DALLAS SEMICONDUCTOR MAXIM

Engineering Journal

Volume Fifty-Eight

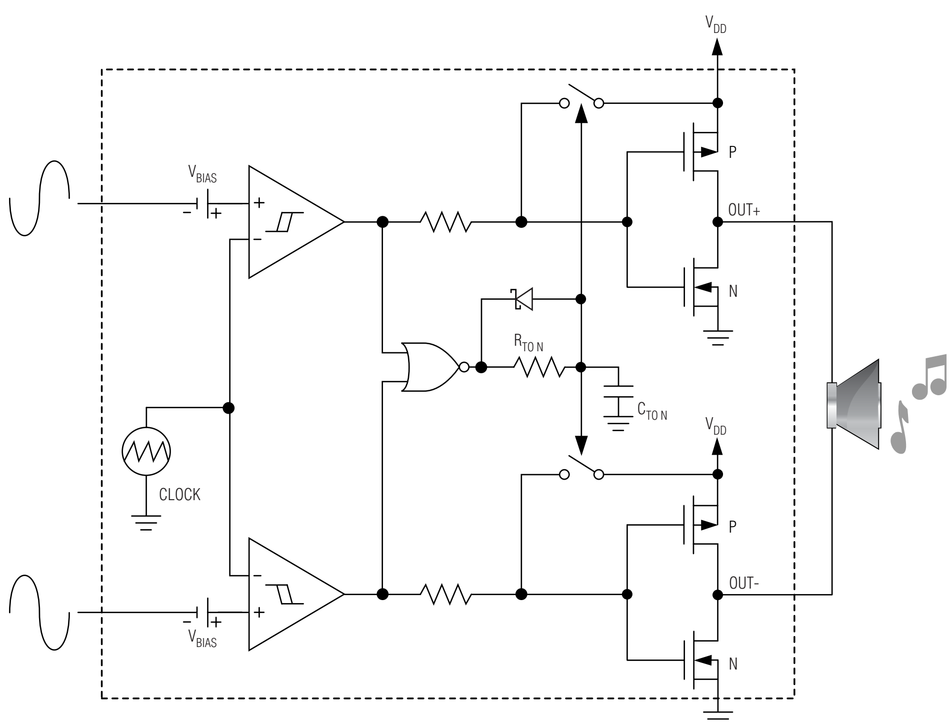

The MAX9705 Class D amplifier has an internally generated sawtooth with a differential input. If a single-ended input is used, a differential input is derived internally. (See article inside, page 13.)

News Brief

MAXIM REPORTS RECORD Q4 ’06 REVENUE, EXCEEDING Q4 ’05 REVENUE BY 27.5%

Maxim Integrated Products, Inc. (MXIM) reported record net revenues of $510.6 million for its fiscal fourth quarter ended June 24, 2006, a 6.8% increase over the previous quarter and a 27.5% increase over the same quarter of fiscal 2005.

Net income for the fourth quarter of fiscal 2006 was $124.3 million, including total stock-based compensation or $0.37 diluted earnings per share. This compares to $120.3 million or $0.36 diluted earnings per share for the third quarter of fiscal 2006, and is consistent with the $0.37 reported for the same period a year ago. Non-GAAP net income, which excludes the impact of total stock-based compensation expense, for the fourth quarter was $158.9 million or $0.48 per diluted share, a 6.7% increase in earnings per share over the previous quarter and a 29.7% over the fourth quarter of fiscal 2005. The total amount of stock-based compensation recorded during the fourth quarter was $51.3 million on a pre-tax basis, as compared to $42.1 million for the third quarter of fiscal 2006.

Gross bookings for the fourth quarter were approximately $557 million, a 4% increase from the third quarter’s level of $537 million. Bookings increased in all geographic locations excluding the U.S. where bookings were down primarily due to the transfer of certain customer order placement activity from the U.S. to Asia and Europe. Fourth quarter ending backlog shippable within the next 12 months was approximately $429 million, including approximately $366 million requested for shipment in the first quarter of fiscal 2007. The Company’s third quarter ending backlog shippable within the next 12 months was approximately $401 million, including approximately $346 million that was requested for shipment in the fourth quarter of fiscal 2006.

Research and development expense was $127.2 million or 24.9% of net revenue. Non-GAAP research and development expense, which excludes $31.9 million of stock-based compensation, was $95.3 million or 18.7% of net revenue. Selling, general and administrative expense was $37.9 million or 7.4% of net revenue. Non-GAAP selling, general and administrative expense, which excludes $10.2 million of stock-based compensation, was $27.7 million or 5.4% of net revenue. Total operating expenses for the fourth quarter were $165.1 million or 32.3% of net revenues. Total non-GAAP operating expenses decreased 0.8 percentage points from 24.9% of net revenues in the third quarter to 24.1% of net revenues in the fourth quarter of fiscal 2006.

During the quarter, the Company repurchased 2.4 million shares of its common stock for $76.7 million, paid dividends of $40.2 million, and acquired $120.8 million in capital equipment. Accounts receivable increased $32.7 million in the fourth quarter to $292.6 million and inventories for the fourth quarter increased $2.8 million to $207.4 million and includes $14.1 million of stock-based compensation.

Mr. Gifford commented: “The Company’s Board of Directors has declared a cash dividend for the first quarter of fiscal 2007 of $0.156 per share, a 25% increase over our previous dividend. Payment will be made on September 6, 2006 to stockholders of record on August 21, 2006. This increases our annualized dividend to $0.624 per share, bringing the yield to 2.1% based on the closing price on August 3, 2006.”

For the complete Q406 press release, including safe harbor information, go to: www.maxim-ic.com/NewsBrief

Meeting the Challenges of Power-Supply Design for Modern High-Current CPUs

The performance capabilities of Intel® and AMD® CPUs have increased dramatically over the past five years. This rise in performance has driven growth in the sophistication and complexity of the voltage regulators that power the CPU. The biggest challenge facing power-supply designers is keeping up with increases in power levels, tighter tolerances, and faster transient requirements, while reducing the overall cost of the power supply.

Rising Performance Requirements and Tightening Cost Constraints

Table 1 demonstrates how the performance capabilities of CPUs have grown over the past five years. Note the dramatic increase in power requirements, while voltage and specifically voltage tolerances have decreased.

1. ***Power*** One dimension that defines a voltage regulator is the number of "phases," or channels, that it can accommodate. Each phase can practically deliver 25W to 40W of power, depending on factors such as available space and cooling. A single-phase voltage regulator was sufficient for the Pentium® III, while current-generation CPUs require 3- or 4-phase regulators.

Table 1. Increased CPU Performance

Features	Pentium III	Pentium 4 Extreme
Year Introduced	2000	2005
Core Speed (Hz)	600M	3.73G
L2 Cache (bytes)	256k	2M
Front-Side Bus Speed (MHz)	100	1066
Voltage (V)	1.75	1.30
Voltage Tolerance (mV)	+40/-80	±19
Power (W)	19.6	150

2. ***Current Balancing*** One challenge in designing a multiphase power supply is ensuring that current (power) is properly shared between phases. A significantly disproportionate amount of current in one phase will stress components and degrade their lifetime. Virtually all multiphase voltage regulators incorporate circuitry to actively balance the current between phases.

3. ***Accuracy*** CPU voltages must be regulated to extremely tight tolerances in order to operate at high clock frequencies. These tight tolerances must be maintained under both static and dynamic load conditions. Static accuracy is achieved by implementing a precise on-chip reference voltage and by minimizing offset voltages and bias currents. Dynamic accuracy is affected by the control-loop bandwidth of the voltage regulator and the amount of bulk capacitance used on the regulator output. Because no regulator can respond instantaneously to a sudden change in CPU current demand, every design requires bulk capacitance. The higher the regulator control-loop bandwidth, the sooner it can 'catch up' to the CPU's demand and supplement the current provided by the bulk capacitors.

The demands placed on the CPU's voltage regulator are not without cost. Both die area and pin count scale upward with the number of phases that the regulator accommodates. High-accuracy voltage references require sophisticated design and calibration techniques. Amplifiers used for voltage and current sensing, basic voltage regulation, and active current sharing must be fast with low offset errors and bias currents and stable over process and temperature.

Perhaps the most significant challenge facing high-power CPU regulator designs is cost. Price per phase for a CPU core voltage regulator has declined 4X or more over the past five years.

Basics of Power-Supply Control

Virtually all multiphase voltage regulators use a form of PWM. Most are fixed frequency, whereby a clock initiates turn-on of the high-side MOSFET (see Q_{HI} in **Figure 1**) and allows the input supply to charge the inductor.

When the control loop determines that it is time to terminate this "on-pulse," the high-side MOSFET is turned off and the low-side MOSFET is turned on (Q_{LO}), allowing the inductor to discharge into the load. This type of PWM control is called trailing-edge modulation, because the leading edge (high-side, turn-on) is fixed (by the internal clock) and the trailing edge (high-side, turn-off) varies based on the control loop and real-time conditions. The percentage of time that the high-side MOSFET is on relative to the clock period is the duty cycle (D) and equals V_{OUT}/V_{IN} under steady-state conditions.

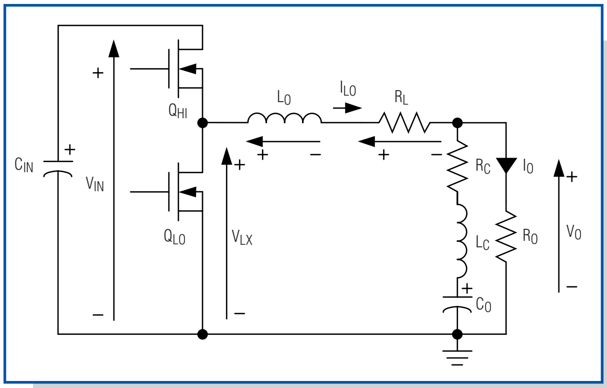

Figure 1. The input supply charges the inductor in this simplified single-phase buck regulator.

Voltage mode (**Figure 2**) compares the output voltage (or a scaled version thereof) with a fixed internal reference voltage. The result is an error signal that is compared to a fixed internal sawtooth (or ramp) signal. The ramp signal initiates coincident with the clock pulse, and the output of the PWM comparator remains high as long as the ramp signal is below the error voltage. When the ramp signal crosses through the error voltage, the output of the PWM comparator goes low and the on-pulse is terminated. The voltage loop maintains the output-voltage constant by suitably adjusting the control voltage (V_C) and, therefore, the duty cycle (**Figure 3**).

Peak current mode (**Figure 4**) adds current information to the control loop by replacing the internal ramp used in voltage mode with the ramp generated by the current in the inductor. As with voltage mode, the fixed-frequency clock turns on the high-side MOSFET, causing the inductor current to ramp up. When the (scaled) peak inductor current is equal to the error voltage, the on-pulse is terminated and the high-side MOSFET is turned off. With this approach, there is an outer voltage loop and an inner current loop. The outer voltage loop maintains the output voltage constant by suitably programming the peak inductor current measured by the inner current loop.

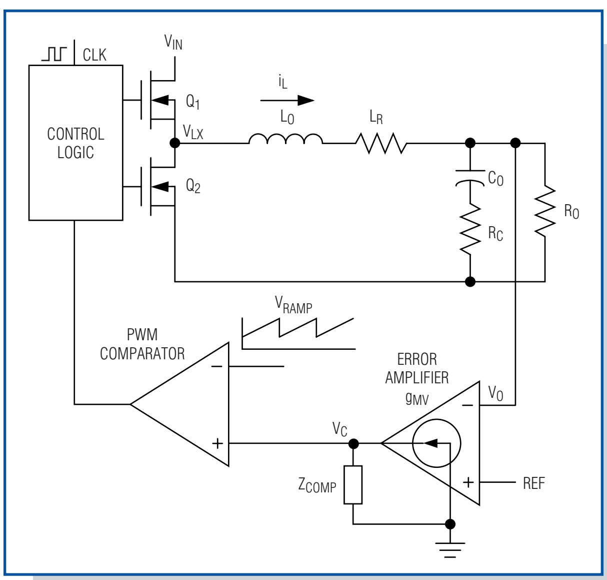

Figure 2. A simplified voltage-mode buck regulator compares the output voltage with a fixed internal reference voltage.

