Reed Electronics Group

*Solid-state relay switches high-frequency test signals Page 13*

# EDN

July 24, 2003 www.edn.com



HE DESIGN MAGAZINE OF THE ELECTRONICS INDUSTRY

## DISTRIBUTED POWER:

# Taming the dragons

pg 36

# DISTRIBUTED POWER: taming the DRAGONS

**NEW ARCHITECTURES AND PACKAGING ARE TEAMING TO TAME THE TINY VOLTAGE DROPS THAT SO EASILY SPELL DISASTER IN SYSTEMS WHOSE SUPPLY VOLTAGES WILL SOON BE MEASURED IN MERE MILLIVOLTS.**

Illustration by Mike O'Leary

Power supplies just may be surrounded by more myths and misconceptions than any other segment of electronics. Within the past year, the very nature of power-system design has changed, and the old rules have pretty much flown out the window. Although some of the following statements may once have been true, most of them, if not out-and-out false, are now accurate only in specialized niches of the power-supply field.

- Power supplies are low-tech.
- Except in unusual cases, any EE—and most electronics technicians—can design a suitable power supply; no special knowledge or experience is required.
- The power-supply field is a technological backwater, which, compared with the rest of the electronics industry, changes at a snail's pace.
- If you integrate power-supply modules into higher level products—as opposed to buying the parts and integrating the power-supply functions—you are wasting a lot of money. Most likely, you are doing so because power-supply vendors' scare tactics have taken you in.
- IC vendors' reference designs for supplies that you can build into your product are much less expensive than the equivalent modules. Moreover, because the IC vendor does all of the hard work, the built-in supply does not increase your development costs compared with those you would incur if you used purchased modules.

The facts are different. Power-supply design is only rarely trivial; it is a highly specialized craft that combines electronic, electrical, thermal, mechanical, and manufacturing engineering. Experienced power-supply designers—the good ones, anyhow—understand as do few others, the extremely complex and subtle trade-offs that can profoundly affect a design's reliability and cost. Even if you plan to integrate an IC manufacturer's power-supply reference design into your product, you would be well-advised to assign the task to someone experienced in power-supply design. Similarly, power-distribution experience can be a big help to engineers who must apply purchased power modules.



That said, the job of designing power supplies into higher level products greatly differs today from what it was even two or three years ago. The whole landscape of power-supply design has changed and continues to change in ways that strongly benefit designers of higher level products. Power supplies are dramatically smaller, more efficient, and less expensive than were their counterparts of just a short time ago. The cost reductions relate more to advances in technology than to the disastrous business conditions that currently plague the industry.

IC manufacturers—not power-supply designers—drove one of the major areas of change onto the power-supply industry. ICs are now designed to make optimum use of semiconductor-manufacturing processes that can support only limited supply voltages. Gone are large pc boards that used just a few supply voltages—often, 5V and ±12V.

Today, such boards often use six or more supply voltages—for example, 5, 3.3, 2.5, 1.8, 1.5, and 1.2V, and sometimes −5V. Moreover, the voltages are now so low and the currents so high that IR drops in the distribution networks (that is, in the power and ground planes) can often make or break a design. And the voltages are going lower: 1, 0.8, 0.6, and even 0.5V are in the offing—generally with no decrease in the power the ICs require. Constant power at lower $V_{CC}$ translates into higher $I_C$ and the need for much lower $R_D$ (distribution resistance).

Also, static (that is, I·R) drops are not the only voltage drops whose importance the new lower voltage ICs has magnified. Minimizing the power dissipated in those ICs has become ever more important—and not just in battery-powered products. Low power consumption is a goal even in large line-powered equipment, because reducing power dissipation reduces cooling requirements and lowers devices' operating temperature, thereby increasing equipment reliability.

Lowering IC power dissipation often dictates the use of power-saving modes. When they enter such modes, ICs demand much less supply current; when the devices return to "normal" operation, however, power supplies must suddenly deliver high currents. These drastic current changes produce voltage transients that the supplies usually must correct. (The formula $e=L\cdot di/dt$ describes the effect. In the formula, e is instantaneous voltage, L is the inductance between the supply and the load, and di/dt is the rate of change in the current delivered to the load.) Achieving tight regulation at the point of load increasingly requires supplies to incorporate remote sensing, a feature that, unless designed with the utmost care, can easily degrade a supply's transient response, dynamic stability, or both.

