# Powering Datacenters with Digital Multi-Phase Controllers

By Yat Tam, Product Marketing Manager at MPS

## Introduction

The explosive growth of cloud services has driven major advances in datacenters, as well as networking and telecom equipment. Innovative demands impact the servers, storage, and networking switches that process these ever-increasing amounts of data. Infrastructure equipment is being pushed to its limits in terms of processing power and bandwidth. For power designers, the main challenge is how to efficiently power and cool datacenter equipment using minimal electricity. Designers must also balance power footprints with thermal performance when using today's advanced CPUs/ASICs and FPGAs.

## Multi-Phase Solutions Improve Efficiency and Size in Datacenters

With the increase in functionality for total end systems, there is a corresponding increase in processing power to address new datacenter design requirements. This processing capability is centralized in datacenters, where high-end CPUs/ASICs and processors run servers, storage, and networking equipment. Servers and networking equipment are then distributed across the network through telecom equipment, and occur at the point of transaction with point-of-sale machines, desktops, and embedded computing systems using CPUs/ASICs and FPGAs.

In the cases above, the CPUs/ASICs and processors have similar digital processing needs with similar power profiles. Despite the shrinking geometries of today's processors, they also have more transistors, which require higher output currents that can easily range from 100A to 500A or more, depending on their complexity. The industry has adapted by integrating lower power states into the digital loads. This allows devices to be idle at lower currents, then peak to full power as needed. This is beneficial to the overall system power budget, but adds another challenge for power designers at the full power end. First, the datacenter power supply must react to the demand for large load steps exceeding 100A in less than 1 microsecond, while maintaining a narrow output regulation window. Second, the thermal performance must be carefully, reliably managed and remain able to sustain the full power range.

A multi-phase voltage regulator module (VRM) addresses these design challenges. VRMs provide power conversion, typically from a 12V input to a 1V (or lower) output. To provide such large load currents, it is easier to design a multi-phase solution that splits the load across smaller stages (called phases), rather than trying to deliver it via a single phase. Attempting to deliver too much current in one phase presents challenges when designing the magnetics and power stages, as well as thermal issues from a power loss perspective. A multi-phase solution offers high efficiency, smaller size, and lower costs than a single stage supplying the high current. This is analogous to multi-core CPUs, which divide the workload for end loads.

## Digital Control to Power Datacenters

Analog control has been the trusted methodology and power system solution for decades. However, analog control has shortcomings when it comes to high-current and high-power applications. For high-end power solutions, power systems should be more intelligent and integrated into the overall solution, and communication between the power solution and the main CPU/ASIC should be a design requirement.

This is why digital control solutions are highly beneficial to datacenter applications. To illustrate this, let's look at the MP2888A from MPS, a digital multi-phase controller that can replace analog-based legacy controllers. The MP2888A is a 10-phase, digital, multi-phase controller with PMBus and PWM-VID interface (see Table 1).

**Table 1: MPS Digital Controller vs. Traditional Analog Controllers**

| | MP2888A | Traditional Analog Controller |
|---|---|---|
| **Communication Protocol** | PWM-VID and PMBus | PWM-VID |
| **Max Phase Count** | 10 | 8 maximum |
| **Telemetry** | Yes | No |
| **Fault Handling** | Yes, registers setting | Requires extra ICs or external components |
| **Package** | QFN (5mmx5mm) | QFN (5mmx5mm) |

To implement PWM-VID, the MP2888A requires only a single pin on the package due to its digital control technique. Compare to analog controllers, which generally require 4-pin solutions and more than 7 high-precision external components to implement the same functionality (see Figure 1). While high-precision components improve functionality, they also drive up the cost of the entire solution.

