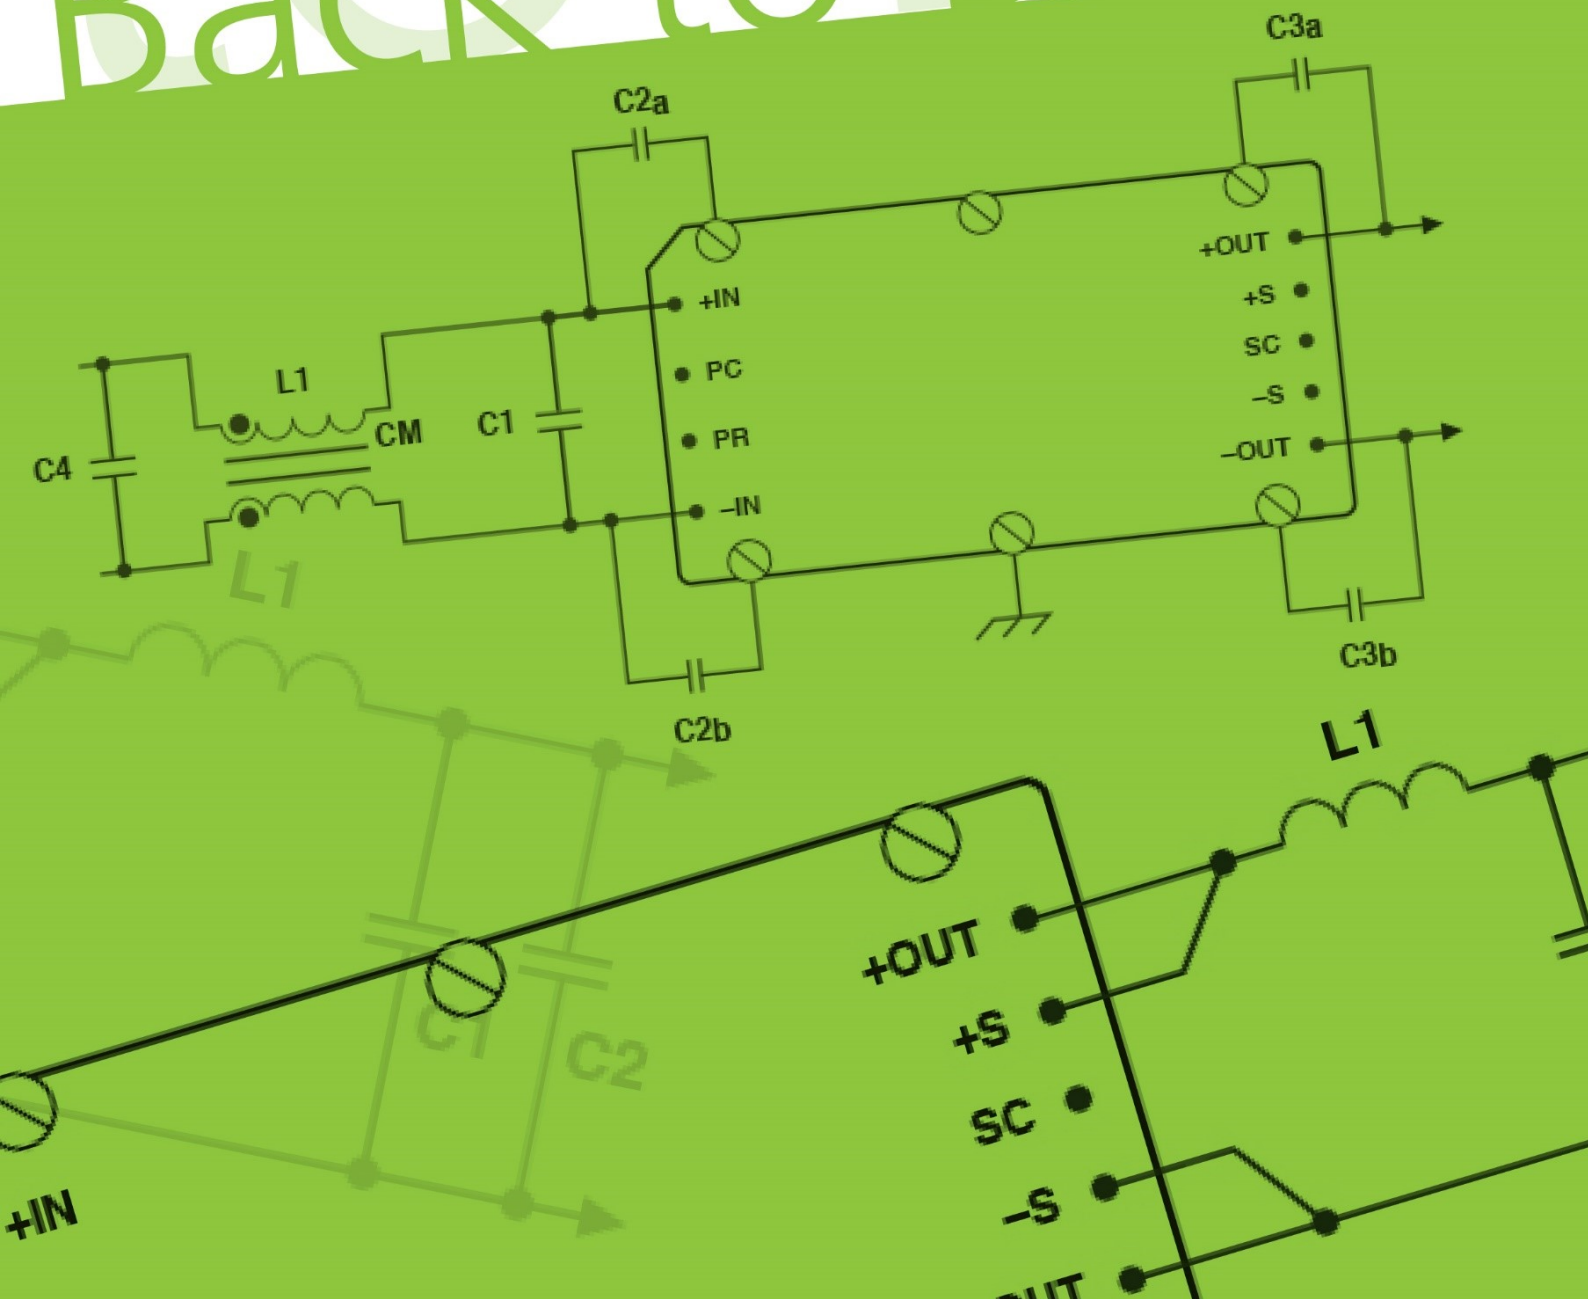


Back to Basics





Back to Basics: An Anthology of Blog Posts

Like most engineering, power system design requires a strong understanding of the fundamentals. A couple of years ago, we began a series of blog posts entitled "Back to Basics", which have become the most popular posts on powerblog.vicorpower.com. Readers range from graduates who are just learning about power, to experienced engineers who want to brush up on a particular topic.


This eBook is a compilation of these Back to Basics posts. It features articles on a range of important topics, such as filtering, design for hot-swap, PFC, EMI suppression and dealing with transients.

We've left the posts in reverse chronological order: the stories haven't been sorted into categories. It's not our intention that you read the book from cover to cover, but rather that you just pick the topics that you find the most interesting (or challenging).

We hope you enjoy reading this book, and are sure that you'll find it both interesting and useful.



Contents



Using Hot Swap to Minimize Downtime	1
Using Adaptive Loop Feedback to Regulate Your Converter Output	4
Using Remote Sensing to Ensure Accurate Point of Load Voltage Control	7
Do I Need Filtering on the Input and Output Side of the Converters?	10
How Do I Ensure Correct Power Rail Sequencing in a Multi-Rail Power Architecture?	14
Why Do We Need Active Surge Suppression?	17
The Importance of Switching Frequency	19
What is Active ORing?	22
Paralleling Power Components for Current Sharing	24
What are Y-Capacitors?	27
Zero Voltage Switching - What it is and Why it's Important for Buck Voltage Regulation	29
What Does Power Factor Mean and Why Must We Correct it?	31
Understanding and Mitigating the Growing Problem of Distribution Losses	36
Handling High Input Transients	41
Meeting EMI for AC-DC Systems	43

Using Hot Swap to Minimize Downtime



Given the critical role of modern telecommunications and data communications equipment, maximizing its online availability must always be the first priority. Clearly, any equipment's availability depends critically on that of its power supply. Fortunately, close to zero power system downtime can be effectively achieved by combining two techniques – redundancy and hot-swapping. In this back-to-basics post we will see how these work together, and how systems can be designed for hot swapping.

Consider for example a rack-mounted telecommunications system comprising modules plugged into the rack's backplane, with a total power requirement of 120 W DC. This could be met in many ways. If, for example, two 60 W power converter units, also mounted on the backplane were to be used a failure of either unit will cause significant downtime while the faulty unit awaits replacement. A better solution is to share the load between three 60 W converters, so that uninterrupted operation continues even if one unit fails. This is known as N+M redundancy, where N units satisfy the load and M provide redundancy.

However, to benefit fully from the uptime 'bought' by the redundant converter, the system must inform an operator about the failed unit, then allow them to replace it without shutting down or otherwise interrupting the power supply. Adding hot swap circuitry to the converter modules facilitates this by protecting the converters and the rest of the system from the problems associated with live insertion and withdrawal.

Removing the faulty unit from the rack is relatively easy, provided the remaining power units can support the step increase in load that will occur if the faulty unit continues to deliver current. Plugging in the replacement unit has more potential for problems, as it will present an uncharged capacitor load and draw a large inrush current. This could cause a momentary, but unacceptable, interruption or sag to the backplane power bus if not limited. Problems can also arise if ordinary power module connectors are used, since the connector pins will engage and disengage in a random and unpredictable sequence during insertion and withdrawal.




One relatively simple – but very manual, and therefore prone to human error – method is known as ‘warm swapping’. This involves plugging a power module into the backplane and then applying the AC or 48 VDC power input, so the power module rather than the supplied load absorbs the step change and inrush current. Problems with random pin insertion sequences are also avoided.