Transient-regulation requirements now often necessitate much more closely controlled dynamic behavior than that of the supplies system designers considered more than adequate just a few years ago. A key to improved dynamic behavior is higher switching frequencies. Most modern POL (point-of-load) converters use MOSFET switches. Without faster high-current switches, the higher switching frequencies would unacceptably limit efficiency.

In this regard, newer ICs' requirements for lower supply voltages have proved to be a boon; in general, increasing a MOSFETs' switching speed reduces both the devices' breakdown voltage and their on-resistance. Lower voltage supplies can tolerate lower breakdown-voltage switches. For any rated output current, lower on-resistance improves a converter's efficiency—even in designs that don't take advantage of the FETs' ability to switch faster.

### NEW MEANINGS

The mandate for multitudes of minuscule supply voltages has given new meaning to time-honored terms such as DPA (distributed-power architecture), IBA (intermediate-bus architecture), IBV (intermediate-bus voltage), and POL (for point-of-load regulation as well as point-of-load converter). BC, IBC, RBC, and BCM stand, respectively, for bus converter, intermediate-bus converter, regulated bus converter, and bus-converter module. Systems that use multiple supply voltages have also raised the importance of such features as supply sequencing, which not too long ago system designers could often safely overlook. Texas Instruments' PTHxx module family exemplifies the advanced sequencing capabilities that supply manufacturers are now starting to offer (**Figure 1**).

DPAs have existed for decades in a variety of forms. For several years, however—until about a year ago—many people understood the term to mean a DPA that originated in the telecom industry and had spread to other applications. This DPA, which remains popular, became the de facto standard for powering systems consisting of multiple boards, each of which usually consumes no more than approximately 100W. In this architecture, each board receives a relatively high IBV, usually a poorly regulated nominal 48V (actually, 36 to 75V).

In telecom and in some high-availability-computing applications, the 48V IBV usually originates either in a line-operated supply separate from the boards or—in the event of an ac-line failure—in a bank of large lead-acid storage batteries whose charge is normally maintained by the line-operated supplies. The wide 36 to 75V input range easily accommodates batteries in any usable state of charge. Because the conversion from ac-line voltage to IBV doesn't occur on the pc boards, some people don't even consider this arrangement an IBA.

#### AT A GLANCE

- Many in the power-module industry agree that the most economical distributed-power architectures concentrate isolation for an entire pc board in one module and locate nonisolated regulators near the points of load.
- Such architectures are well-suited to today's ICs, which use lower and lower supply voltages with each passing month.
- Despite their cost advantages, intermediate-bus architectures impose efficiency penalties, but system designers appear to consider these penalties less important than the cost savings.
- Packaging is playing a key role in the evolution of power modules.

Usually, one or more isolated, closely regulated "sub-brick" dc/dc converters located on the board, near the edge that contains the backplane connector, develop the lower dc voltages that power the board's ICs. Even though, unlike many POLs, these modules incorporate transformer-based input-to-output isolation and they are joined to the loads by pc-board traces that are often longer than 1 ft, these converters technically qualify as POLs.

Within the past year, however, a different DPA has been growing in popularity. This new IBA substitutes nonisolated POLs for the isolated sub-brick devices and concentrates the isolation in the BC. The BC is an isolated sub-brick module whose output is normally unregulated but in some cases is loosely regulated, for example to compensate for input-voltage variations but not for output-voltage variations that result from the BC's nonzero output impedance. The most significant advantage of this IBA is its inherently low cost (see **sidebar**, "It's the money, stupid!"). The most significant disadvantage is reduced efficiency. When you convert the voltage twice—from the voltage supplied to the board to the IBV and from the IBV to $V_{CC}$—you generally lose a couple of percentage points in efficiency. If your load draws 100W, the IBA typically consumes 2W more than would a single-conversion architecture. For many system designers, however, lower cost trumps lower efficiency.

