ISSUE: May 2020

Novel 54-V To 12-V Buck Topology Eases Efficient Power Delivery In Data Centers

by Paolo Sandri, STMicroelectronics, Santa Clara, Calif.

Power consumption from hyperscale data centers poses numerous challenges to the design community to deliver power more efficiently to meet the continuously growing demand. The power needs of the current and upcoming generations of digital CPU, DDR memory, and networking ASICs is skyrocketing, urging the adoption of new ways to deliver power more efficiently. This situation is exacerbated by the increasing deployment of high-power GPUs (graphics processing units) and TPUs (tensor processing units) for artificial intelligence and machine-learning functions.

An important parameter for data centers is the PUE (power utilization efficiency). This indicates the ratio of the total power consumed by the data center divided by the useful (server) power.[1] The ideal PUE is 1, meaning that all data center power is used for IT power (useful power). PUE values range from 1.1 for the most efficient data centers to 2 or greater for the less efficient ones. These are huge opportunities for efficiency improvements.

One key means to address this challenge is to move the rack bus bar power distribution from the traditional 12 V to 48 V (nominal 54 V) as this offers significantly lower distribution losses due to the lower current at the 48 Vdc bus.[2] Changing the dc input bus requires developing new topologies to efficiently manage the conversion from the 54-V/48-V bus and "smoothly merge" with the current 12-V ecosystem in the least invasive way while delivering the best overall power efficiency gain.

Traditionally, a first stage local conversion on the server motherboard steps down the 54 V/48 V from the rack bus bar to 12 V. This "local 12-V" bus feeds the traditional on-board multiphase VRs dedicated to the digital CPUs, GPUs, and ASICs, keeping the benefits of the well optimized 12-V to point-of-load (POL) multiphase architectures. An important aspect of this new 48-V to 12-V power delivery, besides its high efficiency, is scalability. Scalability contributes to an optimized power delivery system based on the ASICs' specific electrical performance, overall component count, and cost.

The transition to a 54-V/48-V bus has nurtured innovative ways to deliver power. Various topologies are available to convert from this bus down to the intermediate 12-V bus. However, the first power system architectures developed were "unregulated," and used a converter topology which reduced the input voltage by a fixed ratio (4:1 was common) in an open-loop fashion. These topologies were generally based on switched-capacitor converter concepts and, during APEC 2017, Google presented a first-stage solution, switched-tank converter (STC), that converted from 48 V to a 12-V unregulated bus, to distribute power across the system board to the second stages' 12 V to POL.[3]

Another option, which is seeing increasing adoption in industry, is to generate a "regulated 12-V" output. The regulated architecture is even less disruptive to the current 12-V onboard power delivery, since it provides the regulated 12 V that is commonly used in a "native 12-V server system." If the conversion from 48 V to 12 V is done in a regulated fashion, it can be seamlessly applied upstream of the existing 12-V power distribution bus without the concern that the wide input voltage variation seen on the 54-V/48-V bus will be passed along to the 12-V bus.

The unregulated voltage on the 54-V/48-V bus can range from a minimum of 40 V to a maximum of 60 V. With a 4:1 unregulated voltage stepdown, that 54-V bus variation translates to an intermediate voltage ranging from 10 V to 15 V. This voltage range might generate some concerns if the downstream load needs a regulated input bus, like for example, the PCIE cards or the storage disks.

ST has been very active in developing innovative topologies to serve the 54 V/48 V to 12-V unregulated and regulated bus introducing several benefits in terms of efficiency optimization, scalability and off-the-shelf component selection. This article presents the details of the newly developed ST Stacked Buck topology (also known as STB), which delivers a high-efficiency, high-density, regulated 12-V bus. The material presented here

will highlight the principles of operation and the fundamental design guidelines for designing and customizing the STB for an application's specific power needs and optimized performance.

The STB Topology

The STB topology derives from the stepdown buck converter approach and minimizes the MOSFETs' losses. In the traditional multiphase approach, two or more buck converters are paralleled at both the input and at the output sides and the equivalent phases are distributed to sustain high output currents and reduce the drain-to-source voltage on the switching elements.

The STB optimizes this strategy and results in a multiphase buck converter topology in which one cell contains a two-phase buck where the phases are "serialized" at the input side and paralleled at the output side. In particular:

- Serialization at the input reduces the drain-to-source maximum voltage swing of the MOSFET to VIN/2.
- Paralleling at the output shares the total output current between two phases (and relative inductors) to reduce conduction losses.