A more automated, true hot swap approach uses an active MOSFET switching device in the power line, and a connector with staggered pins. During insertion, the MOSFET is driven into a resistive state to limit the inrush current; then when the inserted module’s capacitor has charged, the MOSFET becomes fully conductive to avoid voltage drop losses. The staggered pin connector can be configured with the Ground pin longer than all others, to ensure the module remains safely grounded and protected from random pin connect/disconnect sequences during insertion and withdrawal.

Fig. 1 below shows how such a scheme can be implemented using our Picor PI2211 Hot Swap Controller and Circuit Breaker. It demonstrates how the device can be used with an external N-channel MOSFET. The circuit employs the MOSFET’s transient thermal characteristics, to ensure that it operates within its dynamic safe operating area (SOA). Emulation and protection based on the specific MOSFET’s transient thermal performance optimizes the safe operating limits and allows designers to take advantage of the latest power MOSFET technologies. The PI2211 has an operating range of +0.9 to +14 V, and is designed to work with a specific set of Vicor converters (appropriate parts can be selected using the PowerBench™ Solution Selector tool).



Fig. 1



An alternative approach, using Vicor's Input Attenuator Module (IAM48) or now more usually Filtered IAM48 (FIAM48), is available for Advanced TCA (ATCA) and Micro ATCA telecommunications, data communications and other 48 VDC systems. Like the PI2211 example, it uses a staggered pin connector, but configured with a short On/Off pin. The FIAM also provides inrush current limiting and transient overvoltage protection. This part is generally recommended for telecom applications to ensure compliance with EMC (Electromagnetic Compatibility) standards.

Vicor provides other integrated hot swap solutions that comply with Advanced TCA PICMG 3.0 requirements for hot insertion and board level conducted noise limitations. One of these is the QPI-8 QuietPower hot swap controller with active EMI filter, which is designed for applications such as ATCA blades, telecom base stations, IBA and distributed power, network switches and routers, optical line-in cards and TD-SCMA wireless infrastructure installations. And the QuietPower QPI-10 Advanced TCA hot swap SiP with EMI filter, which is designed for use with telecom and ATCA PICMG 3.0 boards using VI Chip technology.

Overall, Vicor offers different solutions for different hot swap applications, but the objective always remains the same; to give systems designers the opportunity to build hot swap capability into redundant power module arrays. This allows telecoms and other mission-critical applications to continue operation without interruption even through the failure and replacement of one or possibly more of their power modules.

Using Adaptive Loop Feedback to Regulate Your Converter Output

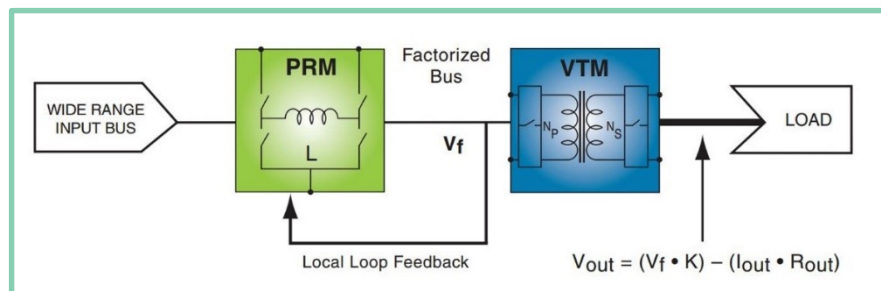


A conventional power supply or regulated DC-DC converter typically uses Remote Sense to achieve true regulation accuracy. For a system designer using, say, a brick converter the task is relatively simple; just connect the + and – sense terminals of the brick to either side of the load, and the brick’s internal circuitry will ensure correct regulation. The sense lines allow the brick to compensate for its own output resistance, as well as the resistance of the power lines and connector pins between its output and the load. With this compensation, the correct voltage is referred to the point of load rather than the brick output terminals.

The brick, like all complete power components, will include an isolation stage; this creates design challenges as the feedback loop must also be isolated. However these issues are solved within the power unit’s design and are therefore transparent to the systems designer

Today, however, the efficiency requirements, transient response times and competition for PCB real estate typical of today’s electronic equipment mean that many engineers are turning to more sophisticated power distribution architectures such as the Factorised Power Architecture (FPA). FPA employs integrated power components such as a PRM feeding one or two VTMs to achieve these expectations of density, efficiency and responsiveness. Fig. 1 is a generic example of a PRM/VTM arrangement.

Fig. 1: A simple FPA implementation with PRM and VTM



The VTM can be considered as a fixed-ratio isolated DC-DC transformer which offers high power density and efficiency, and low output impedance enabling fast transient response. The PRM provides a regulated output – a ‘factorised bus’ – from an



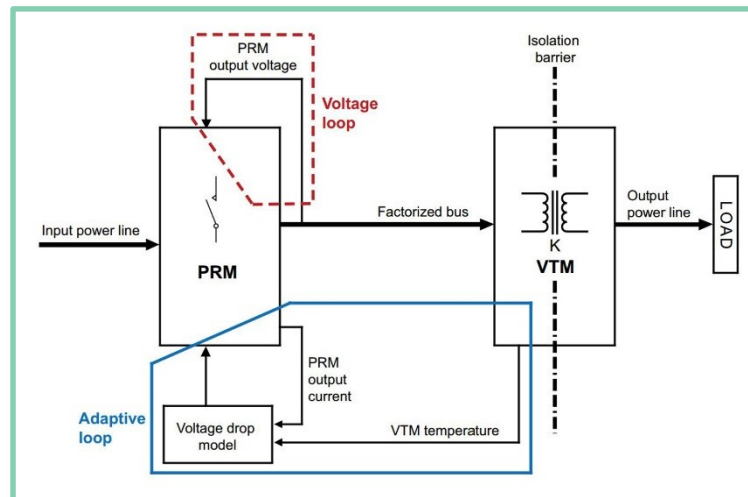
unregulated input source. This PRM-VTM combination creates a complete isolated, regulated DC-DC converter.

Note that Fig. 1 also shows a third wire: the Adaptive Loop connection between the PRM and VTM. This is because the 'intelligent' PRM handles the regulation, yet it cannot directly 'see' the load for which it regulates, because the VTM with its isolation barrier lies in the way. An isolated feedback loop could be designed, but this would be a time-consuming task for systems engineers, with issues of accuracy and tolerance arising. Instead, by providing the adaptive loop feedback, the VTM acts as an eye for the PRM, communicating the information that it needs for accurate regulation.

This information relates to the resistance of the VTM's output as well as the smaller resistance values of the power lines and connector pins between the VTM and load. In theory, these resistance values should be zero, but in practice they are not. Any such resistance will cause a voltage drop which is 'stolen' from the voltage appearing across the load. The PRM can calculate the voltage drop across the VTM output terminals if it knows the current flowing through them, and their resistance, using the simple formula $V = IR$.

Fig. 2 below shows the concept of the adaptive loop (AL). While the voltage loop, also shown, maintains regulation at the PRM output, the AL compensates for the voltage drops that occur from the PRM output to the actual load. The AL is centred on a voltage drop model which presents the PRM with a scaled replica of the total voltage drop in the system. The model is fed directly by current drawn from the PRM output. This can be scaled to represent the VTM's output current by using the VTM's K transfer ratio as the scaling factor. The VTM's third wire feedback contains information about its output resistance, so the model can multiply this value with the scaled PRM current value to obtain the system voltage drop representation.

Fig. 2: Adaptive loop regulation conceptual diagram



Note that the feedback link is referred to as 'VTM temperature'. This is because the resistance across the VTM's output terminals will vary with temperature, so the feedback signal must accordingly also include temperature variance. Such variance is facilitated by a PTC resistor embedded within the VTM module.

The AL is a simple circuit with few components, and avoids the need to transmit signals across the VTM's isolation barrier. Reference to a more detailed VICOR PowerBench Application Note will show that the AL includes a few resistors whose values can be calculated by systems engineers to meet the requirements of the VTM they choose. The App Note includes a Table of VTM part numbers and their associated output resistance and temperature coefficient values. It also provides a set of equations; these are for use in calculating the resistor values necessary to build a voltage drop model accurately matched to the VTM selected.

The App Note covers half-chip as well as full-chip VTM devices. The theory provided is supported with a design example using VI Chip Customer Boards.

Adaptive loop control provides improved regulation over simple local loop control – within +/- 1% – yet requires only a simple, non-isolated, feedback connection between the VTM and PRM.

Using Remote Sensing to Ensure Accurate Point of Load Voltage Control




Today's ASICs and processors typically operate on a low voltage and therefore require high current. This means that if they are placed at any distance from the power supply or DC-DC converter that feeds them, transmission losses can become an issue. Because of resistance and therefore voltage drop in the power path between the converter and load, the voltage across the load will be lower than that across the converter's output terminals. The difference is given by $I \times R$, where I is the current drawn by the load and R is the resistance of the power path.

On a PCB, the power path is usually provided by the PCB's power plane, so R is relatively small. However, as mentioned, I can be large so the voltage drop between the converter and load can be significant. This is even more important in a low voltage/high current environment, because a loss of say 0.5 V is a much higher proportion of a 3 V target voltage than it is of 48 V.

Remote sense, a popular and effective way of overcoming these issues, is used by many power product manufacturers. Most Vicor products have provision for this; the only exceptions are the VI Chip® family, which uses Adaptive Loop Control and the Micro Brick family. Given their power levels, the Micro Brick (Quarter Brick) products simply do not require remote sense for most applications. For the rare cases where it is required, an Application Note (see pages 9 and 10 of this Applications Manual) describing implementation of external remote sensing for these products is available.

Fig. 1 shows the remote sense concept. The S+ and S- Sense wires are connected across the load at the points where the correct voltage should be applied. Note that this is a 'remote sense' configuration; as we shall see, 'local sense' arrangements also have their role to play.