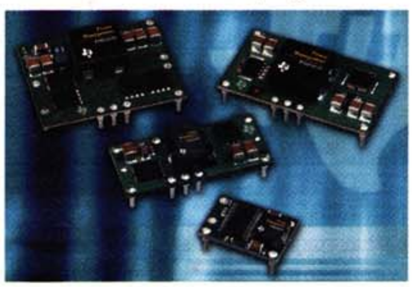

**Although it is best known for its ICs, Texas Instruments also manufactures power modules. Advanced power-sequencing capabilities distinguish the PTHxx series of nonisolated POL converters, which TI expects to gain wide acceptance in systems that require multiple low-supply voltages.**

Even with the cost-effective planar magnetics that most isolated sub-brick dc/dc converters now use, the transformer is a significant cost item. Manufacturers of dc/dc converters are in almost complete agreement that power systems that use only one isolation transformer are less expensive than systems that use many. The key exception to this view is Vicor, which recently introduced what it calls FPA (Factorized Power Architecture). This architecture uses individually isolated POLs, which the company calls VTMs (voltage-transformation modules), and separate regulators called PRMs (preregulation modules) (**Reference 1**). Usually, there is one PRM for each VTM. Vicor won't reveal how it achieves its favorable costs, but the company's initial releases suggest that FPA pricing is often lower than that of systems in which one BC powers several POLs. Moreover, the packaging, called VIC (V·I Chip), delivers exceptionally high power density, and the fixed clock frequency of almost 4 MHz should provide extraordinarily rapid transient response.

Vicor's FPA has raised many questions but so, too, has the aforementioned single-point-of-isolation IBA (see **sidebar** "Do you need an intermediate-voltage bus?"). Because lower cost was the main motivation for creating this IBA, some suppliers have ruled out certain choices that otherwise might make sense. For example, the BC might provide line regulation but not load regulation. A BC with regulation only at its input would allow the familiar 36 to 75V telecom input-voltage range and thus would facilitate battery backup. Also, eliminating the BC's ability to compensate for load changes would save the cost of analog isolation in the feedback path, thus making the BC less expensive than a fully regulated module. Still, if the power-supply manufacturer convinces the system manufacturer that there is no problem with limiting the BC's input-voltage range to, at most, ±10% from nominal, the POLs can do *all* of the regulation. Besides costing less than a unit with input regulation, an unregulated BC increases efficiency and power density and improves reliability.

Another controversy relates to the IBV. Although system designers have made a good case for IBVs of less than 12V, several power-module manufacturers believe that, except perhaps in a few high-volume custom applications, 12V must triumph because of the large infrastructure associated with it.

Module manufacturers say that they are hard-pressed to find customers who design their pc boards with the POLs in immediate proximity to the ICs they power, despite good reasons—such as improved transient response and possible elimination of remote sensing—for so locating the POLs. You might think that, with long ground-return paths and

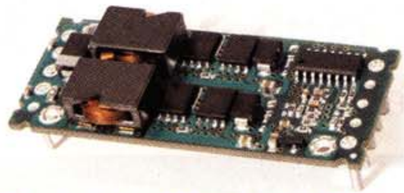

**The nonisolated 25A LEN Series, which Datel offers in SMT and through-hole versions, has extra pins for managing external decoupling capacitors and multilayer ground and power planes.**

$V_{CC}$ values in the vicinity of 1V, mixing different ICs' ground currents would create data-integrity problems, which could motivate designers to provide separate IBV sources either for each POL or for small groups of POLs. Such separate sources would facilitate single-point grounding. So far, though, system designers apparently haven't recognized a need for such schemes. Should the need arise, however, the IBA offers the possibility of a moderate-cost remedy. Power-supply manufacturers can design a BC that produces two isolated 10A outputs in the same-size package as a BC that produces one 20A output. The second output would add minimal cost even if the BC included regulation to compensate for input-voltage variations.

The $64 million (minimum) question is whether IBAs really hold the potential to wipe out the power-module industry. Power-IC manufacturers, of which you will find a partial list in the **sidebar** "For more information," believe that many companies that buy modular dc/dc converters are wasting their money and that these companies could save large amounts by using the IC manufacturers' reference designs to build their own converters—especially POLs.

Module manufacturers now agree that this statement is sometimes true—but not as often as most designers who work outside power-supply design believe. A module company cites the following experience with a large system manufacturer, which finds designing power supplies for use in its systems important enough for it to maintain a power-supply-design department. When the annual quantities of a design exceed 1 million units, and the supplies' output is less than approximately 6A or 20W, the customer

## IT'S THE MONEY, STUPID!

*By Chuck Sabolis, Datel Inc*

A recent e-mail from a company hawking market research revealed (albeit preliminarily, because the group's research is not yet complete) that IBA (intermediate-bus-architecture) usage will grow from 2.8% of power-system architectures today to 17.1% in two years. The market-research company shrewdly does not reveal the more important five-year-growth number. Only paid sponsors get to see that one.

