MAXIM Engineering Journal

Volume Sixty-Six

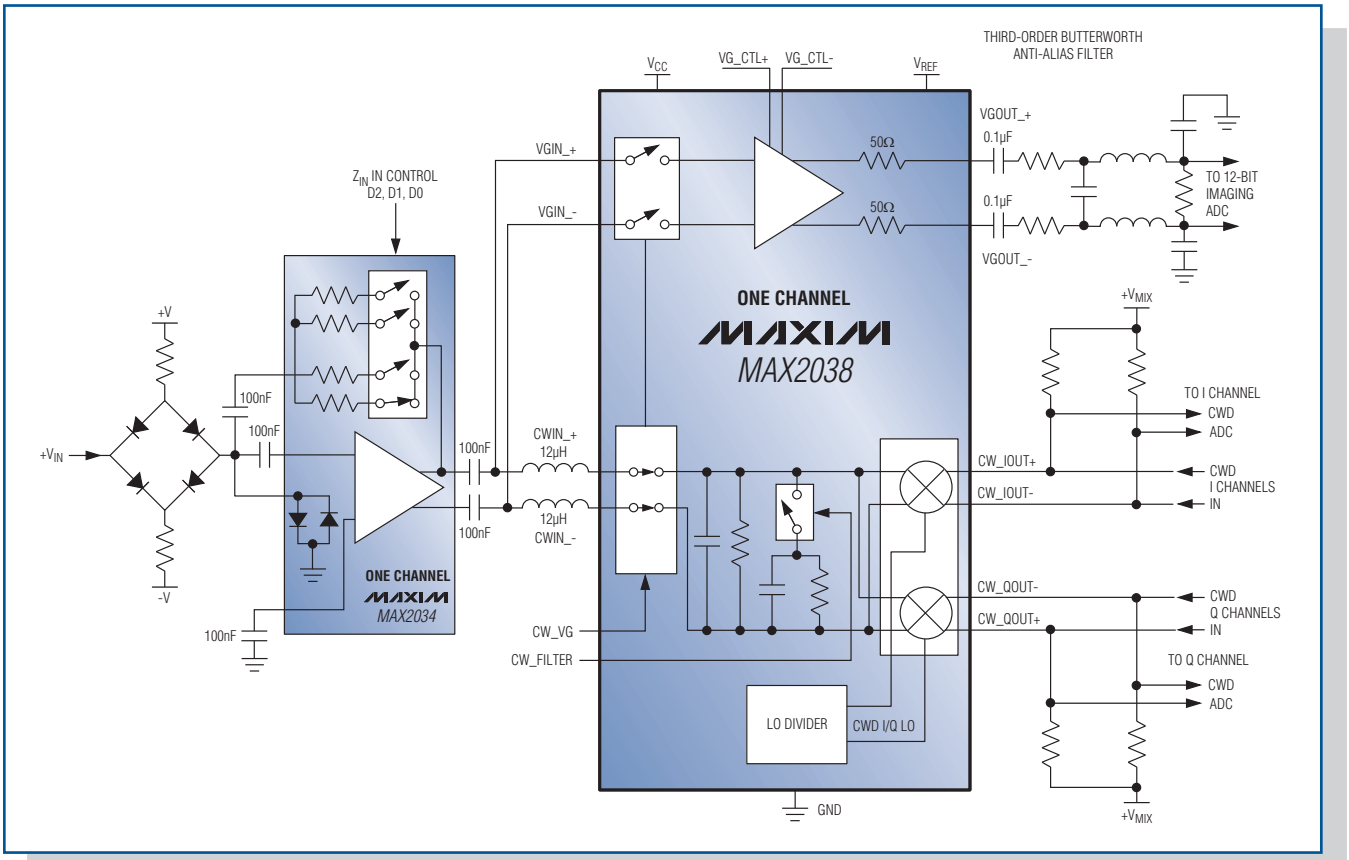

This simplified single-channel ultrasound receiver features the MAX2038 and MAX2034. The MAX2038 integrates eight VGA and CWD mixer I/Q beamformer channels onto a single chip, and the MAX2034 integrates four LNA channels onto a single chip. (See article on page 15.)

Letter from the CEO

Meeting the Packaging Challenges of Today and Tomorrow

Recently, a friend of mine called in a panic. He had dropped his iPhone and it was in pieces. He wanted to know if I could help repair it, since he knew I was an electrical engineer. Unfortunately, the damage was beyond a simple repair and he ended up purchasing a replacement. After we determined that it could not be repaired, my curiosity got the better of me. I asked if I could keep the broken device to do my own "tear down" of the hardware. I was curious.

When I opened the case, I was impressed by the high-density packaging to cram all the functions into one slim, attractive product. Inside were several complex digital ICs and many smaller, mostly mixed-signal devices, housed in chip-scale and near-chip-scale packages. This kind of packaging effort is second-nature to us at Maxim®, where we are always looking at ways to integrate functions and reduce component count. Also on the circuit board were many passive surface-mounted components and even a few discrete transistors—many of these small devices and components were devoted to power management and interface functions.

The "clutter" of discrete components and mixed-signal circuits occupies a significant amount of board space—almost 20% to 30% by my estimation. This scenario, though, is not unique. Just about every handheld product today is space constrained, which is also a prime example of how higher integration can "buy back" some of the limited space on the PC board. In fact, all portable systems will likely benefit from more-integrated solutions.

Many of the products that we design at Maxim are targeted at this same handheld portable application space. Over the years we worked closely with major customers to create highly integrated solutions on a single chip. Such integrated solutions, both custom and off-the-shelf, provide multiple benefits to the customer. Primarily, of course, the chips will reduce the component count and save space, but they will also often lower the power consumption.

Fewer components also translate into higher system reliability—there are simply fewer parts to fail. For the system vendor, fewer components lead to a shorter bill-of-materials, which reduces purchasing and manufacturing overheads.

Lower power consumption is a result of two important factors. First, attention to detail in circuit design minimizes operating currents; and second, integration improves efficiencies because the signals don't have to travel as far and parasitic losses are smaller. Lower power consumption also offers designers two options: the system can deliver longer battery life if the same size battery is used; or the system could be made smaller and lighter and still deliver the same runtime as the previous product if the battery size is reduced.

To implement highly-integrated solutions Maxim leverages proprietary, state-of-the-art process technologies and our extensive library of intellectual property developed during our long history of designing highly integrated, low-power analog and mixed-signal ICs. We also have leading-edge, tiny chip-scale packages (UCSPs™) and wafer level packaging (WLP) technology. This combination lets us deliver solutions that meet today's and tomorrow's design challenges. Talk to us about your system requirements.

We are always at your service,

Tunç Doluca
President and Chief Executive Officer

Add Thermal Monitoring to Reduce Data Center Energy Consumption

Scott Jones, Director, Automatic Information
Erin Mannas, Marketing Engineer
Maxim Integrated Products, Inc.

Precise and adaptable thermal management is imperative for efficiently managing the energy consumption of data centers and other temperature-sensitive environments. This article details how 1-Wire® technology enables double-digit reductions in the power consumption of climate-controlled environments.

Today's Challenge: Reduce Power Consumption in Data Centers

The EPA estimates that data centers accounted for 1.5% (61 billion kWh) of total U.S. electricity consumption in 2006.[1] Astonishingly, IT equipment was responsible for only half of this electricity consumption; power and cooling infrastructure accounted for the other half. Yet to date, most efforts to reduce the power consumption of data centers have addressed only the first half of the problem: the energy efficiency of the chips and components in IT equipment.

To effectively minimize overall data center energy consumption, it is imperative that the energy consumed by the cooling infrastructure be reduced as well. By implementing an intelligent cooling infrastructure, organizations will be able to accommodate higher power densities without building additional data centers. Additionally, they will reduce their electricity bill while minimizing their carbon footprint.

Conventional Cooling Infrastructures Waste Substantial Energy

The ambient temperature in a data center or other thermally controlled environment is not constant throughout the space. Temperature varies significantly across both the vertical and horizontal planes depending on proximity to active or inactive equipment, as well as airflow. Moreover, the cooling zones of computer room air conditioners (CRACs) frequently overlap, resulting in cooling redundancy. Simply stated, data centers waste substantial power by cooling areas that are already cool, because they lack intelligent thermal-monitoring capabilities.

The ability to identify specific areas where ambient temperature needs to be reduced is the key to reducing data center power consumption. However desirable this goal, it has not been simple to implement. The complexity of wiring and interfacing multiple temperature sensors is compounded by the difficulty of managing a vast array of temperature-measurement nodes. Consequently, multipoint temperature monitoring is particularly problematic.

Efficient Multipoint Temperature Measurement

Multiple temperature-measurement nodes, combined with appropriate temperature data processing, enable an optimized environmental temperature-control system that saves power and reduces cost. To implement such a system without adding any complex wiring overhead, designers can draw on various technologies—for example, a network of wireless sensors based on Zigbee® or other wireless standards, or a hardwired approach using a simple serial interface such as Maxim's 1-Wire technology. A 1-Wire system can route all sensor readings over a single wire, possibly an unused wire pair in a CAT5 network cable that is already part of the server farm infrastructure.