Fig. 1: Four-wire remote voltage sensing on a DC-DC converter



S+ and S- are high-impedance inputs, so negligible current is drawn through the Sense lines; voltage drop is therefore negligible. This means that the Sense pins provide the converter with an accurate measurement of the voltage across the load resistance R_L . By contrast, until modified by remote sense control, the voltage at Vout+ and Vout- will be lower than the voltage across the load by an amount $I \times R_P$ where I is the current drawn by the load and R_P is the resistance of the power path. The converter can use this difference between the S and Vout voltage levels to accurately compensate for the voltage drop.

Some applications may not require the remote sense feature. If so, the remote sense lines should not simply be left open, as the power supply could drive upwards to an unwanted and possibly destructive level. Instead, the power supply should be configured in 'local sense' mode, with a low-value resistor connected between the S+ and Vout+ terminals, and another between S- and Vout-.

Westcor products make this particularly easy and reliable to set up, as they have 5 Ω resistors pre-installed in parallel with the remote sense circuits. They also have Westcor's Autosense capability, which allows them to use the remote sense lines if connected, or automatically switch to the 5 Ω resistors if not. This feature protects the Westcor products against open loop mode failure as well as simplifying configurations with no remote sense connections. Similar provision can be made for other Vicor products by using externally-connected low-value resistors.

It should be remembered that remote sense wires are part of the power supply's control loop, so any radiated noise they pick up can introduce instability into the power system. This can be mitigated by using light-gauge wiring – which is viable as the sense lines draw minimal current – and twisting the S+ and S- wires together in addition to shielding them.



Although remote sense allows loads to receive correct voltage levels, it should not be considered as a substitute for good power layout practice. If insufficient care is used in minimising the power path voltage drop, this drop could exceed the compensation range of the power source and make full compensation impossible. In any case, excessive power path voltage drops are to be avoided; for example a drop of 0.5 V in a path carrying 100 A current equates to 50 W losses with associated implications for efficiency and cooling. These issues mean that investment in adequate gauge cabling and power planes and good-quality connectors is easily justified.

In using remote sense, care must be taken to ensure that the power supply is not driven up into an over-voltage protection (OVP) region; this would cause the power supply to shut down, denying power to the load. It is also important that the Vout+ or – lines are not opened while power is applied. This would cause current to flow in the Sense lines. With good design, however, using remote sense is an invaluable technique that ensures the voltage at the load is exactly that value that was intended.

Do I Need Filtering on the Input and Output Side of the Converters?



All switching power supplies generate potentially interfering signals as a result of their high frequency, high power switching. This externally-generated noise can be a threat to any connected electronics PCB, although the source may be outside the control of the PCB's designer. Adequate protection must therefore be provided on the board itself.

The board's power conversion modules are first in line to receive the incoming power. As with all electronic devices, they need protection from high levels of conducted or radiated emissions – and this protection requirement applies equally to the components served by the power converters. Although this implies that the power converter potentially requires both input and output filters, an output filter is seldom actually required for Vicor devices. This is because the Vicor design, to a large extent, addresses the problem by the use of a quasi-resonant, zero-current switching (ZCS) and zero-voltage switching (ZVS) topology. The switching current waveform is a half sine wave that generates far less conducted and radiated noise in terms of both frequency spectrum and magnitude.

A BCM bus converter, for example, generates output noise of less than 1% of its output voltage without any external filtering capacitors. Its 3.5 MHz soft-switching exploits nominal distribution inductance associated with board interconnects, together with small ceramic capacitors at the point of load, to attenuate ripple to below 0.1% at the load. Accordingly, soft switching technology facilitates system-level EMI filtering with reduced complexity, bulk and cost.


There are, however, occasions when the requirements of a particularly demanding application mean that even Vicor converters require external output filtering. These can arise in systems processing low-level signals; for example applications that use ultrasound or other sensors. Switched mode-based solutions for these must achieve the same levels of ripple and transient suppression that a linear PSU alternative could reach.

The LC filter design in Fig. 1 is a comparatively simple solution for reducing ripple on the outputs of converter modules. These components are small and provide significant peak-to-peak noise attenuation. Together with the DC-DC converter's internal output filter capacitor, the external inductor and capacitors make up a



Pi filter. If even greater attenuation of output ripple or additional AC power line ripple frequency rejection is required, active output filters or ripple attenuation modules provide a higher-performance solution.

Fig. 1: LC Output filter example



By contrast to output filters, input filters are often required. They serve to both eliminate noise from the supply, and prevent noise being fed back into the supply. Although Vicor ZCS converters are 20 dB to 40 dB lower in conducted noise than a traditional board-mounted PWM converter, if a specific EMI specification such as FCC or VDE must be met, additional filtering may become necessary. EMI filtering, if properly designed and implemented, can reduce the magnitude of conducted noise by 40 – 60 dB; this in turn reduces the noise radiated by the conductors proportionally.

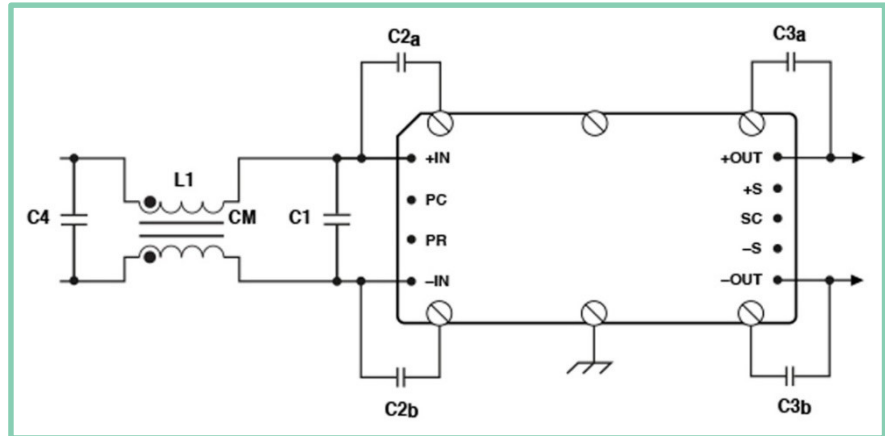
Conducted noise on the input power line can appear as either common-mode or differential-mode noise currents. Common-mode noise, which has mostly high-frequency content, is measured between the converter's input conductors and ground. Differential-mode noise is largely at low frequencies and appears across the converter's input conductors. It comprises the fundamental switching frequency and its harmonics.

The most effective way to reduce common-mode current is to bypass both input leads to the converter baseplate using Y-capacitors (C2) as shown in Fig. 2. The leads should be short to



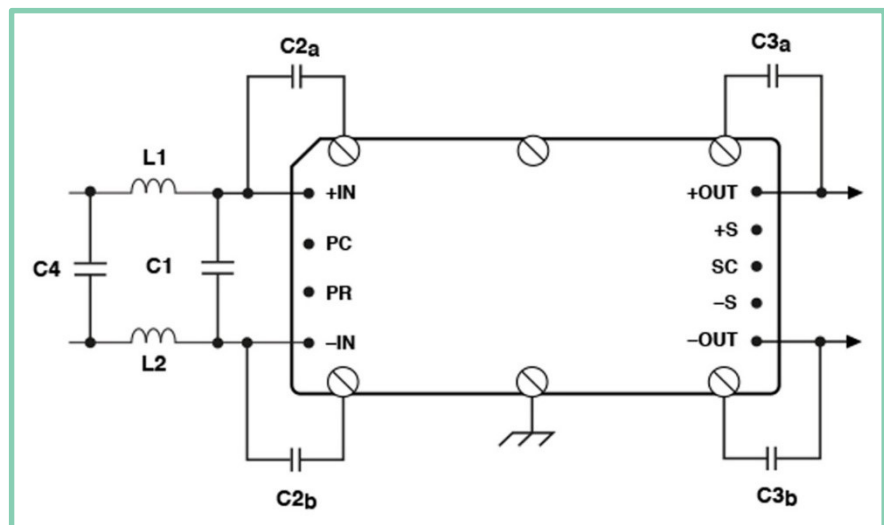
minimize parasitic inductance. Additionally, a common-mode choke (L1) is usually required to meet FCC/VDE A or B.

Fig. 2: EMC common-mode input filter example



All Vicor converters have an internal differential-mode LC filter which, in conjunction with a small external capacitor C1 as shown in Fig. 3 reduces differential-mode conducted noise. The external capacitor should be placed close to the module to reduce loop cross-sectional area of differential-mode current flowing between the supply and C1. Twisting the input leads can cause noise cancellation, and PCB power planes can reduce radiated noise if the traces are on opposite sides of the PCB directly over one another.

Fig. 3: Differential-mode filtering example





In general the common-mode choke is the inductor of choice for EMI filters, as it offers better attenuation. In some cases, such as military applications, both common-mode and differential-mode filters are required; if so they can be used one after another in front of the converter. For more demanding applications, active input filters provide greater attenuation of both common-mode and differential-mode noise.

Vicor's ZCS topology minimizes discontinuities in the switched current waveforms, reducing conducted and radiated noise. This means that in some circumstances, Vicor converters can be used without input or output filtering. Whilst an output filter is usually only required in particularly demanding applications, an input filter is advisable or essential for many applications.

How Do I Ensure Correct Power Rail Sequencing in a Multi-Rail Power Architecture?



For a graphic illustration of why power sequencing is essential, consider what could happen in a large industrial system if its motors and actuators became energized before its control electronics. Clearly, the potential for such unpredictable and dangerous scenarios must be eliminated with carefully-designed power start-up sequences. And shutdown sequences are equally as important.

Incorrect power sequencing can cause problems at the semiconductor as well as at the system level: these conditions could lead to unexpected reverse bias across a CMOS device. This can damage or destroy the device, or latch it into an ON state that can only be released by power cycling. Although blocking diodes can prevent latch-up, they can limit the usable analogue input voltage range, so supply voltage start-up sequencing is the preferred solution. Additionally, multi-voltage devices such as DSPs and microprocessors usually require their I/O voltage to be present before applying the core voltage. Intel processor chip documentation typically specifies start-up and shut-down sequences that must be followed. Similarly, CPU chips must be powered up before systems graphics ICs to prevent uncontrolled outputs to the graphics display.

There are two typical approaches to power sequencing; one based on power and the other on logic.