This new topology concept is depicted in Fig. 1 for a single cell. More cells can be interleaved to achieve a higher power capability with benefits in scalability, modularity, and energy efficiency that a multiphase-like solution inherently provides.

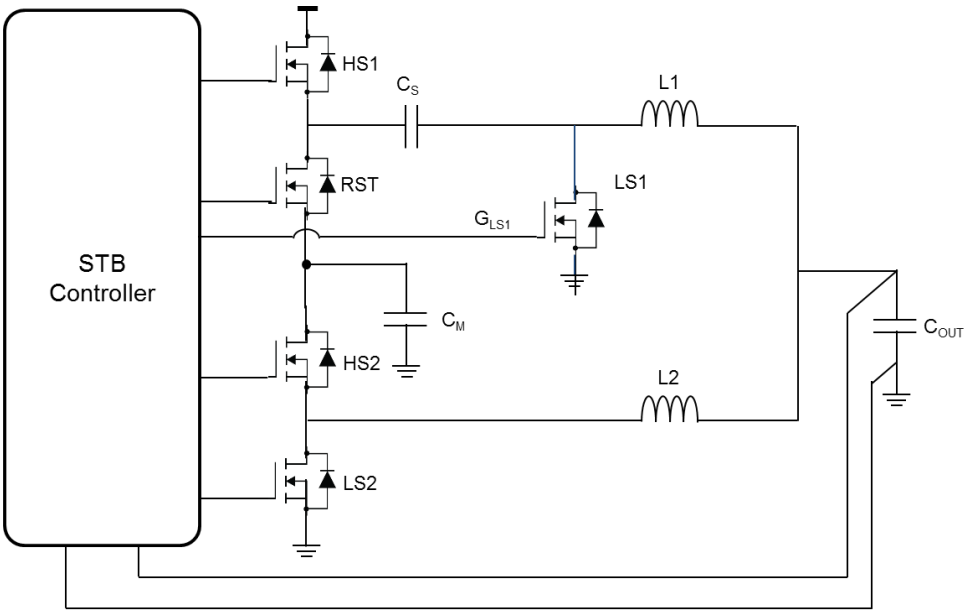

Fig. 1. The STB cell concept.

The STB Concept

The STB topology relies on a multi-cell array where every cell implements a stacked buck approach with a two-phase-like stepdown converter where the phases are serialized at the input and paralleled at the output. Fig. 1 conceptually shows the structure of a one-cell topology, composed as follows.

The upper (floating) phase is implemented by the HS1, LS1 and RST MOSFETs, the L1 output inductor and the C_S stack capacitor. The lower (grounded) phase is implemented by the HS2 and LS2 MOSFETs and the L2 output inductor.

The phases are in series at the input. Thanks to C_M (C middle), the total input voltage is equally divided between the phases, allowing the use of lower voltage-class MOSFETs. In fact, all the MOSFETs have a drain-to-source maximum voltage of $V_{IN}/2$. The phases are in parallel at the output and the total load current can be shared among the phases with the known benefits of conduction-loss reduction.

The STB controller drives the cell, directly or through external high-voltage drivers. The resulting number of cells that can be implemented is equal to N/2 the phases the buck controller can control.

The control scheme drives the phases of the cell with constant on-time T_{ON} control ensuring proper interleaving. A switching frequency control loop (FLL) can modulate the T_{ON} slightly in order to maintain a constant switching frequency across the load.

The system behaves then as a two-phase synchronous buck, controlled in constant-on-time with variable frequency, and the control loop theory of the synchronous buck applies. If necessary, the duty cycle can extend above 50% with HS1 and HS2 turned on simultaneously to sustain the load requirements.

Fig. 2 shows the basic driving signals applied to the cell in cases where the duty cycle is less than or greater than 50%. Generally, an energizing period on the upper phase is followed by a reset period, required to restore the charge into C_S. Then the lower phase can be energized and subsequently the C_M charge is restored.