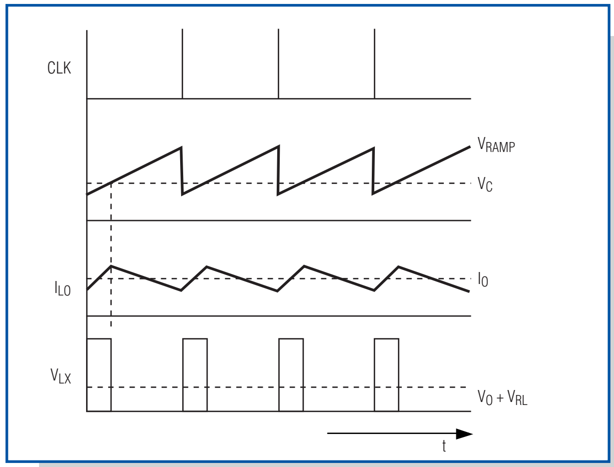

Figure 3. These voltage-mode waveforms show adjustment of the control voltage.

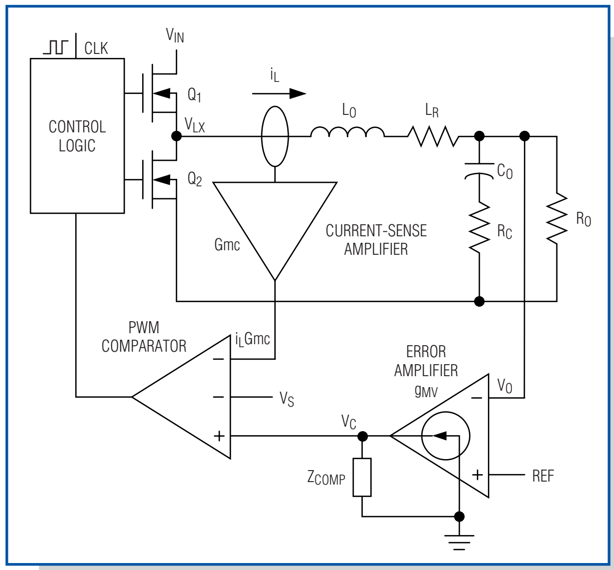

Figure 4. Current information is added to the control loop in this simplified peak current-mode buck regulator.

Tradeoffs and Considerations

As one would expect, there are pros and cons to each approach. The following paragraphs discuss several of these considerations for power-supply design.

Noise Immunity

Voltage mode has good noise immunity because the magnitude of the ramp signal can be made as large as practical when the control IC is designed. The output voltage is the only sensitive signal routed back to the controller, so voltage mode is relatively easy to lay out.

Peak current mode requires that an external current-sense signal be routed back in addition to the output voltage. This is done by sensing across a resistance in the load current path (see *Current Balancing*). To minimize I^2R losses, the resistance is kept as small as possible. Therefore, the signal tends to be an order-of-magnitude smaller than the internal ramp generated in voltage mode. Care must be taken to ensure that the signal is not corrupted by external noise sources. In practical terms, peak current mode is quite common today, and it is not difficult to lay out using standard good practices.

Line Regulation

Voltage mode inherently responds more slowly to changes in the input voltage. Response to an input-voltage change must first be reflected in the regulation error of the output voltage, which must be corrected by the voltage-feedback loop. Therefore, response time is limited by the control-loop bandwidth. Most contemporary voltage-mode regulators incorporate circuitry to sense the input-voltage change and apply “feed-forward” by adjusting their ramp slope accordingly. However, this adds to the complexity of the controller. Recalling that the duty cycle in peak current mode is controlled by the inductor-current ramp, which is a function of both input and output voltage, we see that peak current mode provides inherent feed-forward on a cycle-by-cycle basis. Therefore, response to line-voltage changes is quite fast.

Current Balancing

Voltage regulators composed of two or more phases (multiphase) must actively balance the current between phases to prevent one phase from handling a disproportionate amount of current. Per-phase current sensing can be done by monitoring the current through either the high-side or low-side MOSFET(s), or by sensing the current through a current-sense resistor placed in each phase. The MOSFET methods are inexpensive, as they use existing circuit elements, but they are inaccurate because MOSFET resistances vary significantly over process and temperature. The current-sense resistor method can be very accurate, but adds cost and decreases power-supply conversion efficiency.

Another popular method for extracting per-phase current information uses the inductor’s DC resistance (DCR) as the current-sense element. This approach does not add cost, because it uses an existing circuit element and provides reasonable accuracy depending on the DCR tolerance. A series resistor and capacitor are added across the inductor with the RC time constant matched to the L/DCR time constant. The voltage sensed across the capacitor provides a very good DC and AC representation of the current through the inductor. Both voltage-mode and current-mode CPU regulators use this method quite commonly today.

How voltage mode and current mode use the per-phase current information brings us to another tradeoff. Because voltage mode only uses voltage information in the control loop, it cannot control the individual phase currents in each inductor, which is a requirement for current balancing. Peak current mode provides natural current sharing because it uses the inductor’s current information as part of the control scheme. Modern multiphase voltage-mode regulators must add a secondary control loop to provide current balancing, which increases IC complexity and brings with it yet other trade-offs as discussed in *Voltage Positioning and Transient Response*.

While peak current mode provides inherent current sharing, one artifact does impact current-matching accuracy. Because the inductor current peak is controlled, but not the current valley, any mismatch in inductance between two phases (e.g., due to tolerances) will create inductor-current-ripple signals with different peak-to-peak magnitudes. This creates a DC mismatch in the current between the two phases, and therefore impacts the accuracy with which the phase currents are balanced.

Maxim addresses this limitation through rapid active averaging (RA^2), which averages out the inductor ripple current at each phase. The RA^2 circuitry (**Figure 5**) “learns” the peak-to-peak ripple current of each phase over 5 to 10 switching cycles, and then biases down the peak current signal by half of the ripple current. Because the peak control point has been moved from the inductor-current peak to the DC-current point, we still have all of the benefits of peak current-mode control, but with very accurate DC-current matching. As the RA^2 circuitry is not part of the current-loop path used for regulation, it does not slow down the transient response. This technology is used in the MAX8809A/MAX8810A core regulators designed for Intel VRD 10.1 (and next-generation VRD) and AMD K8 Socket M2.

Voltage Positioning and Transient Response

Modern CPUs have large current transients when processor loading suddenly changes. Voltage tolerances must be maintained under these demanding dynamic conditions, otherwise the CPU is prone to lockup. This can be done by ensuring enough bulk capacitance to absorb or supply sudden changes in CPU current; however, this adds to overall cost.

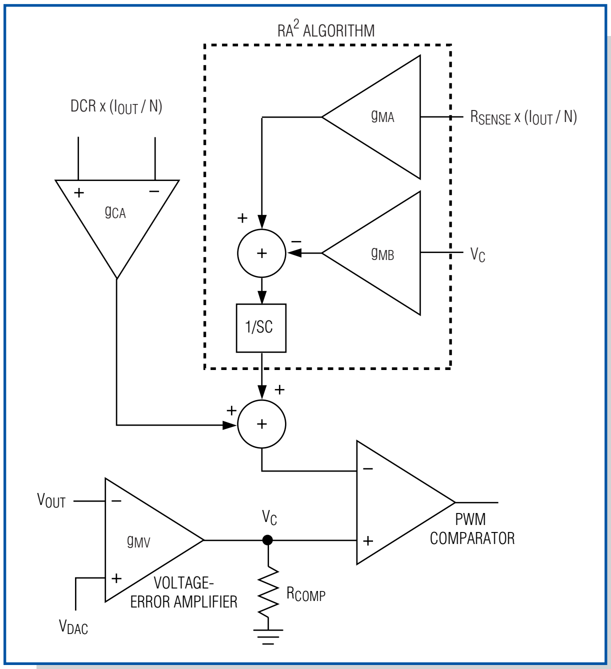

Figure 5. The RA[2] *algorithm is implemented.*

Most high-current CPU-core-regulator designs use a technique referred to as voltage positioning to reduce bulk capacitance requirements. The output voltage is allowed to decrease (droop) as the load current increases, according to a defined slope. The voltage vs. current line is sometimes referred to as a "load-line," and the slope is specified as impedance (e.g., 1mΩ). The benefit is increased voltage margin under dynamic conditions, which reduces the amount of bulk capacitance required for safe operation.

Ignoring voltage positioning considerations, voltage mode does provide a theoretical advantage in terms of voltage-loop response. Theoretical loop bandwidth is a function of (output-voltage) ripple frequency, or per-phase switching frequency multiplied by the number of phases. With peak current mode, voltage-loop bandwidth is a function of per-phase switching frequency only, due to a phenomena referred to as "sampling effect."

There is a practical difference, however, in voltage positioning applications. Remember that voltage-mode control requires a second control loop for current balancing. The loop bandwidth is generally set to 1/5 to 1/10 of the voltage-loop bandwidth to prevent interference with the voltage loop. This is sufficient for current balancing because slow adjustments are usually all that are required. For voltage positioning, however, the ability to respond to a load transient is a direct function of the current-loop bandwidth. For voltage mode, this is quite low (e.g., 5kHz). For peak current mode, the current-loop bandwidth is the same as the voltage-loop bandwidth (e.g., 50kHz to 75kHz) because there is only one loop using both voltage and current information. From the scope shots in **Figures 6** and **7**, it is quite easy to see the difference this makes in transient performance. Both show the response to a 95A load step followed by a 95A load release.

Regulators differ in how they implement voltage positioning. The secondary current loop present in voltage mode usually provides total average current information. This information, usually a scaled version, is forced through a resistor to set up an offset voltage, which is applied to either the reference (desired output) voltage or the actual (feedback) voltage. The resistor value is selected to provide the appropriate load-line impedance.

The MAX8809A/MAX8810A take a different approach, using finite gain to actively set the output load-line (**Figure 8**).

The equation for the error voltage is as follows:

$$V_C = g_{MV} \times R_{COMP} \times (V_{DAC} - V_{OUT})$$

where g_{MV} is the gain of the error amplifier, R_{COMP} is a resistor connected between the output of the error amplifier and ground, V_{DAC} is the desired output voltage, and V_{OUT} is the actual output voltage.

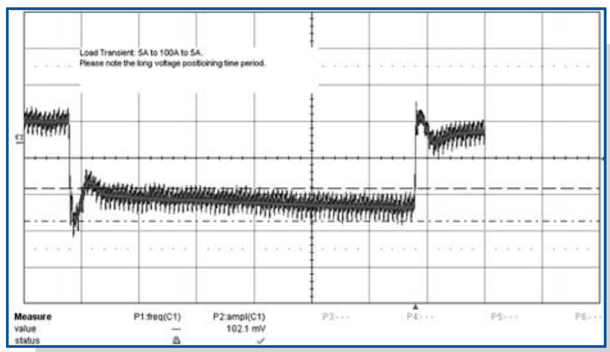

Figure 6. Voltage-mode transient response is shown for a competitive product.

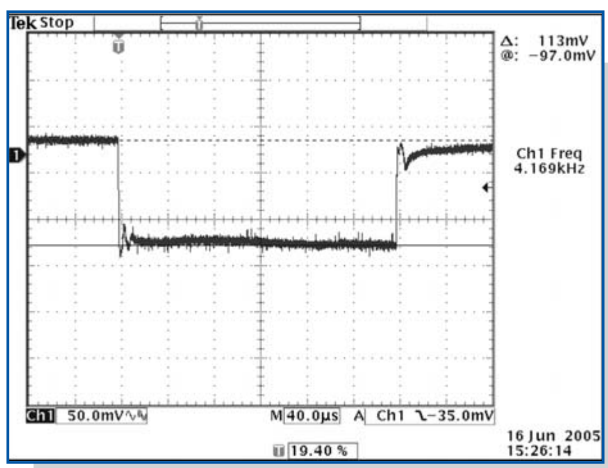

Figure 7. Peak current-mode transient response is shown for the MAX8810A.

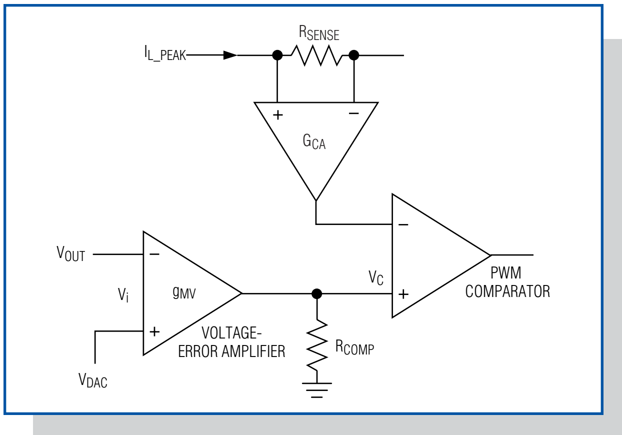

Figure 8. The MAX8810A has peak current-mode control with active voltage positioning.

