ANALOG
DEVICES
AHEAD OF WHAT'S POSSIBLE™

Impacts of Transients on AI Accelerator Card Power Delivery

Hamed Sanogo, End Market Specialist

Abstract

Graphics processing units (GPUs), tensor processing units (TPUs), and other types of application-specific integrated circuits (ASICs) deliver high performance computing by enabling a parallel processing power required to speed up both artificial intelligence (AI) training and inferencing workloads.

AI requires a lot of computational power, especially when it is learning and inferring. This continues to push the boundaries of power delivery networks to new levels never seen previously. These high density workloads have become more sophisticated with higher transient demands pushing every part of the power distribution network to perform at maximum capability. The stringent power requirements of AI accelerator cards are now impacting system performance. In this article, we will review the power distribution network requirements of the AI accelerator card, dissect the impact of the transients, and present a multiphase power delivery solution from Analog Devices that addresses these requirements.

Introduction

AI is revolutionizing computing architectures to replicate neural networks that emulate the human brain. It seems that AI is already everywhere, but in reality, the technology that drives it is still developing. The processor accelerator ICs specialized for AI calculations include GPUs, field-programmable gate arrays (FPGAs), TPUs, and other types of ASICs. This article will collectively refer to all of them as xPUs.

Data centers will continue to buy AI accelerator cards in bulk as deployment of the technology is on a rise. According to Gartner, AI chip revenue totaled more than $34 billion in 2021 and is expected to grow to $86 billion by 2026.[1] xPUs offer a massive leap in AI performance compared to an ordinary CPU by using massively parallel computing implementations. Because of the high number of its small cores, xPUs are well suited for AI workloads, facilitating both neural network training and AI inferencing. However, they typically require relatively large power consumption for the AI computations and data movements. Simply put, xPUs are power hungry ICs. Their stringent power requirements are placing new demands on AI accelerator cards that now affect system performance. In this article, we will review the power delivery network requirements of the AI accelerator card and present a multiphase power delivery solution from ADI that addresses these stringent requirements.

The AI Imposed Power Delivery Challenge

AI is many things, but power efficient is not one of them. When AI is working, especially processing AI workloads such as deep learning and inferencing, it requires extreme computational horsepower. At the system level, AI accelerators play a critical role in delivering the near-instantaneous results that make them valuable. All xPUs have multiple high end cores constructed from billions of transistors and consume many hundreds of amps. The core voltages (V_{CORE}) on these xPUs have been reduced to sub-1.0 V levels. Figure 1 shows a generic block diagram of an AI accelerator card. The article will focus on the multiphase controller and the accompanying power stage ICs proposed for such a system.

The peak current densities encountered on AI accelerator cards have become extremely heavy for any motherboard to handle. The highly dynamic nature of workloads and extremely high current transients are resulting in very high di/dt and spiking voltage transients lasting for several microseconds that are highly disruptive and can potentially be damaging for the xPU. With the average AI workload lasting that long, the decoupling capacitors are not able to provide the energy to meet the instant demand the entire time. The next section of the article will present an ADI multiphase point of load (PoL) solution proposal that will eliminate the typical AI accelerator's transients that create stress across the power distribution network. But first, let's discuss the power design challenges introduced with AI.

