



Demystifying dual winding TLVR inductors

Eaton TLVR inductors



Overview

Data centers rely on powerful CPUs, GPUs, FPGAs, and SoCs in network switches and servers for machine-learning training, validation, test set, inference and acceleration in large-scale data systems (e.g., encryption, compression, virtual switching, filtering, etc). As the technologies for these underlying ASICs and FPGAs have progressed in processing power, so too have their power requirements (where up to a thousand amperes of current can fluctuate rapidly in their power demand). The problem involves an ever-increasing power level requirement with the same, if not less, board space. The traditional solution of the multi-phase voltage regulator (VR) has reached its limit in performance. Meeting the high current output demands requires an increase in the duty cycle of individual phase stages – a process that is excessively long. The newer multiphase trans-inductor voltage regulator (TLVR) was presented by Google in the Applied Power Electronics (APEC) conference. It was released in the technical disclosure commons domain in 2019 to present an alternative circuit topology where the inductors in each of the phases becomes the secondary winding of a transformer and the primary winding is connected in a series loop with an additional compensation inductor (LC). This results in an extremely fast transient response that matches the demands of the load in amperage and bandwidth without sacrificing any other critical parameters (e.g, board space, cost, efficiency, power density, etc.). The TLVR architecture allows end-users to benefit from the advantages of phase coupling, resulting in an extremely fast transient response that scales to the demands of the load in amps and bandwidth.

This white paper discusses the limits of the traditional multi-phase VR module and how the TLVR topology resolves these bottlenecks in performance, while also presenting potential inductor solutions.



Powering Business Worldwide

Limits of the traditional multi-phase VR solution

The multi-phase VR topology

Historically, 5 V power supplies were sufficient for earlier microprocessors (i.e., Intel 8086). However, evolving chip technologies demanded lower voltages and a voltage regulator module (VRM), or DC-DC converter, to supply a more tightly regulated output voltage to the CPU. These single-phase buck converter systems were found insufficient for advancing processing technologies where processing performance/transistor count doubled every two years (Moore's Law). During this time, the power demand of processors put too much thermal strain on the inductors and MOSFETs of single-phase buck converters as the current is divided evenly among all phases. The multi-phase solution came out as a thermally and electrically robust, cost-effective intermediate solution where the high current demands of the ASICs and FPGAs can be met.

The multi-phase VR topology is a buck converter with its multiple parallel sections staggered in phase (Figure 1). The pulse-width modulated (PWM) switching of each phase is symmetrically interleaved with each of the other phases. This is accomplished through an independent controller to adjust the duty cycle of the waveforms based upon the load demand. In an N-phase converter, each stage switches at a 1/N cycle after the previous phase and before the next phase. This way, the switching frequency of the converter is effectively N times that of the individual channels. The output inductor current in each phase is fed to the load. The high output impedance of the topology calls for more output capacitors.

In cases where a load transient inevitably occurs, the duty cycle of each phase must be increased to increase the output current. To meet this demand, multiple phases will then have to increase their output. However, the time it takes to accomplish this can be inordinately slow, causing a major voltage droop in the output voltage of the VR in response to the load transient. The increasingly high output impedance of this topology also calls for the employment of bulky and expensive decoupling capacitors on an already space-constrained board. Several potential solutions have been proposed and implemented to ease the design requirements around multi-phase VRMs.




Figure 1: Conventional multi-phase VR topology


Potential solutions and their shortcomings

To better meet these current and bandwidth load demands with an acceptable bandwidth response, several solutions have been implemented. This typically involves either increasing the number of phases, increasing the switching frequency of the MOSFETs, or reducing the inductance value per phase. Increasing the switching frequency will complicate design and sacrifice efficiency of the converter.

Increasing the phase count will, in turn, lessen the board space for the inductors. Between the output capacitor requirement and controller/switches, the inductor components will grow taller and thinner to fit in the available space (Figures 2 and 3). While this allows design engineers to better meet the stringent amperage and bandwidth demands of modern ASICs, it yields less mechanical stability of the component and increases its profile. A 1 mm distance between pins is typically necessary to mitigate the risk of a solder bridge and short from occurring.

In the multi-phase buck converter topology, the inductor current is ramped up (or down) to match the new load current; a smaller inductance value enables the converter to meet the new amperage demand rapidly and minimizes the output capacitor requirement. The smaller inductance value, however, leads to a large ripple current, which is typically cancelled out by the output capacitor and therefore is minimally transferred to the load. However, the full ripple current still passes through the MOSFETs and the inductor, causing higher losses and high peak current requirements.

The early 2000s brought about transient currents over 100 A with slew rates as high as 500 A/μs – current slew rates go beyond 2 A/ns. If system design parameters are not carefully chosen, the load transients can cause the VR output voltage to go beyond the specified regulation band and manifest as an overshoot and undershoot, limiting the processor's operating speed and/or causing operational malfunctions. Parameters such as specific output voltages, OVP threshold, load current transients of any duty cycle, and switching frequency must be well-managed to prevent this. Various voltage and current control methods are also employed to meet these fast and dynamic requirements of next-generation processors, ASICs, FPGAs and CPUs (e.g., load-line, peak-current mode, V2 mode, current mode, voltage mode, etc). The TLVR topology loosens all of these design restraints caused by the dynamic power demands of modern processors.