To fully analyze the system behavior, four stages of operation can be considered as follows:

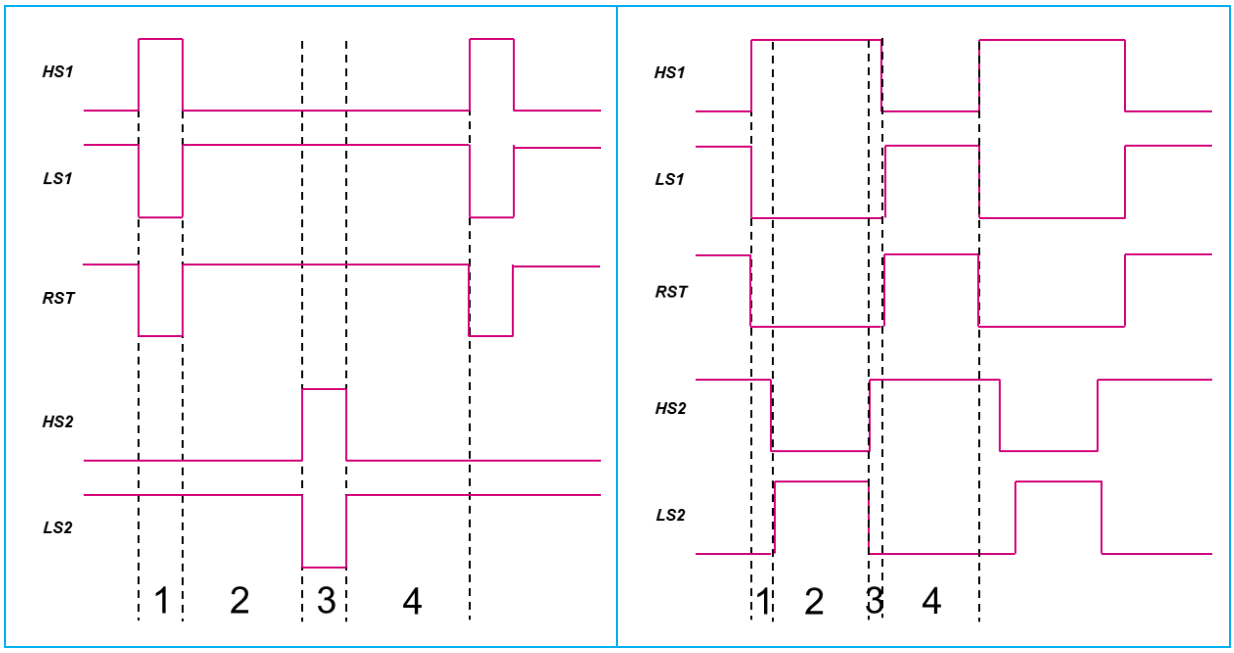

Fig. 2. Basic drive signals for the STB at duty cycles < 50% (a) and > 50% (b).

Circuit Operation For Duty Cycle < 50%

Energizing The Upper Phase—Stage 1

In this stage, HS1 and LS2 are kept on while all the other switches are off (see Fig. 3a). In the upper phase, assuming C_S is initially charged at the mid-bus voltage V_M, the required current flows to the load through the stack capacitor C_S. If C_S is big enough, its voltage drop is negligible. The drain-to-source voltage of LS1 can then be estimated as ($V_{IN} - V_M$).

The upper phase can then be represented as a common buck converter which has an input voltage of $V_{IN} - V_M$ and the output inductor is L1.

The lower phase recirculates the output current from ground through LS2 and L2.

The voltage of the capacitor C_S has varied by $\Delta V_{S1} = (T_{ON1} * I_{L1})/C_S$ where T_{ON1} is the duration of stage 1, i.e. the on-time generated for phase 1. Although the value of C_S is big, it is now necessary to restore its voltage to V_M in order to be ready for the next switching cycle. This is done in the next stage.

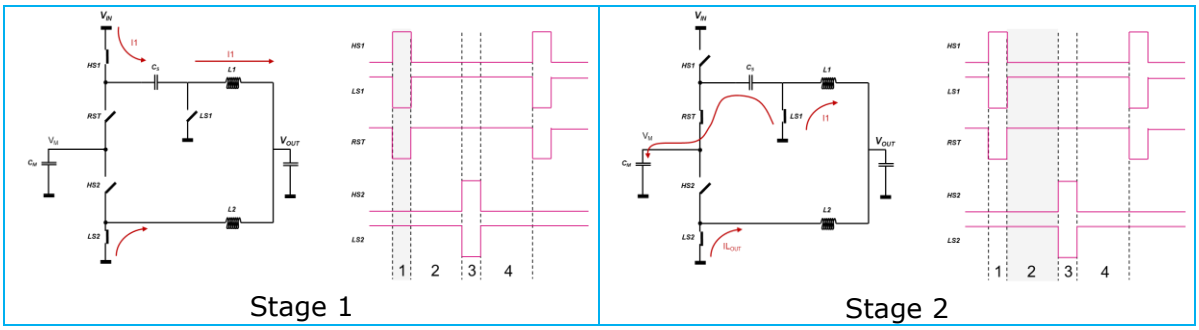

(a) (b)

Fig. 3. Switch states and switching waveforms during stage 1 (a) and stage 2 (b) with duty cycle < 50%.

