Ericsson Microelectronics

The **POWER**BOOK

A designers' guide to distributed power architectures using DC/DC power modules

4th revised edition

ERICSSON

Introduction

In the past years, development of miniaturized switching power supplies has brought on the practical implementation of decentralized power systems utilizing standardized power modules or power components.

Time to market is shorter, with ever increasing pressure on the power system designer to deliver cost effective power systems that are reliable, easy to manufacture, and pass regulatory qualification on the first attempt. No longer is it acceptable to wait several months or even years for the development of customized power system solutions. In short, equipment manufacturers demand denser, more reliable, more manufacturable power systems but with a reduced development schedule and budget. The above can be a difficult challenge. This overview of modern power system design and the technologies and components available to support it, is an endeavor to assist the designer with the above challenge.

The selection of standardized power converters is not always easy due to the wide range of products now available and the sometimes confusing performance claims made for them. Some suppliers emphasize power density, others switching frequency or converter topology, still others efficiency. Which is most important? How does the reliability of decentralized systems compare with conventional designs? What factors determine reliability? How are small DC/DC converters packaged within a system? What provisions are required for cooling? Why can't the maximum rated power of some converters be realized in practical systems? What are the real costs associated with power converter failures? Even though the newly available technologies and products offer exciting benefits, the list of questions such as those above seems to keep growing.

Our purpose here is to answer these questions and others, and provide a practical source of information for the power system designer - information that is based upon experience with actual systems and applications rather than just textbook formulas. It is our hope that the information contained here will be helpful in selecting a power system architecture that meets the needs of the product, selecting the appropriate power modules or components with which to implement the system, and in applying the selected modules and components correctly. The result should be a design that meets the product needs, requires minimal design and qualification time, and has an acceptable manufacturing cost.

We will begin with an overview of power system requirements, starting at the circuit level and extending into system controls and packaging and regulatory requirements. Next, we will address power system architectures in a generalized way before focusing in on decentralized power architectures. Tradeoffs between custom and standard power converters from a design and system management point of view will be discussed followed by an overview of commonly used converter topologies and their characteristics. We then look in more detail at the design and implementation of decentralized power systems including electrical design, thermal design, and other product considerations.

Reliability is ever more important in today's power systems, so we have devoted a section to several aspects of power system reliability, including prediction, design practices, and how reliability is affected by power architecture and hardware choices. The 'bottom line' for most system manufacturers is cost. We devote a chapter to cost analysis techniques useful for decentralized power system design, considering a product's 'life cycle cost', which consists not only of price, manufacturing and installment costs, but also very important indirect costs such as spare parts, service action and time-to-market. We also show how reliability information can be accounted for when doing cost analysis.
We have included a new section that attempts to clear the clouds and uncertainties regarding the design for conducted emissions compatibility, which is one of the more important design challenges. A section describing selection criteria for board-mounted power converters follows this. Finally a list of recent practical references is included for those readers wishing to expand upon the content supplied here.

This book is the fourth revised edition of the Powerbook and the result of a fruitful collaboration between the staff of Ericsson Microelectronics in Sweden and PowerSmith Consulting in USA, specializing in high-density power converter technologies and their markets.

Your reaction to this tutorial is desired. We welcome any feedback, comments, experiences or opinions on the subject matter, as well as any suggestions for improving upon the content that is presented here. You are invited to send any correspondence to:

Mr. Lars Thorsell
Strategic Marketing Director
Ericsson Microelectronics

SE-164 81 KISTA
Sweden
FAX +46-8-757 5482

Contents

1 Circuit Requirements

Power system design encompasses many requirements and criteria. As will be seen, it includes electrical, mechanical, thermal, control, diagnostics, reliability, safety, and regulatory considerations. Before looking at some of these aspects, it is beneficial to review perhaps the most basic of all requirements – the needs of the operating circuitry. The principal demands placed upon the power system by the operating circuitry are discussed below. More detail on how these requirements are accomplished by the power converter(s) and the power system will be provided in subsequent chapters.

Static Voltage

Each load circuit, whether digital or analog, is designed to operate over a limited range of DC voltage, outside of which circuit operation or performance is not guaranteed. One familiar example is the specification for TTL logic circuits, which require the voltage to be within the range of 4.75 to 5.25 V dc for guaranteed operation. This voltage is measured at the input pins of the circuit package, so it is dependent upon several factors, including the load and line regulation of the power converter, voltage setting accuracy of the converter, temperature drift, component aging, and distribution losses. It is important not to confuse this requirement with load regulation, which is a requirement placed upon the converter and not a circuit requirement.

In order to achieve high speed and low power dissipation, recent microprocessor chips require static voltage levels of 1 to 3.3 V dc with a static tolerance of as little as ±50 mV. These types of requirements place severe demands upon the power system. A generalized view of the static voltage and associated tolerance is shown in figure 1.1.

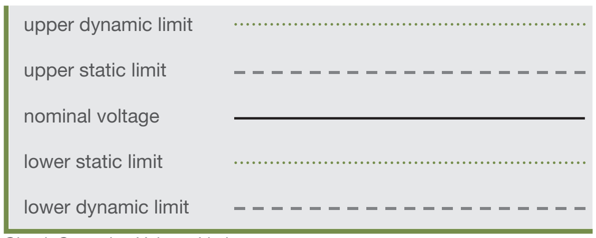

Circuit Operating Voltage Limits

figure 1.1

Static Current

Static current is the steady state load current drawn by the circuitry. The value of this parameter will be a key determinant of the current rating of the power converter. The converter along with the associated power distribution system must be able to supply this current without going outside the boundaries of the static voltage specification.

In practice, it is much more difficult to accurately obtain the operating current information than the static voltage. The current depends upon many variables, including manufacturing tolerances of the components, the operating conditions of the circuit, and environmental conditions such as temperature and supply voltage. In real systems, the maximum static current for each voltage level is typically significantly less than a summation of the specified maximum currents of each circuit element.

Dynamic Response

Many types of circuits demand additional input current for short periods of time over and above the static current requirement. One common example is simultaneous switching activity is allowed to exceed the static limits during these periods of dynamic current demands, as shown in figure 1.1. In most cases, the power system is designed to handle the high frequency dynamic current demands by means of decoupling capacitors located in proximity to the circuit packages. The power converter can handle longer duration dynamic demands. Figure 1.2 depicts how the dynamic current requirement is shared between both the decoupling capacitance and the power converter. Note how the initial dynamic current demand is supplied by current flowing out of the decoupling capacitor. After the output voltage of the converter recovers, the converter recharges the capacitor.

A more severe dynamic response requirement occurs in systems, such as high-speed microprocessors, that incorporate power management functions. To optimize the power consumption over time, circuits are put in a standby state. Upon command they are required to start up immediately which creates high dynamic currents with ramp rates exceeding 30 A/μs, during which the supply voltage should be within a specified limit.

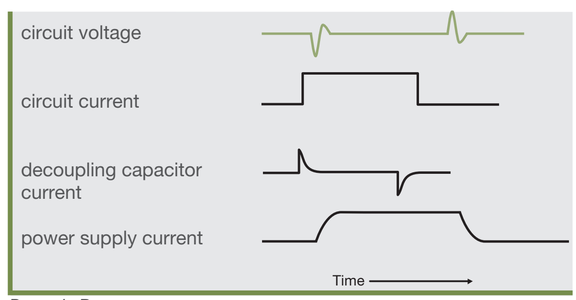

Dynamic Response

figure 1.2

On/Off Control and Powering

In some cases the supply voltage to the circuitry, or parts of the circuitry, needs to be capable of being turned on or off to select or activate circuit functions. The voltage output of the converter must ramp up within a specified time interval but without excessive overshoot. In systems requiring more than one voltage level, the voltages must sometimes be applied and/or removed in a certain sequence to avoid undesired conditions such as excessive power dissipation or latch-up in the load circuitry. This is normally implemented by means of the on/off control of each voltage level under supervision of a power controller.

An example of sequencing of two voltage levels by a power controller is shown in figure 1.3. After the first voltage is verified to be within its static regulation limits, the second voltage is ramped up. A similar approach would be used for turning off the two voltage levels in the desired sequence. Sometimes the required sequencing parameters can be satisfied without delaying the start-up of the second voltage. For example if the circuitry requires that one voltage level must always be more positive than the second voltage level, they can both be started at the same time if proper controls are in place on their ramping characteristics. An example of this is given in figure 1.3.

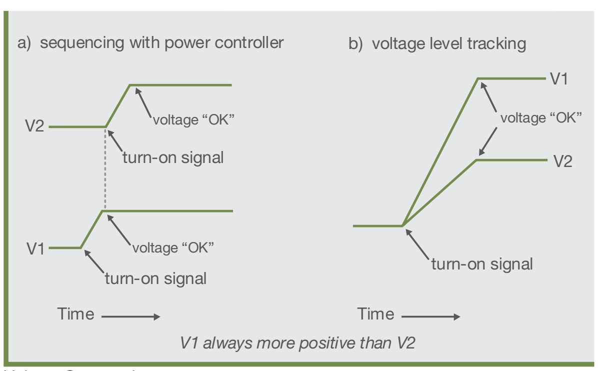

Voltage Sequencing

figure 1.3

Ripple and Noise

Ripple is an AC component that rides on the DC output of the power converter. It is typically larger in a switching regulator than for the older linear regulator designs. It is an artifact of the power converter switching and filtering activity, and has a frequency of some integral multiple of the converter operating frequency, depending upon the converter topology. Ripple values of less than 100 mV peak to peak are commonly achieved with today's converter designs. Ripple is relatively unaffected by load current but can be decreased by external filtering.

Noise occurs at two or more times the converter operating frequency, but it takes the form of short bursts of high frequency energy. It is caused by the need to quickly charge and discharge small parasitic capacitances within the converter during various parts of its operating cycle. Its amplitude is more variable than that of ripple, and can be dependent upon the load impedance, external filtering, the location where the measurement is made, and the measurement technique and bandwidth. Figure 1.4 is an example of the ripple and noise of a typical switch mode power converter.

Ripple and noise are of significance to the circuit in two respects. First, they are a contributor to the total instantaneous voltage present at the circuit package pins, and consequently need to be taken into account in this regard. Secondly, since they are AC components, they sometimes can be coupled into circuits and affect their operation. This is especially true for analog circuits. In many cases, the bandwidth of the circuit and the ripple frequency is such that ripple appears to the circuit as a varying 'DC' power supply voltage.

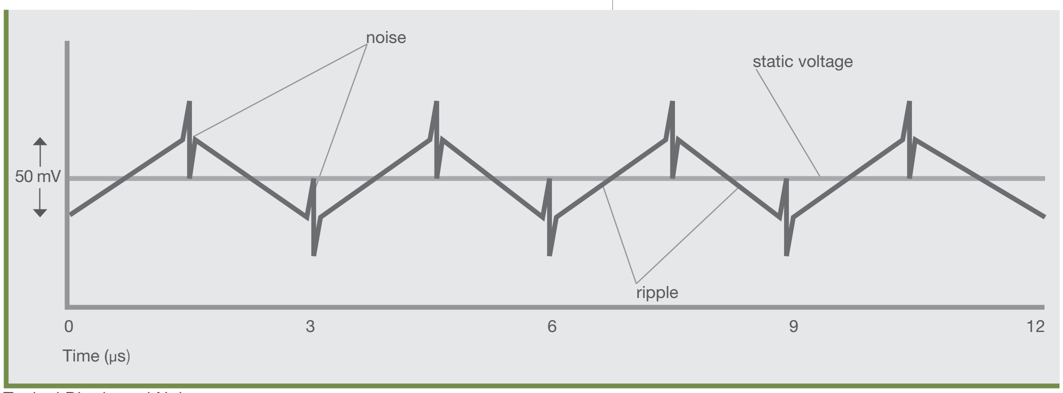

Typical Ripple and Noise

figure 1.4

2 Power System Requirements

The items in the preceding section were required for operation of the circuitry. In order to construct real-world products and systems, however, there are other considerations and requirements that tend to complicate the design process. These requirements evolve from the need to have a product that is safe to operate, reliable, and degrades gracefully in the event of a failure. In addition, every actual product needs to have a physical package and has thermal constraints that need to be satisfied. We include items such as these in the category of power system requirements. The most common power system requirements are listed below.

Current Limiting

Many failure modes in operating circuitry and in power distribution networks result in short circuits between the power source output and return. This can create currents within the system that are limited only by the maximum current capability and internal impedance of the power source. This high current can cause overheating and even danger of fire if it is not limited. Another unwanted result of short circuit faults could be damage to the power converter due to demands on it beyond its capabilities.

Almost all modern power converters incorporate some type of current limiting to address the above problems. The converter will be designed to detect when the output current reaches a certain level, which is above the specified maximum operating current but below the value at which damage to the converter would occur. After detection, the converter is usually designed to either limit its output current to this value or to turn off. It can also be specified that the converter issues a signal to indicate that an overcurrent condition has occurred.

At the system level, especially for multi-board systems, it is often necessary to use fuses to limit currents into sub-assemblies powered by a centralized power source. The fuses also provide the function of electrically isolating failed sub-assemblies so that they do not bring down other parts of the system and propagate faults.

It should be noted that current limiting tends to be a much more difficult problem with traditional centralized power systems than with decentralized systems. Decentralized power with board mounted distributed converters resolves many of these issues and makes for easy implementation of the remaining current limiting requirements. The approaches for accomplishing this will be discussed in more detail in the section on decentralized power system electrical design.

Overvoltage Protection

Except for the most basic unregulated power supply, all power regulators and converters have some fault modes that can result in an increase in output voltage above the desired DC value. This can impose excessive DC voltage on the load circuits and create permanent damage. For example, TTL circuits, which are guaranteed to operate up to 5.5 V, can withstand at least 7 V without incurring any permanent damage. To minimize the risks of circuit damage, most centralized systems incorporate some kind of overvoltage protection. This typically takes the form of a voltage detector in the power converter that turns the converter off in the event of an overvoltage condition. Another approach sometimes used is a ‘crowbar’ zener diode that conducts enough current at the overvoltage threshold to activate the power supply overcurrent shutdown.

As with overcurrent protection, overvoltage protection is a less severe problem with decentralized power system architectures. In decentralized systems, each power module provides power to a small part of the equipment’s total circuitry. In many cases the power module and the circuitry it powers are contained on the same Printed Board Assembly (PBA). In the event of a failure of the circuitry or the power module, the entire PBA is replaced. Since a power module overvoltage fault would not be propagated outside of the PBA, overvoltage protection for the power module is not needed in this situation.

Since the addition of overvoltage protection to a power module requires the addition of components, additional cost, and increased failure rate, its incorporation is not an automatic decision. In general, suppliers of DC/DC converters with power levels in excess of 100 watts tend to incorporate overvoltage protection, since these converters typically power circuitry on more than one PBA. In order to optimize cost and reliability, suppliers of lower power modules may or may not include overvoltage protection. In cases where it is not included, it can easily be added externally by the user if required. A simple voltage detector, the output of which drives the power module remote enable input, can be used to implement overvoltage protection with a user defined trip point.

Galvanic Isolation

In order to guarantee personnel safety, most systems will require that the operating DC voltages be galvanically isolated from the powerline voltage. This is accomplished by means of transformers in one or more power converter stages. In the case of a decentralized telecom system, for example, the off-line rectifiers contain transformers that provide safety isolation of the intermediate bus voltage. Isolation is also often used in DC/DC converters and power modules, both for its safety advantages and for increasing the flexibility for configuration of grounding connections. In other cases, the final stages of DC/DC conversion may be non-isolated, with a common input and output ground connection. Voltage regulator modules for use with high performance computer chips, for example, are usually non-isolated. Figure 2.1 shows a typical decentralized system that contains isolation at the front-end AC/DC converter and both isolated and non-isolated DC/DC conversion devices.

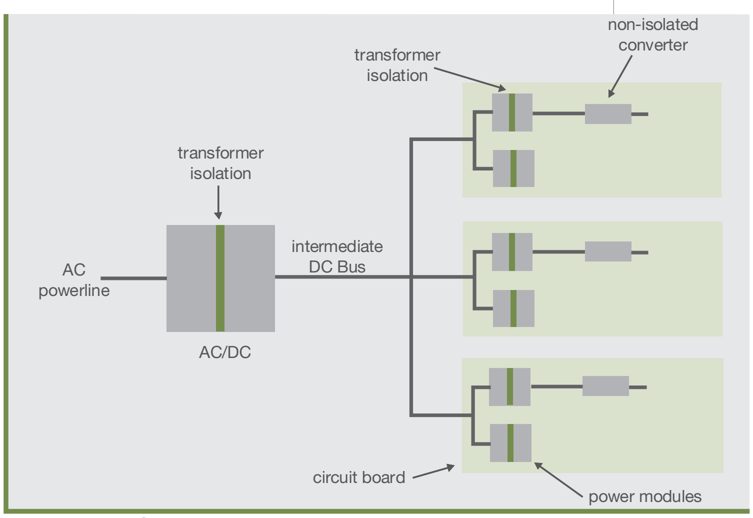

Typical Isolation Scenario

figure 2.1

Diagnostics and Fault Isolation

For very simple products such as consumer electronics, it is assumed that any type of fault within the product, including the power supply, will result in the entire unit either being discarded or brought into a repair center for service. In such products, there is no need for imbedded provision for quickly doing fault diagnostics. The situation is different for industrial, telecom, and data processing equipment and systems. The availability performance of these systems is critical, and down time results in loss of function of attached equipment and subsequent loss of revenue. When failures occur in these systems, it is imperative that they are isolated and repaired very quickly, at the end-user's site.

In order to facilitate rapid repair, the system is usually partitioned into sub-assemblies that are stocked as repair items and designed to be easily and rapidly replaced in the system. These sub-assemblies will be referred to as Replaceable Units (RU). A specific case of RUs are Disablement Units (DU). A DU is defined as the largest permissible unit to be defective in the event of a single failure. The power system design is affected by this partitioning. In systems with one power supply assembly, what is usually needed is a way of determining if the power supply is faulty or if the fault condition is elsewhere in the system. The situation is more complex in larger systems with multiple supply voltages implemented with centralized power system architectures utilizing several power assemblies. Diagnostics and fault isolation in these systems can involve significant amounts of hardware, software, documentation, and repair personnel training.

The problem becomes easier again with decentralized power architectures. In these systems, power conversion is often packaged in the same DU or RU as the load circuit it powers, and the entire function is replaced in the event of a failure. This normally eliminates the need for and cost of hardware and software to differentiate between power converter and load circuit faults.

If more sophisticated diagnostic capability is required, "voltage margining" is a very desirable capability of a converter or power module. Voltage margining allows the output voltage of the converter to be adjusted both upward and downward from its normal nominal voltage setting. This adjustment can be done automatically by means of switched jumpers, resistive programming, or a programming bus. By forcing the converter output voltage beyond its normal range, the operating margin of the load circuitry can be determined. If all is well, this test will verify that the system should exhibit additional robustness over time as both load and power components age.

Efficiency

Every power system specification will contain several references to the required efficiency of the power converters. Efficiency is exceedingly important for several reasons. It determines the losses in the system and the amount of cooling required. It determines now much utility power is "wasted" rather than being used for the desired purpose. It determines the physical package sizes of both the power converters and the final system. It will determine the operating temperatures of the components and the resultant system reliability. All of the above factors will contribute to the determination of the total system cost, both hardware and field support.

It should be clear why efficiency is so important. That is why manufacturers of power converters put so much design effort into maximizing the efficiency of their products. One of the latest trends in this regard is the usage of synchronous rectification in many high-performance power modules in place of the traditional Schottky rectifiers. Even though this approach adds some additional complexity and hardware cost, the total system cost of ownership can actually be improved because of the increased efficiency and reduced power losses.

Thermal Considerations

A very important part of the power system design is the analysis of heat transfer and the design of thermal interfaces to insure that both the power converters and the remainder of the system operate reliably when the product is exposed to its specified environmental extremes. In spite of ever increasing efficiencies, advanced components and packaging are making power converters smaller and smaller so that the thermal density, in terms of watts dissipated per cm^3 can be over an order of magnitude higher than that of older technology. For example, figure 2.2 compares a conventional discrete open-frame AC/DC power supply with the Ericsson PKJ power module. Both devices can supply 3.3 V at 30 A.

Model	AC/DC Generic Open-Frame	DC/DC Power Module Ericsson PKJ 4910 PI
Output Power	99 W (3.3 V @ 30 A)	99 W (3.3 V @ 30 A)
Efficiency	75%	89%
Dissipation	33.00 W	12.24 W
Dimensions (mm)	203 × 107 × 38	61.0 × 57.9 × 12.7
Volume	825.4 cm^3 50.37 in^3	44.9 W/cm^3 2.74 in^3
Thermal Density	0.04 W/cm^3 0.66 W/cm^3	0.27 W/cm^3 4.47 W/cm^3

Thermal Management Problem

figure 2.2

In exchange for its advanced packaging, its thermal density is about 7 times higher in spite of much better efficiency. Unless the amount of power dissipated, the heat transfer mechanism, the thermal impedances, and the ambient conditions are known, the power system design can be jeopardized. Undesired results can include converters that shut down due to overheating, elevated temperatures internal to the product and reduced product reliability. This area has been the source of many problems when attempting to implement systems with high density converters.

Unfortunately, some suppliers of DC/DC converters compound the problem by advertising very optimistic power capabilities without explaining (except perhaps in small print) that very elaborate and physically large heatsinks are required to achieve this performance. The operating temperature range of some converters is also very limited, with derating required beginning at ambient temperatures of 40 or 50 °C. Fortunately, there are power modules available that are designed to operate over wide ambient temperature ranges without extensive external cooling. The section on decentralized power system thermal design gives an overview of the system thermal design process and some examples of thermal analysis using actual parameters from currently available power modules and system implementations.

Packaging Considerations

Another aspect of power system design that can be critical to the success of a product is packaging design and building practices. The selection of power system architecture can affect many of these decisions. More often, the power system design will need to comply with product building practices that are already in existence. Decentralized systems will tend to have a larger number of physically smaller power assemblies than will centralized systems, which tend to have a small number of large and heavy assemblies. These extremes dictate different manufacturing processes, with decentralized systems being more conducive to automated manufacturing.

Indeed, many recent board-mounted converters are now available in SMT versions that allow for very cost-effective manufacturing using the same processes as other components on the board. Power components in decentralized systems can be made to be more independent of the building practices and adapt well to a wide variety of conditions including limited circuit board area and small board-to-board spacing. Centralized approaches cannot offer this flexibility. The weight of the power assemblies is important not only for manufacturing process purposes, but for its impact upon the system's susceptibility to vibration and shock. Heavy power supplies need extensive (and usually manually installed) retention mechanisms to survive the product shock environment.

Finally, the packaging and partitioning of the power system will be influenced by the system diagnostic strategy and maintenance philosophy. Segmentation into the proper DUs or RUs along with other system functions is required. Decentralized power permits more design freedom in this regard, and this additional flexibility allows the product designer to create systems that are very easy to diagnose and maintain.

Regulatory Requirements 3

A product or system marketed in today's environment must comply with an ever-increasing array of standards and regulatory requirements. These requirements were put in place to satisfy concerns about personnel safety, environmental impacts, and compatibility between electronic devices. The power system designer must understand the requirements in order to design compliant systems.

Overview

Regulatory requirements tend to be one of the more difficult areas of electronic equipment design to understand and implement. There are several reasons for this:

- Some requirements are quite complex technically, requiring specialized knowledge to understand and apply.
- The requirements are often written in a form that is difficult to interpret. There are many exceptions and exclusions that are not clearly articulated.
- There are a large number of agencies involved, many of them specific to one country or group of countries. Requirements vary and sometimes conflict from jurisdiction to jurisdiction.

- The requirements are rapidly evolving, with new regulations coming into practice every year.
- There are very few specific product standards for power supplies and no specific standards for on-board DC/DC converters. Power converter requirements are therefore often, by default, based upon system requirements that were never intended for application to individual components such as power modules.

Because of this situation, most companies doing product development work have a person or department that is responsible for only keeping abreast of the latest regulatory requirements, proposals for new requirements, and the procedures for testing and certifying compliance to standards. In spite of this expert assistance, the power system designer is not immune from wrestling with these types of issues. In fact, he or she is often at the center of the action. While most of the requirements do not focus directly on the power system, some of the more difficult standards to meet, such as Electromagnetic Compatibility (EMC), are often highly dependent upon the power converter performance.

Our intent here is not to replace the regulatory requirement expert. This is not possible due to the rapidly changing nature of the requirements and the product-specific nature of many standards. The system designer, assisted by the available standards personnel, should determine the set of requirements that are appropriate for the individual product depending upon the product type and the locations and time frame in which it is intended to be marketed. What will be presented here is in the nature of a 'check list' of the types of requirements that should be investigated, along with some very general commentary on how they impact upon the power system and power converter assemblies. More detailed information on the EMC standards and requirements can be found in the EMC chapter of this book and in Ericsson Applications and Design Note publications.

The following types of regulatory requirements and agencies will typically be encountered at the product and system level in Information Technology and Telecom (IT&T) equipment:

- EC Directives (European Community). If a product is intended to be marketed in the European Community the manufacturer, or the company responsible for the product, has the obligation and responsibility to design and manufacture the product in accordance with the requirements laid down in the relevant Directives. Compliance to the essential requirements is demonstrated by following the procedures for certification, eg. self certification if ISO 9000 certified, to the applicable Harmonized Standards and issue a 'Supplier's Declaration of Conformity'. After this procedure is completed, the manufacturer can apply the CE mark and market the product freely within the European Community (European Union and former EFTA countries). A Harmonized Standard is a technical specification that has been adopted by CEN or CENELEC (European Committee for Electrotechnical Standardization) and has been published in the Official Journal of the European Community (OJ). As of 1 January 1996 the CE mark is required on all apparatus, i.e. finished products with an intrinsic function intended for the final user. A component is defined as any item that is used in the composition of an apparatus and should be excluded from the scope of the Directives. A DC/DC Converter or power module is generally considered to be a component.

- UL (Underwriter's Laboratory) approval. A safety approval that is almost universally obtained to market electrical products within the USA. A UL approval can now also be obtained through the CSA.
- CSA (Canadian Standards Association) approval. A safety approval that is required to market an electrical product within Canada. A CSA approval can now also be obtained through the UL.
- Telcordia (formerly BellCore) standards for telecom equipment in the USA.
- ETSI (European Telecommunications Standards Institute) standards for telecom equipment in Europe.
- Safety Standards. These include EN 60950 'Safety of information technology equipment including business equipment', the new bi-national standard CSA-C22.2 No 950/UL1950 'Safety of Information Technology and Telecommunications equipment' and CSA C22.2 no.234-M90 'Safety of component power supplies', containing requirements to prevent injury or damage due to hazards such as: electric shock, energy, fire, mechanical, heat, radiation and chemical. UL 60950 will soon replace the earlier UL safety standards. As of 1 January 1997 the EC Low Voltage Directive (LVD) 73/23/EEC and the amending directive 93/68/EEC requires the manufacturer to make a declaration of conformity and affix the CE mark if the product is intended to be placed on the Community market. The manufacturer must compile a technical file. The file includes a general description, drawings, diagrams and an operational instruction to demonstrate the means taken to ensure conformity. The Harmonized Standard generally used for the LVD is EN 60950.
- Acoustics. These standards define maximum audible noise levels that may be emitted by the product. The biggest power system contributor to the acoustic level is usually the air-moving device in forced convection cooling systems.
- ESD (ElectroStatic Discharge). Verification that the product is immune from the effects of high voltage low energy discharges, such as the static charge built up on operating personnel.
- Inrush current. The current vs. time waveform imposed on the powerline when a power converter is turned on or plugged in.
- Input transients. Ability of the power supply to survive without damage or operate through temporary variations in powerline voltage. These transients can be in either direction (undervoltage or overvoltage).
- Powerline Standards. Requirements and specifications that are in place to protect the quality of the powerline, including harmonic distortion and phase balance.
- EMC Standards. The most commonly used international standard for emissions is C.I.S.P.R. 22 'Limits and methods for measurement of emissions from ITE', from which the detailed agency requirements are derived. Most of the commonly used immunity standards are contained in various sections of EN 61000. As of 1 January 1996 the EC Directive 89/336/EEC on EMC requires the manufacturer to make a declaration of conformity and affix the CE mark if the product is marketed in the European Community. The presumption of compliance with the Directives is commonly based on the self-certification to the relevant harmonized standards. Examples of harmonized product emission standards are EN 55022 for IT&T equipment and EN 55011 for industrial, scientific and medical equipment. In the USA, FCC part 15 defines the requirements for IT&T equipment.

- Conducted Susceptibility. An EMC requirement demonstrating the ability of the power supply to operate through noise coming into the system from the powerline.
- Conducted Emission. An EMC requirement demonstrating that the product power system does not adversely affect the powerline quality by injecting excessive noise into it.
- Radiated Susceptibility. An EMC requirement demonstrating that the product's operation is not adversely affected by electromagnetic fields imposed upon it from external sources.
- Radiated Emissions. An EMC requirement demonstrating that the product does not radiate excessive electromagnetic fields.

Many, but not all, of the above requirements impact upon the power supply design. Centralized AC/DC converters are affected differently than distributed DC/DC converters. Figure 3.1, in a general way, shows how the regulatory requirements affect both types of power assemblies and the product itself, assuming that the product power system is implemented with a decentralized architecture.

As can be seen from the figure, the power assembly that interfaces to the powerline is subjected to the most requirements and standards. The distributed DC/DC converter or power module is isolated from many of the requirements. Other requirements, such as acoustics and ESD are generally not a problem with power assemblies, but must be verified at the system level for the entire system. One requirement that typically affects both AC/DC and DC/DC power converters, as well as the product itself, is the Radiated Emission EMC requirement. Of all the requirements mentioned, this one will often require the most effort in terms of analysis and verification testing. This requirement must be verified at the system level.

Distributed DC/DC Converters

Product standards for DC/DC converters, particularly standards for usage within decentralized power system architectures, do not exist. This makes it more difficult to specify and design with standard DC/DC converters. However, to avoid nonconformity of the finished end product, the system manufacturer often requires conformance to a set of Safety and EMC standards that are equal to the end product or system requirements. The DC/DC converter manufacturer frequently finds that it is impossible to apply these requirements to the DC/DC converter without additional measures. The consequence is that the DC/DC converter manufacturer sometimes must issue a manufacturer's instruction. This instruction includes the information required to enable the use of the converter in accordance with its intended purpose and

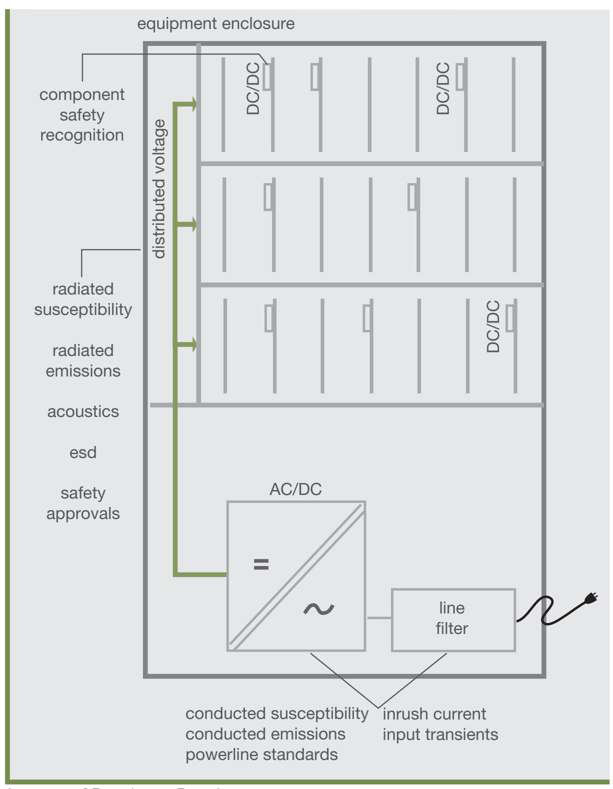

Impacts of Regulatory Requirements

figure 3.1

the necessary additional measures the end product manufacturer has to take to avoid nonconformity of his final product due to the DC/DC converter.

Safety Requirements

Standard DC/DC converters are generally required to be verified and recognized in accordance to EN 60950 and CSA-C22.2 No 950/UL1950 and certified according to CSA C22.2 no. 234-M90 level 0. UL/CSA recognition, although legally not required, is usually obtained for DC/DC converter products. The benefit is that the end use product can obtain UL/CSA listings more easily and quickly if the product is composed of subassemblies that have previously been examined by UL and/or CSA.

In EN 60950 and UL1950 there is a definition of the Safety Extra Low Voltage (SELV) requirement, which affects the system isolation requirements. Basically SELV implies safety isolation from hazardous voltages with Double or Reinforced insulation and an extra low voltage level, not harmful to the human body. Test voltages for the appropriate grade of insulation and working voltage are also specified. Circuits are considered as SELV circuits if the voltage is less than or equal to 60 V dc. There is no isolation requirement, other than operational insulation, if the input and output of the DC/DC converter are SELV circuits. If the input voltage is greater than 60 V dc there is a requirement of Basic insulation in the DC/DC converter or alternatively a requirement of reinforced insulation in the power supply that isolates the input from the AC mains.

The 1,500 V dc isolation voltage sometimes specified for DC/DC converters is required in certain telecom applications where there is a risk of high voltage transients due to lightning interference in the telecommunication network and the Bonding Network (BN) in smaller access nodes. Also, other AC mains disturbances can cause high voltage transients in the BN. A BN connects all metallic parts in the equipment, and is connected to the main system earthing terminal, where the Protective Earth is connected.

EMC Requirements

From a DC/DC converter point of view, the radiated emissions requirements, even though imposed at a product or system level, can represent a difficult design challenge. The high levels of rapidly changing voltages and currents within a switchmode power supply can create EMC problems for the product if the converter is not carefully designed with attention to many circuit design and layout details. In addition to the formal EMC requirements, the ‘near field’ radiation characteristics of DC/DC power should be evaluated for applications where the converter is placed in close proximity to load circuitry. One commonly occurring instance of this type of packaging is the use of DC/DC power modules in “Power per Board” decentralized architectures with small board-to-board pitch. For such applications, it is important to use power modules that have been designed with this type of packaging in mind. Today’s better DC/DC power module designs have an extensive history of successfully meeting the system’s EMC requirements. This is true for even the “open-frame” designs without self-contained shielding.

Generally, in a decentralized power system with distributed DC/DC converters, the overall system design will be optimized if the EMC requirements are complied to at the end product or system level. This can be accomplished by using centralized protection circuitry for input transients and conductive susceptibility, and implementing shielding for radiated emissions at the equipment enclosure level.

The most frequently used references for DC/DC converter EMC conducted emission specifications are EN 50081 and FCC part 15. The emissions standards have two performance levels, curve A and curve B, with curve B being the more stringent. Curve B is almost mandatory in most of the current IT&T equipment unless it is designed to be “hard wired” into the power mains. In order to comply with level B at the system level, DC/DC converters themselves are normally required to meet curve A. Curve B is then met by means of additional filtering within the system. FCC conducted emission limits are equivalent to, or less stringent than, the EN limits defined by curve A, except in the range of 0.45 to 0.50 MHz. Another difference is that EN 50081 covers the range of 0.15 to 30 MHz while the FCC standard covers 0.45 to 30 MHz. In actual practice, designs meeting the EN 50081 requirements will meet the appropriate FCC requirements. EN 61000 contains standards that address conducted immunity of the power converter. These standards define requirements for items such as ESD, EFT (Electrical Fast Transients)/Burst, Surges, and continuous noise.

Ericsson DC/DC power modules are designed to comply with the currently accepted regulatory standards. Ericsson is a leading supplier of telecom equipment, and designs to meet EMC and safety standards worldwide for these demanding applications. As a consequence meeting the product level requirements of your system, whether telecom, information technology, or industrial will be facilitated by the usage of Ericsson’s line of converter products.

Power System Architectures 4

The most basic and important decision to be made in the design of a power system for electronic equipment is the selection of the power architecture.

Introduction

Selection of an appropriate power architecture is key to the success of the end product. Sometimes the power system designer has no choice, and is forced to accept an architecture imposed by previous decisions, system packaging, or schedule. Sometimes the designer can start with a 'clean sheet of paper', selecting the power architecture best suited to the product function. Most commonly, the situation is somewhere in between these two extremes, with some freedom available for partitioning of the power function, but with several constraints imposed by other aspects of the system design. Even in cases where severe constraints exist, the information in this section should be helpful to the system designer.

Sometimes even a small change in partitioning or power distribution can provide significant benefits to the final product.

We will first define and describe the types of power architectures, and then look at some examples of each type. Even though each system has its own unique requirements and limitations, the concepts presented here should prove to be generally useful in making system level power decisions.

Power Architecture Definitions

Centralized Power Architecture

A centralized power architecture is a power system in which all power related functions, from input power source to generation of the DC circuit voltages, are contained within one physical area. Centralized power architectures frequently contain the following elements:

- Multi-output customized AC/DC or DC/DC power supplies.
- Bus Bars or wire to distribute DC power to the circuits.
- Customized safety shielding surrounding the power area.
- Specialized cooling provisions for the power area.

Multi-location Centralized Power Architecture

Some systems contain many of the attributes of centralized power architecture, but are configured with the AC/DC or DC/DC converters, operating with a non SELV input voltage, located at two or more physical locations within the product. These systems are referred to as Multi-location Centralized systems.

Decentralized Power Architecture

A decentralized power architecture (sometimes referred to as distributed power) is a power system that is functionally and physically partitioned such that the final stage of power processing is located in correspondence to load functions and/or packaging. The final stage of power processing operates from a safety-isolated DC voltage. Decentralized power architectures frequently contain the following elements:

- Centralized AC/DC converter to interface with the AC powerline.
- The AC/DC converter provides the functions of safety isolation, DC conversion, noise suppression and power factor correction.
- An isolated intermediate SELV DC voltage (24 to 60 V) is distributed within the product.
- A standardized telecom battery voltage (24, 48 or 60 V) is distributed within the product.
- Individual load converters (DC/DC) are used for each load function or load package.
- The DC/DC converters are physically located at, or very close to, the load.
- The DC/DC converters are small, dense, and may or may not contain isolation.
- The DC/DC converters are standardized modules or components.
- Provision is made for easily upgrading or adding features to the product.
- Provision is made for redundancy in high availability systems.
- Provision for battery back-up or other technique to provide immunity from powerline or AC/DC converter faults.

A system does not need to contain all the above attributes in order to be considered a decentralized power system, but most decentralized systems contain several items on the list.

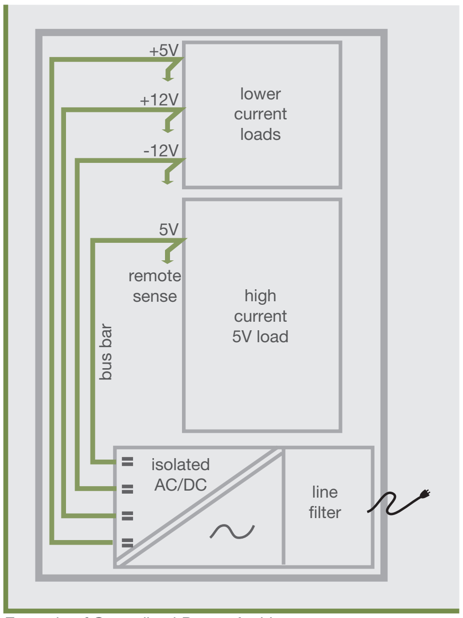

Example of Centralized Power Architecture

figure 4.1

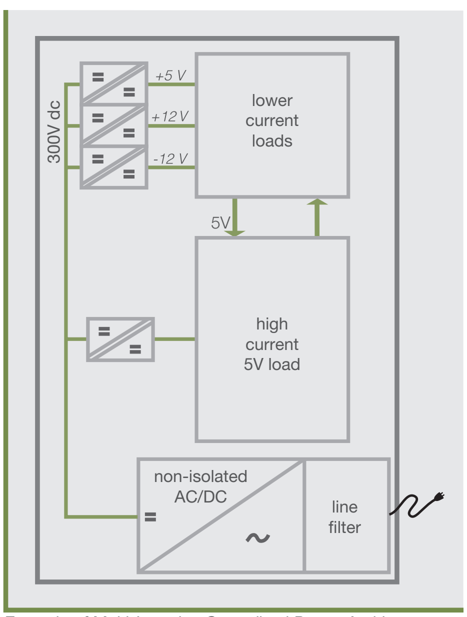

Example of Multi-Location Centralized Power Architecture

figure 4.2

Figures 4.1 through 4.3 show examples of each type of architecture. We have attempted to cover a wide range of Industrial, Information Technology and Telecom systems with the architectural definitions. With the need to encompass both AC and DC input systems and this spectrum of system types, the definitions are somewhat arbitrary and subjective. It is hoped that the discussion here will help the reader in understanding the spirit of the definitions and why the distinctions in architecture were made as they were.

Centralized Power Architecture

Perhaps the least controversial in terms of definition is the centralized architecture. These systems have all the power conversion functions located in one area of the equipment. Figure 4.1, for example could represent a mainframe computer. AC power enters the equipment enclosure, is processed through a line filter, and

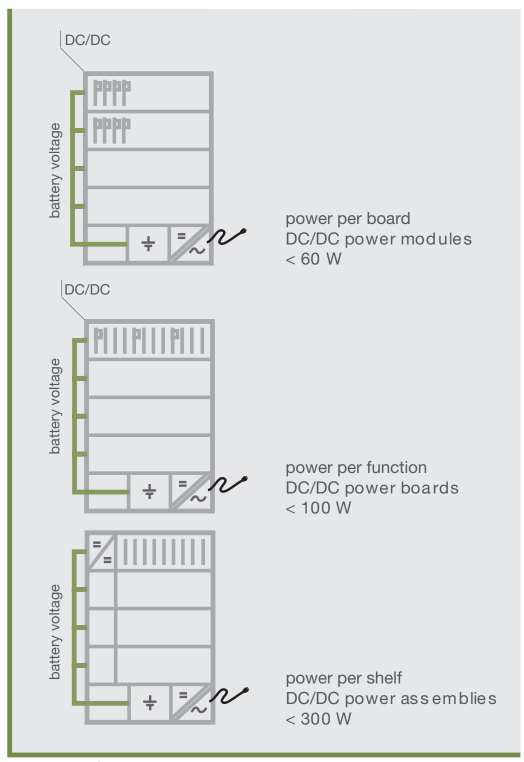

Examples of Decentralized Power Architectures

figure 4.3

then enters a large AC/DC conversion function which provides a total of four output voltages at a power level of perhaps 1,500 W. The AC/DC conversion is done by large, custom designed switching power supplies. The entire power area is enclosed by a safety shield due to the high voltages present in the AC end of the system and the large energy content stored in the output capacitor/filter area. Due to the concentration of power dissipation (about 300 W) in the power area, it is supplied with forced-air cooling that is ducted from fans elsewhere in the enclosure.

This product has a very elaborate distribution system. The high current 5 V load can draw currents of up to 200 A, but requires tight regulation (± 2%) from the power system. To guarantee this regulation, remote sensing is required, which senses the voltage at the load board and compensates for the voltage drop in the distribution system. This sensing adds complexity and possible failure modes to the system. To minimize power loss in the distribution, very heavy copper bus bars are required to bring the 200 A current from the power area to the load board. These bus bars need to be custom designed, manually installed with specific assembly torque requirements, and consequently add expense to the system. This distribution function is even more difficult and expensive with circuit voltages of 3.3 V and below. Distribution to the upper board, with lower current requirements, is done with a cable harness. Remote sensing is required here also, and is included in the cable. The custom made cable harness assembly adds product hardware cost and limits flexibility for changes to the power delivery system.