Likewise the voltage at the inverting input to the PWM comparator is:

$V_C = (I_{OUT} / N) \times R_{SENSE} \times G_{CA}$

where I_{OUT} is the output (CPU) load current, N is the number of phases, R_{SENSE} is the value of the current-sense element, and G_{CA} is the current-sense amplifier gain.

In regulation, these two voltages must be equal. With substitution and some rearrangement, we can solve for:

$(V_{DAC} - V_{OUT}) / I_{OUT} = (R_{SENSE} \times G_{CA}) / (N \times g_{MV} \times R_{COMP})$

The term $(V_{DAC} - V_{OUT}) / I_{OUT}$ is what was previously defined as load-line impedance. The current-sense gain (G_{CA}) and the transconductance of the error amplifier (g_{MV}) are constants fixed by the IC design; the parameters R_{SENSE} and N are determined by the application. Therefore, it is easy to program the load-line impedance simply by selecting the proper value of R_{COMP}, which programs the gain of the voltage-error amplifier.

Loop Compensation

The beauty of the voltage-positioning technique described above for the MAX8809A/MAX8810A lies in its simplicity. The resistor placed at the output of the error amplifier for voltage positioning is also used for loop compensation. Peak current mode only requires single-pole compensation to cancel the zero formed by the bulk capacitors and their ESR. The MAX8809A/MAX8810A only require the addition of a small-value capacitor parallel to the voltage-positioning resistor. The combination of voltage positioning and loop compensation results in fewer error sources that can impact the regulator's output accuracy.

Voltage-mode control is more complex to compensate due to the poles and zeroes formed by the modulator (control loop) and the output filter. Voltage mode usually requires type III compensation, which increases the number of small resistors and capacitors.

Temperature Compensation

One drawback of using inductor DCR for current sensing is that the DCR changes over temperature according to the positive temperature coefficient of copper. This directly affects the accuracy of both voltage positioning and current-limit protection.

To compensate, designs use a resistor with an equal but opposite (negative) temperature coefficient—an NTC. The NTC is generally part of the resistor network that programs the load-line impedance, which ensures that the output voltage vs. current is stable over the operating temperature range. NTCs are not linear over temperature, so the resistor network must include two additional resistors to linearize it over the temperature region of interest.

The drawback to the latter approach is that current-limit information is not temperature-compensated. The current-limit threshold at room temperature must be scaled upward to account for the increased current signal at higher temperatures. At room temperature, the inductor and MOSFETs must be oversized to handle the maximum current at current limit, which in turn leads to a higher solution cost.

The innovative MAX8809A/MAX8810A regulators also use an NTC, but this information is applied independent of the voltage-positioning circuit. Linearization is integrated, saving two resistors. Temperature-corrected current information is then used internally for both voltage positioning and current limiting. Competing products generally require a second NTC to compensate current limit. The MAX8809A/MAX8810A also use the same internal temperature information for a VRHOT function, a signal indicating that the voltage regulator has exceeded a certain temperature. Thus, three temperature functions are accomplished for the price of one temperature-sensing component, which reduces overall costs.

Conclusion

We examined the basics of voltage-regulator control including two popular schemes, voltage mode and peak current mode, used for powering contemporary CPUs. Each approach includes certain tradeoffs that the power-supply designer must consider within the context of a high-current, multiphase design. Features and technology provided by products such as the MAX8809A/MAX8810A core regulators, which implement peak current-mode control with RA^2, simplify the design process and reduce total solution costs. Please visit www.maxim-ic.com/computer-solutions for complete information on more of Maxim's voltage-regulator solutions for desktop-PC and server applications.

To AC-Couple Video or Not to AC-Couple Video?

*Deciding whether or not to design an AC-coupled output for analog video circuits depends as much on company policy and industry standards as technology and cost. An AC-coupled output includes a series capacitor (**Figure 1a**), while a DC-coupled output does not (**Figure 1b**). A designer new to video output circuits can find such a choice confusing because adding a capacitor to the output path increases cost, requires space, and distorts the video signal. However, the choice may have already been made due to historical, technical, or economic reasons.*

AC-Coupled vs. DC-Coupled Outputs

Figure 1a[1] shows the input and output waveforms for an AC-coupled output. Notice how the output waveform “tilts” up and down with respect to the input waveform. Hence, this type of field time distortion is called “field tilt.” The oscilloscope trace in Figure 1b[2] shows a DC-coupled output. In this case, notice that there is no field tilt. The NTSC video test signal used is named “Regulate.” **Figure 2a** shows how the white portion of the test signal appears on a video monitor. **Figure 2b** shows how the black portion of the test signal appears on a video monitor. The Regulate video test signal draws a white border on the edge of the screen during both the white and black portions.

Historical Use of AC-Coupling

Given the drawbacks of AC-coupling, why was it ever used? The simple answer is protection. **Figure 3** shows a simple video-output circuit that might have been created before integrated circuits were widely used. The capacitor prevents the NPN transistor from damaging itself in case the output connector is shorted to ground or a supply voltage.

Contemporary integrated video amplifiers have robust short-circuit protection circuitry so that they are not damaged in the event of a short. Nonetheless, use of a capacitor has become entrenched in some companies, especially those with a long history of making video equipment. Design engineers may be told that they must add a capacitor to comply with company policy.

In addition, industry standards can implicitly force the design engineer to use a capacitor. The Japan Electronics Industry Trade Association (JEITA) has a specification that requires the voltage magnitude to be less than 100mV on an inactive video output connector (Figure 3). If the normal DC bias at the NPN emitter is 4V, then the output connector would also be approximately 4V if the capacitor and bleed resistor were not present. The easiest way to meet the JEITA specification, therefore, is to add a capacitor and bleed resistor to ground.

Technical Concerns

One of the concerns about AC-coupling is that the capacitor is usually large—220µF or higher. This is because the frequency of the pole formed by the capacitor and 150Ω load (the total resistance of the back-termination resistor and the input-termination resistor) should be significantly less than the frame rate of either 25Hz or 30Hz. A 220µF capacitor forms a pole at 5Hz, which is barely adequate. Broadcast equipment typically has capacitors in the 2200µF range. **Figure 4** shows the highpass response of an AC-coupled output connection with a 220µF capacitor.

With the advent of small, portable devices that have video outputs, using a large, AC-coupling capacitor is prohibitive, primarily for space and cost reasons. SAG compensation (**Figure 5**) reduces space and cost while maintaining AC-coupling. The single, large capacitor of the standard connection is replaced by two smaller capacitors. The problem with a single AC-coupling capacitor is that the signal is attenuated at frequencies below the pole frequency—the smaller the capacitor, the higher the pole frequency. SAG compensation boosts the low-frequency response to compensate for the low-frequency attenuation (Figure 5). At low frequencies, the capacitors can be treated as opens; the low-frequency gain is approximately 6. At high frequencies, the capacitors are essentially shorts, and the high-frequency gain is 2.

In the consumer electronics industry, the pressure to reduce cost is intense and, for small devices, the pressure to miniaturize is just as great. Companies, even some with a long history of designing video equipment, are now opting for DC-coupled video connections (Figure 1b). The main difference to note with a DC-coupled output is that the signal has a positive DC bias because most systems have eliminated negative supplies. For the amplifier to remain in linear mode, the output signal must be biased between ground and the positive supply.

[1] The 0.1µF capacitor across the 75Ω input-termination resistor to ground filters out higher frequencies in the video waveform, removing aliasing in the black portion of the video test signal. The 400ms time scale of the oscilloscope shot is very long compared to a horizontal line time (~64µs). Without the 0.1µF capacitor, the aliasing during the black portion of the video test signal would make the black portion of the input signal nearly indistinguishable from the white portion. The black portion of the output signal would be filled in like the white portion.

[2] Just as in Figure 1a, there is a 0.1µF capacitor across the 75Ω input-termination resistor to ground for the same reasons. Without the 0.1µF capacitor, the aliasing during the black portion of the video test signal would make the black portion look the same as the white portion.

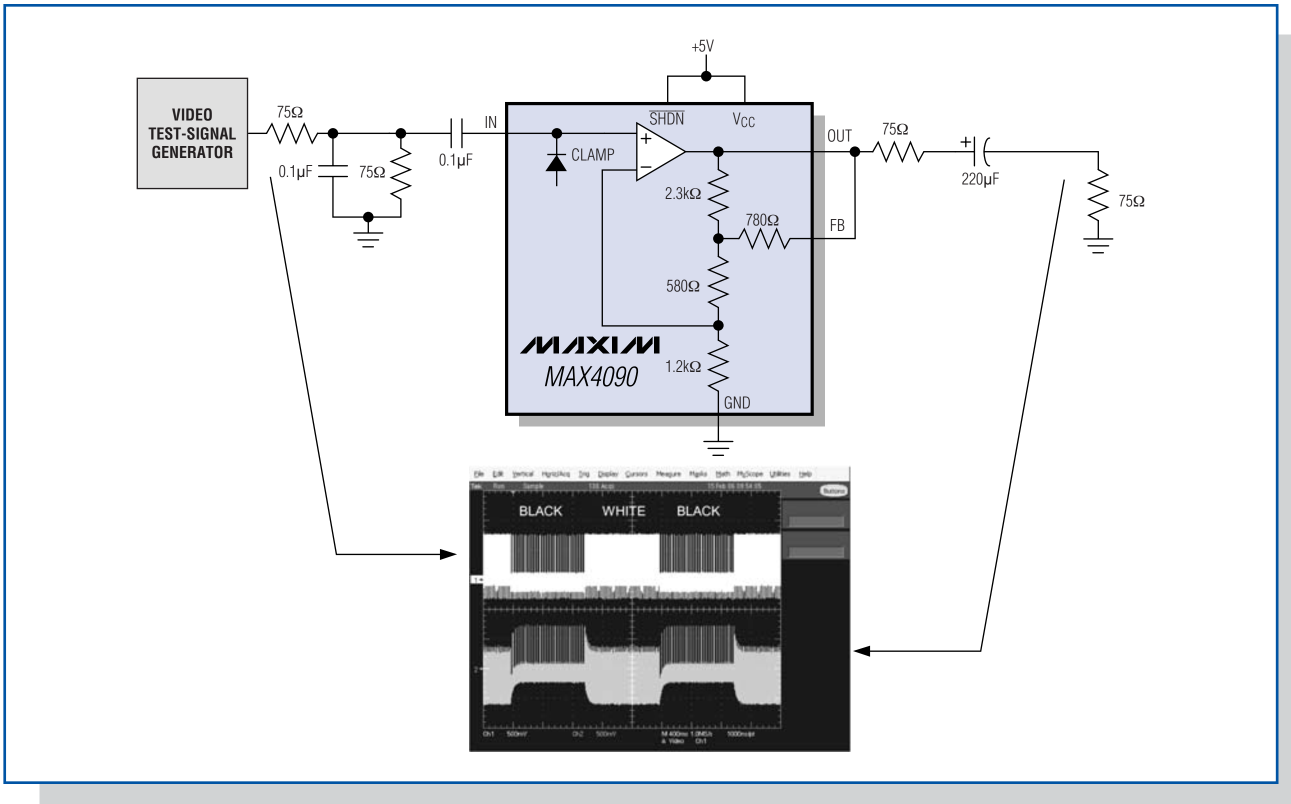

Figure 1a. This AC-coupled output connection includes a series capacitor.

