ISSUE: September 2019

Using Mini E-Loads To Emulate Fast Slew Rates And High Peak Power Levels

by Viktor Vogman, Power Conversion Consulting, Olympia, Wash.

Recently, testing the transient performance of CPU voltage regulators (VRs) and server power supplies has become a challenge due to the increased current slew rates of their loads and the high peak power levels associated with virus conditions. It is not unusual to see dynamic testing requirements for a regulator with an output voltage of 1 V, a load step of 100 A, and a slew rate of 1,000 A/µs,[1] whereas the unit supplying power to the entire platform may have load current spikes that can bring its peak power level to several kilowatts.

Given these requirements, there are different obstacles when using the existing test instruments to perform transient testing on VRs and server power supplies. The limiting factor in most low-voltage VR dynamic loading situations is the series inductance between the unit under test and the comparatively bulky electronic load (e-load). On the other hand, applying high current spikes to an ac-dc power supply output typically demands paralleling and synchronizing multiple e-load modules, which is not trivial to set up.

These difficulties make it worthwhile to develop a new technique of generating transient loading conditions. What's needed is a *miniature* e-load capable of generating high-current slew rates and magnitudes. Such devices are known as slammers, or some variation on this name, and are available for purchase from vendors. You can also build these test tools to address your specific test needs and this article describes the design and operation of one I call a *microslammer*. This particular design addresses the need to test CPU power virus protection. However, variations of this design can also be used for testing power supplies and bus-bar PDNs.

After describing the requirements for the microslammer, the essential elements of its design are explained including methods for determining the maximum current and pulse duration as well as pulse repetition rates that the microslammer can generate, based on the choice of power MOSFET. Methods for operating the microslammer with a laboratory-grade signal generator or a PC control interface are also described. Finally, experimental results are presented that demonstrate the use of microslammer prototypes for testing CPU virus protection.

Replicating Peak Power Events In Operating Systems

A power virus forcing all CPU cores to transition into an excessive power (turbo) mode can have a major impact on computer or server system operation. It could cause a momentary overloading of the CPU VR or power supply, excessive supply voltage sags, or/and a system crash. To prevent these events from happening different protection methods are used for detecting power virus events and suppressing them.[2]

To reliably verify platform immunity to such events and quantify the benefits of such protection techniques, the virus conditions often need to be emulated in an operational system. Emulating a CPU turbo power virus condition by means of software, especially on a multi-socket motherboard, is quite a difficult task. Besides the impossibility of precisely controlling the peak power magnitude with software there is a problem associated with syncing such events on several processors. Sometimes it takes a long time to create a single virus event which may not even represent the worst-case condition.

Hardware tools are much more flexible and usually provide a more efficient solution to this problem. But most of these tools are relatively bulky because they are historically designed to continuously support the CPU's thermal design current (TDC), which is the sustained (dc equivalent) current that the processor is capable of drawing indefinitely.[3] Therefore, these tools do not permit server platform designers to validate the platform's immunity to power viruses in a real environment.

The microslammer described in this article is a miniature, inexpensive and efficient validation transient test tool that can be placed anywhere in the dc power delivery path and is capable of emulating power virus conditions when all system components are installed and operational. This test tool can be connected to the power supply unit (PSU) or VR output.

The microslammer (MS) power stage connection block diagrams are shown in Fig. 1. These diagrams show how the MS power stage can be controlled with power virus detector signals generated by the PSU, VR or CPU itself to interrupt/suppress the peak power pulse as it is supposed to be done in an actual power virus case. Similar to a conventional VR transient test tool (VRTT),[3] the MS must interface its current and voltage waveforms, as shown in the block diagram. Besides smaller size this tool needs to have the following advantages over traditional e-load solutions:

- Provide minimal series parasitic inductance, which allows it to achieve fundamentally higher current slew rates.
- Be able to operate in parallel with installed CPUs and be compatible with external PC or basic signal generator control, allowing for fast and simple test setup, without synchronizing multiple channels to achieve high peak-power levels.
- Be easily configurable for closed-loop control, i.e. reduction of a self-imposed load based on an external virus detector signal fed to their control input.