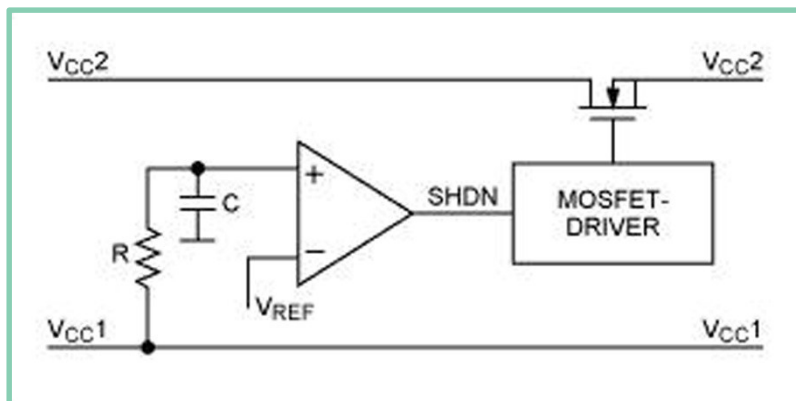
The Power solution uses the voltage rise in one power rail to send a suitably delayed trigger signal to an N-channel MOSFET, which controls the supply to the second power rail. The Logic approach is based on using Enable or other logic pins on power train devices rather than controlling the power to them. Using the logic integrated in power components provides a lower-cost, more compact and simpler solution, but in more demanding applications, MOSFETs permit more accurate control of turn-on time.

Fig. 1 below shows a basic MOSFET sequencing circuit using an N-channel MOSFET and a comparator. VREF sets the threshold for switching VCC2, with VCC2 turning on after VCC1 with a delay set by the RC network. The MOSFET driver includes a charge pump to ensure that VGS exceeds VCC2 by several volts; this fully enhances the MOSFET switch. This circuit also ensures that VCC2 shuts down



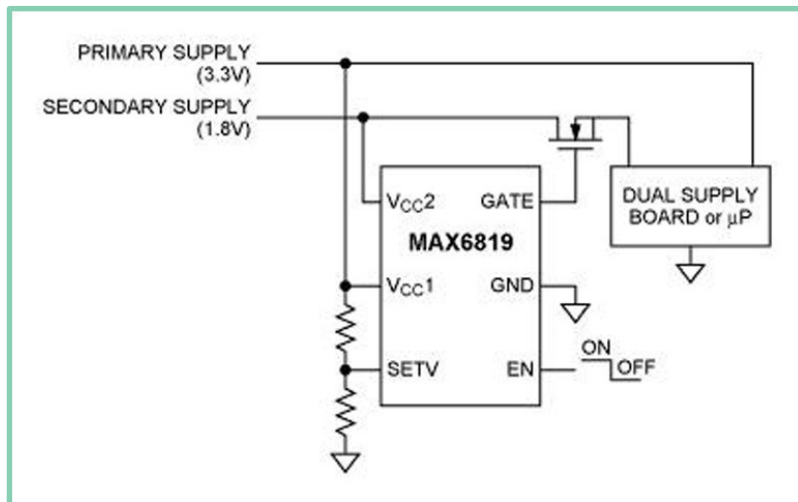
when VCC1 does, however the RC network means that VCC2's shutdown will lag VCC1's. Another issue can arise if VCC1 is not available for use as a reference voltage.

Fig. 1 – Power sequencing using an RC network, comparator and MOSFET driver



More highly integrated and easier solutions are available, such as the example shown below in Fig. 2.

Fig. 2 – Power sequencing with the MAX6819



The sequencer ensures that the MOSFET always has the minimum required VGS enhancement, which minimizes loss in the MOSFET and ensures a low drain-to-source impedance (RDS(ON)). The MAX6819 imposes a factory-set delay time of 200 ms, which occurs after the primary voltage rises above the set threshold and before enabling the charge pump to drive the external MOSFET switch.



Vicor's voltage regulator, VI Chip and Brick products can be sequenced by the MOSFET circuits above. Alternatively, they can often be managed through their control pins and internal logic as mentioned earlier. In one simple arrangement, a 'soft start' output from one device can be connected to the logic Enable/Disable pin of another, effectively controlling its power sequencing. Additionally, many Vicor devices have a Voltage Control pin, or trigger. If a ramp voltage is applied to this pin, the device's output will ramp to track the Voltage Control input, but with a fixed time delay.

Why Do We Need Active Surge Suppression?



It's well understood that unexpected electrical disturbances can cause data loss or even permanent damage to any unprotected electronic equipment in their path. Active as well as passive suppression is essential to provide comprehensive protection from all such events. This post shows why this is so, and how active suppression works to fulfil its role.

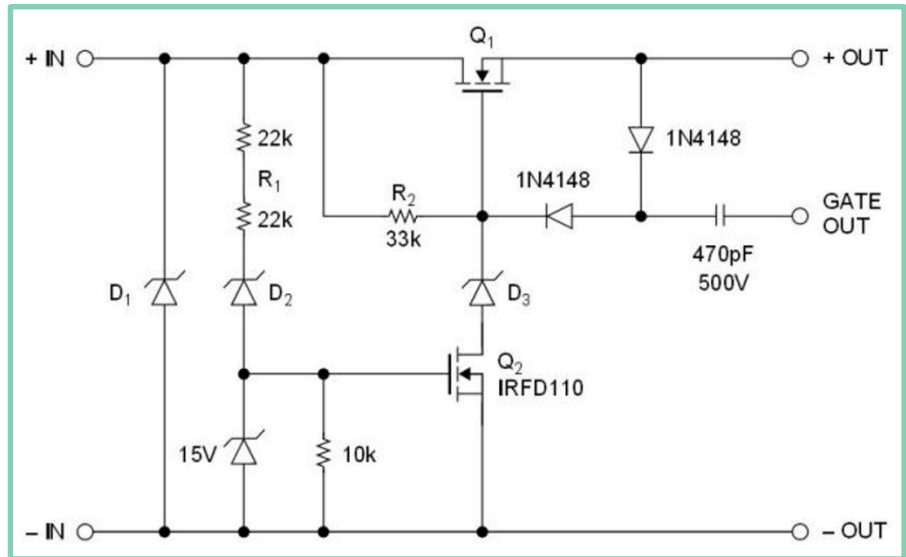
Electrical events threaten electronic equipment by placing excessive electrical energy onto its input power rails in the form of a transient, a surge, or possibly both. A transient is typified as a fast-rising voltage spike of high magnitude – up to, say, 1800 V but of limited duration, typically around 50 μ s. As this short duration limits their energy, they can be successfully managed with a passive Transient Voltage Suppressor (TVS) such as a Zener diode, MOV or varistor. With a clamping voltage chosen to be below the high line voltage or transient limits of the input power converter, this shunts the higher voltage energy away from the protected device. A TVS able to withstand 1.5 J is suitable for many applications.

Transients are caused by events such as lightning strikes, and start-up or shut down of inductive loads close to the protected equipment. Surges are created by train braking, or by fuse clearing. Often found in telecoms applications, this occurs after a fuse has short-circuited, allowing buildup of a large inductive field. When the fuse is cleared (opened), the field collapses, inducing a surge in the adjacent cabling. Similar load dumping can occur in vehicles, if a battery becomes disconnected during high engine speeds. 100 ms or more can elapse before the alternator regains control.

Voltage surges are very different from transients. Although measured in tens rather than hundreds or thousands of volts, they can last for 20 ms or more – and originate from a source impedance as low as 0.2 Ω . This is too much energy for a TVS to handle, so an active suppression circuit must be provided. A simple example is shown in Fig. 1.



Fig. 1: An active circuit to prevent fast transients and surges



Although called an active suppression circuit, its first stage, D1, actually provides a passive suppression function by clamping fast, high voltage spikes. The active circuit is formed around Q1, D2 and D3. MOSFET Q1, which is in series with the critical load, is controlled to behave like a linear regulator during a voltage surge. Protection is provided without need for excessive power dissipation while a known current is passed to the load. Planning and sizing become manageable.

Q1 is kept in full conduction during normal operation by the charge pump circuit comprising the 1N4148 diodes and 470 pF capacitor. The Gate Out signal from the equipment's converter drives the charge pump. When overvoltage occurs, D2 conducts, with current limited by R1; Q2 is turned on. With Q2 in conduction, the gate voltage of Q1 is held to the clamping voltage of D3. The output voltage equals this level minus the gate source voltage of Q1. When overvoltage ceases, the system resumes normal operation.

The FIAM family devices are available as completely integrated packages, including the power MOSFET. FIAMs can be used in MIL, Rail, Telecom and other 48 V power environments. Vicor can also provide advice and reference designs for DEFSTAN and other highly stringent standards where discrete component solutions become essential.

The Importance of Switching Frequency



The importance of switching frequency quickly becomes apparent to systems designers bringing regulated power to on-board semiconductor devices. With the continual demand for higher performance, CPUs, DSPs and other such devices grow more power-hungry. Accordingly, higher power density regulators with minimal PCB footprints have evolved; these use the latest in IC integration, MOSFETs and packaging. However, even these struggle keep pace with the continuing stream of newer, more powerful devices.

With these pressures, it is tempting to increase regulator frequency, as this reduces the size and board footprint requirements of the associated passive devices – inductors, capacitors and resistors. Conventional thinking, based upon classic hard-switching PWM regulators, is that as frequency increases, then so do switching losses. This is because, in these topologies, regulator MOSFETs incur losses every time they switch, so a higher switching frequency leads directly to higher losses. These inefficiencies are primarily due to high side losses during turn on, Miller gate charge and body diode conduction losses. To make things worse, conventional topologies further magnify the losses as higher input voltages are converted or regulated.