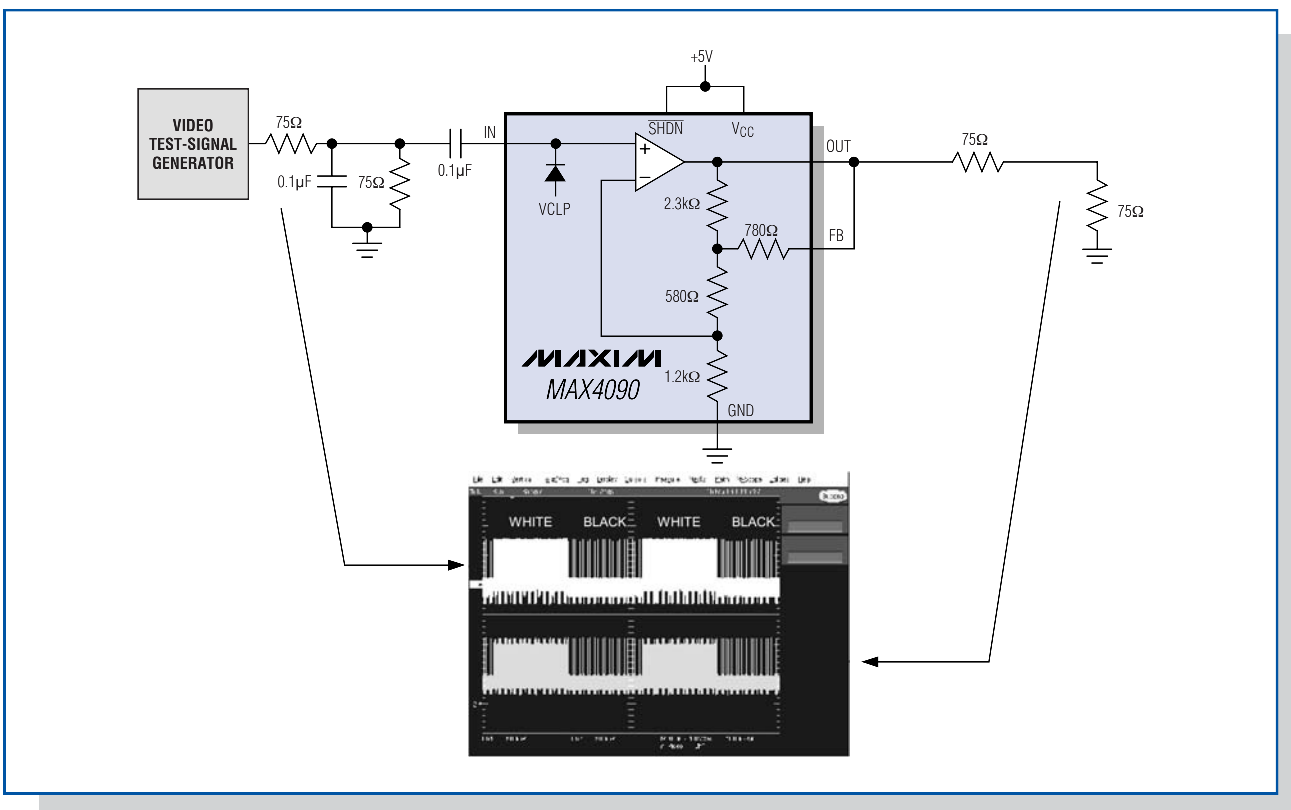

Figure 1b. This DC-coupled output connection does not need a capacitor.

Figure 2a. The white screen of Regulate video test signal is shown.

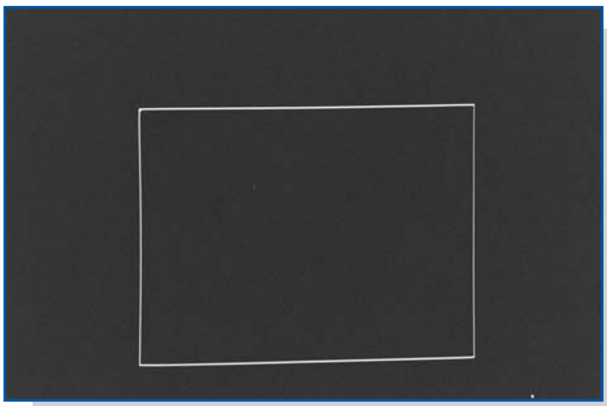

Figure 2b. As in the white screen, the black screen has a white border from the Regulate video test signal.

Design engineers contemplating a change from an AC-coupled output to a DC-coupled output are concerned with compatibility. Most equipment is compatible with either style, but there is still a small percentage that is incompatible with an AC-coupled output, and yet another tiny portion that is incompatible with DC-coupled outputs. **Figure 6a** shows the input stage for a modern television. The video signal is AC-coupled into a DC restoration circuit; hence, the input video signal can have any DC bias. This input circuit is compatible with both AC- and DC-coupled video sources. **Figure 6b** shows the input stage that uses a polarized capacitor. If the DC bias of the incoming video signal is too high, then the polarized capacitor could be damaged. This Figure 6b input stage might have problems with a signal from a DC-coupled source. **Figure 6c** shows an input stage that uses a PNP emitter-follower. If the input signal is too negative, then the PNP emitter-follower could saturate. Therefore, a DC-coupled video source could saturate the PNP emitter-follower, especially if the local ground of the source is lower than the local ground of the receiver.

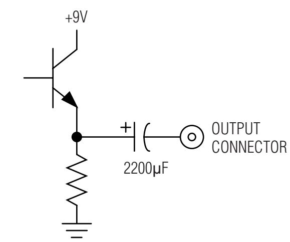

Figure 3. In a simple video-output circuit, the NPN emitter follower drives the video output.

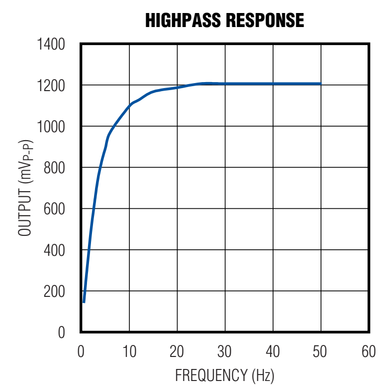

Figure 4. Frequency response is shown for an AC-coupled video connection with 220µF capacitor.

The problem with televisions is that there has never been a commonly accepted way of designing the input stage. Historically, numerous television models have marginal input stages that could have problems with either AC- or DC-coupled sources. It is not possible to maintain universal compatibility when there are so many different kinds of receiving equipment. Most low-end video sources, which represent the greatest volume of transmitting equipment, use DC-coupled outputs for cost reasons.

DirectDrive™ Solution

For design engineers who want to include an AC-coupled video output, Maxim offers DirectDrive technology, which eliminates the need for large output-coupling capacitors. The MAX9503 is the first Maxim part to incorporate DirectDrive technology for video signals (**Figure 7**).

The MAX9503 filters and amplifies standard-definition video signals. The input of the MAX9503 can be directly connected to the output of a video digital-to-analog converter (DAC). An internal reconstruction filter smoothes the steps and reduces the spikes on the video signal from the DAC. The MAX9503 level shifts the video signal to a lower voltage such that the blank level is approximately at ground at the output. DirectDrive requires an integrated charge pump and a linear regulator to create a clean negative power supply to drive the sync pulse below ground. The charge pump injects such little noise into the video output that the picture is seemingly flawless.

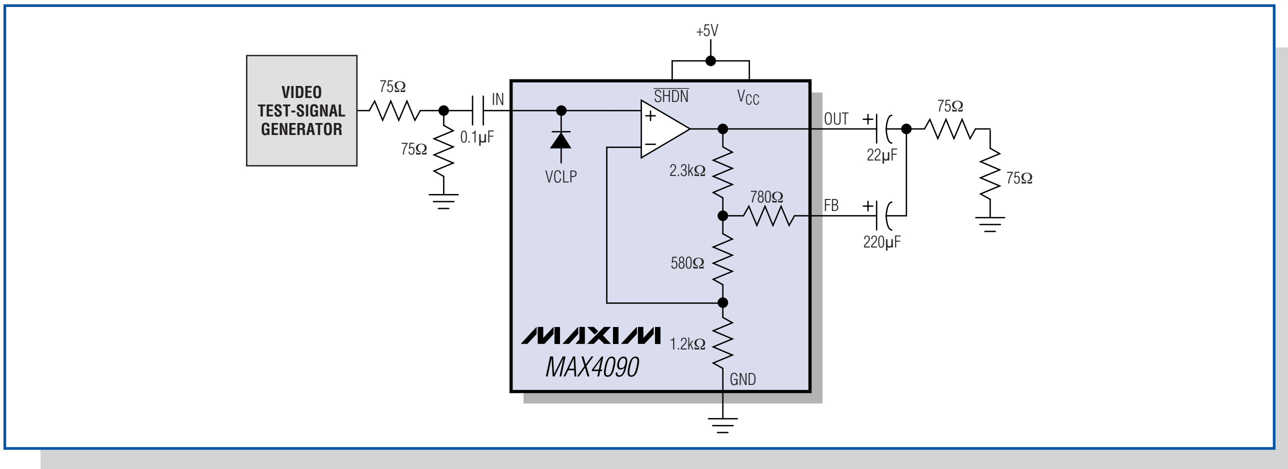

Figure 5. The MAX4090 video driver includes SAG compensation.

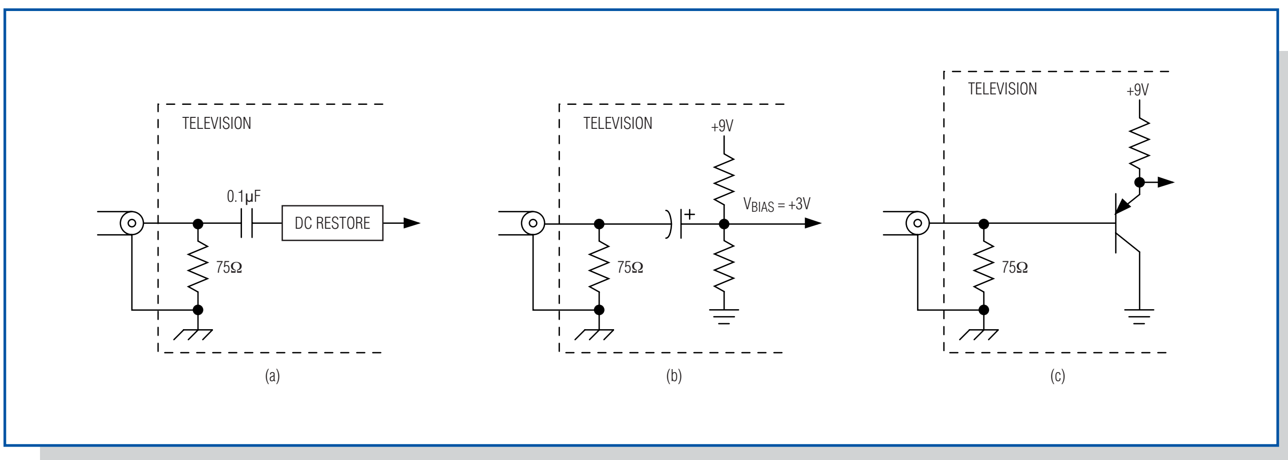

Figure 6. Three types of input stages are shown: a) modern TV, b) polarized capacitor, and c) PNP emitter-follower.

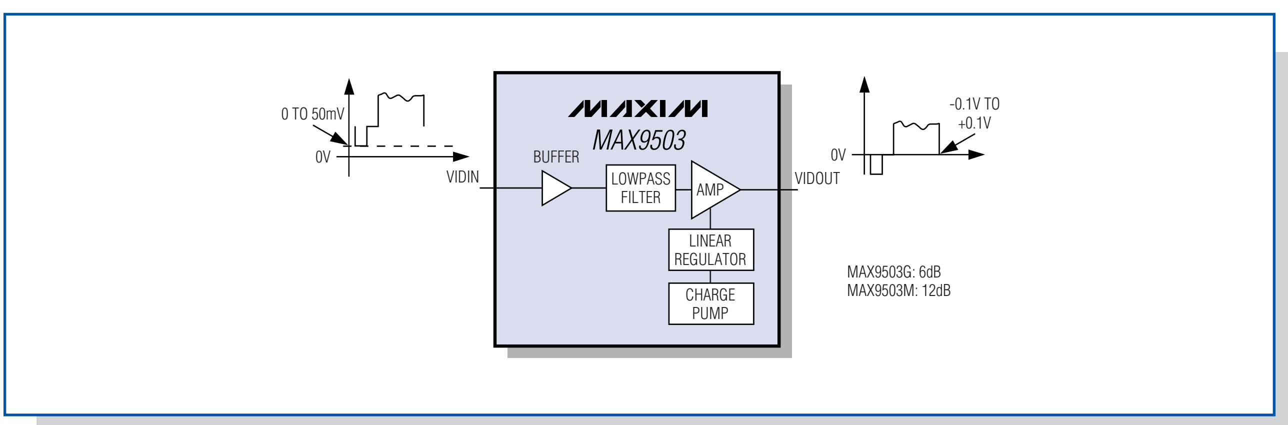

Figure 7. The MAX9503 block diagram is shown with its representative input and output waveforms.