AI Introduces New Power Design Challenges

AI power demands are currently far outpacing traditional power delivery network capabilities. The requirements for the xPU voltage regulators (VRs) are quite different from standard PoL regulators. The industry is seeing greater than 1000 A delivered at <1 V to the xPU in certain applications. It is important that the supply is very stable and produces very little noise while eliminating all possibilities for voltage transients, which can cause false triggering inside the xPU. The design of a high performance AI accelerator VR PoL, with a staggering current demand, must meet certain key requirements.
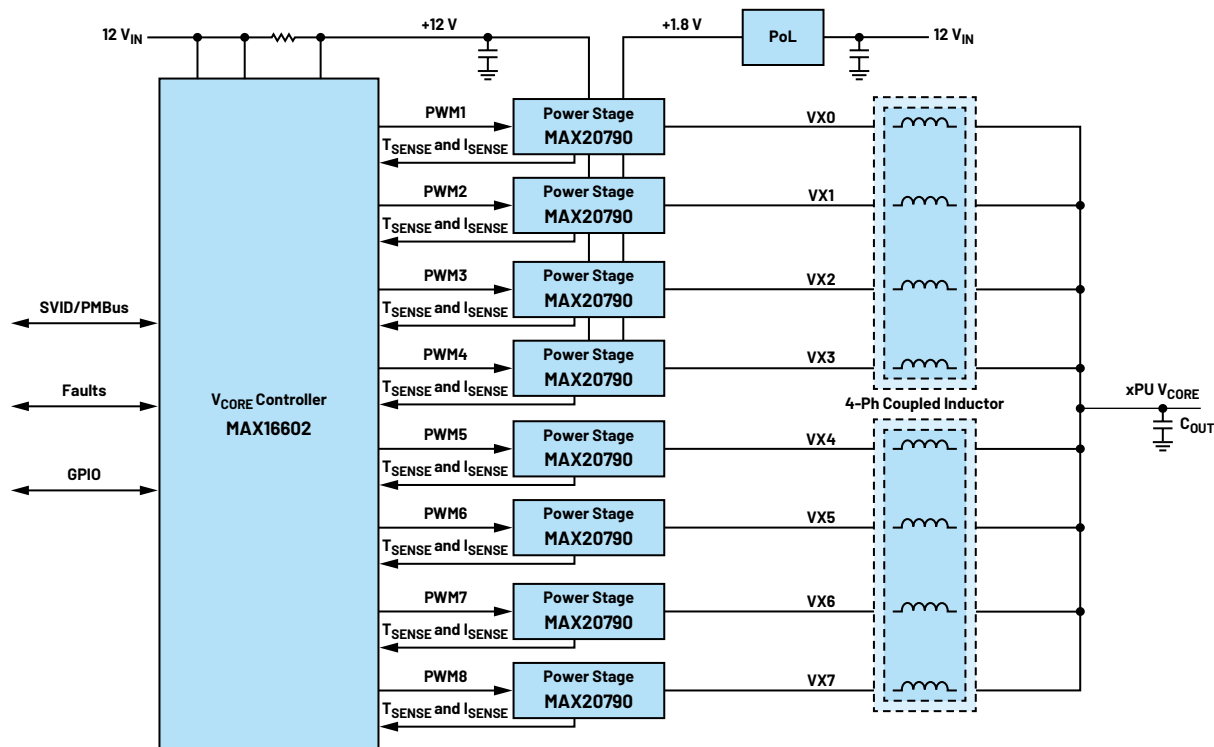
Figure 1. A block diagram of a generic AI accelerator card.

Voltage Spikes and Transients Management

One of the AI accelerator card's key requirements is to have the VR architected in such a way that it offers superior transients voltage management. Delivering kilowatts of power to any system is always a first order challenge. The output voltage including tolerances, ripple, and load transient dips and peaks must stay above the xPU minimum voltage to avoid system hang and must also stay below the xPU maximum voltage to avoid xPU damage. Transient power spikes from these cards can ask for 2× and more than the maximum thermal power target.

What is important here is that the PoL loop bandwidth is flexibly wide enough to deal with the types of faster transients seen. The higher the bandwidth, the faster the loop responds, and with less voltage deviation. One of the most straightforward methods to achieve fast transient power rails is to select regulators that feature fast transient performance. The ADI AI V_{CORE} family of ICs features exceptionally low frequency output noise, fast transient response, and high efficiency. Adding its load line support to this, the ADI AI power chipsets do an excellent job at helping power designers manage their AI workload induced transients and spikes.

I²R Losses in the Long Power Path Traces and Thermal Management

As the current of the AI xPU processor continues to increase, the density of the power delivery solution to the PoL has become a critical element. It is becoming extremely hard to deliver power reliably to every part of the xPU without worrying about the dispersed heat impacting the reliability of the chip and leading to thermal runaway. In other words, thermal management is one of the more significant challenges in designing this high wattage power supply. Traditional power delivery methods place the voltage regulator on the side of the xPU such that power is delivered laterally to the processor. Even the smallest resistance in these traces can lead to unacceptable voltage (I²R) drops. A voltage drop across the PCB power plane resistance is proportionally increased with the xPU current. This represents a few centimeters of PCB power traces between the VR and BGA pins and this is where a significant amount of losses take place. Such losses in the PCB copper power planes have become the most dominant factors in calculating the efficiency and performance of the regulator design. The use of a monolithic power stage IC, with the current and temperature circuit blocks integrated, can greatly reduce the number of high current traces required on the PCB to implement a traditional 3-chip (discrete) power delivery solution.