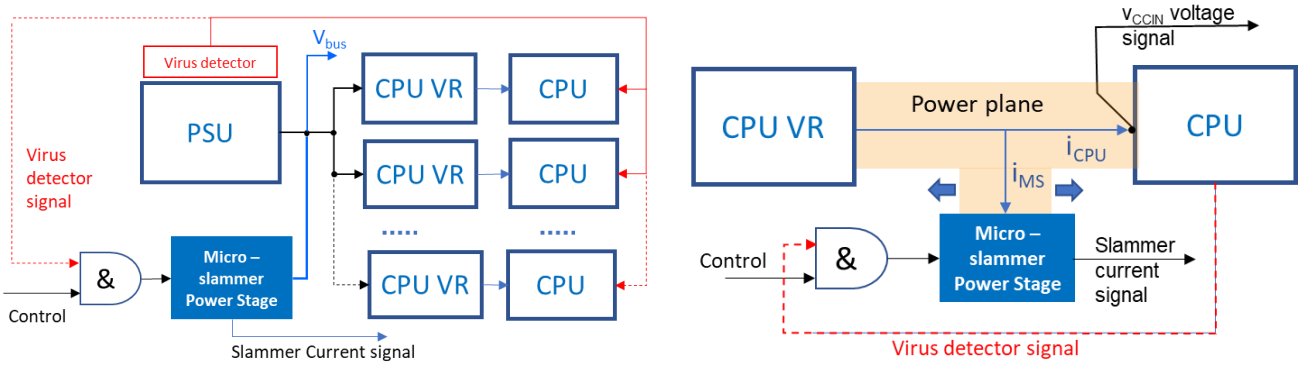

Fig. 1. Microslammer power stage connection options. The power stage's small size and footprint allow it to attach virtually anywhere in the dc power delivery path such as at the PSU output (a) for testing a dc bus protection feature or to the CPU VR–CPU power plane (b) and operate it in parallel with an active self-protected CPU.

Essentially, there are two mandatory conditions that must be met to reliably sink hundreds of amps of current with the miniature power stage shown in Fig. 1:

- All MS power stage components must operate in their safe operating areas
- Power dissipation in the MS must be minimized to guarantee its operation without heatsinks.

Let's consider design techniques satisfying these conditions.

Miniature Power Stage Design Considerations

The simplest analog-controlled power stage topology containing only two components can be realized with a common-drain (source follower) power amplifier, in which the current setting and sensing features are combined in one grounded component—the source resistor. The source follower can actively sink current from the supply voltage source through a transistor drain terminal, while the sunk current magnitude can be controlled by a low-power input signal applied to its gate.

First, let's consider a microslammer operating from a low CPU core supply voltage V_{ccin} although the same concepts can be applied to a power stage attached to the higher V_{bus} voltage generated by the PSU. The power stage's general and practical schematic diagrams are shown in Fig. 2.

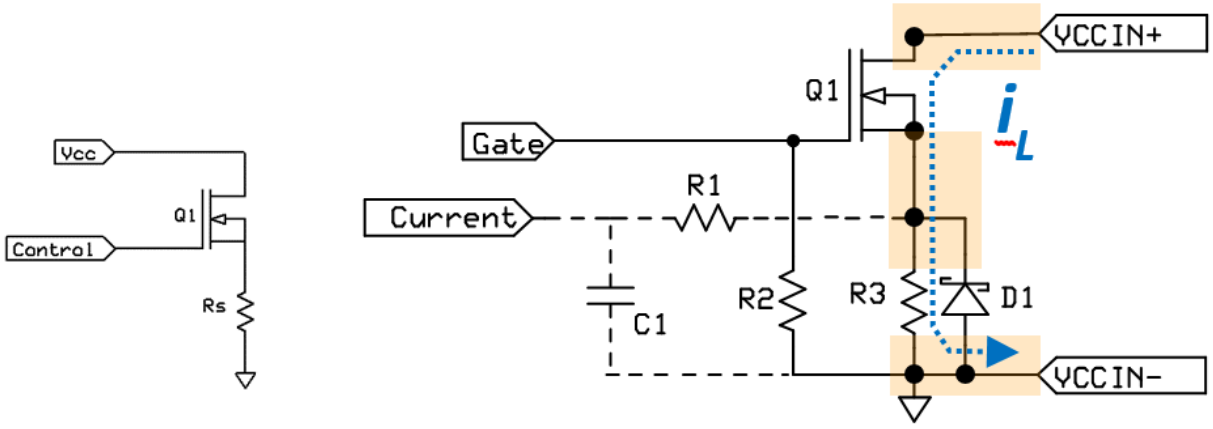

Fig. 2. Simplified schematic (a) and practical implementation (b) of a microslammer power stage. Load current i_L level can be set by the voltage level at Q1 gate terminal and resistance R_s (R_3) in the source network. The actual current signal is sensed with RC network R1, C1 having a time constant matching the ESL/R time constant of the source resistor R_s (R3). Diode D1 clamps the inductive spike across R3 when Q1 turns off.

The MOSFET (Q1) in the power amplifier generating high peak current pulses operates with an ultra-low average duty cycle to achieve low MS power dissipation and to provide its small physical dimensions. Resistor R3 represents the amplifier load and diode D1 acts as an inductive spike clamping component, permitting usage of a low-voltage-rated and low on-resistance active component.

Sinking high currents at low supply voltages can be considered feasible if the combined voltage drop across R3 and fully conducting MOSFET Q1 (Fig. 2b) at the maximum required load current magnitude I_{Lmax} is less than the minimum supply voltage $V_{ccin.min}$: $(R_{DS(ON)} + R_3)\cdot I_{Lmax} < V_{ccin.min}$, or:

$$R_3 < \frac{V_{ccin.min}}{I_{Lmax}} - R_{dson} \qquad (1)$$

The current flowing through R3, which is equal to the MOSFET drain current i_D, can be described by the following equation:

$$i_D = \frac{(V_{in} - V_{thr})\cdot g_m}{R_1 \cdot g_m + 1} \qquad (2)$$

where V_{in} is a pulse magnitude at the gate, V_{thr} is gate threshold voltage and g_m is the MOSFET forward transconductance. Equation (2) allows us to determine the maximum control-signal magnitude required to support the specified maximum load (drain) current $I_{D.max}$:

$$V_{in.max} = I_{D.max} \cdot [(R_3 + 1/g_m(I_{D.max})] + V_{thr} \qquad (3)$$

If PC control is intended for the tool it would be beneficial to keep $V_{in.max}$ below the USB minimum voltage, because this would eliminate the need for an additional voltage boost stage or housekeeping power supply.