Figure 8 shows a Regulate video test signal applied to the MAX9503. Notice how the blank level of the output waveform is held near ground and that there is no field time distortion. This contrasts with the normal AC-coupled waveform shown in Figure 1a. The Regulate test signal at the MAX9503 output maintains a much more well-defined output voltage range than that from an AC-coupled connection.

One reason to use an AC-coupled video output is for protection against shorts to ground and the supply voltage. The MAX9503 typically operates from a 3.3V supply. The MAX9503 application circuit includes a 75Ω back-termination resistor that limits short-circuit current if an external short is applied to the video output. In addition, the MAX9503 features internal output short-circuit protection to prevent device damage in prototyping and applications where the amplifier output can be directly shorted. Hence, the MAX9503 is robust in the face of most common fault conditions.

The major benefit of DirectDrive is that for the addition of just two small, 1µF capacitors to the charge-pump circuit, the design engineer can eliminate the single, large, output-coupling capacitor in a standard AC-coupled video output or two medium-sized output coupling capacitors in a SAG network. Output video quality is improved because field-time distortion is eliminated.

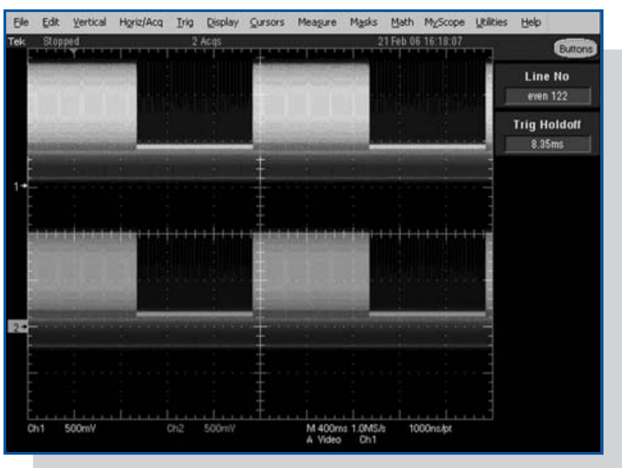

Figure 8. A Regulate video test signal is applied to the MAX9503. The input waveform is on top and the output waveform is on the bottom.

Conclusion

This article explains the historic, technical, and economic reasons for choosing either an AC- or a DC-coupled video output. Maxim's DirectDrive video technology, which has the benefits of an AC-coupled connection with the simplicity of a DC-coupled connection, is introduced. With this information, the design engineer should be better able to decide which style of video output to use in future projects.

EMI/EMC Suppression in Audio/Video Interfaces

All electronic products marketed worldwide undergo EMI/EMC testing before they are offered for sale to prove that they will not create interference, or be interfered with by other devices. For testing purposes, products are grouped into two classes: intentional radiators and unintentional radiators. For example, cell phones and walkie-talkies intentionally radiate energy while TVs, PCs, or laptops should not.

Depending on the class of the product and the agencies involved, the EMI/EMC requirements vary. Regardless, EMI/EMC testing is partitioned into two general categories:

- **Emissions** This limits what amplitude and frequency a product can radiate or conduct so that the device does not interfere with others.
- **Susceptibility (also called Immunity)** This testing category compliments the Emissions requirement by limiting the amplitude and frequency of the radiated and conducted signals that can interfere with a product.

As noted above, the two ways EMI can be emitted by a device are conduction and radiation. These are related because all radiated EMI is due to current flow. However, not all current flow causes radiation. Consequently, radiated interference problems are investigated and suppressed before conducted interference issues. Of the two, radiation is more difficult to predict and suppress. Therefore, it causes most unanticipated EMI testing failures in the unintentional radiator product class. We will concentrate on solving the radiation problem in the audio/video interface found in many products.

You can use several methods to meet the limits specified in EMI/EMC regulations. However, these methods tend to fall into the general categories of shielding and filtering. In actual practice, these are combined with application-specific methods to achieve an overall EMI solution. For example, in most products, a metal chassis acts as a shield to radiation and L-C or R-C filters reduce conducted interference on the input/output wires. In addition, we might dither a clock to spread its spectrum and reduce the amount of filtering or shielding required in a specific application.

When the performance appears adequate, a product is formally tested in an approved laboratory. If it passes, it may be marketed; however, failing is a problem. Even a small change to correct a problem can take time to implement. This may delay introduction because passing EMI/EMC-compliance testing is mandatory in all international and domestic markets.[1] Consequently, the EMI design often degrades video performance to ensure that a product passes these tests. This video degradation is exacerbated in modern designs by the size and cost of the parts required to pass EMI testing.

This design challenge is particularly true for the contemporary audio/video analog interface with its reduced product size and high performance expectations. The first step in solving this problem is to find where most EMI/EMC failures occur and then investigate potential solutions.

Where Failures Happen

EMI/EMC failures occur at the weakest point in a product's design—in this case, where a signal (and interference) enters or leaves a shielded and filtered structure. Within the audio/video interface, these are the cables that connect devices and act as an antenna. Specifically, the cabling that attaches the display and speakers to a PC are particularly vulnerable and frequently cause EMI/EMC problems. Although we might expect this with video because of the large bandwidth involved, audio is of low frequency and would seem to be benign. This was the case when all audio amplifiers were Class A. However, the high-efficiency Class D amplifiers[2] used today have high-frequency switching signals and can have EMI problems if they are not filtered and shielded properly.

In the past, using large external filters and/or shielded cables were the accepted ways to address these problems. These methods added cost, adversely affected performance, and increased product size. As these products have gotten smaller and evolved into today's audio/video players, EMI/EMC solutions have had to decrease in size while maintaining or even improving performance. To achieve this, small devices like the MAX9511 graphics video interface and MAX9705 Class D audio amplifier have been developed to deliver excellent EMI performance. To show how this was done, we will look at the audio and display interfaces[3] of a typical PC, and the EMI performance attainable with these devices despite their diminutive size. However, first we need to understand the different EMI problems that must be solved for an audio/video interface design and the methods available to solve them.

1 FCC Part 15 in the US, VCCI in Japan, and EN55000 in Europe all require EMI/EMC compliance.
2 Application note 1760: "Class D Audio Amplifiers Save Battery Life"; www.maxim-ic.com/an1760
3 "EMI and Emissions," Gerke and Kimmel Associates.

Video and EMI

All computers use a form of video, referred to as "graphics," that is different from that for a TV.[4] Computer video has red, green, and blue (R, G, B) analog video signals and separate logic signals consisting of horizontal and vertical sync and DDC,[5] all of which have fast rise/fall times. The video connector is typically a high-density, D-subminiature type that connects the display to a PC (**Figure 1**). Although this incorporates shielding (coax) of the video signals and a common-mode choke (CMC) to reduce radiated and conducted EMI, additional filtering is needed to ensure that the EMI requirements are met. In broadcast video applications, similar filtering is used to remove aliasing artifacts from the video in a TV. However, this is not done in graphics video, where instead the goal is to reproduce a checkerboard pattern of "on" and "off" pixels at the highest resolution possible. So, for the best display performance, we want as large a bandwidth as possible. However, in reality, EMI and video performance are traded off and video bandwidth suffers for it. This tradeoff happens for several reasons unique to multisignal video interfaces.

For example, when you filter video signals, you introduce a time delay that causes problems like "fringing" at the edge of an image if the individual video channels (R, G, and B) are not closely matched in their timing. To avoid this, the group-delay and group-delay-matching channels[6] must be closely controlled. RGB video is particularly susceptible to both of these parameters.[7] For best performance, the group delay must be constant with frequency, and a minimum group-delay matching of ±0.5 pixel time must be maintained between the channels. Assuming they are matched this closely, the sync signals must also track the channel delay to correctly frame the image. If this is done, we then need to address the issue regarding the multiple video resolutions of which a PC is capable.

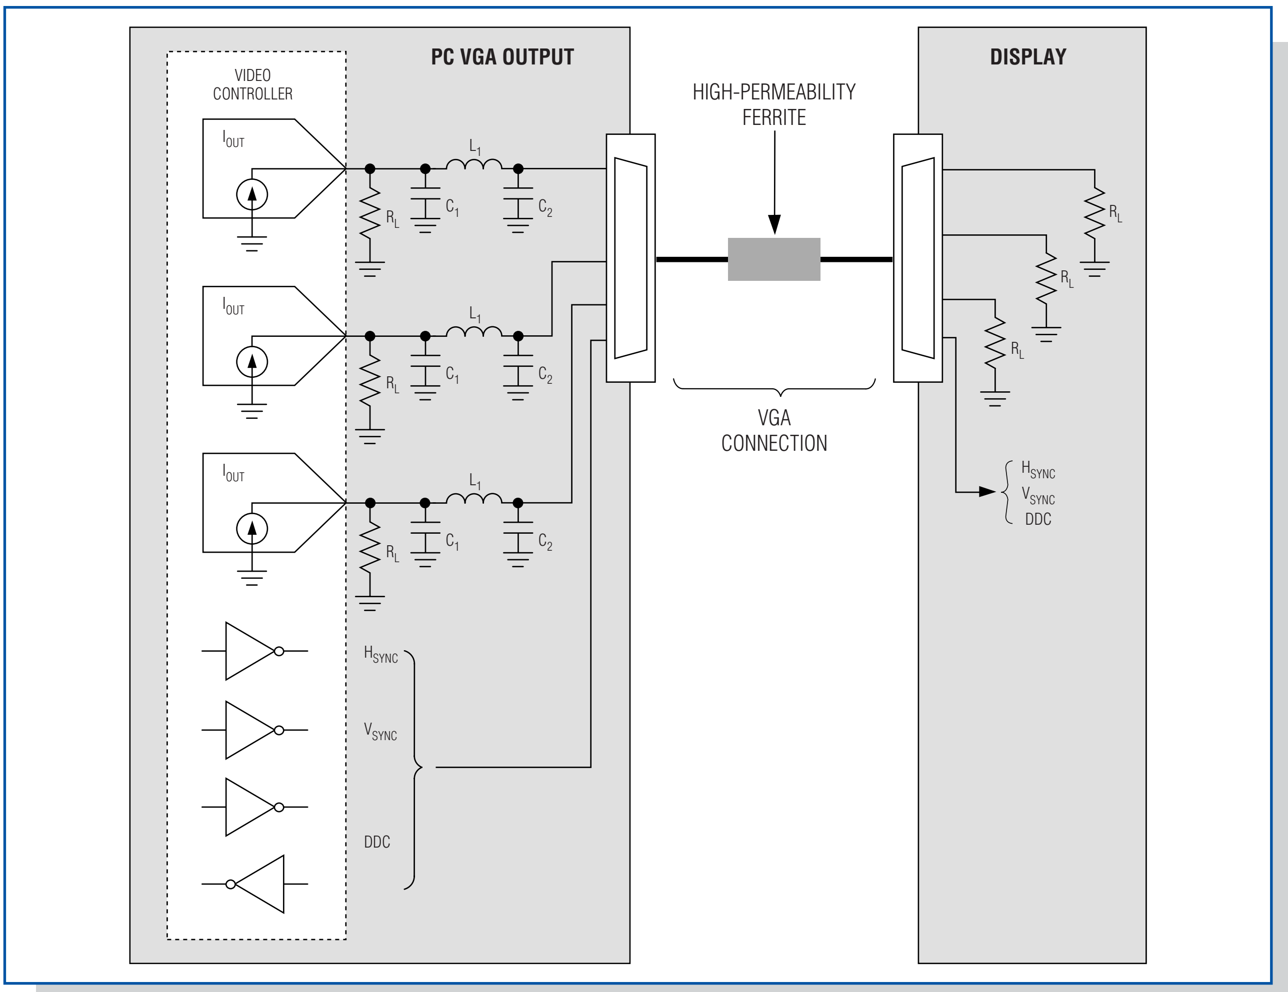

Figure 1. A typical VGA connection is shown with video signals that cause radiated EMI.

4 Application note 1184: "Understanding Analog Video Signals," www.maxim-ic.com/an1184
5 VESA Standard VSIS, v. 1, rev. 2, 12/12/2002.
6 "Design of Analog Filters," Schaumann and Van Valkenburg, Oxford University Press.
7 "Flat Panel Displays and CRT," L.E. Tannas, editor, Van Nostrand.

