

POWER DISTRIBUTION ARCHITECTURES

TRACK THE LOW-VOLTAGE TREND

BY ROBERT MARCHETTI • VICOR CORPORATION

As each new generation of processor, memory chip, digital signal processor (DSP), and application-specific integrated circuit (ASIC) exploits the capabilities of silicon manufacturing processes at ever-smaller feature sizes, the trend to lower supply voltages at higher currents continues. This progression has presented a continuing challenge to power-supply infrastructure and has exposed the limitations of, in turn, established distribution architectures that include centralised power architecture (CPA), distributed power architecture (DPA), and intermediate bus architecture (IBA). The most recent step in this sequence addresses the voltage/current demands of today's largest ICs, under the banner factorised power architecture.

The classic CPA, which is simple and cost-effective, continues to be applied wherever appropriate. Communications systems applications were among the first to reveal the inability of a centralised

architecture to efficiently deliver lower voltages at higher currents.

A centralised power supply contains the entire power supply in one housing, from the front end through the DC-DC conversion stages

(Figure 1). It converts the line voltage to the number of DC voltages needed in the system and buses each individual voltage to the appropriate load. It is cost effective and does not use valuable printed-circuit board (PCB) area close to the point of load for the power conversion function. It is moderately efficient because it avoids serial power transformations, and it confines the thermal and electro-magnetic interference (EMI) issues to one enclosure. In the past, the centralised system – usually a custom design – was the designer's default choice as the least expensive approach. CPA is suited to systems in which power requirements, once defined, are not likely to change and in which space is not overly constrained.

In order to minimise resistive (I^2R) distribution losses, you should locate a central supply as close as possible

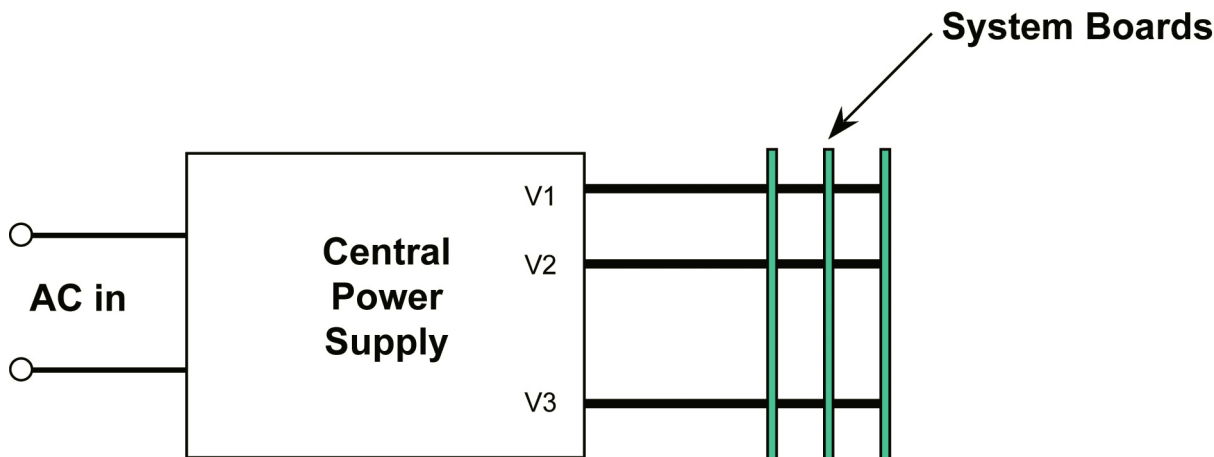


Figure 1 A centralised power supply contains the entire power supply in one housing and buses each voltage to the appropriate load.




Figure 2 Distributed power is a decentralised power architecture consisting of DC/DC converters located near the load they serve.

to the load: conversely, for safety and EMI reasons, you should place it adjacent to the AC entry point. This is often a difficult trade-off.