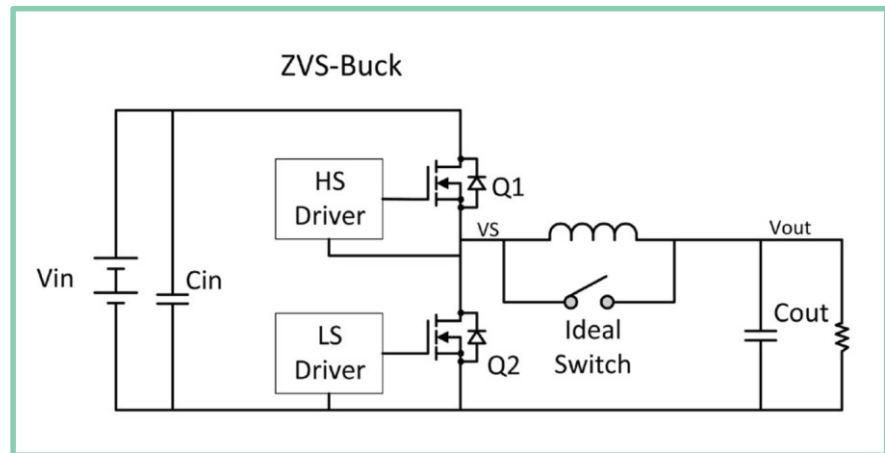
These losses introduce a practical limit for the switching frequency of conventional converters and regulators. There is, however, a solution: devices that use the Zero-Voltage Switching (ZVS) topology do not suffer from losses in the same way as conventional designs, allowing them to operate at higher frequencies, which in turn improves performance and dramatically reduces the size of external filter components. Converters and regulators using ZVS also reduce the additional losses associated with large step-down ratios in PWM topologies.

Unlike conventional regulators that rely on hard switching topology, ZVS uses soft switching; this accounts for its improved efficiency and higher density performance. For example, the Picor PI33xx series operates up to 1.5 MHz and beyond, which is typically 2x to 3x that of conventional high density regulators. Higher frequency operation not only reduces the size of passive components but also reduces the burden on external filtering components and allows for fast dynamic response to line and load transients.



Figure 1 shows the schematic diagram for the ZVS buck technology. Schematically, it is identical to the conventional buck regulators except for an added clamp switch that connects across the output inductor. This switch is added to allow energy stored in this inductor to be used for Zero Voltage Switching.

Fig. 1



Using ZVS addresses the high turn-on losses of the conventional regulator by eliminating high current body diode conduction prior to turn on of the high side MOSFET, bringing the D-S voltage of this MOSFET to zero or nearly zero and producing no high current spikes or damaging ringing. The ZVS action applied to $Q1$ removes the Miller effect at turn on of $Q1$, allowing the use of a smaller driver and lower gate drive at turn on.

Regulation from a higher voltage, at a higher efficiency, and in a smaller form factor is realisable with implementation of an improved switching topology. By utilizing the ZVS topology within the Vicor's Picor PI33XX, the company has developed a buck regulator that offers high performance regulation up to $36 V_{in}$, exceeding the performance found in conventional hard switching, high density regulators. Its high switching frequency allows the external output inductor to be very small, with a total solution size of $25 \times 21.5 \text{ mm}$; yet it produces up to 120 W of output power with a peak efficiency of 98%. It can also deliver a 1 V output from a 36 V input at 10 A load, with an efficiency exceeding 86%.




The same underlying ZVS technology delivers high performance in many other Vicor products, including our AC-DC and DC-DC converters, PRM buck/boost regulators, and the VTM current multipliers.

What is Active ORing?



The uninterrupted availability of server, communications and telecoms equipment is frequently critical to its client applications, so this equipment typically uses two or more power sources in a redundant power architecture. These sources need mutual protection, otherwise a short-circuit fault in one could quickly overload and damage the others.



The diagram above shows a simple solution using diodes. Each diode allows current to flow in a forward direction only, while preventing either supply from drawing short-circuit current. Therefore the system can continue to function if one supply fails. Although simple and fast, this arrangement has a drawback due to the diodes' high forward voltage in their normal state of forward conduction. This creates high power and heat dissipation and an unwelcome need for thermal management and extra board space. As systems have increased in power and component density, these problems have become unacceptable.

Active ORing offers a better alternative; it comprises a power MOSFET and controller IC. The MOSFET has an on-state resistance $R_{DS(ON)}$ which, multiplied by the square of the current through the device, creates an internal power loss. However this loss can be substantially lower than that of a Schottky diode for the same current; a ten times efficiency improvement is typically achieved.

For example, for a 20 A application, a Zener diode with a .45 V forward voltage drop would dissipate 9 Watts of power, while a MOSFET with a 2 mOHM $R_{DS(ON)}$ on would dissipate only 0.8 W, a more than 10x reduction in power loss.

Yet an active ORing solution requires care in its set-up. A MOSFET passes current in either direction when turned on, allowing the possibility of a large reverse current if the protected power source



has a short-circuit failure. This will bring down the system if it persists for long enough. Therefore, the active ORing solution must be very accurate, and able to detect reverse current fault conditions extremely fast.

When a fault is detected, the controller must turn the MOSFET off as fast as possible, isolating the input fault from the rest of the system.

The controller IC senses information across the MOSFET to determine the magnitude and polarity of current flowing through it. Transgression of its reverse current threshold indicates a power source failure; therefore sensing of this must be tightly accurate to provide the consistent, fast fault detection required. Speed of response is also critical, as this determines the magnitude and duration of reverse currents. Higher peak reverse currents increase reliability concerns, leading to larger MOSFETs, higher costs and more demand for real estate. These issues can be exacerbated by some of today's low-impedance power architectures. Conversely, tight, fast control can mean the difference between a damaged system and a safe disconnect of a failed power supply.

To summarise, we can say that the best Active ORing solutions comprise an ORing controller that is fast and accurate in detection and fast in response, together with a MOSFET offering the lowest possible $R_{DS(ON)}$. These attributes allow minimised losses and dependence on thermal management, with smaller size and greater ease of use.

Vicor offers a complete family of Active ORing solutions, including discrete active ORing controllers as well as full-function system-in-package solutions with high-speed integrated low on-state resistance MOSFETs, in high-density thermally enhanced LGA packages that reduce the power footprint.

Paralleling Power Components for Current Sharing



There can be many reasons for paralleling power components such as DC-DC converters for power delivery. One common situation is when the load current exceeds a single module's rating. Another arises when a degree of power redundancy or fault tolerance is required. Paralleling of modules can provide an effective solution in either case, provided they share the load equally. However, most power components will not do this inherently.

In power converter arrays, one or more of the devices will try to assume a disproportionate or excessive fraction of the load. Typically, the converter with the highest output voltage may deliver current up to its current limit setting, and beyond its rated maximum. Its voltage will then drop back, and other converters will contribute current, but the load will not be shared equally. An overloaded module can become temperature-stressed and prone to failure. A module failure will create a significant step load change for other modules, possibly causing system problems such as resets or a shutdown.

Managed current sharing, then, is important because it improves system performance and increases reliability. It is essential to most systems that use multiple power supplies or converters for higher output power or for fault tolerance. Accordingly, a number of current sharing schemes are available: Output current sensing, droop-share, master-slave, analogue current-share and synchronous current-share.

Output current sense

This forces load sharing by sensing each converter's output current using a resistor, and comparing it to the average current. The output current for each converter can then be trimmed until its contribution equals the average. Selecting the right resistor value is a tradeoff between power dissipation and sensitivity or noise immunity.

Droop-share

Droop-share forces equal current by artificially increasing converters' output impedance, based on injecting an error signal into each converter's control loop. It works with any topology, and is fairly simple and inexpensive to implement. However, it does require a



current-sensing device for each power unit. It also incurs a small but often unimportant penalty in load regulation.

Master-slave arrays

These usually contain one intelligent master, and one or more slave modules or boosters. Quasi-resonant converters within these arrays can share a load equally if set up properly, because the per-pulse energy of each converter is the same. The master determines the transient response, and dynamic load sharing is achieved without trimming, adjustments or external components. Sharing is usually within 5%, and the arrays are stable with excellent transient response. However the design is not fault-tolerant, as maintenance of correct output voltage is lost if the master module fails.

Analogue current sharing

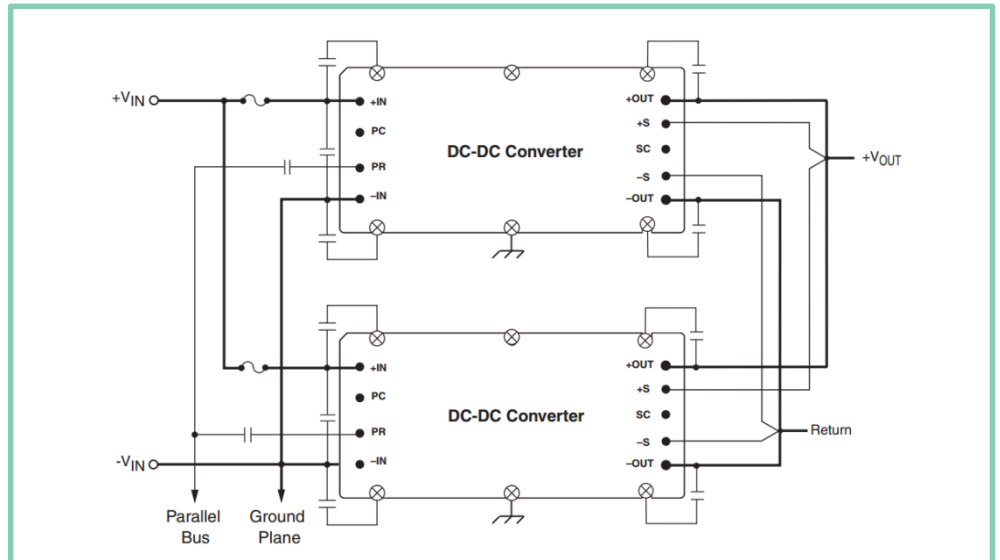
Analogue current-share control involves paralleling two or more identical intelligent modules, with the output voltage of each module being actively adjusted to achieve balanced current sharing. The approach supports a level of redundancy but is susceptible to single-point failures. Also, each converter requires its own voltage regulation loop, current sensing device and current control loop.