ADI Value Proposition: MAX16602 + MAX20790 + Coupled Inductor

The AI voltage regulator accuracy has become even tighter. Efficiency and size are high priorities. Performance and power loss are also under major scrutiny. As was presented in the previous section, solving AI accelerator card VR design problems has become a daunting task. Designers understand too well that generating large steps in demanded current cannot be satisfied without dealing with the unwanted transient effects. Addressing these transient effects also requires some type of high accuracy dynamic voltage positioning or load line scheme. ADI is heavily invested in the AI market and has a complete portfolio of solutions for both 48 V and 12 V systems. This section of the article proposes the ADI AI multiphase power chipsets, the MAX16602 multiphase controller and MAX20790 power stage, along with our patented coupled inductor (CL) technology to help address these AI PoL design challenges. Figure 2 shows the MAX16602, MAX20790, and CL high level block diagram connection for an 8-phase MAX16602CL8_EV design. This relatively clean design achieves a high current delivery capability of ~88 A_{PK} per phase. The internal compensation and advanced control algorithm, along with the integrated current sensing circuits in the power stage and coupled inductor, make it a small solution with best-in-class efficiency.

Figure 2. An 8-phase VR design with ADI's high integration power chipsets facilitate a high density design with less external connections.

Monolithic Smart Power-Stage IC with Higher Level of Integration

The MAX20790 is a feature-rich smart power stage IC designed to work with the MAX16602 (and a few other ADI controllers in the portfolio) to implement a high density multiphase voltage regulator. This is a monolithic integration that nearly eliminates parasitic resistance and inductance between the FETs and driver seen in discrete designs, enabling high switching speeds with significantly lower power losses than traditional implementations. If a switch node (V_x) fault is detected, the power stage immediately shuts down and communicates the fault ID to the controller. This smart power stage IC also includes a current sensor on chip. This current sensing circuit block is far superior to methods that use an inductor's DC resistance. DCR sensing is known to be inaccurate and requires temperature compensation for the current measurement to be trusted.

Controller IC

The MAX16602 is a multiphase controller for xPU V_{CORE} VRs. The IC provides a high density, flexible, and scalable solution to power AI xPUs. The device supports pulse width modulating (PWM) paralleling to control up to 16 phases. The IC's architecture simplifies design, reduces component count, enables advanced power management and telemetry, and increases energy savings over the full load range. Autonomous phase-shedding is implemented to maintain high efficiency across the entire load range. The complete chipset is a highly efficient multiphase buck converter with extensive status and parameter-measurement features. Parameters for protection and shutdown are set and monitored through the serial PMBus® interface, including even faults collected in the power stage ICs.

Here are a couple of other key features supported in the ADI controller that are important for any AI power delivery implementation.

Advanced Modulation Scheme

The MAX16602 includes an advanced modulation scheme (AMS) to provide improved transient response. The modulation scheme allows the turn on and off of phases with minimal delay. Depending on load demand, multiple phases can turn on simultaneously when the load increases or turn off immediately when the load releases. With AMS enabled, the system closed-loop bandwidth can be extended without phase-margin penalty. This allows the PoL a better chance at responding to the type of instant and dynamic current demand experienced with AI VRs.

Load Line Control

The load line allows V_{CORE} to shift between its minimum and maximum based on the output current. It essentially sets V_{CORE} high for light loads and low for heavy loads. The main reason is to allow the control loop to handle higher load current (and this is required to make things work well). The ADI controller provides an accurate output load line control over the entire range of output currents. The output-voltage positioning is performed using the lossless current-sense signals from the power stage IC, which are fed back to the controller. The load line is set in the controller by digitally programming the DC gain of the voltage control-loop-error amplifier. There is a wide range of DC load line profiles presented in the controller's EC table and in Table 6 of the data sheet, from 0.105 mΩ to 0.979 mΩ. Figure 3 shows the transient plot of a 16-phase PoL design for a 40 A to 360 A load step at 800 A/μs slew rate. The result shows minimal overshoot.
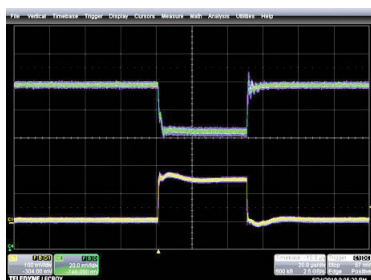
Figure 3. Transient plot of a 16-phase VR for 40 A to 360 A step load at 800 A/μs slew rate.