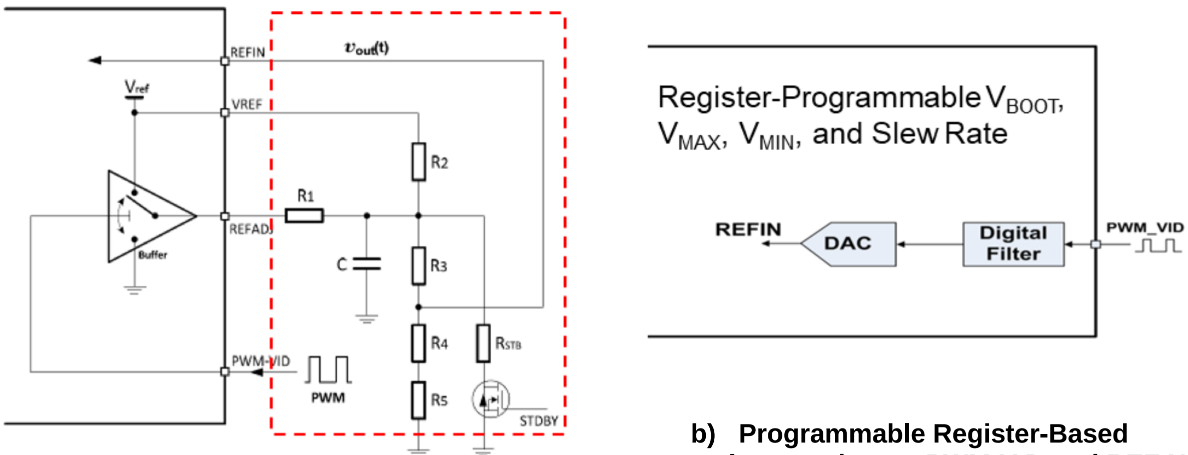


**a) Physical Circuit Implementation on PWM-VID and REFIN for Analog Controller**

**b) Programmable Register-Based Implementation on PWM-VID and REFIN for the MP2888A**

**Figure 1: PWM-VID Implementation (Analog vs. Digital Controller)**

Traditional analog controllers require an RC circuit to form the feedback compensation loop. Optimization of the compensation loop often takes multiple iterations to calculate the correct values of the external components and fulfill various operating conditions. Engineers must also physically change those components before they can test and retest the system.

A digital control solution eliminates repetitive and labor-intensive activities, making the system fine-tuning process hassle-free. Digital control solutions use zero external components due to automatic loop compensation, and can be easily fine-tuned by adjusting the register setting via the PMBus. Calibrating the load-line — which is a useful feature to stabilize the CPU/ASIC during overclocking — is easily accomplished using one resistor to correctly set the tolerance band. Compared to digital controllers, analog controllers require more than 14 external components for loop compensation and load-line calibration (see Figure 2).

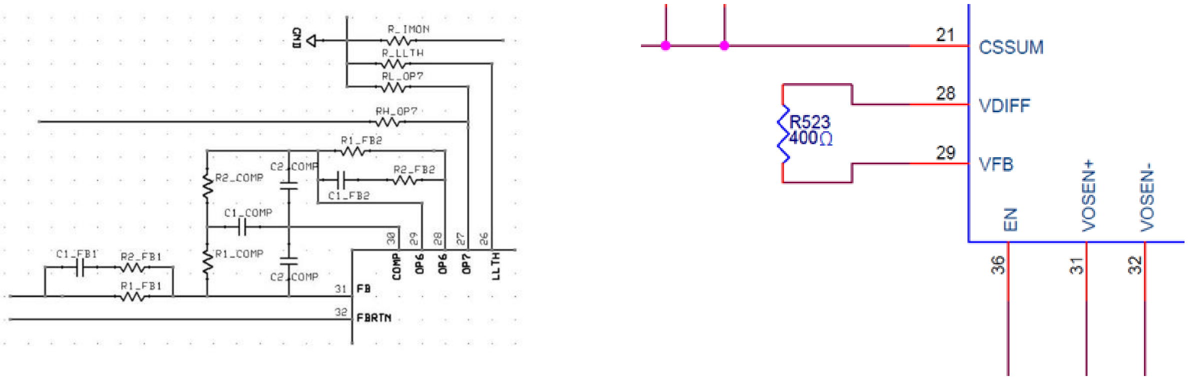


**a) Analog Controller** **b) MP2888A**

**Figure 2: Load-Line and Loop Compensation (Analog vs. Digital Controller)**

Digital controllers simplify system design. For engineers, this eases the complexity of PCB layout design. Figure 3 shows a comparison between an analog and digital controller. Digital controllers can save 36% board area, and require fewer than half of the traditional components.