Although centralised power works well for many applications, it becomes increasingly unsuitable as today’s low voltage loads demand ever-higher currents, into the hundreds of amps. Also, CPA is not scalable: system designers can create highly-configurable products that will ship with a variable mix of function cards. With centralised power, you must rate the power supply to handle the maximum configured system, which represents a cost penalty for the “average” configuration.

Of particular importance for loads such as current-generation microprocessors, remoteness of the supply from the load degrades its transient response. Thermal management can also be a challenge in a centralised architecture, with all of the PSU dissipation – possibly hundreds of watts – located in a single volume. Such PSUs often require large heat sinks and fans, and

can create hotspots that are a source of reduced reliability.

DISTRIBUTED POWER ARCHITECTURE

As low voltage loads proliferated, the concept of the “brick” (sealed power supply function module) first appeared, and distributed power came of age (Figure 2). Distributed power placed DC/DC converter bricks on system boards, close to their loads. Since the 1980s, the bricks of DPA have delivered the classic functions of the DC/DC converter (isolation, voltage transformation and regulation) to the point of load. But as the number of voltages required at the board level continued to increase, DPA began to take up too much valuable PCB area and the cost of duplicating the full converter functionality, many times over, became excessive.

Distributed power is a decentralised power architecture characterised by bussing a “raw” DC voltage, usually 48 or 300 Vdc depending on the power source. On-board DC/DC converters, located near to the loads they serve, complete the chain by delivering the correct bus voltage. You match these isolated DC/DC converters to each load, which improves dynamic response and eliminates the problems associated with distributing low voltages around the system.

A distributed approach spreads the PSU dissipation throughout the system, greatly reducing or eliminating the need for heat sinks or high velocity airflow associated with the power supply function. You can more easily meet reliability specifications in the absence of hotspots. Also, spower provision resides on, and is rated for, each system PCB, configuring system variations and options is much more cost effective than in a centralised

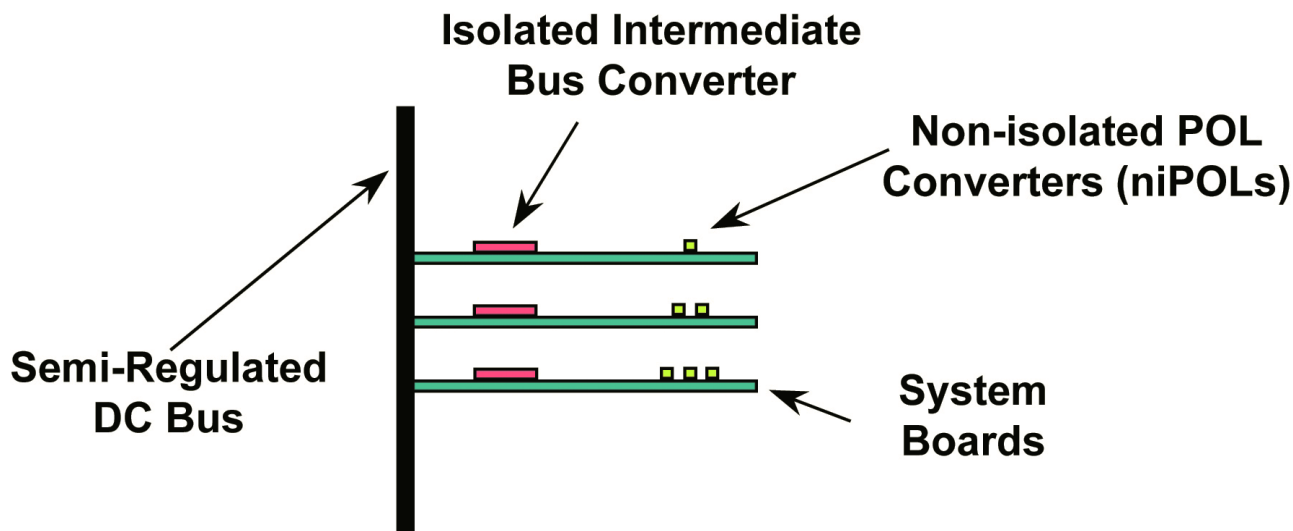


Figure 3 The intermediate bus architecture uses an isolated intermediate bus converter to provide an unregulated voltage to power non-isolated and relatively inexpensive POL converters.

architecture. Redundancy is easy to implement for any critical load; you can simply parallel additional DC/DC converters where required.