In order to service this equipment, a combination of automated and manual diagnostics is required. Failures must be isolated between the load boards, the DC distribution system, and the appropriate power converter. Repair can require replacement of large, heavy, and expensive assemblies such as custom switching converters and bus bar assemblies, which must be stocked on site if responsive repair is required. Even so, diagnosing and repairing a fault in the power system could take several hours of work by highly trained service personnel.

Another example of a centralized system, this one for Telecom applications, is shown in figure 4.4. Here the circuitry is packaged in two card cage assemblies in the upper portion of the equipment rack. Power (several voltage levels) is distributed to these card cages by discrete wire distribution from a centrally located power conversion area. The power converters operate from a 48 V bus that is provided from an off-line rectifier and filter assembly. An internally contained battery assembly backs up the 48 V bus. Even with the isolation from powerline faults provided by the battery backup, this system is still far from ideal from a reliability point of view.

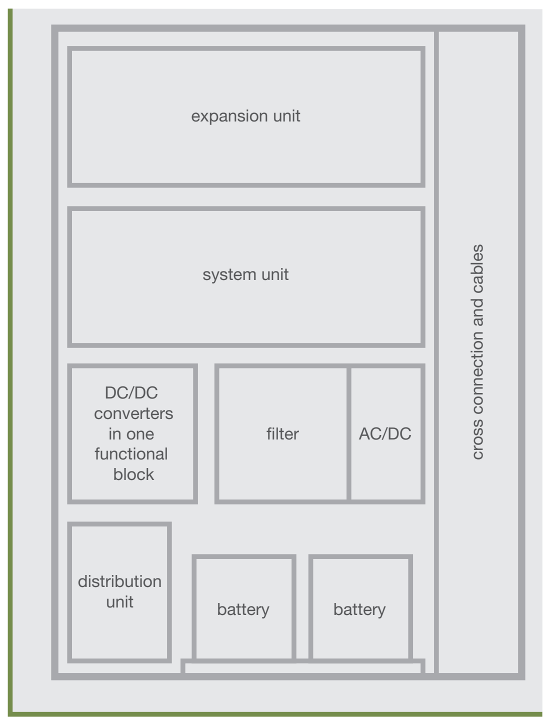

Centralized Telecom System

figure 4.4

Any failure in the DC/DC power converters will disable the entire equipment. The design operates through external power faults, but does not accommodate one of the more likely internal failure modes.

The most ubiquitous example of a centralized power architecture in an Information Technology application is the main power supply in a personal computer. Figure 4.5 shows a typical desktop personal computer and its power supply. These supplies are custom designed multi-level switching regulators, in the power range of 60 to 300 W. They are interfaced to the rest of the system by means of individual cable assemblies that terminate at the planar board, the disk drives, and other major functional areas. Because of the power dissipation in the localized area of the power supply, there is a fan assembly included in the power supply that provides forced convection cooling.

The hardware cost of custom multi-output power supplies can be rather modest. Other costs are not. Due to its custom design, significant development time is often required. The power converter design must anticipate and accommodate the maximum load currents to be encountered over the life of the product as

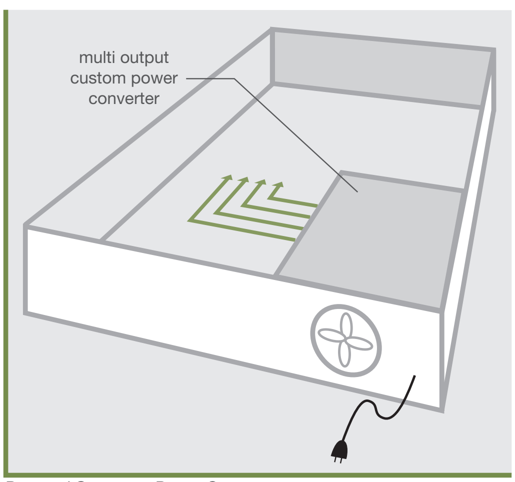

Personal Computer Power System

figure 4.5

the owner installs options. It is consequently overdesigned for the initial product. It is also very inflexible to changes in the system design during the development process. Any changes in voltage, current, or distribution requirements entails a lengthy change process in the custom power supply design. The power supply must be re-qualified for agency approvals after each design modification. The fan that is included within the power converter is a limited lifetime component with a relatively high failure rate. This will degrade the reliability performance, even if the remainder of the power converter is well designed and manufactured with high quality reliable components.

Multi-location Centralized Power Architecture

Unfortunately there is no easily identifiable criterion that separates decentralized systems from centralized architectures, and there are many system designs that are in a gray area between the two. Some of these systems would be considered centralized by some observers and decentralized by others. We have elected to call these systems Multi-location Centralized because they are basically centralized in philosophy, but have some physical partitioning.

Refer to figure 4.2 for an example of such a system. This is basically the same as the system in figure 4.1 except that the DC/DC converters are located at two physical locations in the system instead of just the one location for the centralized AC/DC converter. Note that the intermediate bus voltage is a nonisolated 300 V dc. This choice will require safety shielding of the entire 300 V distribution system. Also note that there is only one converter per output voltage for each main area of the system. If any of these converters fail, the entire system will be inoperable. There is no redundancy created by load/power partitioning as with most decentralized systems. The power converters remain approximately the same total size as they were in the centralized approach, with no finer granularity. It is for these reasons that we have not considered it to be a decentralized system.

The design does offer some advantages over that of figure 4.1. The heat load from the power converters is spread throughout the equipment more evenly. The DC distribution busses are simpler, shorter, and less expensive. However, the potential problems with field diagnostics and service discussed relative to figure 4.1 still exist with this approach.

Decentralized Power Architecture

Decentralized systems include the features of an isolated SELV intermediate voltage bus and some partitioning of the load and power into functional groupings. We will look at four systems that fall into this class of architecture and yet are significantly different in implementation.

Decentralized power is not only applicable to large systems. High performance personal computers contain decentralized power architecture in the form of a voltage regulator module (VRM) that supplies the operating voltage to the processor chip. Figure 4.6 shows how this is accomplished. The VRM is a non-isolated DC/DC converter that operates from one of the outputs of the computer's centralized multi-output AC/DC power supply. The VRM converts this voltage (usually 12 or 5 V) down to the core operating voltage of the processor, which is between 1.3 and 2.5 V. The output voltage is programmable via a bus so that the exact value can be adjusted to the processor chip being used and the operating conditions. The VRM approach is required due to the high current requirements of the processor (up to 50 A), the tight regulation requirements and the large dynamic current transitions as the processor goes between standby and operating states (up to 30 A per ms). This is an example of both functional and physical partitioning. The VRM only serves the processor chip (functional partitioning) and is located next to it (physical partitioning).

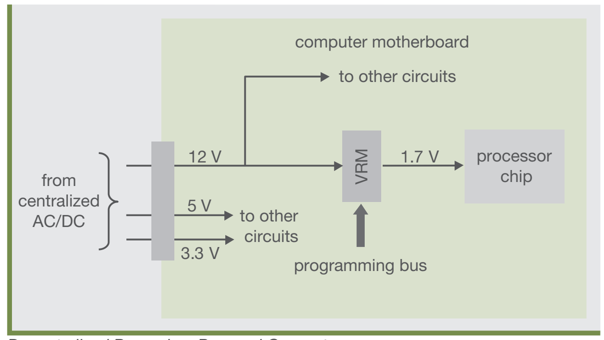

Decentralized Power in a Personal Computer

figure 4.6

The next three systems we will consider are all Telecom products configured in similar rack type enclosures. They exhibit increasing degrees of decentralization.

Figure 4.7 is a decentralized system with the DC/DC conversion function replicated on each shelf. This will be referred to as a 'Power per Shelf' architecture. Each DC/DC converter is probably in the range of 100 to 200 W, and provides the DC output voltages required to operate the entire shelf of electronics. Note that the bus voltage is isolated and is below the SELV limit so that it simplifies the requirements for shielding and safety covers. If there is redundant common function between shelves, the system contains redundancy and can survive failure of any one DC/DC converter without totally disabling the product. The battery back-up provides for operation during AC line power outages.

Diagnostics and service are simplified for this product compared to the centralized systems we discussed. Isolation of DC faults to a shelf is very easy. However, manual intervention will still be needed to isolate short circuits and overcurrent situations to either the DC/DC or to one of several circuit boards.

The next step towards the ultimate decentralized system is referred to as 'Power per Function', and is sketched in figure 4.8.

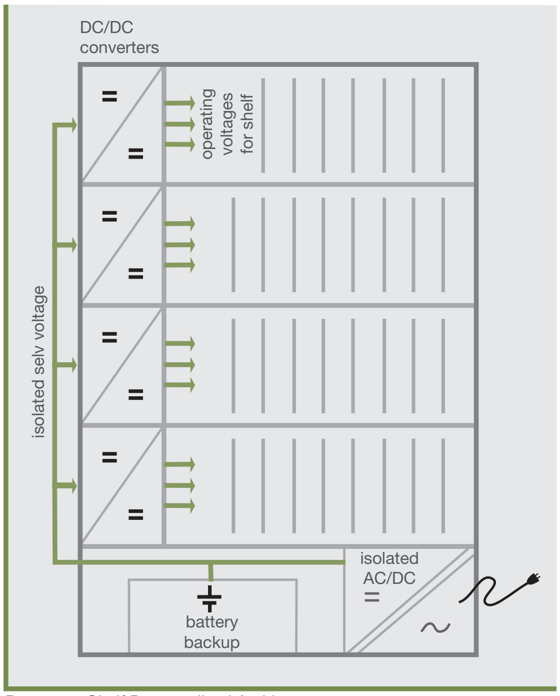

Power per Shelf Decentralized Architecture

figure 4.7

The front end and bus voltage distribution of this system is identical to that of the system described above. In this case, however, the DC/DC converters are mounted on boards similar to the load boards and are plugged into the card cage. The circuit boards are grouped into functional islands, each being powered by an adjacent board containing the appropriate DC/DC converter(s). The output power of the DC/DC converter boards will be in the range of 10 to 70 W. The distributed nature of the power converters will distribute the heat load, and as a result, the cooling environment could be significantly enhanced relative to the system shown in figure 4.7.

Diagnostics and service will be easier than for the previous system. The functional groups are smaller and easier to isolate. Replacement power converter boards are small, inexpensive, and easy to stock and replace. There will be significant replication of converter board part numbers so that a reasonably sized spare parts inventory will service the product very economically. The power converter boards look very similar to the other boards in the system and less like 'power supplies'. This is good news, and an indication that we are progressing towards the goal of functionally partitioned power conversion diffused throughout the system.

The ultimate decentralized system that can be economically implemented with today's technology is shown in figure 4.9. This system architecture, referred to as 'Power per Board', includes the DC/DC conversion function on each load board. The resulting diffused nature of the power dissipation will allow for a large amount of flexibility in the cooling system design. Many systems configured with 'Power per Board' architectures will be ideally suited to free convection environments, eliminating the need for, and negative reliability implications

Power per Function Decentralized Architecture

figure 4.8

of, fans and blowers. The DC/DC converters will be in the range of 2 to 10 W for free convection applications, and up to 100 W for forced convection applications (>1.5 m/s). The diagnostic strategy for the DC part of the system becomes very simple – there isn't any! There are no separate power converters to diagnose or replace. When a circuit board fails, the associated power conversion function automatically gets replaced.

In a system such as this, the power converters essentially become components on the circuit boards, a goal that has been eluding power system designers until recently. For this approach to be economically viable, the modules must be very reliable, very small, and inexpensive. With the availability of the Ericsson's line of power modules, many of them designed for SMT and automated placement, the power designer's dream is now reality. Many successful systems are operating throughout the world incorporating decentralized architectures and Ericsson power modules. These systems exhibit many advantages over previous approaches, some of which will be described in the following chapter.

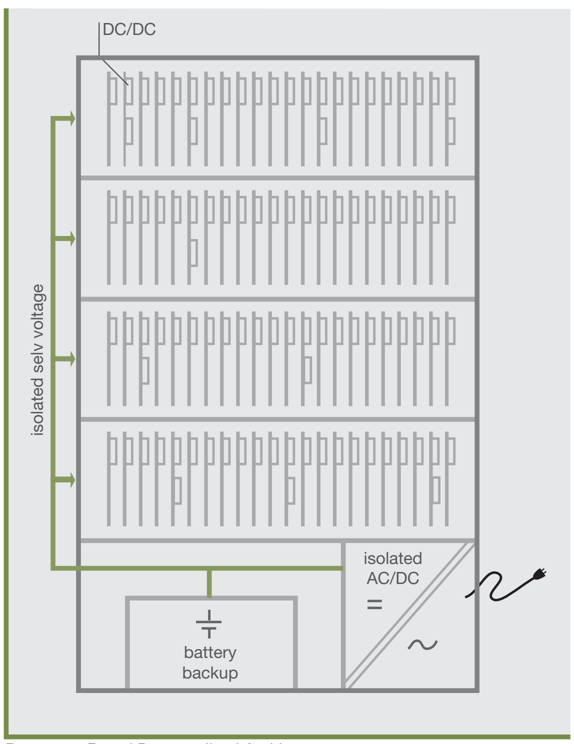

Power per Board Decentralized Architecture

figure 4.9

Decentralized Power 5

In the previous chapter, we defined the various types of power architectures, and gave several examples of how such architectures can be implemented in products. We saw that decentralized approaches resolve many of the problems of the more centralized designs. But the decentralized approach will only be economically viable if commercially available DC/DC converters have the proper characteristics. In this section, we will discuss the required technologies that make this possible. We will also explore in more detail some of the advantages of decentralized architectures. In this discussion, we will assume the 'Power per Board' implementation, which is the ultimate execution of the concept of decentralized power implemented with distributed power modules. Most of the advantages described will also apply, sometimes to a lesser degree, to 'Power per Function' and 'Power per Shelf" decentralized architectures.

Viability of Decentralized Power Architectures

Decentralized power is not a new concept. Many engineers and system designers recognized its inherent simplicity and advantages long ago. Unfortunately, it was not extensively used until recently due to the lack of suitable technology. In order to share circuit board real estate with load circuitry and do so economically, the power converter must possess the following technology-driven attributes: small footprint, low profile, high efficiency, superior electrical performance, manufacturing compatibility, high reliability low cost and low weight.

All of these attributes are now available and are exemplified in several of Ericsson's power module product lines. From the industry standard SMT packaging of the PKF products to the PKL line offering high power capability with extremely high efficiency and power density, there are Ericsson products that address all of these requirements. This product selection allows the construction of high performance decentralized power systems that meet very aggressive cost targets.

Advantages of Decentralized Power Architectures

Better Electrical Performance

Placing the final stage of power conversion close to the load circuitry provides some stunning performance advantages. The DC distribution system becomes much shorter and simpler, eliminating power losses in the distribution network. Better dynamic response performance is also achieved due to the lower inductance between the converter and its load. The proximity to the load also allows for good voltage regulation without the need for remote sensing in many applications. The VRM implementation for powering the processor chip in a personal computer, as described in the previous chapter, is a good example of how decentralized approaches can achieve levels of electrical performance not possible with centralized architectures.

Allows use of Standard Modules

The usage of standardized power conversion hardware is one of the biggest advantages of decentralized architectures. Compared to the custom designs associated with centralized approaches, standard power modules offer lower development cost, faster time-to-market, faster qualification, higher reliability, more flexibility, significantly lower technical risk, and very competitive hardware costs. These advantages are so pronounced that we have included a complete chapter to discuss them in more detail.

Automated Assembly Process

On-board DC/DC power modules are small and light enough to be assembled with automated assembly equipment, eliminating the manual labor traditionally associated with power converter installation. This results in significant cost savings as well as more reliable and dependable interconnections. DC/DC power modules are now available in SMT versions so that mixed packaging implementations are no longer required. It is now possible to achieve automated SMT manufacturing compatibility with power modules up to 30 W.

Better Backpanel Utilization

For rack type systems, where circuit boards are plugged into a backpanel, the decentralized approach results in a significant advantage. Rather than distributing low voltages such as 3 V or 5 V at high current, a higher voltage such as 48 V is distributed. This results in reduction in backpanel currents by a factor of 10 or more, and requires much less copper being allocated to the power function. The saved backpanel areal capacity can then be used for signal trace distribution.

Better Connector Utilization

Applying the same reasoning as above, the currents through the connector pins between the backpanel and circuit boards are reduced by about a factor of 10. This means either that more pins are available for signals, or that a smaller connector can be used with resultant cost savings. The increased popularity of decentralized systems has resulted in the availability of standardized connector systems specially designed for reliable plugging of intermediate bus voltages.

Distributed Heat Load

Rather than concentrating power converters and their resulting power dissipation at one location in the system, decentralized power tends to diffuse it throughout the system. This can reduce cooling air requirements or in some cases allow free convection cooling with no need for fans or blowers. The result is higher reliability. The cooling requirements and interfaces of power modules are well understood and documented. This makes it relatively easy for the power system designer to achieve a reliable and conservative system thermal design without the need for specialized hardware.

Ease of Battery Backup

More and more systems in the Information Technology and Industrial arenas are now recognizing the advantages of having a battery supported bus voltage similar to the approach traditionally used in the Telecom market. To achieve the benefits of operating through temporary loss of the AC powerline, many approaches have been used, including motor generator sets and UPS systems. The use of a battery to support an intermediate bus voltage, such as shown in several examples here, provides these same benefits with levels of reliability and cost far superior to UPS implementations. This approach fits perfectly with decentralized power architectures.

Ease of Regulatory Qualifications

Distribution of an isolated SELV intermediate voltage greatly simplifies regulatory qualification, as does the usage of standard power modules that have been granted agency safety approvals at a module level. This approach eliminates one of the major schedule roadblocks that often occurs with conventional customized centralized architectures – the need for last minute safety approval of custom power converter designs.

Eliminate Problems of Multi-output Converters

Centralized architectures often require the use multi-output power converters. These converters are more difficult to successfully design than dedicated single output power converters. There are often undesired interactions between the outputs, such as cross regulation, noise coupling and dynamic response problems. Changes in the requirements in any one voltage level forces the entire power supply to be redesigned and possibly re-qualified. Decentralized architectures implemented with distributed power modules eliminate these kinds of problems entirely.

Enhanced Reliability

With older technology, replication of DC/DC converters in decentralized architectures would result in unacceptable reliability due to the summation of their individual failure rates. The newer power modules now available offer extremely high levels of reliability, much more than can be achieved with larger conventional power converters. This is due, in part, to the high levels of integration an simplicity that can now be achieved along with the availability of specialized components. The latest power modules offer failure rates over an order of magnitude lower than those of just a few years ago. Because of the importance of reliability in today's system design, we have included a dedicated chapter on reliability. The chapter on total cost of ownership also instructive in terms of understanding the impact of reliability on cost.

Enhanced Failure Isolation Capability

The more decentralized the power system, the easier the power failure diagnosis and isolation becomes. This is due to the close association between power components and circuit functions in Power per Function and Power per Board systems. The spares stocking and field replacement also become easier as the power architecture becomes more decentralized and the power modules become smaller and more granular. In the Power per Board implementation, DC/DC power modules do not need to be stocked at all as a repair part as they are automatically replaced as a part of the load board.

Fault Tolerance

Decentralized architectures lend themselves naturally to providing redundancy of function. Power per Board, for example, allows each board to be completely independent from a functional point of view. Failure of a DC/DC converter will only affect one board, and failure of the electronics on any board (such as a short circuit) will only affect one power converter. This results in dramatic increases in availability due to decreased propagation of failures between system sub-assemblies. It is possible for much of the system to remain up and running in spite of single point failures in either the power or load circuitry.

Flexibility for Upgrades and Features

One of the design goals of many systems is flexibility with regard to its size and performance options. For example, the manufacturer may want to be able to sell an “entry level” system and then later upgrade it with additional capability in terms of more features or enhanced processing power. Sometimes the end user rather than the manufacturer of the product is responsible for this upgrade activity. Ideally, the full range of systems should be accommodated by one basic power system design. These goals can be easily achieved by using decentralized power architectures with standard power modules. The DC power required to operate the new circuitry can be added as part of the upgrade. Also, the base product design is not held up waiting for definition of all possible feature mixes. Power per board is the most flexible in this regard. Very significant cost and schedule enhancements can ensue.

Live Insertion / Hot Plugging

The maintenance and repair philosophy for high availability systems usually requires “live insertion” or “hot plugging” of sub-assemblies. Including the final stage of power conversion on the load assemblies, as in “Power per Board”, greatly simplifies this operation. The intermediate bus voltage is sufficiently high so that the current levels through the connector are modest. The intermediate voltage is also only loosely regulated and can withstand a greater voltage deviation during the plugging activity than would be possible with distribution of the final circuit operating voltage through the connector. The recent availability of “hot plug controller” ICs has also simplified the design of such systems. These products are specifically designed to allow the implementation of decentralized power systems demanding live insertion of the load cards.

Low Cost Entry Systems

With centralized power, the base product contains enough power capacity to power a fully configured system. This imposes a cost penalty on the low end user, and makes the system appear to be less attractive from a cost point of view. With decentralized power, the low end user gets the lowest possible cost of power, and pays for additional capability incrementally as features are added.

Lower Total Cost of Ownership

There has recently been more awareness of the total cost of a power system over the lifetime of the end product. These costs include the areas of product development, hardware procurement of power system components and system field support. The reliability of power converters has a strong influence on the field support costs and consequently on the total cost of ownership. This topic is explored in some detail in the chapter on total cost of ownership. The high reliability of today’s standard power modules results in dramatic reductions in the cost of ownership for decentralized systems.

Reduced Time-to-Market

Decentralized power offers a very significant advantage during the system development process. It is no longer necessary to wait for the complete definition of the power requirements for the system before beginning the design of the power system and its component converters.. Each function or each board can be characterized as it is developed, and appropriate standard converter modules selected. Modular front-end AC/DC converter systems and packaging approaches can be selected with the knowledge that the final configuration can be decided upon later in the system design process. Long lead times and constant redesign for power system development become a thing of the past.

Simpler DC Distribution

The Power per Board approach essentially eliminates low voltage DC distribution, except for the easy to implement on-board low current distribution on each circuit board. Wire harness assemblies and bus bars are eliminated. The need for remote sensing and its associated reliability and diagnostic impacts is eliminated. The intermediate voltage distribution tends to be very non-critical and inexpensive due to the DC/DC converter regulation that occurs later in the system.

6 Standard vs. Custom Power Supplies

Standard modular power conversion components are finally a practical and economical reality, a reality that opens up many exciting opportunities for today's power system designer. But before discussing the benefits of these standardized solutions, it is instructive to first examine the prior art – the traditional customized power system. As will be seen, the customized approach has many problems, financial and administrative as well as technical.

Introduction

Every piece of electronic equipment incorporates one or more power supplies. Many of these power supplies are typically custom designs that are developed to meet the unique voltage and current demands of the system. They are almost always solid-state switching regulators operating at a frequency of between 50 kHz and 1 MHz. Until recently, most all supplies for Information Technology and Industrial applications used the AC powerline as a source of power, but now more are using a battery supported DC bus structure such as the traditional Telecom architecture. For purposes of this discussion of custom power supply development and implementation, an AC powerline source will be assumed. We will follow the development and manufacturing start-up process for a typical custom power supply and observe some of the potential pitfalls along the way.

Product Definition

Before design can begin, requirements and specifications must be developed. The most important specifications are the voltage, current, and regulation requirements for the load circuits within the equipment. A partitioning of the load circuitry into cards, boards, equipment cabinets, etc. must also be known or assumed. This is rarely a smooth process. The designers of the end product have little knowledge about the power requirements during the early design process. Often the product is being built with circuit technology that has not been used previously, so that other products or sub-assemblies cannot be used as a model for the estimate. Also, the development time for the custom power supply is often longer than that for the other circuitry, and power converter design must actually begin before design of the rest of the equipment is well defined. This is especially true for processor design, where automated design systems allow for very rapid design and production of computer and logic circuitry. No such systems exist for the design of custom multi-output power supplies that are used as a power source for this equipment. The most common outcome of this situation is that a power supply design is started based on the best possible estimate of the system requirements. Later, as the system design progresses, many changes are made, and the estimate no longer represents the actual system. The power system design must then be modified or, many times, completely redone. The design process is iterative, with two or more passes required. This has the effect of lengthening what is already a very long and involved (up to two years) development process.

Packaging Interfaces

The challenges are not all electrical. Each custom power supply has a unique mechanical package, some of which needs to be very precisely designed. Connector locations, for example, will need to line up with mating connectors within fractions of a millimeter. Centralized power supplies also create a concentrated source of heat. The thermal density is large enough so that forced convection or even liquid cooling is required to maintain the temperature at the power supply to a reasonable value. The cooling interface can be a very common source of problems for the system designer. For forced convection cooling, airflow direction, volume, and localized ambient temperature must be estimated during the power supply design prior to having a real product to use for testing or analysis. When the resulting power supply is integrated with the actual product, the cooling environment is sometimes different enough to create thermal problems within the power supply or the equipment.

Component Dependencies

Components are also a problem. Since the power converters are custom designs, an attempt is made to optimize the component selection for that particular application. There are hundreds of components in a switching regulator. There are dozens of different types – logic circuits, analog ICs, discrete low level resistors and capacitors, fuses, relays, connectors, power semiconductors, complex high frequency magnetics and electrolytic capacitors. These parts are produced by many different component suppliers, each utilizing unique manufacturing processes. Efficient and reliable operation of the converter may be dependent upon some second-order parasitic characteristic of several of these components – characteristics not documented and guaranteed by the component vendor. Many parts have more than one source. The converter may work with one vendor's part, but be marginal with the same part number from a different vendor. The manufacturer of the power supply usually makes these component substitutions months or years after the design is completed. The exposure for problems is severe. Compare this environment to that of the standardized logic circuitry in the equipment. All the logic usually is from one circuit family manufactured by one or two vendors using very tightly controlled and documented processes. Even though there may be hundreds of logic chips, they are actually just variations of one basic component family and one manufacturing process. Consequently, the exposure for component problems within custom power supplies is much more severe than in the remainder of the equipment.

Manufacturing Start-up

As the custom power supply moves from design into manufacturing, other disadvantages of its uniqueness become apparent. Because each such supply has different dimensions, circuit board sizes, layouts, and test requirements, many of the manufacturing operations must also be unique. Customized handling equipment, process flows, tooling for covers, and operational test programs need to be developed. These items create significant non-recurring manufacturing charges for capital, tooling and programming support. These front-end charges are particularly bothersome because many production runs are relatively low in volume (less than 10,000 units), and the non-recurring costs per power supply are a large percentage of the total manufacturing cost.

Qualification

After the first power supplies come out of the manufacturing process, the power supply usually needs to be qualified. The qualification is done to verify that the design is sound and to gain confidence that the needed reliability levels will be achieved. The qualification process is lengthy and costly due to the high reliability levels specified for today's power supplies. To achieve statistical relevance, the test program will need either a large sample size (high hardware expense) or a long test duration (long test time and support expense). Either of these options is very expensive, and the results are not always available soon enough to mesh with the remainder of the program schedule. In addition to this qualification internal to the equipment manufacturer, power supplies must be submitted to one or more safety agency laboratories to verify compliance to various international standards for safety and Electromagnetic Compatibility (EMC). This process can also be lengthy. Any problems discovered during these tests and qualifications must be addressed. The problem could be related to design, component quality, manufacturing processes or just be a random failure. In any event, the solution to the problem, especially design and component problems, can be very complex and time consuming. Countless product introductions have been delayed for months by a problem discovered during qualification of a custom power supply.

System Integration

At about the same time that the qualification testing is occurring, the final integration of the power supplies with the system is taking place. During this process, more subtle problems with the custom power supplies or incompatibilities between the power supply and the system begin to surface. A common class of problems is EMC. Switching converters generate a broad spectrum of electrical noise due to the fast high current transitions associated with the switching activity. The main switch transistors, the magnetics and the output diodes are the most common sources of such noise. Without very careful electrical and PCB layout design of the power supply, this noise can be coupled into the system either by means of conduction or radiation. This is one of the most difficult aspects of converter design, and even the most skilled and experienced designer sometimes encounters problems during system integration. Even if a power supply meets all the requirements of the international regulations, near-field disturbances can still affect the operation of the product. Again, solving these problems is not easy. It requires re-design or additional shielding, either of which requires re-qualification and time delays.

Supply Pipeline

In order to minimize manufacturing costs, much of the actual manufacturing of custom power supplies is done at assembly vendors, either domestic or offshore. The vendor is often also given responsibility for procuring the components. This presents additional problems and exposures in terms of assuring component quality, understanding the assembly processes, and providing general quality control for the power supply. These problems are potentially severe due to the length of the 'supply pipeline' between the power supply manufacturer and the final product. Power supplies manufactured abroad and shipped ocean freight can take 2 or 3 months to reach the product. By the time a problem is discovered and corrective action taken, there can be 3 or 4 months worth of defective production in transit and in inventory, adding to the expense and delay in fixing the problem.

Flexibility

The custom power supply approach is also limited in terms of its ability to adapt to changes in the product requirements. This is important, as most products incur design changes to allow for additional features to be added after product introduction. If the power system has to be re-designed every time there is a system level change, the resulting costs and schedule impacts can be enormous. This lack of flexibility is one of the biggest limitations of customized power.

Skills Allocation

An engineering team with high skill levels and many years of experience in several engineering disciplines is required to design and support reliable switching power supplies. Some of the skills required include analog and digital electrical design, magnetics design, detailed knowledge of both small and large signal characteristics of switching semiconductor devices, second order effects of magnetics and capacitors, thermal design, mechanical packaging, manufacturing engineering and EMC design. Developing and maintaining a team such as this is not easy or inexpensive. This expense, in conjunction with the long development cycle described above, results in very high non-recurring engineering expenses for custom power supply designs. This is one of the main reasons why alternatives to the custom approach are receiving such acute attention.

This design environment has another disadvantage. Due to the already high expense associated with development, and the shortage of people with these design skills, very little of the resource is assigned to the development of new technology. Consequently, almost all of these custom power supplies are designed with standard topologies and off-the-shelf components. Very little genuine innovation is seen in the areas of developing new circuit topologies, new power semiconductor devices, new components, new packaging techniques including increased integration and new assembly methods. The result is custom power supplies that are often very similar to each other in technology, with relatively little differentiation between manufacturers. Thus, even when everything goes right, the custom power supply is often a non-distinguished (and expensive) product. Even ignoring the very significant real 'costs' of time-to-market and schedule delays for the product, total cost for design, qualification and tooling for a single computer or telecom grade, custom power supply is typically in the 150,000 to 500,000 USD range.

The Standard Advantage

The standard approach resolves almost all of the above exposures. The standard products from the better suppliers are very well designed and have had the benefit of large production runs on closely monitored internal production lines. Due to the large production volumes, economy of scale permits development of unique components, specialized packaging and automated assembly equipment. The assembly area at a board mounted standard converter manufacturer looks more like an integrated circuit (IC) line than a traditional power supply assembly line. This analogy is far-reaching. These board mounted converters or power modules, are components, just like a standard integrated circuit. They are designed and assembled using automated design, analysis, and assembly tools, eliminating the problems of hand assembly prevalent with custom power supplies. This brings to the standard power module the same benefits of repeatability, quality control, and sophisticated packaging that is today taken for granted in the IC industry. Board mounted power modules are now available in SMD packages, making them even more component-like.

Standard DC/DC converters are readily available from distributors and manufacturer's reps with very short lead times. Consequently, parts for prototype systems are immediately available with characteristics identical to the final production units. The development cost is zero. The time-to-market advantages are overwhelming. Even after production has begun, there is no annual purchase commitment as is typical with custom units. DC/DC converters can be ordered cost effectively regardless of product volumes.

Another major advantage of standardized DC/DC converters is that standardization eliminates the need to retain expensive and hard to find specialized power supply design skills. Analog and power design personnel that are already in place can be assigned responsibilities on other parts of the system, where they can contribute directly to the product-level design. Or they can engage in R&D activity. Either of these alternatives is a much more productive use of these valuable skills. As the modern DC/DC converter becomes more and more sophisticated, the days of 'roll your own' are becoming numbered. Putting standard parts on a PC board is no longer sufficient to build a state-of-the-art power supply. The investment needed in component research, specialized packaging and thermal technologies, and automated assembly lines is convincing more and more companies that competitive high density, high reliability DC/DC converters require a standardized solution.

Cost Analysis of Custom and Standard DC/DC Converters

As described in the previous discussion, there are many advantages of using standard converters rather than customized power supplies. Many of these advantages ultimately translate into cost. Ericsson has explored these cost issues in considerable detail. While the details of these studies are too expansive to present here, we will describe some general conclusions that are derived from the cost analyses.

The cost analyses were comprehensive in that they considered the following cost elements:

- Hardware cost.
- Development cost.
- Technical risk cost.
- Time-to-Market cost.

The cost studies considered three types of power architecture, all of which utilized DC/DC converters. The first architecture was "power-per-shelf", with a multi-output DC/DC converter powering each shelf. The second was a "power-per-board" implementation utilizing board mounted power modules. The third alternative considered the usage of converters composed of discrete components assembled into the load PCB by the end equipment manufacturer.

For the shelf mounted multi-output DC/DC converter the studies show that the standard solution is almost always significantly less expensive. Delays to the product development schedule due to the complexities of custom converter development are especially devastating to the overall power system cost. The most striking conclusion is that, conventional wisdom aside, custom solutions are not less expensive in large volume if they cause any significant increase (one month or more) in the development schedule. This is true even at a system volume of one million per year. Figure 6.1 graphically shows this sensitivity to the development lead-time delay. For example, if the custom converter solution entails a three month delay, there is an actual cost penalty of over $285 for each shelf-level custom power supply!

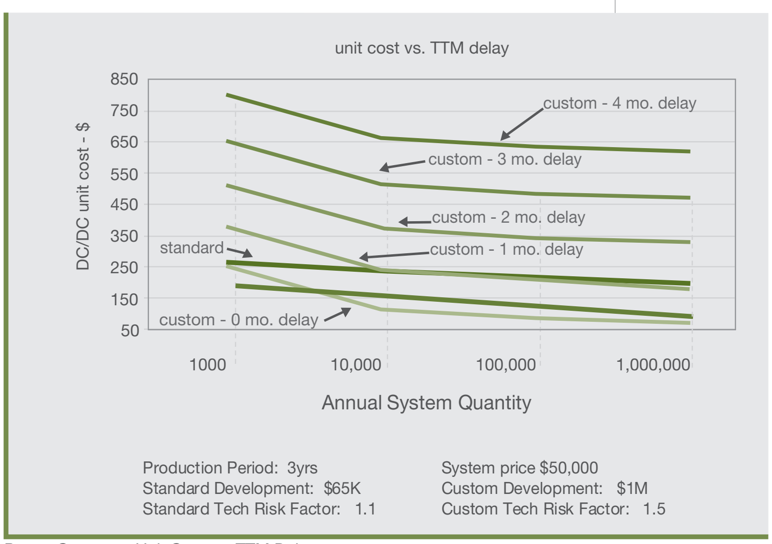

Power Converter Unit Cost vs. TTM Delay

figure 6.1

For the board mounted power module architecture, the study was even more conclusive. The analysis shows that the custom designed DC/DC power module alternative is always more expensive, even at high volume levels. At low and moderate levels of volume, it is considerably more expensive. This is due to the lack of an established market for truly custom power modules and the efficiencies of scale already accomplished by the standard power module suppliers. There is not enough volume for customized power modules to warrant the large expenditure of capitol for the sophisticated integration and packaging solutions and manufacturing resources required to produce a competitive high density power module. As a consequence, standard DD/DC power modules are always the better choice for board mounted power applications.

There is an additional alternative that has recently become feasible that is worthy of mention here. This alternative is to fabricate a DC/DC converter using discrete components and assembling it on the PCB along with the other components. This alternative has actually always been a possibility, but was not often used for the following reasons:

- Need very skilled circuit design resources.
- Power supply design affects the circuit board layout.
- Need specialized components.
- Difficult to make modifications to the converter.
- Only useful for low power.
- Low density consumes significant circuit board area.
- Diagnosis and repair of failures is easier with Standard power modules.
- Lower PCB assembly yield.

What has recently changed is the new availability of control ICs that are specifically designed for these kinds of applications. These ICs are available from the major power analog semiconductor suppliers and are supported by extensive applications assistance including:

- Complete "pre-tested" circuits.
- Component recommendations.
- PCB Layouts.

The latest of such control ICs either include or are compatible with synchronous rectification using power MOSFETs so that the efficiency of the converter can achieve levels comparable with DC/DC power modules. Also, the high level of integration of these devices allows the construction of a converter with relatively few external components. These components usually take the form of one discrete inductor and several resistors and capacitors. For purposes of simplicity, the converters constructed with this approach are normally non-isolated. While this eliminates the need for a transformer, it limits the input voltage range that can be accommodated with the simple buck topologies used. Therefore generating a low voltage in the 1 to 3 V range using the 48 V telecom input is not possible. Instead, an existing DC voltage on the board in the 5 V to 12 V range is used as the input. Figure 6.2 depicts the discrete component approach.

As mentioned above, the discrete converters are designed to operate from a relatively low DC voltage, usually in the range of 5 V to 12 V. This can be convenient if such a voltage exists and it happens to have excess current capability to power the discrete converter. Note that this will add an additional conversion stage so that the overall efficiency for conversion to the final voltage will be less than if a DC/DC power module is used to convert directly from the telecom bus voltage. If the 5 V or 12 V source is developed on the circuit board with a power module, then the power dissipation from both conversions (48 V to 5 or 12 V and 5 or 12 V to the final voltage) will add to the heat load on the printed board assembly. Using an isolated power module to convert directly from the telecom voltage to the final voltage will increase the overall efficiency and reduce the amount of power dissipation that needs to be removed from the printed board assembly.

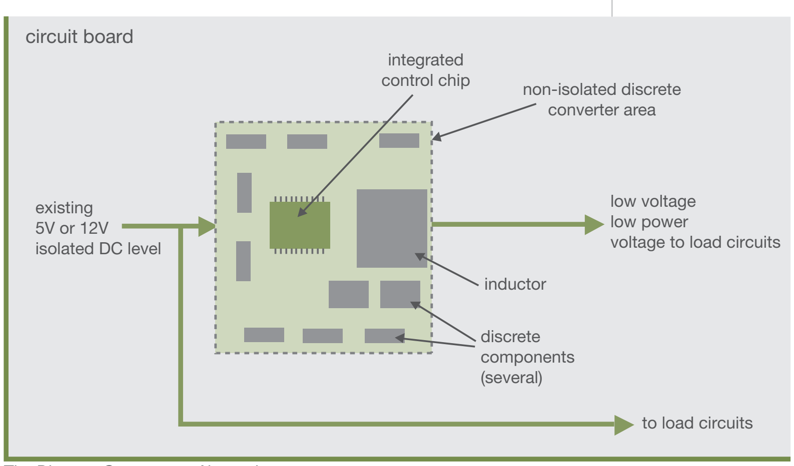

The Discrete Component Alternative

figure 6.2

In order to achieve circuit simplicity and avoid the need for an external transformer, the discrete solutions are almost always non-isolated. Depending upon the application, the lack of isolation may not present a problem. In other cases, the availability of isolation from input return to output return could assist with grounding configurations for the minimization of noise and EMC. Also, with a non-isolated solution, a failure of the switching device can apply the full input voltage to the output load circuitry, resulting in its destruction. With an isolated DC/DC power module, the transformer and fault protection circuits can prevent destruction of the load circuitry. Power modules offer the choice of isolation for increased safety and grounding flexibility.

The discrete solutions offer basic DC/DC conversion, often at very good efficiencies. However, in order to minimize the number of external components, they are fairly simple circuits that do not offer all of the functionality of a complete power module. A power module will typically provide a more complete and rich offering of control and diagnostic features, for example. The better power modules, in the low and mid power range, will also include internal filtering components to help the electronic equipment manufacturer to meet regulatory EMC requirements. To accomplish this with the discrete approach will require the addition of external filtering networks, which are not included in the basic design information.

At moderate or high power levels (above 10 W or so) it takes specialized components, thermal technology and packaging techniques to build an efficient and competitive DC/DC converter. Therefore the basic approach of the discrete converter - putting down conventional components on a normal PCB with standard manufacturing equipment – is limited to lower power levels. Most of the discrete designs are consequently targeted at the 1 W to 10 W power range.

The newer discrete designs that are supported by the control IC manufacturer are definitely more “user-friendly” than past offerings. It is still a power converter design, however, so it will need to be supported by personnel with converter design skills at the OEM user. Such skills will be needed in order to finalize the PCB layout and external component selection, add the required filtering networks to insure compliance to the product’s EMC requirements, and test and evaluate the circuit to make sure that it meets the needs of the load circuit. Most of this effort is eliminated at the OEM if a standard power module is used.

In return for possible cost savings, the user of a discrete converter incurs additional technical risk. This takes several forms. The circuit functionality is dependent upon the selection of the external components, and slight variations in the components used can upset the stability and/or reliability of the circuit. If any changes are required in the converter design it entails a change in the end product’s PCB layout. Such layout changes can be extremely expensive in terms of engineering expense and time-to-market delays. With the discrete approach the OEM is responsible for sourcing and stocking several components rather than just one power module. This will increase procurement support costs and require procurement of somewhat specialized components such as large value capacitors and discrete inductors. These kinds of risks are minimized with the standard module approach. There is only one component to consider, and it is a proven entity with a history of high reliability.

The net of the above discussion is that discrete solutions are indeed attractive from a cost perspective in certain situations. They will offer the greatest advantage for applications that:

- Are 5 W or less.
- Have an existing source of 5 V or 12 V.
- Do not need isolation.
- Need a basic converter without the functionality of a power module.
- Are supported by a skilled converter designer.

Significantly more detail on the above cost analyses can be found in Ericsson Microelectronics Design Note 003 titled “The Economics of DC/DC Power Modules - a Year 2000 Checkpoint”. Readers desiring additional insight into the cost aspects of the custom vs. standard decision are referred to this document.

Converter Topologies 7

Converter topology is a much debated topic between converter designers and researchers who are trying to optimize device-level characteristics, component stresses and hardware costs. There are hundreds of different topologies and variations that can be used to implement DC/DC converters in the power range of 5 to 300 W, so it obviously is not practical to cover them all in this limited space. Rather a high-level overview will be given.

Introduction

Fortunately, with the advent of standardized power modules, the end user does not need to analyze the topological choices in detail. The development of the standard DC/DC converter module as a component for implementation of decentralized power systems allows the power system designer to focus on the final system characteristics of the module, such as reliability, efficiency, building height, and cost. The power system designer can then pay less attention to the internal details of the converter operation. We have, therefore, limited this section on topology to include only high level descriptions and comparisons.

A general classification scheme for topologies is presented, followed by a more detailed description of the topologies most commonly used in distributed converters. These commonly used topologies are then compared from a functionality and application point of view. Synchronous rectification, while not a topological distinction, is discussed as it pertains to the operational and cost characteristics of power converter modules. Finally, there is a discussion of multiphase converters that can enhance the performance of interconnected converter cells of various topologies.

Classification of Topologies

There is no universally accepted classification scheme for converter topologies. This is partly due to the large number of topologies that exist and also to the fact that similar topologies are referred to by more than one name. Rather than attempt to categorize and classify all possible topologies, we will only focus on the characteristics that are most meaningful to the end user and on the topologies that are actually used in commercially available DC/DC converter products.

