



Powering the Next Sophisticated AI System

By Yat Tam, Product Marketing Manager at MPS

Introduction

Artificial intelligence (AI) combines several approaches to problem solving, such as mathematics, computational statistics, machine learning, and predictive analytics. AI systems imitate the human brain's capacity to learn and solve problems. They accomplish this with computer-based "neural" networks consisting of parallel processors that run complex learning and execute software algorithms. Today's AI is revolutionizing computing architecture to replicate neural networks that emulate the human brain. While common models can be trained or developed on a server with traditional central processing units (CPUs), most neural networks require custom, built-in hardware for training.

Graphics processing units (GPUs) and tensor processing units (TPUs) are common accelerators that speed up the neural network training. GPUs and TPUs can handle repetitive and intensive computing, but they are extremely power-hungry. For example, an early AI market dominator, the NVIDIA DGX-1 GPU super computer, contains 8 Tesla P100 GPUs with each GPUs capable of 21.2 TeraFLOPs. This requires 3200W of total system power. The current generation, the DGX-2 super computer, contains 16 Tesla V100 GPUs with each GPUs capable of 2 petaFLOPs, which requires 10kW of total system power. It is no surprise the AI market will grow rapidly to accommodate these rising power demands.

Power Design Challenges

The challenge facing AI power system designers is multi-faceted. Delivering kilowatts of power is the first challenge, and efficiency is absolutely critical. To learn, these computing systems are complex loads that run at full power. As the activity drops, so does the power requirement. The system must remain as efficient as possible throughout the power demand. Every watt of energy wasted dissipates as heat and translates to increased air conditioning requirements in the datacenters. This increases operational costs, as well as the datacenter's carbon footprint.

Real estate is also rising in cost. Modern datacenters contain hundreds or thousands of processing units, and size matters. A size reduction in a single unit, replicated many times over, allows for more devices and a higher concentration of processing power in the same space as larger solutions. However, this smaller size requirement rapidly increases power density and reduces the surface area available for heat dissipation. This makes thermal management one of the significant challenges in designing power for the next generation of sophisticated CPUs, GPUs, and TPUs.

In addition, design resources have been stretched thin by increasing system complexity and shortening design cycles, with resources primarily being allocated to developing the key intellectual property of the system. This often means that the power scheme related circuits are ignored until later in the development cycle. With little time, and perhaps limited power design resources to address the challenges described above, the ideal overall power solution would be space-conscious yet efficient, scalable, flexible, and require minimal design effort.

Digital Control vs. Analog-Based Solution

Analog-based solutions are no longer a viable approach to address the rapidly growing power demand in the AI market. As power systems become more intelligent and integrated into the overall solution, communication between the power solution and the main CPU/GPU/TPU is a design requirement. When designing high-end power solutions for the AI market, a digital control solution is highly beneficial.

The ideal control solution is compatible with multiple products (e.g. Intel, AMD, PMBus), and is be easy to use due to scalable and flexible configuration. Companies including MPS offer these ideal advanced controllers (see Table 1). They provide broad and accurate system control while offering detailed, precise monitoring. The voltage, current, frequency, and faults are configurable over a broad range. These values are accessible in real-time to encourage comprehensive visibility into the solution’s performance. Empowered engineers can optimize the run time through predictive analysis, and minimize downtime by having more data available when repairs become necessary.

Table 1: Advanced Controllers vs. Analog-Based and Digital Control Solutions

	Advanced Controller Solution from MPS	Other Analog Controller Solutions	Other Digital Controller Solutions
Communication Protocol	<ul style="list-style-type: none"> • PMBUS with or without AVS • SVID • SVI2 • I²C • PWM-VID 	<ul style="list-style-type: none"> • SVID • SVI2 • I²C • PWM-VID 	<ul style="list-style-type: none"> • PMBUS with or without AVS • SVID • SVI2 • I²C
Max Phase Configuration	<ul style="list-style-type: none"> • Dual or Single-Rail with flexible phase assignment • 10-phase or less 	<ul style="list-style-type: none"> • Dual or Single-Rail • 7-phase or less 	<ul style="list-style-type: none"> • Dual or Single-Rail with flexible phase assignment • 8-phase or less
Package	QFN 5mmx5mm or smaller	QFN 7mmx7mm or smaller	QFN 7mmx7mm or smaller

Power Stage: Integration is the Key

Obviously, no power solution works without a power stage, and the power stage would conventionally choose a discrete solution. The building blocks for discrete solutions consist of a driver IC and a pair of external MOSFETs, which creates a three-chip solution. Another approach is a multi-chip driver-MOSFET (DrMOS) co-packaged into one IC solution. As addressed earlier, the ever-shrinking system board area makes the three-chip solution less than ideal, as it increases the number of components on a limited board real estate. The co-packaged multi-chip solution is smaller and requires fewer components; however, parasitic inductance inside the package is still high and contributes to efficiency loss, which is not ideal for high-power applications like AI.