The voltage v_{DS} across the MOSFET can be defined as a difference between the VR output voltage V_{CCIN} and a voltage across source resistor R3:

$$v_{DS}(i_D) = v_{CCIN}(i_D = 0) - (R_3 + R_{LL}) \cdot i_D \qquad (4)$$

where R_{LL} is the VR load line resistance at the remote sense points.

Expressions (1) through (4) are also suitable for a "high voltage" MS connected to the PSU output if V_{ccin} is replaced with a PSU output voltage acting as a supply voltage V_{cc} for the e-load power stage. Next, using equation (4) and the MOSFET's safe operating area graphs let's establish the microslammer safe operating conditions.

MS Maximum Current And Pulse Duration

The time elapsed from the application of a load step until the supply voltage reaches a new steady-state level, determined by its load line, or so-called settling time, is always greater than zero. That is why for safe MS operation in the dynamic mode we need to assume in equation (4) that R_{LL} = 0.

Plotting the $v_{DS}(i_D)$ graph over the MOSFET dynamic safe operating area (SOA) chart allows us to determine safe boundaries for Q1's operating point. Since the SOA essentially represents energy dissipation limits as a function of operating conditions (v_{DS}, i_D) it allows us to find a maximum single pulse- or burst duration, which would guarantee reliable operation of the MS power stage.

The y-coordinate of the highest crossing point of the MS load line with the dynamic SOA plot boundary determines the absolute maximum MS current I_{Dm} that can be achieved with a selected MOSFET and source resistor value. The maximum instantaneous power dissipation in the MOSFET can be derived from the Q1 power equation. Since the MOSFET represents a variable resistance $R_Q = v_{DS}(i_D)/i_D$ controlled by the gate signal, we can write the following equation for instantaneous power dissipated in it:

$$P_{Q1}(R_Q) = \frac{V_{cc}^2 \cdot R_Q}{(R_Q + R_s)^2}$$

where V_{cc} is the e-load supply voltage and R_s is the value of the resistor connected to the MOSFET source terminal (Fig. 2a). Setting the first derivative of function $P_{Q1}(R_Q)$ with respect to variable R_Q equal to zero, we find that the power dissipated in the MOSFET reaches its maximum ($P_{Q1max} = V_{cc}^2/4R_s$) when $R_Q = R_s$ or when $v_{DS} = v_{Rs} = V_{cc}/2$ and $I_D = V_{cc}/2R_s$. This point on the SOA chart area determines the maximum pulse duration guaranteeing safe MOSFET operation.

The process for determining MS maximum current and pulse duration is illustrated by the examples in Fig. 3.

For a "low voltage" (V_{cc} = 1.8 V) application (Fig. 3a), selecting a power MOSFET with 400-A pulsed drain current rating for Q1, R_s= 2.5 mΩ yields the maximum MS current I_{Dm} =400 A and max pulse duration of 1 ms. In practice, the max current magnitude will need to be limited to some lower level such as 350 A to provide design margin. Furthermore, the pulse duration could be increased to 10 ms, for example, if i_{Dm} was limited (by the gate control signal) to 58 A (Fig. 3a).

Usually a 1-ms pulse duration is sufficient for a CPU VR application. However, if a longer (e.g. 10 ms) pulse duration is required, then several stages with a lower current limit need to be paralleled to achieve the 350-A magnitude.

For a "high voltage" (V_{cc} = 12 V) application (Fig. 3b), the selected MOSFET SOA yields I_{Dm} = 200 A and a maximum pulse duration of 1 ms at R_s = 50 mΩ and I_{Dm} =110 A or 10 ms at R_s = 100 mΩ. As it can be seen from these examples, the pulse duration can be increased to 10 ms if the R_s value is doubled (blue line in Fig. 3b), but to achieve the same (200 A) current magnitude would require paralleling of two stages with R_s = 100 mΩ.

