BY PETER JAMES MILLER • TEXAS INSTRUMENTS

Powering DDR memory and SSTL

ALTHOUGH DDR MEMORY IS POPULAR FOR MEETING TODAY'S DEMAND FOR LARGE AMOUNTS OF HIGH-SPEED MEMORY IN SMALL FORM FACTORS, PROVIDING POWER TO DDR MEMORY CAN POSE SOME DIFFICULTIES.

Most modern electronics require some form of DRAM. By far, the most common DRAM is SSTL (stub-series-terminated-logic)-driven DDR (double-data-rate) memory. Conventional logic and I/Os use standard system bus voltages; however, DDR-memory devices need the precision that only local POL (point-of-load) regulators can provide. For sufficient noise margin, two of the five system supply voltages must reference other voltages. These voltages are the I/O V_{DD} (drain-to-drain voltage), the V_{DDQ} (drain-to-drain core voltage), the V_{DDL} (drain-to-drain logic voltage), the V_{TTREF} (termination-tracking reference voltage), and the high-current-capable midrail V_{TT} (termination-tracking voltage).

V_{DDQ}: THE SIMPLEST SUPPLY RAIL

Most DDR-memory devices use a common supply for core, I/O, and logic voltages; these terms are commonly combined and referred to simply as V_{DDQ}. Current standards include 2.5V for DDR and DDR1, 1.8V for DDR2, and 1.5V for DDR3. DDR4, which should debut in 2014, will have a voltage of 1.05 to 1.2V, depending on how far the technology advances before the release of the standard.

DDR memory's V_{DDQ} is the simplest supply rail. A variety of POL power sources can supply most DDR-memory devices because they allow 3 to 5% tolerance. Single-chip, onboard memory for small embedded systems might require only a linear regulator to provide 1 or 2A of current. Large multichip systems or small banks of DDR modules typically require several amps of current and demand a small switch-mode regulator to meet efficiency and power-dissipation needs. Large multimodule banks, such as high-performance processing systems, large data-logging applications, and testers, may demand 60A or more of V_{DDQ}, driving designers to develop processor-core-like, multiphase power supplies just to meet memory needs.

Although a conventional converter can typically support V_{DDQ}, it generally requires prebias support and the ability to regulate through high-speed transients as the memory switches states. No defined standard exists for prebias support, but it implies that the POL converter providing the V_{DDQ} must prevent sinking current from the V_{DDQ} supply if any voltage is stored on the V_{DDQ} bypass and output capacitors during V_{DDQ} power-up. This requirement is critical because SSTL devices commonly contain parasitic and protection diodes between V_{DDQ} and other supply voltages; V_{DDQ} can damage these diodes if it sinks current through them during start-up.

High-speed memory cells rapidly switch states. A memory chip or a module may change from a low-intensity sleep state, a standby state, or a self-refresh state to a highly demanding read-write cycle in just a few clock cycles. This rapid switching places another strong demand on the POL supply providing the V_{DDQ}. V_{DDQ} supplies should switch from only 10% of their maximum load current to 90% in 1 to 2 µsec. An array of small, local bypass capacitors near each V_{DD}, V_{DDQ}, and V_{DDL} input typically provides faster, cycle-by-cycle transitions to the memory device, and a combination of large output capacitors and high-speed control loops provides for sustained mode transitions and meets the tight accuracy requirements of DDR memory.

SMALL MEMORY SYSTEMS TYPICALLY USE A SMALL RESISTOR DIVIDER TO MINIMIZE ANY VARIATION IN THE THRESHOLD VOLTAGE.

V_{TTREF} REALIZES WIDER NOISE MARGINS

Whereas V_{DDQ} is a high-current supply that powers the core, the I/O, and the logic of the memory, V_{TTREF} is a low-current, precision reference voltage that provides a threshold between a logic high (one) and a logic low (zero) that adapts to changes in the I/O supply voltage. By providing a precision threshold that adapts to the supply voltage, V_{TTREF} realizes wider noise margins than those possible with a fixed threshold and normal variations in termination and drive impedance. Specifications vary from device manufacturer to manufacturer, but the most common specification is 0.49 to 0.51 times V_{DDQ} and draws only tens to hundreds of microamps.

Small memory systems using one or a few ICs typically use a simple resistor divider, employing the low leakage currents of the reference input voltages to minimize any variation in the threshold voltage and achieve the 2% tolerance necessary to realize the best possible noise margins. Large systems using multiple memory modules, such as standard DIMMs (dual-in-line-memory modules), typically elect a less sensitive,

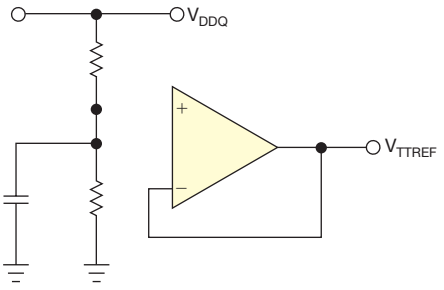

Figure 1 Large systems using multiple memory modules, such as standard DIMMs, typically elect a less sensitive, active approach to V_{TTREF}, such as an operational-amplifier buffer after the resistor divider or a voltage from a dedicated DDR memory, such as Texas Instruments' TPS51116, TPS51100, or TPS51200 switchers and low-dropout regulators.

active approach to V_{TTREF}, such as an operational-amplifier buffer after the resistor divider or a voltage supplied by a dedicated DDR memory (**Figure 1**), such as Texas Instruments' TPS51116, TPS51100, or TPS51200 switchers and low-dropout regulators.