Wireless solutions would be easy to install, since few new wires would be needed (just a tap into the local power supply to power the module). However, today's per-module cost for each sensor could be $10 (U.S.) or more, thus adding thousands to tens of thousands of dollars to the data center overhead. Alternately, a hardwired system using a simple 1-Wire serial interface can run off parasitically derived power (less than 1mA per node) and keep costs per node to just a couple of dollars.

The 1-Wire interface can thus solve the wiring complexity challenge of large, multipoint temperature-sensing arrays (**Figure 1**).[2] Using 1-Wire technology, sensors are multidropped throughout an enclosure or space from a single, low-cost twisted pair like that available in CAT5 cables. Temperature data from individual sensors, or a group of sensors, is communicated digitally to a host processor through the 1-Wire protocol.

Location-Aware Temperature Sensors

Monitoring many temperatures in a data center is challenging. How will the measurement of a given sensor be associated with its physical location? Any solution that requires manual configuration is time-consuming, expensive, and prone to error. Chain-mode operation, a new feature offered on some 1-Wire temperature sensors, solves this problem (**Figure 2**).

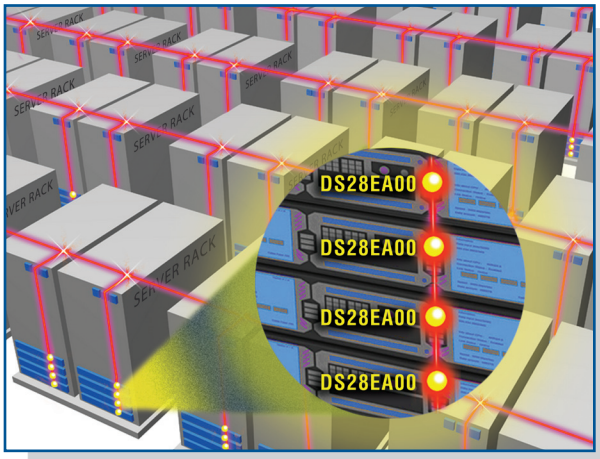

Figure 1. Energy management in a server farm can be improved by adding a network of 1-Wire digital temperature sensors, such as the DS28EA00, to monitor multiple locations for more efficient, controlled cooling.

All 1-Wire temperature sensors are factory programmed with a unique and unalterable 64-bit ID value that can be electronically read from the 1-Wire master device. Unique among Maxim's 1-Wire-based thermal sensors, the DS28EA00 has two additional pins to implement a sequence-detect function (i.e., chain mode). Using the DS28EA00's chain mode, the 1-Wire master reads and associates the sensors' registration numbers according to each device's physical location in a multidrop-configured connection. If the sequence-detect function is not needed, these two pins can be used as general-purpose inputs or outputs.

The chain-mode function, therefore, gives the host controller a computer-controlled and fully automated method to determine the sequential ordering of the unique ID for each device in a multidrop temperature-sensor configuration. This sequential ID information directly correlates to the physical construction of the multidrop temperature array and, therefore, the physical location of each sensor in the enclosure.[3] This correlation between the DS28EA00's 64-bit ID and physical location in the enclosure enables the application to use its environmental control algorithm to measure temperature in a storage tower at different heights and/or location.

Using Parasitic Power

Powering an array of temperature sensors can add complexity and cost to a distributed temperature-sensing application. Fortunately, 1-Wire temperature sensors can operate from parasitic power supplied by the single data line. This parasitic mode of operation is, in fact, a fundamental feature of 1-Wire products.

At the IC level, individual 1-Wire devices, or sensors, have internal circuits to capture and store electronic charge from the 1-Wire data waveforms. Whenever the 1-Wire line is in a logic 0 state (i.e., transmitting a data value of 0), sensors derive their operating power from this captured charge. Note that the actual temperature-measurement and digitization process requires more power than can be stored by this internal parasitic power circuit. Thus, when operating in a parasitic power mode, the 1-Wire line must remain at the logic 1 state during a temperature-conversion cycle. There is an alternative to parasitic-power mode that would remove this restriction: provide an external power source by tapping into a power supply near where the sensor is positioned. (The DS28EA00, for example, has a V_{CC} pin that can optionally be used to supply power to the IC.)

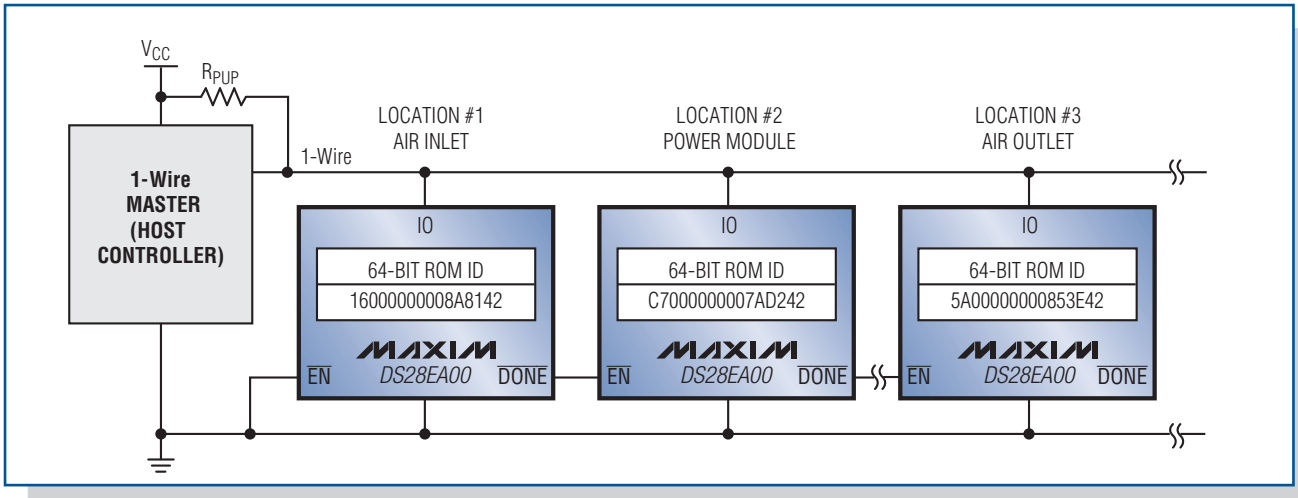

Figure 2. The DS28EA00 employs the chain-mode function in a typical 1-Wire network.

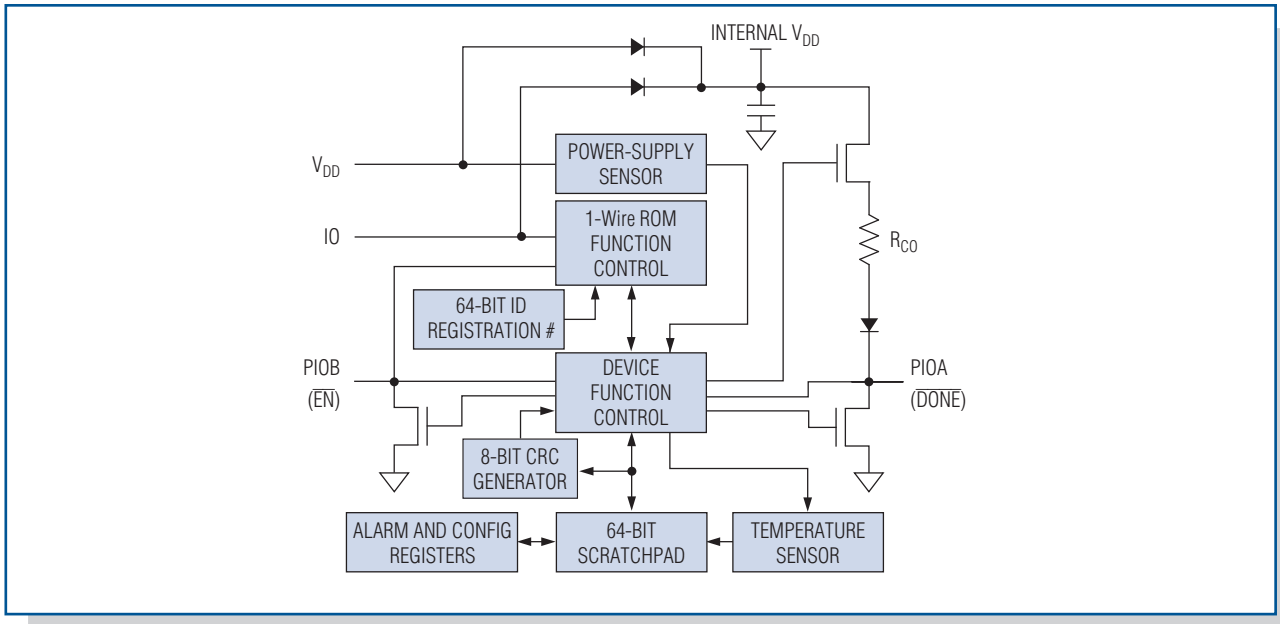

Figure 3. Able to run from either parasitically derived power or from an external supply, the DS28EA00 combines a temperature sensor, GPIO, alarm capabilities, and the 1-Wire interface in an 8-pin micro-small-outline package (µSOP).