Energy Restore (Reset)—Stage 2

In this stage, HS1 is switched off, while LS1 and RST are switched on (see Fig. 3b). In the upper phase, C_S is connected between ground and the mid-bus V_M: C_S charge is then shared between C_S (whose voltage has increased from the initial value of V_M) and C_M (still charged at V_M).

This phase is partially restoring into C_S the charge lost during the previous cycle; the two caps C_M and C_S are connected in parallel and share the charge. Therefore, the mid-bus cap V_M will increase its voltage by $\Delta V_{M1} = (T_{ON1} * I_{L1})/(C_M + C_S)$. The lower phase still recirculates the output current from ground through LS2 and L2.

Energizing Of Lower Phase—Stage 3

In this stage, HS1 remains off and LS1 and RST switches remain on. The lower phase changes state and HS2 is switched on while LS2 is switched off (See Fig. 4a). While the upper phase keeps restoring the charge into C_M, this same energy is used by the lower phase to provide energy to the load through inductor L2. Additionally, the upper phase also recirculates the output current from ground through LS1 and L1.

It should be noted that this is the same equivalent buck converter described in stage 1 when assuming symmetry conditions where $V_M = V_{IN}/2$ and L1 = L2.

The lower phase, providing energy to the output, is now extracting charge from C_M and C_S which results in voltage being reduced by $\Delta V_{M2} = (T_{ON2} * I_{L2})/(C_S + C_M)$ where T_{ON2} is the duration of stage 3 and corresponds to the on-time of the lower phase.

When driving the same T_{ON} to both phases, the total amount of charge given to and extracted from the C_S and C_M capacitors sums to zero. The upcoming stage 4 will be used to further stabilize the voltage of C_S and C_M.

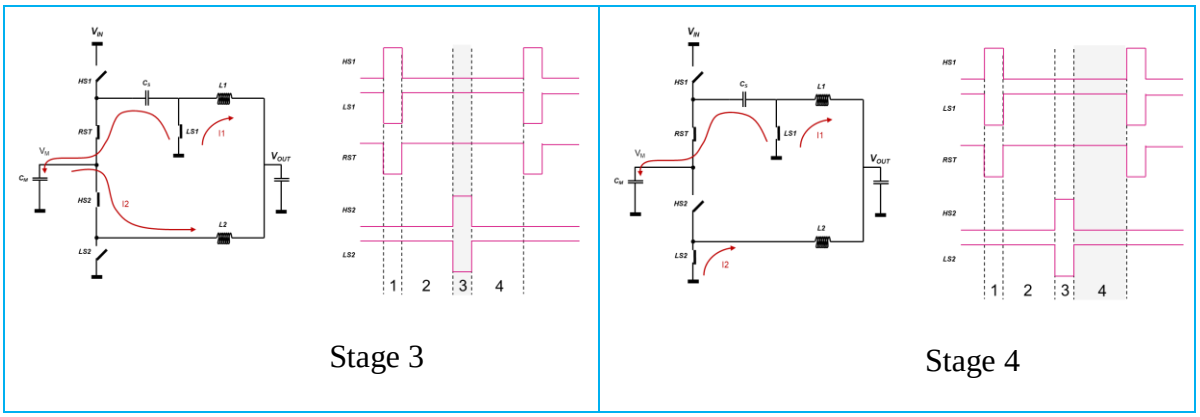

Fig. 4. Switch states and switching waveforms during stage 3 (a) and stage 4 (b) with duty cycle < 50%.

Energy Restore (Reset)—Stage 4

This stage is configuring the same switches in the same positions as stage 2 (see Fig. 4b) and confirms the reset of the voltage across C_S at the same value of the voltage across C_M. The cell is now ready for a new cycle.

Circuit Operation For Duty Cycle > 50%

In this case, the overlap between HS1 and HS2 being on simultaneously produces a slightly different sequence as depicted in Fig. 5.