The categorization of the most common topologies is shown in figure 7.1. The must fundamental distinction between topologies is the basic operating mode - either Pulse Width Modulated (PWM) or resonant.

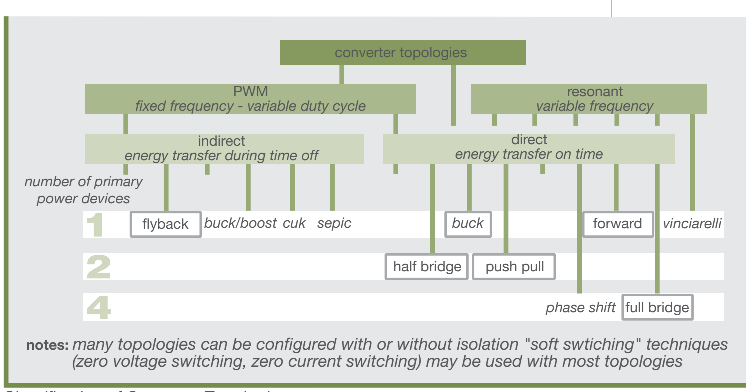

Classification of Converter Topologies

figure 7.1

In a PWM converter, the operating frequency remains fixed at a constant value, typically between 50 kHz and 1 MHz. The regulation function is accomplished by changing the duty cycle of the converter, the percentage of time that the converter's power switching devices are active. The fixed frequency has the benefit that any subsequent filtering, both inside the converter and at the load, will be done at a known frequency. The constant frequency is also convenient from an EMC point of view in that the fundamental frequency will not change and can more easily be predicted and designed for. In a PWM converter, there are dynamic switching losses when the power devices are turned on and also when they are turned off. These transition losses tend to limit the maximum operating frequency at which reasonable efficiency can occur. With good design and packaging techniques, PWM converters can be very efficient up to an operating frequency of about 2 MHz.

A resonant converter operates by changing its operating frequency, and regulates by means of the frequency dependent impedance characteristics of a series or parallel resonant LC circuit. This approach, in theory, removes the constraint of lower efficiency at higher operating frequencies. Because of this, there is continuing research being done on resonant approaches and new resonant topologies and implementations are being proposed in great numbers. For practical converters, the upper operating frequency remains below 5 MHz due to present limitations in components and packaging. The actual efficiency of commercially available resonant converters is in the same region as that of the better PWM designs utilizing conventional rectification. PWM designs utilizing synchronous rectification can exhibit even higher efficiencies. Consequently resonant designs as yet exhibit no practical efficiency advantage to the end user. In the future, as the component and packaging limitations are resolved, it is possible that high frequency resonant approaches will become increasingly popular and more commonly used due to the higher packaging density that may result. On the other hand, the use of the multiphase techniques that will be described later may raise the effective operating frequency of PWM converters and obviate this potential advantage for the resonant approach.

For PWM converters, power transfer to the converter secondary can occur either during the power device on time or during off time. In the first case, the converter is referred to as a direct converter. In the second case, it is called an indirect converter. There are several examples of each that are commonly used in available DC/DC modules, and these are indicated in figure 7.1. The topologies that we will discuss in more detail are outlined with a box. The topologies are listed in the figure according to the number of primary switching devices each normally uses. In general, a larger number of devices implies additional complexity and cost, but also a higher power handling capability.

DC/DC converters are available in both isolated and non-isolated versions. Isolation, as referred to here, is the absence of a DC circuit path from the converter input voltage source to the output return, and is accomplished by means of some form of transformer within the converter. Both isolated and non-isolated converters are useful devices, with isolated varieties being more popular due to their greater flexibility for system grounding and safety design. Many of the topologies shown here can be implemented as either isolated or non-isolated converters.

One benefit of the research into resonant topologies has been the development of 'soft switching' techniques. Using these approaches, some of the switching losses of PWM converters can be minimized by taking advantage of resonance effects at the switching time and turning the power devices on or off under controlled conditions. Zero Voltage Switching (ZVS) and Zero Current Switching (ZCS) are two commonly used terms that apply to these kinds of techniques. In return for additional complexity of design, slightly improved efficiencies are possible, especially at higher power levels. It should be noted that most converters using these techniques are PWM converters rather than resonant, and as such operate at a constant frequency.

Description of Common Topologies

The most widely used topologies are the buck, the flyback, the forward, the push-pull, the half bridge and the full bridge. We will give an overview of each of these approaches. The information provided is very general and only the fundamental elements of the topology are shown. There is no attempt to include all of the features and details required for a practical implementation of the topology. For simplicity and consistency, the power switch is shown as a MOSFET transistor. While MOSFETs are presently the most commonly used device in the range of power levels under consideration, bipolar transistors or IGBTs could be substituted without loss of generality in describing the topologies.

The simplest converter topology is the buck. This topology can be implemented either as an isolated or non-isolated converter, but the non-isolated version as shown in figure 7.2 is by far the most popular and will be described here. The buck is a forward converter with energy transfer to the output occurring during the on time of switch Q_1. Setting the output voltage as a function of the input voltage is done by changing the duty cycle (or duty ratio) of the converter. This is done with a feedback loop from the output that controls the converter duty cycle to maintain a fixed output voltage. As with the discussion of the other topologies, we will not include the feedback and control mechanisms in the simplified topology schematics shown here.

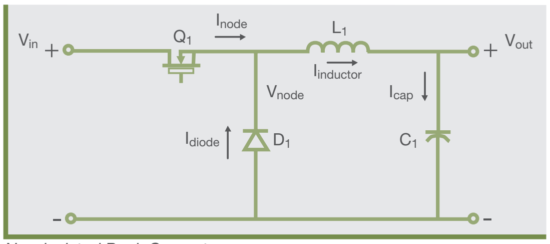

Non-isolated Buck Converter

figure 7.2

The fundamental operating voltage and current waveforms for the buck converter are shown in figure 7.3. During the Q_1 on time, energy from the input voltage is transferred to the inductor L_1, with diode D_1 reversed biased and conducting no current. The inductor current supplies the output current as well as current to charge the output capacitor C_1.

When Q_1 switches off, the output section of the converter receives no additional energy from the input side. Load current continues to flow, however, supplied by the stored energy in L_1

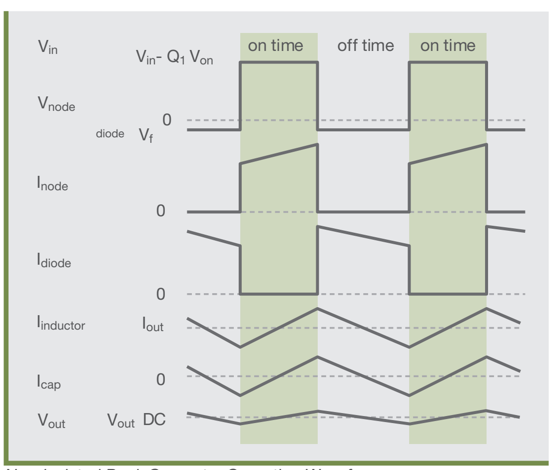

Non-isolated Buck Converter Operating Waveforms

figure 7.3

and C_1. During the off time, the load current is re-circulated through L_1 through the diode that is now forward biased. Consequently, the voltage at the Q_1, L_1, D_1 node during the off time is set at one diode forward drop below the output ground. During the on time, this node voltage is equal to the input voltage less the on resistance voltage drop through Q_1.

The inductor current has an average value equal to the converter output current, but supports a triangular AC component as the inductor current increases via Q_1 during the on time and depletes through the load during the off time. The output capacitor current has an average value of zero, but has an appreciable AC ripple component that is equal to the inductor ripple current. The converter output ripple voltage also has a triangular waveform that it largely determined by the ESR of the output capacitor.

The buck topology is limited to down conversion, with the output voltage less than the input voltage source. The buck converter is primarily used in low power applications (less than 25 W) and is often used in battery powered applications to develop voltages less than the battery voltage. Many of the self-contained switching regulator ICs utilize the buck topology.

The flyback converter is shown in figure 7.4. This is the only indirect converter that we will be discussing here. The uniqueness of this approach is that the transformer acts as an energy storage device during the converter operating cycle. During the switch on time, the output diode is reverse biased so that no current can flow into the secondary filter. During this time, the converter output current is provided by energy previously stored in the converter output capacitors and inductor. When the switch is turned off, the transformer polarity reverses, or 'flies back', and the energy stored in the transformer is released to the secondary.

The flyback is a very simple topology in terms of number of components required. It is limited, however, in power handling capability due to the use of a single switching device and the voltage stress levels imposed on the switch. Most flyback converters are designed for applications of less than 100 W. The flyback topology is used in Ericsson's PKF series of 3 to 10 W DC/DC power modules.

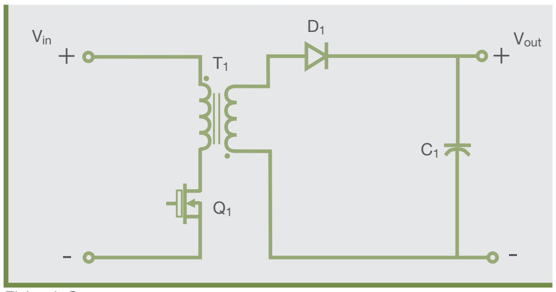

Flyback Converter

figure 7.4

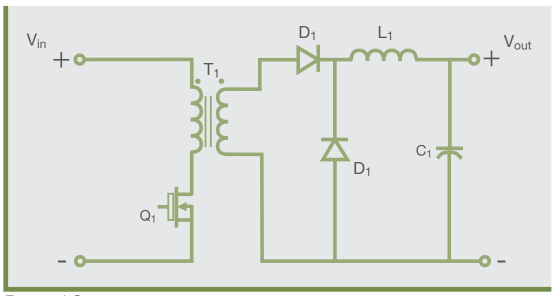

Forward Converter

figure 7.5

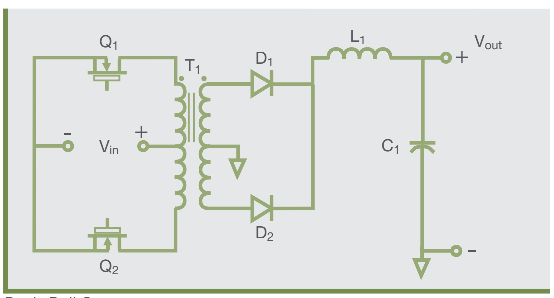

Push-Pull Converter

figure 7.6

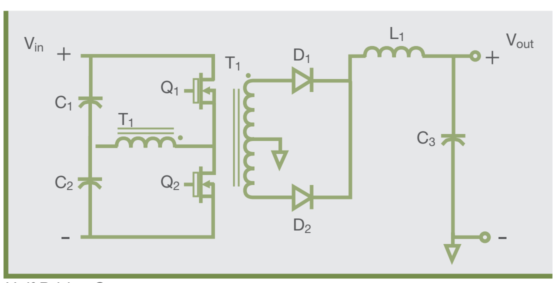

Half Bridge Converter

figure 7.7

The forward converter, the simplest form of which is shown in figure 7.5, is a very common and versatile topology. It is a direct converter, so the energy from the input to output is transferred during the on time of the switch device. During this time, secondary diode D_2 is reverse biased and current flows to the load through the secondary inductor. During the switch off time, the transformer primary voltage reverses polarity due to the change in primary current. This forces the transformer secondary to also reverse polarity. Secondary diode D_2 now becomes forward biased, and conducts current through the load driven by the stored energy in the output filter inductor. The simple topology shown in the diagram is not practical for power levels above perhaps 100 W. However there are many extensions of the forward topology that permit cost-effective operation at higher power levels. Some of the more commonly used variations and extensions of the forward topology are the resonant reset forward and the two-transistor forward. Ericsson's PKG series of power modules uses a forward topology.

The push-pull converter is shown in figure 7.6. It is a two transistor topology that utilizes a tapped primary on the converter transformer. As each power switch conducts in turn during the operating cycle, the direction of current in the primary changes resulting in a bipolar secondary current waveform. The push-pull is most useful for lower input voltages, since each of the power switches is exposed to a voltage stress of two times the DC input voltage due to the tapped transformer primary. The 24 V input versions of Ericsson PKA and PKC converters use the push-pull topology.

The half bridge topology is shown in figure 7.7. This is a two switch direct converter. A voltage level of half the input voltage is generated by the two stacked capacitors on the input. The transformer primary is alternatively switched from this voltage to either V_{in} or input return, so that the transformer primary voltage is V_{in} / 2. In return for the additional complexity of the input capacitors, this topology exposes the power switches to a maximum voltage stress of V_{in} rather than $2V_{in}$ as with the push-pull. This allows the half bridge to be useful at higher power levels. The half bridge is used in the Ericsson 48 V input PKA, PKC, PKE, PKM and PKN series converters.

The full bridge converter is a direct converter using four switching devices, and is shown in figure 7.8. In this topology, diagonally opposite switches are simultaneously conducting, imposing the full input voltage across the primary winding of the transformer. During each half cycle of the converter, the pair of switches used changes, so that the polarity of the primary reverses. As in the half bridge, the maximum switch voltage stress is equal to V_{in}. At a given power level, the primary current and switch current is half that of the half bridge due to the higher primary voltage. This makes the full bridge suitable for higher power levels. It is perhaps the most common topology for converters in the 400 to 2000 W power range. Ericsson's PKJ and PKL converters use a full bridge topology.

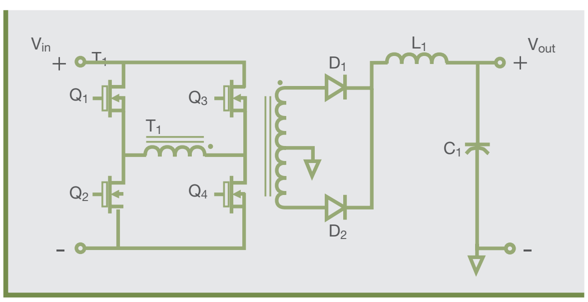

Full Bridge Converter

figure 7.8

As will be discussed in the chapter on converter selection, topology should not be a major criterion during selection of a standard DC/DC module. There is no one topology that is always best. For a given power level and other design considerations and constraints, there are one or more topologies that are suitable. The module supplier will select a topology based upon efficient design practices, availability of highly reliable components, good performance characteristics, and reasonable cost. Figure 7.9 provides a high-level comparison of the topologies we have presented.

Synchronous Rectification

One of the biggest recent improvements in efficiency of DC/DC power modules has been due to the utilization of synchronous rectification in their design. Synchronous rectification does not constitute a new or different topology. It can be used with all of the topologies discussed above. We will discuss it in this chapter, however, because it is an important topic and can be a significant criterion for the selection of a standard power module.

Synchronous rectification is simply the substitution of a MOSFET device for a conventional silicon or Schottky rectifier. This substitution can enhance efficiency because the DC losses for the MOSFET will be lower over a fairly broad range of forward current. The forward DC losses for a conventional diode will be the forward current times the forward voltage drop (0.4 to 0.5 V for a Schottky device). The power dissipated by a conducting MOSFET will be equal to the device's on resistance times the square of the forward current. Modern low voltage MOSFET devices can have very low Rds-on ratings – 10 mΩ or less. Consequently, their forward conduction losses will be lower than that of even a Schottky diode for reasonable levels of current. These characteristics are shown in figure 7.10. Note

Topology	Required number of power semiconductors	Device voltage stress (V)	Optimal power Level (W)	Comments
Flyback	2	V_{in} + N V_{out}	5 to 50	N = transformer turns ratio low complexity, low cost
Forward	3	2 V_{in}	15 to 150	many variations
Push-Pull	4	V_{in}	15 to 150	
Half bridge	4	V_{in}	15 to 200	
Full bridge	6	V_{in}	150 to 1500	

Comparison of Common Converter Topologies

figure 7.9

that the synchronous approach achieves a lower conduction loss until the crossover point with the linear losses of the diode. This crossover point will vary with the available MOSFET device technology at any given time, but is presently in the range of 80 A or so for single devices. Note that this analysis only addresses DC forward losses. The switching losses are equally real and significant for both diodes and synchronous rectifier MOSFETs. These losses are very design-dependent and will not be accounted for here.

The synchronous rectification approach involves some increased complexity and cost as will be described shortly. Consequently, it is not often used at the very low end of the output current spectrum, say 10 A or less. The efficiency advantages of synchronous rectification are greatest at low output voltages. As a result, it is mostly used for converters with output voltages of 5 V or less with output currents between 10 and 80 A.

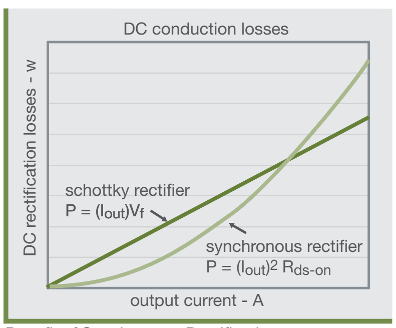

Benefit of Synchronous Rectification

figure 7.10

The additional complexity of synchronous rectification referred to above arises from the need to provide gate drive signals to the MOSFET(s) to control the conduction time. Unlike diodes, they are not self-commutating two terminal devices. The overall efficiency is determined by both the control of the on time of the MOSFET(s) relative to the operating cycle of the converter and the minimization of dynamic switching losses during turn on and turn off of the device(s). In general, achieving the highest levels of efficiency requires the use of more complex gate drive arrangements.

The two general categories of gate drive approaches are self-driven and external control. The self-driven approach is simpler and requires the use of less external components. The external control approach gives the converter designer greater flexibility in the turn on and turn off times and the generation of optimal gate-drive signals. An example of each approach is shown in figure 7.11. The synchronous non-isolated buck converter shown is an example of the usage of external control. The rectifier MOSFET, Q_2, is driven by control and drive logic in a similar way to the main switching device Q_1. This approach is easier to implement in non-isolated converters since the rectifier control and drive circuits do not need isolation from the primary-side control functions for the main switching MOSFET. The second example shows a basic isolated forward converter with a self-driven synchronous rectification design. In this design, the drive for the rectifier MOSFETs Q_2 and Q_3 is achieved by the polarity reversals on the transformer secondary. This achieves the ultimate simplicity of design, but with no control over the exact switching times or gate drive characteristics. Similar designs sometimes use an auxiliary winding on the transformer dedicated to the generation of the drive waveforms.

The DC/DC converter user should be aware of another characteristic of converters using synchronous rectification. In general, they cannot be directly paralleled without the use of isolation diodes. Thus, in architectures that require paralleling for increased current demand, the efficiency advantage of the synchronous rectification is at least partially offset by the losses inherent in the isolation diodes. One notable exception to this general rule is the Ericsson PKL series of power modules. This design uses patented architecture and technology that allows the direct paralleling of these modules without isolation diodes. Ericsson uses synchronous rectification in its PKF-B, PKJ, PKL, PKM and PKN families of power modules.

Multiphase Conversion

Multiphase converters are another recent trend that relates to topological issues. They are not a circuit topology, as they can be configured with arrays of converters of the topologies we have discussed above. They can be thought of as "Meta Topologies", or arrangements of converter cells at the sub-system level to achieve advantages not possible with a single converter.

A multiphase converter is an arrangement of two or more identical converters with their inputs and outputs connected in parallel and operating with fixed phase shifts relative to each other. The number of interconnected converters, n, is typically between 2 and 8. An example for n equal 3 is shown in figure 7.12. The operating cycle phase shift between converters is set to 360° / n, or 120° for the example shown. The three converters will share the total output power and current essentially equally, so that 33 W converters would be used for a 100 W load. Note that the effective output ripple frequency is multiplied by a factor of 3 and the ripple amplitude is reduced by the same factor. These are desirable and useful characteristics for the power system designer. The advantages of the multiphase approach include:

- Increased ripple frequency allows for smaller filtering components.
- Less ripple amplitude.
- Better dynamic response.
- Smaller converter components allow for increased integration and lower packaging profiles.
- Converter power dissipation is spread out over several components, allowing enhanced thermal performance on a PCB without heatsinks.

The potential disadvantage of this approach is the replication of converter cells, each composed of several components. Unless highly reliable components and manufacturing techniques are used, this could increase the failure rate for a given total converter output power. Fortunately, today's components and automated

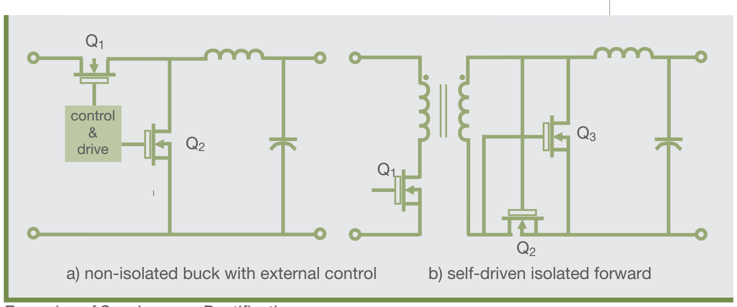

Examples of Synchronous Rectification

figure 7.11

manufacturing are up to this challenge, and the multiphase approach will be seeing increasing usage in the near future. Presently, the largest use is in the latest non-isolated voltage regulator modules for high performance micro-processors. These devices use synchronous buck topologies in a multiphase arrangement with n equal 4.

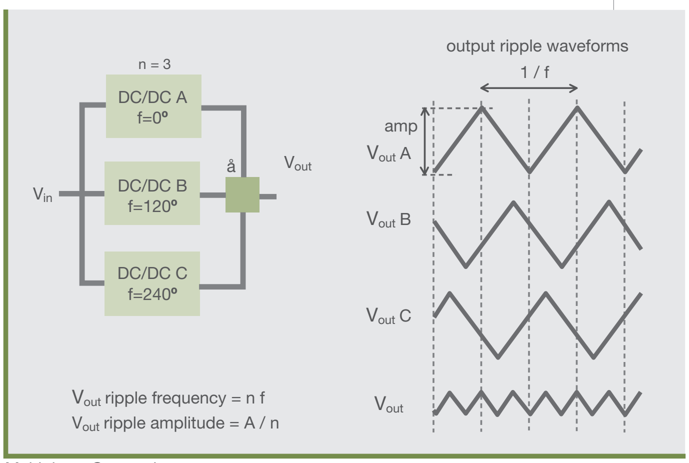

Multiphase Conversion

figure 7.12

8 Decentralized Power System – Electrical Design

The advent of standard modular power converter solutions has provided a substantial benefit to the power system designer. No longer must there be concerns about such things as converter circuit topology selection, circuit design, manufacturing sourcing, component selection, and a host of other issues related to design and production of a custom switching regulator. Instead, the designer can choose with confidence from a selection of standard power modules that are already proven designs with agency approvals and qualifications in place.

Introduction

The power system designer's task is to select and apply standard power modules in order to implement a power system that will provide the desired electrical and thermal performance as well as meet the system needs in terms of cost and reliability. A very important aspect of this overall objective is the electrical design, which consists of determining the required module specifications from the system electrical requirements and designing the electrical interfaces between the module(s) and the remainder of the system. An overview of the system electrical design process will be presented here.

Perhaps the most fundamental electrical design task is determining the DC voltage and current requirements for the system and translating them into appropriate ratings for the power modules. We will investigate methodologies for accomplishing this and examine some of the trade-offs inherent in selecting module current ratings. Another area that is key to the success of the system is load partitioning. We will look at some partitioning strategies and see how they affect the power module selection. Power DC distribution and decoupling is much easier and simpler to design with decentralized power architectures, but still is an important component of the overall product design. We will address the requirements for decoupling and establish some guidelines for implementation. Several aspects of fault protection will be addressed and recommendations given for simple filtering and fusing which can add considerably to the ruggedness and robustness of the system. Paralleling of DC/DC converters will be discussed, along with the merits and disadvantages of using this approach. Finally, DC/DC converter controls and diagnostics will be considered.

Converter Power Sizing

The first task that a power system designer must undertake is to determine the power demands of the system. The mixture of circuit families used to implement the design usually defines the needed operating voltages. The designer's focus here is to minimize the number of unique voltages required. Additional voltage levels can add cost, complexity, the possibility of interaction, more difficult diagnostics, and reliability impacts. Once the minimal set of operating voltages has been decided upon, the next step is to determine the current demand at each of the DC operating voltages. This is often done in two stages – first estimation and then measurement.

Early in the system design process, even prototype hardware does not exist. Yet it is important at this stage to do a preliminary power system design so that the feasibility of the resulting power system concept can be proven. This first iteration of power design is accomplished by using estimates of the expected current demands for each voltage level. For newly designed custom discrete circuitry, the current estimate must be obtained either from a specification for the circuit function or from the judgment of the circuit designer. These types of estimates can vary widely in accuracy. Although it cannot be guaranteed, most designers tend toward conservatism in their estimates. It is not uncommon to find initial estimates that overstate the current demand by factors of 2 to 3. The best way for the power system designer to probe or challenge these estimates is to ask for the assumptions and operating conditions for which they were generated. If the answers are very vague, the validity of the estimates should be questioned.

Since the advent of both digital and analog standard ICs, the current estimation process has gotten much easier and more accurate. If the new designs are based on using standard IC technology, the IC supplier will be able to provide data on which to base the system current estimates. This data tends to be much more reliable than estimates for custom discrete circuitry, but the power system designer must still understand the assumptions implicit in the supplied data and also make some system level judgments as to how to apply the data. For digital logic, the largest source of uncertainty is what assumption to use for worst case simultaneous switching of the input and output driver circuits. If all output drivers are assumed to be active simultaneously in their maximum power configuration, some very large current estimates can result. Some designers, attempting to be 'safe' or conservative, use this type of estimate. There may be systems for which this is a valid approach, but experience has shown that for the vast majority of systems this approach will give a result that is significantly overstated. This is especially true for large systems. As the size of the system increases, and the number of IC packages grows, the law of large numbers becomes dominant and there is significant 'averaging' of current demand. For complex systems, the probability that every circuit package is simultaneously in its worst-case power condition is very small. Even though quantitative rigor may be difficult to achieve, the experienced power system designer will make allowances for this phenomenon when estimating the expected actual current levels.

For most logic families, the IC supplier provides both typical and maximum currents on the device datasheet. Typical values are usually based on assumptions of nominal operating voltage and +25 °C ambient temperature. Maximum current is usually specified for maximum operating voltage and for the temperature that gives the highest current. These estimates are typically very good for devices without significant high current drive capability. For devices with several high current outputs, the test conditions used in the specification must be examined to determine if they are similar to those expected in the system application. The variation between the specified typical and maximum current is often large, the maximum being as much as 50% higher than the typical. Which of these should the system designer use?

For digital logic, there is much historical evidence that actual system level currents end up somewhere between the two values. With today's sophisticated power converters and distribution systems, operating voltages are usually maintained very close to nominal as opposed to maximum. This is especially true for decentralized power architectures where distributed DC/DC power modules are used. The large potential board to board DC distribution drops that sometimes occur with centralized approaches are no longer a factor. Consequently, circuitry is usually operated very close to its nominal voltage. Switching and driver activity, along with the system temperature profile, can increase the current somewhat from the typical estimate. For most systems, the following result will be valid:

$$\Sigma I_{Typical} < I_{Actual} < \Sigma I_{Maximum}$$

where: $I_{Typical}$ = Device typical current specification

$I_{Maximum}$ = Device maximum current specification

I_{Actual} = Actual system current

A more sophisticated estimation technique is to develop a computer model that contains a database of characterization data for each device. The device level current for a given operating condition may be modeled as a normally distributed variable with the mean at the typical value. By utilization of Monte Carlo statistical analysis, the program can then generate estimates of system current for any desired level of confidence. These types of techniques are very useful and convenient for designers who implement many systems with a given family of circuitry.

After prototypes of the actual hardware are available, it is advisable to perform measurements of the power supply currents to enhance the accuracy of the estimates. These measurements should be made for every expected system operating condition, including fault conditions that can reasonably be expected to occur. Start-up and shut-down conditions should receive special attention. If the circuitry is powered with laboratory power supplies during this investigation, it is important to simulate the same type of start-up voltage profile as will be seen with the power modules in the final product. For example, some power modules come up to final voltage as quickly as possible and others 'ramp up' to the final voltage over a period of several milliseconds. These characteristics can significantly affect the start-up behavior in terms of maximum power module current and possible system level problems such as latch-up and interaction between different voltage levels. Time spent on this type of activity early in the design process will pay large rewards in terms of understanding the operation of the power system, generating operating profiles for the power converters, and preventing potentially significant problems later in the product development cycle.

After completion of the above estimates and measurements, current levels for each operating voltage must be specified and power modules selected to power the system. All known system conditions, such as start-up currents, must be taken into consideration when creating this specification. If the distribution and decoupling design (to be discussed later in this chapter) is done correctly, the power module should be isolated from most of the short duration current transient demands of individual circuits, and see only an 'averaged' current with few fast transition time increases or decreases. The resulting specification should represent the best possible estimate of the actual system current demand including the effects of load dynamics and system start-up.

When it is time to select a power module, the designer should allow for some margin between the system specification and the output current rating of the power module. With custom DC/DC converters, this margin can be an arbitrary value selected by the designer. With standard power modules, the supplier's catalog will determine the available choices. The standard power module market is now mature enough that, in the most popular power range of 5 to 200 W, there are products available in nicely spaced power increments that enable the system designer to achieve the desired operating margin. The amount of margin to use is one of the more important trade-offs that the power system designer makes, and one that makes full use of their accumulated experience and skill. There is no one right answer. Each specific system requirement and development situation will influence the amount of margin selected. Some of the trade-offs are shown in a generalized fashion in figure 8.1. As can be seen, there are several variables to be considered.

Most designers end up with margins between 15 and 40%. It should be noted that the penalty for inadequate margin is more severe with customized DC/DC converter designs. With custom converters, any change in the design to increase current levels beyond the selected margin can impose very lengthy schedule delays into the system development schedule. With standardized DC/DC power modules, the impact is much less severe, as no new DC/DC converter design must be done. It is often just a matter of procuring a different part number standard power module. With highly decentralized architectures, such as 'Power per Board', the modest load circuitry complement supplied by each power module makes determination of margin a much easier task.

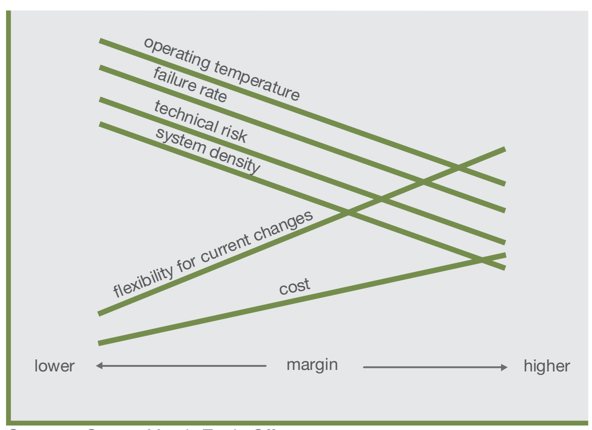

Converter Current Margin Trade-Offs

figure 8.1

Load Partitioning

The power system designer's responsibility for load partitioning essentially relates to the selection of the power system architecture. A more centralized architecture results in fewer DC/DC converters with larger functional groupings and circuit counts serviced by each converter. Using a more decentralized architecture results in a larger number of distributed DC/DC converters with each converter servicing a more restricted set of functions or circuits. The commonly used architectures were discussed in the chapter on power system architectures, and the advantages of using a decentralized approach were identified in the decentralized power chapter. These same arguments will apply to the task of load partitioning.

It is understood that the power system designer does not have complete freedom in selecting the architecture and the resulting load partitioning and DC/DC converter complement for the system. The power designer does, however, usually have some influence on decisions affecting system architecture and packaging. To the extent that this influence exists in any given situation, the following summary of the benefits of decentralized approaches is provided so that the designer may maximize the effectiveness of the load partitioning:

- Automated assembly process.
- Distributed heat load in system.
- Ease of battery backup.
- Ease of regulatory qualifications.
- Enhanced diagnostic capability.
- Enhanced reliability.
- Fault tolerant designs.
- Flexibility for upgrades and features.
- More effective utilization of backpanels and connectors in rack systems.
- Capability of hot-plugging.
- Reduced time-to-market.
- Reduced usage of multi-output converters.
- Simpler, less expensive DC distribution.
- Utilization of standard power modules.

More detailed information in each of these areas is available in the decentralized power chapter.

Intermediate Voltage Selection

Assuming that some form of decentralized architecture will be used in the power system, a nominal value for the intermediate bus voltage must be selected. While in theory any value could be used, there are practical constraints that have resulted in a few voltage values that are widely used. We will discuss some of the trade-offs associated with this decision.

In terms of DC distribution, power loss, and conductor size, the higher the intermediate voltage the better. This is because of the constant power nature of the DC/DC converter load on the intermediate bus. The higher the bus voltage, the lower the current, and the bus can be made with less copper and thus be lower in cost. One important constraint tends to limit the adoption of this strategy – safety. Every country has some kind of safety standard or requirement that limits the maximum value of voltage exposure for an equipment operator or service personnel. This limit is generally referred to as SELV (Safety Extra Low Voltage). The most commonly accepted value for SELV is <60 V. Consequently, if the bus voltage is designed to be less than 60 V, the product gains advantages from a safety shielding and regulatory compliance point of view. There is some high-end data processing equipment that use a decentralized approach with an intermediate voltage obtained by rectifying and filtering the AC line voltage. The intermediate bus voltage in these applications is 300 V and higher. The vast majority of decentralized systems, however, are configured with the bus voltage below the SELV limit.

Assuming that we will be using an intermediate bus voltage below 60 V, what is the best choice? The most common choice is a 48 V system. This voltage has been used in the telecom industry since its inception, and was the battery voltage at the central office. Even today it still has three big advantages:

- It is the most 'standard' voltage. There is a large variety of available components, including DC/DC converters, AC/DC converters, filtering, and distribution hardware and control components that are designed to operate with the 48 V bus. The power system developer thus has much more available to choose from and less need for costly custom components.
- Many decentralized systems utilize a battery-backup feature. For such systems, the bus voltage should be centered at the battery voltage. The 48 V standard is easily implemented with lead-acid batteries, and extensive standard battery charging and monitoring equipment is available.
- Any bus voltage will extend above and below the nominal value under various system conditions such as battery charging and load switching. In the 48 V telecom standard ETS 300 132, the maximum voltage is specified at 60 V, which makes the 48 V bus the highest voltage alternative if SELV is a design objective. The nominal voltage is 54 V; equal to the battery floating voltage.

For the above reasons, most decentralized systems are being designed using a 48 V bus. There are some other values that are used in more specialized applications, the most common of which are described below.

Parts of Europe use a 60 V telecom bus, with a nominal voltage of 68 V and a 75 V maximum steady state value. Due to the SELV exposure, it is not recommended that this system be used other than for those telecom applications where it is mandatory.

Many industrial control and mobile phone systems use a 24 V bus with a nominal voltage of 27 V. This allows for a battery back-up with fewer cells and consequently lower battery cost. In return, the bus currents are higher and there are more costs and power losses associated with the DC distribution than with a 48 V system. Due to the growth of the mobile phone industry, the availability of standardized power modules for 24 V applications is increasing. For equipment with modest power requirements, this may be a viable alternative.

There is occasional usage of 36 V. This used to be a standard in some Scandinavian telecom systems. It is used rarely today. Figure 8.2 graphically summarizes the main trade-offs and considerations that have lead to the adoption of the nominal 24 and 48 V intermediate busses as the most popular choices.

The different bus voltages are designated according to their typical battery discharge voltage, i.e. 24 V, 48 V and 60 V. Both the 48 V and 60 V telecom systems are configured with a negative intermediate bus voltage relative to system earth. Many Datacom systems utilize a positive bus voltage. The majority of the DC/DC converter modules presently available contain galvanic isolation

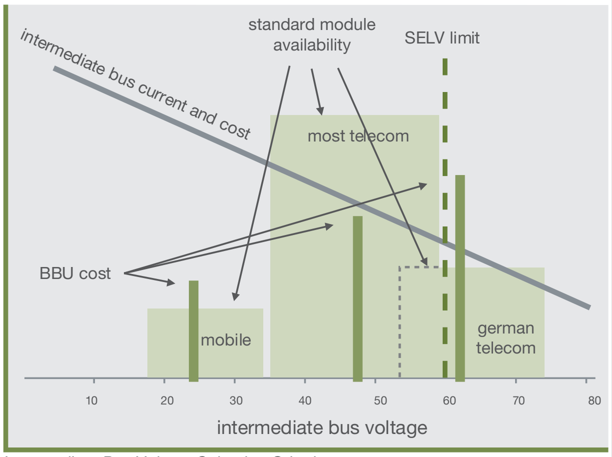

Intermediate Bus Voltage Selection Criteria

figure 8.2

from input to output allowing them to be used with either polarity input voltage. The isolation provides excellent flexibility in this regard and also some other advantages.

In addition to the ability to operate from either polarity input voltage, isolated converters allow the user a high degree of flexibility in system earthing. Input and output can be tied together or left isolated. The isolation also is an advantage for some regulatory approvals, such as CSA certification. Some costs can be saved in the converter if a non-isolated design is used, and there are situations in which this is an appropriate design. With careful attention to system architecture and earthing practices, non-isolated converters can perform equally as well as isolated versions. There is speculation that European telecom systems will be allowing for use of non-isolated converters in the future. This should provide incentive for DC/DC power module suppliers to provide offerings of both varieties. Ericsson offers both isolated and non-isolated power modules, and will provide application assistance with determining which choice is appropriate for your system.

DC Distribution

With decentralized power architectures, the DC power distribution task is considerably simplified relative to traditional centralized architectures. The customized bus bars to conduct DC operating voltages from centralized converters to load boards are not needed. Remote sensing and analysis of its effect on converter stability are not needed. The size of the DC distribution network for each DC/DC converter is considerably smaller, which eliminates most of the load-to-load interaction, and greatly simplifies the analysis of the earthing system. The ideal configuration is with a 'Power per Board' type of decentralized architecture where the complete DC distribution system is contained on the same circuit board with the power module. In such systems, the board size is typically small, and the number of load circuits is very manageable. For our discussion here, we will assume this type of implementation.

The DC distribution system on the circuit board must accomplish three design goals:

- DC voltage requirements of circuits must be satisfied.
- Dynamic current requirements of circuits must be satisfied.
- Provide suitable environment for control of high frequency noise.

Fortunately, all of these goals are easily attainable with very simple and cost effective structures when using a 'Power per Board' type of system. The structure most often used is the multilayer printed circuit board (PCB), with at least one layer utilized as an earth plane. Unless otherwise noted, we will assume this type of board is used.

Of the three requirements identified above, the DC voltage is the easiest to design and analyze. The DC current levels with a 'Power per Board' architecture are limited to reasonable values on each board. With typical circuit voltage tolerances of ± 5%, currents of this magnitude are easily distributed using circuit traces of modest width. The variables involved are the width and length of the circuit trace, the thickness of the copper, the current being distributed, and the allowable DC voltage drop. To simplify the design process, the relationship between these variables is shown in figure 8.3.

The curves in figure 8.3 are normalized for 35 µm (1 oz/ft2) copper, 1 A of current, and a voltage drop of 1% of the DC output voltage. It is assumed that a ground plane return is used that offers no voltage drop. If a return path similar to the voltage trace is used instead of a ground plane, the voltage drop will be twice that as determined from the figure. Curves are plotted for output voltages of 2, 3, 5, 12, and 15 V. Here are some examples of how to use the curves for estimation of DC distribution performance and requirements:

- Assuming the normalized conditions and a 3 V output, what is the maximum length of distribution for a trace width of 2.5 mm? Using the 3 V curve, the maximum length is 15 cm.
- Assuming the normalized conditions and a 5 V output, how wide must the trace be if the length is 50 cm? Using the 5 V curve, the width needs to be 4.9 mm.

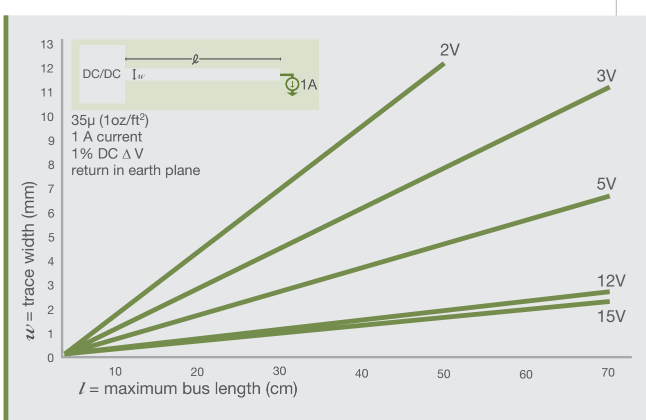

Voltage Distribution Guidelines

figure 8.3

The normalization allows easy modification of the results to account for actual system requirements. If the result is a length, then it is modified as follows:

$$l_{modified} = l \times \frac{[\%\ DV]\ [Cu\ Thickness]}{[A]}$$

If the result is a width, then it is modified as follows:

$$w_{modified} = w \times \frac{[A]}{[\%\ DV]\ [Cu\ Thickness]}$$

A couple of examples will help clarify how this is done:

A 5 V system needs to deliver 3 A over a distance of 30 cm. A 2% voltage drop is allowable. What is the minimum trace width for 35 µm (1 oz/ft2) copper?

First, use figure 8.3 to obtain the width for 1 A and 1% drop. We obtain a trace width of 2.9 mm. Next, we modify this result to account for the current and voltage drop requirements relative to the normalized values:

$$2.9\ mm \times \frac{(3)}{(2)\ (1)} = 4.35\ mm$$

A 3 V system uses a 3 mm wide 70 µm (2 oz/ft2) copper voltage distribution trace and must supply 5 A of current at 1% maximum drop. What is the maximum length?

Solving, using normalized conditions, we obtain a length of 18.3 cm. Correcting for the actual conditions,

$$18.3\ cm \times \frac{(1)\ (2)}{(5)} = 7.3\ cm$$

Figure 8.3 assumes that the entire current is lumped at the end of the distribution system.

This is a worst case assumption, as most systems have current being pulled out of the voltage bus all along the extension of the bus, so that the average bus current is less than the maximum current. The user can account for this in two ways. One approach is to follow the assumptions presented here with the understanding that the results will be conservative, and contain some extra margin. The other possibility is to estimate the average bus current and use this value instead of the maximum current. This method has the most validity if the current distribution coming out of the bus is rather uniform in nature. For actual board layouts, a practical approach is to vary the distribution trace width as a function of current level, with the widest traces close to the converter and narrower traces at the far end of the distribution network.

The next function of the DC distribution system is to satisfy the dynamic current requirements of the load circuits. It is difficult to generalize, as each type of application will have different requirements. For digital circuits, higher speeds usually imply more stringent dynamic requirements. Circuits that switch significant levels of current, such as line drivers, tend to have greater need for good dynamic response. The designer of the board must be familiar with the dynamic current demands of the circuitry. The supplier of the ICs used to implement the design will usually have application data available to help with this determination.

As a simple example, we will examine the dynamic behavior of a system consisting of just one circuit supplied with an on-board DC/DC converter. The circuit draws a DC current of 1 A and has a switching behavior that requires an additional 100 mA of current from the DC distribution. The rise and fall times of the 100 mA requirement are 5 ns. The system is depicted in figure 8.4. It is required that the DC voltage change during the switching activity be less than 4% of the nominal 5 V value, or 200 mV. Assuming a 35 µm copper 5 V distribution trace, and an earth plane on the reverse side of the board, the nominal resistance of the DC distribution network is 42 mW. This resistance results in a nominal DC distribution drop of 42 mV at 1 A (less than 1%) and only a 4.2 mV change in DC voltage during the 100 mA dynamic transient. Things look fine from a strictly DC point of view. We must, however, also examine the dynamic behavior of the system.

