



Demanding Dynamic Loads: Can Power Devices Keep Up?

With each new generation of processor, the trend is toward lower voltages, higher currents, and faster dynamic loads. As a result, power-system designers are challenged to provide ever-faster transient response. They also have to do it using less board area while providing cost-effective and efficient power-systems solutions that offer the requisite performance. The question is if power devices can keep up.

Power designers traditionally responded to the need for fast, dynamic loads by putting very simple energy storage—one or more capacitors—right at the point of load. This often addresses the issue of delivering energy very quickly to a device, but there are tradeoffs. Capacitor technology hasn't moved forward as quickly as many other technologies, so a large amount of real estate is required to accommodate this capacitance.

The reliability of the capacitor, especially when there may be dozens of them in high-performance applications, is another concern. Simply adding more parts affects mean time before failure (MTBF) calculations, decreases available board space (most often in critical areas of the board), and adds cost. Capacitor failure could compromise the system.

DESIGN REQUIREMENTS

Often while addressing severe dynamic loading conditions, you may need to maintain the voltage within a certain regulation band. Also, you may need to be able to deliver energy instantaneously, or absorb load dumps, where the current at full load drops down to no load at all.

Typically, power-supply control loops are limited in bandwidth to a frequency significantly less than the switching frequency. Nevertheless, power-conversion manufacturers have been trying to achieve faster control loops and higher frequency to achieve better transient response.

One of the first steps in improving the transient response of "brick" converters was removing most of the internal capacitance within the converter and locating it at the load. The parasitic inductance between the converter and the point of load capacitance could then appear as part of the output inductor.

This allowed a great increase in the closed loop bandwidth of the converter while permitting it to be some distance away from the point of load. For some applications, this approach provided the desired faster dynamic performance, but it was gained at the expense of significant additional design and implementation complexity.

A lot of work is going on with digital power conversion right now, and there eventually could be some significant improvement in faster dynamic load response. So far, there have been small increments of improvement.

While some designers are trying to improve control loops and develop better controllers, the most prevalent way to improve converter transient response is with multiphase-type buck converters. Using multiple power trains controlled by a control chip, the apparent frequency of the combination can be multiplied, resulting in a number of favorable advantages over the single power-train approach. Most important, the frequency is increased without also increasing switching losses, helping to achieve a faster load transient response.


Bulk energy storage is still needed with this approach, and more phases add complexity, both in terms of control and more components. The control device not only manages the phasing of the individual voltages, it also is needed to ensure current sharing. The additional components play into the reliability equation as well. The multiphase solution has survived, at least in part, because the intense competition has driven costs down. At this point, multiphase is an acceptable approach.

Rethinking the architecture/topology of power conversion has produced significantly faster transient response. The Factorized Power Architecture employs a separate voltage transformation stage that enables a module with no external control loop. It's a fixed-ratio converter. This, and an effective switching frequency of 3.5 MHz, creates a very powerful platform for delivering energy very quickly to a dynamic load.

WHAT'S NEXT?

The Factorized Power Architecture was a rather significant leap. Over time, there will be improvements in multiphase buck converters or other solutions. It's inevitable. There'll be better products. More clever design will go on.

Finally, there must be a practical limit on processing speed. The single-core microprocessor demanded higher and higher currents: over 100 A, with talk about well over 150 A. Manufacturers recognized that microprocessors demanding such high currents could be at the mercy of fewer power-supply makers who could achieve that performance and name their price.

So, they rethought the concept and made a conscious effort to divide and conquer, yielding the birth of dual-core and multicore processors. When they split the cores up within the processor, they can split up the power supply as well. Now you may only need 50 A. Of course, they had to make the cores work together, and they did! As a result, fast dynamic load response isn't quite the problem that it could have been with the single-core processor following Moore's Law. 

TOM CURATOLO holds a BSEE from Worcester Polytechnic Institute, Worcester, Mass. Look for more columns from him in print and online.