Nevertheless, growth from 2.8 to 17.1% in two years is a 611% increase and a 247% compound-annual-growth rate. If IBAs keep growing at that rate, they will represent a whopping 257% of the total market in five years.

Still, a span of two years is mere milliseconds compared with the glacial pace of change characteristic of power-supply markets. IBAs are sweeping the market, and they should. Their principal driver is the most powerful of all incentives: They save money.

What's surprising is that many sales and marketing types are reluctant to openly tell customers about IBAs' potential cost savings. Publicly, at least, power-supply vendors persist in maintaining a death grip on the high ground of "improved technical performance" (better load regulation, lower noise, quicker transient response, and so on).

Although the claims are true, the power-supply community apparently would rather pretend that performance is the *whole* IBA story rather than acknowledge that most engineers pursue IBAs under orders from all-powerful program- and supply-chain managers whose goal is to cut–by as much as half–the cost of onboard power.

Why do many presentations and trade-journal articles that compare the costs of "old" and "new" distributed power use "going-rate" prices for the old isolated dc/dc converters and published prices for the new bus and POL (point-of-load) converters? Years of fierce competition and savvy purchasing managers have beaten down the going-rate prices, whereas the published prices are still enforceable because the new products have yet to experience commoditization. Such comparisons shortchange IBAs, whose low cost will become even more obvious once market forces take effect.

Is the power-supply industry's reluctance to position IBAs as a money-saver the fallout from the painful price erosion that maturation and commoditization have brought to the industry? Is the industry trying to delay the falling of the second shoe?

Alas, the developers of the whole IBA scheme invented it to save money. Sure, some, such as those who have found no other way to deal with inductance between the regulator and the load, truly require its improved performance, but even they want the lower costs that go with the architecture.

The fact is: An IBA has the potential to save money whenever you require more than one voltage. (However, simply deriving one voltage from another may not qualify as a full-blown IBA.) The relentless pressure to reduce costs will inevitably drive every engineer to discover the architecture's merits.

Shoveling sand against the tide has always been a poor investment of time and energy. For most dc/dc-converter users, IBAs save money. It's high time for the manufacturers to promote them that way.

AUTHOR'S BIOGRAPHY
*Chuck Sabolis is director of marketing at Datel Inc.*

designs and builds its own supplies as part of larger pc boards. The module maker says that the customer builds these supplies for approximately one-third of the price the module maker would have to charge.

Even though it has a power-supply-design department, the system manufacturer must see large cost differentials before it can justify the expense of developing a supply in house. This company buys supplies for which its annual requirements exceed 100,000 pieces—even when these supplies' output does not exceed the 6A/20W limit. In short, this system manufacturer has found that economics favor its building supplies only when the output power is modest and the quantities are large. If this company is representative of system manufacturers in general, the power-module business may never return to its former size, but it will remain sizable nevertheless.

Packaging is an essential part of the DPA picture. Unlike isolated converters, in which—at power levels higher than approximately 20W—half-, quarter-, and eighth-brick modules dominate, noniso-


**Manufacturers such as SynQor provide POL modules that comply with Intel's specs for its Itanium2 processors. This HyperQor unit operates from 48V dc and delivers as much as 130W.**

## DO YOU NEED AN INTERMEDIATE-VOLTAGE BUS?

*By Conor Quinn, PhD, Artesyn Technologies*

Depending on the power they require at the subsystem level, most **computer systems** use one of two dc-distribution voltages: 48V or 12V. These voltages are isolated from the ac line, and both are reasonably well-regulated (typically to less than ±5%). The breakpoint between 12 and 48V usually falls at 1000 to 1500W. Below the breakpoint, 12V dominates; above it, 48V dominates. A few large, high-end systems distribute 400V dc.

Although 12V systems facilitate the use of today's POL (point-of-load) converters, the most heated debate is taking place within the community that uses 48V. To take advantage of low-cost POL products that are manufactured in the largest volumes, some companies are moving toward using a bus converter to provide a 12V intermediate bus. Others insist on staying with single converters to convert directly from 48V to the low IC-supply voltage. Architectural considerations based on redundancy and fault tolerance typically drive this choice.