The 100 mA current increase must be supplied very quickly – in less than 5 ns. This requires that both the DC/DC converter and the on-board distribution network be capable of supplying 100 mA in 5 ns. First, let's assume that the DC/DC converter is an ideal component and can provide unlimited dynamic response. With this assumption, we need to determine the capability of the distribution network. For fast transitions such as this one, the dynamic impedance of the distribution network is needed. Even though the DC resistance is reasonably low, the inductance inherent in the 5 V distribution trace will limit the ability to conduct a changing current without resulting in inductive voltage fluctuations, given by the relationship:

$$\Delta V = - L \frac{dI}{dt}$$

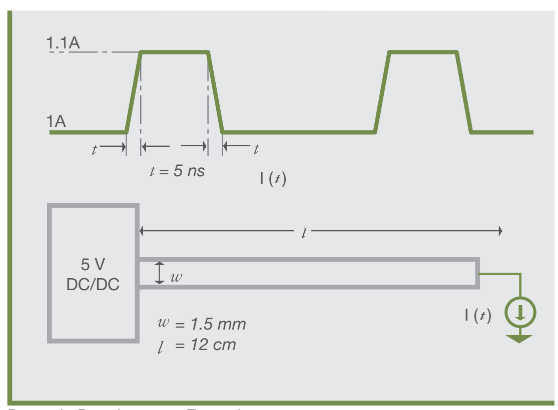

Dynamic Requirements Example

figure 8.4

The above relationship assumes a lumped model with only one inductive component. The actual situation with board distribution is more closely approximated by a distributed model, with the beneficial effects of the capacitance between the distribution trace and the earthed backplane included. The dielectric constant and thickness of the board insulating material will determine this capacitance. This model is shown in figure 8.5. If this concept is extended to a large number of circuit elements, it approaches a transmission line model, and a resulting characteristic impedance or dynamic impedance can be determined. The dynamic impedance will be a function of the distribution trace geometry and the PCB material and thickness.

Some typical dynamic impedance ranges for common PCB voltage distribution implementations are shown in figure 8.6. For our example using a 2 layer board, a value of 50 Ω would be a good estimate. This is vastly different from the 42 mW DC resistance! It also explains why the on-board distribution network cannot respond immediately to demands for transient currents. Even with a 'perfect' DC/DC converter, the response time of the system is limited by the inductive nature of the distribution system. The result will be a dip in voltage during the 100 mA current increase and a corresponding increase in voltage during the current decrease.

The above effect is independent of the DC/DC converter. The board designer has some control over the value of the dynamic impedance. It can be improved (made lower) by using a wider distribution trace. This will increase the amount of distributed capacitance within the circuit board. This alone, however, will not normally result in a system capable of meeting the dynamic current demands of the circuitry.

We have so far assumed an ideal DC/DC converter with instantaneous transient response. It should be noted that all real switching converters are not perfect have a finite ability to supply transient currents. This is due to two factors:

- The operating frequency determines the time required to sense and respond to a change in load current. For example, a converter operating at a frequency of 500 kHz will require 2 µs to begin to supply additional energy for an increased transient load.
- The amount of energy delivered to the load per operating cycle has a limit determined by the converter topology and circuit implementation. Several operating cycles may be required to deliver the increased power required during a demand for dynamic current.

The transient load capabilities of a DC/DC converter are characterized by the dynamic response specifications supplied by the converter manufacturer. The format of the specification

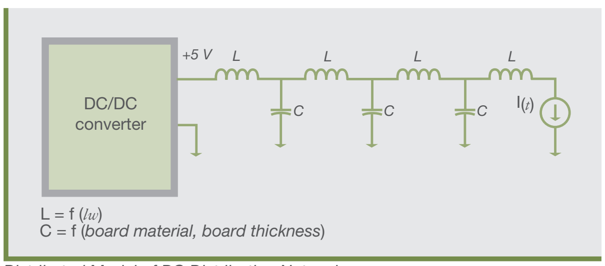

Distributed Model of DC Distribution Network

figure 8.5

Circuit board Configuration	Dynamic Impedance Ω
MultiLayer board with dedicated plane for each voltage plus earth plane[1]	1-5
2 layer board with complete earth plane[2]	20-60
2 layer board with incomplete earth plane[2]	50-100
Single layer board	>100

[1] assumes 0.8 mm glass epoxy dielectric
[2] assumes 1.6 mm glass epoxy dielectric
Assumes 2.5 mm typical voltage bus width

Typical Dynamic Impedance Values

figure 8.6

varies from supplier to supplier, but most data sheets will include this information in some form. The more customer-oriented suppliers of DC/DC converters will also work with the user to assist in determining the dynamic response that should be expected within the actual application.

We have determined that, even with ideal DC/DC converters, the DC distribution model shown in our example cannot meet the dynamic current requirements of many systems. Fortunately, there is a solution that will result in easily implemented systems with excellent performance. The solution is the inclusion of decoupling capacitors within the distribution network. The purpose of the capacitor is to act as a localized reservoir of energy to supply the dynamic current demand of the system during the time period required for the power converter and the distribution system to deliver the increased power level to the circuit location. To maximize the effectiveness of the decoupling capacitor, it should be physically located as close as possible to the load circuit. This is done to minimize the amount of inductance between the capacitor and the load. The addition of decoupling capacitance to our simple example is shown in figure 8.7.

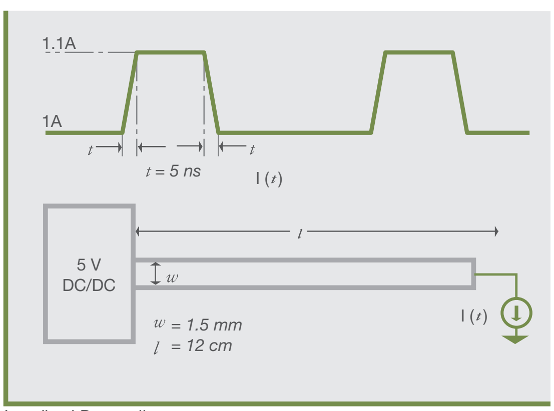

Localized Decoupling

figure 8.7

The next step is to determine the amount of capacitance required. This will be a function of the load transient recovery time of the DC/DC converter and the distribution network. A typical value for the recovery time is in the range of 40 µs. This means that the decoupling capacitance must supply the increased load current for a period of 40 µs, after which the converter and distribution trace will be supplying enough energy to meet the 4% dynamic voltage requirement. After a longer period, perhaps 100 to 200 µs, the 1.1 A current will be treated as a 'DC' condition by the power system. To determine the capacitance value, we use the expression:

$$I = C\frac{dV}{dt} \quad \text{or} \quad C = I\frac{dt}{dV}$$

In our example, $I = 100\text{ mA}$

$dt = 40\ \mu s$

$dV = 200\text{ mV}$

$$C = \frac{(0.1)(40 \times 10^{-6})}{(0.2)} = 20\ \mu F$$

The above example was a very simplified case where only one very high current circuit was assumed. In actual systems, the more typical condition is to have a larger number (10 to 100) smaller circuits on a board that is supplied by a DC/DC converter module. Each of these circuits will have some dynamic activity, typically much smaller in magnitude than the 100 mA assumed in the above example. It is important to have localized sources of decoupling for these circuits so that their immediate needs for transient current can be satisfied. An added benefit of this is that the interaction between circuits is reduced and the effects of switching noise minimized. For typical digital circuits, the highest dynamic current requirements occur during the switching time of the output circuits, and are a result of the charging and discharging of capacitance in the signal distribution network. Typical values of dynamic current would be

10 mA for 10 ns. These requirements can be handled easily with smaller ceramic capacitors. Using these values, and allowing for a voltage deviation of 200 mV, results in:

$$C_{min} = \frac{(0.01)\,(10 \times 10^{-9})}{(0.2)} = 0.0005\ \mu F$$

In actual practice, this type of decoupling is accomplished with ceramic capacitors in the range of 0.01 to 0.1 µF. Several are scattered throughout the board, with non-critical IC packages in the same physical area sharing a common capacitor. Circuit devices with higher dynamic requirements, such as line drivers and microprocessors, should have a dedicated decoupling capacitor located as close to the IC as possible. IC suppliers typically suggest a decoupling plan similar to the following:

Circuit Type	Decoupling
Standard Logic	1 0.01 to 0.1 µF cap per 5 ICs
High Speed Logic	1 0.01 to 0.1 µF cap per 3 ICs
Driver or Receiver	1 0.01 to 0.1 µF cap per IC

These high frequency capacitors should be located as close to the IC packages as possible and be mounted with short lead lengths to reduce series inductance. There will be some internal equivalent series resistance (ESR) and inductance (ESL) within each capacitor. The strategy of using several such capacitors scattered around the board will minimize the effects of the series resistance and inductance since all the capacitor ESRs and ESLs will be in parallel with each other and reduce in value, while the parallel capacitive values will add. The result will be a much higher equivalent ratio of C to R and L.
For example, consider a case where 10 capacitors are used, each with a capacitance C, ESR of R, and ESL of L:

	C	R	L	C/R	C/L
Individual Cap	C	R	L	C/R	C/L
10 Caps	10C	R/10	L/10	100C/R	100 C/L

The ratio of C to R or L is much better than it would be with a single capacitor of larger value. This benefit is in addition to the benefit of the localized capacitors helping to isolate the circuits from each other in terms of interaction.

In addition to the high frequency ceramic capacitors, each board should contain at least one higher value bulk decoupling capacitor for each DC voltage to handle longer-term demands for transient current. The location of this capacitor on the board is less critical, as long as the ceramic capacitors are well distributed physically to handle the short duration current demands. The bulk decoupling capacitor should typically be in the range of 5 to 100 µF/A of load current. A good quality (low ESR) tantalum or computer grade aluminum electrolytic should be used. Some DC/DC converters may not limit the rate of voltage rise at turn-on and exhibit an overcurrent condition if large values of capacitance are applied to the output. If the 100 µF/A value is not exceeded, this should not generally be a problem. We thus have a '3 tier' decoupling system:

- Bulk electrolytic capacitor for low frequency.
- Distributed ceramic capacitors for mid frequency.
- Circuit board dielectric capacitance for very high frequency.

Figure 8.8 shows a board layout for a 'Power per Board' decentralized approach. Note how the design guidelines presented here have been incorporated. The resulting layout should provide for reliable operation, minimal interaction between circuits, and clean switching waveforms.

One of the more demanding applications for decoupling and dynamic regulation is the newer personal computer processors, such as the Pentium®. These chips currently operate at supply voltages of 3 V and below and at currents up to 50 A. They are specified to include a 4 A transient that occurs within a 20 ns time. This results in a ramp rate of 200 A/µs that requires careful attention to decoupling and board layout, since only 75 mV maximum deviation is allowed. Intel® recommends a combination of six 100 µF low ESR tantalum capacitors and 25 1 µF ceramic capacitors to decouple the processor chip.

The Energy Star power conservation requirements also place dynamic demands upon the power system. To meet the Energy Star requirements, equipment manufacturers either selectively put parts of the circuitry into a 'sleep mode' or vary the operating frequency to achieve lower power operation when lower performance is acceptable. Coming out of 'sleep mode' can result in a 10% to 90% load change on the converter powering the processor at a rate of more than 10 A/µs. Processor operating frequency changes can result in converter load excursions of 20% to 50% with similar ramp rates.

When designing to meet these kinds of requirements, it is important to use a distributed decoupling approach with several capacitors rather than attempting to use one very large capacitor. Large values of capacitance (over 100 µF/A) can actually be counterproductive. Some of the negative effects can be:

- Longer settling times.
- Lower converter bandwidth.
- Reduced phase margin.
- Possible instability.

With careful attention to details, very good performance can be achieved. This requires the combination of a converter with a high operating frequency for good bandwidth, intelligent board layout with wide traces, and a combination of the correct kinds of decoupling capacitors. As circuit operating voltages migrate downward to achieve higher performance operation, the trend for more stringent regulation and dynamic performance from the power converter will continue. Decentralized power architectures will be a definite requirement to handle these requirements.

One recent technique to provide additional dynamic response performance is the usage of "droop regulation" on low voltage high current converters, such as the voltage regulator modules for high performance processor chips. The droop regulation technique takes advantage of the fact that increasing current transients create negative-going output voltage responses and vice versa. By positioning the output voltage of the converter at the highest extreme of the regulation band at low output current, maximum headroom is provided for the negative voltage transient. Conversely, a lower output voltage setting at high current gives the maximum amount of room for a positive voltage transient. This is illustrated in

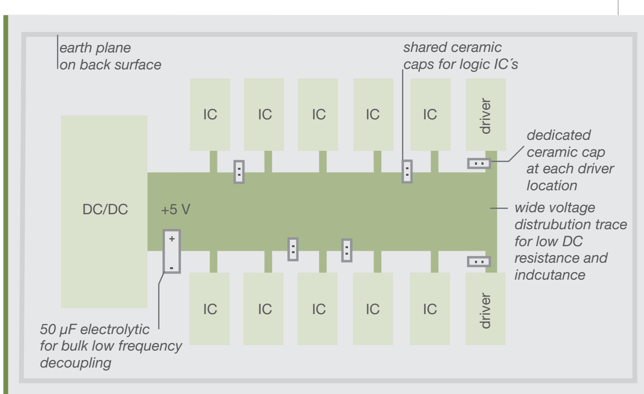

Typical Board Layout

figure 8.8

figure 8.9. Droop regulation can be thought of as trading off DC regulation for dynamic performance.

The third requirement of the DC distribution system is to provide for control of high frequency noise. Some of this noise is generated by the switching activity within the DC/DC converter. Some of it may be the result of coupling (conducted or radiated) from other boards or other systems. Most of it will be generated by the switching activity of the load circuits on the board. In general, the same design practices that we discussed above for the purpose of supplying dynamic current also provide a favorable environment for noise control. The most important aspect is to use as much of a complete earth plane as possible. This will provide significant benefits. Also, the high frequency decoupling with ceramic capacitors and the wide voltage distribution traces will help suppress high frequency switching noise conducted from the power supply outputs.

In extreme cases, such as highly sensitive low-level analog circuits, it may be necessary to do some additional filtering of the DC/DC converter output. This usually would take the form of an R-C filter or an L-C lowpass filter, as shown in figure 8.10. With modern DC/DC power modules operating at frequencies of 200 kHz and above, the size and cost of the filter elements are very modest, and the resulting filters are easily implemented with standard manufacturing techniques. Additional information on output filtering can be found in Ericsson applications and design notes. It should be noted that it is rare for filters such as this to be required. With proper DC/DC converter selection and board layout, almost all circuitry is compatible with on-board DC/DC power modules without additional filtering.

Fault Protection

In spite of the enhanced reliability of today's power systems relative to those of the past, it is still possible for failures to occur. In the event of a failure, it is very desirable that the following two objectives are achieved:

- The failure should be contained to the smallest possible number of replaceable sub-assemblies – ideally one.
- The failure should not result in any personal safety hazard to the operator of the equipment.

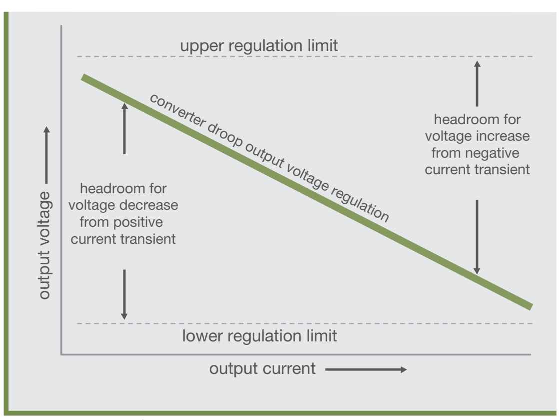

Droop Regulation for Increased Dynamic Margin

figure 8.9

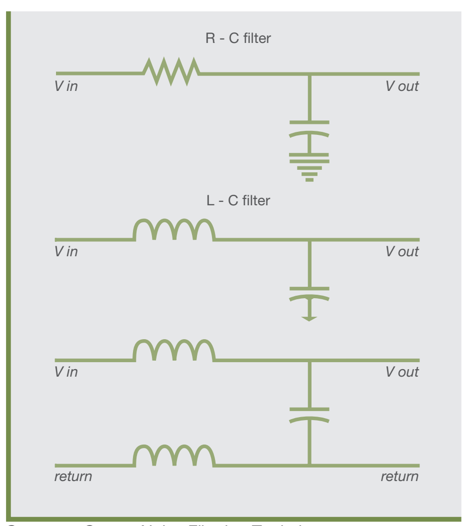

Converter Output Noise Filtering Techniques

figure 8.10

Meeting these design goals falls under the category of fault protection. We will examine the most common power system failure mechanisms and discuss how the system can be designed so that such failures do not adversely affect the remainder of the system elements or the safety of the operator. Unless otherwise noted, we will be assuming that the DC/DC converter is a DC/DC power module utilized in a decentralized power system architecture.

In the event of a failure internal to the DC/DC power module, there are four possible resulting conditions, each of which will affect the system differently:

- The power module will degrade in performance.
- The power module will stop operating altogether.
- The power module output voltage will rise above the upper DC regulation limit.
- The power module will impose a short circuit condition across its input terminals.

The first two conditions are not 'dangerous' failures, in that they will not damage other circuitry or create hazardous conditions. They are handled by system diagnostic techniques, some of which will be discussed in section "Dignostics". The third condition, commonly referred to as an overvoltage fault, can cause damage to load circuitry connected to the output of the power module. This fault is handled by some type of overvoltage protection. The fourth condition can cause excessive current to flow into the power module. This current can cause overheating of the intermediate voltage distribution system if precautions are not taken. Proper sizing and fusing of the distribution system should prevent this type of fault from occurring.

A very similar condition can be caused by failures in the load circuitry that result in excessive current demand from the DC/DC power module. These faults are referred to as overcurrent conditions, and are handled by proper DC distribution design along with some form of overcurrent protection in the power module.

Overvoltage faults are most commonly caused by a failure in the regulation feedback loop that results in the power module regulating to a higher voltage or running 'open loop' without any regulating function. In some older designs, the voltage output under these conditions could be quite high. With centralized power systems, this high voltage could wipe out all of the load circuitry in an entire piece of equipment. Consequently, some very elaborate overvoltage protection schemes have been used, often encompassing the same amount of circuitry as the primary regulator control itself. High power zener diodes have been used as 'crowbars' across the converter output to activate the converter's overcurrent function in the event of excessive voltage on the output. These diodes, besides having accuracy problems due to their temperature coefficients, added significant cost and failure rate to the system.

With decentralized power architectures, overvoltage protection is becoming much less of a problem. With 'Power per Board', for example, the converter is typically replaced as a unit with its load circuitry in the event of a failure in either. No longer is there an exposure for an overvoltage fault in one power supply to destroy all of the circuitry in an entire piece of equipment. Also, DC/DC converters have become much more reliable and integrated, so that overvoltage type failures are becoming increasingly rare. With intelligent circuit design, the maximum converter output voltage in the event of an 'overvoltage' failure can be very modest – on the order of 7 V for a 5 V output. In cases like this, the addition of a separate overvoltage control function is not normally a reasonable trade-off. For users who demand a separate overvoltage control loop, it can be implemented easily external to the converter module, and used to turn off the module by means of the converter's control line in the event of an overvoltage condition. Circuit application assistance is available from Ericsson for those customers who desire to add this function to those Ericsson power modules that do not have internal overvoltage protection.

Some form of current limiting is used in virtually all converters. This limiting will prevent damage to the converter in the event of shorts occurring in the load circuitry. Depending upon the implementation, the limiting can be done by sensing the converter's secondary current or by sensing its primary power. The power system designer should be aware of the maximum possible output current from the converter, since this value will determine the amount of current that the DC distribution must support in the event of an overcurrent condition. The DC distribution should then be sized to operate safely with this level of current. There are various ways in which the converter can react to an overcurrent condition. It can shut down and require recycling of the input power to restart it. It can automatically attempt to restart, and operate in a 'hiccup' mode until the short circuit condition is removed. It can continue to operate as a 'constant current' converter. The DC/DC converter datasheet should define the overcurrent implementation and operation. The DC/DC converter supplier's applications group can also be used as a resource for assistance with understanding system and converter failure mechanisms and their interactions.

The final type of failure is a short circuit within the converter that results in excessive current demand from the intermediate voltage bus. This current could be very high, especially in systems with battery back-up of the bus. In such cases only the internal impedance of the battery limits the current unless some form of external fusing is used. This high current could cause overheating of the distribution network and possibly even be a fire hazard if the system is not properly designed. It is also important, in the case of a rack type system using a 'Power per Function' or 'Power per Board' type of architecture, that a DC/DC converter fault affects only the function or board associated with the converter. The fault should not propagate and disable other converters or other functions.

These objectives can be achieved by using a fuse on the input of each unit, as shown in figure 8.11. These fuses should be sized to conduct the maximum input current for each DC/DC converter, but to open for a short circuit fault condition. This will limit the current in the intermediate distribution bus to a safe level, and also isolate the failed unit and its load from the rest of the system. The remaining functions will continue to operate.

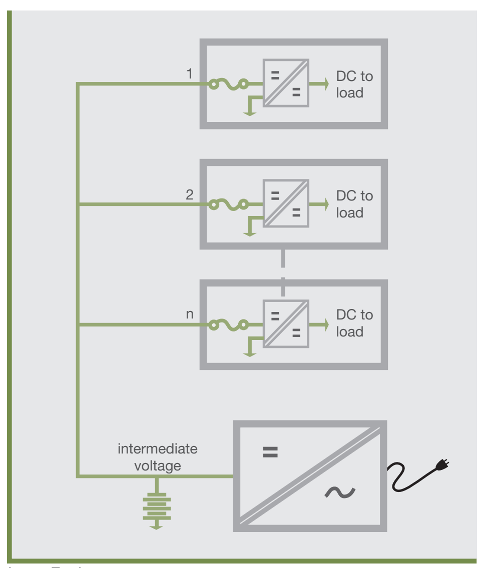

Input Fusing

figure 8.11

Input Considerations

For many applications, DC/DC power modules can be connected directly to an intermediate voltage bus within the input operating voltage range of the power modules. The power modules will function very adequately in this type of configuration. There are other applications, such as telecom, where reliability, availability, and fault tolerance are of paramount importance. For such systems, it is common to enhance the above characteristics by adding external components between the intermediate voltage bus and the converter input. We saw, for example, how the addition of fuses to the DC/DC converter inputs allowed for the isolation of a faulty function from the remainder of the system. Similarly, the system can be made even more robust by considering other possible sources of disruption of function.

Transients on the intermediate voltage bus represent one class of disturbance that need to be considered for telecom systems. There are many possible sources of transients, the most common being load switching on the intermediate voltage bus. If a load is rapidly switched, such as when a fuse opens, any inductance in the distribution network will generate a voltage pulse that can have a significant energy content. All Ericsson DC/DC power modules have built in transient protection on the input to absorb some transient energy. The datasheet will specify the amount of this protection. In order to provide maximum flexibility to the user, additional protection can be added externally, either for each power module or for functional groupings of power modules. Figure 8.12a shows one such network. The RC network provides a lowpass function, with the large value capacitor absorbing excess transient energy. In the event of transient reductions of the intermediate source voltage (ie, voltage 'drop outs'), the diode will isolate the converter from the undervoltage, and the converter will operate during the transient duration from the energy stored in the capacitor.

If there is concern about higher frequency noise, or so-called 'Electrical Fast Transients' (EFT) affecting the operation of the converter or circuitry connected to its output, an EMC filter can be added to the system as shown in figure 8.12b. This should not normally be required if the DC distribution and decoupling design is well executed.

With the exception of input fusing, the power system designer may choose none or all of the techniques presented here, depending upon the system requirements. To avoid unwanted replication of components, it is possible to partition the system so that the larger, more expensive, components can be shared between multiple DC/DC power modules. Figure 8.12c shows one such partitioning, where the transient protection and EFT filtering functions are shared between three power modules.

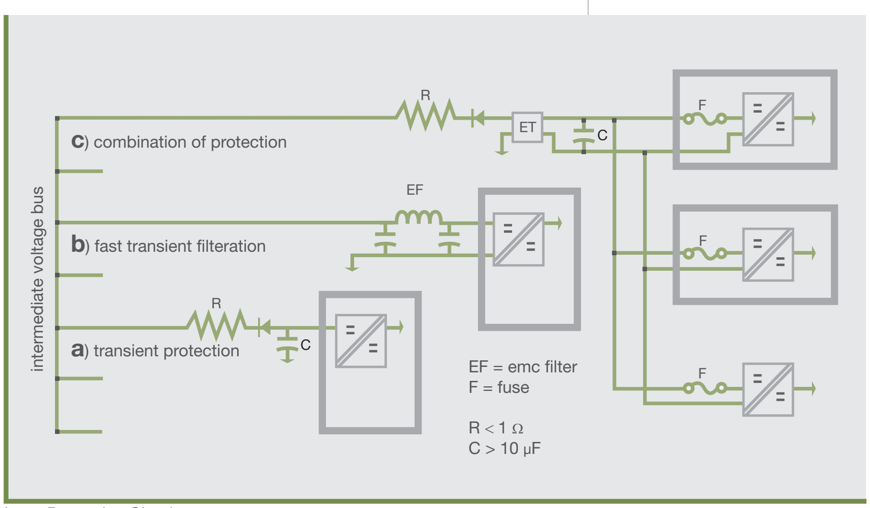

Input Protection Circuits

figure 8.12 a,b,c

Controls

Power system controls have had an interesting and varied history. Initially, they were simple or non existent. As systems became more complex and sophisticated, the power control function followed suit. Perhaps the best example is the high-end mainframe computer, where it is not uncommon to find a power control system that requires its own system enclosure, contains one or more computers, and is supported by literally hundreds of sensors and thousands of lines of programming code. These types of control systems are driven in complexity by the centralized philosophy of the power architecture. For such systems, there needs to be a centralized repository of information about what is happening from a power point of view at many different places, and mechanisms available to take appropriate actions based on this information. The result is very expensive power systems.

In today's environment, this trend is changing. We are now seeing more and more usage of decentralized power architectures, which inherently resolve many of the issues and complications addressed above. The merging of power conversion with the other functionally related electronics has resulted in diagnostics and service on a functional basis, rather than on a total system basis. The trend toward smaller, self-contained, single output DC/DC converters instead of large multi-output units has simplified many of the power sequencing needs. No longer is there a need at a centralized location for detailed information about current levels or overvoltage protection status of a specific converter in a remotely located function. The availability of small, cost effective and highly reliable DC/DC converters has resulted in system control designs with higher levels of simplicity and 'user friendliness'. There are, however, some remaining control functions that will find widespread usage. They are discussed below. The most basic and useful control function is the ability to turn the converter on and off by means of an external signal. This can be used to do power sequencing in a very simple and cost effective fashion. It can also be used to implement additional control functions designed by the user. All Ericsson DC/DC power modules are configured such that connecting one pin to ground disables the converter. This allows users who have no need for the control function to leave this pin unconnected, and the converter will automatically start up when the input power is applied. To disable the converter, the pin can be switched to ground either with a mechanical switch or by an external logic gate.

Another type of control that is very useful for telecom systems is the ability of the DC/DC converter to sense the input voltage and to turn itself off if the input voltage goes below a certain value. This function is sometimes referred to as "Input Undervoltage Lock Out". If the converters are operating from a battery supplied bus (48 V, for example), as the batteries discharge the input voltage to, the converters will be reduced. Since all switching converters are constant power devices (input current goes up as input voltage goes down), the battery load current will increase as the batteries discharge. The effect of this will be to increase the rate of discharge. Many systems will include battery management functions that disable the system when the battery is discharged to a predetermined point. For those systems lacking this function, the input voltage monitoring function internal to the DC/DC converters can prevent excessive battery discharge by automatically turning off the converters and removing load from the battery when the battery voltage drops to a specified level.

Some newer converter designs allow the user to program the voltage levels at which the converter will turn on and off. The turn-on level is set higher than the turn-off level. This hysteresis is built into the circuit so that the interaction between the converter input current and the battery voltage will not result in an oscillating condition, with the converter turning on and off as the battery voltage goes up and down. This implementation is shown in figure 8.13. The shaded areas represent the range of possible turn-on and turn-off due to the effects of circuit tolerances. Datasheets for individual converters should be consulted for specific voltage levels at which the input voltage sensing is activated. The user should also be aware that this function is provided mainly for the protection of the battery, and that the output regulation of the DC/DC converter may be guaranteed over a smaller range of input voltage. The minimum input voltage for regulated output at full load is provided in the datasheet.

Figure 8.14 depicts a system in which input undervoltage shutdown is implemented by means of the detector internal to each DC/DC converter. In this case, rather than using the default values built into the converter, the levels are custom tailored to the application by means of the programming resistors, R. The system designer should be aware that, due to circuit tolerances, each converter will turn on and off at slightly different input voltages, resulting in a range of battery voltages for which not all converters will be operational. Also, if the converters are heavily loaded, it is possible that their output voltages will drop below the regulation limits before the converter is shut down. This will occur if the shutdown voltage is less than the minimum input operating voltage. Another approach that could be used if decentralized control is desired is to shut down each converter by means of an

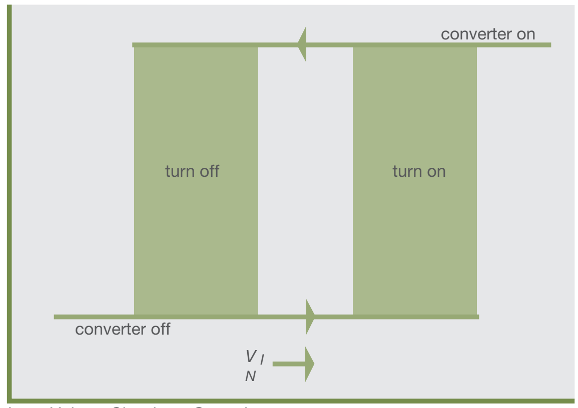

Input Voltage Shutdown Operation

figure 8.13

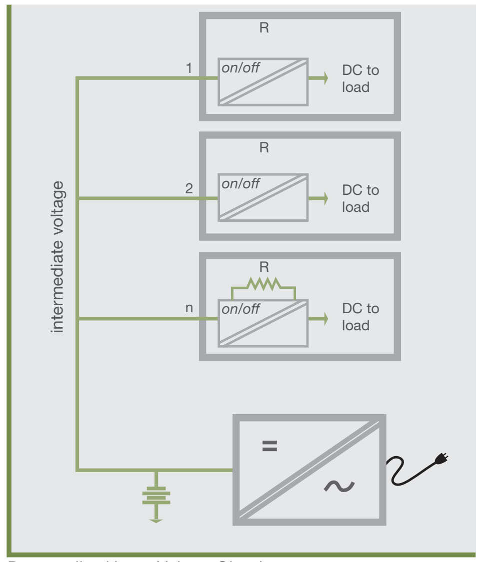

Decentralized Input Voltage Shutdown

figure 8.14

undervoltage detector on the DC output of the converter. This will make the input voltage shutdown point dependent upon the converter load current, but have the advantage that the minimum DC voltage at the load prior to shutdown can be guaranteed.

Another implementation of input voltage shutdown is shown in figure 8.15. In this system, the shutdown is directed from a centralized detector, and all of the DC/DC converters can be turned on or off at the same time. Depending upon its design, it could

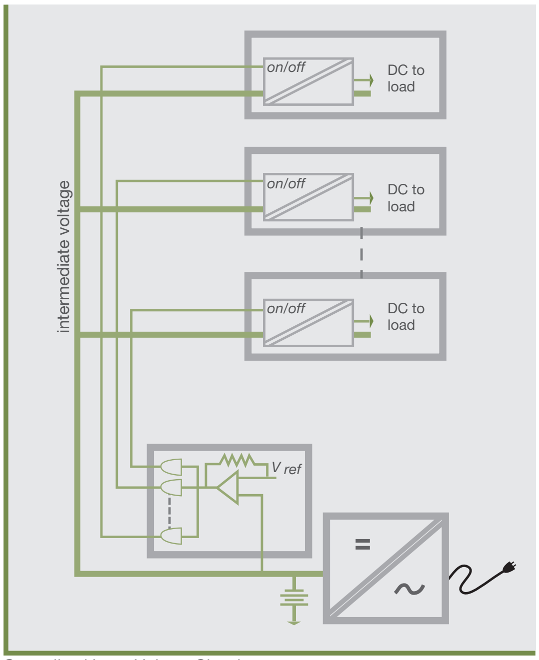

Centralized Input Voltage Shutdown

figure 8.15

also provide for tighter tolerances and greater flexibility on the voltage levels at which the converters are turned on and off. It does, of course, require additional complexity in terms of control lines, cabling and connector pin usage.

The ability to adjust the DC output voltage of a converter is often useful, and is called "output voltage trimming". The adjustment range does not need to be large. Rather, its most common usage is to compensate for the DC distribution voltage drop within the application, and to avoid the need for remote sensing. All Ericsson DC/DC power modules with provision for output voltage adjustment are designed to be set to their nominal output voltage without the need for any external components. They can be programmed either upward or downward from their nominal setting with a single external resistor. The power module datasheets will provide information on the implementation of this control function and on the selection of the proper value of the programming resistor.

The voltage trimming capability can also provide other benefits. With the trend to lower output voltages below 3 V, the performance of some circuitry is very sensitive to the applied voltage. With the trim function, the power module output can be adjusted to the exact value desired by the circuit designer. This approach allows for voltage adjustment late in the product design and production schedule without the need for changing the part number of the power module. "Voltage margining" is another common use of the trimming capability. Margining is a technique used in the product's final test cycle to verify that the design is robust and will tolerate the range of possible voltage variations over its lifetime. With the resistor controlled trim implementation of Ericsson power modules, the automated test equipment can easily adjust the output voltage both downward and upward from its nominal value to verify the circuit performance.

Remote Sensing

Remote sensing is a method of locating the point of optimal voltage regulation external to the power converter. It has been used in both linear regulators and switching converters for at least five decades and therefore is well understood by the power system designer. Its usage became popular during the period when centralized power systems were the dominant architecture. With centralized converters, remote sensing is almost mandatory for any voltage of 5 V or less with tight regulation requirements. With the increasing usage of lower operating voltages, remote sensing would be a universal requirement without the advent of decentralized architectures.

Remote sensing was a valuable tool for centralized systems, but has its share of problems. For each voltage level, two additional sense conductors must run between the load and the power converter. These conductors add complexity and cost to the system as well as requiring additional connector pins. Remote sensing also adds an additional failure mode to the power system. If either of the sensing lines opens, the converter will operate at a higher voltage level, which may or may not trip the overvoltage detector. This failure is an especially difficult one to diagnose correctly. Power converters are often needlessly replaced when the problem actually resides with the sensing system. This adds to repair time and cost.

Fortunately, remote sensing is not often needed with decentralized architectures, especially "Power per Board". Each converter powers a reasonably sized group of load circuits at low to moderate power levels. The converter is located physically close to the load. Consequently, remote sensing is rarely needed with the possible exception of voltages of 3 V and below at high current levels. Even in this circumstance, it is recommended that voltage trimming be considered as a method of setting the output voltage at the load rather than remote sensing. Unless the voltage set point needs to be maintained over a wide range of output current, voltage trimming will provide similar performance to remote sensing without the additional complexity and failure mode.

Paralleling

As the degree of decentralization increases, the current and power demands upon individual DC/DC converters becomes less. With 'Power per Function' and 'Power per Board' architectures, individual DC/DC converters in the 5 to 100 W power range are normally sufficient to provide the entire DC operating current for each voltage level in the function. In some cases, however, there may be a current demand that is higher than the output of a single DC/DC converter that can be accommodated within the building height restrictions of the system packaging. In such cases, the power system designer may consider the option of paralleling two or more DC/DC converters to obtain output currents higher than that of an individual module. Another circumstance in which paralleling is sometimes used is for the implementation of 'n+1' redundant designs, where one additional DC/DC converter in addition to the number required to power the system is used to provide for uninterrupted system availability in the event of one converter failure.

Paralleling adds complexity to the system and typically entails accepting some performance or cost compromises. It is suggested that other alternatives be investigated before the decision is made to parallel converters. Some of these alternatives are:

- Selecting a converter with a higher output rating.
- Re-partitioning the load so that no paralleling is required.
- Achieving redundancy on a functional basis rather than just with power – ie, multiple boards in a 'Power per Board' architecture.

If it is decided to parallel DC/DC converters, there are three techniques to consider. We will examine each of them, in order of increasing complexity and performance. The discussion will assume that two converters are paralleled, but the concepts can easily be extended to larger numbers of paralleled units.

Direct paralleling is the most straight-forward approach to paralleling DC/DC converters. With this technique, the DC outputs of the converters are directly connected to the common load. There are no other connections between the converters. The most common difficulty with this approach is load current sharing. Even though each converter will be set at the factory to the same nominal output voltage, there will be some variation in the absolute setting – typically on the order of 0.5 to 1%. If the output regulation of the converters is very good, the converter with the highest output voltage setting will supply most of the current, and may begin to go into its overcurrent region before the other converter supplies appreciable current.

In addition to this dependency on output voltage setting, the degree of current sharing is also affected by the following factors:

- The output regulation characteristic or 'stiffness' of the converters.
- The output voltage vs. temperature characteristic of the converters.
- The impedance of the distribution system between the converters and the common load.

The first two parameters are fixed once a DC/DC converter family is selected. The user has some control over the effect of the temperature coefficients. It is better to have all the converters mounted in a common thermal environment so that they experience as close to the same temperature transitions as possible. This will use the effects of the temperature coefficients to maximum advantage.

The distribution impedance is under the control of the power system designer. Increased DC resistance between the converters and the load will improve current sharing performance. The designer can control the degree of current sharing by changing this resistance value. It is important to keep the resistance between each converter and the load as equal as possible.

Figure 8.16 depicts how the distribution resistance affects load sharing. As RD is made larger, the voltage drop across it due to current from the higher voltage converter will bring the load voltage to the point where the lower voltage converter will supply current. The power dissipated in RD represents a loss of efficiency, so RD must be selected carefully to balance the need for current sharing with the need to minimize system power losses. Typical values for RD range from 20 mW for 5 A output converters to 100 mW for 1 A devices. In many cases, the required resistance can be integrated into the board distribution traces, and a separate discrete resistor is not needed.

The second technique is voltage matching. As described above, it is the slight differences in the output voltages that results in the non-equal current sharing performance. The output voltage of most DC/DC converters can be trimmed by means of a resistor external to the unit. By installing a variable resistor external to one of the converters, it can be trimmed so that it provides the same voltage as the other unit within the application. If the inclusion of a potentiometer and the need for manual adjustment are acceptable to the user, this

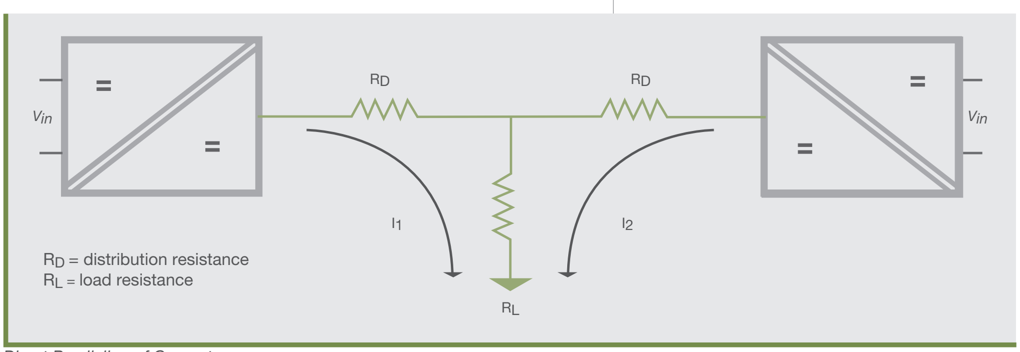

Direct Paralleling of Converters

figure 8.16

is an acceptable method for enhancing the paralleling performance and reducing or eliminating the need for current sharing resistors on the outputs. More information on this technique can be found in Ericsson application materials.

The third technique is to use active current sharing, if it is supported by the design of the DC/DC power modules being used. Active current sharing can be either "democratic" or "master-slave", depending upon the specific design of the converters. Active current sharing is the highest performance approach, but does require external components or interconnections. It is also possible to configure active current sharing by means of external circuitry for converters without internal provision for current sharing. More information on active current sharing can be obtained from Ericsson datasheets and applications notes.

If paralleling is done to enhance reliability (n+1 redundant configurations), converter failure modes must be taken into account. Because Ericsson uses only very high quality capacitors in the output filter, their power modules are very reliable. Most converter failures will not result in a short on the output, and the failed power module will 'isolate' itself from the remaining functioning parts of the system.

If the very small failure rate of the output capacitors remains a concern, the DC/DC power modules can be isolated in the event of failure by inclusion of a forward biased diode on the output of each unit. If this approach is taken, remote sensing may be required to achieve the desired load regulation performance. There will also be a reduction in efficiency. It is suggested that the failure rate of the diodes be compared with that of the power module output filter before adapting this approach.

One of Ericsson's most successful products, the PKF series DC/DC power modules, is ideally suited for applications where paralleling is required. It has been designed with a 'soft' output regulation characteristic that will normally allow for current sharing using the direct paralleling technique without the need for external resistors.

The information presented here has been general in nature since each type of DC/DC power module has unique criteria and considerations relative to paralleling implementation and performance. Datasheets and application notes for individual power module families contain more detailed information and suggestions for paralleling. Ericsson applications support personnel are also an excellent source of information and experience on paralleling of power modules.

Diagnostics

In the section on controls, we saw that as power systems are becoming more decentralized, the control systems are becoming less complex. The same is true for diagnostics. If the power architecture is planned correctly, the power diagnostic needs become very simple or even non-existent. With Power per Function and

Power per Board decentralized approaches, the need for specialized power diagnostic hardware and procedures is very minimal. Power per board even eliminates the need to stock DC/DC converters as a separate field spares inventory item!   8.17 summarizes the diagnostic implications of various power architectures.

The most important consideration is to keep the diagnostics as simple as possible. This will minimize the system complexity in terms of cables, connectors, and other sources of cost and failure rate. As shown in figure 8.17, the more decentralized architectures will provide significant advantage in this regard.

Architecture	Failure detection	Isolation procedure	Spares inventory
Centralized	Voltage sensors at each board	Extensive software or lengthy manual test procedure	Large power converters
	Monitoring of converter overcurrent		Distribution buses
			Diagnostic hardware
Power per shelf	Voltage sensor at converter	Simple test procedure	Modest power converter
Power per function	None - use funktion diagnostics	Simple test procedure	Small power card
Power per board	None - use funktion diagnostics	Replace board	None

Diagnostic Implications of Power Architecture Selection

figure 8.17

Decentralized Power System – Thermal Design

9

Decentralized power systems provide a vastly different thermal situation than conventional centralized power systems. The thermal design of systems utilizing high density board-mounted DC/DC converters is often a source of confusion for many first time users.

Introduction

There are several reasons for the confusion when first using decentralized power approaches. First, decentralized power architectures are a somewhat new practice in terms of implementation of actual systems, and many power designers are faced for the first time with the need to configure a design with power conversion and load circuitry on the same PC board. Secondly, there is some confusion in the industry on the definitions of some key thermal parameters. This can lead to false assumptions and sometimes to inappropriate designs. Perhaps the most commonly misused term is 'ambient temperature'. Finally, in an attempt to be as competitive as possible in the marketplace, some suppliers of power conversion modules advertise and promote

power handling capabilities that can only be achieved with extensive (and sometimes impractical) provisions for external cooling. A system designed with the 'headlines' from the DC/DC converter specification sheet may, in practice, be vastly undercooled and consequently
unreliable.