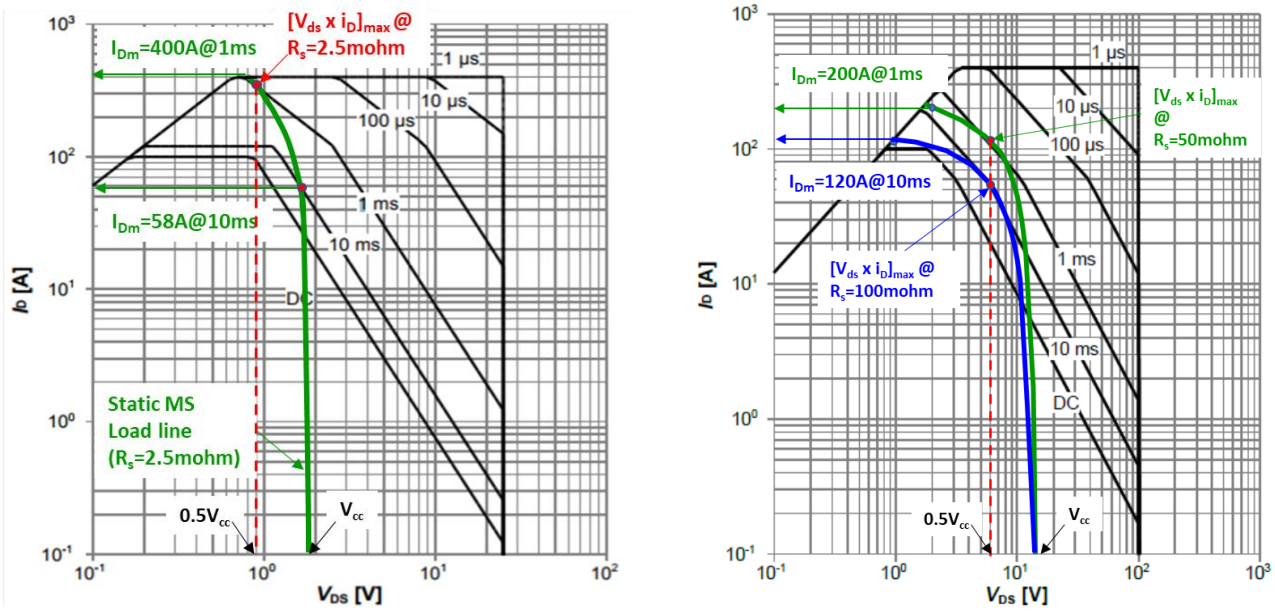

Fig. 3. $v_{DS}(i_D)$ and slammer MOSFET SOA plots. Interposing the load line for a given source resistance R_s over a selected MOSFET SOA diagram allows us to determine a MS maximum peak current and pulse/burst duration. As graph (a) reveals, for a "low voltage" (V_{cc} = 1.8 V) application the selected MOSFET for Q1, R_s= 2.5 mΩ yields the maximum MS current I_{Dm} = 400 A and maximum pulse duration of 1 ms. Graph (b) shows that for a "high voltage" (V_{cc} = 12 V) application the selected MOSFET SOA yields I_{Dm} =200 A and maximum pulse duration of 1 ms at R_s = 50 mΩ. In both cases, the pulse duration can be increased to 10 ms if, i_{Dm} is limited to a lower value.

The above study demonstrates that increasing R_S results in a reduction of the MS maximum current magnitude, which allows us to extend the active pulse duration. Achieving higher current magnitudes with larger R_S values may require paralleling power stages. Extending the allowed pulse duration with a low (comparable with $R_{DS(ON)}$) R_S value (Fig. 3a) may also require paralleling of power stages and limiting the current magnitude in each stage by limiting the control voltage level.

For a given V_{cc}, peak current magnitude I_{Dm} and pulse duration, the following procedure can be followed to select the microslammer's main components and their key operating points:

1. Select a power MOSFET with pulsed drain current i_{Dmax} and voltage $V_{DS.max}$ ratings exceeding the required peak current magnitude I_{Dm} and V_{cc}, respectively, with sufficient margin.
2. Using Eq. (4), plot several load lines over the MOSFET SOA graph and determine max current magnitudes as y-coordinates of the highest crossing points of the MS load line and dynamic SOA plot boundary. Consider only the R_S range that provides maximum currents exceeding the I_{Dm} level, meeting condition (1).
3. On the SOA determine the max allowed pulse durations at the points on the load lines corresponding to the v_{DS} = V_{cc}/2 x-coordinate. Maximum pulse duration can also be determined by positioning on the SOA graph the $P_{Q1.max}$ data points for each of the R_S values and using V_{cc}/2 as an x-coordinate and $V_{cc}/2R_S$ as a y-coordinate.
4. Select an R_S value that provides both the max current magnitude and pulse duration needed.

Pulse Repetition Rates

If operating conditions are determined based on the load line and selected MOSFET SOA, as described above, a pulse duty cycle within each burst can vary within the 0 to 100% range. This would allow us to conduct a search for the worst-case VR transient response as recommended in references [4] and [5].

To provide the maximum possible di/dt (given its source resistor equivalent series inductance), the MS physical dimensions need to be minimized to provide minimum parasitic inductance in the MS self-current path. The smallest MS dimensions will be achieved if none of the MS's components need heatsinking. To meet this requirement a long-time average ("active window") duty cycle D_{max} must be set to guarantee that the "heatsink-less" power ratings of Q1 and R_S, which correspond to MOSFET P_{Q1r} and resistor P_{Rr}, are not exceeded.

As it was stated above, the maximum instantaneous power dissipation for Q1 will be achieved when V_{cc} is split equally between the MOSFET and its source resistor and will be equal to $P_{Q1max} = V_{CC}^2/(4 \cdot R_s)$, the maximum power dissipated in R_S at $i_D = I_{Dm}$: $P_{Rs.max} = I_{Dm}^2 \cdot R_s$. The $P_{Rs.max}$ value needs to be compared against the resistor pulse rating at a given pulse duration. Typically, at 1-ms pulse durations, the surface-mount (SM) resistor pulse power rating exceeds its continuous rating by two orders of magnitude.

