Power Systems Design EUROPE

Empowering Global Innovation

June 2009

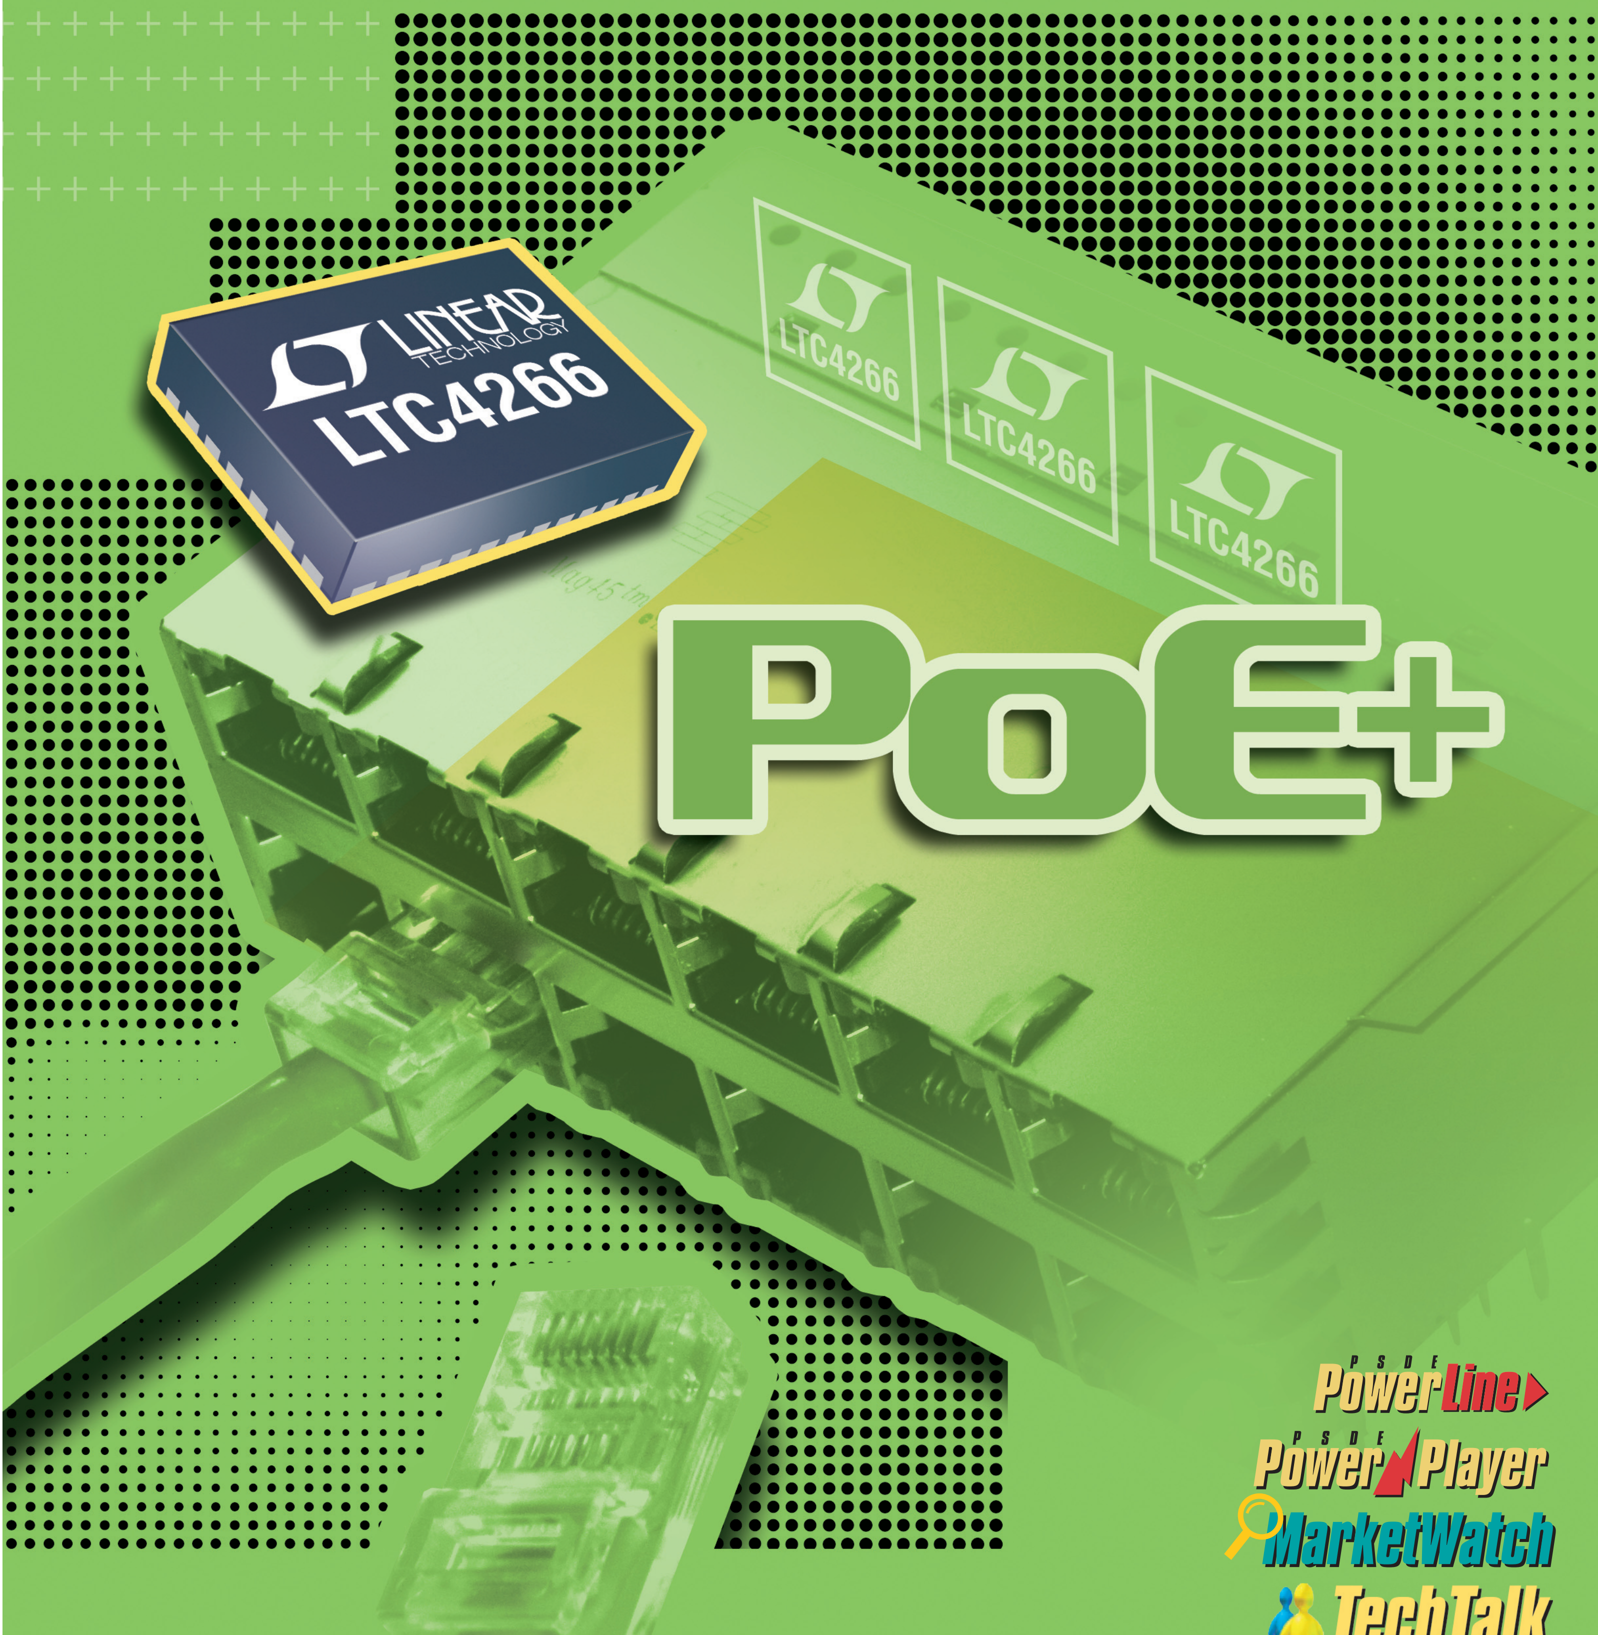

PSDE PowerLine

PSDE PowerPlayer

MarketWatch

TechTalk

Design Tips

PSDE Green Power ENERGY EFFICIENCY

Special Reports - GreenPower Leadership Awards & Green Power

ISSN: 1613-6365

The Challenges of PoE+

Standardized PSE module simplifies switch design

The IEEE is close to completing the PoE+ standard and network equipment makers are rushing to upgrade their designs, but making the transition can be challenging. A new industry-standard PSE module makes the job a lot easier, reduces time-to-market, and simplifies testing.

By Alison Steer, Product Marketing Manager,
Mixed Signal Products for Linear Technology Corp., Milpitas, CA.

PoE+ is nearly here

The Power over Ethernet (PoE) market has grown tremendously over the past few years. PoE has become almost ubiquitous, with millions of PoE-enabled switches installed all over the world.