As an example, consider a DC/DC converter that is promoted as '200 watts'. If this unit is used in a free convection cooled environment with an equipment internal ambient temperature of 80 °C, its actual usable power output could be only in the vicinity of 4 W, whereas a competing product, rated at 25 W, but specified with board-mounted applications in mind, could actually produce its entire 2 W rating. This example will be examined in more detail later during the discussion of free convection cooling.

Parameter	Symbol	Description
Environmental Temperature	T_E	Air temperature that the end product is exposed to. Used for all methods of cooling. For equipment installed indoors it is equal to room temperature, TR.
Ambient Temperature	T_A	Air temperature external to power module or PBA. This will be higher than the environmental temperature due to heat dissipation within the product enclosure. Used for free convection and forced convection cooling.
Heatsink Temperature	T_H	Average temperature of heatsink attached to the power module. Will typically be slightly lower than the case temperature. Used for all types of cooling if heatsink is present.
Case Temperature	T_C	Temperature of the power module case. Used for free convection and forced convection cooling as well as conduction cooling if the case is the principal thermal path.
Pin Temperature	T_P	The average temperature of the power module pins. Will be very close to the temperature of the circuit board that the power module is mounted on. Used for conduction cooling when the power module pins are the principal thermal path.
Component Surface Temperature	T_{CS}	Temperature of a component within the power module, measured at the component's external surface. Used for all methods of cooling.
Component Core Temperature	T_{CC}	Temperature of the component interior. Used for all methods of cooling. For semiconductors it is equal to the Junction temperature, TJ.

In this chapter on thermal design, we hope to help rectify this lack of understanding about the thermal design of decentralized power systems. We will begin with some definitions of thermal parameters along with the corresponding nomenclature. We will then show detailed examples of the three most commonly used cooling mechanisms for decentralized power systems. To conclude, we will address some general trade-offs and options that the designer has available to balance the system thermal design with other important product characteristics.

Thermal Definitions

The following temperature definitions and symbols will be used. It is especially important to distinguish between environmental temperature and ambient temperature. The parameters are listed in order of generally increasing temperature. Unless otherwise specified, the unit of all temperatures referred to will be in degrees Celsius (°C).

In addition to temperature definitions, even basic thermal analysis will require the usage of thermal resistance. Fortunately, thermal resistance is an easy concept to understand and actual values easy to obtain. They are the thermal analog to electrical resistances, and indicate the ability of a thermal path to con-duct power (heat). A good thermal path or interface will have a low thermal resistance, and will need a small temperature gradient across it to conduct a given amount of heat or power. The same amount of heat flowing through a higher thermal resistance will create a larger temperature drop, analogous to voltage drop in an electrical system as shown in figure 9.1. Unless otherwise specified, the unit of all thermal resistances will be degrees Celsius per watt (°C/W).

With the seven temperature locations defined previously, a total of 21 thermal resistances can be defined. The following will be the most generally used.

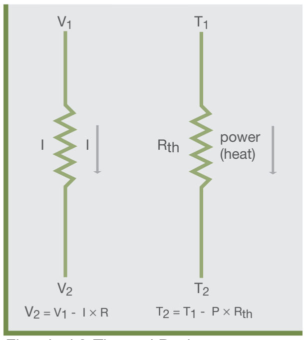

Electrical & Thermal Resistance

figure 9.1

Parameter	Symbol	Description
Case to Ambient	$R_{th\ C\text{-}A}$	Thermal resistance from power module case to air surrounding module. Used for free convection and forced convection cooling if no heatsink is present.
Case to Heatsink	$R_{th\ C\text{-}H}$	Thermal resistance from power module case to heatsink. Used for all types of cooling with heatsink. Heatsink to $R_{th\ H\text{-}A}$ Thermal resistance from the heatsink Ambient to air surrounding the heatsink. Used for all types of cooling with heatsink.
Core to Surface	$R_{th\ CC\text{-}CS}$	Thermal resistance from component core to the component surface. Used for all types of cooling. For semiconductor devices it is equal to resistance from Junction to Surface, $R_{th\ J\text{-}CS}$.
Component Surface to Case	$R_{th\ CS\text{-}C}$	Thermal resistance from component surface to the power module case. Used for all types of cooling.

As with the case of electrical resistances, thermal resistances are additive. In the case of a forced convection cooled DC/DC power module with a heatsink, for example, the thermal resistance from the junction of a semiconductor device to the ambient would be:

$$R_{th\ J\text{-}A} = R_{th\ CC\text{-}A} = R_{th\ CC\text{-}CS} + R_{th\ CS\text{-}C} + R_{th\ C\text{-}H} + R_{th\ H\text{-}A}$$

The remaining parameter to be defined is power, which will be expressed in watts (W).

Parameter	Symbol	Description
Input Power	P_I	Input power to power module
Output Power	P_O	Output power from power module
Power Dissipated	P_d	Power dissipated within the power module. Equal to input power less output power.

Conduction Cooling

There are three primary cooling mechanisms that could be used to cool electronic equipment – conduction, convection, and radiation. Radiation turns out not to be an appreciable source of cooling for the temperature ranges encountered in normal electronic systems. Convection will be covered in the next section. Here, we will address conduction, which is the mechanism that is the easiest to understand and analyze.

Conduction is the transfer of heat energy through a material or materials by means of transfer of energy between adjacent atoms or molecules. Unlike convection cooling, conduction cooling involves no net movement of any material or fluid. The ability of a material to act as an efficient conduction cooling medium depends upon the material's thermal resistivity, ρ, expressed in the units of °C-cm/W. Here are some thermal resistivities for materials commonly used in electronics:

Material	Resistivity (°C-cm/W)
Copper	0.25
Aluminum	0.48
Silicon	1.2
Alumina	6.0
Air	3050.0

Once the thermal resistivity of the material is known, the thermal resistance can be calculated from the physical geometry of the material using the relationship:

$$R_{th} = \frac{\rho\ l}{A},$$

where:

l is the length of the heat flow path, and

A is the cross sectional area.

See figure 9.2.

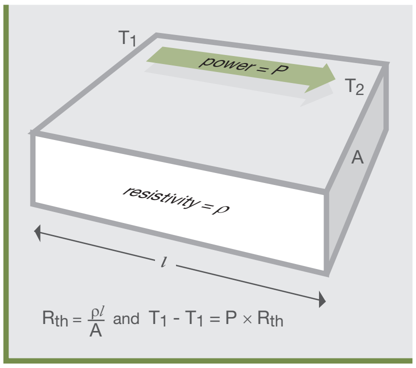

Calculations of Thermal Resistance

figure 9.2

Parameter	Symbol	Description
Case to Ambient	Rth C-A	Thermal resistance from power module case to air surrounding module. Used for free convection and forced convection cooling if no heatsink is present.
Case to Heatsink	Rth C-H	Thermal resistance from power module case to heatsink. Used for all types of cooling with heatsink. Heatsink to Rth H-A Thermal resistance from the heatsink Ambient to air surrounding the heatsink. Used for all types of cooling with heatsink.
Core to Surface	Rth CC-CS	Thermal resistance from component core to the component surface. Used for all types of cooling. For semiconductor devices it is equal to resistance from Junction to Surface, Rth J-CS.
Component Surface to Case	Rth CS-C	Thermal resistance from component surface to the power module case. Used for all types of cooling.

As an example, let's calculate the thermal resistance and temperature rise for an axial lead conductively cooled power diode, using the following assumptions as shown in figure 9.3

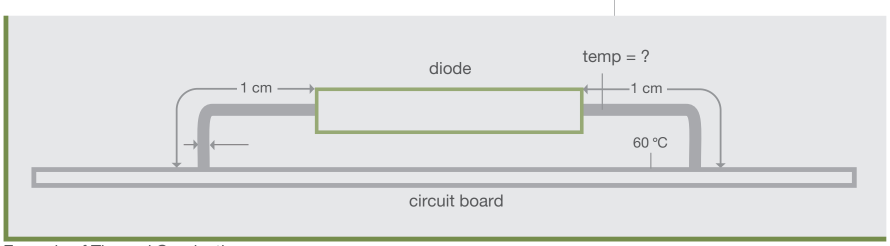

Example of Thermal Conduction

figure 9.3

The lead cross-sectional area will be

$$pr^2 = (3.14)\,(0.05)2 = 0.0079 \text{ cm}^2$$

The thermal resistance of each lead can be calculated from:

$$R_{th} = \rho\ l/A = (0.25)\,(1)\,/\,(0.0079)$$
$$R_{th} = 31.6\ °C/W$$

The total thermal resistance will be half of this value, since both leads will act in parallel, giving a net thermal resistance of 15.8 °C/W.

The temperature rise through the leads can now be calculated:

$$\text{Temp rise} = (2\ W)\,(15.8\ °C/W) = 31.6\ °C$$

The diode end of the lead will therefore be about 92 °C. The junction of the diode will be further elevated dependent upon the diode package construction. In practice, convection cooling of the diode body would slightly improve the thermal performance.

Conduction cooling is used at both extremes of the DC/DC converter power range. It is sometimes utilized for very high power DC/DC converters in computer mainframes that supply hundreds or even thousands of amps. These converters are cooled by circulating a cooling liquid either within the converter or within a 'cold plate' to which the converter is mounted. The other extreme is low power DC/DC converters, 10 W or less, which are designed to be cooled via conduction through the leads. We will use such a converter for an example of a conduction cooled power module.

The Ericsson PKF series is a popular Ericsson line of high performance board mounted DC/DC power modules. The series provides 3–15 watts of output power at a typical efficiency of 83% for the 5 V output version. It is packaged in a very low profile 18 pin package that looks something like a DIP IC package (see figure 9.4). The 18 metal pins are soldered to the ceramic substrate that forms the component mounting surface for the power module. These pins are designed to effectively remove all heat from the power module without resorting to any convection or conduction cooling to the other package surfaces. In practice, there will be some additional cooling via these other mechanisms, so the actual thermal performance will be somewhat better than that predicted by the analysis presented here.

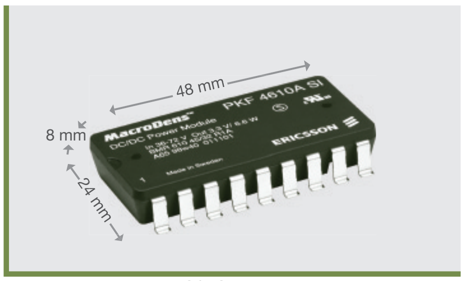

Ericsson PKF series DC/DC power module MacroDens™

figure 9.4

We will assume that a PKF 4611 DC/DC power module is soldered to a PCB and is used within a convection cooled product. We will discuss both free convection and forced convection systems in more detail later, but this information is not needed for this example. Because this DC/DC power module is designed for conduction cooling via the pins, it is specified for operation at a pin temperature rather than at an ambient or case temperature. The absolute maximum operating temperature specification for the Ericsson PKF is +95 °C pin temperature. As long as the pins are 95 °C or less, the converter will not be damaged and will provide basic functionality. All the datasheet specified parameters are guaranteed over the pin temperature range of -30 to +85 °C. We will now assign some product temperature values. Since this is an arbitrary example, we will use values that would be typical in practice for many systems. In an actual application, the temperatures shown below would be based on measurements made in a prototype or model of the system, or could be estimated based upon convective modeling techniques. The approximate power dissipation of the DC/DC power module should be included in the measurement or estimate if it is a significant percentage of the total power dissipated by the board. If we assume that the PKF is putting out its full 6 W with 83% efficiency, the power dissipated by the DC/DC power module can be calculated by subtracting the output power from the input power:

P_O = 6 W P_I = 6 W / 0.83 = 7.23 W
P_d = 7.23 - 6 = 1.23 W

A drawing of the system is shown in figure 9.5.

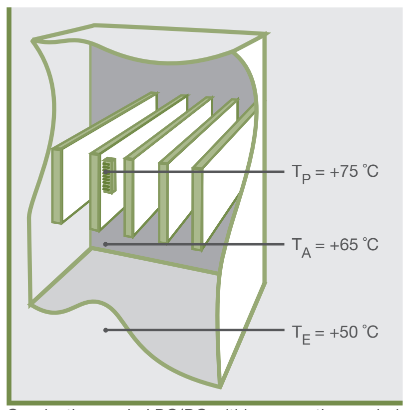

Conduction cooled DC/DC within convection cooled equipment

figure 9.5

T_E = 50 °C (Room environment)
T_A = 65 °C (Air surrounding the circuit board)
T_P = 75 °C (Converter pin and circuit board temp at converter location)

This is actually all that is required to apply this conductively cooled module. As long as the board and pin temperature is less than 85 °C, which it is in this case, the thermal design will be successful. One circuit board temperature measurement may be all that is required to verify a valid design. For completeness, we will look inside the DC/DC power module and show how conductive heat transfer is handled as part of the design and how the resulting system assures that component temperatures are within safe limits for reliable operation. Again, it should be noted that the following analysis is not required when designing with these DC/DC power modules. Only the circuit board temperature at the interface to the power module will normally be needed.

During the DC/DC power module design process, every component must be addressed to assure proper operating temperature. We will look at one critical component, the output rectifier diode. This component has one of the higher power dissipations in the converter, and will consequently tend to operate at a higher temperature than most other internal components. Assuming an averaged diode current of 1.2 A and a 0.6 V forward voltage drop, the average power dissipation in the diode is 0.72 W. This diode is contained in a surface mount package that is soldered to the ceramic substrate of the DC/DC power module. The diode is conduction cooled, with the main heat path through the heavy SMD leads. The diode manufacturer guarantees a thermal impedance from junction to lead of 20 °C/W. We can now calculate that the junction temperature is elevated from the substrate temperature by:

$$(0.72\ \text{W})\ (20\ \text{°C/W}) = 14.4\ \text{°C}$$

By making detailed temperature measurements at different locations on the ceramic substrate with known amounts of power dissipated on the substrate, Ericsson can char-acterize the thermal performance of the substrate. The data indicates that the average temperature differential from the diode to the edges of the substrate is 3.7 °C.

Ericsson has also characterized the thermal impedance of the DC/DC power module pins from the point they are soldered to the substrate to the solder connection to the PCB that the module is mounted on. This impedance is about 2.7 °C/W for the total of the 18 power module pins. The entire 1.23 W of maximum power module dissipation flows through this resistance, giving a temperature drop of:

$$(1.23\ \text{W})\ (2.7\ \text{°C/W}) = 3.3\ \text{°C}$$

We can now construct a thermal model of the entire system, extending from the external environment to the actual silicon junction of the DC/DC power module output rectifier diode. This model is shown in figure 9.6. Note how the temperature drops calculated above are added to the power module pin temperature imposed by the system to allow temperatures internal to the power module to be determined.

We find that the diode junction temperature is 96 °C. This is well below the recommended maximum temperature for good reliability, which is 120 °C. Using this model, we can easily extend the analysis to an application where the power module pin temperature was equal to the absolute maximum rated 95 °C

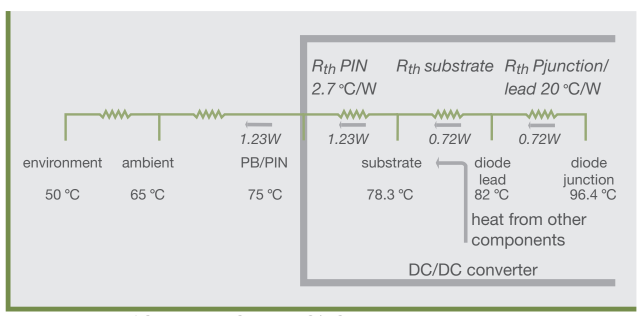

Thermal Model of Conduction Cooled DC/DC converter

figure 9.6

value rather than 75 °C. The 20 degree extra differential would extend into the model all the way to the diode junction, raising the junction temperature to:

$$T_{Jmax} = 96.4 + 20.0 = 116.4\ °C$$

Note that this is still well below 120 °C, indicating that this power module design is extremely conservative and should result in very high levels of reliability. The conservative thermal margins inherent in the design make the PKF series of power modules very applicable to equipment located in remote where T_E can extend to +65 °C. Reliability performance during normal conditions, with T_E in the +5 to +35 °C range is exceptional.

Free Convection Cooling

In convection cooling, heat is removed from the body being cooled by means of physical motion of a fluid, typically air. In free convection applications there is no external force applied to create the air movement, the movement being caused only by the natural thermal movement set up by the heat being dissipated. In forced convection systems the air is continuously moved by means of a fan or blower. Airflow rates in free convection applications are low, often between 0.2 and 0.3 m/s (1 m/s ≈ 200 lfm). For practical forced con-vection systems the rate of air movement is between 0.5 and 4.0 m/s.

Convection cooling is more difficult to model and analyze than conduction cooling. The ability of a convective interface to remove heat is dependent upon the surface area of the interface, the temperature differential between the body being cooled and the cooling fluid, and a parameter called the film coefficient of heat transfer. This coefficient, in turn, depends upon the temperature of the interface and the velocity of the fluid. Consequently, no simple equation exists that can completely model the cooling behavior. Instead, empirical measurements are often used to characterize the con vective cooling, and converter manufacturers typically publish this data in their datasheets. More recently, cost-effective PC based thermal simulation software has become available. These programs provide a very useful tool for the analysis of convectively cooled systems. In general, convective cooling can be enhanced by:

- Reducing the temperature of the cooling fluid.
- Increasing the surface area of the body being cooled.
- Increasing the velocity of the cooling fluid.

Free convection cooling has the advantage that it is very simple and reliable; the acoustic noise, preventative maintenance, cost, and reliability impact of fans that would be needed for forced convection are not present. On the down side, significantly less power dissipation per board area is possible than with forced convection. Also, free convection systems are orientation sensitive, with vertical board orientation required for maximum convective cooling effect.

As an example of applying a DC/DC converter in a free convection environment, we will look at a typical telecom system that uses 'power per function' distributed power. We will assume that each function is powered by a medium power DC/DC converter mounted on a circuit board in a rack and shelf arrangement. Each converter will power several adjacent boards. An Ericsson PKE series DC/DC power module will be used for this example, (see figure 9.7). The PKE is designed to be convection cooled, and is specified for a maximum case temperature of 115 °C. The thermal resistance from case to ambient, $R_{th\ C\text{-}A}$ is characterized by Ericsson and specified at 5 °C/W.

The key parameter that the system designer must know to apply the PKE DC/DC power module is the ambient temperature (T_A) within the equipment. T_A will be elevated above the room environmental temperature (T_E) by the internal power dissipation in the equipment cabinet. For this example, we will use a T_E toward the upper end of room environments encountered in actual telecom applications for non temperature controlled environments: 60 °. The internal temperature rise will be determined by the natural or free convective air movement set up within the equipment. A typical value of temperature rise for such a system, which will be used in this example, is 20 °C. Thus T_A will be 80 °C.

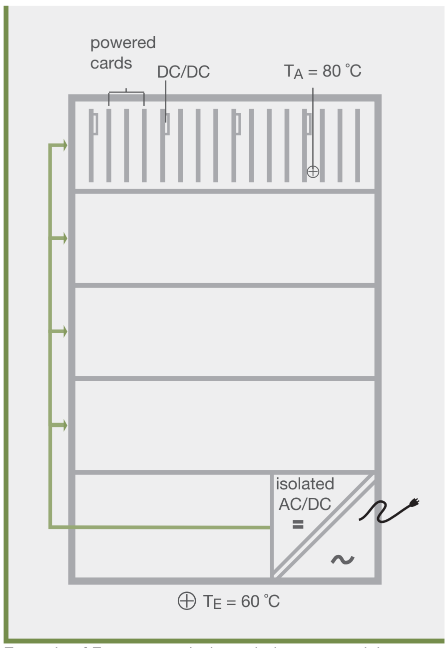

Example of Free convectively cooled power module

figure 9.7

The PKE 4211 PI can supply 5V at 5A, giving a maximum output of 25W. The efficiency is specified at 80%, but is typically better. We will assume that the PKE is putting out its maximum rated power of 25W. The input power and power dissipated within the PKE 4211 PI can be calculated as:

$P_I = P_O / \text{eff} = 25 / 0.8 = 31.2 \text{ W}$
$P_d = P_I - P_O = 31.2 - 25 = 6.2 \text{ W}$

The temperature rise of the case above ambient can now be calculated to be:

Temp Rise = ΔT = (6.2 W) (5 °C/W)
= 31 °C

The maximum case temperature will be:

$T_C = T_A$ + Temp Rise = 80 + 31 = 111 °C

This is below the maximum 115 °C case temperature, and the thermal design is satisfactory. The PKE 4211 PI should provide reliable operation even at its maximum rated 25W output. This ease of application is due to the fact that the PKE was designed for free convection cooling.

It is instructive to explore the performance of another DC/DC converter in this same application. We will use a popular ‘200 watt’ series converter that incorporates a heavy aluminum baseplate as a cooling interface. The converter is rated at 80% efficiency and +85 °C maximum baseplate or case temperature. The $R_{th\ C\text{-}A}$ is rated at 5.1 °C/W. If we were to put this converter in the free convection environment described above with T_A equal to 80 °C, we would get some surprising results.

The allowable temperature rise is only 5 °C, and the allowable converter dissipation can be calculated as follows:

$$P_d = \text{Temp rise} / R_{th} = \frac{5\ °C}{5.1\ °C/W} = 0.98 \text{ W}$$

Using the 80% efficiency we can calculate the maximum output power:

From $P_I = P_d + P_O$ and $P_I = P_O/eff$,
$P_O = P_d (eff)/(1 - eff)$
$P_O = (0.98)(0.8) / (1 - 0.8) = 3.9$ W

Here is a situation where a ‘200 watt’ converter supplies only 4 W, while the PKE 25 W power module, designed for free convection, supplies its full rated 25 W. In addition, the ‘200 watt’ converter is larger and more expensive. This example highlights the importance of using a converter in the environment for which it was designed. It also illustrates the benefit of high baseplate or case temperature ratings. In the following section an example will be shown where the ‘200 watt’ converter is utilized in an environment more suitable to its design.

Forced Convection Cooling

In principle forced convection cooling is identical to free convection, the only difference being that the moving air creates a more rapid interchange of the air at the interface to the body being cooled. This interchange results in much greater cooling efficiency, making the convective thermal resistance between the converter or heatsink and the cooling air effectively much lower. This effect increases with increasing air flow, but begins to have diminishing returns at air velocity values above 4 m/s. For example, figure 9.8 shows the effect of air flow on the thermal resistance of the baseplate of a ‘200 watt’ converter without additional heatsinking.

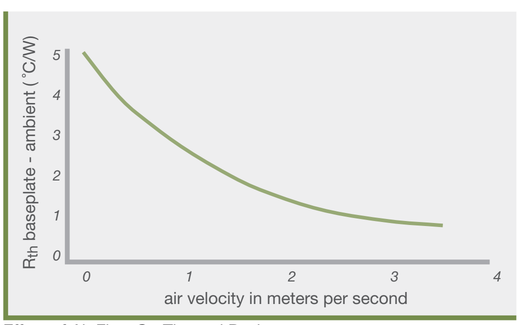

Effect of Air Flow On Thermal Resistance

figure 9.8

Even though forced convection and free convection are similar in concept, their practical implementation is significantly different. A fan or blower is needed to create the air movement for forced convection systems. In many systems the power requirements, acoustic noise, and service requirements for fans are significant. The fan or blower also represents a reliability concern since the motor and fan or blower bearings have mechanical wear mechanisms and a finite service lifetime. In return for the above inconveniences, forced convection offers greatly improved thermal performance. In typical systems, for a given board area and environmental temperature, forced convection can cool about 4 times the power per board compared to free convection cooling. Forced convection can also significantly reduce the required size of equipment enclosures.

As an example of forced convection cooling we will use a power-per-shelf type of telecom system. We will assume that each shelf contains a ‘200 watt’ DC/DC converter that supplies operating voltage to each board location on the shelf. The system includes a blower that supplies an air velocity of 3.5 m/s through a vertical plenum supplying the area containing the DC/DC converters. A sketch of the system configuration is shown in figure 9.9.

We will assume a room environment temperature of 50 °C. With forced convection the temperature rise internal to the equipment is typically less than for free convection due to the rapid interchange of cooling air. In this example we assume that the DC/DC converter on each shelf causes a 2 °C temperature rise in the cooling air. This will result in slightly different ambient temperatures for the converter in each shelf. The worst case location will be the upper shelf, with a T_A of 56°C. The exhaust air temperature will be 58°C. The DC/DC converter efficiency is 80 %, and we will assume an output power from each converter of 150 W. The maximum converter baseplate temperature is rated at 85 °C.

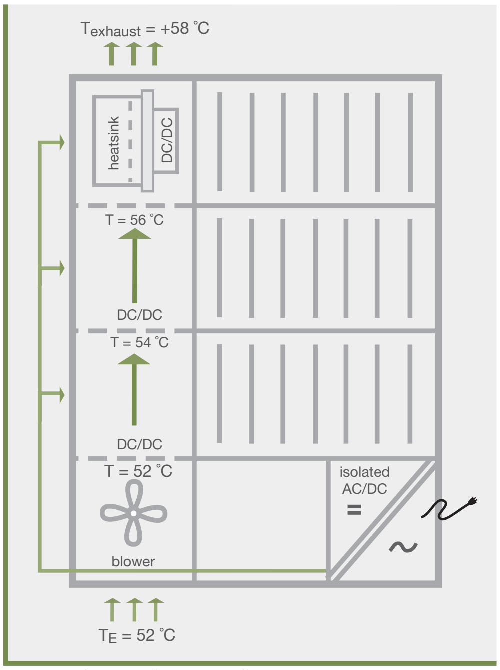

Example of Forced Convection System

figure 9.9

We will first check the thermal performance of the system without any external heatsinks on the converters. Referring to figure 9.8, we see that at an airflow of 3.5 m/s, the converter baseplate to ambient thermal resistance is 1.0 °C/W. The power dissipated in the converter is next calculated as:

$$P_I = P_O / eff = 150 / 0.8 = 187.5\ W$$
$$P_d = P_I - P_O = 187.5 - 150 = 37.5\ W$$

Using the above result and the thermal resistance obtained previously, we can determine the temperature rise of the baseplate above ambient:

$$\text{Temp Rise} = \Delta T = (1.0\ °C/W)(37.5\ W)$$
$$\Delta T = 37.5\ °C$$
$$\text{Baseplate Temp} = 56 + 37.5 = 93.5\ °C$$

This is well above the 85 °C maximum rating, and the design is grossly unacceptable.

At this point, we have determined that a heatsink is required to lower the effective thermal resistance from baseplate to ambient. We will now calculate how much heatsinking is required. The maximum allowable temperature rise is:

$$T_{BASEPLATE} - T_A = 85 - 56 = 29\ °C$$

Using this result along with P_d, we can find the required thermal resistance:

$$R_{th} = \text{Temp Rise} / P_d = 29\ °C / 37.5\ W = 0.77\ °C/W$$

The thermal interface between the converter baseplate and the heatsink isn't perfect, and has some thermal resistance that is effectively in series with the thermal impedance of the heatsink itself. This interface resistance can be assumed to be in the range of 0.2 °C/W. Taking this into account, the required resistance of the heatsink is:

$$0.77 - 0.2 = 0.57\ °C/W$$

We can then consult heatsink catalogs to find a heatsink that will interface with the converter and achieve a heatsink to air thermal resistance of less than 0.57 °C/W at an airflow of 3.5 m/s. There is a commercially available heatsink that is rated at 0.48 °C/W with an airflow of 3.5 m/s. This unit should work, and achieve a slight safety margin, which can now be calculated:

$$R_{th} = 0.48 + 0.2 = 0.68\ °C/W$$
$$\text{Temp Rise} = (0.68\ °C/W)(37.5\ W)$$
$$= 25.5\ °C$$
$$\text{Baseplate Temp} = 56 + 25.5 = 81.5\ °C$$

We have a 3 or 4 degree safety factor, which should be considered barely acceptable. A sketch of the resulting power converter assembly is shown in figure 9.10.

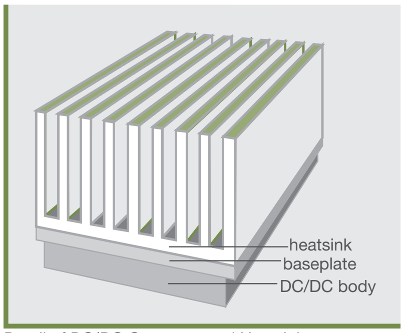

Detail of DC/DC Converter and Heatsink

figure 9.10

It should be noted that the heatsink selected above is approximately 3 times the volume of the DC/DC converter itself and provides barely acceptable performance, even with a very high airflow. This is an example of why volumetric density claims for converters must be approached cautiously. The system's thermal design, including heatsinks and cooling air requirements, must be understood before deciding on the appropriate DC/DC converter.

The above example used an established DC/DC converter from another supplier's product line. To show how newer technology can vastly improve system designs from a thermal point of view, we will also look at placing an Ericsson PKL power module into the above application. The PKL 4110 PI provides 3.3 volts at up to 165 watts of output, so one of these modules can easily source the 150 watts needed in this example. The PKL uses synchronous rectification to increase its efficiency to 90% and also has a maximum case temperature rating of +100 °C. Both of these characteristics will help dramatically in this application. By referring to the PKL datasheet we find that the module will source well over 150 watts at an ambient temperature of 56 °C at 3.5 m/s without any heatsink! This is accomplished in a converter volume of 2.42 in^3. The alternative solution shown above with a heatsink requires a volume of approximately 22 in^3. This almost ten times improvement in power density dramatically illustrates the influence of operating efficiency and high case temperature ratings.

Indeed, there are many trade-offs to be made in terms of reliability, density, cost, and packaging before a cooling technique and conv erter technology can be intelligently selected. In the next section, we will explore some of these trade-offs.

Combined Cooling Techniques

There are some applications and converter types that combine both conductive and convection cooling. Using both cooling methods enhances the overall effectiveness of the cooling system at the expense of complicating the analysis required to estimate the thermal performance. One such situation is the Ericsson PKG series of 60 W DC/DC power modules. These power modules have two effective cooling paths: conductive cooling through the pins to the circuit board similar to the PKF series and convection cooling through the case or auxiliary heatsink similar to many other DC/DC con verters. The combination of the two cooling paths allows for exceptional thermal and reliability performance at power levels up to 60 W. We will develop a thermal model for the PKG series of power modules and show how estimates of thermal performance can be obtained for various application conditions. Figure 9.11 shows the thermal model for the PKG series. The power dissipated by the power module, P_d, flows out through two parallel thermal resistances. $R_{th\ sub\text{-}P}$ is the thermal resistance from the power module ceramic substrate to the module's pins. This thermal resistance is a function of the module's internal components, materials and physical structure, and is not affected by ambient air temperatures or airflow rates. The typical value of $R_{th\ sub\text{-}P}$ is 2.5 °C/W.

The thermal path from the power module's ceramic substrate to ambient air is modeled by $R_{th\ sub\text{-}A}$. This thermal resistance combines the thermal conductivity from the ceramic substrate to the power module's case and the effective thermal resistance from the case to ambient air. $R_{th\ sub\text{-}A}$ is dependent upon the rate of the airflow across the power module. Ericsson offers an auxiliary heatsink for the PKG series of power modules to further enhance the thermal path to ambient air.

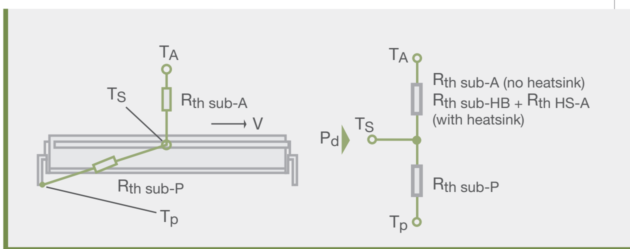

Thermal Model of PKG Power Module

figure 9.11

When this heatsink is used the thermal model is modified by including the thermal properties of the heatsink. In this case, the total thermal resistance from the substrate to the ambient air is lower, and is modeled by $R_{th\ sub\text{-}HS} + R_{th\ HS\text{-}A}$. The dependencies of the thermal resistances on the air velocity with and without the heatsink are shown in figure 9.12.

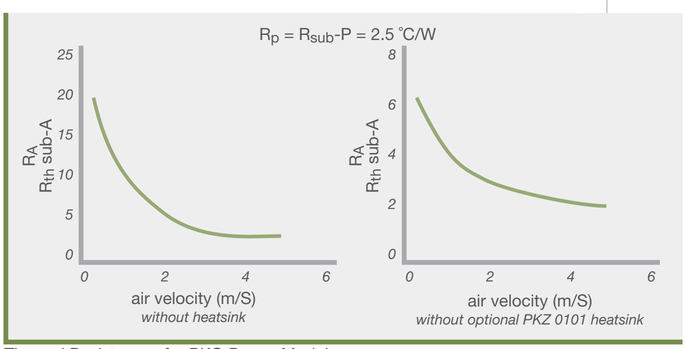

Thermal Resistances for PKG Power Module

figure 9.12

In order to obtain a generalized equation for situations with and without the heatsink, and to simplify the nomenclature, the term R_A will be used to designate the effective total thermal resistance from the power module substrate to ambient air. For a given airflow, the value of R_A will be smaller when the heatsink is used, as shown in figure 9.12. Similarly, the term R_P will be used to designate the thermal resistance $R_{th\ sub\text{-}P}$.

Solving the thermal model of figure 9.11 for the power module substrate temperature, T_S, yields the following expression:

$$T_S = T_A + \frac{R_A}{R_A + R_P} \{T_P - T_A + (P_d)(R_P)\}$$

Where:
T_S = Power Module Substrate Temperature
T_A = Ambient Air Temperature
T_P = Module Pin Temperature
P_d = Power Dissipated by Power Module

Since $R_P = R_{th\ sub\text{-}P} = 2.5$ °C/W, the expression can be further simplified to:

$$T_S = T_A + \frac{R_A}{R_A + 2.5} (T_P - T_A + 2.5\ P_d)$$

As an example, we will now use the equation presented above to calculate the substrate temperature of a PKG DC/DC power module for a given set of operating conditions. Assume that:

- Module is PKG 4611 PI (5 V, 60 W).
- Output power is 50 W.
- Power module efficiency is 86%.
- Pin temperature is 60 °C.
- Ambient air temperature is 50 °C.
- Airflow is 1 m/s.
- No heatsink is used.

We first calculate the power dissipated in PKG 4611 PI as:

$$P_d = P_I - P_O = (P_O/\text{efficiency}) - P_O$$
$$P_d = 50/0.86 - 50 = 8.14\ W$$

Using the left curve of figure 9.12, we find that R_A at 1 m/s is 10 °C/W.

Inserting these numbers into the equation for T_S yields:

$$T_S = 50 + \frac{10}{10 + 2.5}[60 - 50 + 2.5(8.14)] = 74.28\ °C$$

This is below the 100 °C maximum rating, so reliable operation will be achieved.

How much would the addition of the heat-sink reduce the substrate temperature? To answer this question, we use a value of R_A obtained from the right hand curve in figure 9.11.

We find that R_A has a value of 4.2 °C/W at an airflow of 1 m/s. Using this value in the equation for T_S results in the following:

$$TS = 50 + \frac{4.2}{4.2 + 2.5}[60 - 50 + 2.5(8.14)] = 69.03\ °C$$

We find that the addition of the heat-sink reduced the power module substrate temperature by about 5 °C. This temperature reduction would further enhance the reliability of the power module within the application.

Design Trade-offs

As we have seen in the preceding pages, there appears to be no one best cooling method for DC/DC converters. Conduction, free convection, and forced convection all have their place in the power system designer's toolbox, with the best choice being dictated by a variety of influences, some under the control of the designer and others imposed by system design constraints. In figure 9.13 we have attempted to summarize the typical usage, advantages and disadvantages, and some design trade-offs associated with each cooling methodology.

Perhaps some commentary on the design trade-offs would be useful. Control of temperatures (from room environment to the power module interface) is universally helpful in achieving good thermal performance no matter what cooling technique is used. The cooler the thermal interface, whether the circuit board and pin temperature for a conduction cooled power module or the ambient air temperature for a convectively cooled power module, the lower the internal temperature of the power module will be, and the better the reliability of the system will be. Of equal importance is the maximum temperature rating on the power module, T_P for conduction cooling, and T_C for convection cooling. This should be as high as possible. We saw in a previous example the limitations of using a power converter with a 85 °C maximum base-plate temperature. The object is to maximize the temperature differential between the maximum temperature rating of the power converter and the source of cooling. The temperature differential will in some cases determine if external heatsinks are required on power converter modules, and if so, how large the heatsinks need to be. Maximum pin or case temperature should be a very key criterion in the selection of the power converter module.

	Conduction	Free Convection	Forced Convection
Typical Module Usage	Low Power Board Mounted	Low to Medium Power Board Mounted	Medium to High Power Board Mounted or 'Centralized'
Advantages	Low Cost High Reliability Ease of Thermal analysis Not orientation sensitive	Low Cost High Reliability	Higher Density
Disadvantages	Limited selection of module suppliers	Low density Orientation sensitive	Less reliable More Complexity Field setrvice required More costly
Design Trade-offs	T_E, T_A, T_P Reliability (T_{CC}) Efficiency Maximum Pin Temperature	T_E, T_A Reliability (T_{CC}) Efficiency Board Area Maximum Case Temperature	T_E, T_A Reliability (T_{CC}) Efficiency Board Area Air Flow Maximum Case Temperature

Comparison of Cooling Techniques for DC/DC Converters

figure 9.13

In forced convection systems the amount of airflow can be increased to improve the cooling performance. Practical limits are soon approached, however. As shown in figure 9.8, increasing airflow above 4 m/s starts to have diminishing returns. Also, at high airflow rates the acoustic noise generated by the fans, blowers, and moving air create a problem for people in the vicinity of the equipment. Most equipment manufacturers have a maximum sound level that they will accept for the product, and this limit will impose restrictions on the amount of airflow that can actually be accommodated.

Reliability is also a trade-off. The design examples discussed here are predicated upon keeping the semiconductor junction temperatures within the power module below 120 °C at max ambient temperature. The basis for these kinds of guidelines will be explained in the chapter on Reliability. Even greater reliability can be achieved at lower junction temperatures. So for very critical systems, the designer may choose to operate the power modules (as well as other electronics) at reduced temperatures to improve reliability. This will need to be done at the expense of at least one of the other variables, such as board area, airflow, or volume needed for heatsinks.

Board area used for circuitry or power conversion functions is a critical factor in system design. For obvious reasons designers want as much density as possible. They must be aware, however, that density generally goes hand-in-hand with higher temperatures and more difficult thermal solutions. For a fixed temperature rise, a given board area and board spacing will only support a certain amount of power dissipation. It is not easy to quantify the amount of power supportable by a convection-cooled board due to the large number of variables involved. We have, however, included a rough 'rule-of-thumb' estimate based upon the popular double extended Euroboard used for many telecom type systems. This board is 220 mm deep and 233 mm high, giving a board area of 5.13 dm^2. Figure 9.14 indicates that in typical applications this board supports an average power dissipation of 5 W for free convection cooling and 20 W or more for forced convection. A power dissipation of 5 W per board in free convection corresponds to a temperature rise of 20 °C. The amount of power able to be cooled on a given board also depends upon the height of the components and upon the total surface area they expose to the cooling air. Very planar structures will not cool as well as higher components. This estimate assumes a normal mix of surface extensions and building heights.

The aspect ratio of the board is also a factor in addition to the total area. For a given area, boards with a larger dimension in the direction of the cooling airflow will be hotter than boards extended in the other dimension. This effect is due to the 'pre-heating' of the cooling air by components upstream of those being cooled, and is shown in figure 9.14. Components located at the lower edge of the board have a better cooling environment than those at the upper edge. This should be taken into account when doing the board layout, placing components requiring the most cooling lower on the board. Power system designers should be very aware of the location of the power

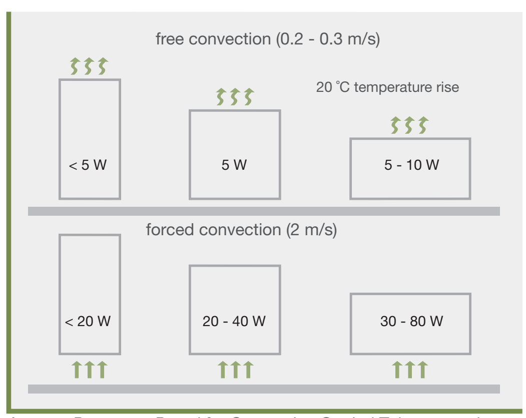

Average Power per Board for Convection Cooled Telecom equipmen
Board Area is 5.13dm^2 (equal to double Euroboard)

figure 9.14

module relative to other components on the board. If it is 'downstream' from other circuitry, make sure that the preheating effects due to the power dissipation of this circuitry is taken into account when defining the ambient temperature for the power module.

For systems implemented with board-mounted power modules, with the power module and load circuitry sharing the same board, it is possible to estimate the amount of board area allocated to the power conversion function by considering the respective thermal performance of the two types of technology. The areal power density (W/cm^2) of power modules is higher than that of logic and load circuitry. For the PKE and PKF Ericsson power modules operating at 80% load, for example, the power dissipated by the power module per board area is about 0.06 W/cm^2. The figure also includes the additional area required for connections and mounting purposes. If additional filtering is required it will decrease this figure. For the board as a whole, the areal power density can be estimated from the values shown in figure 9.14. Using the standard Euroboard aspect ratio, we obtain the following estimates for free and forced convection:

Free convection 5 W / 513 cm^2 = 0.01 W/cm^2
Forced convection 20 W / 513 cm^2 = 0.04 W/cm^2

We can now compare the power density of the power module with the board:

Free convection	0.06 / 0.01 = 6.0
Forced convection	0.06 / 0.04 = 1.5

We have found that the power module has between 1.5 and 6.0 times the areal power density of the board as a whole. We must now take into account the percentage of power dissipated in the power module. If we assume a power module efficiency of 82% for PKF and 84% for PKE, then 18% and 16% respectively of the power will be dissipated in the power module and the rest in the remainder of the board. Combining this with the power areal density ratios obtained above, we can now estimate the percentage of board area allocated to power conversion:

Free Convection	18% / 6 = 3%
Forced Convection	16% / 1.5 = 10%

These values can be considered rough 'rules-of-thumb' for use in initial design sizing of board layouts. They indicate that with today's board-mounted power modules, a very small percentage of board area is required to provide the benefits of power per board architecture.

The final trade-off parameter is efficiency. It is controlled by the DC/DC power module manufacturer, but the system designer must be very aware of the efficiency, and it should be a very important selection criterion when choosing the converter to use. Efficiency has a very large impact on all aspects of thermal design. For 25 W power modules with 5 V output, the range of available efficiencies is roughly 75% to 90%. At first glance, this may appear to be a small variation. In actuality, it is a very significant difference. Let's calculate the power dissipated by the converter, assuming a 25 W output:

for 75%, $P_I = 25 / 0.75 = 33.3$ W
$P_d = 33.3 - 25 = 8.3$ W

for 90%, $P_I = 25 / 0.90 = 27.8$ W
$P_d = 27.8 - 25 = 2.8$ W

Comparing the two power dissipations,

8.3 / 2.8 = 2.96

The 75% efficient converter dissipates 3 times more power than the 90% efficient converter!