Industry observers often ask whether the emergence of server-blade architectures will significantly affect power architectures. So far, the answer is undetermined, because no standard blade architecture has yet emerged. Hence, form factors and power levels cover a broad range, and the issues discussed in the previous paragraphs apply.

The computing industry is also discussing alternatives to 12V. For example, 12V buses present major problems when the POL converters must produce voltages that approach 1V. The industry has proposed a 7 to 8V range as technically superior to 12V, but the large infrastructure associated with 12V, which is also used to power fans and disk drives is holding back support for 7 to 8V.

In general, **telecom equipment** uses as its power source an unregulated supply of −48V (typically, 36 to 75V) or 24V (18 to 36V). Mechanically, card cages that allow little space between cards characterize the equipment. Because it limits the flow of cooling air, this physical architecture dictates that individual cards dissipate only modest amounts of power.

At typical power levels of 50 to 100W per card, even supplies that deliver voltages quite a bit lower than 12V produce manageable currents. Hence, in many telecom applications, it is common to use "brick" modules to generate 3.3 or 5V for distribution. POLs that run from the 3.3 or 5V rail produce the lower voltages. Although 2.5V is now common, the associated high currents and problems in designing POL converters to run from such a low voltage limit its use as a bus voltage.

Higher power telecom cards can also use a 12V intermediate bus. In this case, the bus converters usually provide some degree of regulation to overcome the 2-to-1 input-voltage range. Manufacturers have also discussed lower intermediate-bus voltages, such as 7 or 8V for these systems. Because the fans and disk drives don't usually operate from 12V, fewer reasons exist for using a 12V bus in telecom than in computing.

Some system designers are passionate about the use of IBAs (intermediate-bus architectures) and are equally passionate about the choice of an IBV (intermediate-bus voltage). Here are the key arguments:

***Do you need an IBV between 48V and the low voltage?*** **Yes**, if you want to leverage the low cost of high-volume, 12V POL products. **Yes**, if you want to pay for isolation only once. **No**, if you want the highest efficiency and you don't want to waste power by converting twice. **No**, if you believe that 48V architectures can provide superior levels of redundancy. **Yes**, if you believe that 12V POLs are simple and reliable enough to eliminate the need for redundancy.

***Should you consider an IBV other than 12V for enterprise systems?*** **Yes**, because converting 12V to approximately 1V is inherently inefficient and difficult to control with simple low-cost topologies, for example, those that don't use step-down transformers; 8V would be preferable and would probably result in lower cost POL designs. **No**, if your system already depends on 12V for running disk drives and fans; the last thing you need is another bus voltage. **Yes**, if your systems don't use disk drives and either provide airflow externally or can provide fan power from another source.

***Should you stick with 3.3V for lower power cards?*** **Yes**, if you can get all the power you need from 3.3V quarter- and eighth-brick modules and can scatter small POLs around the board for other voltages. **No**, if you find, as many designers do, that the number of ICs requiring 3.3V is diminishing and that more ICs are using 2.5V. POLs that use 2.5V as an input are impractical, and some boards need as many as 10 voltage levels, including 5V. For these boards, you can use loosely regulated 12 or 8V) from a small, inexpensive bus converter and use POLs to produce the lower voltages the ICs require.

AUTHOR'S BIOGRAPHY

*Conor Quinn holds PhD and MS degrees from the University of Minnesota–Minneapolis and a BE from University College (Cork, Ireland).*


**It may or may not be an MCM (multichip module); Vicor won't say. It's small enough to be one, though. The unit produces 200W in a 21.5×32×6 mm package.**

lated converters have yet to coalesce around a de facto packaging standard. For a long time, it looked as if SIP (single-in-line-package) modules that mount perpendicular to the plane of the large board might become the norm among nonisolated POLs. Now, however, it is difficult to make the case that such packaging will dominate. Still, betting against the emergence of *some kind* of de facto packaging standard appears unwise.

Among the factors that are entering the packaging picture is the trend toward lead-free soldering. Ericsson was the first major power-module supplier to offer lead-free dc/dc converters. Lead-free solder melts at higher temperatures than do conventional tin-lead solders. In addition, the range of melting points that lead-free soldering processes can use is smaller than the corresponding range for tin-lead solders. If you use a lead-free process to attach small power modules to a large board, the solder within the modules is likely to reflow as you solder the modules in place. Thermally shielding the small modules can mitigate problems that would otherwise result from this reflow, however. Eventually, all board assemblers will understand how to use lead-free processes to perform feats that are currently uncommon with such processes—for example, soldering small SIP modules to large boards.