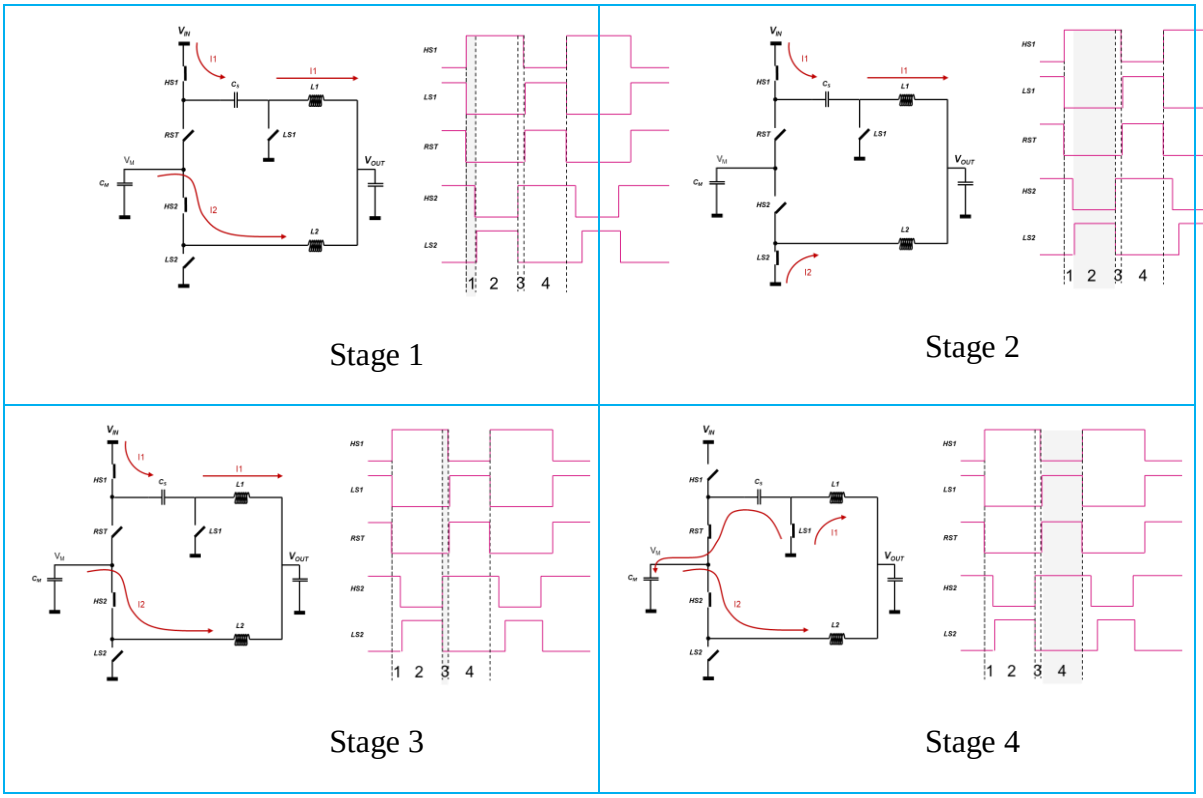

Fig. 5. Stage summary with duty cycle > 50%.

Energizing The Upper And Lower Phase— Stage 1

In this stage, the configuration for the upper phase is the same as in the previous stage 1. On the other hand, the lower phase is extracting charge from C_M in order to provide energy to the load (see Fig. 5 again).

This stage lasts only while the two phases overlap.

Energizing The Upper Phase—Stage 2

In this stage, the configuration for the upper phase is the same as in the previous stage 1. Meanwhile, the lower phase recirculates the output current from ground through LS2 and L2.

Energizing The Upper And Lower Phase—Stage 3

This stage is configuring the same switches in the same position as stage 1. Both the upper and lower phase are providing energy to the output.

Energy Restore (Reset)—Stage 4

In this stage, the configuration of the upper and lower phase is the same as in the previous stage 3. While the lower phase is still providing energy to the output, the RST switch is restoring energy into C_M and equalizing the voltages between C_S and C_M.

Current Sharing And Control Loop Requirement

Regardless of the duty cycle, it is clearly observable that the energy stored in C_S during the energizing stage is used by the lower phase to deliver energy at the output when commanded. If considering the same $T_{ON1} = T_{ON2}$ there is an automatic balancing of the currents between upper and lower phases.

$\Delta V_S = \Delta V_M \rightarrow T_{ON1} * I_{L1} = T_{ON2} * I_{L2}$

Since $T_{ON1} = T_{ON2}$ results in the currents being equalized, there is no need for any current sharing circuit between the phases of the same cell.