This 20'0% increase will have a large impact on requirements for board area, heat sinking, and ultimately on system reliability and cost. The benefits of converter efficiency to the system designer cannot be over emphasized. Efficiency, along with reliability, should be at the top of the system designer's priority list. A good power module supplier will work with the system designer to help him or her understand the actual efficiencies to be expected from the system implementation and operating conditions.

The intent of this chapter was to explain the basics of power converter thermal performance and thermal design of the power system. It has, by necessity, been a brief overview. The reader is encouraged to refer to the thermal specifications in the datasheets of specific Ericsson power modules. The detailed information required for designing with that module will be found there. Ericsson also has a variety of application notes and design notes that touch on thermal design.

Perhaps the most useful tool for the power system thermal designer is the "thermal calculator" provided on Ericsson's website. This is an interactive tool that allows the user to enter the specific operating conditions for an Ericsson power module. Variables under control of the user include ambient temperature, output power, airflow velocity, board-to-board spacing and PCB board size. The calculator then does an analysis of the proposed cooling environment and returns values of temperature at various locations in the system as well as a prediction of the failure rate of the power module when operated under those actual conditions. The thermal calculator can be a very powerful and easy way to conduct "what if" studies to determine how different cooling designs affects the design margins and reliability.

Manufacturing and Packaging Considerations 10

As standard power modules become more sophisticated, packaging, materials and manufacturing processes are becoming a major factor in the selection process. Manufacturing and packaging efficiency are equally important for the end user of the modules, with emphasis on automated environmentally friendly assembly.

Introduction

When constructing systems using decentralized power architectures and distributed power modules, it is helpful to think about power converters differently than in the past. It was formerly assumed that a power supply was large and heavy. It was installed manually and retained with bolts or screws. Because of the mass, it had to be mechanically tied to the frame or structure of the product. If it were mounted on a board, when the product was subjected to the shock and vibration environments encountered during shipping and usage, the power supply could overstress the circuit board and result in mechanical fatigue, fracture, or failure. The same can be true of some of today's larger and heavier high density DC/DC converter products.

Fortunately, there are available modules with extremely light weight, which are specifically designed for circuit board mounting. The Ericsson PKF series, for example, weighs less than 20 grams. This offers significant advantage in terms of mechanical integrity of products constructed using this series of converters. It can be easily board mounted along with other components without any need for special mechanical retention. The solder joints (either pin-through-hole or SMD) provide all the mechanical support that is needed. The lower mass achieves much better board-level performance during vibration and shock, further enhancing the reliability of the product.

The industry trend is for SMD solutions to extend to higher levels of output power so that system manufacturing efficiency can be maximized. The recently introduced Ericsson PKD series, for example provides over 25 watts in a SMD package. This trend creates challenges for both the module manufacturer and for the system manufacturer. We will discuss issues that relate to both, such as module package designs, lead structures, module interconnection technologies, pin-in-paste attachment, building height and lead-free solder requirements.

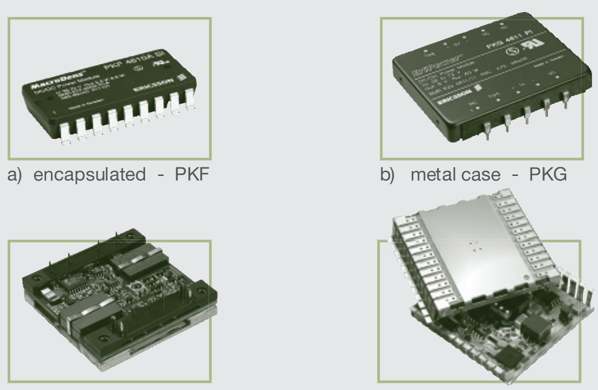

a) encapsulated - PKF

b) metal case - PKG

c) open frame with baseplate - PKL

d) open frame planar - PKD

figure 10.1

Power Module Considerations

We will first examine some of the mechanical design, packaging and manufacturing options available to the power module manufacturer. The most visible characteristic of a power module is the external package, which can take many forms in today's DC/DC converters. Ericsson employs several different approaches so that the design is optimized to the most common applications for each power module series. Figure 10.1 shows some examples of the types of packages used.

Many of the first generation DC/DC converters were encapsulated. That is, they were filled with a material (usually some form of thermally conductive epoxy) that surrounded the internal components and provided a path for heat to the external surfaces. This can be a useful technique when the application provides a cold plate for heat removal or for cases in which there is room for a heatsink to be attached to the converter. Encapsulation provides a more isothermal environment than the other approaches we will be describing. This can sometimes be a disadvantage if heat sensitive components such as optical isolators become heated by power dissipated from the power components. It also adds extra mass to the converter, which is a disadvantage in terms of automated assembly and mechanical integrity. Ericsson uses an overmolding technique in the low power PKF series of power modules. In this design, the thermal path is via conduction through the pins to the PCB.

A second approach is to not use encapsulation so that the space above the components is free air, but to provide an external case on the converter. This approach is used on several of Ericsson's designs including the PKA, PKC, PKE, PKG and PKN. With proper design, this results in lighter weight, lower cost and better thermal management than the encapsulated approach.

11 Reliability

Reliability and availability is one of the key selling points for electronic equipment.. This was always true in the telecom and military marketplaces, but it is now also a priority in the datacom, industrial and even consumer markets.

Introduction

Present day analog and digital ICs, when properly applied, are very reliable, and most equipment manufacturers use similar components from many of the same suppliers. Consequently, the power system for the equipment often is one of the few ways for the manufacturer to differentiate the product from those of its competitors. Because of their design complexity, component stress levels and variation in power dissipated, today's high power density converters can vary widely in reliability. Some of them will be the most unreliable part of a system. The better ones are capable of demonstrating truly astounding levels of reliability and availability. In this treatment of reliability, we will define some terms and concepts, show examples of how to use the concepts, and examine various ways of predicting and measuring power module reliability. We will also address the design and procurement practices available to the power system designer and see how they affect the overall reliability level of the product. Finally, we will consider the concept of power converter lifetime or wearout and see how it relates to failure rate and MTBF.

Definitions and Concepts

The most fundamental concept in electronic reliability theory is the so-called 'bathtub' curve that depicts how the failure rate of a component (or assembly of multiple components) varies as a function of time. This curve is shown in a general way in figure 11.1. The values on both ordinates will vary widely depending upon the types of components and systems that are being evaluated, and we will address some appropriate typical numbers for power modules and systems later in this chapter. For now we will discuss the effects shown on the curve in a qualitative way.

The area at the beginning of the curve is called the infant mortality period. This time period represents the first few hours (typically less than 200) of a component's or product's life. The failure rate is generally higher during this period due to fallout of manufacturing defects and marginal components. As we will see, there are often actions taken by manufacturers on the component, sub-assembly and product levels to prevent these failures from occurring after the product is in the customer's hands. The most common technique for achieving this end is a 'burn-in' process.

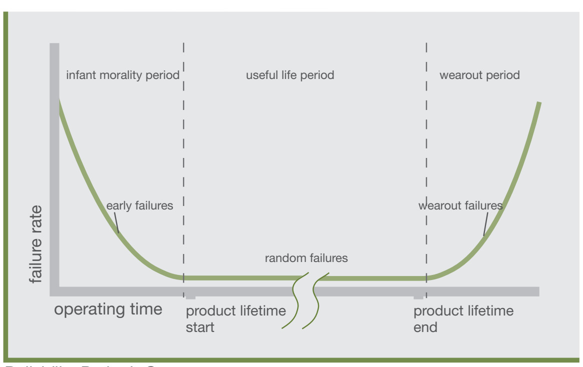

Reliability Bathtub Curve

figure 11.1

At the other end of the time spectrum is the area labeled "wearout period". After extremely long periods of time some electronic components begin to fail due to known long-term wearout mechanisms. Some examples are certain types of electrolytic capacitors, batteries and fan motors. For other components the long-term failure mechanisms either do not exist or are not yet known. In any case a good product design will have no known wearout mechanisms occurring before the expected product lifetime ends, which is normally determined by obsolescence or economic replacement intervals. With today's better power module offerings this goal can be easily achieved.

The period between the infant mortality period and the wearout period, referred to as the "useful life" period, is much more interesting and relevant. This period is of significantly greater duration than either of the other time intervals and is the time during which the equipment or component is expected to operate in a reliable fashion. As can be noted from figure 11.1, there are two important characteristics of the failure rate during the useful life period:

- The failure rate is constant.
- The failure rate is low.

This low and constant failure rate is due to random component failures, and is influenced by operational stress and temperature conditions. We will be discussing this type of failure rate in great detail since it is what determines the real-world reliability of all electronic products.

As the name implies, the failure rate has the units of failures per unit of time. We designate failure rate with the symbol λ. Therefore a low value of λ implies a low failure rate and a high reliability. λ is referred to as the "intrinsic failure rate", or IFR. This relationship is expressed mathematically by using the exponential reliability function that defines the probability of a component operating after time t as:

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

where: λ is the constant failure rate.

Using this expression, what is the probability of a capacitor operating after 300,000 hours if its failure rate is 5 × 10-7 failures per hour?

with $\lambda = 5 \times 10^{-7}$ and t = 3 × 105
$\lambda t = (5 \times 10^{-7})(3 \times 105) = 0.15$
$R(t) = e^{-0.15} = 0.861$
and the probability of no failure after 300,000 hours is 86%.

Similarly, we can calculate R(t) as a function of t with the result as shown in figure 11.2. Note that the probability of failure remains very low until the operating time becomes quite large (105 hours) and that the probability of failure becomes very large when the operating time is above 107 hours. If the component sample size is large, most of the units failing will fail within this period of time (see also section "Failure Rate, MTBF, and Lifetime").

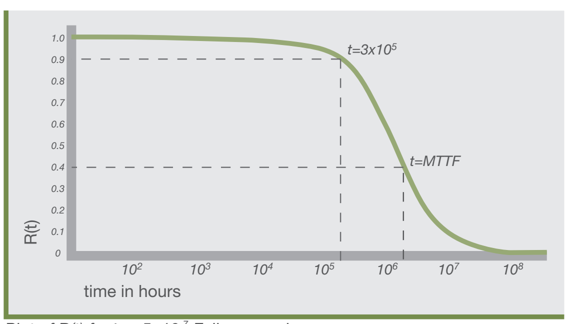

Plot of R(t) for $\lambda = 5\times10^{-7}$ Failures per hour

figure 11.2

There is sometimes confusion when using failure rates due to the lack of standardization of units for λ. If highly reliable components are being considered expressing λ in units of failures per hour results in very small numbers. Because of the awkwardness of using failures per hour, there are other units used in practice in various corporations. Some of the more commonly used units are:

- Failures per million hours.
- % failures per thousand hours.
- Failures per billion hours , referred to as "FIT", an acronym for "Failures in Time".

All of the above are equally valid ways to represent the IFR. With the increasing levels of reliability and the corresponding reduction in the size of λ for today's electronics it is recommended that the FIT terminology be used. This will result in numbers that are easier to record and manage. Failure rates expressed in any of these units can be converted to any other unit by using the conversion factors shown in figure 11.3.

Given failure rates for individual components, the failure rate of an assembly of these components is obtained simply by summing the individual failure rates. This approach

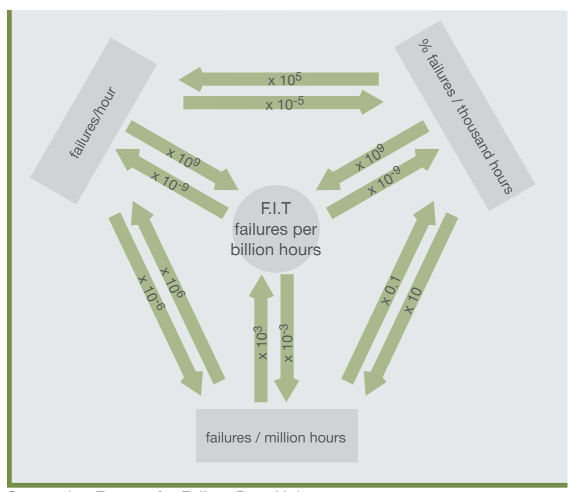

Conversion Factors for Failure Rate Units

figure 11.3

assumes that the failure of any component results in failure of the assembly. This assumption, in practice, is close enough to reality that the resulting mathematical simplicity is well worth the very slight difference in the reliability calculation. Also IFR calculations are based on many assump-tions about the components and their operating conditions and the results are accurate to at best ± 10%. The effect of the above assumption will normally be less than this.

We will now calculate the failure rate for an assembly composed of several components. The overall IFR is the summation of all the individual component failure rates. Normally there are many components with the same approximate failure rate (i.e. – all thick film resistors of a given power rating), so a commonly used expression to calculate the IFR of an assembly is:

$$IFR = \lambda_A = n_1 \lambda_1 + n_2 \lambda_2 + n_3 \lambda_3 + \ldots + n_i \lambda_i$$

where: λ_A = IFR of Assembly
n_i = Number of component type i
λ_i = IFR of component type i

Using this approach we will calculate the assembly IFR of a multi-component assembly as follows:

Type of Component	Number Used	Component IFR (FIT)	S IFR (FIT)
Resistors	10	100	1000
Capacitors	5	500	2500
Transformer	1	50	50
Power Transformer	2	1000	2000
Total			5550

The assembly failure rate is 5550 FIT or 5.55×10^{-6} Failures/h.

Even though we cannot predict exactly when each component will fail, if there are a large number of identical components operating under the same conditions we can predict the average or mean time that a component will operate before failing. In the case of figure 11.2, for example, most failures occur between 10^5 and 10^7 hours, and we would expect the average time to failure to be somewhere within this range. Solving the reliability function for the mean, we find that the mean time to failure occurs when t has a value of $1/\lambda$. The mean time to failure is often abbreviated as MTTF, and the units are hours to failure. MTTFs can be found by taking the inverse of the IFR, or λ. Some MTTFs from our previous examples are:

Example	λ	MTTF (hrs)
Capacitor	5×10^{-7}	2,000,000
Assembly	5.55×10^{-6}	180,180

Another term commonly used in reliability analysis is MTBF, or Mean Time Between Failures. MTBF is found by using the MTTF and adjusting the result for the length of time required to replace failing units and restore operation to the end product. Present day technology, with very low failure rates and efficient field service practices, normally negates the need to adjust for restoration time as it tends to be insignificant compared to the MTTF. Consequently, except for very special situations, we can define MTBF to be essentially equal to MTTF. Thus:

$$MTTF = 1 / \lambda \approx MTBF$$

Even though MTTF is mathematically the more correct terminology for items such as components and power modules that are not repaired in the field, MTBF is more commonly used than MTTF and we will sometimes use it here. Mostly, however, we will use IFRs, as they are easier to work with mathematically and tend to avoid the confusion that sometimes occurs between MTBF and product lifetime.

If we plot the same relationship as shown in figure 11.2, but this time on a linear scale centered on the MTTF, we obtain the result shown in figure 11.4. This shows us that, for

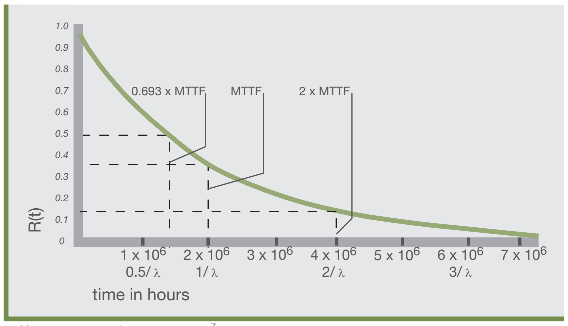

R(t) and MTTF for λ = 5×10^-7 Failures per hour

figure 11.4

a large sample size, 37% of the units will be operational after time MTTF, and that 50% of the units will be operational after time 0.69 × MTTF, if the IFR remains constant over this time.

Reliability Prediction and Measurement Techniques

We now know how to express and use failure rates and how to convert them to MTBFs. We have not, however, addressed the source of the IFR values we used in the preceding section. There are two main categories of IFR determination - prediction and measurement. Prediction techniques are used to estimate the IFR of an assembly before it can be tested. Measurement techniques are used when access to actual operational hardware data exists. In this section, we will discuss the most commonly used IFR prediction and measurement techniques:

- **Prediction**
  Military handbook
  Bellcore/Telcordia
  Supplier database
  Field history of previous designs
- **Measurement**
  Field history
  Accelerated life testing

As we will see later, measurement techniques are very time intensive, often requiring months or years of effort to arrive at meaningful results. Product development time-to-market considerations typically do not allow for extensive measurement of reliability during the development period. In spite of this, we must have a way to determine the reliability of new designs. To solve this problem reliability prediction techniques have been developed.

The most widely known and used prediction method is the so-called 'Mil-Handbook' approach. Because of the military's concern with the reliability of its electronic products, they developed a very detailed methodology for calculating predicted reliability. This methodology is documented in a handbook referred to as 'MIL-HDBK-217'. As of mid 2001, this document is at revision level F, notice 2. The MIL-HDBK-217 approach is basically identical to the technique we used in section "Definitions and Concepts" to derive the IFR for an assembly from the summed IFRs of its constituent components. The only difference is that the MIL-HDBK-217 technique modifies the base failure rate of each component in an attempt to account for the quality of the component and the operational conditions for the component. In general, each component failure rate is modified for the following parameters:

- Component quality and sourcing
- Operating environment
- Circuit density
- Temperature

The modified component failure rates are then summed to arrive at the predicted IFR for the entire assembly. This approach has merit as a concept, but in practice has several shortcomings that have limited its direct usage to military systems. Some of the problems with applying the MIL-HDBK-217 method to non-military systems are:

- The MIL-HDBK is dependent upon a database of component types and their corresponding failure rates. This database takes time to be generated and focuses on the component types developed for military systems. As a consequence most newer commercial technologies and components are not available in the database. Reliance on the MIL-HDBK for reliability prediction results in designs lacking innovation and no usage of newer components with better performance and reliability.

- The MIL-HDBK imposes a very harsh reliability penalty to non-military components. In reality, many of the better commercial components come down the same manufacturing line at the component vendor as the equivalent military part and are physically identical. The only difference is the extra documentation and extended testing that the military part is exposed to. Several independent tests have shown that, with intelligent procurement, commercial parts with reliability performance equal to military equivalents can be obtained.
- The MIL-HDBK assigns failure rates to some components that are not consistent with their actual performance. For example, transformers and magnetic devices have a very low actual failure rate but the MIL-HDBK IFR prediction is very high. An even better example is integrated circuits. We have all experienced the growth in reliability of systems configured with a smaller number of ICs replacing large quantities of discrete components. Yet the MIL-HDBK assigns very high IFR values to ICs, and penalizes integrated solutions. The result can be that designing to minimize MIL-HDBK IFR values can sometimes result in a less reliable product.

In cases where MIL-HDBK predictions have been compared with actual field performance, it has been found that for power converter assemblies the IFR predicted by MIL-HDBK techniques is in the range of 3 to 10 times higher than actual field failure rates.

Another commonly used prediction technique is based upon Bellcore reliability test methodology. Bellcore (Bell Communications Research) was a spin-off from AT&T and is the R&D organization of the Bell operating companies in the US. Bellcore now supports and continues development of a reliability prediction methodology originally developed by AT&T Bell Labs in the mid 1980s. Bellcore is now transitioning to the name "Telcordia" for at least a portion of their operations. The document that defines their reliability prediction technique is TR-332 "Reliability Prediction Procedure for Electronic Equipment". This document was at "issue 6" revision in mid 2001.

The Bellcore/Telcordia methodology is conceptually very similar to the MIL-HDBK approach, but is optimized for telecom systems. It includes component types used in telecom systems that are not covered by the MIL-HDBK system. It also includes provisions for the usage of burn-in and field reliability data and is therefore much more flexible in terms of incorporation of historical data in addition to theoretical predictions. Bellcore/Telcordia prediction methods are widely accepted in the US and are now gaining more acceptance worldwide. The failure rate prediction results using Bellcore/Telcordia techniques are typically lower (and therefore closer to actual field performance) than those resulting from MIL-HDBK procedures.

In an attempt to retain the very valid concepts and methodologies of the MIL-HDBK and Bellcore/Telcordia approaches, but remove the shortcomings of the component database, many power supply manufacturers utilize their own database of component IFR performance. These suppliers are usually large corporations that have the benefit of doing some of their own component development, have access to failure analysis data from field failures, or have accumulated significant experience with manufacturing very large quantities of power supplies. These ‘supplier databases’ contain up-to-date information on the types of components actually used by the supplier. Variation in IFR between component vendors is tracked. The resulting information is more accurate than the MIL-HDBK data and represents a very powerful and valuable competitive advantage for the power module supplier. Of the prediction methods presently available, this is perhaps the most accurate and consistent predictor of operational performance.

Another predictive approach that can be used by a power module manufacturer is to estimate the IFR of a power module based on the assembly level field history of similar products. This can be a useful and easy to apply technique if the new design is similar to an existing design. For example, if only the output voltage of a converter is changed, it is reasonable to assume that its failure rate will be very similar to other previous converters using the same topology and components. This technique is less useful for prediction of IFR performance of designs that are significantly different from past designs.

It should be noted that field history is not easy to obtain. A big percentage of failed units are sometimes discarded and replaced without notifying the power supply manufacturer, so that returns to the manufacturer are only a subset of the total field failures. The best field history information exists within vertically integrated companies that not only manufacture power converters but also use these converters within their own products. These companies can tap into the product repair records and be aware of all field incidents involving the power modules.

We will now address methodologies for measuring reliability as opposed to predicting it. By far the most accurate method of reliability measurement is analysis of actual field data. It measures large quantities of units under actual user conditions. Probably the most difficult aspect of collecting field history is to accurately assess the number of operating hours on each power module. For telecom and some datacom applications that operate close to continuously, the operating hours can be obtained fairly easily. For other types of equipment such as personal computers and consumer electronics, it is much more difficult to assess operational time. With good data for number of failures and number of operating hours, it is an easy calculation to obtain the overall IFR for the power module assembly. If the power modules are returned to the manufacturer for failure analysis, failures can be traced back to individual components and component level IFR data developed. Even though the field history method is the most accurate, it has one very big disadvantage – the data isn’t available for months or years after the product is released. It is of no use in developing new designs or in predicting IFRs for newly introduced equipment. To remedy this major shortcoming, the accelerated life testing method of reliability measurement was developed.

Life testing consists of accumulating a significant number of operating hours on a product in a reasonable length of time by operating multiple units at the same time. Since it is the IFR, or random failure rate, that is desired, it is important to use units that have passed the burn-in test and have survived past the infant mortality period. Life testing is a measurement technique since it uses a sample of actual product rather than just analytical calculations. For inexpensive products or assemblies with fairly large failure rates, practical life tests can be done without acceleration. For example, if a DC/DC converter has an IFR of 1×10^{-4} failures/h and a quantity of 100 of these converters are operated continuously for a period of 30 days (720 hours), the expected number of failures would be:

$$(100)\ (1 \times 10^{-4}\ \text{fails/h})\ (720\ \text{h}) = 7.2\ \text{Failures}$$

The actual IFR would then be calculated based on the actual number of failures occurring within the 30 day test period. Here is a case where an unaccelerated life test is fairly economical and time efficient. This is not the case for present day power module technology that can exhibit IFRs in the range of 300 FIT. Using this number, the expected number of failures for the above test would be:

$$(100)\ (300 \times 10^{-9}\ \text{fails/h})\ (720\ \text{h}) = 0.022\ \text{Failures}$$

This number is much too small to be useful for a valid life test. In fact, to obtain 10 expected failures, one of the following test plans would be required:

- 100 units for 333,333 hours (38.1 years)
- 46,296 units for 30 days

The first alternative is unacceptable in terms of time, and the second unacceptable from a cost and implementation point of view. It is because of these problems that the concept of accelerated life testing was developed. Accelerated life testing essentially simulates very long test times by means of increasing the stress applied to the units and using a more reasonable and practical test time.

The stress level is increased by operating the power modules at a temperature that is elevated relative to the normal operating ambient temperature. The greater the ratio of the test temperature to the operating temperature, the greater the acceleration factor. The acceleration factor also is affected by a variable called the activation energy. The activation energy varies from component type to component type because the failure mechanisms differ greatly, but for power converter life testing an average value in the range of 0.6 eV to 0.8 eV is commonly used. Temperature ratios and the activation energy is used in the Arrhenius equation to determine the temperature acceleration factor. It is tempting to use very large acceleration factors to reduce the length or sample size of life tests, but large acceleration factors increase the uncertainty of the resulting data. In order to obtain data that will be widely accepted as valid, it is best to limit the acceleration factor to 15 or less.

If the life test of the 300 FIT converter described above is redesigned as an accelerated life test with two expected failures, an acceleration factor of 10, and 200 units under test, the required test time would be in the order of 3000 to 4000 hours or 5 months. This is certainly a more reasonable test in terms of both time and expense than the available alternatives without acceleration.

Ericsson uses a combination of all the prediction and measurement techniques described above. For initial designs and estimates prior to production a MIL-HDBK type approach is used, but with an enhanced Ericsson database containing up-to-date information on the actual types of components used in the design. New designs are subjected to an accelerated life test at the beginning of production to verify predicted reliability. Field history is also tracked in great detail using the large quantity of operating hours accumulated within Ericsson telecom equipment. Power module failures that do occur are failure analyzed to gain knowledge of the component at fault. Ericsson external customers also provide feedback on reliability levels achieved using our power modules in other types of applications.

Failure Rate, MTBF, and Lifetime

As we have seen previously, MTBF is approximately the inverse of λ, the failure rate. A low failure rate implies a long MTBF. These are both valid statements, mathematically and in practice. We have used failure rate rather than MTBF in most of this text for two reasons. First, failure rates can be directly added when computing cumulative failure rates for assemblies with multiple parts, while MTBFs need to be converted to failure rates, added and then converted back again. Secondly, MTBF is a concept that sometimes gets misunderstood and consequently misused.

MTBF is sometimes confused with lifetime even though they are the result of completely different failure mechanisms. MTTF and MTBF are statistical measures determined by random failures that occur during the useful life period of the bathtub curve (figure 11.1). Component lifetime is driven by wearout mechanisms that are often very different in nature and occur in the wearout period of the bathtub curve. The misunderstanding centers around the very large MTBF values calculated from the IFRs of today's very reliable electronics.

Let us look at some numbers for power supplies. These can be considered typical reliability indicators for the best available power supplies for the respective years.

Year	Converter	IFR (FIT)	MTBF (h)	Telecom * MTBF (yr)	Office ** MTBF (yr)
1975	AC/DC Early Technology	50,000	20,000	2.3	9.6
1985	AC/DC Better Technology	10,000	100,000	11.4	48.1
2000	Board Mounted DC/DC Power Module	200	5,000,000	571	2,404

* Assumes Continuous Operation (8760 hours per year)

** Assumes 8 hours per day, 5 days per week, 52 weeks per year (2080 hours per year)

The above analysis also assumes that the power-off failure rate is zero, which is not strictly true.

As the IFR has improved with time the MTBF calculation has resulted in some very impressive numbers. The difference between MTBF and Lifetime can then be shown by the following example:

The reliability function, $R(t) = e^{-\lambda t}$, can be used to determine the probability of operation after t = MTBF

$$R(t{=}MTBF) = e^{-(MTBF/MTBF)} = e^{-1} = 0.368$$

This means that for a large population of the board mounted DC/DC converters, 36.8% will survive the first 571 years of operation if the IFR remains constant over this time. It does not mean that on average a DC/DC converter will operate for 571 years.

However, it would be unrealistic for the IFR to remain constant for that long. Components that exhibit wearout mechanisms determine the useful lifetime of the DC/DC converter and other power supplies. Some wearout mechanisms are:

- Long term chemical interactions between materials.
- Shock and vibration effects on components, wire bonds, solder joints.
- Temperature cycling effects such as stresses due to thermal coefficient of expansion mismatches.
- Long term penetration of moisture into device packages.

One component type whose lifetime has been suspect is electrolytic capacitors. Recent research has shown that the lifetime is highly dependent upon the type and quality level of the capacitor and the circuit application and the stress levels imposed on the capacitor. Specifically, solid tantalum capacitors have demonstrated service lives in excess of 20 years if properly derated and not exposed to high dV/dt stresses. The output filter of a DC/DC converter is an application that meets these requirements. The dV/dt applied to the capacitor is limited by the converter ramp-up circuit and also by the other elements in the output filter. Also, the converter limits the maximum energy that can be applied to the capacitor. If a good quality capacitor with a rated voltage of greater than 2 times the converter output is selected, extremely reliable and long life operation is achieved. It should be noted that rated voltage changes with temperature.

It is important to realize that the input filter of a DC/DC converter does not meet the above conditions. Very high values of dV/dt with extremely high energy levels can be imposed on the capacitor by transients on the input voltage as described in chapter 8. For this reason, the reliability and lifetime of DC/DC converters utilizing an electrolytic capacitor as part of the input filter should be questioned.

The probability of the above type of problems can be minimized and the projected converter lifetime maximized by selecting the proper converter. Below are some guidelines for making this selection.

Minimize

- Number of different materials used.
- Number of process chemicals used.

Maximize

- Usage of stable & inert materials.
- Silicon integration.
- Thermal performance of passive components by mechanical integration.
- Knowledge of components.
- Manufacturing process controls, including cleaning.

To summarize, maximum lifetime is achieved when a well understood and proven design is manufactured by a supplier who places emphasis on a c ontrolled and clean manufacturing process. Successful designs tend to be highly integrated and manufactured in a manner that resembles process lines in an integrated circuit facility. The better suppliers will be proud of what is inside their product and share with you the criteria they use for component vendor selection. They will also offer tours of their process lines so that you may see the care and controls that are included in the manufacturing process.

A way to approach the issue of lifetime is to compare the operating lifetime of components and sub-assemblies that are determined by wearout mechanisms to the useful lifetimes for products that are driven by economic and marketing considerations. Below are some typical product lifetimes:

Product Type	Lifetime (years)
Personal Computer	5
Mainframe Computer	10
Telecom System	20

Ericsson's DC/DC Power modules are designed for product lifetimes exceeding 20 years. This is verified through extensive long term component testing and DC/DC power module temperature cycling.

A realistic application of MTBF numbers for power supplies is in predicting the percentage of units that can be expected to survive to a given time. If the DC/DC converter in the previous example has a useful life of 20 years, the percentage of units that will survive the useful lifetime is calculated by:

$$R(20year) = e^{-(20/571)} = 0.966$$

Therefore, 1 - 0.966 = 3.4% of these units will fail before wearout occurs. Using the 200 FIT IFR, let us also calculate the expected number of failures per year experienced by the manufacturer of the office equipment assuming the following implementation:

Number of DC/DCs per product	3
Number of Products in field	10,000

Failures/yr =
$(200 \times 10^{-9}$ fails/h)(2080 h/yr)(3)(10,000) = 12.5

MTBF = 1/12.5 = 0.0801 yr = 29.2 days

For the 30,000 converters in the field, the manufacturer can expect failures at average intervals of 29 days, and should structure the field service strategy and spares stocking in accordance with this estimate.

Design Practices for Maximum Reliability

Power module manufacturers have made significant progress in enhancing the reliability of the available power modules, with some of today's offerings being capable of failure rates as low as 200 to 400 FIT. In order to achieve this kind of performance within the product application, the user of these modules should be aware of the external factors that influence the power module reliability. Intelligent design of the power system can contribute greatly to system reliability by insuring that the stresses imposed on the power modules are consistent with long life.

By far the single most important factor under the control of the power system designer is the temperature of the power module. Ericsson products are designed to operate without derating at high pin or case temperatures in order to satisfy the demand for easy to apply products that meet the needs of the telecom, industrial and datacom markets. When operated within the specified limits the modules will be safe and reliable, with all components below their maximum allowable operating temperatures. But as we saw in the section on accelerated life testing there is a definite relationship between operating temperature and expected reliability. The module user can take advantage of this fact and improve the system reliability to even greater levels by controlling the module temperature. The module reliability predicted by Ericsson is predicated on an average ambient temperature of +40 to 45 °C. In applications with prolonged exposure to higher ambient temperatures the IFR will be degraded. If the average ambient is less than +45 °C, the failure rate can be even lower than the projected datasheet values (see figure 11.5). There are two main ways of achieving this end – derating the power module and controlling the ambient temperature.

The thermal design of the power module and the failure rate prediction assume that it is operating at its maximum rated output power. If it is operating at less than full load, internal power dissipations will be less as will all the temperatures internal to the module. For the components inside the power module, reducing the load current will have the same effect as lowering the external ambient temperature – a lower temperature rise between ambient and the internal component temperature. Running

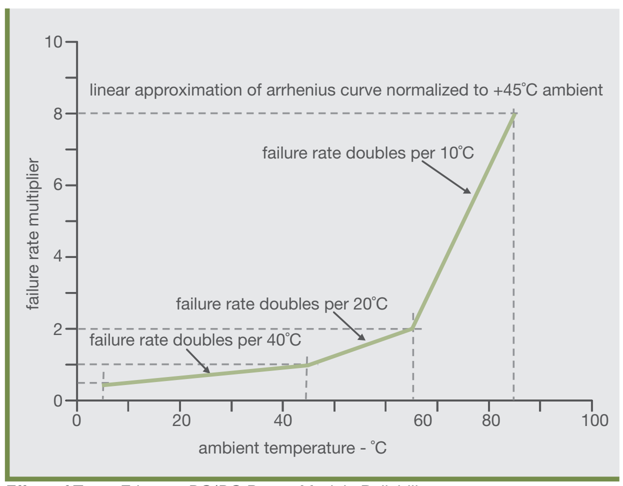

Effect of T_A on Ericsson DC/DC Power Module Reliability

figure 11.5

a power module at half load, for example, will reduce internal power dissipation by half and the internal temperature rise will be half of the full load value. This could result in component temperature decreases of up to 20 °C, and substantially improved reliability. Even derating to 70 or 80% can result in very meaningful and measurable reliability enhancement.

Control of ambient temperature is equally powerful as a tool for reliability design. The lower the case temperature (for convection cooled power modules) or the pin temperature (for conduction cooled power modules), the better the reliability. In either case, lowering of the system's internal ambient temperature will enhance the reliability. If this is not possible, some of the techniques discussed in the thermal design chapter will be helpful in lowering the temperature at the module. Some of the possibilities are heatsinking, forced convection cooling, careful placement of power modules in free convection systems, and thoughtful board layouts. The bottom line is:

- **Cooler Modules = Better Reliability**

The above hints will be very successful in controlling thermal stress to the power module. Electrical stresses should also be considered. The most common electrical stress that can degrade the power module reliability – or even cause immediate failure – is voltage transients on the input. Ericsson DC/DC power modules are designed to operate over a wide range of DC input voltage to allow for convenient application to telecom systems. In real-world systems, there are often transients on the DC input bus that can exceed the DC static voltage limits. This transient activity can be the result of load switching in other parts of the system, fault conditions in the system or external to the system, or external influences such as lightning strikes, etc. The inductive elements of the system DC distribution network can contribute to making these transients even worse. Ericsson DC/DC power modules contain circuitry to absorb a reasonable amount of this kind of transient energy, but the amount of energy that can be handled is limited as specified on the module datasheet. System designers can prevent problems in this area by understanding the DC distribution system. The system should be modeled with equivalent resistive, inductive, and capacitive elements, and the effects of load switching and other system effects studied. If required, the internal power module transient suppression capability can be supplemented with additional components external to the power module. This topic is covered in greater detail in the chapter on power system electrical design.

Procurement Practices for Maximum Reliability

We have seen above how the power system designer can improve the reliability performance of the system by how he/she applies power modules. Of equal importance to achieving the system reliability goals is selecting the proper power module and supplier. The design of the power module, the manufacturing process, and the supplier's testing process are all of critical importance and vary widely from supplier to supplier. The better manufacturers will have a detailed understanding of the reliability impacts of all aspects of their design and manufacturing environment and demonstrate a willingness to share this knowledge with their customers. Below are some of the criteria to explore when selecting a supplier for high reliability power modules:

- Experience with application of its power modules in actual high reliability systems.
- Access to long-term field performance data for tracking of reliability.
- Intensive knowledge of components, including database of actual failure rates as a function of stress for components used in their designs.

- Design integrity – proven designs and well conceived new designs.
- Conservative thermal design.
- High baseplate or case temperature rating.
- High efficiency.
- Accelerated life testing of new designs.
- Care in component vendor selection.
- No electrolytic capacitors in input filter designs.
- Simple designs with high levels of integration and low component count.
- Hybrid-like packaging.
- Thick film resistors rather than discrete.
- High degree of automation in manufacturing process.
- Tight manufacturing process controls.
- Emphasis on process cleanliness.
- 100% burn-in to deliver units with lowest possible failure rates.
- Use of environmental stress testing to increase assurance of reliable long-term performance when exposed to temperature, humidity, shock, vibration, temperature cycling, power cycling, etc.

Examples of Power System Reliability Calculations

Now that we have explored some of the concepts of power system reliability, we will apply them to some examples, showing how the reliability assessments will help in making system architecture and logistics decisions.

Example 1 - Expected failures and spares strategy

A system is configured with distributed power modules on a power per function basis (each power module is mounted on its own PCB and supplies power to adjacent boards). There are a total of 12 identical power boards in the system, and each board assembly has a failure rate, l, of 500 FIT. The system is a telecom product and operates continuously. What is the expected number of power board failures per year for each system? How many spare boards should be stocked?

With continuous operation, there are 8760 operating hours per year. For the 12 units, the expected number of failures per year is given by:

$$(12)\ (500 \times 10^{-9}\ \text{fails/h})\ (8760\ \text{h}) = 0.053\ \text{Failures}$$

We find that the probability of a failure within the one year period is very small. Stocking one spare power board assembly should be sufficient.

Example 2 - Effect of power module failure rate on repair actions

We saw in the previous example how highly reliable power modules can result in very reliable power systems with minimal exposure to field replacements. We will now expand upon this example. Assume that the product manufacturer has a total of 2000 of these systems in the field under warranty. What are the total yearly module replacements as a function of module reliability? Assume a range of reliability consistent with presently available power modules (10,000 FIT to 200 FIT). For the 500 FIT unit, we obtain:

$$(2000)\ (12)\ (500 \times 10^{-9}\ \text{fails/h})\ (8760\ \text{h}) = 105\ \text{units}$$

Repeating this calculation for other values of λ we obtain the following result:

Power Module Failure Rate - l (FIT)	Yearly Replacements
10,000	2,102
5,000	1,051
2,000	420
1,000	210
500	105
200	42

The more reliable power modules offer a very significant reduction in yearly replacements that will drastically decrease the costs for replacements and labor. Further, it will simplify the field service strategy and result in a much better reliability reputation for the manufacturer's products.

Example 3 - Comparison of power architectures

Consider a rack and shelf telecom system, where the designer can either use a 'power per shelf' concept or a 'power per board' approach. The shelf DC/DC would be a 100 W unit with a failure rate, lPS, of 1500 FIT. For the board mounted version, each board would contain a 25 W output DC/DC power module with a lPM of 500 FIT. See figure 11.6 for a sketch of each power architecture implementation. Each of the functions is identical and serves the end users. The electronics of each function has a lFU of 2500 FIT. It is desired that this be a high availability system, with the maximum number of functions operational at all times. Compare the two power architectures in terms of the total reliability per function.

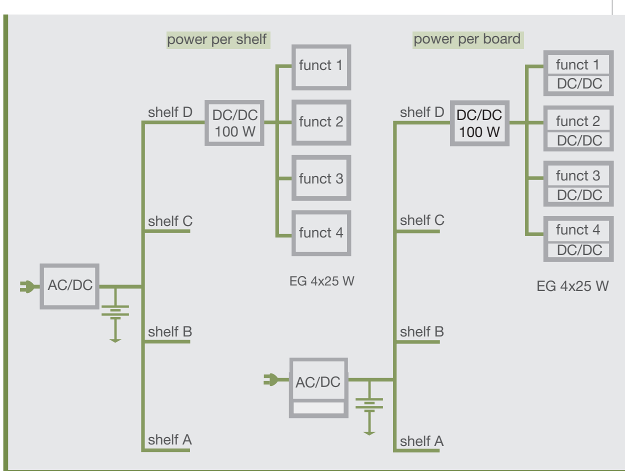

Comparison of Power Architectures

figure 11.6

Since the front ends of both power systems (AC/DC and battery) are identical and have the same failure rates we can ignore the front ends and focus on the DC/DC conversion alternatives and the function boards. For the power per board approach, the total failure rate for the function is the sum of lFU and lPM, or 3000 FIT. For the power per shelf alternative, failure of either the DC/DC or the function board electronics results in failure of the function. Thus the total failure rate of any given function is the sum of lFU and lPS or 4000 FIT. We see that the reliability per function is higher with the board mounted power alternative.

There are additional considerations. In the power per shelf design, failure of the DC/DC converter results in loss of 4 functional units vs. only one unit for the power per board design. Furthermore, any power per board unit could be exchanged without interrupting the operation of the other functions. Thus, the availability advantages of the fully distributed approach are greater than the 1.33 times improvement in the individual functional element reliability.

A recent trend in DC/DC converter design is to use some form of "open frame" design with no external case. These designs may or may not have a baseplate structure for heat spreading and heatsink attachment. The higher power offerings often include a baseplate for applications that need a heatsink or have low airflow environments. These designs are still considered "open frame" if the circuitry on the internal PCB is exposed to the ambient airflow. This approach tends to keep the control circuitry at a much lower temperature than the power components. Ericsson uses this approach in the PKJ, PKL and PKM power modules.

A planar open frame power module package contains no baseplate and depends upon conduction through the pins and the direct exposure of components to ambient airflow for cooling. This design will have a lower power density than heatsink compatible designs, but offers several advantages. It is the simplest and lowest cost package. It is very light in weight and can provide the lowest height above the PCB for applications with low board-to-board pitch. The Ericsson PKD series power modules are constructed with a planar open frame package.

While not as visible as the external construction, the packaging technology used internal to the module is equally important. Many DC/DC converter manufacturers use conventional FR4 printed circuit boards with component selection and placement not optimized for efficient heat removal. The internal construction of Ericsson power modules varies from series to series, but uses some very high technology techniques for the enhancement of thermal design and module reliability. Many of these products use ceramic substrates as an interconnection platform in conjunction with highly reliable screened thick film resistors. Direct bonded copper is used for interconnection in some cases. A recent Ericsson offering, the PKD series, is compliant with the lead-free solder initiatives that are now becoming more commonplace as a design requirement. The PKD modules themselves do not contain lead and are designed to withstand a soldering profile appropriate for the usage of non-lead solders as they are assembled by the OEM.

System Manufacturing Considerations

We will now turn our attention to manufacturing considerations encountered by the end user of the power module. Today's DC/DC converters are connected to the user's PCB by either pin-in-hole or SMD attachment, and both approaches require attention to detail if an efficient manufacturing process is desired.

Most pinned Ericsson DC/DC power modules can be mounted either with sockets or directly soldered into a circuit board. The solder method is recommended and preferred in most applications. It provides a higher reliability lower inductance connection, as there is no mechanical interface through the socket. Due to contact oxidation, sockets can be a reliability problem, especially over long periods of time without insertion activity. It is not uncommon for the socket to have a higher failure rate than the converter itself. The socket also adds cost. In addition to its actual cost, the socket approach will require a hand insertion operation to install the power supply. This labor cost can be significant in large volume applications.

There is sometimes reluctance to solder the power supply rather than using a socket. This position is almost always related to the historical assumption that the power converter will need to be replaced more often than the rest of the circuitry. With decentralized power architectures and reliable power modules, this is no longer reality. When board mounted power modules are used in a power per board architecture, the field maintenance strategy is to replace the entire board in the event of failure of any part of the board. The failure is most likely to be other than a power module fault, but the low cost and high reliability of today's power modules allow for this to be a very successful replacement strategy. There is no more need to socket the power converter than to socket most of the other components.