Whether MCMs (multichip modules) will enjoy more success in the power-module arena than they have enjoyed elsewhere is another subject of lively debate. Three companies, International Rectifier, Philips, and Power-One offer POLs fabricated with MCM technology. Vicor offers several products, including the PRMs and VTMs of its FPA family as well as a family of BCs in the very small VIC format. Vicor won't say whether it based VIC on MCM technology.

Power-One points to the fact that its postage-stamp-sized MaXyz MCM POLs fit neatly on the underside of large boards, where board designers can position them immediately beneath the ICs they power, reducing to an absolute minimum the lead length between the dc/dc converter and its load. Because Vicor's VICs connect to the host board via solder bumps on a thin section around the device periphery, you can make a rectangular cutout in the host board to accommodate the thick part of the VIC. This "in-board" mounting technique can limit the thickness of the loaded host board to just 6 mm, which is the thickness of the VIC package itself.

If you decide to design your own POLs, you need to contemplate several interesting packaging alternatives. The most basic question is whether you should mount the components that make up the dc/dc converter directly on the host board or create a subassembly similar to a purchased module. Although using a separate small board or substrate may initially appear to add unnecessary cost, after you consider the differences between the test and burn-in processes for dc/dc converters and large digital boards, you may conclude that the approach saves money.□

**The Austin SuperLynx and MicroLynx families are popular POL converters. Tyco manufactures SMT versions that mount parallel to the host board and through-hole-mounted SIP versions.**

REFERENCE

1. Strassberg, Dan, "Tiny packages, new architectures rock distributed power," *EDN*, June 12, 2003, pg 18.

AUTHOR'S BIOGRAPHY

*Before he joined* EDN, *Senior Technical Editor Dan Strassberg, who holds BSEE and MSEE degrees and is a registered Professional Engineer, designed the power subsystems of large automatic-test systems and managed the development of small, linear power supplies for a manufacturer of analog-circuit modules.*

## FOR MORE INFORMATION...

For more information on bus converters, non-isolated POL (point-of-load) converters, and ICs used in POL converters contact any of the following manufacturers directly, please let them know you read about their products in *EDN*.

**Analog Devices Inc**[2]
1-800-262-5643
www.analog.com

**Artesyn Technologies**[1]
Tel: 1-888-283-3122
www.artesyn.com

**Astec Power**[1]
1-888-412-7832
www.astec.com

**Cherokee International**[1]
1-714-544-6665
www.cherokeellc.com

**Datel Inc**[1]
1-800-233-2765
www.datel.com

**Ericsson Inc Power Modules**[1,3]
1-877-374-2642
www.ericsson.com/products/powermodules

**Galaxy Power Inc**[1]
1-508-870-9775
www.galaxypwr.com

**International Rectifier**[2,3]
1-310-252-7105,
www.irf.com

**Intersil Corp and Elantec Product Group**[2]
1-888-352-6832
www.intersil.com

**Lambda Electronics**[1]
1-800-526-2324
www.lambdapower.com

**Linear Technology**[2]
1-408-432-1900
www.linear.com

**National Semiconductor**[2]
1-800-272-9959
www.power.national.com

**Philips**[2,3]
1-800-447-1500
www.semiconductors.philips.com

**Power-One**[1,3]
1-805-987-8741
www.power-one.com

**RO Associates Inc**[1]
1-800-443-1450
www.roassoc.com

**STMicroelectronics**[2]
1-781-861-2650
www.us.st.com

**SynQor**[1]
1-888-567-9596
www.synqor.com

**Texas Instruments**[1,2,3]
1-800-477-8924
http://power.ti.com

**Tyco Electronics Power Systems Inc**[1]
1-800-843-7497
www.tycopower.com

**Vicor**[1,4]
800-735-6200
www.vicorpower.com

**Vishay Siliconix**[2]
1-408-988-8000
www.vishay.com

**XPiQ Inc**[1]
1-508-429-9883
www.xp-iq.com

**Notes:**
[1]Modules assembled from discrete devices.
[2]ICs.
[3]Multichip modules.
[4]The company does not reveal the packaging technology used in its latest products, but they are small enough to be multichip modules.