and fault isolation mechanisms to differentiate between load faults and power faults. Either or both are easily corrected by replacing the failed card. When the repair strategy calls for hot-plugging - either of the AC/DC converters or of the circuit cards - provision must be made for controlling the surge currents that can result. Fortunately, this is easily accomplished with either simple discrete circuitry or with one of the many available 'hot plug controller' ICs. Figure 3.13 shows the overall fault-tolerance approach for a card-level DPA system.

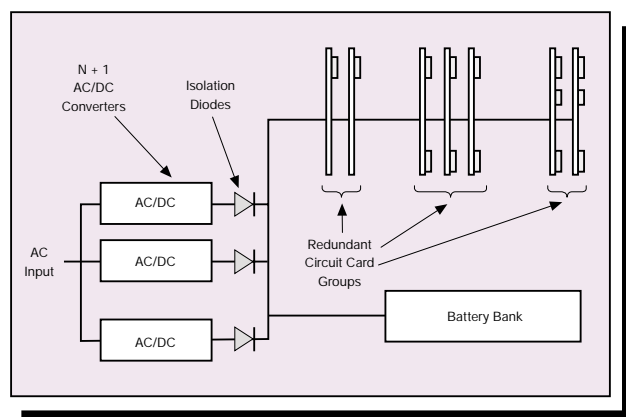


Figure 3.13 - Fault Tolerant Design for Card-Level DPA System