V_{TTREF} should always be locally generated, referencing the V_{DDQ} at the source device to provide the most accurate threshold voltage and the widest possible noise margin. That requirement dictates that the memory's V_{TTREF} must reference the processor's V_{DDQ} and that the processor's V_{TTREF} must reference the memory's V_{DDQ}.

V_{TT} CLOSELY TRACKS V_{TTREF}

V_{DDQ} and V_{TTREF} are relatively straightforward power supplies. Most power converters source current to their output to maintain a regulated voltage or current at their output. This scenario is true even of a termination voltage for conventional logic, which is equal to the logic's I/O voltage.

The midrail V_{TT} in SSTL and DDR-memory devices is different. When the SSTL circuit generates a zero, an active pulldown device sinks current from the termination rail, and the termination supply acts as a conventional supply voltage, sourcing the required current to maintain the desired termination voltage. However, when the logic circuit generates a one, a pullup device sources current into the termination rail, and the termination supply must suddenly become a load, sinking current from the memory output. This sink-and-source requirement significantly increases the complexity of the V_{TT} design but provides a valuable feature to the memory device. Each logic one sources current into the termination, and each logic zero sinks a similar current from the termination. Therefore, the termination supply needs to support only the differential current—zeros minus ones—normalizing the load currents and improving signal integrity over rail-terminated logic.

Further increasing its complexity, V_{TT} must closely track V_{TTREF}. SSTL voltages are so low that small variations between V_{TTREF} and V_{TT} could quickly erode the noise margin and degrade signal integrity. Conventional regulators compare a divided version of the output voltage to a high-precision reference voltage to set an output voltage. V_{TT}, on the other hand, must compare an output voltage to the V_{TTREF} to ensure the widest possible noise margin, largest eye windows, and most accurate data transfer. This tracking requirement dramatically limits the range of available devices that can perform this task.

V_{TT} experiences more severe transients than most other power supplies that support similar current. V_{DDQ}, for example, might switch from 10 to 90% of nominal load in a few microseconds. If the data and address lines of a DDR memory switch from all zeros to all ones, however, the load seen by the termination supply rapidly changes from sourcing its maximum load to sinking its maximum load. This 200% load step makes transient performance critical for DDR-memory power.

For simplicity and load balancing, V_{TT} is usually generated using a tracking sink-source linear device, such as a high-current op amp or a dedicated sink-source low-dropout regulator. In such applications, V_{DDQ} commonly generates V_{TT}. This approach minimizes the power loss in a linear device by providing a low source voltage and normalizes the load current on V_{DDQ}. Any address line that does not draw current from the I/O function of V_{DDQ} by generating a one is sinking current from V_{TT}, which is drawing the same current from V_{DDQ}. Although not the most efficient approach, it does provide a consistent load equal to one-half- to one-times the V_{TT} current, I_{TT}, on the V_{DDQ} supply.

In large systems that terminate hundreds of lines, such as the long recording arrays in digitizing test equipment, or in extremely power-sensitive systems, such as battery-powered systems that must operate for extended periods without recharging, you can use a tracking synchronous buck switcher, such as TI's TPS40042 (**Figure 2**). A synchronous buck converter can draw current from its output, returning the recovered energy to its input voltage, much like a boost converter. This sink/source capability, along with the efficiency of a synchronous buck converter, makes the synchronous buck converter an ideal choice for high-current or high-efficiency termination. When using a tracking syn-

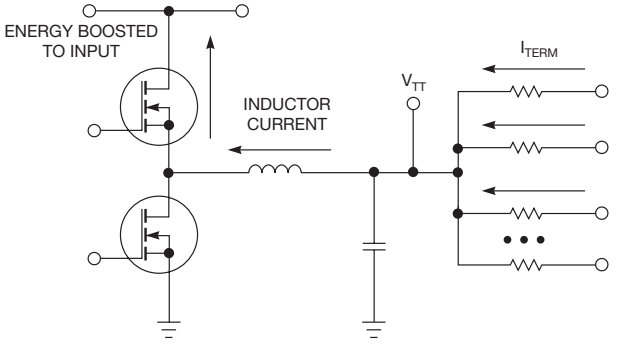

Figure 2 In large systems that terminate hundreds of lines, such as the long recording arrays in digitizing test equipment or in extremely power-sensitive systems, such as battery-powered systems that must operate for extended periods without recharging, you can use a tracking synchronous buck switcher, such as Texas Instruments' TPS40042.

chronous buck converter to realize the termination voltage, consider the source voltage for the termination converter. Although it is often a good idea to operate the V_{TT} regulator from V_{DDQ}, the low V_{DDQ} might make this approach impractical. Further, the cascade effect of double conversion can sacrifice some efficiency benefits that you might be able to realize with alternative schemes.