Accurate and Dependable Measurements

A fundamental requirement for an accurate environmental temperature-control system is precise temperature measurement. For environments like a data center, temperature-measurement accuracy typically needs to be ±1°C or better. Exceeding this basic performance requirement, 1-Wire temperature sensors like the DS28EA00 provide ±0.5°C measurement accuracy.

The DS28EA00 also includes thermostatic-type functions with user-programmable temperature-measurement alarms which can be polled electronically. Additionally, the device's general-purpose I/O pins can be used to turn on and off visual or audible indicators, such as LEDs or buzzers, for notification of out-of-range temperature conditions (**Figure 3**).

In a typical server farm, each rack of multiple servers might incorporate several temperature sensors. These sensors, along with sensors from other racks in the farm, would send their temperature readings back to the host system that monitors the server farm. That host system would then control overall system cooling. Meanwhile, using its internal alarm capabilities, each 1-Wire sensor could locally control indicators, alarms, and/or temperature subsystems to enable localized cooling.

Multipoint Temperature Monitoring Reduces Cooling Energy by 30%

One such system implementation by a leading manufacturer of large-scale server systems leverages the DS28EA00. In an initial implementation, the company reported a 30% reduction in cooling energy in their data center. They are now saving over two million kWh/year, which equates to a reduction of over 1300 tons/year of CO_2 in the atmosphere.

Summary

Cooling infrastructure accounts for half of the electricity used by data centers. Consequently, intelligent thermal monitoring is critical for effectively minimizing data center power consumption.

1-Wire temperature sensors like the DS28EA00 provide a cost-effective solution for implementing the advanced thermal monitoring needed in twenty-first century applications such as energy-efficient data centers. Data shows that more efficient monitoring of many temperature sensors provides better management of the entire system and significantly reduces overall power consumption.

References

1. "Report to Congress on Server and Data Center Energy Efficiency, Public Law 109-431," U.S. Environmental Protection Agency ENERGY STAR Program (August 2, 2007).
2. For a basic overview of 1-Wire concepts, see Maxim application note 1796, "Overview of 1-Wire Technology and Its Use", and a Flash® presentation at: www.maxim-ic.com/1-WireFlash.

3. For information on chain mode, see Maxim application note 4037, "Regain Location Information by Leveraging the 1-Wire Chain Function—A Simple Signaling and Protocol Method Determines Device Physical Location."

The authors would also like to thank Matthew Lewsadder, a Maxim editor, for his contributions to this article.

Assess Power-Supply Noise Rejection in Low-Jitter PLL Clock Generators

Sharon Wang, Principle Member of Technical Staff
John Abcarius, Senior Member of Technical Staff
Maxim Integrated Products, Inc.

This article discusses the effects of power-supply noise interference on PLL-based clock generators, and describes several measurement techniques for evaluating the resulting deterministic jitter (DJ). Derived relationships show how frequency-domain spur measurements can be used to evaluate timing-jitter behavior. Laboratory bench-test results are used to compare the measurement techniques, and demonstrate how to reliably assess the power-supply noise rejection (PSNR) performance of a reference clock generator.

Clock generators that employ PLLs are widely used in network equipment for generating high-precision and low-jitter reference clocks or for maintaining a synchronized network operation. Most clock oscillators give their jitter or phase-noise specification using an ideal, clean power supply. In a practical system environment, however, the power supply can suffer from interference due to on-board switching supplies or noisy digital ASICs. To achieve the best performance in a system design, it is important to understand the effects of such interference.

First we will examine the basic power-supply noise rejection (PSNR) characteristics of a PLL-based clock generator. We will then explain how to extract timing-jitter information from measurements taken in the frequency domain. These techniques are then applied and several different measurement methodologies are compared using lab bench testing. Finally, we will summarize the merits of the preferred approach.

PSNR Characteristics of PLL Clock Generators

A typical PLL clock generator is shown in **Figure 1**. Since the output driver can have very different PSNR performance for different types of logic interfaces, the following analysis will focus on the impact of supply noise on the PLL itself.

Figure 2 shows the PLL phase model. The model assumes that the power-supply noise, V_N, is injected into the PLL/VCO, and that the divide ratios, M and N, are set to 1.

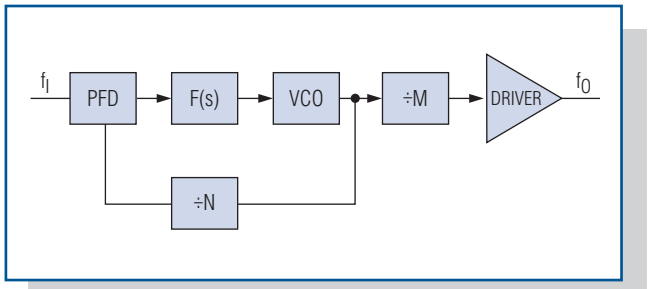

Figure 1. A typical topology for a PLL clock generator.

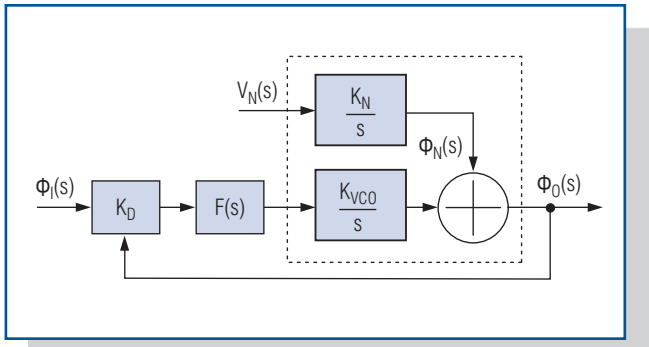

Figure 2. The phase model for a PLL.

The PLL closed-loop transfer function from $V_N(s)$ to $\varnothing_O(s)$ is given by:

$$\frac{\phi_O(s)}{V_N(s)} = \frac{K_N}{s + K_{VCO} \times K_D \times F(s)} \qquad \text{(Eq. 1)}$$

For a typical 2nd-order PLL:

$$F(s) = K_N \times \frac{(s + \omega_Z)}{s} \qquad \text{(Eq. 2)}$$

$$\frac{\phi_O(s)}{V_N(s)} \cong \frac{s \times K_N}{(s + \omega_Z) \times (s + \omega_{3dB})} \qquad \text{(Eq. 3)}$$

Here ω_{3dB} is the PLL 3dB bandwidth, ω_z is the PLL zero frequency, and $\omega_z << \omega_{3dB}$.

Equation 3 demonstrates that in a PLL clock generator the power-supply noise is rejected by 20dB/dec, when the power-supply interference (PSI) frequency is greater than the PLL 3dB bandwidth. For PSI frequencies between ω_z and ω_{3dB}, the output clock phase varies with the PSI amplitude as:

$$\frac{\phi_O(s)}{V_N(s)} \cong \frac{K_N}{2 \times \pi \times f_{3dB}} \qquad \text{(Eq. 4)}$$

As an example, **Figure 3** shows the PSNR characteristics of a PLL for two different settings of the PLL's 3dB bandwidth.

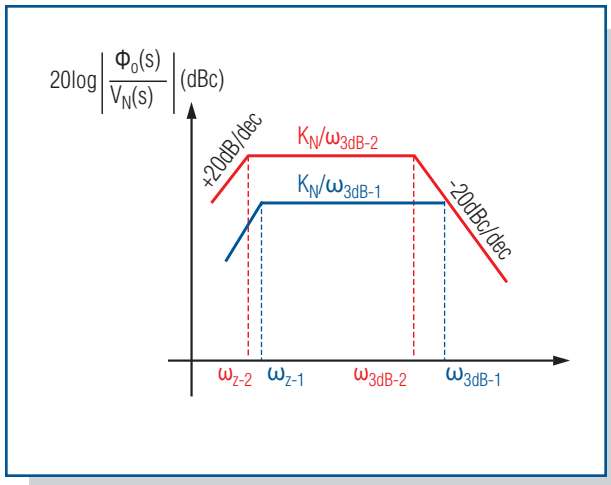

Figure 3. Typical PLL PSNR characteristics.

Conversion of Power Spectrum Spurs to DJ

When a single-tone sinusoidal signal, f_M, is applied to the power supply of a PLL, it produces a narrowband phase modulation at the clock output. This phase modulation can be generally described using Fourier series representation:

$$V(t) = \sum_{n=-\infty}^{n=+\infty} V_O \times J_N(\beta) \times \cos[(\omega_O + n \times \omega_M) \times t] \qquad \text{(Eq. 5)}$$

Here β is the modulation index representing the maximum phase deviation. For a small index modulation (β << 1), the Bessel function can be approximated as:

$$J_N(\beta) \approx \frac{\beta^N}{2^N \times n!} \qquad \text{(Eq. 6)}$$

Here n = 0 represents the carrier itself. When n = ±1, the phase-modulated signal is given by:

$$V(t) = V_O \times \frac{\beta}{2} \times \cos[2 \times \pi (f_O \pm f_M) \times t] \qquad \text{(Eq. 7)}$$

When measuring the double-sideband power spectrum $S_V(f)$, if variable x represents the level difference between the carrier at f_O and the fundamental sideband tone at f_M, then:

$$x = 20 \times \log\left(\frac{\beta}{2}\right) \quad [\text{dBc}] \qquad \text{(Eq. 8)}$$

Since β is the maximum phase deviation in radians, the peak-to-peak DJ caused by this small index phase modulation can be derived as:

$$DJ = \frac{2 \times 10^{-\left|\frac{x}{20}\right|}}{\pi \times f_O(\text{Hz})} \times 10^{12} \quad [ps_{P\text{-}P}] \qquad \text{(Eq. 9)}$$

The above analysis assumes that there is no amplitude modulation contributing to the tone at f_M. In reality, both amplitude and phase modulation can be generated, thus reducing the accuracy of this approach.