DPA can, however, also be more costly. Isolation, regulation, transformation, EMI filtering, and input protection are repeated at every load so, as the loads proliferate, both the costs and PCB area for power conversion increase.

INTERMEDIATE BUS ARCHITECTURE

To deal with a multiplicity of low voltages more cost-effectively, IBA relies on non-isolated point-of-load regulators (niPOLs), reducing the POL function to regulation and transformation. The niPOLs operate from an intermediate bus voltage that upstream isolated converters provide. IBA can be a more cost-effective solution because niPOLs, being non-isolated, are less expensive than complete DC-DC converters. However, with typical niPOL buck converters there is invariably a trade-off between efficient power distribution, and efficiency in the power conversion duty cycle.

The intermediate bus architecture differs from the distributed power architecture in that it converts the raw DC voltage (for example, 48 or 300 Vdc) to an intermediate voltage, typically 9.6 or 12 Vdc, to feed non-isolated and relatively inexpensive POL converters (**Figure 3**). The

niPOLs are also likely to be smaller and lighter than full-function DC/DC converters, easing PCB-area concerns. Eliminating isolation and high voltage-transformation ratios improves their cost-effectiveness.

The niPOLs of IBA depend upon a bus converter to provide isolation and voltage step-down from the raw DC bus. This is accomplished by the intermediate bus converter (IBC), which is usually either a complete DC/DC converter operating from a wide-range DC source, or an unregulated IBC operating from a narrow-range input. The addition of the conversion stage to the intermediate bus voltage intrinsically reduces efficiency of the system. Also, it becomes even more essential that you position the IBC intermediate bus converter as close as possible to the load because, even with a 12-V intermediate bus, four times the current needs to flow around the board as compared to a 48-V distributed power system. The inevitable consequences include wider PCB traces, heavier copper, or shorter connection runs.

The 12-V intermediate bus is also, still, too high for efficient conversion to low voltage outputs (less than 2 Vdc) as the transformation ratio becomes too high, and the switch duty cycle becomes too low. Lowering the bus voltage to overcome this limitation simply exacerbates the problems associated with providing suitable paths for large currents.

Non-isolated POL Converters (niPOLs)

System Boards

Since the niPOL includes regulation, it needs an inductor in series with its output. These same low voltage loads (typically, microprocessor cores) generally need fast transient response, but the introduction of that inductance places inertia where agility is most needed. These are the fundamental limitations of IBA when it comes to powering today's sophisticated low-voltage, high-speed loads.

One further disadvantage of niPOLs is inherent in the absence of isolation: loads are vulnerable to deadly faults and the entire system is susceptible to ground-loop and noise coupling problems.

FACTORISED POWER ARCHITECTURE

FPA reorganises the basic power conversion functions — voltage transformation, isolation, and regulation — and implements them in IC-style packages (**Figure 4**). A buck/boost pre-regulator module (PRM) provides a stable voltage from an unregulated DC bus, and a voltage transformation module (VTM) steps the voltage up or down and once again re-locates isolation, placing it back at the point of load. The VTM is a fixed-ratio DC/DC converter. High-frequency, high-integration FPA modules using zero-current/zero-voltage soft switching topologies combine attributes such

as small size, high efficiency, low noise and fast transient response, and can achieve very high power density of over 1kW/in³, at the point of load.