When V_{DDQ} is not used as the source for an active, switcher-based termination regulator, the termination regulator should share a common source with the V_{DDQ} regulator. This approach ensures that, when sinking current from logic ones, the active switcher cannot source more energy into its supply than the V_{DDQ} regulator is drawing. This step eliminates the need for the V_{TT} sourcing supply to also sink load current and prevents a dangerous overvoltage condition at the source of the V_{TT} regulator.

HOW MUCH POWER DO I NEED?

It's easy to figure out how much power your design needs, even if it's unclear after reading a data sheet. Consider, for example, the specifications provided for a 256-Mbyte SDRAM. The data sheet lists I_{DD} (drain-to-drain current), along with the expected input current for each state of operation. You could analyze the full operational cycle of the memory device, considering the worst possible rolling 10-µsec average, and design a power solution that can deliver only that much current. It is typically advantageous, however, to design a power solution that can provide continuous current equal to the maximum expected from the memory device.

When examining the data sheet's I_{DD} value, consider two key test conditions: output current and on-die termination. For example, a configuration with an output current of 0 mA with on-die termination disabled assumes no load and unterminated outputs. Any usable system, however, must have terminated outputs, and V_{DDQ} must be able to supply the full termination current when all outputs are high. The V_{DDQ} regulator must thus be able to source this maximum I_{DD} plus the full V_{TT} current.

A device with 16 data lines and as many as eight differential strobe lines has a maximum of 20 terminations, each of which can drive an output voltage as high as 1.8V into a typical 50Ω on-die or discrete termination resistor to a 0.9V termination voltage for 360 mA of current. This 256-Mbyte DDR2 memory chip might require as much as 680 mA of 1.8V supply current to maintain its I/O, logic, and core operations. For termination current on the V_{TT} supply, it is also necessary to terminate address, bank-address, clock, chip-enable, and other memory-logic lines for 38 terminations and 684 mA of termination current. You can ignore differential pairs because the termination current sourced by one will automatically sink into the other.

The V_{TTREF} leakage current, when V_{TTREF} is at a valid level—that is one-half of V_{DDQ}—is always less than 2 µA. If you plan to use a simple resistor divider, it is important to note that 2 µA of current across just 9 kΩ of resistance would introduce 18 mV of error, the total allowable error in the programming of V_{TTREF}, so relatively low resistor values should be used.

Power requirements become more complex as memory ICs are added. You need to consider whether the system allows multiple memory ICs or just one IC to operate in the high-power interleaved-read state at the same time. Typically, only one memory device sharing common address and data lines can be in this active state at a time. All other shared devices are in a lower-power state, such as a burst refresh.

After reviewing the worst-case operational state of each shared-memory device, add the I_{DD} for these states for all of the memory devices. For example, a 1-Gbyte memory system might use four of these 256-Mbyte ICs sharing the same address and data lines, with added logic to select which of the four memory devices the memory is accessing, effectively increasing the number of available address lines by two. After reviewing the operational states, you'll find that one device will be effective in continuous interleaved read, drawing 320 mA of current, but the other three devices will remain in burst refresh, drawing 180 mA each for a total of 860 mA of I_{DD} current. Because all of the outputs will share a common data bus and only one device drives these lines, 360 mA of additional V_{DD} current requires the same 20 output currents for a total of 1.22A of current from the 1.8V supply.

YOU CAN TRANSFORM MEMORY POWER FROM A DAUNTING TASK TO A VALUABLE ADDITION TO DESIGN BY MAXIMIZING NOISE MARGIN.

Termination current can share the termination lines, using common data, address, and bank selection. Chip-enable and command lines must be separate so that each chip can be accessed separately; therefore, each chip requires four additional terminations, bringing the total terminations to 50, for 0.9A of total termination current. With four chips, each able to sink 2 µA of V_{TTREF} current, even lower resistor-divider values must be used. An active V_{TTREF} buffer or a separate divider for each memory IC might also be desirable.

With attention to detail and care about the needs of each supply voltage, you can transform memory power from a daunting design task to a valuable addition to design by maximizing noise margin and improving accuracy. By understanding the function of each supply voltage, designers can more confidently select the power designs that best meet their overall design goals and balance among size, power, efficiency, performance, and cost. EDN

ACKNOWLEDGMENT

This article originally appeared on EDN*'s sister site,* Memory Designline, *http://bit.ly/n9oLAo.*

AUTHOR'S BIOGRAPHY

Peter James Miller is an applications engineer at Texas Instruments, where he supports technical applications for dc/dc controllers and evaluation-board design and develops internal and external application training for dc/dc controllers. Miller received both a bachelor's degree and a master's degree in electrical engineering from Worcester Polytechnic Institute (Worcester, MA).