Conversion of Phase-Noise Spectrum Spurs to DJ

There is a way to avoid the amplitude modulation effect when measuring the power spectrum $S_V(f)$. One can instead calculate the DJ by measuring the spur in the phase-noise spectrum, while applying a single-tone sinusoidal interference on the supply. With the variable y (dBc) representing the measured single-sideband-phase spurious power at frequency offset f_M, the resultant phase deviation ΔØ (rad_{RMS}) can be derived as:

$$y = 10 \times \log(\Delta\phi^2) - 3 \quad [\text{dBc}] \qquad \text{(Eq. 10)}$$

$$\Delta\phi = \sqrt{2} \times 10^{\frac{y}{20}} \quad [rad_{RMS}] \qquad \text{(Eq. 11)}$$

$$\Delta\phi = \frac{2 \times 10^{\frac{y}{20}}}{\pi \times f_O(\text{Hz})} \times 10^{12} \quad [ps_{P\text{-}P}] \qquad \text{(Eq. 12)}$$

It should be noted that the single-sideband-phase spectrum in the above analysis is not the folded version of the double-sideband spectrum. That is the reason for the 3dB component in Equation 10. **Figure 4** shows the relationship between the DJ and the phase spurious power given by Equation 12.

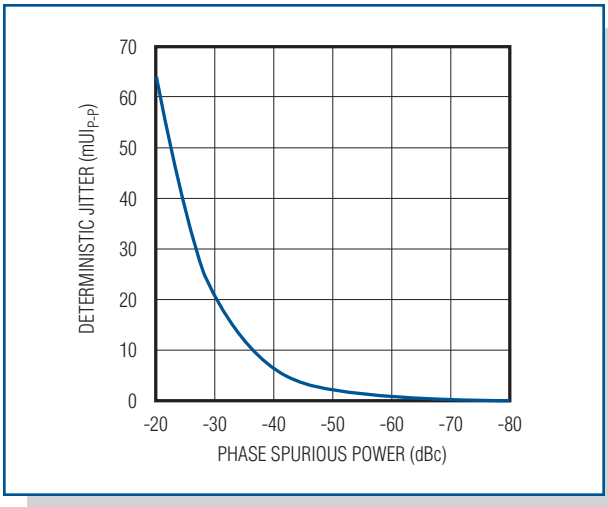

Figure 4. DJ vs. phase spurious power.

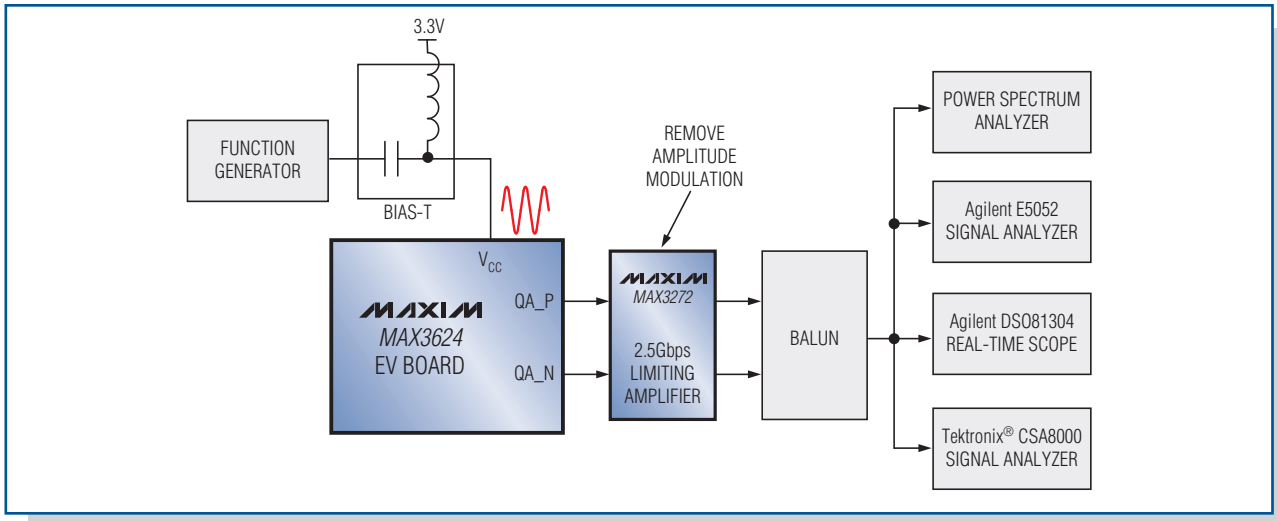

Figure 5. PSNR measurement setup.

PSNR Measurement Techniques

This next section demonstrates five different ways of measuring the PSNR of a clock source. The MAX3624 low-jitter clock generator serves as an example. The measurement setup shown in **Figure 5** uses a function generator to inject a sinusoidal signal onto the power supply of the MAX3624 evaluation (EV) board. The amplitude of the single-tone interference is measured directly at the V_{CC} pin close to the IC. A limiting amplifier, the MAX3272, is used to remove amplitude modulation; it is followed by a balun that converts the differential output into a single-ended signal for driving the different test equipment. To compare the results from different tests, all the measurements were done under the following conditions:

- Clock output frequency: f_O = 125MHz
- Sinusoidal modulation frequency: f_M = 100kHz
- Sinusoidal signal amplitude: 80mV_{P-P}

Method 1. Power Spectrum Measurement

When observed on a power spectrum analyzer, the narrow-band phase modulation appears as two sidebands around the carrier. **Figure 6** shows this case when viewed using the spectrum monitor function of the Agilent® E5052. The measured first sideband amplitude relative to carrier amplitude is -53.1dBc, which translates to 11.2ps_{P-P} DJ, according to Equation 9.

Method 2. SSB Phase Spurious Measurement

On a phase-noise analyzer, the PSI will manifest itself as a phase spur relative to the carrier. The measured phase-noise spectrum is plotted in **Figure 7**. The phase spurious power at 100kHz is -53.9dBc, which translates to 10.2ps_{P-P} DJ using Equation 12.

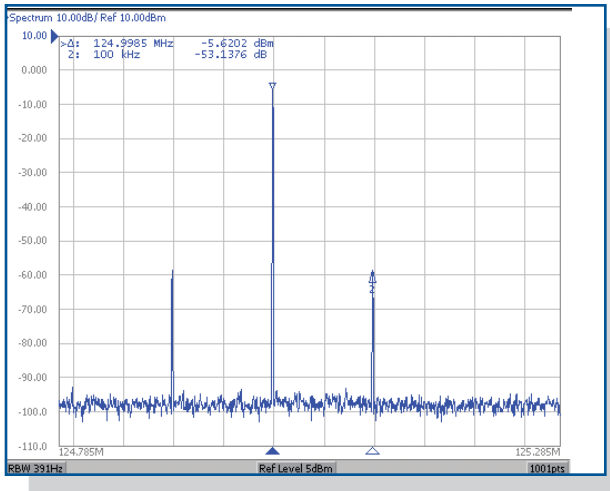

Figure 6. Measured power spectrum.

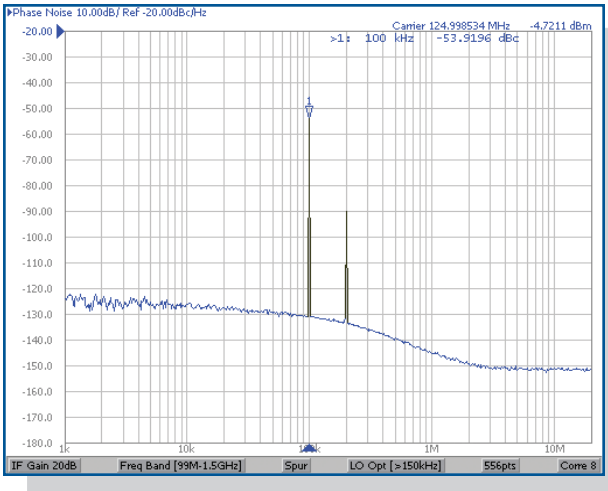

Figure 7. Measured SSB phase noise and spur.

Method 3. Phase Demodulation Measurement

Using the Agilent E5052 signal analyzer, the phase-demodulated sinusoidal signal at 100kHz is measured directly as shown in **Figure 8**, which gives the maximum phase deviation from its ideal position. The peak-to-peak phase deviation is 0.47°, which translates to 10.5ps$_{P-P}$ at an output frequency of 125MHz.