Synchronous current sharing

Designs using variable frequency zero-current-switching converters can use synchronous current sharing. This allows a truly 'democratic' array in which any converter module can take control. All modules are interconnected via a bidirectional PR pin as shown below. The module in command transmits a pulse on PR which synchronises each converter's high-frequency switching. The PR pulse becomes a current-sharing signal for power expansion and fault-tolerant applications. If the lead module relinquishes control, another module can change from Receive to Transmit mode, transparently assuming control of the system.



Fig. 1: Synchronous power architecture simplifies current sharing and enhances fault tolerance



This pulse-based design allows improved fault tolerance: transformers or capacitors can be used for DC-blocked coupling, containing certain failure modes within their module of origin. Transformer coupling also provides a high level of common-mode noise immunity while maintaining SELV isolation from the primary source.

Synchronous current sharing in democratic arrays allows power architects to achieve simple, non-dissipative current-share control. It eliminates the need for current-sensing or measuring devices on each module and provides uncompromised load regulation with excellent stability and transient response.

This method applies to quasi-resonant, variable frequency converters with the necessary intelligence, such as the Vicor Maxi, Mini or Micro high-density DC-DC converters with fixed-energy pulses.

What are Y-Capacitors?



When electronic equipment is connected to the AC mains, it has the potential to generate common-mode electrical noise. If this is allowed to flow back on to the mains supply line, it can disturb other equipment also connected to the same line.

What's the solution?

Manufacturers design capacitor-based power line filtering into their systems to decouple any such common-mode noise produced by the equipment's power supply, preventing it from reaching other equipment through the mains power line. The reliability of these capacitors is critical to the safety of the equipment's users.

When does capacitor reliability become critical to safety?

Line filter capacitors are classified either as X-capacitors or Y-capacitors. X-capacitors are connected between line and neutral, to protect against differential mode interference. Their failure does not create conditions for dangerous electric shock, although it can create a fire risk. However Y-capacitors are designed to filter out common-mode noise, and are connected between line and chassis; if they short-circuit, they create a risk of shock to the user.

How are Y-Capacitors designed and deployed to ensure safety?

Y-capacitors are designed to enhanced electrical and mechanical reliability standards. Capacitance values are also limited to reduce the current passing through the capacitor when AC voltage is applied, and reduce the energy stored to a safe limit when DC voltage is applied. Capacitors must be tested to applicable standards to qualify them for use as Y-capacitors.

Which European standards are applicable?

The EN 132400 Standard was issued on 26 June, 1995 replacing all the European National Standards in force up to that date. This was identical to the International Standard IEC 60384-14 2nd Edition 1993. Since then, to make CENELEC and IEC standards identical in name as well as specification, the European Standard EN 132400 has been superseded by EN 60384-14 that is identical to the International Standard IEC 60384-14. Any European national body can issue approvals, with validity recognized by the bodies of all the other CENELEC member countries with no need to repeat the tests.



How about applicable standards in other regions?

USA: UL 1414 for across the line applications and UL 1283 for Electromagnetic Interference filters

Canada: CAN/CSA C22.2N°1 and CAN/CSA 384-14

China: GB/T14472

Are there any sub-classifications for X and Y-capacitors?

EN 60384-14 defines sub-classifications for both types. X1 capacitors are used for high pulse applications, while X2 and X3 types are used for general-purpose applications with different peak pulse operating voltages and peak value surge voltages. Y-capacitors, which are used to bridge operational insulation, are classified as Y1, Y2, Y3 or Y4 according to type of bridged insulation as well as AC and peak voltage ratings. Y1 class capacitors are rated up to 500 Vac, with a peak test voltage of 8 kV. Y2 capacitors have 150 to 300 Vac ratings and a peak test voltage of 5 kV. Y3 capacitors are rated to 250 Vac with no peak test voltage specified. Y4 capacitors are rated to 150 Vac with a peak test voltage of 2.5 kV.

What important tests are related to IEC/EN 60384-14?

These include impulse voltage, endurance and active flammability tests. Application and parameters for these tests depend on the capacitor classification and sub-classification.

What capacitor types are used in line filters?

Two common types are metallised paper/film and ceramic. For Y capacitors, ceramic types are less expensive than metallised film, but unstable over time and temperature and less mechanically stable. Ceramic failure mode also tends toward short circuit, whereas metallised paper and film types tend towards open circuit.

How are Y-capacitors used with Vicor products?

Vicor products including FARM, ARM and VI Brick AC Front End have filters that include Y-Capacitors. Refer to the FARM Design Guide and Application Manual for an example of an integral input filter comprising a common-mode choke, Y-capacitors and X-capacitors

Zero Voltage Switching – What it is and Why it's Important for Buck Voltage Regulation



Buck voltage regulation, an essential part of distributing power from a DC supply to its points of load (POLs), has typically been implemented with a PWM circuit. The PWM duty cycle is varied to accommodate the required voltage reduction.

However a couple of factors have applied increasing pressure to regulator design. Firstly, power density is being forced up as continuously more powerful devices are designed onto boards without corresponding increase in footprint. Secondly, DC power supply voltage levels are tending to rise to minimize distribution losses, while device voltages are reducing to increase internal speed and efficiency. These trends combine to increase the voltage drop and associated switching losses across the regulator.

For example, a process control system may call for regulation from 24 V to 3.3 V – a gap that would typically be covered using two regulation stages; increasing board space, cost and reliability issues. The regulator's switching frequency is also limited, as increased switching activity incurs more losses. This in turn limits the use of smaller passive components for filtering, penalising the density of the total solution.

Although high-density PWM regulators have evolved with improved IC integration, MOSFETs and packaging, their design is no longer sufficient to meet the power demands they face. This is mainly due to switching losses within the regulator MOSFETs. These must be overcome or avoided to achieve any significant boost to regulator performance.


A better solution uses Zero Voltage Switching (ZVS) topology, which allows for operation at a higher frequency and at higher input voltages without sacrificing efficiency. While still PWM based, a separate phase is added to the PWM timing to allow for ZVS operation. Fig. 1 compares a conventional buck regulator with a version modified for ZVS operation. Utilizing the added phase, the ZVS type uses the clamp switch and circuit resonance to operate the high side (Q1) and synchronous (Q2) MOSFETs efficiently with soft switching, avoiding the losses they incur during conventional PWM operation and timing.



For example in the conventional circuit, as Q1 is turned on and Q2 turned off, a very high current flows through the MOSFET pair, since Q2's body diode appears as a short circuit during its reverse recovery time. Other losses arise due to discharging Q1's output capacitance, and to reverse recovery in Q2. These losses increase as the switching frequency or input voltage increases.

By contrast the ZVS design addresses the high turn-on losses of the conventional regulator by eliminating high current body diode conduction prior to turn on of the high side MOSFET, bringing the D-S voltage of the high side MOSFET to zero or nearly zero and producing no high current spikes or damaging ringing. The ZVS action applied to Q1 removes its Miller effect at turn on, allowing the use of a smaller driver and lower gate drive.

Vicor has utilized ZVS topology within our Picor Cool-Power buck regulators. These provide regulation up to $36 V_{in}$ at a higher efficiency and from a smaller form factor than achievable from conventional hard switching, high density regulators.



What Does Power Factor Mean and Why Must We Correct it?



Today's commercial, industrial, retail and even domestic premises are increasingly populated by electronic devices such as PCs, monitors, servers and photocopiers which are usually powered by switched mode power supplies (SMPS). If not properly designed, these can present non-linear loads which impose harmonic currents and possibly voltages onto the mains power network.

Harmonics can damage cabling and equipment within this network, as well as other equipment connected to it. Problems include overheating and fire risk, high voltages and circulating currents, equipment malfunctions and component failures, and other possible consequences.


A non-linear load is liable to generate these harmonics if it has a poor power factor. Other loads can present poor power factors without creating harmonics. This post looks at these issues, the circumstances that can lead to damaging harmonic generation, and practical approaches to reducing it.

The two causes of poor power factor

At the simplest level, we could say that an electrical or electronic device's power factor is the ratio of the power that it draws from the mains supply and the power that it actually consumes. An 'ideal' device has a power factor of 1.0 and consumes all the power that it draws. It would present a load that is linear and entirely resistive: that is, one that remains constant irrespective of input voltage, and has no significant inductance or capacitance. Fig. 1 shows the input waveforms that such a device would exhibit. Firstly, the current waveform is in phase with the voltage, and secondly both waveforms are sinusoidal.



Fig. 1: Input voltage and current waveforms for a device with PF = 1.0



In practice, some devices do have unity power factors, but many others do not. A device has a poor power factor for one of two reasons; either it draws current out of phase with the supply voltage, or it draws current in a non-sinusoidal waveform. The out of phase case, known as 'displacement' power factor, is typically associated with electric motors inside industrial equipment, while the non-sinusoidal case, known as 'distortion' power factor, is typically seen with electronic devices such as PCs, copiers and battery chargers driven by switched-mode power supplies (SMPSs). We shall look briefly at the displacement power factor before moving on to the distortion case, which is of more immediate concern to electronic power system designers. However it is important to be aware of both cases. For example, some engineering courses discuss the power factor issue only in terms of motors, which causes confusion when their students later encounter poor power factor as exhibited by an SMPS.

Electric motors and displacement power factor problems

Electric motors create powerful magnetic fields which produce a voltage, or back emf, in opposition to the applied voltage. This causes the supply current to lag the applied voltage. The resulting out of phase current component cannot deliver usable power, yet it adds to the facility's required supply capacity and electricity costs. Fitting capacitors across motors reduces the phase lag and improves their power factor.




SMPSs and distortion power factor problems

Whereas displacement power factor loads do not cause harmonics and their associated problems, distortion power factor loads such as SMPSs will do so unless their power factor is improved.

An SMPS's AC front end typically comprises a bridge rectifier followed by a large filter capacitor. This circuit only draws current from the mains when the line voltage exceeds that across the capacitor. This causes current to flow discontinuously, resulting in the non-sinusoidal current waveform shown in Fig. 2.

