ISSUE: June 2022

Advanced PMICs Can Flatten The Processor Power Curve For Data Centers

by Trey Roessig, Empower Semiconductor, San Jose, Calif.

The growth of data-intensive and high-data-volume applications, ranging from artificial intelligence through social media to the internet of things (IoT) has created an insatiable need for ever-more data center processing capacity. As well as increasing the number of servers, an element of this need will be met by new energy-hungry processors, which, in turn, will drive up the power usage per rack.

However, with data center operators under pressure to keep energy consumption and costs to a minimum, they are looking for the most efficient ways to address these increased power demands and to break the linear relationship between increased processor power and increased rack power.

Here we look at the factors driving increased rack powers, and the importance of voltage regulator transient response and dynamic voltage scaling (DVS) in maximizing server energy efficiency. We then introduce a PMIC technology that can significantly reduce the overall power consumption needed to drive next-generation processors.

This technology, which is implemented in Empower Semiconductor's IVR family of devices, reduces the usual tradeoff of efficiency for higher switching frequency, enabling development of highly compact and very fast voltage regulators. In addition to describing the characteristics of these devices, including their ability to implement DVS, we'll discuss their benefits in powering a typical DSP application.

Data Centers And Power Density

Data center traffic is experiencing exponential growth, around 30% year-on-year, fueled by 5G, video streaming and social media use in the public arena, and in industry by machine learning, AI and the Internet of Things. The global size of the data center market by investment is predicted by Arizton Advisory and Intelligence to exceed $250B by 2026[1] and the extra storage and processing required is being built out rapidly, along with the associated power systems, cooling and back-up infrastructure.

However, costs are high—the largest data centers consume more than 100 MW and build costs are estimated at $10M to 12M *per megawatt*. So operators look to recoup their investment as quickly as possible and then keep ongoing energy costs low, while maximizing data processing throughput. Overall, in 2020, data centers consumed around 1% of the global energy supply or 200 to 250 TWh with as much again in the data transmission networks,[2] so there is also every reason to limit power consumption as far as possible to reduce environmental impact.

Data centers have a performance metric—capacity per square foot utilized—which creates pressure to pack more processing power into each rack. A 2020 survey by the Uptime Institute found that an average rack consumes 8.4 kW with some as high as 30 kW.[3] This extreme power density poses problems for cooling systems and limits further expansion of rack capacity.

In response, power designers have worked to squeeze ever-better efficiency out of existing ac-dc and dc-dc conversion technologies, aided by advances such as wide-bandgap semiconductors with their lower losses. This has remarkably resulted in overall energy consumption remaining relatively flat over the last few years despite the exponential growth in traffic and associated hardware. However, incremental efficiency improvements in power conversion are becoming more and more difficult and expensive to find, and the overall power usage per rack is beginning to "turn up" in the same exponential way as processing capacity.

Moreoever, any efficiency improvements available in a power converter are often achieved at the expense of size. For example, dynamic losses can be lower if switching frequency is kept low, but this means larger magnetic components and a larger footprint. This pushes the converter physically further from the end load (typically a processor or FPGA) degrading regulation and noise performance.

There is therefore an increased focus on improving the "innate" density in the system power architecture—reducing the size while holding efficiency constant. This can create a virtuous circle of smaller power components, a reduced cooling requirement and more space for processors and storage.

The power management ICs (PMICs) that enable efficient dc-dc conversion in a smaller footprint are key to this pursuit of higher power density, and there is now a new voltage regulator technology from Empower Semiconductor that can yield dramatic improvement, flattening the curve of rack power usage with time (Fig. 1). But before describing the characteristics of the PMICs based on this technology, let's consider some of the challenges normally encountered when trying to shrink voltage regulators.