Figures 2 & 3: Typical VRM for CPU showing smart power stage (SPS) modules, inductors, and decoupling capacitor matrix

Multi-phase trans-inductor voltage regulator (TLVR)

The TLVR topology

The trans-inductor voltage regulator (TLVR) developed in the technical disclosure is, in fact, very similar to the conventional multi-phase buck converter topology. There are two TLVR schematics presented (Figure 4 seen on pg. 3):

- The single-secondary TLVR
- The dual-secondary TLVR

Replacing the magnetizing inductors (L_m) in the conventional N-phase buck converter with TLVR inductors and a singular compensation inductor (L_c) would allow an engineer to switch from the traditional multi-phase VRM topology to the single-secondary TLVR topology. The TLVR involves multiple phases with an output inductor that is the secondary winding of a transformer. The separate primary winding is tightly coupled to the secondary with negligible leakage, or a coupling coefficient (k_{ps}) close to unity. The primary windings are all connected in a series loop with the compensation inductor (L_c) and grounded. The primary to secondary turns ratio is typically 1:1 (or higher). When a change in duty cycle occurs from a load transient, it is immediately reflected in the secondary winding of each phase. The coupled primary windings note this change instantaneously, allowing all the phases to adjust their output current rapidly. In a single-secondary TLVR topology, each magnetic core can have only one primary winding and one secondary winding for each phase. In the dual-secondary TLVR topology, each magnetic core can have one primary winding and multiple secondary windings for multiple phases. In this case, the inter-secondary coupling coefficient (k_{ss}) is much lower than k_{ps} .

The PWM waveforms drive each phase where, at steady-state, the phases are interleaved equally. Similar to the N-phase conventional buck topology, a magnetizing current is carried by the magnetizing inductor (L_m) of each phase. The current through the compensation inductor (L_c) and the “primary” windings of each phase is a factor of N times larger than the switching frequency. The current through L_c and primary winding is reflected to the secondary winding of each phase, where the output current is the superposition of each phase stage’s magnetizing current (current through secondary winding of each transformer) and the compensation inductor’s current. Unlike the TLVR inductor, the compensation inductor matches the traditional single-turn inductor and exists to optimize the tradeoff between transient performance and system stability. It is a parameter that can be optimized to better adjust to the rise/fall of the converter’s output current to match the demand of the controller.




Figure 4: N-phase conventional buck, N-phase single-secondary TLVR, and N-phase dual-secondary TLVR topologies (left to right)

Advantages of TLVR

TLVR eliminates the massive bottleneck of a slow transient response to the load in amperage and bandwidth while maintaining the thermal benefits of the multi-phase VRM. It relaxes the multi-phase VR issues of ripple current and an increasing switching frequency to optimize transient response. This allows designers to bring the switching frequency of the MOSFETs down to the 500 kHz to 600 kHz range and loosen the design requirements around the VRMs which, in turn, increases the efficiency of the converter. Coupled inductors have been used in multi-phase topologies for more significant ripple cancellation; however, engineers have often resorted to discrete inductor and integrated non-coupled inductor solutions to avoid infringing on existing patents for integrated coupled inductors.

The TLVR is an open solution that allows designers to shift away from the non-coupled inductor without encountering any troubles down the line. Coupled inductors also offer an intrinsically small output inductance, allowing the TLVR topology to exhibit an excellent transient response at high phase counts.

Performance, cost, and board space are the base concerns for implementing the TLVR solution. While the TLVR inductors themselves may be more expensive than single-turn inductive solutions, the decrease in bulky and expensive decoupling capacitors allow for cost savings. These factors enable the TLVR topology to be a cost-competitive solution with more optimal performance. The risk of implementing a single-secondary TLVR structure is also greatly reduced because of its similarity to the traditional N-phase conventional buck topology. Both of these topologies exhibit an identical PCB layout outside of the differing inductors. As long as inductors are footprint compatible, or, if the two different inductors can be co-layered, the engineer can switch between a standard inductor and TLVR inductor without requiring two different PCBs. This eliminates any risk with time/cost investment as engineers begin building prototypes.

A look at TLVR power inductors

Eaton offers a series of power inductors that are designed specifically to aid in the shift from the common N-phase buck converter to the TLVR circuit. The TL and TLP series of TLVR inductors feature low DCR, dual tight coupling windings and utilize a low loss ferrite core. The single-turn SL series of inductors is footprint compatible with the TL series, allowing engineers to switch between the standard multi-phase VRM and the TLVR topology. The LC series of compensation inductors can also be employed to implement the TLVR topology with a 6 MHz operating frequency and small footprint.