Fig. 2: Non-sinusoidal current waveform drawn by SMPS with poor power factor



It is possible to use Fourier transforms, a mathematical process, to analyse this waveform and break it down into a set of sinusoidal components. These comprise the fundamental frequency – 50 Hz in Europe, 60 Hz in America – and a set of predominantly odd multiples of the fundamental, known as harmonics. The third harmonic is 150 (or 180) Hz, the fifth, 250 (300) Hz and so on. Fig. 3 shows a typical harmonic spectrum for an electronic SMPS load. The fundamental component is usefully consumed by the SMPS, while the harmonics are reactive and create the problems described above. The ratio of the fundamental amplitude to the sum of all the harmonic amplitudes gives the device's power factor.

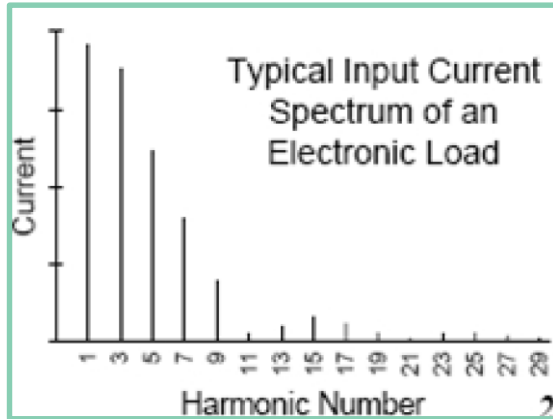
International Standard

An international standard exists to describe and set acceptable limits for a product's mains harmonics generation. Within the EU, its reference is IEC 61000-3-2, covering equipment power levels from 75 W to 600 W. The standard assigns equipment into four Classes –



A, B, C and D. Class D covers personal computers, personal computer monitors and television receivers.

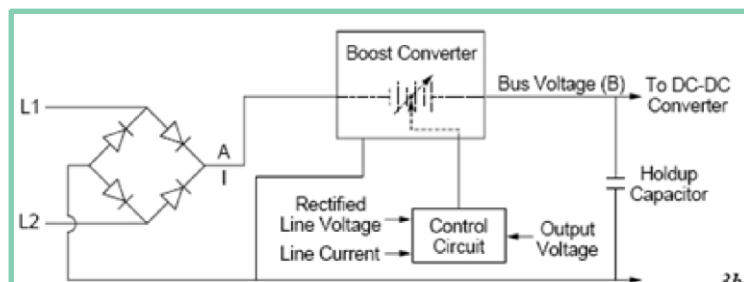
Fig. 3: Typical harmonic spectrum for an electronic SMPS load



Established and innovative PFC solutions

Although passive power factor solutions exist, the general industry view is that active designs offer the best power factor improvements. These are typically based on boost converter technology, as in the example shown in Fig. 4.

Fig. 4: Active power factor correction circuit using voltage boost



In this active power factor correction circuit the incoming line voltage passes through a bridge rectifier, which produces a full wave rectified output (Figure 5–A). Since the peak value of the line is less than the bus voltage, no current will flow into the holdup capacitor unless the line voltage is boosted above that present on the holdup capacitor. This allows the control circuit to adjust the boost voltage (5–B–A) to maintain a sinusoidal input current.

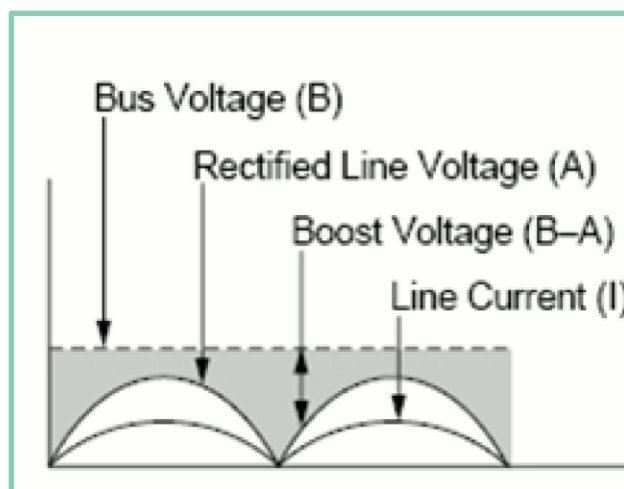


To do so, the control circuit uses the input voltage waveform as a template. The control circuit measures the input current, compares it to the input voltage waveform, and adjusts the boost voltage to produce an input current waveform of the same shape (5-I). At the same time, the control circuit monitors the bus voltage and adjusts the boost voltage to maintain a coarsely regulated DC output (5-B). Since the primary function of the control circuit is to supply a sinusoidal input current, the DC bus voltage is allowed to vary slightly.

The use of an active power factor correcting circuit results in few discontinuities in the input current and consequently low distortion and harmonic content imposed on the input current drawn from the line. However Vicor has recently introduced a modular VI Brick® AC Front End, based on their new dynamic converter architecture, called Adaptive Cell.

The AC Front End offers a number of improvements for systems designers. In particular, it provides 85 V to 264 Vac universal input, high efficiency and high power density, especially considering that it is a complete solution including isolated and regulated DC output as well as rectification and power factor correction. The device reduces propagation of AC line harmonics, improving overall power quality at system and facility level. Total Harmonic Distortion exceeds EN61000-2-3 requirements, while high switching frequency and resonant transitions simplify external filtering and EMI standards compliance.

Fig. 5: Voltage and current waveforms for active boost circuit of Fig. 4



Understanding and Mitigating the Growing Problem of Distribution Losses



There are many articles available describing how to overcome power distribution losses with suitable power conversion devices and architectures. This post goes back to basics: explaining the underlying issues of distribution losses and the principles of their solution.

If DC-DC converters and power conversion devices are any part of your working life, you will be aware of their key role in conducting electrical power from the edge of a printed circuit board (PCB) to its onboard components efficiently. You may also be aware of Vicor's Intermediate Bus Architecture (IBA) and Factorized Power Architecture and other technologies designed to maximize this efficiency. The objective is to minimize distribution losses as the power flows from the PCB connector to its onboard point of load.


However you may not be aware of the fundamentals – what causes distribution losses, why does advancing technology exacerbate their impact, and how do power devices and architectures work to mitigate this? We provide answers to these questions below as we explain the underlying problems that FPA and other power architectures are designed to solve.

Even with mitigating technologies, distribution losses – or, more definitively as we shall see, I^2R losses – are an inevitable part of transmitting power from one place to another. This is equally true of mains power transmitted across the National Grid or of low voltage DC power applied to a PCB. The reason is that the conducting medium, whether it is an overhead cable or a PCB track, always has a finite resistance, and power is lost, or dissipated, in overcoming this resistance. The remaining power is dissipated by the device being supplied; a microprocessor chip, for example.



To understand more about the nature of electrical power dissipation, let's start with Fig. 1:

Fig. 1: Voltage applied across a load



In Fig. 1, a DC voltage V volts is being applied across a resistive load R ohms. The load is drawing a current I amps. At this point, we need to introduce a couple of formulae; however they're the only two we have to understand. We can derive everything else from them.

Firstly, there's Ohm's Law, which says

$$V = IR$$

Where V is the voltage applied across (say) a device, R is its resistance, and I is the current it draws. For a fixed voltage supply, we can see that the lower the device's resistance, the more current it draws.

The other formula is Watt's Law which states

$$P = VI$$

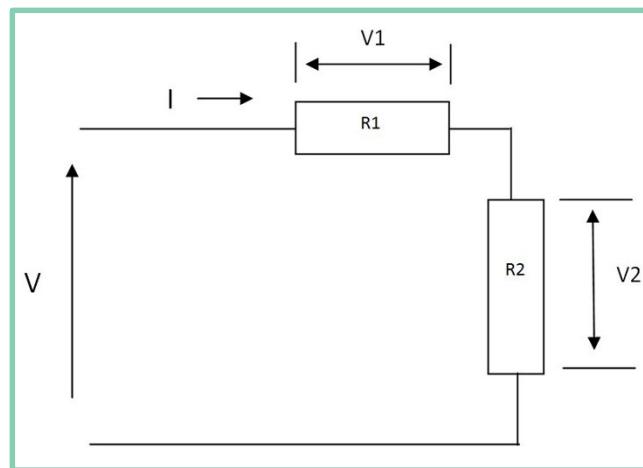
This says that in Fig. 1, the power (P) dissipated by the device is the product of the voltage (V) across it and the current (I) that it draws. Note that P depends equally on V and I . Note also that if a device takes, say, 100 W, this could be delivered either as 100 V at 1 A, or 1 V at 100 A, or any other VI combination that multiplies out to 100. However in considering these variations we must not forget the device's maximum input voltage rating.

At this point, let's imagine that the device in Fig. 1 is a microprocessor, and that it draws 100 A at 1 V. These are not real figures, but they are representative of what could be expected from



a desktop PC chip. From Ohm's law we can see that it has an internal resistance of $V/I = 1/100 = 0.01$ Ohm. As we move from Fig. 1 towards reality – but without yet giving ourselves the benefit of modern power train technology – we recognise that the chip's 1 V supply has to travel along copper tracks connecting the board's power connector to the chip's power pins. These copper tracks have a finite resistance. If they are 5 mm wide and 5 cm long this could be approximately 0.005 Ohms. Note that the track resistance is directly proportional to the track length. Fig. 2 represents the electrical view of this scenario.

Fig. 2: Microprocessor load on a PCB



In Fig. 2, $R1$ represents the PCB track resistance of 0.005 Ohms while $R2$ represents the microprocessor internal resistance of 0.01 Ohms. Accordingly, if we now apply 1 V to the board edge connector, Ohm's law shows that the current would be

$$\frac{1}{0.01+0.005} \text{ or } 66.7 \text{ Amp}$$

not the 100 Amp the chip may require. It also shows that there would be a voltage drop across the track of $66.7 \times 0.005 = 0.33$ V – leaving only 0.67 V across the chip. Track resistance has deprived the chip of both voltage and current.

Consider also the power lost through dissipation in the track. If $P = IV$ and $V = IR$, then by substituting for V we can see that $P = I^2R$. So, the power dissipated in track resistance $R1$ is $66.7^2 \times 0.005 = 22$ W. In this case the Distribution Loss, or I^2R Loss, is 22 W.