Optimizing performance is very difficult with a fixed-frequency filter in this application. If we design a filter to suppress EMI from the lowest resolution, the filter's stopband may intrude into the signal bandwidth of higher resolution formats, thus compromising their performance. Design for the highest resolution and you may not meet the EMI requirements. Clearly the best solution is a "tunable" filter with a frequency response that tracks the display resolution being used, but this approach increases cost and possibly size. Secondary, but still important to EMI performance, are the fast rise/fall times of the sync and DDC drivers. A designer must include a way to filter these rise/fall times in any complete EMI solution. There are also legacy issues like video DAC load detection to satisfy plug-and-play requirements.

The MAX9511[8] performs all of these functions. **Figure 2** shows typical before and after performance of a high-resolution graphics board output with a MAX9511 compared to an L-C filter and the raw output.

Complete EMI Solution (MAX9511)

The MAX9511 graphics video interface shown in **Figure 3** provides a matched, triple-channel, tunable EMI filter for the RGB video over the range of VGA to UXGA resolutions with channel-to-channel skew error of < 0.5ns. Tuning is accomplished by changing a single resistor (R_X). **Table 1** shows the resistor value vs. the slew rate for different VESA resolutions and their sampling clock range. In the **Figure 4** example, the MAX5432[9] I^2C-controlled potentiometer provides 32 individual steps of filter control. However, as can be seen in Table 1, only three or four levels are required in most applications. This allows the modification of a product's EMI profile without making any mechanical or electrical changes during final EMI/EMC testing.

The RGB video outputs are of low impedance (Z_{OUT} < 1Ω) and, with the 75Ω back terminations, provide 45dB to 50dB of isolation between the remote monitor and docking station. Previously, driving two different outputs this way required a switch to avoid having a long, unterminated stub attached to the L-C filters' output. Figure 4 shows how the output load is detected and reflected to the input as an apparent change in the DAC termination impedance. The video controller driving the RGB inputs can sense this and, if no load is present, shut off the video and sync outputs through the shutdown pin. The DDC is always on to support plug-and-play applications, and the drivers feature voltage-level translation to convert low-voltage controller levels to standard 5V interface levels. The sync drivers have 50Ω (typ) of output impedance that can be configured to filter the edges with a single external capacitor (Figure 4). The sync jitter (without any capacitors) is typically less than 0.5ns. Video performance includes +6dB of gain with a SNR of 50dB, a linearity error of 0.036%, and < 1% overshoot/undershoot with a well-damped response.

Audio and EMI

Audio has a different set of problems for achieving efficiency and performance without causing EMI. In portable applications, we want to maximize battery life and we do not want heat caused by an inefficient design, so Class D audio amplifiers are widely used. The problem with Class D amplifiers is that they use PWM to attain high efficiency, much like a switching power supply. Attaching unshielded speaker wires to the output causes them to act as an antenna and radiate EMI. Although the clock frequency is above the audio spectrum, typically 300kHz to 1MHz, it is a square wave with sizeable harmonic content. A filter that can remove this harmonic content is large and expensive. In portable applications like laptops, this is not an option due to size alone.[10]

Table 1. Slew Rate, Bandwidth, and R_X Value for the MAX9511

R_X (kΩ)	MAX9511 Slew Rate vs. R_X		
	Slew Rate (V/ms)	Pixel Clock Frequency (MHz)	VESA Resolution
7	1408	160 to 230	UXGA (1600 x 1200)
10	1255	160 to 230	UXGA (1600 x 1200)
12	1050	100 to 150	SXGA (1280 x 1024)
15	810	100 to 150	SXGA (1280 x 1024)
20	613	45 to 95	XGA (1024 x 768)
25	470	45 to 95	XGA (1024 x 768)
30	368	45 to 95	XGA (1024 x 768)
35	298	35 to 50	SVGA (800 x 600)
40	255	35 to 50	SVGA (800 x 600)
45	203	35 to 50	SVGA (800 x 600)
50	158	25 to 30	VGA (640 x 480)
> 50	< 150	< 25	QCIF

8 MAX9511 data sheet; www.maxim-ic.com/MAX9511
9 MAX5432 data sheet; www.maxim-ic.com/MAX5432

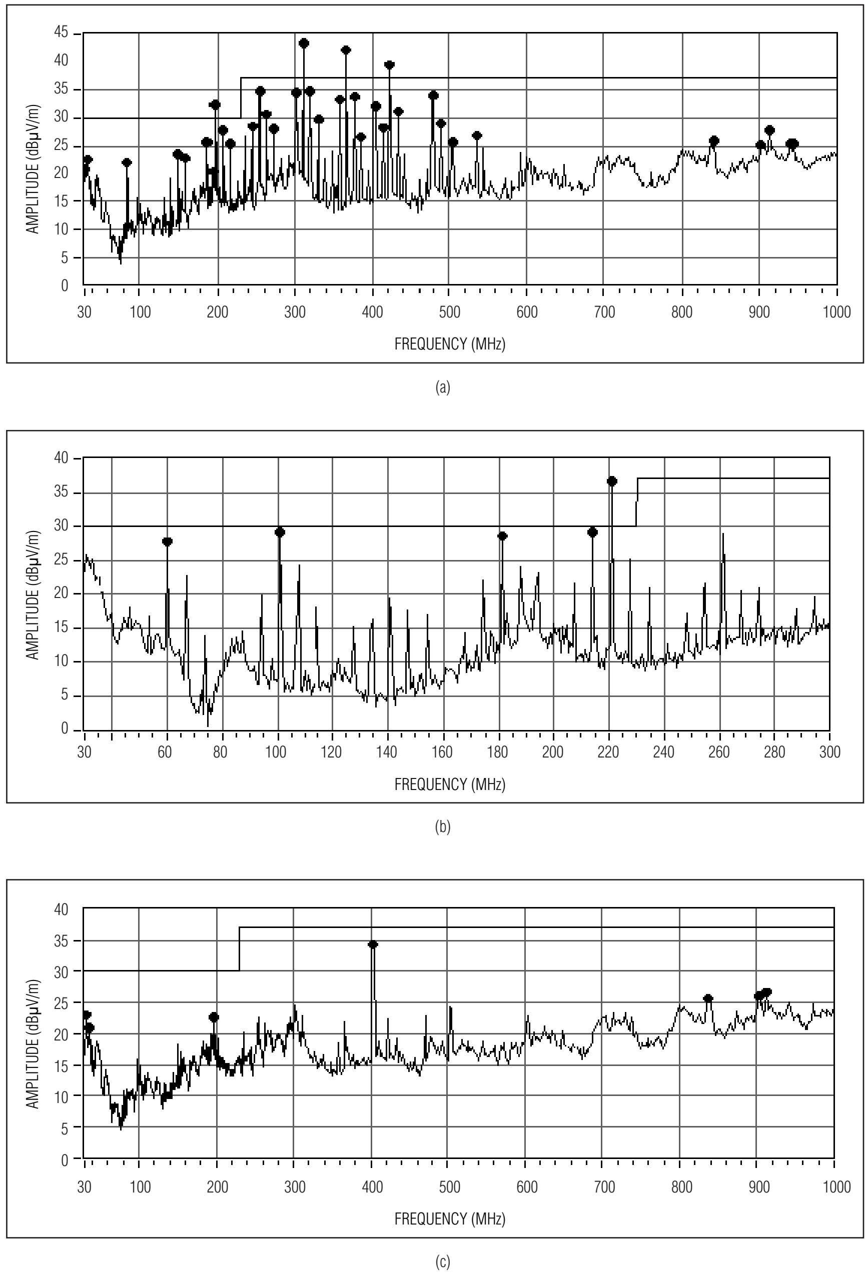

Figure 2. Radiated EMI is shown a) without filtering, b) with a passive LC filter, and c) with the MAX9511.

The typical design topology does not help either. To maximize the output audio power, portable applications use an output connection referred to as a bridge-tied load (BTL) in which both speaker wires are actively driven (**Figure 5**). With Class D, a comparator monitors the analog input voltage and compares it to a triangular clock waveform. The comparator trips when the input magnitude of the triangular waveform exceeds the audio input voltage, and an inverter generates the complementary PWM waveform to drive the other side of the BTL output stage. Because of this BTL topology, the output filter actually requires twice as many parts as a single-ended audio output: two inductors (L1 and L2) and two capacitors (C1 and C2). Because the inductors need to handle the peak output current, they are large and take most of the space.

Class D amplifiers can be made to appear to run without filters by using the speaker voice-coil inductance and discrete capacitors to implement a filter. However, you are limited to internal speakers as the wires still radiate a

[10] Application note 1760: "Class D Audio Amplifiers Save Battery Life," www.maxim-ic.com/an1760

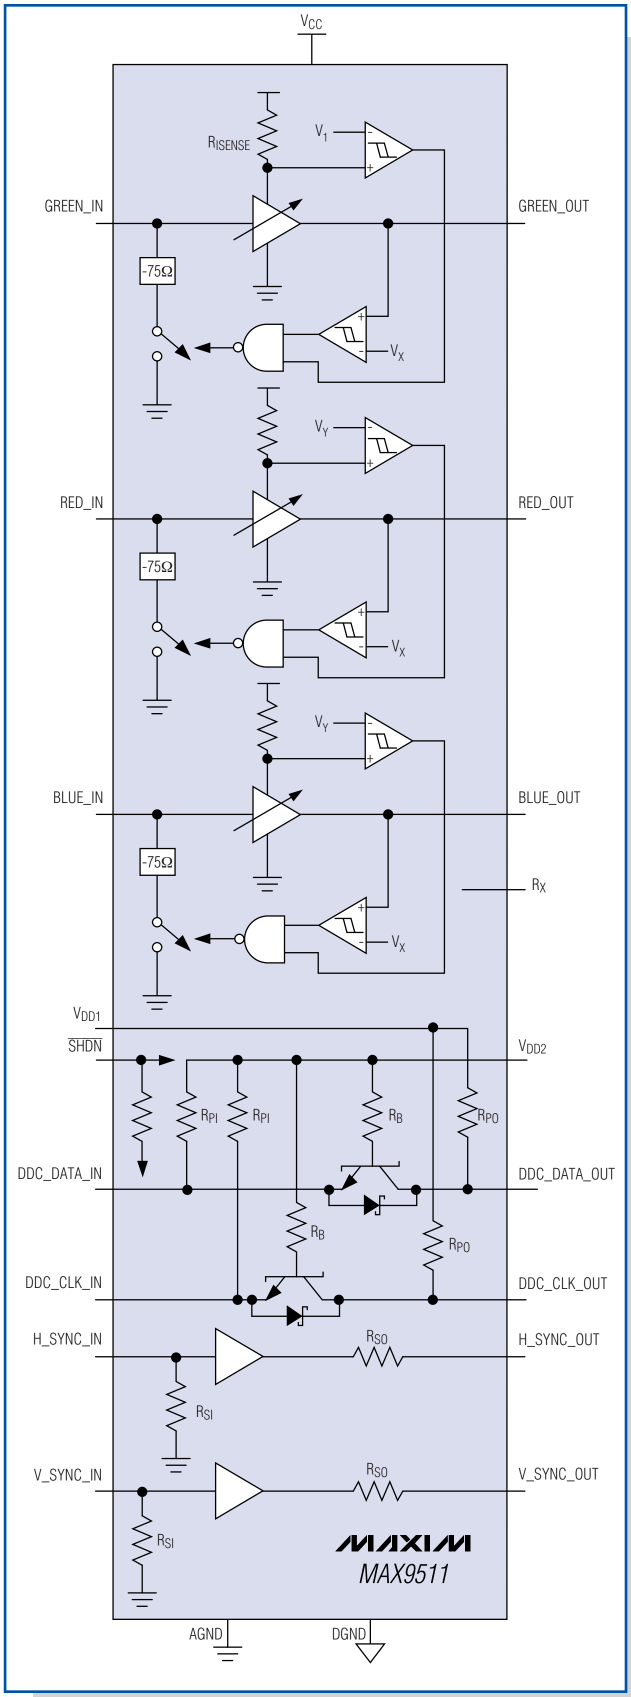

Figure 3. The MAX9511 VGA interface has EMI suppression.

considerable amount of energy. An alternative is to modify the switching process so the amplifier remains highly efficient, but exhibits less EMI and therefore requires a smaller filter. One way to do this is to modulate the frequency of the clock to reduce the energy on a per-hertz basis.[11] This is called spread-spectrum modulation,[12] or dithering of the clock frequency. However, the spectrum can be spread only so far before the returns begin to diminish. **Figure 6** shows the effect of this technique on a typical emissions profile.