The primary application of PoE is still to remotely power IP telephones and wireless access points. Engineers have dreamed of using PoE for many other applications, but too often these dreams have been frustrated by the small amount of available power. Under the original IEEE 802.3af standard, a Powered Device (PD) could only draw up to 12.95W.

The IEEE is about to improve the situation with the eagerly anticipated 802.3at revision (sometimes called PoE+) that is nearing completion. This latest revision will increase the power limit, allowing a PD to draw up to 25.5W, and will open the door to a host of high volume applications such as Pan-Zoom-Tilt (PZT) cameras, multimedia kiosks, industrial controllers, and laptop battery chargers.

Making the transition to PoE+

The challenge now is for Power Sourcing Equipment (PSE) manufacturers to get those high-power PoE+ ports into the field rapidly. High-power PDs won't become commonplace until high-power PSE ports are widely available.

Upgrading an existing PSE design for PoE+ requires:

- Improved Ethernet magnetics that can take more bias current without increased bit error rates at full gigabit line rate.
- New PSE controller chips with higher cutoff current thresholds.
- Depending on which controller chip is used, larger MOSFETs with larger Safe Operating Areas (SOA) may be needed.
- Larger main power supply.
- Miscellaneous components such as connectors, fuses, common-mode chokes, transient voltage suppressor diodes, current-sense resistors, and EMI filters may need to be upgraded for higher currents.

These components are already available and vendors have tried to make their new PoE+ magnetics and chips simple drop-in replacements for 802.3af components as much as practical. But unfortunately, upgrading a PSE design for PoE+ will rarely be as simple as changing the bill of materials; usually, significant PCB layout changes are needed.