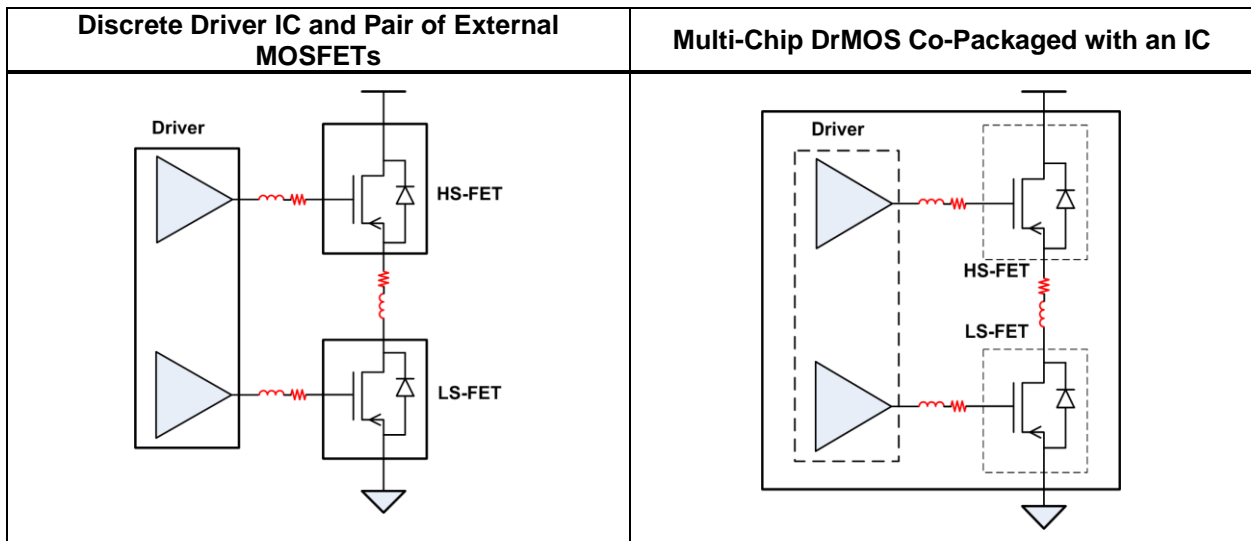


Figure 1: Conventional Method of Approach for Implementing Power Stage

Unlike conventional discrete and multi-chip solutions, MPS implements a monolithic power stage solution. This power stage has a low quiescent current, synchronous buck gate driver, and a pair of high-side and low-side MOSFETs on a single die. With all key elements integrated in one package, the driver/MOSFET is easily controlled and able to minimize ringing at the switch node. In addition, the parasitic inductance between the package and board level is tremendously reduced. This design enables higher efficiency at lower output voltages required by cutting edge CPU/GPU/TPU designs.

The monolithic power stage requires a minimal number of external components, which simplifies the schematic and PCB layout. The base design can be completed in two steps:

1. Choose the appropriate amount of input and output capacitors to satisfy voltage and current ripple requirements.
2. Select an inductor to fulfill the total load current demand.

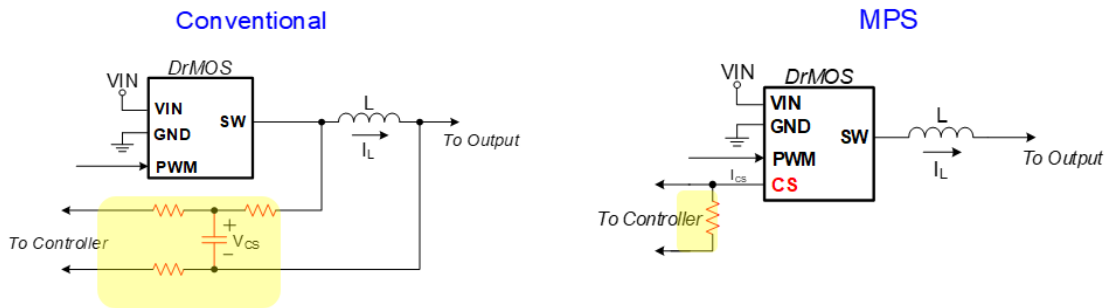


Figure 2: Conventional Solution Compared to Advanced Solution

This device’s typical power stage can achieve excellent current-sense accuracy ($\pm 2\%$) throughout the entire load range in various temperatures, and operates with as high as a 3MHz switching frequency (see Figure 2). Configurable fault protections like over-current protection (OCP), phase fault detection, and IC temperature reporting in a space-conscious system offer designers a small, powerful solution without compromising efficiency or transient response.

Conclusion

Today’s AI systems are enabled through several high-performance computer systems that are challenging power designers on many fronts. The traditional datacenter designs are rapidly migrating from general purpose CPU-only solutions towards combinations of CPUs, GPUs, and TPUs, which bring new and more stringent demands on power design solutions.

Digital controllers and their power stage solutions bring flexibility and adaptability as well as precise control, telemetry, and protection features. This enables power designers to create state-of-the-art power solutions with high efficiency and power density to meet both the current and future high-power needs for the rapidly expanding AI market.