Thus, maximum long-time average duty cycle can be determined as:

$$D_{max} = min\left\{\frac{P_{Q1r}}{P_{Q1max}}; \frac{P_{Rr}}{P_{Rs.max}}\right\}.$$

Normally, at 1-ms durations, the SM MOSFET and resistor power packages yield maximum long-time average duty cycles no less than 1%.

Tool Control Options

With the externally accessible MOSFET gate, the microslammer control can be made really simple and provided by a laboratory-grade signal generator having a low enough impedance to drive the gate. To make the MS setup fully automated and autonomous, additional circuitry needs to be used to interface a PC to the tool. A block-diagram of the circuitry that can be used in this case is shown in Fig. 4. The detached power stage in this block diagram can be used for restricted space applications.

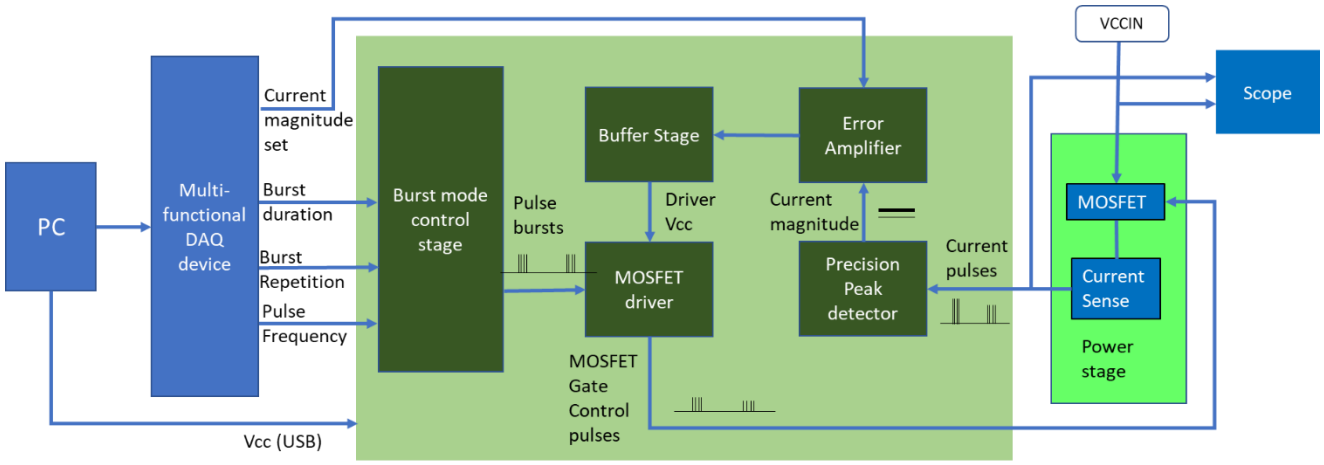

Fig. 4. Block diagram of the microslammer with PC control interface circuit and detached power stage for restricted space applications.

The tool consists of a control circuit that uses the power from a USB port and a power stage energized from the main CPU VR output generating V_{CCIN}. It operates in a burst mode which allows it to repeatedly generate high current pulses of the required frequency within the active window while keeping average power dissipation low. Such an operating mode allows the duty cycle within a burst to reach 100%, which makes it possible to apply dynamic load profiles similar to much larger-size e-loads. The timing diagram of the slammer control signal is shown in Fig. 5.

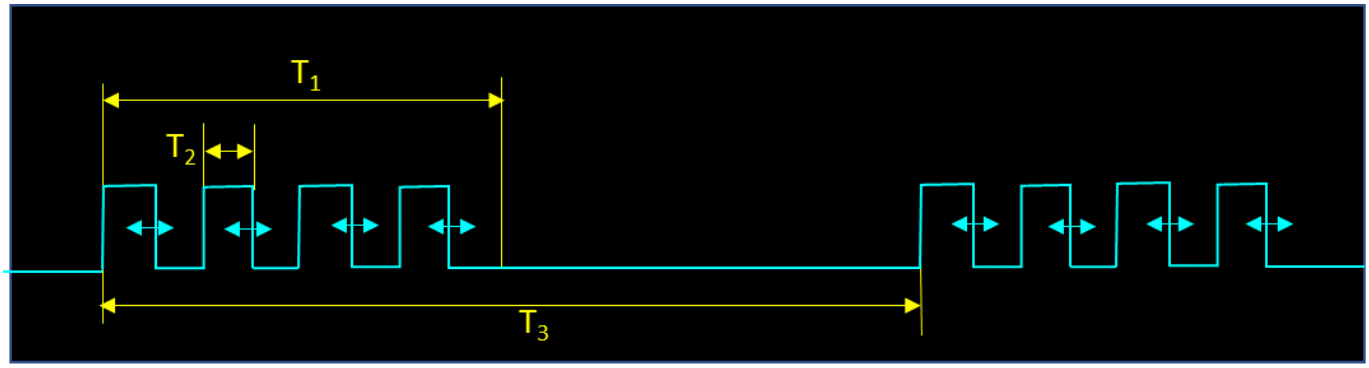

Fig. 5. Microslammer control signal. Three analog signals, fed to the burst-mode control stage (Fig. 4) by a multifunction DAQ device, set the duration of the burst (T_1), pulse duration/frequency within each burst (T_2), and the burst repetition cycle (T_3).