For example, designs that use discrete Ethernet magnetics may need layout changes. The traces that carry gigabit Ethernet data and power from the RJ45 connectors to the transformers must have controlled impedances, but also must be heavy enough to carry

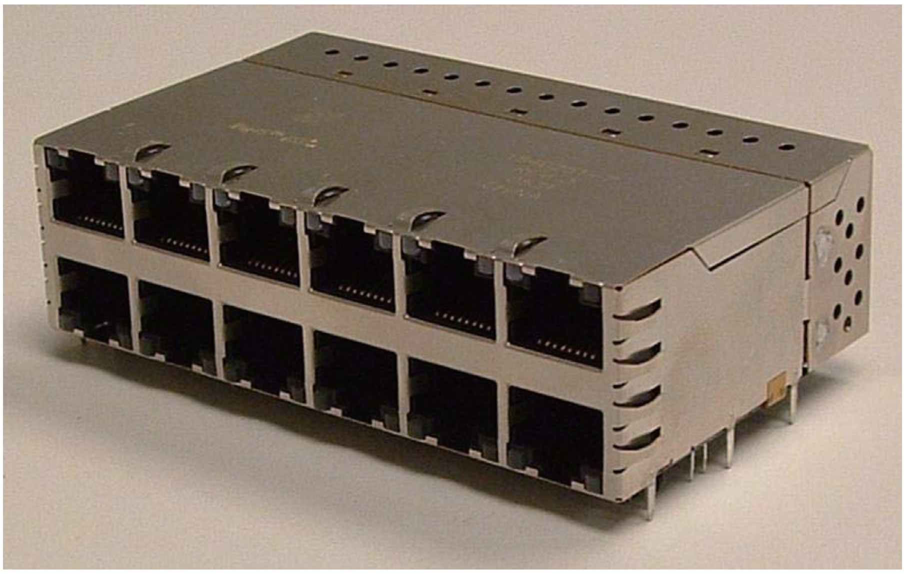

Figure 1: 12-Port PSE-ICM, COURTESY OF Tyco Elecronics.

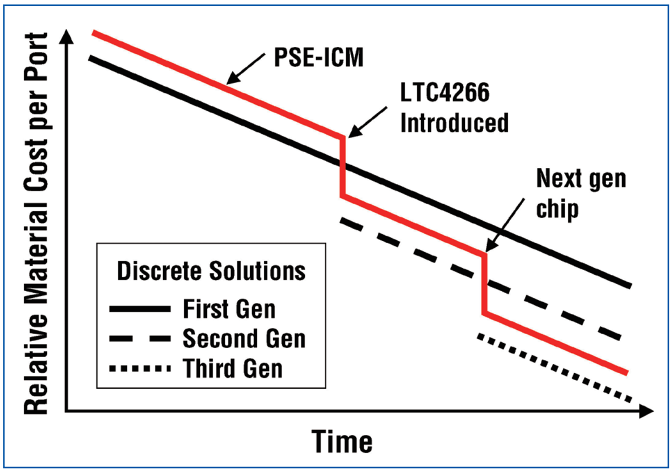

Figure 2: Block Diagram.

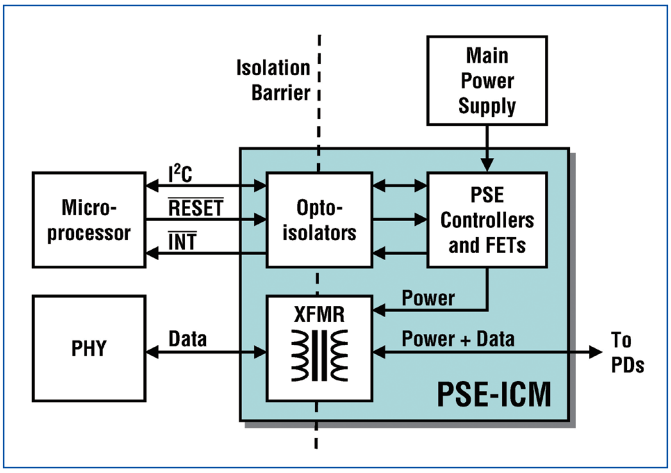

Figure 3: PSE Cost vs. Time.

the increased current. Many existing designs use single-ended 50Ω traces on inner layers which are typically 6 mils wide in 0.5 oz. (14g.) copper; some designs use closely coupled differential traces where the widths are even narrower. While these layouts may have worked at 400mA for 802.3af, there could be a serious risk of overheating at 600mA for 802.3at. Therefore, a PoE+ switch may need to route these traces on the outer layers where 50Ω line widths are typically 8 to 10 mils. What is more, surface layers are typically 1.5 oz. (42g.), after plating.