Construction of the TLVR power inductor

The TL series of TLVR coupled inductors can feature either a side-by-side or inside-out winding configuration (Figure 5). In the side-by-side winding construction, the secondary winding is separated from the primary winding with a layer of insulation, while the inside-out winding configuration has the primary winding fitted below the secondary (tin-plated silver copper winding). In terms of ease-of-integration, the side-by-side winding configuration offers better coplanarity control and therefore balances on the PCB better for proper soldering to the board. This prevents solder failures where the component is subject to falling over. The side-by-side winding structure is easier to inspect and debug on all terminals after PCB assembly. The isolation between the primary and secondary windings is also higher at 100 VDC to 200 VDC. However, there is an inconvenient co-layer with the single-turn inductor solution, which is where the inside-out configuration excels.




Figure 5: Side-by-side and inside-out winding configurations for TLVR inductors

Things to consider when implementing the TLVR power inductor

There are several factors to consider when choosing a TLVR inductor. Some are as follows:

- Convenient co-lay
- Coplanarity
- Dimensions of the inductors
- Ease of simulation

Engineers looking for a convenient co-lay between the traditional single-turn inductor and the TLVR inductors can turn to the SL (one winding) and TL (dual tight coupling windings) series of inductors (Figure 6). Both of these models are footprint compatible, allowing engineers to easily switch between the N-phase traditional buck and N-phase single-secondary circuits. A poor lead coplanarity can lead to weak and open solder joints during the solder reflow process, yielding poor manufacturability. The standard SL inductor, as well as the TLVR inductors (TL and TLP), offer good coplanarity control for ease of assembly and solderability.

TLVR inductors are dimensionally similar to their single-turn inductor counterparts, offering 5 to 6.4 mm widths for high power density applications with a high phase count. The profiles of the inductors are also relevant to ensure more mechanical stability of the component. For ease of simulation, Eaton offers the coupling coefficient (kps) and an equivalent circuit model of their TLVR inductors. All of these factors combined lower barriers for engineers to begin deploying the TLVR solution into their VRM design to meet the increasing dynamic load requirements of next-generation ASICs.



Figure 6: Convenient co-lay between TLVR inductor solutions (TL series) and single-turn standard inductor solution (SL series)

Conclusion

The trans-inductor voltage regulator (TLVR) allows engineers to improve on the transient response of their VRMs and meet the demanding load requirements of CPU, FPGAs, and ASICs in current and bandwidth without hurting other critical parameters. For engineers looking to rapidly implement a TLVR prototype, the single-secondary TLVR topology involves minimal risk given that the footprint of the TLVR inductors match with the standard single-turn inductor. This new topology enables engineers to meet changing system requirements without sacrificing on cost, boardspace, and manufacturability.

Eaton offers an extensive range of product offerings geared towards the TLVR solution.

Eaton has collaborated with several IC houses and large-scale OEM/ODM customers to assemble a solid initial lineup of nine TLVR families with matching footprint compatible standard single turn inductors. See the chart below for the TLVR family that best suits your requirements:

Family	Footprint	Inductance	Saturation Current	Standard single turn footprint inductors
TL1011V2-RXXX-R	9.6 x 6.4 x 11.0 mm	70 nh to 17 0nH	163 A to 67 A	SL1011V2-RXXX-R
TL1012V2-RXXX-R	10.0 x 6.0 x 12.0 mm	70 nh to 170 nH	157 A to 64 A	SL1012V2-RXXX-R
TL1012V3-RXXX-R	10.0 x 5.0 x 12.0 mm	70 nh to 170 nH	121 A to 50 A	SL1012V3-RXXX-R
TL1111V3-RXXX-R	11.0 x 4.8 x 11.0 mm	70 nh to 180 nH	150 A to 58 A	SL1111V3-RXXX-R
TL1211V1-RXXX-R	12.0 x 6.0 x 11.1 mm	70 nh to 200 nH	180 A to 62 A	SL1211V1-RXXX-R
TL1211V2-RXXX-R	12.0 x 6.0 x 11.0 mm	70 nh to 200 nH	170 A to 59 A	SL1211V2-RXXX-R
TL1212V2-RXXX-R	12.0 x 6.0 x 12.1 mm	70 nh to 170 nH	186 A to 76 A	SL1212V2-RXXX-R
TL1212V3-RXXX-R	12.0 x 5.0 x 12.0 mm	70 nh to 170 nH	145 A to 60 A	SL1212V3-RXXX-R
TLP1013V1-RXXX-R	9.6 x 6.4 x 13.15 mm	105 nh to 170 nH	108 A to 66 A	SLP1013V1-RXXX-R

If none of the above families meets the requirements of your application, please inquire about Eaton's custom TLVR solutions.

Eaton
Electronics Division
1000 Eaton Boulevard
Cleveland, OH 44122
United States
Eaton.com/electronics

© 2022 Eaton
All Rights Reserved
Printed in USA
Publication No. ELX1160 BU-ELX2201
March 2022