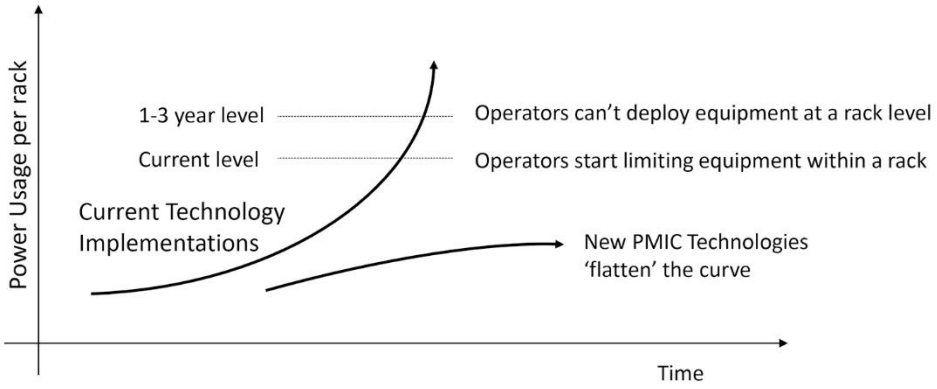

Fig. 1. Power usage per rack with existing PMIC technologies is beginning to follow the same exponential increase as data center traffic. A new approach can "flatten" the curve.

Passive Components Offer Limited Power Density Improvements

The goal of higher power density is simply to achieve more load current capability, typically from a point-of-load dc-dc converter (or POL for short) in a smaller space. Better efficiency is part of the solution, with knock-on savings all the way upstream in preceding stages of power conversion from the line rectifier, PFC stage and ac-dc converter through system bus converters.

With lower losses, a POL can be physically smaller for the same temperature rise and potentially fitted closer to the end load. This is important with modern processors operating from sub-1-V power rails at currents of hundreds of amps, demanding short connections to the power source, to avoid voltage drops. Equally, the POL must be small enough to fit in around the load heatsinking arrangements and data and address buses in the surrounding PCB.

POLs are switched-mode regulators that have zero conduction and switching losses with ideal components, unlike linear types that inherently lose power. Semiconductor technology advances have enabled step improvements in switch and diode performance, approaching the ideal, so that residual losses are low, even with tiny die sizes. And in cases where they are still not low enough, cost and performance can be traded by paralleling multiple devices, for example.

However, passive components, the necessary capacitors and inductors, have not followed the same trend as the ICs and have become a limit to better power density. A solution is to switch at higher frequency with these components scaling down in size, but at some point, the extra dynamic losses in semiconductors become the limit again, along with parasitic effects in the passive components. For this reason, there is a "sweet spot" in frequency that gives the lowest overall losses for a given size and cost. As a consequence, current POLs operate at a few megahertz at most.

With inductors, the frequency limit is related mainly to losses in the ferrite core needed to achieve a given inductance value with a reasonable number of winding turns. However, the inductance needed and therefore winding turns decrease with frequency, so there is some frequency over which a core is not needed, opening up the possibility of using an air core, which then becomes the basis for a low-loss inductor fabricated as a tiny spiral of copper tracking.

Capacitor technology is also improving at high frequency. So if the remaining problem of increasing dynamic losses in the POL's switches can be solved, the door is open to much higher operating frequency with maintained high efficiency and a consequent step improvement in power density.

Breaking Through The Frequency Barrier

The latest technology developed by Empower Semiconductor[4] has broken the frequency barrier with a patented, high efficiency, multi-phase conversion approach which is digitally configurable. Fabricating in advanced CMOS processes and designing the controller, package and passives in non-mainstream technologies and in unison helps overcome the dynamic loss problem. These devices can switch in the low hundreds of megahertz while offering relatively flat efficiency across the load range with efficiency peaks of up to 92%.

Combining up to four voltage regulators in a single compact package, the IVRs feature an input of 1.8 V or 3.3 V and output voltages that are programmable from 0.5 V to 2.5 V depending on the part chosen. An FcCSP package with a 5-mm x 5-mm PCB footprint and a height of just 0.75 mm means that IVRs can be fitted directly at the load—or even on an SoC substrate. Furthermore, unlike other PMICs that claim a small form factor, the Empower parts require no external components for full rated performance, just a core-less inductor fabricated in the motherboard PCB, typically directly beneath the IVR. In a like-for-like performance comparison, the FcCSP devices represent a 10x reduction in footprint compared with competitive parts.