These losses and inefficiency are clearly unacceptable, so what can we do? To see the answer, let's suppose that our microprocessor is a very flexible or even magical device that can work on a 50 V input supply just as well as on 1 V. We know that it draws 100 A at 1 V, so according to Watt's Law it consumes 100 W. Therefore it will require 2 A at 50 V. Importantly, in this new scenario, the device's internal resistance would have increased from 0.01 Ohm to $50/2 = 25$ Ohm.

So if we apply 50 V across the connector in Fig. 2, the actual current will be

$$\frac{50}{25+0.005} = 1.996 \text{ A}$$

a very small loss in current to the device. The voltage drop across the track is just 1.996×0.005 or 0.0099 V. The chip is enjoying nearly all of the available voltage and current. The dissipation across the track is 0.02 W. By transmitting power at a higher voltage and lower current, we have delivered it where it is needed, with minimal transmission losses.

All of this prompts the question: If higher voltage levels are so beneficial for power transfer efficiency, why don't chip makers produce devices with higher input voltages? The answer is that the manufacturers are driven by their own priorities, which are to deliver ever more processing performance while curbing the associated increase in power demand. Reducing the chip's supply voltage contributes to both these objectives. So, over the years we have seen chips demanding steadily increasing current at steadily reducing voltages; a formula for increasing distribution losses! Pressure on power system designers to find a way of supplying the chip efficiently has grown accordingly.

An obvious solution would be to place a DC-DC converter physically close to the microprocessor chip. This can receive power at a high voltage and low current and convert it into low voltage and high current suitable for the chip. I^2R losses are minimized because the track length and resistance between the converter and chip are minimized. In fact one of the earliest responses was to use exactly this approach. However continued development in electronic device and PCB design rendered it increasingly unsatisfactory. In reality, DC-DC converters had to handle isolation and voltage regulation as well



as conversion, and their own size and electrical efficiency became increasingly important issues. This was especially true as the number of devices on a PCB, and the supply voltages they demanded, proliferated steadily.

For these reasons, more advanced approaches such as Intermediate Bus Architecture (IBA) and Factorized Power Architecture (FPA) as previously mentioned, have appeared. They are sophisticated tools that tackle these complex issues, yet the fundamental, underlying problem they solve remains the same: How to deliver increasing levels of power to increasing numbers of ever lower voltage devices, while minimizing distribution I²R losses en route to the point of load.

Handling High Input Transients



Although perhaps not as obvious as the EMI requirements mentioned in an earlier post (Meeting EMI for AC-DC Systems), the protection of the system against transients is not only a fundamental part of the development but can also become a major challenge during the design cycle.

Differential Mode

These are transient voltages between the supply lines itself and can be generated by sudden load changes like a load on/off switch, bouncing of relays, motors, short circuits or instabilities of the source itself. Imagine a powerful motor connected to the same supply voltage as the power supply of your system. With full load this motor is taking a lot of energy from the supply source. When the load disappears – or even worse, if the motor has to go into break mode – this energy will be fed back into the source and can generate high overvoltages with a significant amount of energy.

The same can happen if inductive loads are switched off. According to the laws of physics an inductor will maintain the current due to the stored energy. However, if disconnected all of a sudden, the voltage goes up very rapidly. The higher the impedance in the system the higher will be the voltage spike.

One may think that this only happens in a heavy industrial environment or in railway applications, where you can even see the arcing between the pantograph and the catenary during short interruptions of the connection. But even in domestic applications a fuse may blow due to overload or a short circuit. Suddenly, your TV set gets a high transient voltage at the input due to the wires and some other inductive components in the electric installation.

Common Mode

These transients occur between the input lines and the ground or chassis of the system. Typical sources are electrostatic discharging, which can happen when someone touches the metal chassis of a system. This normally does not do any harm to the system and, in improperly designed systems, it may cause a short interruption. But there are other sources like very strong electromagnetic fields which worst case are caused by lightning strikes near the equipment.



This brings us to another point of consideration, which is the energy content, the transient source impedance and the duration of the transient voltage itself. While in most cases electrostatic discharge has low energy combined with high source impedance, one can imagine that it's just the opposite when lightning strikes a system or the supply lines. In the case of differential mode transients, a 10 W motor certainly generates transient pulses with lower energy compared to one with a power of 100 kW.

Therefore it is necessary to look up the values defined for these transients in the applicable standards. They are based on long term experience and real measurements in various applications and under different operating conditions. It's not only the maximum voltage but also the duration and the energy content or source impedance that can be found in these standards. Due to their experience Vicor's application engineers can assist in finding the right standard.

The next step is to select the correct strategy to protect the system. First of all the converter itself has to work over the full rated nominal input voltage as every transient suppression circuit can only protect the system for a limited period of time. Transients with the duration of a few ns or μ s and lower energy content may be blocked by passive filters or simple active protection devices like varistors or transzorb diodes. Some examples are shown in this video: Input Overvoltage Protection.

For higher energy pulses with voltages up to 3-4 times the maximum nominal supply voltage, and sometimes a duration of several ms, this simple approach is no longer possible, and more complex active clamping circuits are needed. During the transient pulse the power supply is either electronically disconnected from the source or a linear regulator limits the current and voltage at the front end. By using the appropriate complete filter solution, like the FIAM110 for railway applications the system will be protected. Another option is the use of a discrete solution as shown in this application note: MIL EMI and Transient Solutions.

Meeting EMI for AC-DC Systems



EMI is seen as one of the key challenges of power conversion design and system integration, particularly for AC-DC systems. Our latest technical tip by David Fletcher, Principal Engineer from Vicor's Westcor Division, looks at the causes of EMI and its impact on a system, before going on to make some practical suggestions on how to minimize its effect.

Causes of noise

We have to start with what engineers affectionately call noise. This is generated whenever rapid voltage (dv/dt) and/or current (di/dt) transitions take place. AC-DC power converters typically utilize a number of power switching topologies, including fully resonant, quasi-resonant and PWM (Pulse Width Modulation), the most typical control in the front-end section of the converter.

PWM-controlled converters use a rectangular control signal with a continuously varying pulse width in response to the operating conditions of the converter. The result is typically a "white noise" energy distribution spectrum. If this was not filtered and shielded it would interfere with consumer electronic equipment using the same AC mains.

"Noise" currents exiting the power converter through the AC power lines and/or via output power cables are known as "conducted emissions." The noise manifests itself in two forms: differential and common-mode. The definition of differential-mode is noise that is only on the power lines and is not present on the earth ground lead, and can be measured with respect to the power lines. The definition of common-mode is noise that can only be measured from earth ground to one of the power lines.

AC-DC power converters employ EMI filters within the power converter enclosure. These filters contain noise suppression topologies containing inductive and capacitive components. These components are called "X" and "Y" elements. The "X" components are placed across the power lines and filter the differential-mode noise; the "Y" components are placed between the power lines and earth ground and filter common-mode noise.



How to minimize EMI filtering

AC-DC power converters are designed to meet various regulatory and safety standards, including various EMI standards. The individual details and the standards met are generally covered thoroughly in the product literature. However, during a system integration of various components, including an AC-DC power converter, a system design engineer may find it necessary to add additional EMI filtering to enable the integrated system to comply with all relevant agency standards. If this approach becomes necessary, it is recommended to minimize the additional filter components required by empirically gathering test information. It is recommended that a few areas be given special attention:

1. Test the system without any extra filters. This will give the system engineer a good baseline.
2. Identify the problem frequencies on the EMI plots.
3. Try to relate these frequencies to areas of the integrated hardware within the system.

Re-check all points of earth grounding and re check cable routing — all these should be in accordance with good engineering practice. These include but are not limited to separating signal cables from power cables, using twisted pair techniques (3 to 5 twists per inch if possible) on cable carrying signals and /or high current and their ground returns. If using this technique is deemed impractical then strip lining the cable as close to one and other is an alternative to minimize noise. Consider shielding cables with fast voltage and/or current transitions. In integrated systems that have several grounding points a careful study of the system integration block diagram is advisable to detect the presence of grounding loops. If these are present they must be eliminated.

What to consider when adding external filtering

If EMI problems still persist, then adding an external filter may be necessary. There are areas to consider when considering an off-the-shelf filter:

1. The filter must be able to handle the full rated worst case current of the system.



2. The filter must provide enough attenuation at the noise frequencies of concern to provide enough margin for system-to-system manufacturing variations.

3. The amount of extra leakage current (touch current) caused by the additional "Y" capacitance inherent in the filter should be calculated. If the leakage current is not stated on the filter data sheet the following formula can be used to calculate the leakage current:

$$2 \times \pi \times (\text{AC-Volts-max}) \times (\text{AC-frequency-max}) \times (\text{Y capacitors}) = I \text{ leakage}$$

20% to 25% should be added to this value before you add this to the overall leakage current. If the total leakage current exceeds the allowable regulatory level you may want to consider a filter that has an internal series inductive component in the ground lead. These filters tend to have no "Y" capacitors and may be referred to as medical filters.

4. The filter should have an internal resistor between the power lines to ensure that upon disconnection of the AC line input the internal "X" capacitors discharge properly – typically within a second, as recommended by VDE 0806 and IEC380.

5. The effect of adding an external filter means additional inductive elements are now in series with the power converter's AC input. Care should be taken to prevent a phenomenon called "ring-up." This occurs when the forward current through a series inductor is high enough to saturate the filter inductor, so that when the resultant magnetic field collapses a very high peak may occur, with a voltage amplitude up to twice the original impulse. These peaks may be caused by lightning strikes, external power company grid switching, or heavy inductive transitions due to motor loading in a building. Commercial converters have input protection schemes and comply with regulatory standards. However with the addition of an external filter through "ring-up" it is possible to exceed the protection components and cause damage to the power converter. External protection components, such as MOV (Metal Oxide Varistors) may be needed across the external filter to add robustness to the overall system.



6. The filter should be placed as close as possible to the AC input of the system enclosure. This will maximize the EMI filtering. All ground connections to the filter should be as short as possible to minimize series inductance and impedance.

Back

• +IN

• PC

• PR

-IN

+OUT

+S

SC

-S

-OUT

Phone:

Europe 00 800 8426 7000

USA 1 800 735 6200