Adding active voltage positioning to the design eases off the requirements on load transient response and makes better use of the total xPU tolerance window. Load line control helps lower peak-to-peak output voltage deviation for a given step load while making it possible to reduce the amount of bulk capacitance on the output rail. The total voltage fluctuation is reduced, lowering the risk of xPU crashing or getting damaged. Note that the load line circuit block can be disabled in the MAX16602.

Benefits of Coupled Inductor

ADI has been investing its patented CL technology for more than a decade. The technology enables higher density, larger bandwidth, faster transient solutions, and as compared to a discrete implementation 50% higher efficiency and 1.82× smaller magnetics. CL effectively works as large inductance in steady state, small inductance in transient, enabling C_{OUT} savings in addition to smaller inductor size.[2] Figure 4 shows a series of coupled inductors commonly used in ADI's multiphase VR designs.

Figure 4. A series of coupled inductors commonly used with ADI's multiphase VR designs.

Depending on the design specifications and priorities, the benefit of current ripple cancellation of coupled inductors can be traded for either smaller size or higher efficiency.[1] The big system benefit and ADI's differentiator is that AI PoL designers can use CL to help them achieve a small total VR footprint solution relatively easily. Several well-known and popular magnetics vendors have a free CL license from ADI and can provide us multiple sources for the needed parts.

Top-Side Cooled Packages

Top-side cooling provides an alternative heat dissipation route for surface-mount packages. Both the MAX16602 and the MAX20790 are flip chip quad flat no-lead (FCQFN) packages with exposed top-side thermal pads. FCQFN is an advanced packaging that provides best-in-class thermal performance that designers will appreciate. This leadless package not only reduces parasitic inductances but also allows the heat dissipation directly from the device's junction to the ambient environment. The MAX20790 has a junction-to-case top ($\theta_{JC\text{-}TOP}$) thermal resistance of 0.25°C/W. Taking advantage of top-side cooling configuration with AI power designs, system thermal performance and design flexibility can be improved.

Altogether, ADI's multiphase power conversion and PoL products deliver high efficiency and high power density. Figure 5 shows the efficiency plot with bias and inductor losses of our 16-phase MAX16602 + MAX20790 + CLH1110-4 evaluation board. ADI offers voltage regulator and other power conversion solutions for various AI accelerator applications. Our multiphase controller and integrated power stage solutions allow ADI customers to address the most stringent dynamic xPU power requirements and design challenges present with today's AI applications.
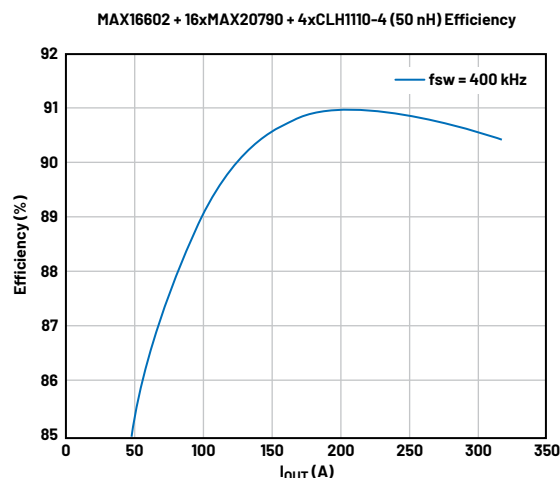

Figure 5. Efficiency plot of a 16-phase AI VR evaluation board design.