Increasing switching frequency does not just scale down the size of passive components, it also allows the internal regulation control loop to have a wider bandwidth. This is limited to around one fifth to one tenth of the switching frequency to avoid instability, and in traditional designs, this leads to transient response times on the order of tens or hundreds of microseconds for switching at frequencies up to a few megahertz.

At these lower frequencies, the output inductor also limits the output current slew rate significantly. The Empower IVR design however, with its high switching frequency, has a loop response time of 100 ns, typically 100x faster than existing voltage regulator designs. The low inductor value allows an output current slew rate of 0 A to 10 A in 500 ns, with just 15-mV excursion, compared with around 50 mV for traditional technology. This is achieved without the need for any external capacitors. An example of this kind of transient response is shown in Fig 2.

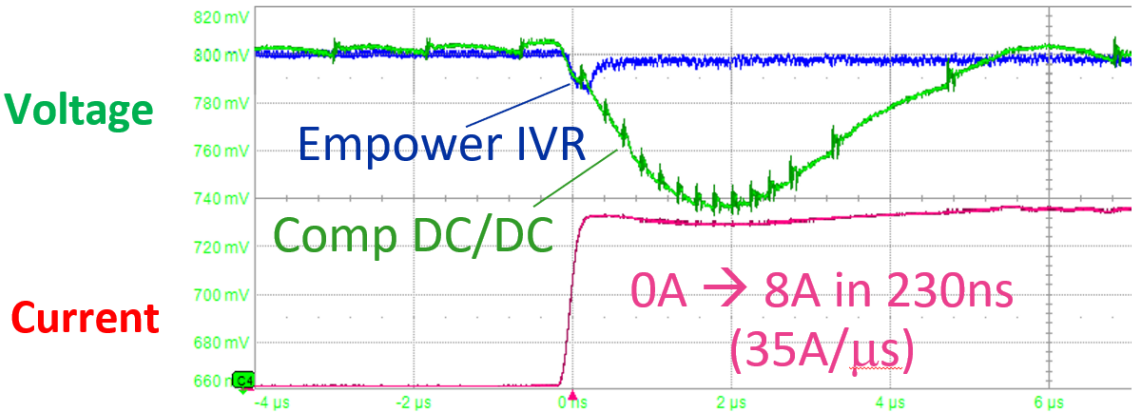

Fig. 2. Comparing Empower Semiconductor IVR transient response to the competition.

A consequence of the very low transient voltage excursion is that the nominal output voltage can be set lower than traditional POLs, which must be set so that the larger transient dip in voltage with load steps does not fall below the minimum load operating voltage. Setting the nominal load voltage lower can represent significant energy savings.

Dynamic Voltage Scaling At Speed

Dynamic voltage scaling or DVS is a technique whereby a power rail is adjusted on-the-fly, to minimize power dissipation, depending on the load conditions. With traditional POLs, this has been of limited use, as load demand changes can be rapid and the traditional converter control loop has not been fast enough to track. However, the fast control loop of Empower IVRs opens up new possibilities.

One example is when a processor is waiting for cached data. The Empower part can rapidly slew its output voltage downward to reduce dissipation in a few nanoseconds. Older technology POLs with response times around 1000x slower just cannot react quickly enough to obtain any advantage. The result could be a 65% energy saving during the wait cycle while clocking continues.

Fig. 3 shows the situation. V_{PMIC} is the voltage necessary from previous generation POL converters, with its nominal value set to 0.82 V to allow for the higher voltage excursions with load transients, compared with the Empower part set to a nominal of 0.74 V. During the wait cycle, this can rapidly drop under control of the digital interface to 0.57 V to achieve the energy saving. The beneficial effect is amplified by power dissipated in the load during switching transitions being proportional to the square of the voltage, as device capacitances are charged and discharged.