This condition leads to the requirement for a constant on-time (CoT) control topology in order to avoid current sharing issues. As is always the case with CoT control, T_{ON} is always set to the same value, and the switching frequency varies according to the load and natively stretches during light-load conditions.

A frequency control loop can be additionally implemented, slightly correcting the T_{ON} duration in order to maintain the switching frequency at the programmed value. Nevertheless, the switching frequency temporarily boosts during load application to deliver the current requested by the load and lowers during load removal.

Capacitor Selection

The proper choice of the stacked (C_S), input (C_{IN}) and mid-bus (C_M) capacitors is essential to obtaining the best efficiency and performance possible from this topology. The three main considerations driving the selection of C_S, C_{IN} and C_M are explained below. The system designer must take these issues into account.

Stacked Capacitor Equivalent Impedance And Bias Effect

Considering the voltage drop created across C_S during the energizing stage (see stage 1) the resulting equivalent phase-node voltage of the upper buck is not a perfect square wave. Rather, it assumes a trapezoidal shape because of the C_S discharge (an effect of the inductor L1 ripple, which is neglected here for simplicity). But because of this effect, the equivalent square-wave of the phase voltage is a square-wave with lower voltage. Instead of being V_M by itself, it is in practice:

$$V_M - \frac{\Delta VC_{S1}}{2} = V_M - \frac{T_{ON1}}{2 \cdot C_S} \cdot I_{L1}$$

An equivalent series resistance

$$R_{eq} = \frac{T_{ON1}}{2 \cdot C_S}$$

can be considered in series with HS1 and therefore affects the equivalent phase node excursion, and in turn, also the efficiency.

It is therefore clear that a bigger value of C_S leads to a lower R_{eq}, which then positively affects the system efficiency.

Based on the above consideration, it is also necessary to pay attention to the "bias effect" of the MLCC in the equations above. In other words, it is good to use the real value of C_S at the biasing voltage (typically $V_{CS} = V_M = V_{IN}/2$) for proper estimations of R_{eq}.

Stacked And Mid-Bus Capacitor Recharge Time-Constant

After the energizing stage, the C_S voltage is recharged to V_M during the subsequent stages of the switching cycle. Considering the equivalent circuit of stage 2, it is possible to identify a time-constant driving this recharge process (see Fig. 3):

$$\tau_{CS} = (C_S + C_M) \cdot (R_{RST} + R_{LS1})$$

For an optimum recharge of the capacitors it is recommended that C_S be chosen such that the time constant τ_{CS} is much shorter than the remaining time to complete the switching cycle. A good rule is to have τ_{CS} lower than one-fifth of (T_{sw} - T_{on1}). This will avoid additional ripple being created by the next cycle being initiated before the C_S/C_M recharge process is completed:

$$\tau_{CS} \leq \frac{T_{SW} - T_{ON1}}{5}$$

Mid-Bus Capacitor As Input Capacitor Of The Buck Phase

The mid-bus capacitor C_M is, in practice, the input capacitor for the buck phase. In the case of duty cycles lower than 50%, it is also supported by C_S. But in the worst case (transient response or high duty cycle conversions) C_M is alone in supplying the energy for the buck phase conduction zone. The maximum drop of the C_M capacitor can then be quantified as:

$$\Delta V_{2\_max} = \frac{T_{ON2} \cdot I_{OUT2\_max}}{C_M}$$

Selecting the right C_M to minimize this ripple and to satisfy the time-constant condition above allows us to maximize the efficiency of the overall system.

Expanding The Basic Concept

The basic structure proposed in Fig. 1 can be extended in two main directions. The first is further partitioning the input voltage (see Fig. 7a), i.e. adding more stacked phases in series. The result is a converter in which all the inner phases are effectively supplied by V_{IN}/N (where N is the total number of phases) further benefitting from the lower voltage class of the MOSFETs.

Additionally, more phases in parallel can carry higher currents. As per the basic cell structure described already, no current sharing control is required to balance the current among the phases of the same cell. This is accomplished by the symmetry of the driving signals T_{ONx}.

Another direction for expansion is adding (packing) more cells in parallel (see Fig. 7b), which allows an increase in the overall current capability. In this case, a multi-cell topology can be considered in which every basic cell is composed of the basic structure defined in Fig. 1 or by an extended-stacked structure mentioned above. A resulting matrix of phases can be considered.