But the list of headaches for the switch designer is even longer: the increased heat dissipation may require additional or stronger fans; the whole industry is under pressure to make network equipment more energy efficient; and the upgraded PSE design will have to repeat qualification and certification testing. All of the above mentioned tasks can put a significant burden on switch designers who, in many cases, are already overburdened.

One solution is to use a multiport PSE module. There have been some module assemblies available for 802.3af PSE in the form of DIMM cards, or power supplies with PSE port circuitry built in. But these types of modules leave the designer some significant challenges because they don't include the Ethernet magnetics or RJ45 connectors. The Ethernet signals must be carefully routed from the connectors through the magnetics to the PHY chips, and power must be routed from the transformer center-taps to the PSE port circuits. As mentioned above, this can be tricky: maintaining controlled impedances, maintaining clearances for high potential (hipot), and making the traces heavy enough to carry maximum current under worst-case thermal conditions is not trivial.

PSE Integrated connector modules

Probably the most elegant approach is to put all of the PSE circuitry and Ethernet magnetics inside a ganged connector assembly. This really simplifies the task of laying out a board because all of the Ethernet signal pins are on the PHY side of the transformers; they don't carry DC currents, so you don't have to worry about the ohms per square of the traces, and you don't have to worry about maintaining clearances for hipot. Just route these signals directly to the PHY chips as normal impedance-controlled traces.

Some PSE modules like this have been available for 802.3af switches but one of their main drawbacks has been a lack of standardization. Each vendor has a specific footprint and electrical characteristics. Once you pick a vendor you're locked into their design.

But that's changed now, thanks to PoETec. PoETec is a consortium of leading manufacturers of network equipment and components, dedicated to advancing and promoting PoE technology. PoETec has developed, and will soon publish a specification for the industry's first standardized PSE module, which they call the PSE Integrated Connector Module (PSE-ICM). The specification defines all aspects of the PSE-ICM characteristics including footprint, signal functions, and the internal register set. So a PSE design that uses one brand of PSE-ICM can simply drop in another brand without changing the board layout or the system software.

Figure 1 shows a 12-port PSE-ICM from one vendor. At the time of this writing, two PoETec member companies (Molex and Tyco Electronics) are shipping PSE-ICMs and two more companies are about to start shipping. Presently, there are 12-port and 8-port PSE-ICMs available; 16-port PSE-ICMs may be coming in the near future. There are also versions with and without LEDs. Passive modules without the PSE circuitry, just the magnetics, are also available.

Figure 2 shows a simplified block diagram of a 12-port PSE-ICM. It includes an isolated I²C interface for control and monitoring of PSE functions. Also included, but not shown, are common-mode chokes and terminations. All that's needed is the main power supply and an external microprocessor to run the power management software. The PSE-ICM can also be configured for AUTO mode in which standalone operation is achieved. In this mode, the external microprocessor is not required.

The PSE-ICMs were made practical by some pretty advanced technology, and probably couldn't have been built at a reasonable cost just a year ago. One key enabling technology is the new

LTC4266 quad PSE controller chip from Linear Technology Corp. The LTC4266 has the smallest package (5x7mm QFN) and lowest power dissipation of any quad PSE controller in the industry: just 165mW/port at 600mA, including current-sense resistors and MOSFET on-resistance. What is more, the LTC4266 has a unique non-linear foldback feature that protects the module from short circuit faults; without this feature, larger MOSFETs with bigger Safe Operating Area (SOA) would be needed to support the higher current levels reliably.

PSE-ICM advantages

These days, switch designers face many tough engineering challenges. All of the digital and software tasks are difficult enough, without having to worry about analog issues. For example, hipot, EMI, UL certification, lightning surge protection, and heat dissipation are several areas where problems often crop up near the end of a project, when it's most costly and time-consuming to fix them. In fact, these are probably the most common reasons why products miss their launch dates.