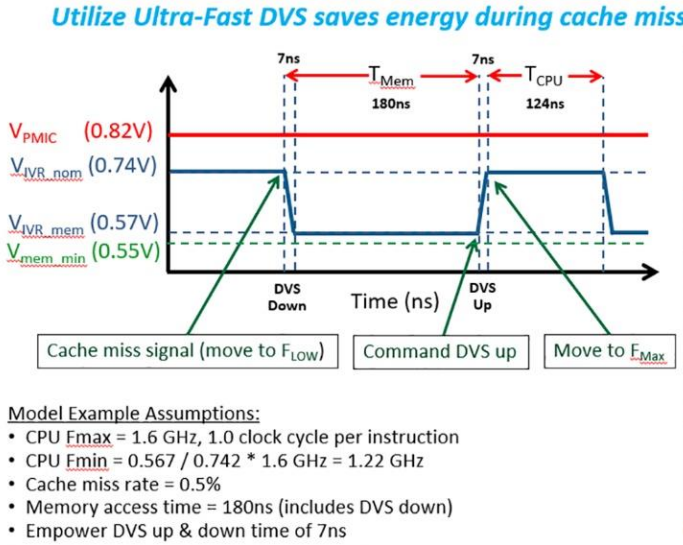

Metric	VR Type	Operating Segment: Mem Wait Cycle	Operating Segment: DVFS (Up & Down)	Operating Segment: CPU Cycle	SoC Total per Cycle
Time (ns)	RIVR	173	7 + 7	124	311
	PMIC	180	0	124	304
Avg Voltage (V)	RIVR	0.57	0.65	0.74	0.64
	PMIC	0.82	0.82	0.82	0.82
Current (A)	RIVR	4.50	5.20	7.71	6.0
	PMIC	8.5	8.5	8.5	8.5
Avg Power (W)	RIVR	2.55	3.40	5.72	3.9
	PMIC	7.0	7.0	7.0	7.0
Energy (µJ)	RIVR	0.44	0.05	0.71	***1.20***
	PMIC	1.26	0	0.87	***2.13***

65% savings during wait cycle

Fig. 3. Empower IVRs slew from $V_{IVR\ nom}$ to $V_{IVR\ mem}$ in 7 ns, saving energy.

While dedicated voltage regulators for microprocessors have used DVS techniques for more than a decade, these techniques have not really been available to the broad market, in part because of the speed of the regulators themselves. With the speed available in the Empower IVR, this now becomes a possibility.

But even I3C speeds become a bottleneck when talking about DVS times <100 ns, so the Empower family has implemented a single-pin DVS function, allowing the user to program two different output voltage registers and allowing a hardware pin to select between the two. This reduces the system overhead for DVS usage to near zero. An example of the DVS speed achievable is shown in Fig. 4.

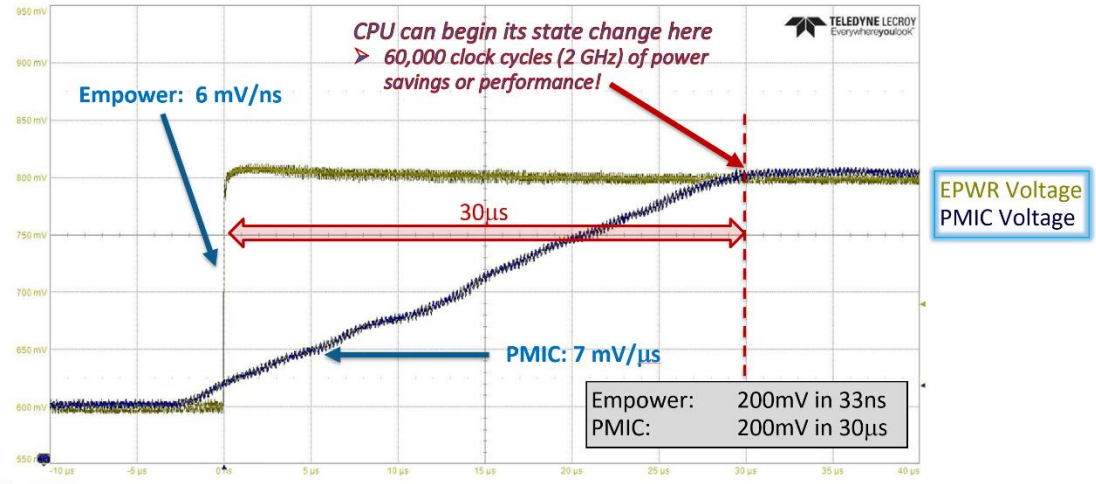

Fig. 4. Empower IVRs slew from VIVR nom to VIVR mem in 7 ns, saving energy.