In parts with spread-spectrum modulation alone, speaker wires longer than a few inches radiate too much energy above a few hundred milliwatts of output power. Increasing the clock frequency does not help because, as the frequency increases, the Class D amplifiers' output spectra drop off. However, the speaker wires become more efficient as antennas, effectively canceling any improvement in performance. To further improve the EMI requires modifying the PWM waveforms used in the Class D amplifier itself. This is done using an application-specific method called active emissions limiting.

Active emissions limiting circuitry sets the minimum pulse width in the amplifier, which is not bounded in designs like that of Figure 5. Together with control of the overlap, rise/fall time, and clock frequency, this bounds the power spectra that the process causes[13] for a given output power level. The objective is to reduce the spectra to a level where the part can operate without any external filtering with up to 24in of external speaker wire and still meet the radiated emissions limits.

We also want audio performance, and to do that we need > 2W peak power out. At the same time, we want to minimize heat and maximize battery life. Therefore, we want high efficiency while operating from a low-voltage single supply and a low-power shutdown mode for headphone use. The THD+N must be low, the SNR should be high with click-and-pop suppression, and the input must be compatible with either single-ended or differential inputs. The MAX9705 can accomplish these tasks and more, as we shall see next.

Active Emissions Limiting (MAX9705)

The technique used for active emissions limiting in Maxim's Class D amplifiers is shown schematically in **Figure 7**. What is not obvious from this figure is how the switching is done. By carefully tailoring the

[11] Application note 3503: "Clock Generation with Spread Spectrum"; www.maxim-ic.com/an3503

[12] "Spread Spectrum Systems," R.C. Dixon., J. Wiley & Sons, 1976.

[13] "High-Speed Digital Design," Graham and Johnson, Prentiss Hall, 1993.

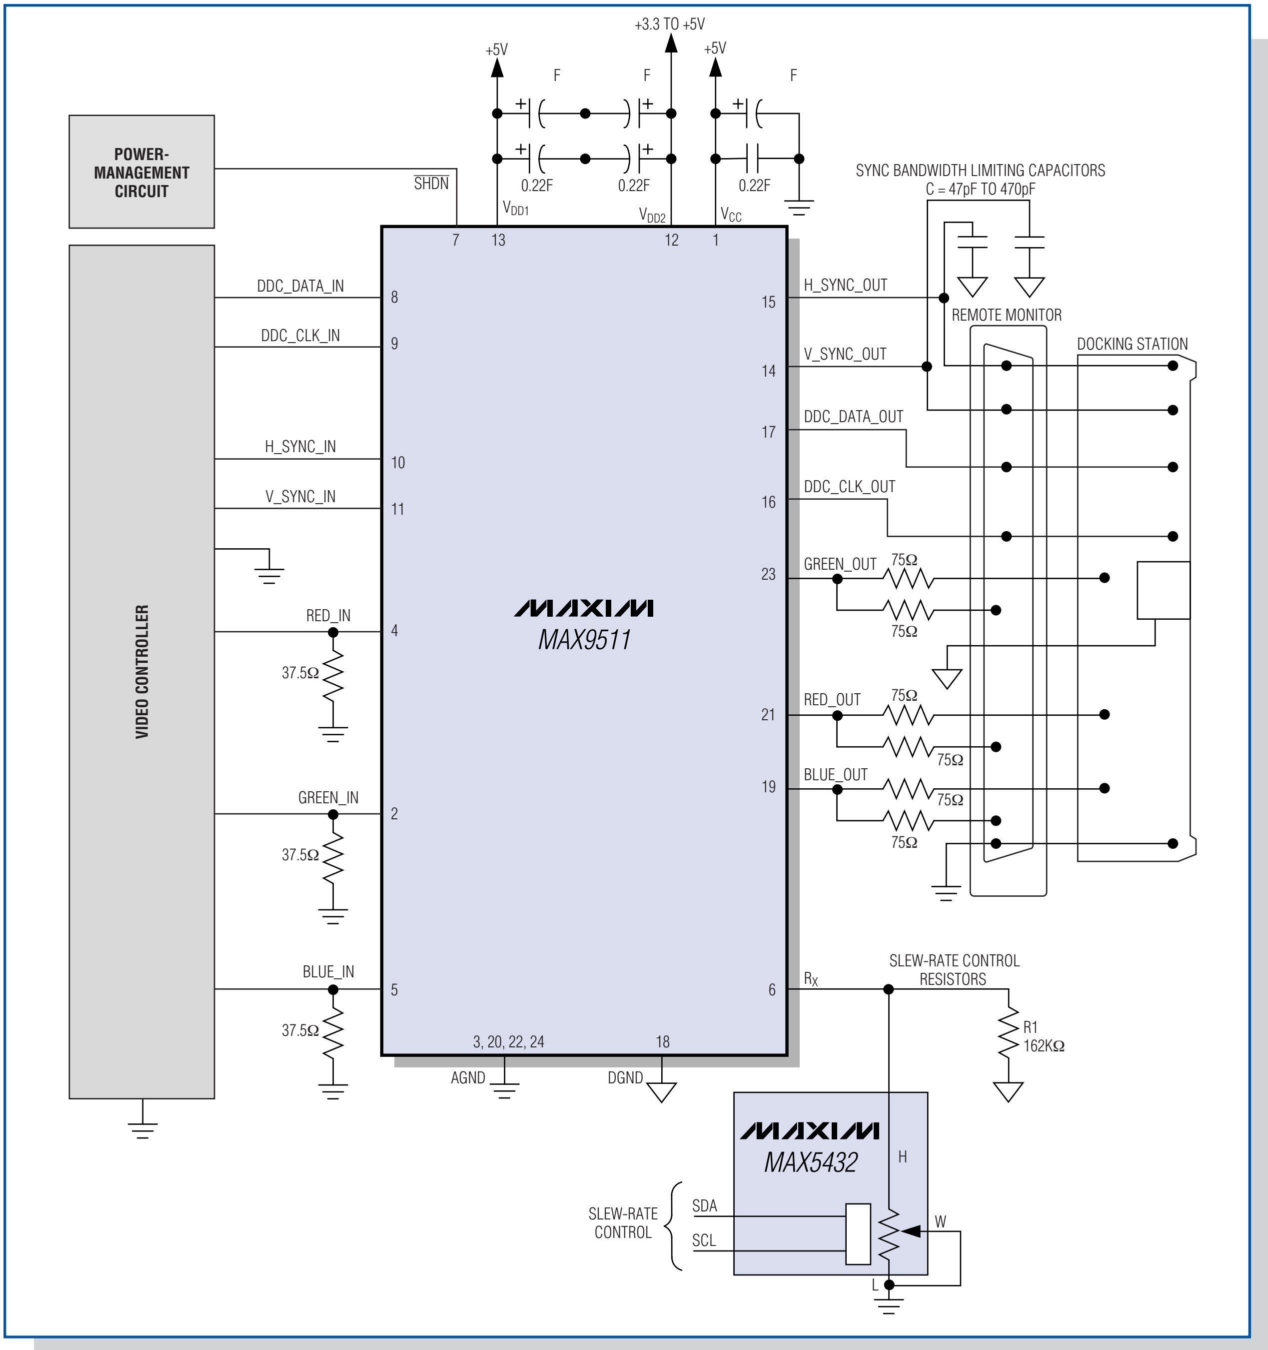

Figure 4. The MAX9511 drives multiple outputs. Adjustable filtering is controlled by the MAX5432 I^2C-adjustable digital potentiometer.

drive and using zero dead-time control, the efficiency of the MAX9705 Class D amplifier is > 85%. The unique, spread-spectrum modulation mode flattens the spectral components, reducing the EMI emissions radiated by the cable and speaker. For stereo or multichannel operation, a sync input locks the amplifiers to a common clock range of 800kHz to 2MHz to minimize intermodulation products that would otherwise occur with multiple free-running sources. The combination of these two unique technologies, spread-spectrum modulation and active emissions limiting, allows Maxim's Class D audio amplifiers to run "filterless" with up to 24in of unshielded speaker wire before exceeding the EMI limits of FCC Part 15 (**Figure 8**).

Besides EMI, the audio performance is excellent with a THD+N of 0.02% at 1W, increasing to 1% at 2.3W, and an SNR of 90dB. The input can be differential or

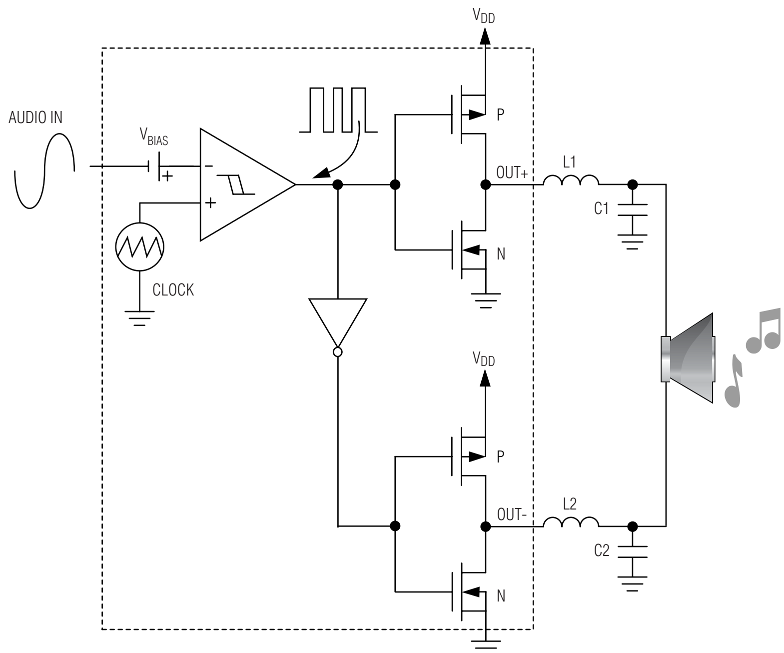

Figure 5. The active-emissions limiting technique is shown in a typical Maxim Class D audio amplifier.

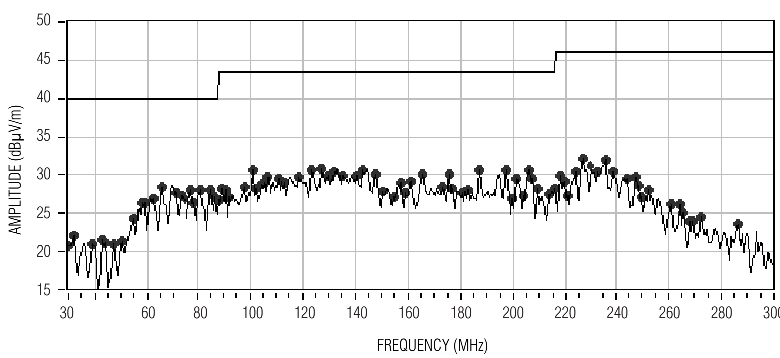

Figure 6. The MAX9705 radiated emissions data, obtained using a MAX9705EVKIT (12-inch, unshielded twisted pair), shows the effect of spread-spectrum modulation.

single ended with a fixed gain of +6dB, +12dB, +15.6dB, or +20dB available to address any application (Figure 7). A shutdown minimizes power. Also, a sync input allows the MAX9705 to provide monaural, stereo, or multichannel high-performance audio and still meet EMI radiation requirements with external speakers, but without filters.

Conclusion

The MAX9511 and the MAX9705 represent a modern approach to EMI/EMC control. They actively reduce EMI in the products in which they are used. Instead of relying on large external filters and shields with inherent increases in cost and size, as done in the past, they use state-of-the-art technology to virtually guarantee electromagnetic compatibility and performance.