The control pulses generated at the input of the MOSFET driver are controlled by three analog signals. These three signals, fed to the burst-mode control stage (Fig. 4), set the duration of the burst (T_1), pulse duration within each burst (T_2), and the burst repetition cycle (T_3). The driver controls the power stage (a common-drain/source follower amplifier as shown in Fig. 2) so that the slammer current magnitude is proportional to the driver's supply voltage generated by the buffer stage (Fig. 4).

The current-sense signal from the power stage load resistor (shown in Fig. 2b) is fed to the precision peak detector, producing dc voltage equal the current signal magnitude. This magnitude is being compared with the fourth analog signal representing a reference current magnitude fed to the error amplifier (EA). The EA output provides the amplified error signal that controls V_{cc} of the MOSFET driver through a buffer stage, making the current magnitude follow the reference set by the DAQ device.

Control of the tool can be done manually by providing the four analog input signals directly to the interface board (the DAQ device or an MCU) or through a Software/LabView controlled DAC, setting repetition, durations, and magnitudes of the generated current pulses.

Experimental Results

A low-voltage MS power stage prototype was built using one OptiMOS power MOSFET (part number BSC010NE2LS) with a 400-A pulsed drain current rating. The MS is able to reliably sink load currents up to 270 A at a 1-ms pulse duration. There are two versions of the MS eval board. One contains an integrated power stage, while the other has a separate power stage, which is attached to the motherboard. These two versions are pictured in Fig. 5a and 5b, respectively.

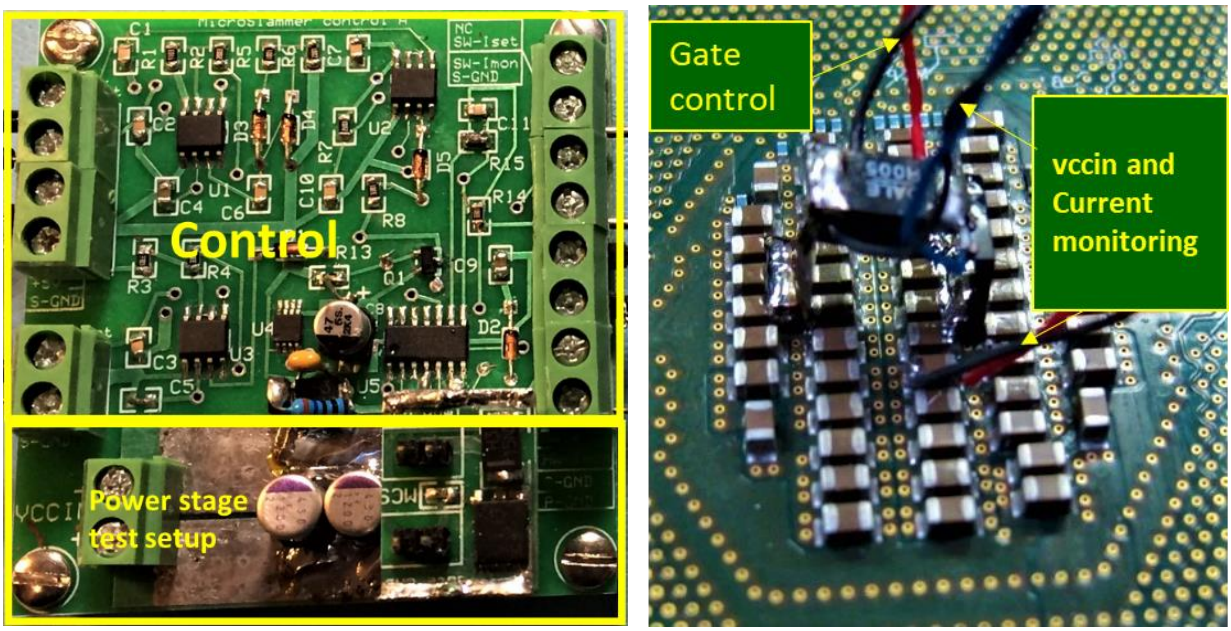

(a) (b)

Fig. 5. Microslammer hardware. A microslammer eval board with an integrated power stage (a) and a microslammer power stage attached to the motherboard (b).

Fig. 6 shows current waveforms for the microslammer control circuitry and power stage. Fig. 6a illustrates burst-mode operation with three pulses. Fig. 6b shows the gate-burst-duration set pulse on a long time scale, demonstrating the ultra-low long-time average duty cycle. Finally, Fig. 6c offers a zoomed-in view of the gate-control- and power-stage current signals.

Note that with hard gate drive (high current, high voltage and short rise/fall times) required for the highest load current di/dt, the current slew rate at turn-off gets slightly higher because the R_S-ESL time constant during discharge is larger than the time constant during charge.