Vertical Power

The industry has witnessed a dramatic increase in power consumed with the advent of xPUs' processing complex AI functions. VRs with the capability of up to 650 A continuous current and over 1000 A peak current delivery have become common. The challenges of powering AI processors lie in maintaining efficiency. Conventional power architectures are not going to keep pace with these power hungry AI xPUs. VR chip makers and architects are fundamentally looking at different approaches to power delivery. A new trend to powering AI xPUs that the industry is talking about is called vertical power, also known as backside power delivery.

VRs must be located as close as possible to the load input xPU power pins for high current power deliveries. We cannot get there with traditional lateral power delivery methods. Vertical power delivery moves and relocates the power regulator directly underneath the processor itself, eliminating all the losses that one would have on the PCB. The structure consists of placing the power converter, power stages, capacitors, and magnetics on the back side of the PCB and delivering power vertically through vias to the xPU. In other words, current delivery takes place vertically from underneath the xPU BGA array. This is a reduced length vertical path that significantly reduces impedance and eliminates losses. Figure 6 shows the vertical power module architecture mounted underneath the xPU on the other side of the PCB. This is for illustration purposes only.
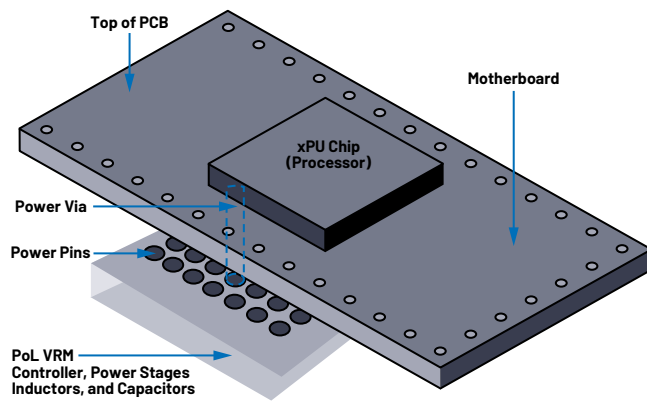
Figure 6. A vertical power module architecture (for illustration purposes only).

The challenging part of building a vertical power solution includes being able to solve the module weight and mounting problems. Underneath the xPU on the other side of the PCB is also the prime location for the high frequency decoupling capacitors that are required for energy storage to meet instantaneous energy demands. Vertical power delivery paired with ADI's CL technology achieves a higher current density, power density, and faster transient performance. Vertical power gives PoL makers like ADI new opportunities to innovate and continue, in their own way, to support the advancement of Moore's law.

Conclusion

With machine and deep learning, accelerator cards have taken AI from theory to mainstream by enabling the parallel processing power required to speed up both training and inferencing workloads. Designing a VR PoL for a high performance AI accelerator card is a complex task, especially with the ever-increasing power requirements in terms of current levels and voltage accuracy required by the current advanced xPUs.

It has been shown in this article that the requirements for the xPU VR are quite different from standard PoL regulators. xPU rails have extremely fast load changes, require dynamic voltage positioning or load line, and must be small.

ADI has a wide portfolio of AI xPU V_{CORE} solutions to solve these problems today. Our power solutions enable best-in-class efficiency in the smallest form factor. The proposed combination of the multiphase controller MAX16602 and smart monolithic power stage MAX20790 offers the highest power conversion efficiency, fastest transient response, and most accurate telemetry reporting in the industry. For more detailed information about these power chipsets or to purchase the MAX16602CL8 evaluation kit, please visit here.

References

[1] "Forecast: AI Semiconductors, Worldwide, 2021-2027." Gartner, April 2023.

[2] "Utilizing the Benefits of Coupled Inductors." Analog Devices, Inc.

About the Author

Hamed M. Sanogo is an end market specialist for cloud and communications in Analog Devices' Global Applications Group. Hamed graduated with an M.S.E.E. degree from the University of Michigan-Dearborn and later earned an M.B.A. degree at the University of Dallas. Following graduation, Hamed worked as a senior design engineer at General Motors and a senior staff electrical engineer and Node B and RRH baseband card designer at Motorola Solutions before joining ADI. Hamed has spent the past 17 years in different roles, including FAE/FAE manager, product line manager, and currently an end market specialist for communications and cloud.

Engage with the ADI technology experts in our online support community. Ask your tough design questions, browse FAQs, or join a conversation.

ADI *EngineerZone*™
SUPPORT COMMUNITY

Visit *ez.analog.com*