Method 4. Real-Time Scope Measurement

In a time domain measurement, the DJ caused by PSI can be obtained by measuring the time interval error (TIE) histogram. On a real-time scope, the clock-output TIE distribution will appear as a sinusoidal probability density function (PDF) when a single-tone interference is injected into the PLL. The DJ can be estimated using the dual-Dirac model[1] by measuring the peak distance between the mean of two Gaussian distributions from the TIE histogram. **Figure 9** shows the measured TIE histogram using the Agilent Infiniium DSO81304A 40Gsps real-time scope. The measured peak separation is 9.4ps.

It should be noted that the memory depth of the real-time scope may limit the low sinusoidal modulation frequency that can be applied to the PLL supply. For example, if the test equipment has a memory depth of 2Msps when the sample rate is set to 40Gsps, it would only be able to capture jitter frequency components down to 20kHz.

Method 5. Sampling Scope Measurement

When a sampling scope is used, a synchronous trigger signal is required for analyzing the clock jitter under test. Two triggering methods can be used for TIE measurements.

The first method is to apply a low-jitter reference clock to the input of the PLL clock generator; use the same clock source as the trigger for the sampling scope. **Figure 10a** shows the measured TIE histogram, which gives a peak spacing of 9.2ps. The advantage of triggering with a reference clock is that the measured TIE histogram peak separation is independent of the horizontal time delay from the trigger position. However, the measured TIE histogram can be affected by the triggering clock jitter. Therefore, it is important to use a clock source that has much lower jitter than the clock generator device under test.

The alternate approach uses self-triggering to eliminate the impact of triggering clock jitter. In this case, the output of the clock generator under test is separated into two identical signals using a power splitter. One signal is applied to the data input of the sampling scope, another signal to the trigger input. Since the triggering signal contains the same DJ as the test signal, the histogram peak separation varies when the horizontal position of the scope's main time base is swept through one period of the sinusoidal modulation frequency. At a horizontal position of one-half period of the modulation signal, the peak separation on the TIE histogram will be twice the DJ from the test signal. **Figure 10b** shows the measured MAX3624 TIE histogram when the horizontal time delay is set to 5µs. The estimated TIE peak separation is 19ps, which gives an equivalent DJ of 9.5ps$_{P-P}$.

Figure 10c shows the measured TIE histogram peak spacing at a different horizontal time delay from the trigger point. For comparison, the TIE result is also shown when the sampling scope is triggered by a reference clock input.

Measurement Summary

Table 1 summarizes the measured DJ at the MAX3624 125MHz clock output. Data was gathered using the different methods discussed above. It should be noted that the DJ measured using a dual-Dirac approximation from the TIE histogram is slightly smaller than the DJ obtained from the frequency-domain spectral analysis. This

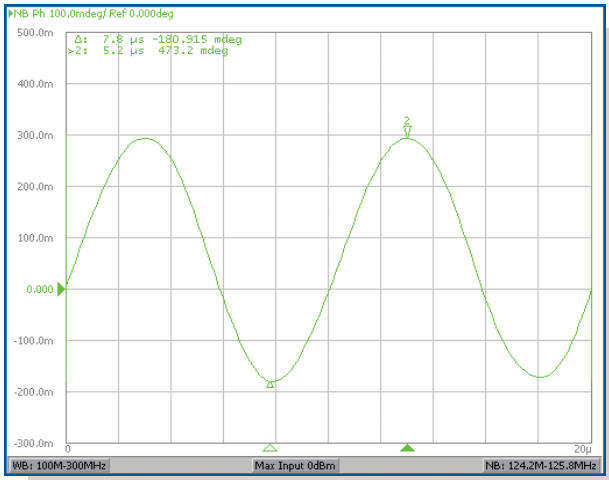

Figure 8. MAX3624 phase demodulation signal.

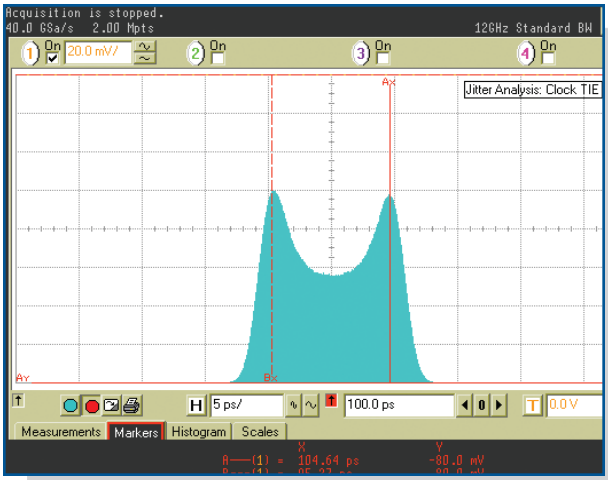

Figure 9. Measured TIE histogram.

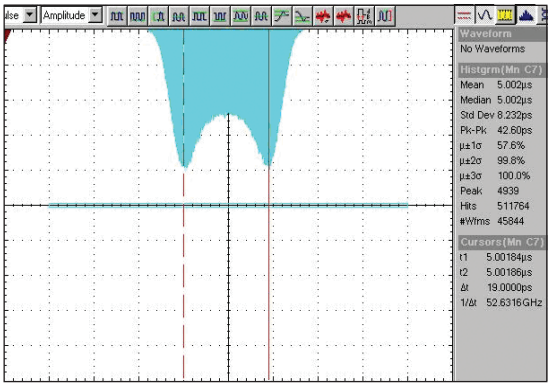

Figure 10a.

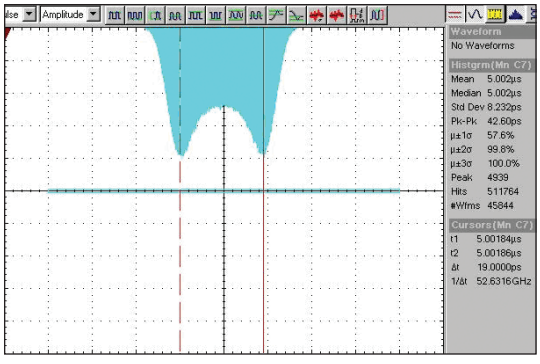

Figure 10b.

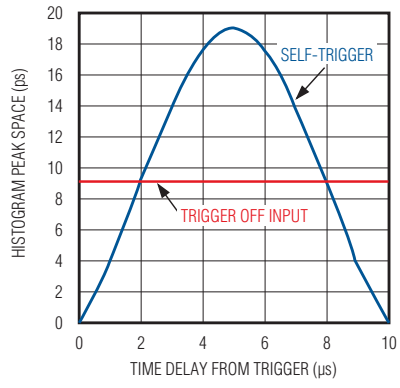

Figure 10c.

Figure 10. TIE histograms are shown for various trigger conditions: triggered by REF_IN (a); self-triggered, t_d = 5µs (b); and peak spacing vs. time delay from trigger (c).

difference is caused by the process of convolution of the sinusoidal jitter (SJ) PDF with Gaussian distribution of the random jitter component.[1] Therefore, the DJ extracted from the dual-Dirac model is only an estimation; it should only be applied when the standard deviation of the random jitter is much smaller than the distance between the two peak separations of the jitter histogram.

Table 1. DJ Comparison*

Measurement Methods	DJ (ps_{P-P})
Power Spectrum	11.2
SSB Phase Spurious	10.3
Phase Demodulation	10.5
Real-Time Scope	9.4
Sampling Scope (Reference Triggered)	9.2
Sampling Scope (Self-Triggered)	9.5

*80mV_{P-P}, 100kHz sinusoidal signal on the supply.

Conclusion

For the relatively large interference used in the examples, the results were well correlated. However, when the level of interference drops relative to the random jitter, the time-domain methods become less accurate. Furthermore, if the clock signal is corrupted by amplitude modulation, measurements using a power spectrum analyzer become unreliable. Therefore, of all the methods presented, the phase spur power measurement using a phase-noise analyzer is the most accurate and convenient way to characterize the PSNR of a clock generator. The same method can be extended for evaluating the DJ aspect caused by other spurious products appearing on the phase-noise spectrum.

Reference

1. Agilent white paper, "Jitter Analysis: the dual-Dirac Model, RJ/DJ, and Q-Scale."

A version of this article was published online by Electronic Design *magazine, March 27, 2009.*

Understand Thermal Derating Aspects of PWM ICs to Ensure the Best System Performance

Suresh Kariyadan, Senior Member of Technical Staff*
Nitin Kalje, Director, Product Definition, System and Power Management Business Unit
Maxim Integrated Products, Inc.

Modern DC-DC converters use PWM controllers with integrated MOSFETs to achieve the highest power density for DC-DC modules. Since the power MOSFETs are inside the PWM chip, they can significantly affect the thermal performance of the device. For optimal device performance, it is thus critical to establish the thermal-derating aspects of PWM ICs under realistic operating conditions. This article describes how to construct and calibrate a thermal-derating box for evaluating the thermal performance of PWM controllers.