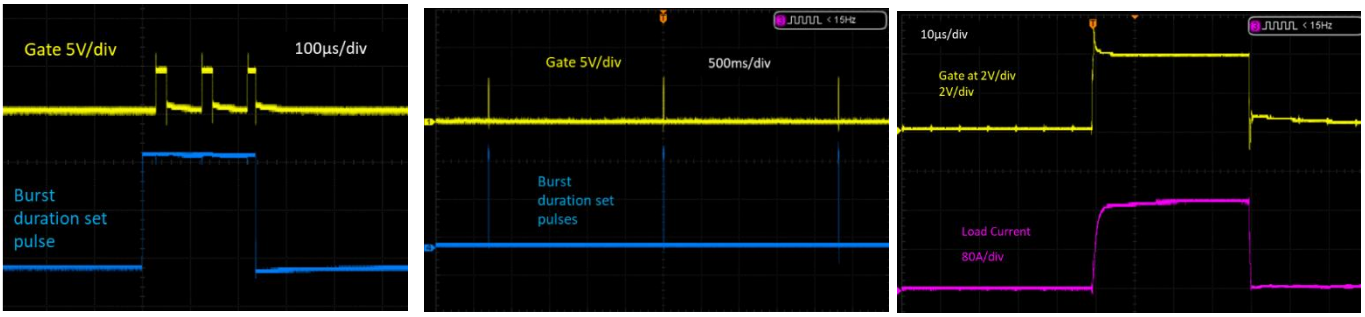

(a) (b) (c)

Fig. 6. Microslammer control and power-stage current waveforms. The gate signal and burst-duration set pulse are measured in (a) on a time scale of 100 µs/div. Then, the gate signal and burst-duration set pulse are captured in (b) on a much longer time scale of 500 ms/div. Finally, zooming in, the gate-control and compensated current signals are shown on a 10-µs/div scale.

Fig. 7 gives examples of using the microslammer for testing a CPU virus protection feature in an operating system for single pulse- and burst mode. In this case, the MS enables measurement of the threshold and the

time delay of the CPU protection, as well as verification of system immunity to the power virus when CPU virus detector is activated.

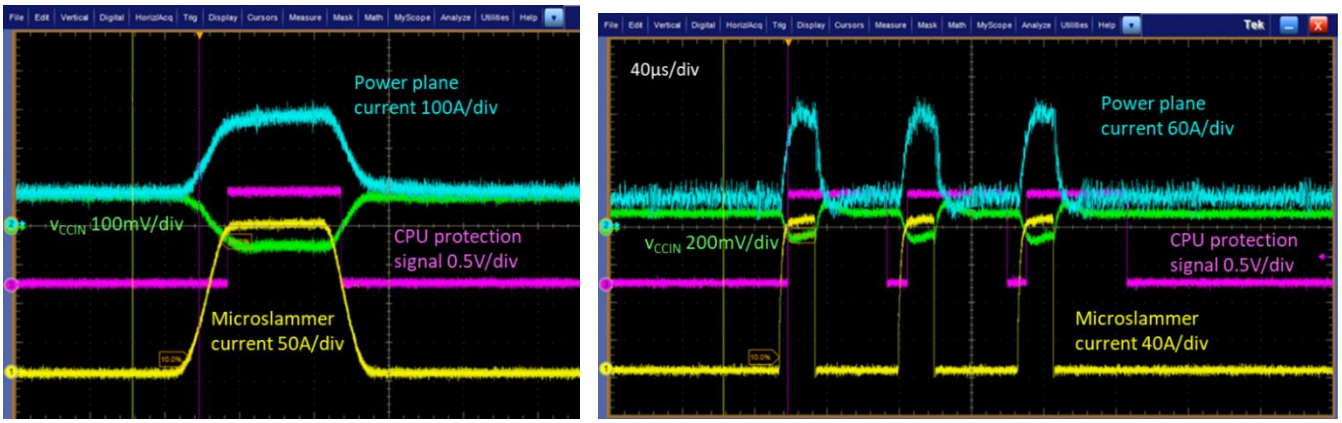

(a) *(b)*

Fig. 7. Using the microslammer for testing CPU virus protection. The microslammer's single-pulse mode enables measurement of the threshold and time delay of the CPU power virus protection signal (a). Alternatively, a burst mode allows measurement of the power virus protection signal hold times with multiple power virus events (b).

In a second experiment, a "high voltage" (12-V) microslammer power stage prototype was built from three parallel channels using an n-channel enhancement-mode power MOSFET, the FQP65N06. The prototype is able to sink load currents up to 200 A at a 10-ms pulse duration. Such an e-load can be used for validation of system protection features, such as closed-loop system throttling, as shown in the block-diagram in Fig. 1a. It can also be used to simplify setups for PSU transient loading testing or generating high power spikes for PDN overpower protection testing.

Fig. 8 demonstrates the example of using the microslammer for virus protection testing of a 12-V dc bus that uses a power virus detector signal to provide closed-loop control of the bus's power consumption. If the power virus detector signal fed to the MS, as shown in Fig. 1a, is not enabled (Fig. 8a) then a dc bus designed to support a real application power not exceeding 2.4 kW crashes when peak power exceeds this level for longer than 100 µs. When the detector signal is activated, the MS emulates CPU throttling with the intended time delay and protects the dc bus from crashing. A similar test setup in a conventional case would require paralleling of multiple traditional e-loads controlled by external circuitry.