Reducing Component Count

A typical voltage regulator power requirement for a DSP chip might consist of a VCSEL driver requiring a 3.3-V input, a transimpedance amplifier (TIA) operating at 3.3 V and 0.27 A and an MCU and EEPROM both requiring 1.8 V and 0.05 A. For the DSP itself, the respective voltage and current requirements for the various elements will be in the region of 1.8 V and 0.25 A for the digital I/O, 0.6 to 0.8 V and 3 A for the core, 1 V and 3 A for the analog circuitry and 1.2 V and 0.05 A for the PLL. Fig. 5 outlines how this might be implemented using conventional PMIC technology in which there are 36 discrete components taking up around 360 mm^2 of board space.

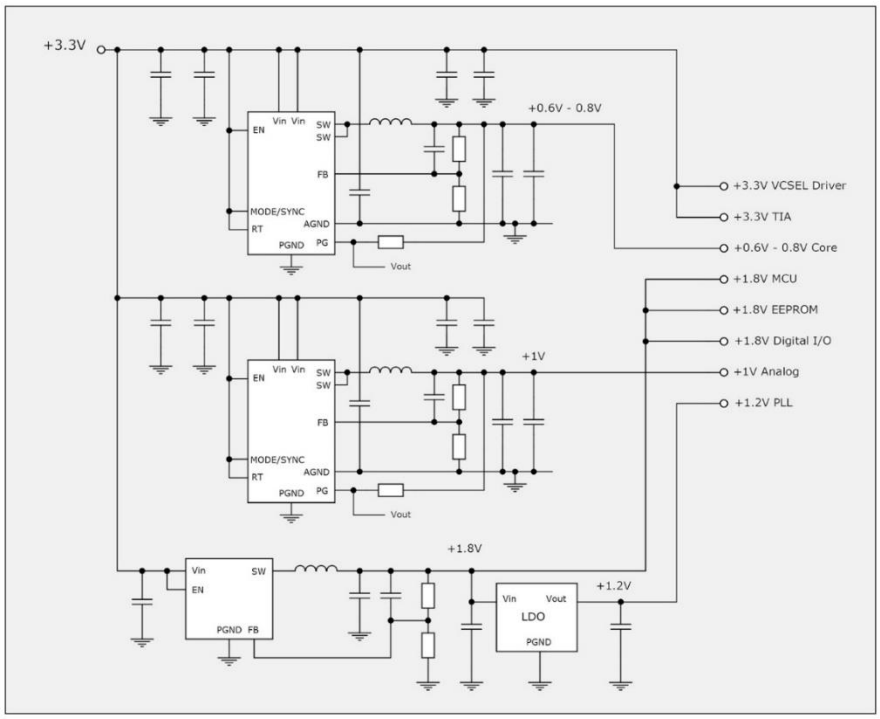

Fig. 5. Implementing a typical DSP power scheme with conventional PMICs.

Now consider what happens if we employ Empower's EP71xx series of quad-output IVRs in a similar application (Fig. 6). The EP71xx IVRs operate from a single 3.3-V input and offer output current options up to 12 A. This family comprises a variety of devices that can provide the available 12 A from one to four independent outputs. Output voltages are programmable from 0.5 V to 2.5 V and output voltage set-point accuracy is ±1.0%.