In this case a current-sharing control circuit is required to balance the current into the two or more cells (but the current between stacked and buck phases of the same cell is still guaranteed by the symmetry of the T_{ON} driven by the CoT control loop). In the Fig. 7b diagram, APWM0 and BPWM0 could eventually be the same electrical signal. But for current-sharing purposes, they need to be different and driven differently because of the current balance control.

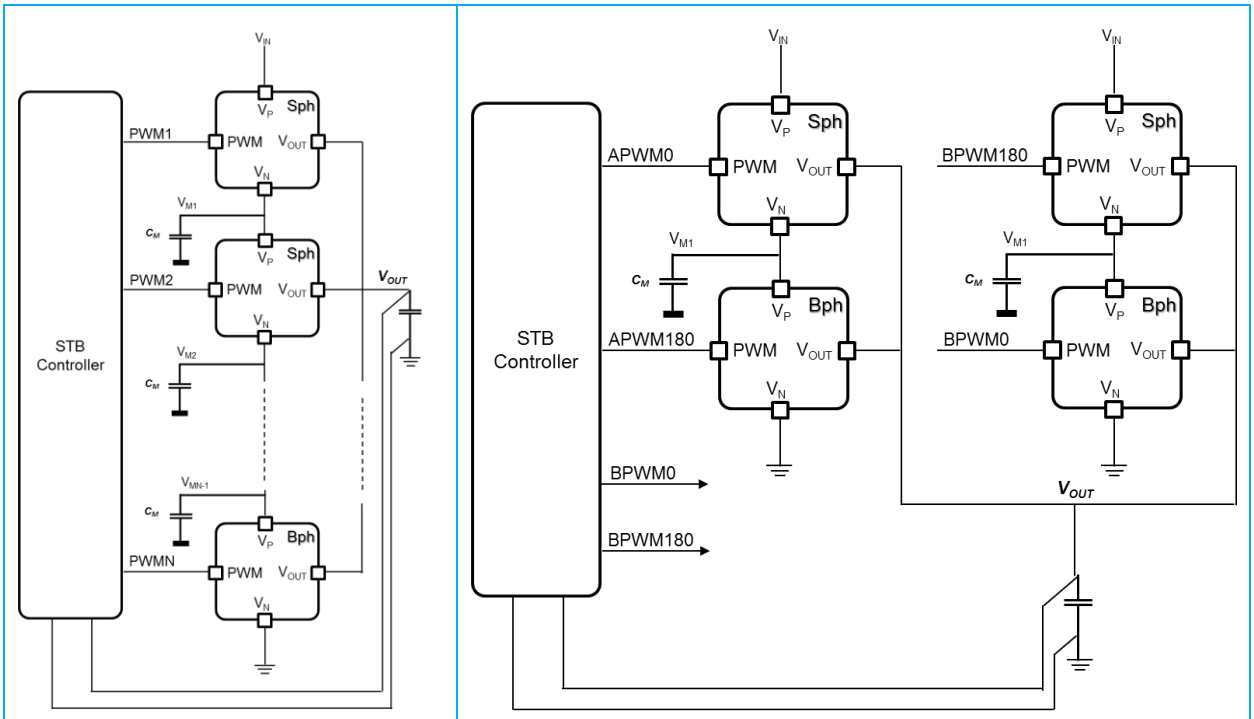

Fig. 7. Expanding the basic concept: array (a) and matrix (multi-cell, (b)).

Reference Design And Experimental Results

Various reference designs have been implemented with a focus on maximizing efficiency and power density as well as scalability. Furthermore, these reference designs have been developed to allow designers to effectively and easily customize the STB to the specific power needs of an application. Fig. 8 presents a picture of the 800-W TDP (Thermal Design Power) reference design, and its efficiency curves are shown in Fig. 9.