Introduction

The final operating environment for a chip is often not well known during the IC's evaluation stage, and that environment can vary considerably from application to application. It is especially critical to understand thermal derating in portable and nonportable applications, since end systems often depend on forced airflow for cooling. This is why thermal-derating graphs are included in data sheets for pulse-width-modulated (PWM) controllers with internal MOSFETs. The derating curves show how much power can be drawn from the chip under given airflow and ambient temperature conditions.

To ensure that the application's thermal environment does not overload the PWM controller, a thermal-derating graph can be used to select proper PWM ICs. A test system with a thermal-derating box provides a practical way of evaluating the thermal-derating performance of the PWM chip.

This article describes how to construct and calibrate a thermal-derating box. Test results for two PWM controllers show a close match to actual operating conditions.

Standardizing Thermal Testing

For a given module size (including heatsink) and for a given airflow, the maximum power that can be dissipated is limited by physics. Thermal performance can, therefore, be estimated. However, without some type of standardized control of the IC test environment, the thermal-derating results from different vendors and for different packages will not be consistent.

One approach for generating standard device-derating data is to actually measure the thermal performance of the module under different airflow and ambient temperature conditions. A thermal-derating box will have an evaluation (EV) board with the PWM IC mounted inside the box, and a fan that can be calibrated to obtain uniform airflow. The test results would be published and used by designers as a basis for choosing the right kind of module for the end application.

Test Apparatus and Calibration

A standard thermal-derating box (**Figure 1**) was created. Box dimensions should be adjusted so that the largest EV board fits inside and sufficient clearance remains around the box. The minimal linear dimensions of the box should be 1ft x 1ft. The height of this test box should be 3in to accommodate a fan. Two or more fans can be used to achieve uniform airflow inside the box. The box material should have low thermal conductivity; polycarbonate, polypropylene, or a glass epoxy sheet can be used. The thickness should be at least 3.2mm to provide rigid sides.

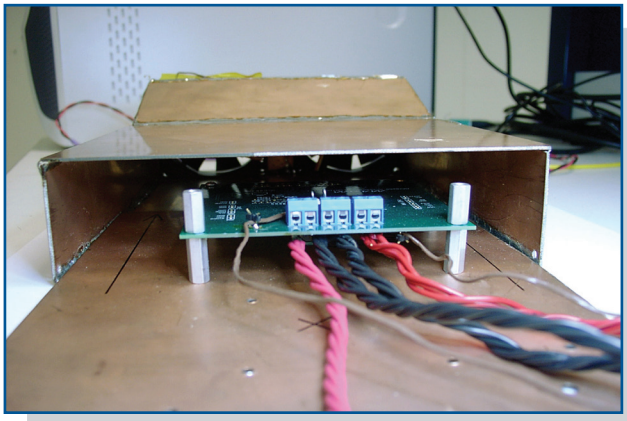

Figure 1. The thermal derating box surrounds the EV board and keeps the airflow predictable to ensure reproducible measurements.

Place the box horizontally inside a thermal chamber. Calibrate the box with an average airflow meter (**Figure 2**) to measure airflow at the open end of the box at left, center, and right, and to confirm that airflow is uniform inside the box. When a module is placed in the box, make sure that the fan is at least 2in from it. Airflow can be adjusted by using a voltage-controlled variable-speed fan and increasing or decreasing the voltage.

Start with a standard EV board, which can be mounted on a bigger board for convenient handling and placement. Elevate the bigger board at least 1in from the bottom of the box to ensure uniform airflow under the board.

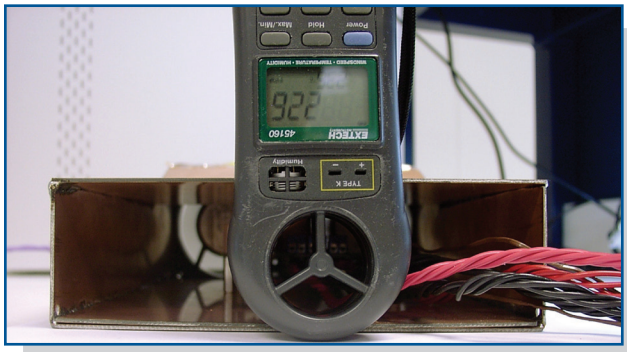

Figure 2. An airflow meter at the open end of the box measures the airflow to calibrate subsequent thermal measurements.

Once the box is inside the oven, you must determine if the oven's airflow will affect the airflow inside the derating box. If the oven's airflow does affect the box's airflow, then place a larger box over the test enclosure to maintain uniform airflow inside the box. To minimize interference, consider mounting the derating box with its fan airflow at a right angle to the oven's airflow.

Common fasteners and adhesives can be used to assemble the box—just make sure that they can withstand higher temperatures in the oven. A thermocouple can be placed at the geometric center of the IC. However, the attachment of the thermocouple to the IC is very critical, since you must be sure that the thermocouple does not act as a heatsink. The best approach is to use a minimal amount of thermal epoxy to attach thermocouple wire to the IC.

The meter measuring the thermocouple output should float electrically so that temperature readings are not affected by any voltage applied to the IC. Thermocouple wire size can be 30 AWG or smaller, and the wires should be routed to minimize interference with the airflow. The same holds true for the power and any other wires coming from the board under test.

Before actual device testing, the thermal-derating box must be calibrated (see Figure 2). **Table 1** shows the calibration data for an empty test box, while **Table 2** shows data for the test box with the EV board. Data plots are shown in **Figures 3a** and **3b**. For both configurations, almost uniform airflow is measured inside the box.

With the box ready, mount the EV board on the test board and put everything inside the thermal chamber. Run the chamber for some time so that thermal equilibrium is reached. Then the EV board can be loaded and run until temperature readings stabilize. If two readings taken 5 minutes apart do not change more than 0.2°C, then it can be assumed that thermal equilibrium has been reached.

Note, finally, that thermal derating with no airflow is harder to determine since natural convection is unstable, especially at higher IC temperatures and higher power levels.

Table 1. Airflow Calibration Data for the Empty Test Box

	Airflow (LFM)		
Fan Supply (V)	**Left**	**Center**	**Right**
3.3	102	98	91
4	215	207	187
6	339	337	323
8	449	447	445
10	565	585	581
12	685	709	673

Table 2. Airflow Calibration Data for the Test Box with EV Board

	Airflow (LFM)		
Fan Supply (V)	**Left**	**Center**	**Right**
3.3	91	90	89
4	189	160	205
6	333	321	325
8	455	445	443
10	561	581	565
12	673	689	671

Test Results

To verify the performance of the test box, experiments were done to measure the thermal performance of two Maxim PWM ICs, the MAX15035 and MAX8686. The test results are shown in **Figures 4a** and **4b**, respectively. The MAX15035 is a 15A step-down regulator with internal switches. The MAX8686 is a single/multiphase, step-down, DC-DC converter that can deliver up to 25A per phase. The MAX15035 EV board has four layers measuring 2.4in x 2.4in with 2oz copper, while the MAX8686 EV board contains six layers and measures 3.5in x 3.0in with 2oz copper.

By capturing this real-world data, designers can more accurately determine the necessary thermal derating for their application. From these thermal-derating graphs, the end user can choose the right part for their application. For a given ambient temperature and airflow, the graph shows how much power can be handled by a PWM IC without exceeding the chip's safe operating region. If a PWM IC does not meet the safe operating criteria for an application, then the user either has to increase the airflow available for cooling or improve the thermal design.

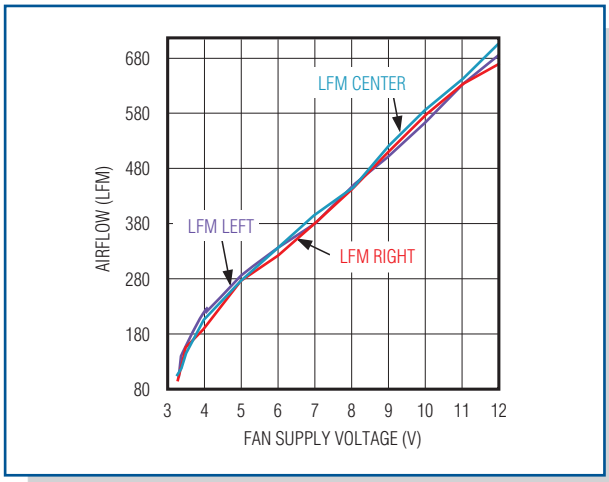

Figure 3a. With the test box empty, three airflow measurements provide baseline information.

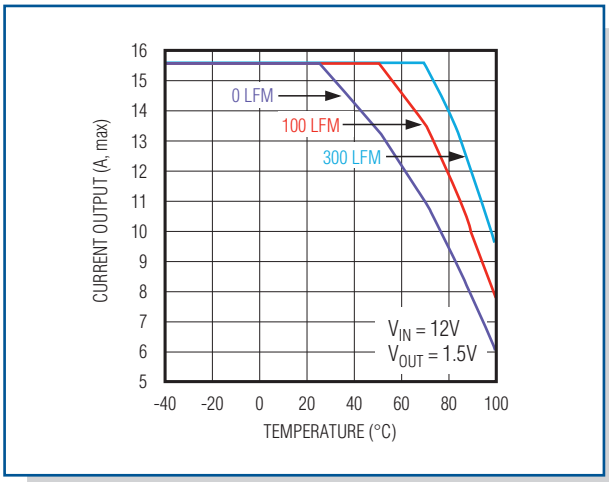

Figure 4a. Maximum output current vs. ambient temperature for the MAX15035. With the MAX15035 EV board in the test box, the thermal derating curves are determined for three different levels of airflow.