The primary advantage of the PSE-ICM is that it's already been through all these tests. Therefore it not only reduces the switch designer's workload, but also reduces the risks of last minute problems. Of course, it's still possible for a switch to fail hipot or EMI because of layout issues outside the PSE-ICM, but the chances of that happening are reduced.

The PSE-ICM also reduces risk because there are multiple sources. This not only creates price-competition, but reduces the risk of late delivery that sometimes occurs when using sole sources for components.

Poor technical support can also lead to project delays and cost overruns. Suppose you're testing a new PSE prototype and one of the MOSFETs overheats. Was the failure due to a bad MOSFET or a bad controller chip? The two vendors will likely point their fingers at each other, while your project slips further behind schedule. But with the PSE-ICM there is no finger pointing; if one PSE-ICM brand seems unreliable, you can simply switch to one of the others with no changes to your board layout or software.

The cost issue

Of course cost is king in the network equipment industry. Some designers may take one look at the PSE-ICM and say it's too expensive, but a smart designer won't be so hasty. The real objective is to reduce overall costs; that's rarely as simple as picking the cheapest parts.

Figure 3 shows a qualitative cost comparison between a discrete design – where the designer places all the PSE components (controller chips, FETs, magnetics, etc.) on the main board – and a design that uses the PSE-ICM.

The graph is intended to illustrate three points:

• The costs of both alternatives decrease over time.

• There are sudden drops in cost when new technologies are introduced.

• The material cost of the discrete approach will always be slightly lower.

Think about the second bullet. One of the advantages of the PSE-ICM is it allows you to more easily keep up with advancing technology and take advantage of the savings. For example, when a new chip that reduces cost is introduced, it may be difficult to use that chip in the discrete design because a new PCB layout would be needed. But when that chip is incorporated into a new PSE-ICM, you can just drop it on your board because the PSE-ICM footprint hasn't changed.

Now think about the third bullet. The discrete approach has a slightly lower material cost, but the PSE-ICM offers a lot of added value that makes up the difference.

For example:

• Shorter time to market. Some products have market windows only 8 to 12 months before they're obsolete. If the launch of such a product were delayed just 2 months, due to PCB layout problems for example, then the overall revenue from that product would be severely reduced. The PSE-ICM can greatly reduce the risk of these delays, and that has economic value that should be counted.

• Lower assembly and test costs. The PSE-ICM obviously reduces assembly cost, but its benefit of reducing test costs should not be overlooked. A designer who goes with the discrete approach must develop test setups and software sufficient to catch all the defects that might occur during the assembly process: at a minimum one would need to verify basic functions such as detection, classification, and disconnect sensing. All these functions are pre-tested in the PSE-ICM; the only testing required would be to verify there are no bent pins or bad solder joints when the PSE-ICM is stuffed on the board. You get almost complete coverage just by running Ethernet traffic on all the ports and verifying the PSE-ICM acknowledges when it is addressed via the I^2C bus.

• Lower configuration management costs. For example, PSE-ICMs are available without the PSE electronics in the same footprint. This means a switch maker can design two products (a PoE-enabled switch, and a switch without PoE) that use the same main board; the only difference being which type of PSE-ICM is stuffed.

Conclusions

As the new IEEE 802.3at standard nears completion, many companies are preparing to launch their new PoE+ products. In this environment, where designers have a long list of technical challenges and a wave of competing products is imminent, shortening the design cycle time can be critical for success.

The PSE-ICM can greatly reduce time to market, but offers many other advantages described in this article. The two main drawbacks of previous modules were lack of standardization and high cost, but both of these are addressed by the new PoETec industry standard: Multiple sources are already on line, with more coming. The new PSE-ICMs are already lower cost than their predecessors, and over time competition and technological advances will drive costs down even further.

www.linear.com