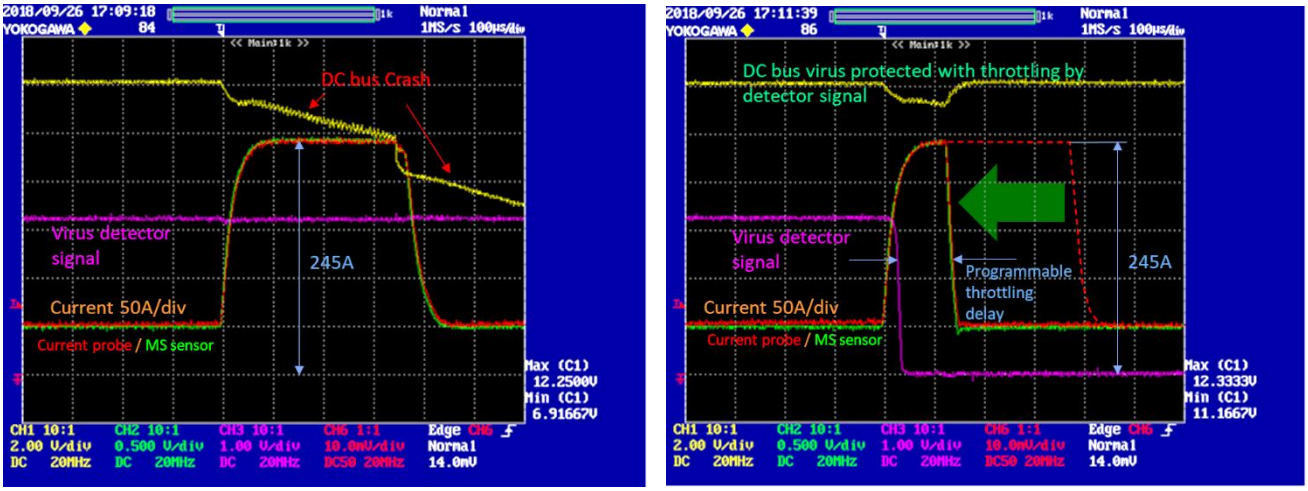

(a) *(b)*

Fig. Using the microslammer to test the virus protection of a 12-V dc bus with a power virus detector signal. If the detector signal fed to the MS is not activated, the 2.4-kW dc bus crashes when peak power exceeds the rated level for longer than 100 µs as shown in (a). However, when the detector signal throttles the MS, emulating CPU throttling with a projected delay, the power bus does not crash and rides through the virus condition as shown in (b).

Conclusions

The initial experimental results obtained with the microslammer prototypes described here prove that the tools can be used for applying virus power levels in single- and multiprocessor systems to verify the effectiveness and robustness of their virus protection techniques. The microslammer can be considered the universal instrument for testing CPU voltage regulators and power supplies in real application environments.

The MS's key distinctive features are its ability to generate high slew rates due to its minimal series inductance and to operate in parallel with active CPUs. The small size and footprint of the MS power stage allow it to attach virtually anywhere in the dc power path, or to add layout placeholders for using the MS at various stages of motherboard development.

The low-voltage microslammer is capable of generating CPU VR load current pulses in the range of 20 A to 270 A at slew rates up to 300 A/µs, while the high-voltage (PSU) microslammer can produce 20 A to 200 A at slew rates up to 30 A/µs. The miniature e-load technique described above allows testing CPU VR solutions and system virus immunity "in action" i.e. while system is running the required software and all its components including CPUs are active and operating in their normal modes.

Future Work

Future work could be focused on making the tool more user friendly by developing a graphical user interface (GUI) control for it. Interfacing the MS signal with a USB digital storage oscilloscope would allow display of the resulting waveforms on the same device, thus providing full test setup, control and data collection with a single PC.

References

1. "Power Tips #79: Load transient testing with high slew rates" by Robert Taylor, EDN, Aug. 21, 2017.

2. "Maximum Power Specification Management feature in Intel Xeon CPUs," by Nazar S. Haider and Aman Sewani, 2017 IEEE 26th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS), pages 66-68.
3. Voltage Regulator Test Tool (VRTT), Intel demo video.
4. "Optimizing Regulator Output Capacitance (Part 2): Using Load Slammers" by David Baretich, How2Power Today, September 2018 issue.
5. https://loadslammer.com/category/slammers.

About The Author

Viktor Vogman currently works at Power Conversion Consulting as an analog design engineer, specializing in the design of various power test tools for ac and dc power delivery applications. Prior to this, he spent over 20 years at Intel, focused on hardware engineering and power delivery architectures. Viktor obtained an MS degree in Radio Communication, Television and Multimedia Technology and a PhD in Power Electronics from the Saint Petersburg University of Telecommunications, Russia. Vogman holds over 50 U.S. and foreign patents and has authored over 20 articles on various aspects of power delivery and analog design.

For more information on power supply related test and measurement, see How2Power's Design Guide, locate the Design Area category and select Test and Measurement.