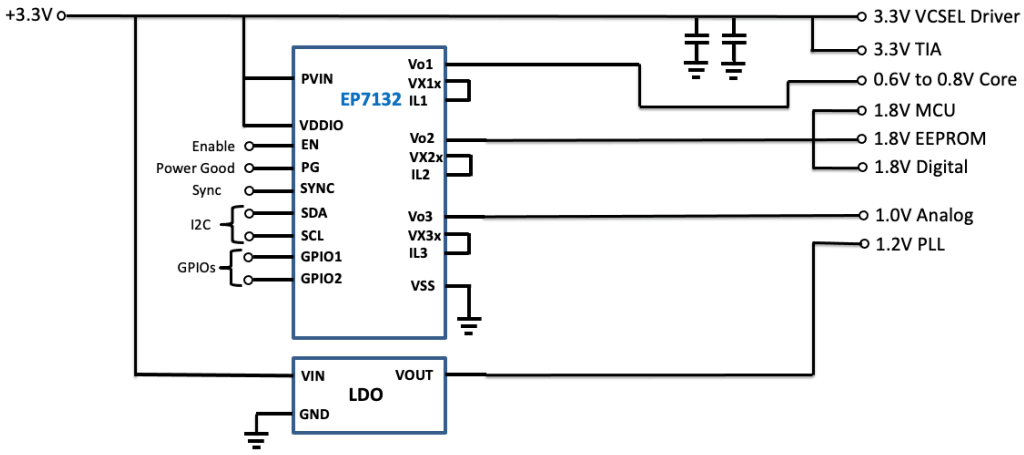

Fig. 6. A DSP power scheme with Empower EP71xx quad-output IVR.

Built-in fault protection and warning capabilities cover OVLO, UVLO, OVP, OCP and short circuits, while the programmable ExpressV DVS-on-demand is capable of operating at up to 6 mV/ns. By using the Empower technology we can reduce the total number of discrete components and reduce board space by as much as 50% to 90%, as shown in Fig 7.

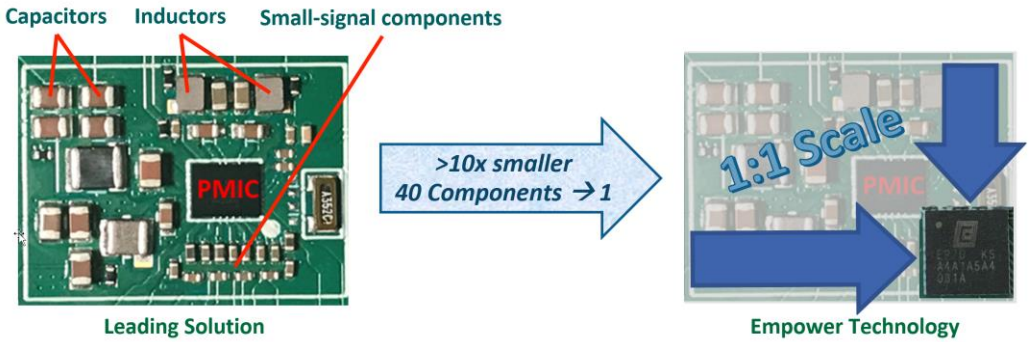

Fig. 7. A comparison of a leading three-output solution to the EP70XX family footprint.

The EP71XX devices incorporate a full suite of telemetry functionality to report voltage, current and die-temperature, and soft-start and DVS ramp rates are programmed via the digital interface. Advanced features such as user-programmable power-up/power-down sequencing, user-configurable logic levels and programmable GPIOs to enable customized system functionality and comprehensive diagnostics are available through a standard I^2C interface.

In addition to POL power for DSPs, CPUs, GPUs and accelerators, target applications for the new series will include optical transceivers, active cables, memory modules and storage devices, networking equipment such as GPON and Photonic Service Switches.

References

1. https://www.prnewswire.com/news/arizton-advisory-%26-intelligence/

2. "Data Centres and Data Transmission Networks, More efforts needed," IEA Tracking report, November 2021.

3. "How Rack Power Densities Influence Data Centre Design," Server Room environments, 02/01/2021.

4. Empower Semiconductors website.

About The Author

Trey Roessig brings over 20 years of experience and 25 U.S. patents in power and mixed-signal design to the Empower Semiconductor engineering team. In his role as CTO & SVP engineering, Trey manages a team of analog, digital, power, and system architects.

Previously, he has served in both technology and business roles in companies of various sizes, from co-founder of Integrated Micro Instruments (acquired by Analog Devices) to director of engineering at Volterra Semiconductor to general manager of ADI's Power Management business unit. He has a Ph.D. from the University of California, Berkeley in Mechanical Engineering/Controls & Dynamic Systems.

For more on designing dc-dc converters and switching regulators, see the How2Power Design Guide, locate the Power Supply Function category and select "DC-DC Converters".