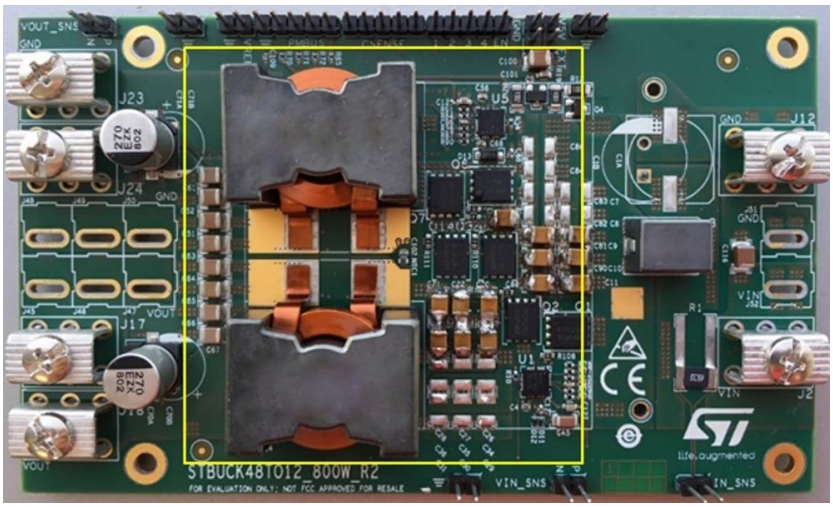

Fig. 8. 800-W reference design board

The 800-W reference board delivers an efficiency curve with a peak of 98.5% at 350 W and an efficiency always above 98% from 150 W up to 650 W with a very flat profile.

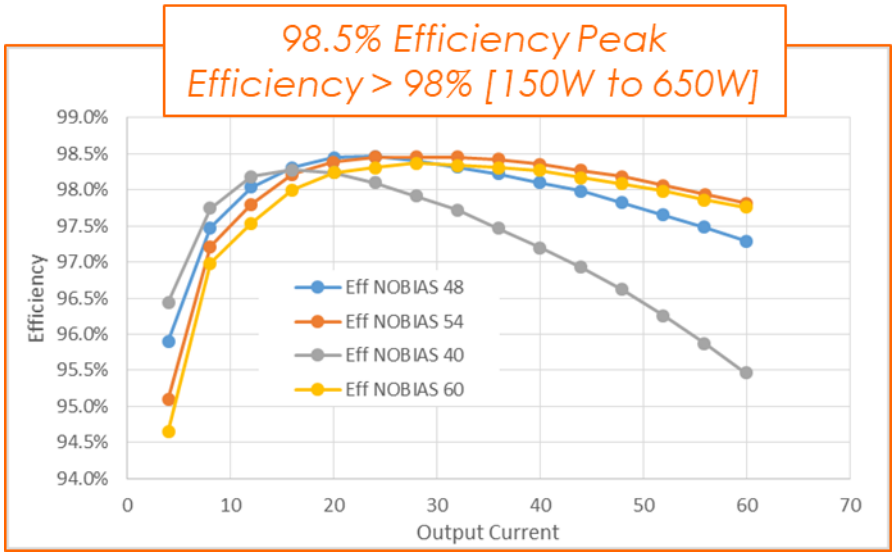

Fig. 9. 800-W STB efficiency results.

One interesting reference design shown in Fig. 10 leverages an interleaved design with more cells in parallel at the output side. In this reference design the basic cell is an 800-W STB, with an eighth-brick equivalent form factor, that can be interleaved in up to four cells capable of delivering up to 3.2 kW at the output. The advantage of this design is the scalability and the high density offered by the eighth-brick form factor.

Fig. 10. Four 800-W cells interleaved in a 3.2-kW STB.

Conclusions

ST's Stacked Buck is a newly developed topology that offers efficiency maximization, high power density and wide scalability to design customized solutions for stepping down the 54-V/48-V bus to a 12-V regulated intermediated bus. Effective scalability and easy-of-use configuration of the multi-cell topology allows an almost ideal customization of the design, enabling it to deliver from 500 W up to 3.2 kW of output power.

References

1. "The Datacenter as a Computer. An Introduction to the Design of Warehouse-Scale Machines Second Edition" Luiz André Barroso, Jimmy Clidaras, Urs Hölzle. Morgan & Claypool Publishers. 2013

2. "Google 48V Power Architecture" by Xin Li and Shuai Jiang, APEC 2017 plenary presentation.

3. "Switched tank converters" by Shuai Jiang, Chenhao Nan, Xin Li, Chee Won Chung, Mobashar Yazdani, 2017 IEEE Applied Power Electronics Conference and Exposition (APEC).

About The Author

Paolo Sandri is responsible for specialized Power Management Products at STMicroelectronics. Paolo joined ST as an analog IC design engineer, where he developed several power conversion ICs. He took on market development responsibility for power ICs and is focused on power management for data centers, where he contributes to improving data center efficiency and power delivery through innovative solutions. Paolo holds a BSEE and MSEE from University of Padua (Italy) with specialization in power conversion.

For more information on dc-dc converter design, see the How2Power Design Guide, locate the Power Supply Function category and select "DC-DC Converter". Also locate the Popular Topics category and select "Buck Converters".