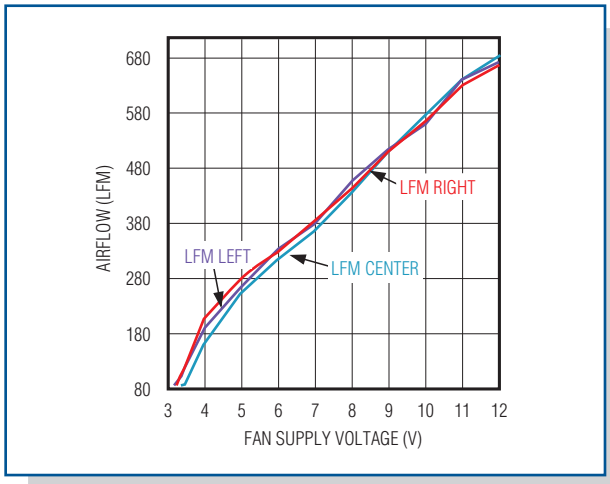

Figure 3b. With the EV board inserted into the test box, the airflow measurements are repeated to identify differences in the airflow vs. the empty box.

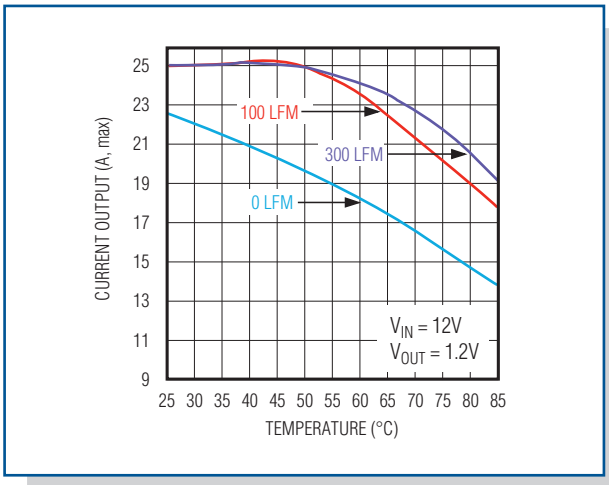

Figure 4b. Maximum output current vs. ambient temperature for the MAX8686. The thermal derating graph of the MAX8686 shows that at +50°C ambient the controller can handle its rated 25A current with as little as 100 LFM of airflow.

*Suresh Kariyadan is no longer with Maxim Integrated Products, Inc.

A version of this article was published in Power Electronics Technology *magazine in February 2009.*

Continuous-Wave Doppler (CWD) Design Challenges in Compact Ultrasound-Imaging Equipment

John Scampini, Director, Equipment Management
Maxim Integrated Products, Inc.

Recent advances in highly integrated, low-power bipolar amplifiers and continuous-wave Doppler (CWD) mixer/beamformer chips now allow designers to achieve high-end CWD performance for the next generation of compact ultrasound-imaging equipment.

One of the most demanding clinical diagnostic tools employed in ultrasound systems is continuous-wave Doppler (CWD). Unfortunately, design tradeoffs required to minimize space and cost have resulted in CWD receiver implementations with less than optimal sensitivity. By analyzing why current-generation CWD receivers have limited performance, designers can leverage recent advances in integration that yielded highly integrated, low-power bipolar amplifiers and CWD mixer/beamformer chips. These next-generation circuits allow new CWD solutions to provide uncompromised diagnostic performance.

CWD Basics

In a typical phased-array CWD implementation, the 64 to 128 available ultrasound transducer elements are split into approximately equal halves about the center of the transducer aperture. Half of the elements are used as transmitters to produce a focused acoustic CWD transmit beam, and the other half as receivers to produce a focused receive beam. The signals applied to the transmit elements are square waves at the Doppler frequency of interest, typically 1.0MHz to 7.5MHz. The transmit beam is focused by properly phasing the signals applied to the transmit elements. In a similar way, the CWD received signals are focused by phasing and summing the signals from each receive element.

The resultant "beamformed" CWD received signal is a combination of strong signals from stationary tissue, commonly referred to as clutter, and the much weaker Doppler signals from moving blood. The typical clutter signal on the input of an individual phased-array receive channel can be as large as $200mV_{P-P}$, whereas the noise floor of the receiver referred to input can be as low as $1nV/\sqrt{Hz}$. This suggests that a per-channel SNR of approximately 157dBc/Hz is required for optimal receiver performance.

For a typical hypothetical CWD receiver with 64 channels, it should be noted that the SNR requirements are extreme. The noise from each of the receive channels is not coherent, and, as a result, the noise floor of the summed 64-channel, beamformed signal will be approximately 18dB above the noise floor of an individual channel. The CWD signal on each channel, however, is coherent, and the beamformed CWD signal should be approximately 36dB larger than the CWD signal on an individual channel. This "summing gain" effect results in a required beamformed SNR approximately 18dB higher than an individual channel, or about 175dBc/Hz! To make matters even more difficult, low-velocity Doppler signals of interest can be within 1kHz or less of the clutter signal. One can easily appreciate why this ultrasound modality is so challenging to implement effectively.

Delay-Line-Based CWD Beamforming

Current compact ultrasound systems have typically implemented CWD using an analog delay-line receiver (**Figure 1**). Input signals from the ultrasound receive elements are buffered and amplified by approximately 20dB using LNAs. The outputs of these LNAs are converted to currents, which are then beamformed at the incoming RF frequency using a combination of crosspoint switches and analog delay lines.

This architecture is relatively easy to implement in a compact system since it relies on easy-to-integrate voltage-to-current converters, analog switches, a few passive delay lines, and a single I/Q mixer pair. The required delay for each receiver is achieved by programming the crosspoint switches to sum and route current signals through the appropriate taps in the delay line.

The beamformed RF CWD signal is then mixed to baseband I and Q audio frequency signals, so the I and Q signals can be bandpass filtered and then converted to a digital format using high-resolution audio ADCs for digital spectral processing. The RF-to-baseband mixing process is generally the SNR bottleneck in this receiver lineup, and is where significant CWD performance degradation occurs. The I/Q RF mixers for the beamformed signal require a dynamic range of approximately 175dBc/Hz at a 1kHz offset in this 64-channel example.

Mixers with such a performance capability are extremely difficult to find or design. In addition, local-oscillator drive signals must have extremely low jitter performance. Unfortunately, no readily available logic family has performance near this level. Although delay-line CWD