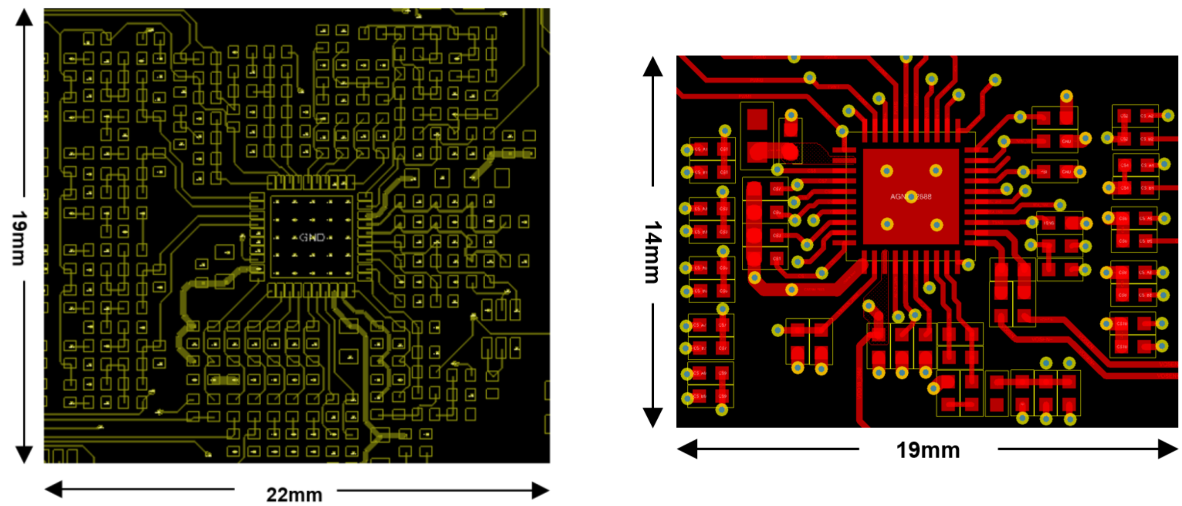


**a) Analog Controller** **b) MP2888A**

**Figure 3: Complexity of PCB Layout (Analog vs. Digital Controller)**

Typical analog controllers use one PWM signal to drive a single power stage. For high-power applications such as datacenters, in which processors demand at least 500A or above for the load current, the typical approach is to:

1. Choose the controller that has the highest phase count (e.g. 20-phase).
2. Use a phase-doubler that doubles the number of phases by generating two interleaved signals that are formed by the original signal.

However, an actual 20-phase analog or digital controller does not yet exist in the market, and a phase-doubler adds components and cost while increasing system design complexity.

By contrast, digital controllers such as the MP2888A can drive two power stages with one common PWM signal (see Figure 4). This type of digital controller eliminates the need for a phase- doubler to ensure current balancing and sharing between the two power stages. In this case, a 20-phase system design can be realized with a 10-phase, digital, multi-phase controller.

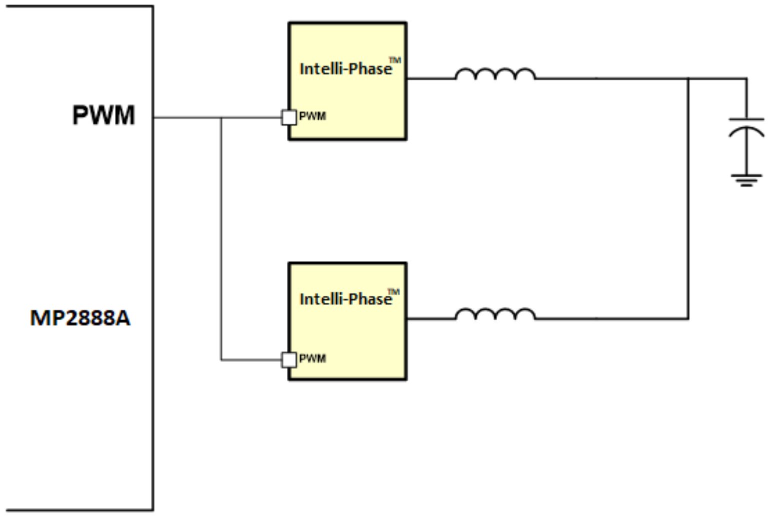


**Figure 4: The MP2888A Drives Two Intelli-Phase™ Devices with One Common PWM Signal**

## Conclusion

The multi-phase solution has evolved into digital control that enables greater capability in solving the challenges of powering high-current/power applications such as datacenters. With digital control solutions, the design burdens of component selection, loop/performance optimization, and layout are greatly reduced. These control solutions shorten the overall design time, system troubleshooting, and ultimately improve time-to-market.