Figure 5 shows the FPA modules in a basic arrangement, but the PRM and VTM can be operated alone, together, open loop, local loop, adaptive loop, remote loop, co-located, separated, paralleled, or combined with conventional power conversion devices (that is, DC/DC converters, point-of-load converters, or charge pumps). The underlying technology of the VTM is the class of power conversion topologies called sine amplitude conversion (SAC), which has seen increasing application since its first appearance in around 2003 – it also forms the basis of the intermediate bus converter itself. A sine amplitude converter uses, as the name implies, a high-frequency controlled oscillator with high spectral purity that runs at the resonant frequency of the power train, including the internal transformer. A completely symmetrical circuit configuration (all switching each side of the transformer is mirror-image) minimises common-mode noise. As a fixed-ratio device, the SAC-based VTM has no internal regulation mechanisms, and therefore avoids issues of stability and loop delay while still providing fast transient response. Output stability derives from the voltage produced by the upstream PRM, and you can

apply feedback signals to that device in a variety of configurations.

Some of the characteristics of the SAC that contribute to its role in overcoming the limitations of the earlier IBA include the fixed switching frequency mentioned above. In present implementations, this is at 3.5 MHz, a value that minimises size of all reactive components, is very easy to filter, and decreases the response time. “Soft” switching at zero-voltage and zero-current points minimises switching losses and cuts — by as much as an order of magnitude — dV/dt and dI/dt, resulting in low noise. FPA also eliminates the output inductor of the IBA niPOL: it has minimal serial energy storage, with no associated power loss, and there is no penalty imposed on the dynamic response.

SAC maintains 100% switch duty cycle at any transformation ratio, resulting in efficient power train utilisation. It also maintains



Figure 4 Vicor recently added the PFM isolated AC/DC converters to its range of low-profile modules, enabling the company to offer a complete factorised power system in 9.5-mm-profile packaging. The PFM – these deliver 330W at 48V – converts AC line voltage to the 48 Vdc that the pre-regulator module in Figure 5 accepts.

the analogy of the VTM’s function as a “DC transformer” in that it supports bi-directional power processing, recycling load dump energy to the input, improving transient response. SAC also exhibits capacitance reflection and multiplication; you get a high effective point-of-load capacitance without the physical presence of bulk capacitance. **EDN**

AUTHOR’S BIOGRAPHY

Robert Marchetti is Senior Manager of Product Marketing for Vicor’s Brick Business Unit in Andover, MA

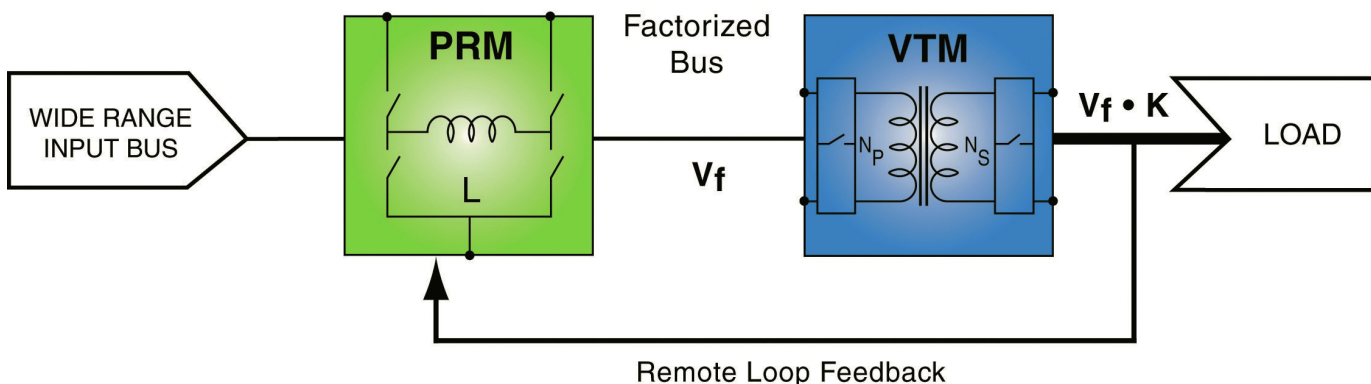


Figure 5 In this basic FPA configuration, feedback to enable the PRM to regulate the intermediate bus voltage comes from the final point-of-load (“remote-loop”). You might also derive a control signal from the VTM itself (“adaptive”); or have the PRM monitor its own output (“local”); or rely on the inherent stability of the separate modules (“open-loop”).