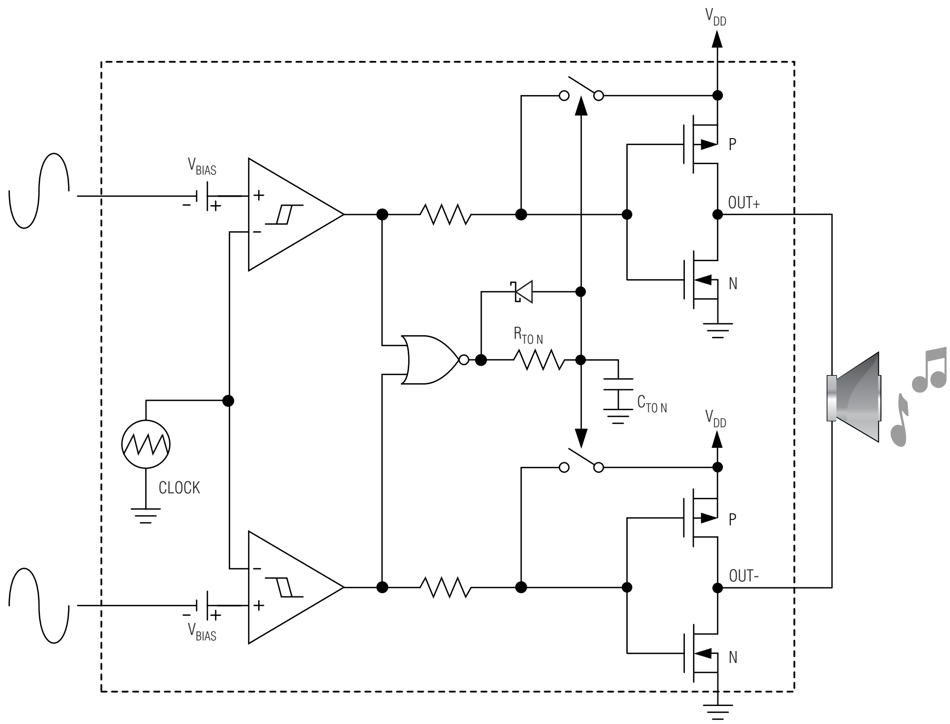

Figure 7. The MAX9705 Class D amplifier has an internally generated sawtooth with a differential input. If a single-ended input is used, a differential input is derived internally.

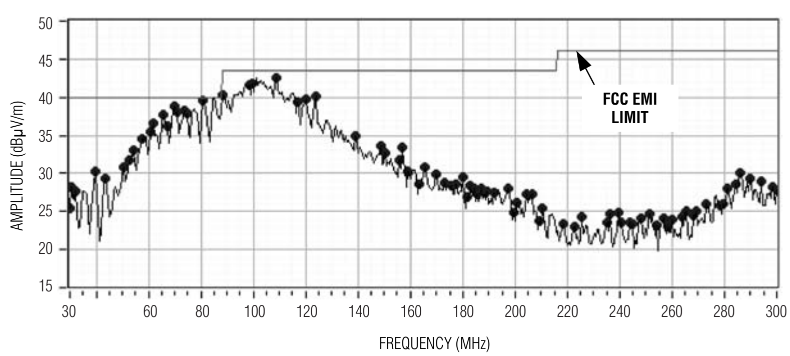

Figure 8. The MAX9705 radiated emissions data is shown for a 24in, unshielded twisted pair in spread-spectrum modulation mode.

DESIGN SHOWCASE

Divide and Conquer the Resistive Divider

The resistive voltage-divider is a basic circuit that is taught in every introductory electronics course, but problems can arise in its implementation. If you have ever stood in front of a resistor kit while punching values into your calculator, this article is for you. Typical resistor kits do not include every value, so finding an elusive ratio for which both values are commonly available can be a challenge. The small spreadsheet in **Figure 1** not only provides a table of 1% resistor values, but also makes it easier to find the ratio you need using two available resistors.

Standard 1% resistor values are logarithmically spaced in increments of exactly 1/96th of a decade. Figure 1 shows a spreadsheet table of 1% resistor values, rounded to three significant digits. The formula in cell B5, for example, is =10^((B$4+$A5)/96), which is copied to B5 through G20. Only the highlighted values are commonly available in resistor kits. You may notice that a typical kit contains slightly more than one quarter of the values.

	A	B	C	D	E	F	G
1	1% Standard Resistor Dividers						
2		Vout	Vfb	Steps			
3		1.5	0.6	16.9			
4		0	16	32	48	64	80
5	0	**1.00**	1.47	2.15	3.16	4.64	**6.81**
6	1	1.02	**1.50**	**2.21**	3.24	**4.75**	6.98
7	2	1.05	1.54	2.26	**3.32**	4.87	7.15
8	3	1.07	1.58	2.32	3.40	**4.99**	7.32
9	4	**1.10**	**1.62**	2.37	3.48	**5.11**	**7.50**
10	5	1.13	1.65	**2.43**	**3.57**	5.23	7.68
11	6	1.15	1.69	**2.49**	3.65	5.36	7.87
12	7	1.18	1.74	2.55	3.74	5.49	8.06
13	8	**1.21**	1.78	2.61	3.83	**5.62**	**8.25**
14	9	1.24	**1.82**	2.67	**3.92**	5.76	8.45
15	10	1.27	1.87	**2.74**	4.02	5.90	8.66
16	11	**1.30**	1.91	2.80	4.12	6.04	8.87
17	12	1.33	1.96	2.87	4.22	**6.19**	**9.09**
18	13	1.37	**2.00**	2.94	**4.32**	6.34	9.31
19	14	**1.40**	2.05	**3.01**	4.42	6.49	9.53
20	15	1.43	2.10	3.09	4.53	6.65	9.76

Figure 1. This spreadsheet not only lists 1% resistor values, but it also lets you find the divider needed using two commonly available resistors.

In a standard resistive divider application (**Figure 2**), the divider ratio R2/(R1 + R2) provides external feedback to the 4MHz regulated stepdown converter. The feedback threshold at the FB terminal is 0.6V, and the desired output is 1.5V.

To calculate values for R1 and R2, enter the regulator's desired output voltage (voltage at the top of the divider) in cell B3 and the regulator's feedback threshold (voltage at the divider midpoint) in cell C3. Then, the spreadsheet formula in cell D3 calculates =96*LOG(B3/C3-1), which returns a value of +16.9 steps for this example. This value is the number of 1% resistor-value steps separating R1 from R2. Therefore, if you use R2 = 100kΩ (cell B5), you would move +17 steps down the list to R1 = 150kΩ (cell C6). This is a good choice because both resistors are highlighted (commonly available).

You can see very quickly from the spreadsheet table that R2 = 110kΩ is not a good choice for this example because R1 = 165kΩ is not commonly available. You can also quickly identify an exhaustive list of commonly available resistor values that would be suitable: 1.00:1.50, 1.21:1.82, 1.62:2.43, 1.82:2.74, 2.00:3.01, 2.21:3.32, 3.32:4.99, and 4.99:7.50. When the number of steps (D3) is negative, R1 is less than R2 and you should therefore move in the opposite direction on the value list. In either direction, you may wrap around the list from 9.76 to 1.00, which means that you have moved to the next decade of resistance values.

The Figure 1 spreadsheet is available for download through links in application note 3646 at www.maxim-ic.com/AN3646. One link is for a spreadsheet in Excel® format, and the other is in Pocket Excel format for Pocket PC, formatted to fit the Pocket PC screen at 60% zoom. If you use resistive dividers for purposes other than regulator feedback, you may wish to rename cells B2 and C2 as "Vtop" and "Vmid," respectively.

DESIGN SHOWCASE

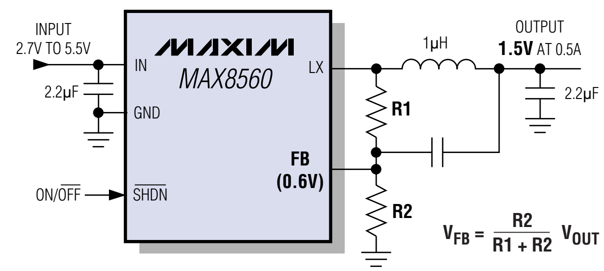

Figure 2. The resistor-divider ratio R2/(R1+R2) sets the output voltage for this 4MHz stepdown regulator. The spreadsheet of Figure 1 lets you select R1 and R2 from commonly available values.

A similar design idea appeared in the April 13, 2006 issue of Electronic Design.

DESIGN SHOWCASE

White-LED Driver Includes Overtemperature Protection

High-output LEDs are used in an ever-widening range of applications because of their advantages over conventional lighting sources. Among those LED advantages are higher efficiency, longer lifetime, and improved reliability. However, these high-output LEDs produce more heat than traditional LEDs.

The lifetime of any IC is reduced when it operates above its maximum operating temperature range. Irreparable damage occurs when the junction temperature exceeds a certain point. Philips Lumileds LUXEON® high-output LED modules reduce some of this heat because they are fabricated on a thermally-enhanced substrate. This material improves thermal performance and allows continuous operation at the high-current levels necessary for high-brightness lighting. However, in applications such as camera flashes, additional thermal protection is required to prevent damage from power dissipation during continuous operation.

The **Figure 1** circuit details a charge-pump regulator (IC1) designed to supply regulated current for up to eight white LEDs (WLEDs) in a camera-flash application. By paralleling all eight drivers, you can supply up to 480mA for a single 1W, LUXEON Star high-output WLED module. IC1 enters shutdown mode when the open-drain input, EN, is pulled to ground.

A compact, cost-effective thermal shutdown circuit can be implemented with a thermistor and a dual open-drain comparator with internal reference (IC2). Comparator A is configured to pull EN to ground when V_{THERM} falls below the internal 1.2V reference voltage. This occurs when the thermistor (R2) is hot. Comparator B is configured as an open-wire fail-safe that drives EN low when the thermistor connection is broken. In that latter case, V_{THERM} is pulled high by R1, causing comparator B to pull EN to ground. Resistive divider R3-R4 sets the open-wire threshold voltage, and resistor R1 with thermistor R2 sets the thermal-shutdown threshold.

A similar design idea appeared in February 2006 on the EE Times website.

LUXEON is a registered trademark of Philips Lumileds Lighting Company.

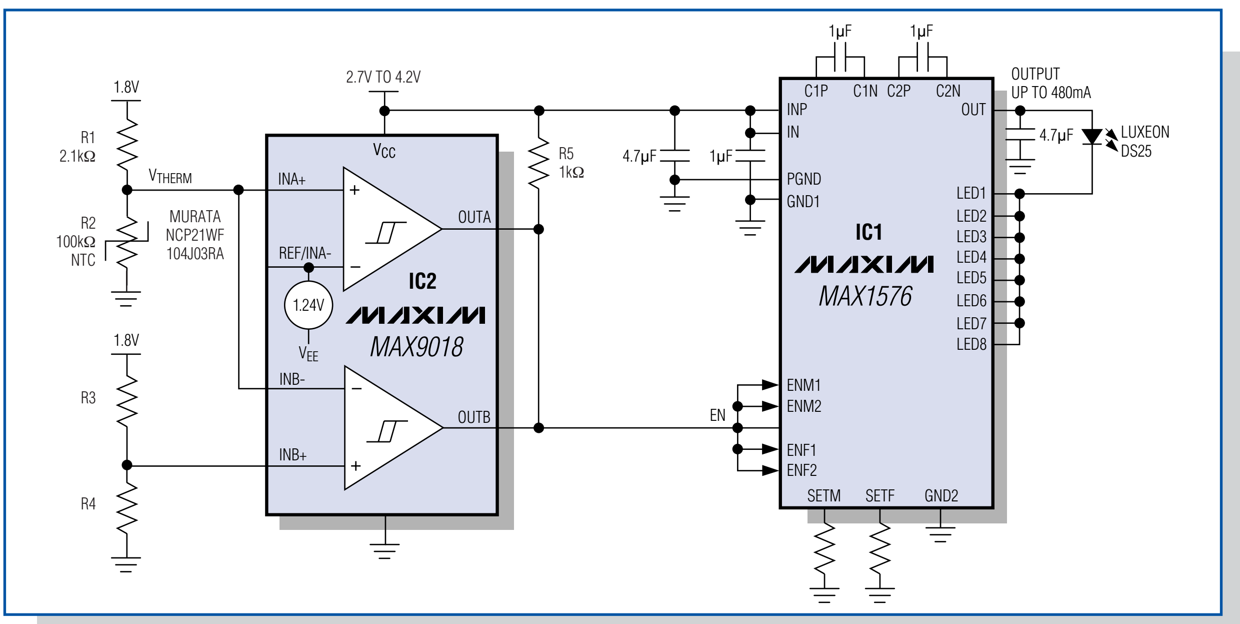

Figure 1. This circuit drives a high-output WLED. It shuts off the LED current in response to excessive temperature and open- or short-circuit thermistor faults.