Ericsson DC/DC power modules are designed and manufactured to allow them to be handled and treated as components. The pins are designed to be solderable with a standard process. Compare this with the pin size and thermal mass encountered in some other DC/DC converters. Some of them will require very specialized manual operations, with negative impact upon cost and reliability. Ericsson DC/DC power modules are also fully compatible with the normal chemicals, solvents, washing operations and other process environments of standard circuit board process lines.

SMD manufacturing is now becoming the norm, and there is considerable pressure to extend the range of SMD compatible DC/DC converters to higher levels of output power. The incentive is the usage of a standard automated assembly process that does not involve costly manual operations or specialized process modifications. Low mass board mounted power modules are now available with SMD pinning so that they can be installed using normal SMD pick and place equipment. The circuit board can then be put through the same solder process as the other SMD components. This elimination of the need for special processing for the power module is a significant advancement in the concept of power becoming a component. Ericsson's PKD and PKL series of power modules are fully SMD compatible, including high temperature lead-free solder attachment in the case of the PKD.

Making a reliable line of SMD DC/DC converters requires much more than terminating the devices with a SMD header. The entire module design is affected. The converter must be capable of being reliably attached using a standard SMD soldering profile while not thermally overstressing any internal component during the solder operation. Ericsson has done considerable work in this area and supports its SMD power modules with solder profile curves and user manufacturing engineering support. Years of experience in high quality telecom SMD manufacturing in its own facilities has resulted in a unique and comprehensive knowledge base that benefits the users of Ericsson's power modules. For example, coplanarity of the SMD leads is an important manufacturing consideration, especially for large footprint components such as power modules. Ericsson places special emphasis on this aspect of the design, and delivers power modules meeting tight coplanarity specifications.

The popularity of SMD assembly processes has created an incentive for development of techniques for utilizing it with pinned components. The Ericsson PKG series power modules, for example, are small and light enough for compatibility with automated assembly equipment, but are not available with SMD pinouts. Ericsson has developed a technique and qualified the pinned PKG modules so that they can be assembled to the PCB with a standard SMD process. Ericsson refers to this assembly option as a "pin in paste" process. Other names used for the same technique include paste in pin, paste in hole reflow soldering, reflow of through hole, single center reflow soldering and intrusive reflow soldering. The first common usage of this technique was for the assembly of pinned connectors on a standard SMD assembly line, eliminating the wave soldering operation.

While process details vary from user to user, the basic idea is to apply normal SMD solder paste to the holes in the PCB using standard stencil technology. The PKG module is then inserted into the paste filled holes. The populated board is then passed through an oven on the SMD conveyer, reflowing the PKG solder connections. All PKG modules produced since the fourth quarter of 2001 are compliant with the pin in paste, along with some semi-standard parts previous to that date. The maximum allowable PKG pin temperature during the process is 210 ºC, with a recommended typical value of 200 ºC. Temperature and process profiles are available upon request from Ericsson.

Most systems that are implemented with decentralized power architectures use several circuit boards in a rack type configuration. This approach places boards in a vertically or horizontally stacked arrangement with board-to-board spacing (sometimes referred to as board pitch) of as little as 20 mm. In order to accommodate designs such as this, the height of the module above the board (sometimes referred to as building height) becomes very important. If allowances are made for tolerances, pin extension through the board, board thickness, and safety spacing, the module building height must be as small as 12 mm in many cases. Figure 10.2 depicts some of the considerations that determine the allowable module building height for pinned converters. SMD converters offer the advantage of no pin protrusion through the board and, in the case of planar open frame designs, lower module height. As a consequence, these newer designs are more flexible when applied to systems with small board-to-board pitch.

In summary, the best power module design and package is one that is optimized to the user's building practice. A high power DC/DC converter designed to be used in a power per shelf architecture with an attached heatsink

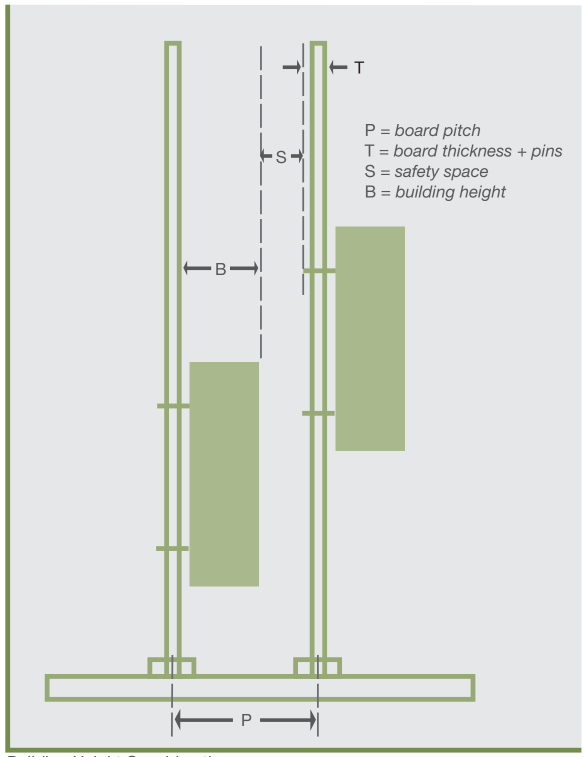

Building Height Considerations

figure 10.2

will be packaged considerably differently than a low power SMD converter used in a power per board application. Products should be available to allow compatibility with either through hole or SMD assembly processes. Low building height is becoming very critical for many applications. Ericsson's design philosophy is to provide as many of these criteria as possible so that the power system designer can enjoy the flexibility and freedom that comes with such a family of power conversion products.

12 Cost Analysis

Cost is one of the most important issues in equipment design. It is invariably one of the first three criteria on the list of priorities when selecting a power supplier. In spite of this very valid focus on cost there has been very little information published in the industry on cost analysis techniques for power systems. In this chapter we introduce methodologies for making informed decisions when comparing the costs of alternative power system implementations.

Introduction

We will first address the cost implications of reliability. For the past several years reliability has been identified as a very important issue in power system design. Most of the focus has been on issues such as system availability and customer satisfaction, with little attention on how reliability affects the bottom line costs of the system manufacturer. We will explore how power module reliability affects total lifetime system costs and find that the reliability performance can have a profound effect on total system cost for systems of even moderate size.

Next, we will look at the costs of power system development. Often power system cost analysis only compares the power supply procurement cost. There are many other significant costs involved during the development process. This section attempts to quantify these costs for a typical system and show how these non-recurring costs can have a significant impact on decisions relative to the power system implementation.

We will conclude the chapter with an overview of how the information presented here can be combined in an overall cost analysis for the power system and how this analysis can assist in making decisions about the power system strategy. It should be noted that our intent here is to suggest analysis methodologies. In order to accomplish this and show examples we needed to make assumptions about several parameters that are in the domain of the system developer. The reader will have better knowledge about these parameters than we do. Our intent is not to defend the validity of the numbers used in the analysis examples but rather to encourage each reader to do a similar analysis for their system that contains numbers they are comfortable with. Readily available spreadsheet software allows this kind of analysis to be done easily and cost effectively. In addition, Ericsson has developed cost analysis tools that are available from either a CD or from the Ericsson website.

Reliability Costs

Decentralized power systems have become increasingly common in the past few years and are now supported by availability of DC/DC power modules from several suppliers. Since their inception these power modules have seen dramatic improvements in efficiency, power density, thermal management and reliability. Some manufacturers are now producing third and fourth generation designs that are highly integrated and produced with automated manufacturing technologies. These converters achieve extremely high reliability performance. The Ericsson PKF series of DC/DC power modules, for example, have a predicted MTBF (Mean Time Between Failure) of 4.9 Million hours or over 550 years. The first reaction of some users to these types of ratings is to consider them to be either meaningless or "overkill", since their particular system is designed to have a lifetime of perhaps 5 or 10 years. We will show that reliability ratings in this range are indeed meaningful and provide substantial cost benefits for all types of products, even low cost, short lifetime products and systems. Since the enhanced reliability is achieved largely with higher integration level, lower component count, and less manual manufacturing intervention, it can be done without increasing the initial cost of the power modules. Indeed, some power modules with the highest actual reliability are also some of the most cost effective from an initial procurement point of view. In other cases, a higher initial procurement cost can be more than offset by the significant savings over the life of the product due to fewer repair actions.

After reviewing how reliability is associated with failure rate and end product lifetime, we will look at reliability ratings of some actual power modules. We can then apply these specifications to a typical system configuration and explore the costs over the life of the end product. As will be seen, the cost benefits of reliability are real, significant, and predictable.

As we saw in chapter 10, the most fundamental expression of reliability is failure rate. This is the expected rate at which a component, power module, sub-assembly or system will fail in terms of failures per unit of time. The failure rate is considered to be constant with time over the useful life of the end product as indicated in the flat portion of the "bathtub curve" (see figure 11.1). The initial "infant mortality" failures should be removed by means of component and power module burn-in and the "wearout" failures are not seen in most systems if proper component selection and design practices are utilized in the power module. The failure rate is most commonly expressed in terms of failures per hour, failures per million hours, or failures per billion hours. The latter measure is sometimes referred to as FIT. For example, 100 failures per billion hours = 100 FIT. It is often used because it results in easily managed numbers for the levels of reliability achieved by today's DC/DC power modules.

MTBF is the reciprocal of failure rate and is expressed in units of time, most commonly hours or years. One FIT equates to an MTBF of 109 hours or 114,155 years. Using failure rates more typical for DC/DC converters gives numbers in the range shown in figure 12.1.

FIT	MTBF - Hours	MTBF - Years
100	10,000,000	1,142
1,000	1,000,000	114
10,000	100,000	11
100,000	10,000	1 (approximately)

Reliability Range for DC/DC Converters

figure 12.1

The time equal to MTBF corresponds to the mean value of the reliability function:

$$R(t) = e^{-t/MTBF}$$

If t = MTBF, then:

$$R(t) = e^{-1} = 0.368$$

This means that in a population of DC/DC converters with a failure rate of 1000 FIT, 36.8% would survive the first 114 years of operation if the failure rate were constant over this period. In reality, however, the failure rate would not remain constant for that long because components will exhibit wearout mechanisms before 114 years and determine the useful lifetime for the end product. A MTBF of 114 years does not mean that an individual power module will operate for that long. A high MTBF means that a population of DC/DC converters will have good reliability, i.e. low probability of failure, during the useful lifetime.

For example, if a system manufacturer is using converters with a failure rate of 1000 FIT and has 10,000 such DC/DC converters in the field, the expected number of failures per year can be calculated as:

$$\text{Fails per year} = \frac{10{,}000 \times 8760 \text{ hours/yr}}{1{,}000{,}000 \text{ hours/fail}} = 87.6$$

About 88 failures per year can be expected using DC/DC converters with an actual MTBF of 114 years. Consequently, DC/DC converter reliability is extremely important even for systems with lifetimes much shorter than the DC/DC converter's MTBF.

When comparing the reliability ratings of DC/DC converters it is important to note the assumptions and conditions of the specification. Reliability is strongly related to operating temperature as shown in figure 11.5, with failure rate doubling with each 20 °C increase at moderate ambient temperatures. In a typical system the normal ambient air temperature in the vicinity of the DC/DC converter is usually in the area of 45 °C due to heating from other system elements within the equipment enclosure. The baseplate or case temperature is typically elevated another 30 °C to about 75 °C due to the thermal impedance between case and ambient. Ericsson specifies the reliability of its power modules at these temperatures so that the specification will be meaningful under typical system conditions. Other manufacturers specify their products at various lower temperatures, either case or ambient. These specifications can be converted to an equivalent 45 °C ambient or 75 °C case rating by multiplying their MTBF rating by a derating factor of 0.5 for each 20 °C between the rated and target temperatures.

Figure 12.2 summarizes reliability data for currently available DC/DC power modules in the 25 to 60 W range. They are all from major suppliers and are intended for convection cooled systems.

The adjusted MTBF figures will be used for the remainder of this study. They reflect predicted reliability at either 45 °C ambient or 75 °C case temperatures.

Analysis Methodology and Assumptions

It is convenient to use spreadsheet software to examine the effects of reliability over the life cycle of the products containing the DC/DC power modules. There are two main sources of reliability cost other than the initial procurement of the power modules - replacement cost for the failed units themselves, and the costs of doing the diagnosis, repair, rework, customer downtime, and customer satisfaction. Customer downtime is a real cost to the end user (his system is not available for productive work), and customer satisfaction is a real cost to the system provider (the customer will go elsewhere next time). In order to focus upon the reliability related costs it is initially assumed that the procurement cost of the DC/DC power module from each supplier is the same. The only variables are the actual adjusted reliability specifications and warranty terms. Later we will examine the effects of procurement cost. Since many of the costs occur several years in the future, the time value of money needs to be taken into account in order to get the results in terms of today's dollars. This was done by reflecting future costs back to the present by means of a net present value (NPV) calculation that includes an assumed rate of return (interest rate).

We will use an Ericsson PKG power module for our example, and compare its reliability cost with the other competitors shown in figure 12.2. Several assumptions needed to be made to show the expected results for a typical system. They were:

- 10 DC/DC power modules per system.
- Production of 5000 systems per yea for a 5 year period.
- 5 additional years of field support after end of production.
- All systems operate 24 hours per day.
- Procurement cost of replacement DC/DC power modules is $50.
- Cost of repair is $250 (includes travel, diagnosis, repair time, rework, customer downtime, and loss of customer satisfaction).
- Time value of money (rate of return) is 7% per year.

Consequently, the total product life cycle is 10 years. The results of the spreadsheet calculations are shown in figures 12.3 through 13.6.

Summary of Results

The most instructive way to compare the effects of reliability on total cost is to focus on the NPV of each alternative studied, as shown in

	SPEC. MTBF MHR	RATING CONDITIONS	DERATING FACTOR	ADJ. MTBF MHR	WARRANTY YEARS
Supplier A	0.33	25 °C Amb.	0.50	0.17	2
Supplier B	1.0	35 °C Case	0.25	0.25	2
Supplier C	2.6	40 °C Case	0.33	0.87	3
Ericsson PKG	2.0	45 °C Amb.	1.00	2.00	5

Reliability of DC/DC Converters

figure 12.2

figure 12.7. As can be seen, even the most reliable competitor will result in over 2 million dollars of additional cost over the life of the system. The least reliable competitor incurs a cost penalty of over 17 million dollars!

Effect of Warranty Period

An additional analysis was done to determine the effect of the warranty period on the total cost. The warranty period for the Ericsson PKG was changed from the actual 5 year value to a period of 2 years for the system example. As can be seen from figure 12.8, the NPV of the reliability associated costs increased by over $100,000. Warranty period is not as important as reliability and how it is specified, but better warranty terms can result in significant savings to the end user.

Offset of Initial Procurement Costs

To avoid the need to make assumptions about initial procurement costs and to simplify the analyses, most of this study assumed that all four suppliers had the same initial costs. We will now examine the case where the initial costs are not the same and the higher reliability power module has a higher procurement cost. The system designer needs to know how much of a "cost premium" during initial procurement is allowable in exchange for this additional reliability. This can be easily determined by using the spreadsheet's "goal seeking" function

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
New DC/DCs	50 000	50 000	50 000	50 000	50 000	0	0	0	0	0	250 000
Total in Field	50 000	100 000	150 000	200 000	250 000	250 000	250 000	250 000	250 000	250 000	250 000
Expected Fails	2576	5 153	7 729	10 306	12 882	12 882	12 882	12 882	12 882	12 882	103 056
Non-Warranty Fails	0	0	2 576	5 153	7729	10 306	12 882	12 882	12 882	12 882	77 292
Cost -Replace DC/DCs	$ -	$ -	$ 128 800	$ 257 650	$ 386 450	$ 515 300	$ 644 100	$ 644 100	$ 644 100	$ 644 100	$ 3 864 600
Cost - Repair	$ 644 000	$ 1 288 250	$ 1 932 250	$ 2 576 500	$ 3 220 500	$ 3 220 500	$ 3 220 500	$ 3 220 500	$ 3 200 500	$ 3 220 500	$ 25 764 000
Cost - Total	$ 644 000	$ 1 288 250	$ 2 061 050	$ 2 834 150	$ 3 606 950	$ 3 735 800	$ 3 864 600	$ 3 864 600	$ 3 864 600	$ 3 864 600	$ 29 628 600
NPV @ 7%	$ 19 355 258										

Lifetime Cost Numbers for Supplier A

figure 12.3

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
New DC/DCs	50 000	50 000	50 000	50 000	50 000	0	0	0	0	0	250 000
Total in Field	50 000	100 000	150 000	200 000	250 000	250 000	250 000	250 000	250 000	250 000	250 000
Expected Fails	1 752	3 504	5 256	7 008	8 760	8 760	8 760	8 760	8 760	8 760	70 080
Non-Warranty Fails	0	0	1 752	3 504	5 256	7 008	8 760	8 760	8 760	8 760	52 560
Cost -Replace DC/DCs	$ -	$ -	$ 87 600	$ 175 200	$ 262 800	$ 350 400	$ 438 000	$ 438 000	$ 438 000	$ 438 000	$ 2 628 000
Cost - Repair	$ 438 000	$ 876 000	$ 1 314 000	$ 1 752 000	$ 2 190 000	$ 2 190 000	$ 2 190 000	$ 2 190 000	$ 2 190 000	$ 2 190 000	$ 17 520 000
Cost - Total	$ 438 000	$ 876 000	$1 401 600	$ 1 927 200	$ 2 452 800	$ 2 540 400	$ 2 628 000	$ 2 628 000	$ 2 628 000	$2 628 000	$20 148 000
NPV @ 7%	$ 13 161 947										

Lifetime Cost Numbers for Supplier B

figure 12.4

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
NewDC/DCs	50 000	50 000	50 000	50 000	50 000	0	0	0	0	0	250 000
TotalinField	50 000	100 000	150 000	200 000	250 000	250 000	250 000	250 000	250 000	250 000	250 000
ExpectedFails	503	1 007	1 510	2 014	2 517	2 517	2 517	2 517	2 517	2 517	20 136
Non-WarrantyFails	0	0	0	503	1 007	1 510	2 014	2 517	2 517	2 517	12 585
Cost-ReplaceDC/DCs	$-	$-	$-	$ 25 150	$ 50 350	$ 75 500	$ 100 700	$ 125 850	$ 125 850	$ 125 850	$ 629 250
Cost-Repair	$ 125 750	$ 251 750	$ 377 500	$ 503 500	$ 629 250	$ 629 250	$ 629 250	$ 629 250	$ 629 250	$ 629 250	$ 5 034 000
Cost-Total	$ 125 750	$ 251 750	$ 377 500	$ 528 650	$ 679 600	$ 704 750	$ 729 950	$ 755 100	$ 755 100	$ 755 100	$ 5 663 250
NPV@7%	$ 3 691 649										

Lifetime Cost Numbers for Supplier C

figure12.5

to vary the procurement cost of the power module until the NPV of the higher reliability power module is equal to that of the other supplier.

Such an analysis was done for the previous system example. The initial procurement cost of the power modules was added to the spreadsheet model, and was assumed to be $50 for the competitive converters. The Ericsson PKG (the power module with the highest effective rated reliability) was then compared with each of the other suppliers. The procurement cost of the Ericsson PKG was varied until the overall NPV of the Ericsson PKG was equal to that of the other supplier. This result is the allowable procurement cost for the Ericsson DC/DC power module that results in the same overall total life cycle cost. The detailed results are shown in figures 12.9 to 12.11, and summarized in figures 12.12. For equal costs over the life of the product, the Ericsson PKG could support a procurement cost premium of from $10 to $86 each!

Conclusions

Although this study made several assumptions that may not apply to your application it hopefully demonstrates the ease with which a reliability cost study can be done. It can easily be modified or extended to accommodate the system or product with which you are working. We have shown that:

- Reliability has a dramatic effect on life cycle costs of the product.
- Even power modules with seemingly high reliability can adversely affect your system cost.
- The way in which the power module reliability is specified is extremely important, especially the assumed case or ambient temperature.

Dc/Dc Source	Reliability Affected Costs	Extra Cost Vs. Supplier D
Supplier A	$ 19 355 258	$ 17 820 131
Supplier B	$ 13 161 947	$ 11 626 820
Supplier C	$ 3 691 649	$ 2 156 522
Ericsson PKG	$ 1 535 127	

Summary of Results

figure 12.7

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
NewDC/DCs	50 000	50 000	50 000	50 000	50 000	0	0	0	0	0	250 000
TotalinField	50 000	100 000	150 000	200 000	250 000	250 000	250 000	250 000	250 000	250 000	250 000
ExpectedFails	219	438	657	876	1 095	1 095	1 095	1 095	1 095	1 095	8 760
Non-WarrantyFails	0	0	0	0	0	219	438	657	876	1 095	3 285
Cost-ReplaceDC/DCs	$-	$-	$-	$-	$-	$ 10 950	$ 21 900	$ 32 850	$ 43 800	$ 54 750	$ 164 250
Cost-Repair	$ 54 750	$ 109 500	$ 164 250	$ 219 000	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 2 190 000
Cost-Total	$ 54 750	$ 109 500	$ 164 250	$ 219 000	$ 273 750	$ 284 700	$ 295 650	$ 306 600	$ 317 550	$ 328 500	$ 2 354 250
NPV@7%	$ 1 535 127										

Lifetime Cost Numbers for Ericsson PKG

figure12.6

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
NewDC/DCs	50 000	50 000	50 000	50 000	50 000	0	0	0	0	0	250 000
TotalinField	50 000	100 000	150 000	200 000	250 000	250 000	250 000	250 000	250 000	250 000	250 000
ExpectedFails	219	438	657	876	1 095	1 095	1095	1 095	1 095	1 095	8 760
Non-WarrantyFails	0	0	219	438	657	876	1095	1 095	1 095	1 095	6 570
Cost-ReplaceDC/DCs	$-	$-	$ 10 950	$ 21 900	$ 32 850	$ 43 800	$ 54 750	$ 54 750	$ 54 750	$ 54 750	$ 328 500
Cost-Repair	$ 54 750	$ 109 500	$ 164 250	$ 219 000	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 2 190 000
Cost-Total	$ 54 750	$ 109 500	$ 175 200	$ 240 900	$ 306 600	$ 317 550	$ 328 500	$ 328 500	$ 328 500	$ 328 500	$ 2 518 500
NPV@7%	$ 1 645 243										

Effects of 2 Year Warranty on Ericsson PKG

figure12.8

	Supplier A	Ericsson PKG
MTBF-MHr	0.17	2
Warranty-Years	2	5
Cost-New DC/DCs	$ 50	$ 136.15
Cost-Repair	$ 250	$ 250
Interest Rate-%	7	7

Supplier A

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
New DC/DCs	50 000	50 000	50 000	50 000	50 000	0	0	0	0	0	250 000
Total in Field	50 000	100 000	150 000	200 000	250 000	250 000	250 000	250 000	250 000	250 000	250 000
Expected Fails	2 576	5 153	7 729	10 306	12 882	12 882	12 882	12 882	12 882	12 882	103 056
Non Warranty Fails	0	0	2 576	5 153	7 729	10 306	12 882	12 882	12 882	12 882	77 292
Cost-Replace DC/DCs	$-	$-	$ 128 800	$ 257 650	$ 386 450	$ 515 300	$ 644 100	$ 644 100	$ 644 100	$ 644 100	$ 3 864 600
Cost-Repair	$ 644 000	$ 1 288 250	$ 1 932 250	$ 2 576 500	$ 3 220 500	$ 3 220 500	$ 3 220 500	$ 3 220 500	$ 3 220 500	$ 3 220 500	$ 25 764 000
Cost-Procurement	$ 2 500 000	$ 2 500 000	$ 2 500 000	$ 2 500 000	$ 2 500 000	$-	$-	$-	$-	$-	$ 12 500 000
Cost-Total	$ 3 144 000	$ 3 788 250	$ 4 561 050	$ 5 334 150	$ 6 106 950	$ 3 735 800	$ 3 864 600	$ 3 864 600	$ 3 864 600	$ 3 864 600	$ 42 128 600
NPV@7%	$ 29 605 751										

Ericsson PKG

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
New DC/DCs	50 000	50 000	50 000	50 000	50 000	0	0	0	0	0	250 000
Total in Field	50 000	100 000	150 000	200 000	250 000	250 000	250 000	250 000	250 000	250 000	250 000
Expected Fails	219	438	657	876	1 095	1 095	1 095	1 095	1 095	1 095	8 760
Non Warranty Fails	0	0	0	0	0	219	438	657	876	1 095	3 285
Cost-Replace DC/DCs	$-	$-	$-	$-	$-	$ 29 817	$ 59 635	$ 89 452	$ 119 270	$ 149 087	$ 164 250
Cost-Repair	$ 54 750	$ 109 500	$ 164 250	$ 219 000	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 2 190 000
Cost-Procurement	$ 6 807 624	$ 6 807 624	$ 6 807 624	$ 6 807 624	$ 6 807 624	$-	$-	$-	$-	$-	$ 34 038 120
Cost-Total	$ 6 862 374	$ 6 917 124	$ 6 971 874	$ 7 026 624	$ 7 081 374	$ 303 567	$ 333 385	$ 363 202	$ 393 020	$ 422 837	$ 36 392 370
Npv @ 7%	$ 29 605 751										

Equal NPV Analysis for Supplier A and Ericsson

figure 12.9

	Supplier B	Ericsson PKG
MTBF-MHr	0.25	2.00
Warranty-Years	2	5
Cost-New DC/DCs	$ 50	$ 106.21
Cost-Repair	$ 250	$ 250
Interest Rate-%	7	7

Supplier B

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
New DC/DCs	50 000	50 000	50 000	50 000	50 000	0	0	0	0	0	250 000
Total in Field	50 000	100 000	150 000	200 000	250 000	250 000	250 000	250 000	250 000	250 000	250 000
Expected Fails	1 752	3 504	5 256	7 008	8 760	8 760	8 760	8 760	8 760	8 760	70 080
Non Warranty Fails	0	0	1 752	3 504	5 256	7 008	8 760	8 760	8 760	8 760	52 560
Cost-Replace DC/DCs	$-	$-	$ 87 600	$ 175 200	$ 262 800	$ 350 400	$ 438 000	$ 438 000	$ 438 000	$ 438 000	$ 2 628 000
Cost-Repair	$ 438 000	$ 876 000	$ 1 314 000	$ 1 752 000	$ 2 190 000	$ 2 190 000	$ 2 190 000	$ 2 190 000	$ 2 190 000	$ 2 190 000	$ 17 520 000
Cost-Procurement	$ 2 500 000	$ 2 500 000	$ 2 500 000	$ 2 500 000	$ 2 500 000	$-	$-	$-	$-	$-	$ 12 500 000
Cost-Total	$ 2 938 000	$ 3 376 000	$ 3 901 600	$ 4 427 200	$ 4 952 800	$ 2 540 400	$ 2 628 000	$ 2 628 000	$ 2 628 000	$ 2 628 000	$ 32 648 000
NPV @ 7%	$ 23 412 441										

Ericsson PKG

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
New DC/DCs	50 000	50 000	50 000	50 000	50 000	0	0	0	0	0	250 000
Total in Field	50 000	100 000	150 000	200 000	250 000	250 000	250 000	250 000	250 000	250 000	250 000
Expected Fails	219	438	657	876	1 095	1 095	1 095	1 095	1 095	1 095	8 760
Non Warranty Fails	0	0	0	0	0	219	438	657	876	1 095	3 285
Cost-Replace DC/DCs	$-	$-	$-	$-	$-	$ 23 260	$ 46 520	$ 69 780	$ 93 040	$ 116 301	$ 348 902
Cost-Repair	$ 54 750	$ 109 500	$ 164 250	$ 219 000	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 2 190 000
Cost-Procurement	$ 5 310 528	$ 5 310 528	$ 5 310 528	$ 5 310 528	$ 5 310 528	$-	$-	$-	$-	$-	$ 26 552 639
Cost-Total	$ 5 365 278	$ 5 420 028	$ 5 474 778	$ 5 529 528	$ 5 584 278	$ 297 010	$ 320 270	$ 343 530	$ 366 790	$ 390 051	$ 29 091 540
NPV @ 7%	$ 23 412 441										

Equal NPV Analysis for Supplier B and Ericsson

figure 12.10

	Supplier C	Ericsson PKG
MTBF-MHr	0.87	2.00
Warranty-Years	3	5
Cost-New DC/DCs	$ 50	$ 60.43
Cost-Repair	$ 250	$ 250
Interest Rate-%	7	7

Supplier C

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
New DC/DCs	50 000	50 000	50 000	50 000	50 000	0	0	0	0	0	250 000
Total in Field	50 000	100 000	150 000	200 000	250 000	250 000	250 000	250 000	250 000	250 000	250 000
Expected Fails	503	1 007	1 510	2 014	2 517	2 517	2 517	2 517	2 517	2 517	20 136
Non Warranty Fails	0	0	0	503	1 007	1 510	2 014	2 517	2 517	2 517	12 585
Cost-Replace DC/DCs	$-	$-	$-	$ 25 150	$ 50 350	$ 75 500	$ 100 700	$ 125 850	$ 125 850	$ 125 850	$ 629 250
Cost-Repair	$ 125 750	$ 251 750	$ 377 500	$ 503 500	$ 629 250	$ 629 250	$ 629 250	$ 629 250	$ 629 250	$ 629 250	$ 5 034 000
Cost-Procurement	$ 2 500 000	$ 2 500 000	$ 2 500 000	$ 2 500 000	$ 2 500 000	$-	$-	$-	$-	$-	$ 12 500 000
Cost-Total	$ 2 625 750	$ 2 751 750	$ 2 877 500	$ 3 028 650	$ 3 179 600	$ 704 750	$ 729 950	$ 755 100	$ 755 100	$ 755 100	$ 18 163 250
NPV @ 7%	$ 13 942 142										

Ericsson PKG

	Year 1	Year 2	Year 3	Year 4	Year 5	Year 6	Year 7	Year 8	Year 9	Year 10	Total
New DC/DCs	50 000	50 000	50 000	50 000	50 000	0	0	0	0	0	250 000
Total in Field	50 000	100 000	150 000	200 000	250 000	250 000	250 000	250 000	250 000	250 000	250 000
Expected Fails	219	438	657	876	1 095	1 095	1 095	1 095	1 095	1 095	8 760
Non Warranty Fails	0	0	0	0	0	219	438	657	876	1 095	3 285
Cost-Replace DC/DCs	$-	$-	$-	$-	$-	$ 13 233	$ 26 467	$ 39 700	$ 52 933	$ 66 166	$ 164 250
Cost-Repair	$ 54 750	$ 109 500	$ 164 250	$ 219 000	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 273 750	$ 2 190 000
Cost-Procurement	$ 3 021 292	$ 3 021 292	$ 3 021 292	$ 3 021 292	$ 3 021 292	$-	$-	$-	$-	$-	$ 15 106 458
Cost-Total	$ 3 076 042	$ 3 130 792	$ 3 185 542	$ 3 240 292	$ 3 295 042	$ 286 983	$ 300 217	$ 313 450	$ 326 683	$ 339 916	$ 17 460 708
NPV @ 7%	$ 13 942 142										

Equal NPV Analysis for Supplier C and Ericsson

figure 12.11

	Procurement Cost	MTBF (MHr)	NPV Including Initial Procurement	Allowable Converter Cost For Ericsson Pkg	Allowable 'Cost Premium'
Supplier A	$ 50	0.17	$ 29 605 751	$ 136.15	$ 86.15
Supplier B	$ 50	0.25	$ 23 412 441	$ 106.21	$ 56.21
Supplier C	$ 50	0.87	$ 13 942 142	$ 60.43	$ 10.43

Procurement Cost Comparison

figure 12.12

- The warranty period, while less important than reliability, can also have a substantial cost impact.
- A procurement "cost premium" can be justified for power modules with higher reliability.

We have demonstrated cost differences of as much as 17 million dollars in a typical system just due to the reliability differences between some of today's popular DC/DC power modules. It is not "overkill" to use a power module with a several hundred year MTBF in your system - it is good common economic sense. The next time you need to select a power module for your power system, do so carefully. Determine the expected reliability and the temperature at which it is specified. Remember that you are making a multi-million dollar decision.

Development Costs

When doing cost comparisons between alternative power supply implementations it is often only the recurring hardware procurement costs that are compared. These are certainly relevant and vital numbers that need to be included in the analysis, but are not the only cost elements involved with the power supply sourcing decision. Power supply development cost, especially for custom approaches, needs to be considered as part of the overall cost analysis. Included within the domain of 'development cost' are all costs associated with specifying the power requirements, selecting a supplier, developing hardware and doing system design to interface with and accommodate the power supply assemblies. In this section we will identify these cost elements in more detail and offer examples showing what we believe to be typical values for two types of power implementations. It is important to stress again that the reader will gain maximum benefit from these approaches by substituting his/her own estimates of each parameter for their particular system.

We will begin by defining the elements associated with development cost:

Develop Specification – Engineering time to decide upon ratings for each powersupply assembly. This includes electrical, mechanical, thermal and control/diagnostic interface definition and generating and maintaining custom specification documentation if required.

Procurement – Procurement engineering and buyer time and costs associated with selecting a supplier, qualifying the supplier, maintaining quality control, vendor visits, negotiations and supporting problems with the supplier.

Hardware Development – Cost charged by the power supply vendor to develop or modify a custom power supply.

Tooling – Cost charged by the power supply vendor for unique tooling or manufacturing start-up costs.

Prototypes – Cost charged by the power supply vendor for delivery of proof-of-concept prototype power supply hardware.

Qualification – Cost to qualify the power supply vendor and unique design, including reliability demonstration testing.

Approvals – Cost to obtain agency approvals for unique power supply designs.

Distribution Design – Engineering time required to define and design cables, bus bars and harnesses required to distribute power within the system.

Other System Design – Engineering time required to define and design system interfaces to the power supply assemblies. This includes mechanical, thermal and control/diagnostic aspects of the system.

Rework Factor – Anyone who has experienced power system development knows that it seldom goes entirely smoothly. There is usually some rework required and often more than one iteration of a power supply design. This element captures that reality in the form of a confidence factor. For example, a rework factor of 1.0 assumes only one design iteration is required, and a rework factor of 2.0 implies an equivalent cost associated with two complete designs.

We now present two examples that show some typical values for these cost elements. The first example, shown in figure 12.13, estimates the development cost for a power system consisting of one 1000 W custom centralized power supply with four output voltages. The second example, shown in figure 12.14, is a corresponding estimate for a similar system (1000 W, four outputs) configured with a decentralized architecture utilizing a standard front-end AC/DC converter and distributed standard DC/DC power modules. In each case, three estimates (low, nominal, and high) are presented to bracket the expected costs. Hourly rates include overhead and burden. For those without experience in power system development using standardized power modules, some comments relative to figure 12.14 may be useful:

		Cost (K$)		
Cost Element	**Comments**	**Low**	**Nom**	**High**
Develop Specs.	160, 250, 600 Hr @ $60/Hr	9.6	15.0	36.0
Procurement	200, 300, 600 Hr @ $40/Hr	8.0	12.0	24.0
Travel costs		15.0	20.0	40.0
Hardware Development		20.0	40.0	70.0
Tooling		10.0	30.0	50.0
Prototypes		10.0	30.0	50.0
Qualification		50.0	100.0	200.0
Approvals		5.0	10.0	15.0
Distribution Design	80, 120, 250 Hr @ $60/Hr	4.8	7.2	15.0
Other System Design	120, 250, 600 Hr @ $60/Hr	7.2	15.0	36.0
Sub Total		139.6	279.2	536.0
Rework Factor		1.2	1.5	2.0
Total Development Cost		167.5	418.8	1 072.0

Development Cost for 1 kW Centralized Custom Power System

figure 12.13

		Cost (K$)		
Cost Element	**Comments**	**Low**	**Nom**	**High**
Develop Specs.	30, 60, 120 Hr @ $60/Hr	1.8	3.6	7.2
Procurement	20, 30, 60 Hr @ $40/Hr	0.8	1.2	2.4
Travel costs		1.0	2.0	3.0
Hardware Development		N/A	N/A	N/A
Tooling		N/A	N/A	N/A
Prototypes		1.0	2.0	3.0
Qualification		2.0	5.0	10.0
Approvals		N/A	N/A	N/A
Distribution Design	30, 40, 60 Hr @ $60/Hr	1.8	2.4	3.6
Other System Design	80, 150, 400 Hr @ $60/Hr	4.8	9.0	24.0
Sub Total		10.2	18.2	40.2
Rework Factor		1.0	1.2	1.5
Total Development Cost		10.2	21.8	60.3

Development Cost for 1 kW Decentralized Power System using Standard Power Modules

figure 12.14

Develop Specs – This task reduces to selecting standard power modules that will satisfy the needs of the system. No detailed power supply specifications need to be written and maintained, as is the case when using custom power supplies.

Procurement – Since standardized power modules are used much of the procurement can be greatly simplified. Selecting a supplier with a proven record of reliability negates the need for multiple factory visits, specialized quality control techniques and supplier problem solving. Since the supplier is already producing and pricing the part in quantity, minimal cost negotiations are required.

Hardware Development/Tooling – Not required with standard power modules.

Prototypes – Procured from distributor at market price or from module supplier. No lead-time concerns.

Qualification - Reliability data already exists for standard power modules so no testing is required. Qualification entails only review of the supplier's reliability data and manufacturing quality plan.

Approvals – Not required at power supply level. Standard power modules already have worldwide safety agency approvals.

Distribution Design – Simplified due to low current distributed intermediate voltage instead of high current cables and bus bars. DC/DC output distribution can be imbedded into boards and backpanels, eliminating the need for many cables.

Other System Design – Simplified due to better defined and proven power module interfaces. Proof-of-concept model can be con structed immediately with off-the-shelf power modules without having to wait for power supply design and prototypes. Power module supplier application notes and design assistance will make all these design tasks easier and faster.

Rework Factor – The modular nature of this approach allows for much easier 'tweaking' of the power system as the system loads and demands change. These modifications can be done without incurring power supply redesign costs and schedule delays.

One variable in the design and production process that we have not yet addressed is the volume, or quantity, of systems produced. As the volume increases, non-recurring costs associated with development of customized power supplies can be amortized over larger number of systems, reducing the impact on a per system basis. Indeed, the most attractive application for custom power supplies is for equipment with very large production runs of completely identical systems. We show in figure 12.15 the impact of development cost on a per system basis for the two examples presented above. This type of information will be needed

System Quantity	Centralized Custom System			Decentralized Standard Power Modules		
	Low	Nominal	High	Low	Nominal	High
100	$ 1 675.00	$ 4 188.00	$ 10 720.00	$ 102.00	$ 218.00	$ 603.00
500	$ 335.00	$ 837.60	$ 2 144.00			
1 000	$ 167.50	$ 418.80	$ 1 072.00	$ 10.20	$ 21.80	$ 60.30
5 000	$ 33.50	$ 83.76	$ 214.40			
10 000	$ 16.75	$ 41.88	$ 107.20	$ 1.02	$ 2.18	$ 6.03
50 000	$ 3.35	$ 8.38	$ 21.44			
100 000	$ 1.67	$ 4.19	$ 10.72	$ 0.10	$ 0.22	$ 0.60

Power Development Cost per System

figure 12.15

for the total cost analysis discussed in following section.

Total Cost Analysis

In order to compare the actual costs of alternative power system implementations, it is important to include all the relevant elements. These fall into four categories:

- Hardware procurement costs.
- Development costs.
- Reliability costs.
- Time to Market costs.

Hardware procurement costs are always included when doing cost analyses but the other items are sometimes ignored. We have seen in sections “Reliability Costs” and “Development Costs” how estimates of reliability costs and development costs can be developed. In this section we will show how this type of information can be combined with the other cost elements to obtain a total cost estimate for the power system.

Time-to-Market costs have not yet been addressed. These are costs (or more accurately lost opportunities for profit) that are incurred due to the length of the development process. For example, if a custom power supply development program is undertaken and the delivery and qualification of the power supply causes a delay in the shipment of the end product, there is a loss of sales and profit. This loss of revenue should then be recorded as a cost associated with the power system. Another scenario that often occurs is the need to modify the power system after the end product has been in production to accommodate system upgrades or design enhancements. Delays in the introduction of such enhancements while waiting for modifications in a custom power supply represents loss of revenue and time-to-market costs for the power system. Time to market costs are often very difficult to estimate since they depend upon marketing forecasts rather than on more quantifiable data. In spite of this difficulty, it is important to do some type of estimate as these effects and their associated costs are real and should not be ignored.

As an example of how the time to market estimate can be quantified, we will use the 1000 W system previously discussed. The following assumptions are made:

- Custom centralized power system causes a delay of 6 months in shipping the end product due to either a problem with design/qualification or longer lead-time inherent with custom design vs. standard modules.
- Sales projection for end product is 1000 systems per year for 5 year period.
- Lost sales at beginning of program are not recovered.
- End product sales price is $10,000.
- Profit is 15% ($1500 per product sold).

Using the above assumptions, the time to market cost can now be estimated:

- Lost product sales = (1000 products/yr) × (0.5 yr) = 500.
- Lost Profit = (500) × ($1500) = $750,000.
- Time to market cost = Lost profit / system shipped = $750,000 / 4500 = $166.67.

While this type of lost profit is incurred for a 6 month delay, it should be kept in mind that this can often be a conservative estimate. Missing the introduction of a product can affect sales more than linearly as assumed here. Loss of customer confidence and opening the door to competitors and their products can result in real cost impacts that dwarf the estimate shown here.

Figure 12.16 depicts a methodology for doing overall power system cost analysis. This form can be useful as a 'checklist' to ensure that all cost elements are considered when estimating power system costs. The reliability cost referred to can be obtained by dividing the NPV developed in section "Reliability Costs" by the quantity of systems.

An example of how this methodology might be used for a specific system is shown in figure 12.17, where the cost elements of the 1000 W system discussed above are summarized. Since this is a fictitious system, it should be noted that many of the numbers in figure 12.17 are arbitrary estimates rather than actual values. The bottom line costs show that the two options result in roughly similar total costs. Consequently there is shown to be no cost penalty for the following benefits of decentralized power that are not accounted for in the cost estimate. Refer to chapter 5 for additional discussion of these advantages.

- Flexibility for upgrades and features.
- Compatibility with battery backup for enhanced availability.
- Possibility of second sourcing of standard modules.

It is hoped that the methodology presented here will be useful in understanding and comparing power system costs. Developing estimates of these cost elements for your particular system will yield not only cost data but also an increased appreciation for how power system design can affect the profitability of the end product.

To make this task more manageable, Ericsson has published additional information and design tools relating to cost analysis. Design Note 003 offers more detail on the methodologies of cost analysis. An interactive system lifetime cost calculator tool is also available both on CD-ROM and on the Ericsson Microelectronics website. This tool allows the power system designer to enter appropriate assumptions for his/her system and easily obtain the cost impacts over the life of the system.

Hardware Costs
Power Supplies
AC Distribution
DC Distribution
Controls / Diagnostics
Power Supply Cooling
Mechanical
Development Cost
Reliability Cost
Time To Market Cost

Elements of Total Power System Cost

figure 12.16

1000 W Four Output System - Total Quantity = 5 000 Units				
Cost Element	**Cost Per System - $ Centralized Custom**		**Cost Per System- $ Decentralized Standard**	
Hardware Costs				
AC/DC	395.00		450.00	
DC/DCs	N/A		525.00	
Distribution	250.00		45.00	
Controls	30.00		30.00	
Cooling	40.00		15.00	
Mechanical	20.00		20.00	
Total Hardware Cost	$ 735.00	$ 735.00	$ 1 085.00	$1 085.00
Development Cost		$ 83.76		$ 4.36
Reliability Cost		$ 140.00		$ 20.00
Time To Market		$ 166.67		N/A
Total Cost		$ 1 125.43		$ 1 109.36

Power System Cost for 1000 W Four Output Example

figure 12.17

13 EMC Design for Decentralized Power

EMC has historically been one of the more difficult system design criteria for switching power converters. The required standards have been uncertain. The distinction between converter performance and system performance has been cloudy. Design techniques for meeting EMC requirements have sometimes seemed like "black magic". We hope to eliminate the confusion and offer sound design approaches that will result in successful system EMC design

Introduction

EMC (Electromagnetic Compatibility) design is one of the more important challenges for the electronics system designer. Most all equipment types need to meet one or more EMC standards at the system level, and there are many international standards and controlling agencies to deal with. All DC/DC converters contain one or more switching stages which generate a noise spectrum that is capable of leaving the module by conduction and radiation. As we will see later, conduction is usually the more troublesome mechanism. Even though the applicable standards are imposed at the system level rather than at the power module, understanding and controlling the conducted emissions performance of the power module is a good beginning to meeting the system-level specifications.

Unfortunately for the system designer, the available standard DC/DC converters have a wide range of conducted emissions performance. Some contain internal filtering and will meet many conducted emission standards without any external components. Others may contain no internal provision for EMC filtering and may require an external filter at some point in the system in order to meet the applicable standards. Many suppliers of standard DC/DC converters complicate the system designer's dilemma by providing poor or non-existent specifications for the EMC performance of the converter. Our intent here and in other Ericsson publications is to provide useful information to ease the task of the power system designer.

We will first define some of the basic mechanisms involved with the generation of EMI (Electromagnetic Interference) and its propagation between the converter and the system. The source of EMI from a DC/DC converter is the rapid switching of high current levels. Modern converters operate at frequencies of 150 kHz and above with rapid switching transitions to optimize conversion efficiency. Today's low voltage high current system requirements result in the switching of 10s of amps at several locations within a DC/DC converter. High current in conjunction with rapid transitions creates voltage disturbances due to the parasitic inductance present in components and interconnection traces in the converter and power system. For modern DC/DC converters, most of the EMI noise is generated in the primary switches, the high frequency magnetics and the output rectifiers.

EMC design encompasses both emission and susceptibility. Emission addresses electrical noise coming from the converter and possibly interfering with the remainder of the system. Susceptibility (sometimes referred to as immunity) addresses electrical noise coming from the system (or external systems) and possibly interfering with the operation of the DC/DC converter. For DC/DC converters, emissions are almost always more of a concern than susceptibility.

While the converter design goal is to minimize the amount of escaping noise, some of the generated EMI can exit the converter by means of two mechanisms – conduction and radiation. Conducted noise can exit from any of the converter's connections – input, output or control lines. The input connections usually have the highest levels of conducted EMI because the power conversion switching occurs on the primary side of the converter. Radiated noise can escape from the converter itself or from wires and board traces connected to it. With the prevalence of multilayer PCB packaging with ground planes, radiated EMI is not normally a problem when using high quality DC/DC converters. We will focus mostly on the conducted path in the remainder of this chapter.

A final concept that must be understood is the distinction between differential mode and common mode noise, which applies to conducted emissions from the DC/DC converter. A differential mode (sometimes called single-ended) signal represents the classical two-wire input to or output from a converter. By definition, the current going into one terminal will be equal to the current coming out of the other terminal as shown in figure 13.1. The definition applies to both AC and DC currents, but we are considering only AC currents when working with EMI. As an example, measuring the output ripple and noise of a converter is an attempt to measure the differential mode output noise voltage.

The real world is more complex than the above assumption, however. When working with high frequency AC signals, the capacitive coupling between points in the circuit and surrounding circuitry can cause AC currents to flow in conductors other than the intended

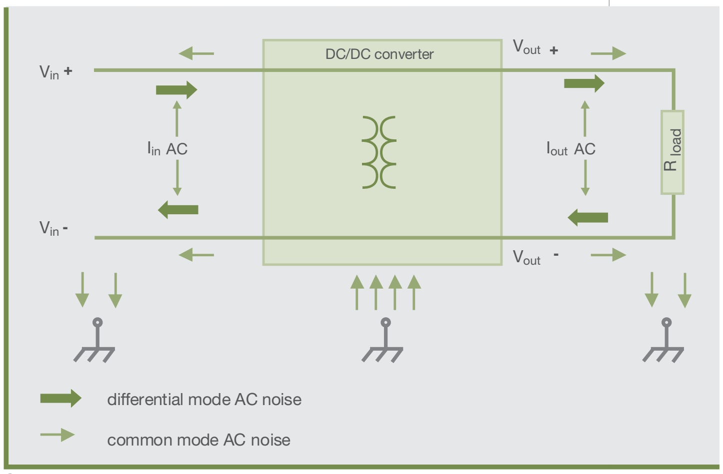

Conducted Noise

figure 13.1

two wires. These are referred to as common mode currents. When working with distributed power systems, the most common path for the common mode currents to complete their circuit is through chassis ground as shown in figure 13.1. Note that a DC connection to chassis does not need to exist in order for these currents to flow - capacitive coupling is sufficient. For the input or output terminals, common mode noise will be in phase, whereas differential mode noise will be out of phase.

The power system designer must understand the distinction between differential mode and common mode noise, as techniques for noise suppression are different for the two types of noise. Noise suppression filter designs, for example, are significantly different for the two noise modalities. Another important distinction is that differential mode noise is much more predictable and reproducible. Common mode noise, on the other hand, can vary significantly from system to system when using the same DC/DC converter as a function of the PCB layout, external component choices, and system grounding design.

Standards and Requirements

The first step in achieving a sound EMC design is to understand the requirements. There are no recognized conducted emissions standards that apply specifically to board mounted DC/DC converters or power modules. The applicable standards that the system designer must meet are system level standards. That is, the end-use equipment must meet a selected set of conducted emission standards, depending upon the equipment's usage and the country into which it is sold. We discussed these requirements in a general way in chapter 4, and will explore them more fully here.

Most equipment has one interface with the power utility, and this is where the conducted emissions standards apply. Individual power modules and DC/DC converters are isolated from the power line by such items as EMI filters, circuit breakers and fuses, transient protection networks and AC/DC front-end power converters. Consequently, any noise conducted from the DC input of a power module does not appear directly at the power line – the noise is modified (usually reduced) by the mentioned components. The result is that the equipment may meet its requirements without any of the several DC/DC converters or power modules meeting the EMC standard that is applicable to the equipment.

Many systems will meet all applicable EMC standards by only using the EMC filter normally associated with the front-end power supply or the powerline entry filter – without any specific provision for filtering the intermediate DC bus voltage. In other cases doing some filtering at the DC bus voltage level can help the system meet its overall requirements – often saving cost and complexity by sharing one filter between several DC/DC converters or power modules.
Even though individual converters and power

modules do not need to meet the conducted emissions EMC standards, it can be constructive to measure their performance relative to such a standard. This provides a common, easily understood measurement methodology that can be applied using the same measurement equipment as is used for the end equipment. It also gives a relative measurement of the EMC spectrum of each converter or module so that the equipment designer can make the proper decisions regarding how to meet the system-level requirement.

Even though there are many different conducted emissions standards in existence in various countries, the situation is not as complex in practice as it may seem at first glance. Consolidation of requirements within Europe has made the equipment designer's task much more manageable. The two most commonly used requirements are the IEC specifications for the European market (and adopted by many other markets) and the FCC specifications within the US – and these two standards are very similar to each other. Both of them will be discussed here.

The European requirements derive from the International Electrotechnical Commission (IEC). The IEC body responsible for EMC requirements is the International Special Committee on Radio Interference (C.I.S.P.R.), and the requirement that defines the strictest limits on conducted emissions is C.I.S.P.R. 22. These limits are described in EN 55022 and EN 55011, which are the corresponding European Economic Community requirements. The European standards then define two different conducted emission limits as a function of equipment type and usage:

- EN 50081-1 European generic emission standard for domestic equipment (refers to basic standard EN 55022 class B for conducted and radiated noise).
- EN 50081-2 European generic emission standard for industrial equipment (refers to basic standard EN 55011 class A for conducted and radiated noise).

Domestic equipment is typically portable and can be moved around and plugged into various power outlets. Home and office equipment falls into this category. Industrial equipment is more typically "hard-wired" in place and is not portable. The regulations are tighter and more restrictive for domestic equipment. EN 50081-1 requirements are shown in figure 13.2 and EN 50081-2 requirements in figure 14.3. In practice, the industrial equipment standard is sometimes referred to as "Class A" while the domestic standard is called "Class B" to correspond with the FCC nomenclature to be discussed later.

The requirements are expressed in terms of the noise voltage impressed upon the powerline as a function of frequency, 0.15 to 30 MHz for the European requirements. This is only meaningful if the powerline impedance is known. Consequently, a Line Impedance Stabilization Network (LISN) is specified and used when making the measurement to set the line impedance to 50 ohms. The limits are expressed in terms of voltage decibels relative to 1 microvolt. For example, a measurement of 250 mV would be expressed as
20 log 250 = 48 dB.

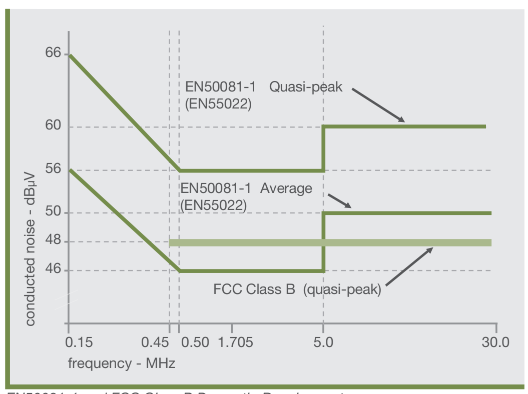

EN50081-1 and FCC Class B Domestic Requirements

figure 13.2

The European standards give two limits as a function of the type of detector used for the measurement. A higher limit is provided for a quasi-peak detector (less smoothing) and a lower limit for use with an average detector. Both limits must be met for the equipment to pass the requirement.

For purposes of future awareness, it should be mentioned that there is a proposed annex to EN 50081 that would specify conducted emission limits for DC input systems (i.e. -48 V input telecom systems) using a different test methodology than the existing EN 50081 requirements and the FCC specifications. Because this standard is at this point just a proposal and would require testing with different equipment (current measurement rather than voltage), we will not include it in the remainder of the discussion here.

Another standard that is sometimes relevant is EN 300 386-2. This standard applies only to telecommunications network equipment, with either AC or DC inputs. For AC input systems, the requirements in this standard are the same as the EN 50081 standards. For DC input systems, the test limits when using an averaging detector are identical to the test limits in EN 50081-2. When using a quasi-peak detector for DC input systems, the test limit between 0.15 MHz and 30 MHz is identical to the EN 50081-2 requirement, however, the 79 dBmV limit is extended downward in frequency from 0.15 MHz to 0.02 MHz. Since Ericsson Power Modules have a fundamental operating frequency higher than 0.15 MHz, they easily meet this extended lower frequency requirement.

The other most commonly used requirement is the US Federal Communications Commission (FCC) regulation for unintentional radiators as defined in part 15. This standard calls for test methodologies very similar to the European standards, and since the same 50 Ω LISN is used, the same measurement data can be applied to either set of standards. The FCC conducted emission standards begin at 0.45 MHz (rather than the 0.15 MHz of the European standard) and continue to 30 MHz. The limit is expressed in terms of a quasi-peak measurement. As with the European standard, there are two sets of requirements. The "Class A" requirement corresponds to the European industrial specification, and the more restrictive "Class B' is used for equipment that would be measured to the European domestic standard.

The FCC limits are also shown in figures 13.2 and 13.3, and are a few dB lower than the corresponding quasi-peak European requirements. Be aware, however, that language in the FCC standard allows for reduction of some of the quasi-peak measurements if the quasi-peak reading is 6 dB or more higher than the reading with an average detector, making the FCC requirement not as stringent as it appears at first glance. There is also language in the FCC regulation allowing use of the test methodology and limits of EN 55022. Thus, as a practical matter, equipment meeting the European standards will be compliant with the FCC requirements. For this reason, we will focus on the EN 55022 standards as the baseline requirements for the remainder of this discussion.

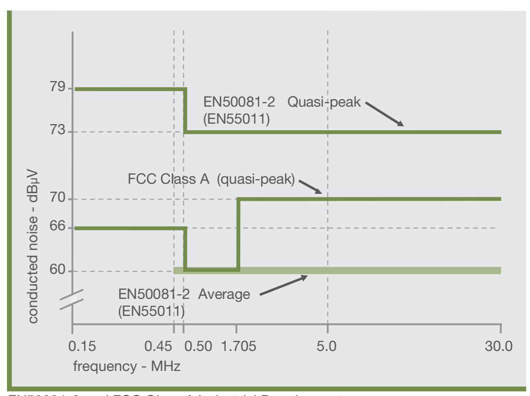

EN50081-2 and FCC Class A Industrial Requirements

figure 13.3

The above has been a generalized overview. When determining the appropriate standards and test limits for your system make sure that you consult with the most knowledgeable EMC and regulatory agency people at your company so that the latest regulatory requirements are addressed.

Compliance Strategies

From an overall system design point of view, all aspects of EMC must be considered, not just conducted emissions. While the power converters in the system will not represent the limiting factor for most EMC requirements, it is instructive to summarize their EMC behavior as it relates to an overall system strategy for regulatory compliance.

Radiated susceptibility is almost never a problem with power converters. This is because power converters operate with very high levels of current and consequently have low impedances at most points in their circuitry. The modest levels of radiated fields imposed on the system during a radiated susceptibility test will only affect low power high impedance circuitry that is unshielded. Consequently, a practical approach is to assume that radiated susceptibility will not be a problem with either centralized AC/DC converter or distributed DC/DC converters.

Radiated emissions are almost never a problem with DC/DC converters, but compliance may require some attention to proper system design. Most equipment contains a metal enclosure that will effectively shield the internal converters from the outside environment where the emission measurement is made. A more valid concern is radiated interference from the converter affecting other circuitry internal to the equipment. With usage of multilayer printed circuit boards containing ground planes, this exposure is minimal. Even the newer DC/DC converters constructed with "open frame" packaging without a separate metal shield around the converter pose no problem in this regard with proper design. Perhaps the most common radiated emission problem with power converters arises in cases where the input or output distribution to/from the converter is done with discrete wires rather than on circuit board traces. Without proper decoupling and/or shielding, these wires can radiate noise and become a concern for radiated emissions. If proper decoupling is used with multilayer printed circuit board construction, however, it can be assumed that high quality DC/DC converters will not create a radiated emissions problem for the end equipment.

Conducted susceptibility requirements are also not a problem for DC/DC converters. The conducted susceptibility test waveforms are applied at the input to the equipment, and will not propagate to the DC/DC converters with proper system design. The EMC filter at the input to the equipment along with the bus holdup capacitance on the intermediate bus as described in chapter 8 will absorb the transient energy so that the wide input voltage range DC/DC converters will not be affected.

This leaves conducted emissions which, from a practical matter, will be the most likely source of a system level EMC problem arising from power converters. But with the selection of proper DC/DC converters and reasonable system design practices, this too is a very manageable situation. In the following section we will describe the benefits of Ericsson power modules in this regard. The most common overall system compliance strategy for conducted emission is to design for the European domestic emission limits. Equipment meeting these limits should be suitable for regulatory qualification and sales in all but the most unusual worldwide markets. Note that individual DC/DC converters need not meet the domestic (Class B) limits in order for the system to comply.

Power Module Performance

While EMC compliance is ultimately a system level effort, the selection of the DC/DC converters internal to the system can have a large influence on its success or failure. The priority of EMI performance relative to other design goals varies from manufacturer to manufacturer. Some converter suppliers will not include any additional internal components for EMI suppression in order to minimize their production cost. Others may not invest the time to completely understand the EMC performance of their products and to publish accurate data along with suggestions for overall compliance at the system level. Ericsson has put a big effort into understanding and characterizing the EMC performance of its power modules and in supporting the system designs of its customers. Ericsson's experience with telecom system design is very helpful in this regard.

Ericsson power modules are designed with compliance to conducted emission specifications as a high priority. Many power modules include internal filters that will even allow the module to meet some conducted emission standards as a "stand alone" module without external filtering or the positive influence of other system elements. Some power modules with a metal case include provision for connecting to the case via a module pin for grounding purposes. Datasheets for Ericsson's power modules now contain plots of the EMI performance vs. the European test limits. In some cases, suggested external filter designs are given to allow for full compliance to class B requirements at the individual module level. A more concentrated treatment of the EMI performance of Ericsson power modules can be found in "DN009 - Conducted Emission Performance of Ericsson Power Modules, Characterization and System Design". This document, which is available on the Ericsson Microelectronics website, includes the following useful information:

- Standards definition.
- EMI provision internal to Ericsson power modules.
- Definition of test setup for measurement of converter EMI performance.
- Measured EMI spectra of power modules.
- Design data on tested external filter networks for low emission requirements.
- System EMC design suggestions.

System Design Guidelines

The EMC performance of the power module itself will not guarantee a successful system. Equally important are the layout and grounding practices used by the system designer for the power module and the power distribution system. Some items that will enhance the system performance include:

- Use adequate decoupling capacitors in the DC distribution system to the load. Include both low frequency bulk capacitance and high frequency ceramic capacitors.
- Use short leads on all filter and decoupling components.
- Minimize inductance by using wide distribution traces over ground planes.
- Avoid usage of discrete wire for power distribution, especially outside of equipment enclosures.
- Use separate input and output ground planes, and return common mode noise to the input ground.

It is sometimes required to construct small filters within the system for the purpose of reducing the degree of conducted emissions from the end product. The Ericsson documents referenced above are a good source for these filter designs, including component values. Components used are mostly miniature SMD inductors and ceramic capacitors. The capacitors may be referred to as either “X caps” or “Y caps”. X caps are installed across the two input conductors to the converter for the purpose of filtering differential mode noise. Y caps are installed from both input conductors to ground for the purpose of providing a convenient path for the common mode currents to return without exiting the equipment.

As examples, two types of commonly used filters are shown in figure 13.4 - Pi and common mode. The Pi filter is effective with differential mode noise and has the advantage of being simpler and less expensive. The common mode filters will attenuate noise that is common to both input conductors as well as differential mode noise. These filters tend to be more effective, but have the disadvantages of additional cost and complexity and the need to terminate the “Y capacitors” to a nearby earth ground point.

These filter designs can be physically implemented in a very small board area using normal printed circuit board packaging techniques in conjunction with miniaturized SMD resistors and inductors. In practice, one filter network, with appropriate attention to current levels, can be used to filter more than one power module. For example, many “power per board” distributed systems contain one filter near the DC input to the board that provides the filtering function for several power modules on the board.

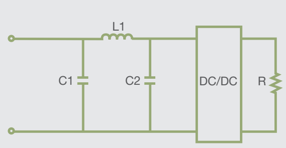

filter	compopnent	value	rating	type	suggested part
B	C1, C2 L1	0.68 µF 10 µH	100V 0.5A	ceramic DC choke	

a) Pi filter

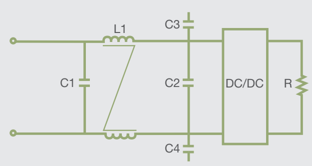

filter	compopnent	value	rating	type	suggested part
B	C1, C2 C3, C4 L1	0.68 µF 10 nF 1 µH	100V 1500V 0.5A	ceramic ceramic common mode	Siemens B 82790-S0513

b) common mode filter

Typical EMI Filters

figure 13.4

14 DC/DC Power Module Selection

The purpose of this chapter is to summarize and prioritize the criteria that we feel should be used by power system designers when selecting DC/DC power modules. We will also highlight some specified parameters that are widely advertised and promoted but which do not result in measurable benefit in the system applications.

Introduction

Every application is somewhat unique, of course, but there tend to be some common criteria that most users find to be relevant and important. These criteria will be summarized here, in rough order of importance. Some of this information can be found on the converter datasheet, while some will require communications with, or research into, the supplier of the power module.

Datasheet Specifications

Efficiency

Efficiency has many benefits. It is good for its own sake in terms of energy conservation and utilizing the minimum amount of power from the AC powerline or battery utility. It is important because it determines the amount of heat dissipated in the power module that must be removed by the product cooling system. This dissipated heat will increase the internal product ambient temperature and affect the reliability of other system components. Efficiency is important because the dissipated power determines the temperature rise internal to the power module and consequently the ultimate internal component temperatures and the reliability. Both minimum and typical efficiencies are important specifications. The typical data from reputable suppliers are believable and can be very useful for determining average system thermal load and for reliability estimations. For applications that require a variation in the output load current from the converter, or if the converter is operated at a derated or light load, look for curves in the specification sheet that show typical efficiency vs. output load. There can sometimes be surprises here. Some power modules on the market have severely degraded efficiency at low load current. Unless the power system designer is aware of these conditions the system design may be jeopardized.

Operating Temperature

The maximum operating temperature is a very critical specification. It is usually specified as maximum baseplate or case temperature before derating is required. For conduction cooled power modules it can be specified as maximum pin temperature.

In many cases this specification determines the actual power that can be supplied by a power module in a given application. The resulting number is often a shock to the user after seeing the advertised power density. Refer to chapter 10.3 for an example of how seemingly high density power modules can deliver very little useful power in some applications. This is especially true for convection cooled systems in uncontrolled environments where room temperatures of up to +65 °C are commonplace. The maximum operating temperature is an overall indication of the merit of the power module thermal design. For a fixed component or junction temperature, power modules with higher operating temperature ratings will deliver significantly more power to the load than those with lower temperature ratings. This will be a big factor for those users who are after maximum economy in terms of dollars per watt. High reliability power modules are now available with full load maximum operating case temperatures up to +115 °C.

Reliability

Reliability appears on every system designer's list of important criteria and it is now becoming more and more commonplace on the specification sheets for DC/DC power modules. It is a somewhat unique specification in the sense that it cannot be directly measured. Consequently some judgment must be made about the credibility of the module supplier, the techniques and conservatism with which they arrived at the reliability prediction, and their willingness to assist and support the customer in the application dependent aspects of reliability design and prediction. When obtaining and comparing reliability specifications there are several assumptions that should be determined. What power module temperature is the estimate valid for? How was the estimate arrived at? MIL-HDBK prediction? Supplier failure rate database? Has life testing been performed on the power module design? What is the field history? Are there any known wearout mechanisms in the design or manufacturing process?

In actuality, all three of these criteria – Efficiency, operating temperature, and reliability – are very interrelated. They reinforce each other. Unlike the more typical situation in engineering where desired properties need to be traded-off against each other and cannot be simultaneously optimized, here we have a synergistic relationship between three very desirable properties that can all be obtained with a proper design (figure 14.1).

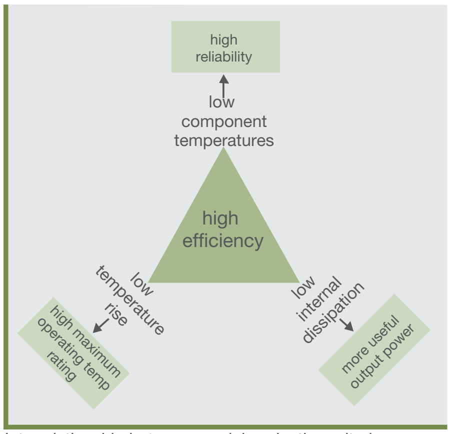

Interrelationship between module selection criteria

figure 14.1

Input Voltage Range

The input voltage operating range should be carefully examined. For telecom systems the 48 V nominal system in common usage actually will range from 38 V to 60 V under some conditions. A power module designed to operate over this entire range will offer significant advantages in terms of availability than one with a more limited range. Some European telecom systems utilize a nominal 60 V battery with maximum voltages up to 75 V. With some DC/DC converter families it is necessary to use a different power module to cover this high voltage range. As a result, two part numbers must be kept in inventory, which has a negative impact on procurement economies. A power module that covers the entire 38 V to 75 V range will provide the advantage of being usable for both types of products.

Weight

This is a specification that has not had much attention in the past. With board mounted distributed power modules and SMD assembly it takes on increased importance. Heavy power modules require special mechanical retention mechanisms that add complexity and cost. Excessive module weight can also impose excessive stress levels to the circuit board when the product is exposed to vibration and shock, such as during shipment of the product. Heavy power modules can cause the board to flex and fracture. The lower the mass of the power module, the less mechanical stress is imposed and the better the design from a cost and reliability point of view.

Power Density

Power density is one of the most frequently promoted parameters of DC/DC power modules and also the most useless. A '50 W/in^3' converter may provide only 5 W/in^3 in a realistic product application, especially if the maximum operating temperature rating is low. Power densities should only be compared after heatsinks and other thermal design require-ments are taken into account. This is very important for convection cooled equipment. Of much more importance than the power density rating is the maximum operating temperature and efficiency.

Topology

Some suppliers promote the benefits of one particular topology. Given that designs are well executed, topology is of little actual consequence for the end user. While it is true that the topology selected by the power module designer can affect cost, reliability, and some aspects of performance, all of these effects should be visible in other areas of the datasheet. Across the power range of today's DC/DC power modules (5 W to 300 W), there is no one optimal topology. Intelligent designers and suppliers will select different topologies as a function of the power requirement and other considerations. This selection and implementation will be relatively transparent to the user of the power module. Perhaps the one exception is resonant converters, where the variable operating frequency often needs to be a consideration in the system design.

Operating Frequency

Operating frequency is another parameter that is more important to power module designers than to users. In spite of this, it sometimes is heavily promoted due to the general relationship between higher frequency and smaller size, and higher density. As we have previously seen, however, the resulting density is often illusionary in terms of real-world density in product applications. Almost all present day DC/DC power modules operate at 100 kHz or above and thus achieve significant size benefits relative to older 20 to 50 kHz designs. Additional miniaturization is due as much to intelligent package design, innovative component development, and increased integration as it is to increasing the operating frequency. So, for pulse width modulated fixed frequency power modules, the operating frequency tends not to be an important selection criterion. For variable frequency topologies the effects upon the system of the range of operating frequencies must be considered.

Regulation

Regulation tends to be over emphasized as a DC/DC power module parameter. Regulation (and associated parameters such as drift, initial setting accuracy, and temperature coefficient) certainly needs to be tight enough to guarantee that the output DC voltage satisfies the requirements of the load circuitry. These requirements are most often fairly reasonable however. A converter with a 0.05% regulation specification is, in practical terms, no more useful than one with a 0.2% specification. In fact, there are situations in which the more loosely regulated converter offers advantages. When two or more converters are operated in parallel, tight regulation sometimes requires elaborate external circuitry to achieve current sharing. More loosely regulated devices can automatically current share due to their output voltage vs. current characteristic. This technique is described in greater detail in chapter 9.

Power Module Supplier Criteria

The better DC/DC power module manufacturers try to make their datasheets as complete, informative, and useful as possible. However, the datasheets contain mostly technical criteria that are measurable and quantifiable. When selecting a power module supplier there are many factors to consider that are not included in the specifications. These factors can be equally as important as the technical data.

Design Philosophy

The long-term performance and reliability of a product is only as good as the care that goes into its design. Even though the power module is small and contains relatively few parts, designing such a unit is a very sophisticated task. The designers must be knowledgeable in many areas, including electrical, mechanical and thermal stress mechanisms. Selecting appropriate topologies, components and packaging techniques, when done correctly, requires high skill levels and commitment of resources for extensive periods of time. Design criteria are important. Components should be derated to enhance reliability. Very complete thermal modeling and analysis should be done so that the designers know exactly the thermal environment and stresses on every component, including profiles of internal component temperatures vs. operating conditions.

Silicon semiconductor devices can operate with junction temperatures up to +150 °C, but with negative impacts to long-term reliability.

Some module suppliers offer products with junction temperatures approaching this value under some operating conditions. Significant reliability improvement can be achieved by keeping junction temperatures under +120 °C at all times. This may require selection of a more costly device with larger die area or development of special packaging or thermal designs to limit the temperature rise. Commitment to details like this rather than taking short cuts with the design is what differentiates the manufacturers' design philosophies and resulting products.

Testing of new designs is also important. In spite of the economic incentives to place new designs on the market as soon as possible, a supplier who cares about the integrity of its designs will do extensive testing of the power module before releasing it for general availability. This testing will include stress screening and evaluation, life testing, and internal application experience.

Component Sourcing & Controls

Part of the design process is the selection of appropriate components. The final product is only as strong as the weakest component, so this process is critical to the success of the converter module. The power module manufacturer should select component suppliers that have the same commitment to quality and reliability as the manufacturer itself. The components must be understood in great detail, not only their performance parameters, but also the component's manufacturing process and all materials used in the component. The power module manufacturer should have a database with actual failure rates for each component under the operating conditions they will be exposed to in the converter application. Components with known reliability limitations, such as electrolytic capacitors used in input filter designs, should be avoided.

When appropriate components are not available, suitable devices need to be developed – either in conjunction with an external supplier or internally by the power module supplier. This type of work can be very time consuming and costly but the result is better designs and advances in available products.

Integration

One of the keys to achieving high reliability, high packaging density, and lower cost, is integration. Using a custom IC can eliminate dozens of discrete components and the associated interconnections with very dramatic improvement in reliability. Utilization of thick film resistors rather than discrete units also has a similar result. Applying these integration concepts to a DC/DC power module can result in a unit with very minimal parts count, enhanced functionality, extreme reliability, and low cost.

Manufacturing Process

Of equal importance to design is the manufacturing process. The key to successful production of large volumes of units with repeatability and cost effectiveness is automation. Automation is capital intensive and requires a large up-front monetary commitment on the part of the manufacturer, but the rewards are extensive. The resulting cost efficiencies are required to make board mounted decentralized power conversion a viable reality.

The successful supplier of power modules will have a manufacturing environment that is closer in concept to a semiconductor process line than to a traditional power supply manufacturer. Indeed, the ultimate goal of decentralized power architecture is to make the power module appear to the user as a component, just like ICs. Some of the techniques used will include building the power module on a ceramic substrate as a hybrid, with thick film resistors, automatic placement equipment, and automated real-time testing (ATE). Every material, chemical, and process used during manufacturing must be very tightly characterized, understood, and controlled. Cleanliness is of utmost importance in order to achieve the very extended operating lifetimes that many applications demand.

Reliability Focus

There has been considerable treatment of reliability issues in this tutorial. This is due to the emphasis that Ericsson places on reliability in the design and application of its products. A good power module supplier should have this type of reliability focus in all areas of its operations, from design to customer support. They should have extensive component-level reliability data, conduct stress and life testing on the power modules, begin with conservatively rated and highly integrated designs, and have a modern and very clean manufacturing process line. 100% burn-in testing should be conducted to assure that the modules are delivered with the lowest possible intrinsic failure rate. Application field history data should be available so that reliability under actual real-world conditions is understood. With close attention to details such as the above, DC/DC power module failure rates as low as 200 FIT are possible, making decentralized power architectures a practical reality.

Application Knowledge

A 'vertically integrated' power module supplier, who is also a user of the power technology, offers advantages over companies that manufacture nothing other than power supplies. The applications knowledge gained by building and maintaining products is vital to understanding the true needs and requirements for power modules. The unrestricted access to field reliability data helps with the development of components and power modules that are designed to contend with the real-world stresses encountered in product environments. The reputation of the supplier for the reliability of his end products is very important, and the power modules that go into these products are a big contributor to the overall equipment reliability. Thus the supplier has strong incentives to make the power modules extremely reliable, and not to cut any corners with the design. They are interested in more than selling power modules – they are motivated to supply power modules that will retain high level performance over extended product lifetimes.

Customer Support

Customer support takes many forms. Applications assistance is one important area. The product-level knowledge obtained by the manufacturer from his own product operations should be available to his customers. Assistance with electrical interfaces, packaging solutions, system level thermal design, reliability enhancement and prediction, and power module selection are only a few of the types of inquiries that are commonly handled.

The better suppliers also help the customer understand as much as they want to about the internal design and manufacturing of the power module. They should share component selection criteria, thermal data, and other required design details. A manufacturer who is proud of their product and manufacturing operations will offer tours of the manufacturing process line so that the customer can observe first hand the care and cleanliness inherent in the construction of the product.

Second Sourcing

The second sourcing procurement concept was developed as a way to avoid reliance on a single supplier of an item, in this case power modules. Inherent in the second sourcing philosophy are two assumptions:

- The power module suppliers are unreliable.
- The second-sourced part is identical to the main source.

In our experience, both of these assumptions must be challenged. The better suppliers of high density DC/DC power modules have built their entire business around the ability to consistently manufacture and supply large quantities of power conversion products. In many cases, the power modules are needed for their own internal products also, so there is even additional incentive to have an uninterrupted flow of product from the manufacturing line. The reputation of the supplier is at stake, and the commitment a company such as Ericsson makes to their customers is a meaningful one.

The second assumption is typically not true. In the board-mounted DC/DC power module marketplace there are many claims for 'pin compatibility'. The products are generally pin compatible in the sense that they can plug into the same board layout and inputs and outputs end up connected to the right places. But there is sometimes an assumption that pin compatibility implies identical performance and reliability. This is not often true. Internal designs, components, derating criteria and manufacturing processes are different. The result is that the power module will not offer the same levels of performance and reliability.

There is one exception to this general situation. That is where a particular power module design becomes so much of an 'industry standard', that the original manufacturer licenses the design to others for manufacturing. If this licensing is done properly, the design content, component sourcing criteria, and manufacturing controls can be transferred intact, and the resulting product be comparable in quality. An analogy outside the world of power supplies would be your ability to go into a McDonalds® anywhere and know with a high degree of certainty what you will receive when you order a Big Mac®. Design and manufacturing licensing, when done correctly, can have similar results. But the user must differentiate this type of arrangement from the more usual situation where a supplier just copies a pin layout and uses a unique design. In this instance, there are no guarantees as to what the results will be.

The most important procurement decision is selecting the right main supplier. If this is a company such as Ericsson, the user should feel comfortable without a second source. In our view, it is better to have one reliable and dependable source than two or more questionable ones. In the case where the converter selected is licensed by the manufacturer to provide alternate manufacturing sites, these controlled second sources can be a way of enhancing procurement logistics. But these considerations should not overshadow the main purpose – getting the best possible DC/DC power module into the end product.

References 15

By necessity, this treatment of power systems and DC/DC power modules has been of an overview in nature. For readers who wish to expand upon this content with additional information, the following references are recommended.

Power System Architecture

1 'Distributed Power Architecture Delivers Solid Advantages' (Electronics Products and Technology, March, 1995).

2 C.D. Smith, Electronics Packaging Forum, Vol 2, Chapter 4, J.E. Morris Editor, (Van Nostrand Reinhold, New York, 1991).

3 L. Thorsell, 'Mini DC/DC supplies simplify redundancy in parallel systems', EDN, April 28, 1988.

4 L. Thorsell, 'Will Distributed On-board DC/DC Converters become Economically Beneficial in Telecom Switching Equipment?', Proc. 12th Annual INTELEC Conference, 1990.

5 P. Lindman and L. Thorsell, 'Practical Considerations using Distributed Power Modules in Decentralized Power Systems', Proc.15th Annual INTELEC Conference, 1993.

6 D. Staffiere, ‘Power in the Year 2000’, (Proc. 10th annual IEEE Applied Power Electronics Conference,1995).

7 A. Fury, ‘Low Voltage ICs Force Distributed Power and then What?’, (Proc. 10th annual High Frequency Power Conference, 1995).

8 J. Flood, ‘Regulatory Changes in the European Union and the impact upon the US Power Supply Design Industry’, (Proc. 10th annual High Frequency Power Conference, 1995).

9 J. Prak, ‘Projection of Microprocessor Voltage and Power for the Year 2000’, (proc. 10th annual High Frequency Power Conference, 1995).

10 A. Rozman, K. Fellhoelter, ‘Circuit Considerations for Fast, Sensitive, Low-Voltage Loads in a Distributed Power System’, (Proc. 10th annual IEEE Applied Power Electronics Conference, 1995).

11 G. Suranyi, ‘The Value of Distributed Power’, (Proc. 10th annual Applied Power Electronics Conference, 1995).

12 C. Lee, R. Williams, ‘Power Supply Control ICs for Pentium Microprocessor Applications’, (Proc. 10th annual High Frequency Power Conference, 1995).

13 J. Burns, J. Riel, T. DiBene, ‘An Intelligent, Fault Tolerant, High Power, Distributed Power System for Massively Parallel Processing Computers’, (Proc. 9th annual IEEE Applied Power Electronics Conference, 1994).

14 F. Kunzinger and M. Jacobs, ‘Powering the next generation of electronics equipment with distributed architectures’, Powertechniques, Dec 1988.

15 C. Smith, J. Cassani, ‘Distributed Power Systems via ASICs using SMT’, Surface Mount Technology, Oct 1990.

16 S.D. Cogger, ‘Simplifying system power distribution with high power-density DC/DC converters’, Powertechniques, Sept 1988.

17 A. Bindra, ‘Toward distributing DC/DC power’; Electronic Engineering Times, March 26, 1990.

18 ‘The Economics of DC/DC Power Modules’, Ericsson Design Note.

19 C. Lazarovici, ‘Fault-tolerant power systems’, Powertechniques, Jan 1989.

20 W. Tabisz, M. Jovanovic, F. Lee, ‘Present and Future of Distributed Power Systems’, Proc. 7th Annual IEEE Applied Power Electronics Conference, 1992.

21 ‘FAST, Advanced Schottky TTL Logic Databook’, National Semiconductor, 1990.

Power Converter Technology

22 J. Kassakian, M. Schlect, G. Verghese, Principles of Power Electronics, (Addison Wesley, Reading, Mass, 1991).

23 M.H. Rashid, Power Electronics, (Prentice Hall, Englewood Cliffs, New Jersey, 1988).

24 A.I. Pressman, Switching and Linear Power Supply, Power Converter Design, (Hayden, Rochelle Park, New Jersey, 1977).

25 P. Wood, Switching Power Converters, (Van Nostrand Reinhold, New York, 1981).

26 E.R. Hnatek, Design of Solid-State Power Supplies, 2nd Edition, (Van Nostrand Reinhold, New York, 1981).

27 R.D. Middlebrook and S. Cuk, Advances in Switched Mode Power Conversion, Vol 1, (TESLAco, Pasadena, California, 1981).

28 R.D. Middlebrook and S. Cuk, Advances in Switched Mode Power Conversion, Vol 2, (TESLAco, Pasadena, California, 1981).

29 I. Jitaru, ‘The Impact of Low Output Voltage Requirements on Power Converters’, (Proc. 10th annual high Frequency Power Conference, 1995).

30 S. Delaney, ‘The Impact of the Next Generation ICs on Power Conversion Semiconductors’, (Proc. 10th annual High Frequency Power Conference, 1995).

31 G. Mulcahy, ‘PWMs versus resonant converters’, Electronic Engineering Times, March 26, 1990.

32 E. de la Cruz, S. Ollero, J. Rodriguez, J. Uceda, J. Cobos, ‘Review of Suitable Topologies for On Board DC/DC Converters in Distributed Power Architectures for Telecom Applications’, Proc 8th Annual IEEE Applied Power Electronics Conference, 1993.

High Frequency DC/DC Converters

33 P. Lindman, L. Thorsell, 'Applying Distributed Power Modules in Telecom Systems', (Proc. 9th annual IEEE Applied Power Electronics Conference, 1994).

34 R. White, 'How to Select DC/DC Converters', (Proc. 10th annual IEEE Applied Power Electronics Conference, 1995).

35 P. O'Farrell, 'Market trends in power supplies', Electronics Industry, May 1985.

36 D. Haynes, 'DC/DC converters are getting smaller with higher performance and lower cost', PCIM, Mar 1990.

37 T. Harbert, 'DC/DC converters work in varied environments', EDN Product News, Oct 1986.

38 J. H. Mayer, 'Compact DC/DC converters serve distributed power applications', Computer Design, Nov 1, 1988.

39 E. Marchese, 'High-density converters', Powertechniques, March 1990.

Thermal Management

40 M.E. Jacobs, 'Practical limits of forced-air cooling', Powertechniques, Feb 1990.

41 A. Beane, 'Change...The New Paradigm of Thermal Management', (Proc. 10th annual High Frequency Power Conference, 1995).

42 R. Hock, L. Pitzele, 'Comparison of Usable Power Density in Board Mounted DC/DC Converters', (Proc. 9th annual IEEE Applied Power Electronics Conference, 1994).

43 R.W. Johnson, R.Weeks, D.C. Hopkins, J. Muir and J.R. Williams, 'Plated copper on ceramic substrates for power hybrid circuits', IEEE Trans. Components, Hybrids & Manuf. Technol., 12 (4), Dec 1989.

44 H. Fick, 'Thermal management of surface mount power devices', PCIM, Aug 1987.

45 H. Fick, A. Pshaenich and D. Hollander, 'Metalbacked boards improve performance of power semis', PCIM, Sept 1989.

46 R. Ridley, M. Reynell, S. Kern, 'Thermal Simulation Cuts Power Supply Design Time', PCIM, Nov 1992.

47 R. Ridley, M. Reynell, S. Kern, 'Thermal Considerations for Distributed Power Converter Systems', Proc. 8th Annual IEEE Applied Power Electronics Conference, 1993.

Reliability

48 L. Thorsell, P. Lindman, 'Reliability Analysis of a Direct Parallel Connected n+1 Redundant Power System Based on Highly Reliable DC/DC Modules', Proc. 10th Annual IEEE INTELEC Conference, 1988.

49 P. Lindman, C. Smith, 'Reliability is a Matter of Course', (Electronic Engineering Times, 25 October, 1993).

50 B. Jacobson, J. Newbold, 'Reliability of Fault-Tolerant Distributed Power Systems', (Proc. 10th annual High Frequency Power Conference, 1995).

51 P. Hunter, T. Johnson, 'Fault-Tolerant, Hot-Pluggable Power System Design', (Proc. 9th annual IEEE Applied Power Electronics Conference, 1994).

52 P. Rostek, 'Power System Design for Massive Parallel Computer Systems', (Proc. 9th annual IEEE Applied Power Electronics Conference, 1994).

53 V. Kane, 'Calculation of Test Time Required to Remove Infant Mortality Failures', Proc. 7th Annual IEEE Applied Power Electronics Conference, 1992.

54 'Reliability Aspects on Power Supplies', Ericsson Components, Power Modules, Design Note.

Ericsson is shaping the future of Mobile and Broadband Internet communications through its continuous technology leadership. Providing innovative solutions in more than 140 countries, Ericsson is helping to create the most powerful communication companies in the world.

Ericsson Microelectronics is a supplier of world-class microelectronic components for wireless applications, broadband communications, fixed access and communications via fiberoptics as well as solutions for distributed power architectures. With its global design, development, manufacturing and sales network Ericsson Microelectronics is a leading supplier of technology solutions to meet the customer demand for increased speed, bandwidth and mobility.

Ericsson Microelectronics
SE-164 81 Kista, Sweden
Telephone: +46 8 757 50 00
www.ericsson.com/microelectronics
For local sales contacts, please refer to our website
or call: Int + 46 8 757 47 00, Fax: +46 8 757 47 76

Powerbook
EN/LZT 146 116 R1A
(Replaces EN/LZT 137 12 R1