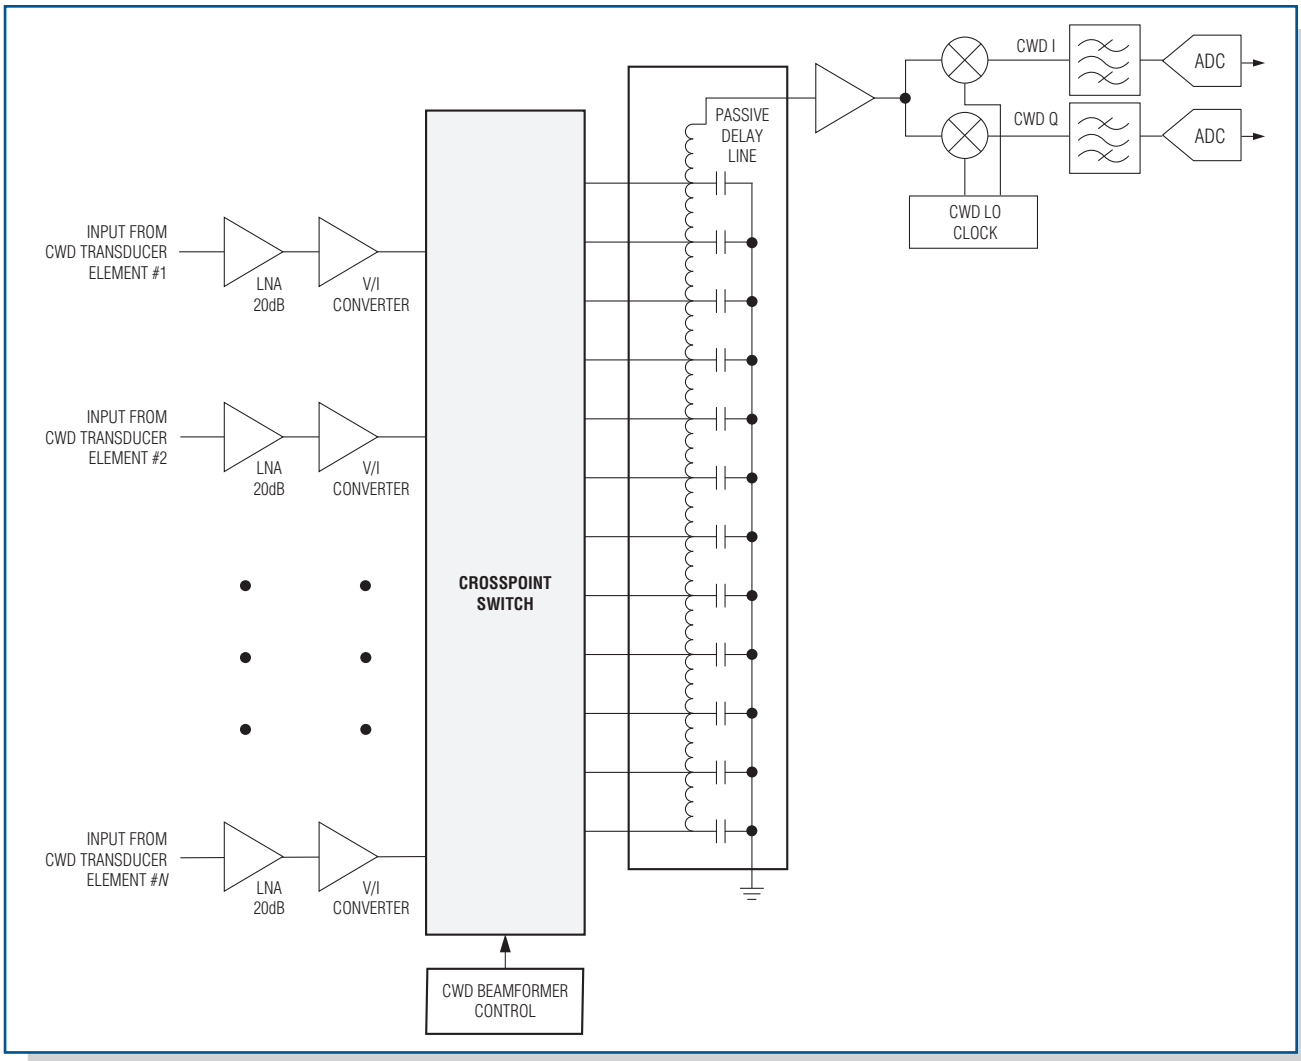

Figure 1. A simplified diagram of a CWD delay-line-based receiver.

beamformer implementations meet the minimal needs of compact ultrasound systems, the performance limitations caused by these issues can be significant.

Mixer-Based CWD Beamforming

A higher performance approach to craft a CWD system employs a CWD mixer/beamformer, shown in simplified form in **Figure 2**. In this implementation, an I/Q mixer is provided for each channel, and the beamforming summation occurs at baseband instead of at RF. The LO phase for each I/Q mixer in this implementation is programmable to one of n = 8 to 16 phases. Changing the phase of the LO changes the phase of the received signal, and beamforming is achieved.

Because the mixers are implemented on a per-channel basis, the SNR requirements for each mixer can be relaxed to 157dBc/Hz at a 1kHz offset. This SNR is still very demanding, but is achievable using bipolar mixers and standard logic families. Since the mixer outputs are currents and passively summed at audio baseband frequencies, the required beamformed CWD SNR can be achieved.

Mixer-Based CWD Beamforming Solutions

In the past, implementation of this superior type of CWD beamformer architecture was not practical for a large number of ultrasound systems because of inadequate integration. This is no longer the case. For less power-sensitive applications requiring uncompromised CWD and imaging performance, there are integrated bipolar octal VGAs with programmable CWD mixer/beamformer channels. The MAX2038 VGA is shown in the receiver lineup in **Figure 3**.

For applications where power and space are at a higher premium, there are newer, more highly integrated and lower power solutions like the MAX2078 in **Figure 4**. This chip contains fully integrated octal receivers incorporating LNAs, VGAs, anti-aliasing filters, and programmable CWD mixer/beamformer channels in a single bipolar IC. Now a wider variety of ultrasound systems can achieve

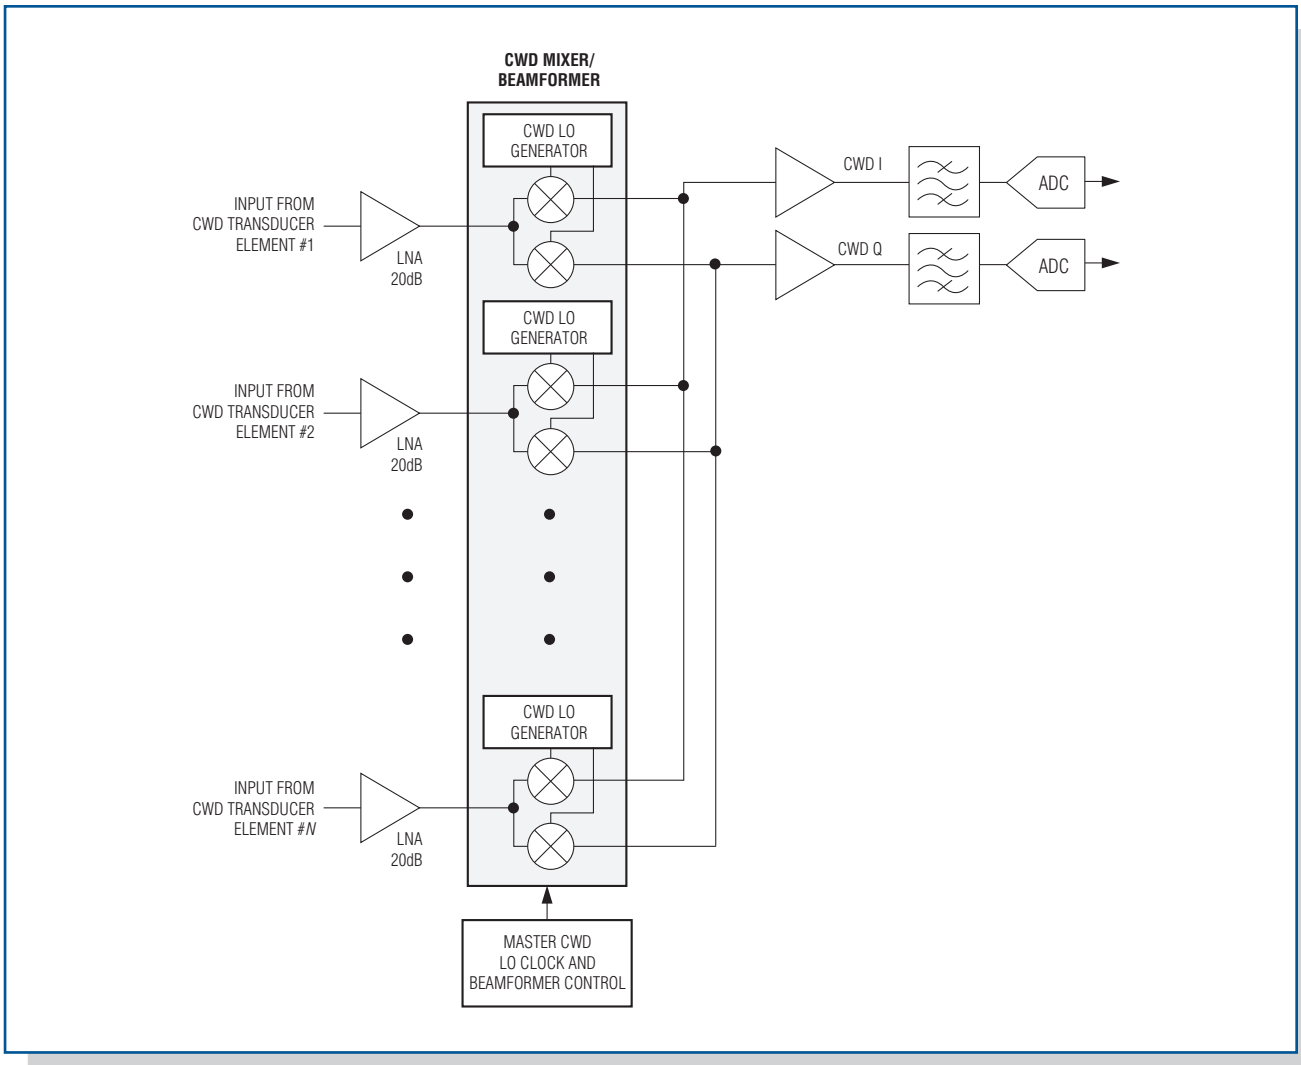

Figure 2. Low-power bipolar LNAs and a CWD mixer/beamformer can implement a simple, high-performance CWD receiver.

excellent CWD performance without the previous limitations associated with delay-line CWD architectures.

It should be noted that an additional potential problem in the implementation of any CWD receiver is the SNR performance of the LNA amplifiers. Many ultrasound system designers have chosen CMOS LNAs to reduce size and power. While these devices may seem appropriate for the application, they can limit CWD performance. This is especially true of amplifiers fabricated in CMOS geometries less than 0.35µm. Circuits implemented in those smaller process nodes typically have higher 1/f noise, and 1/f noise can cause low-frequency modulation of the LNA gain—a very undesirable effect.

A large RF CWD clutter signal passing through an LNA of this type will produce significant low-frequency-modulated noise skirts that can corrupt SNR performance and reduce CWD sensitivity. As a result, low-power bipolar LNAs like the 4-channel MAX2034 are generally preferable for this type of application.

A version of this article was published in Chinese in EEPW *magazine in February 2009.*

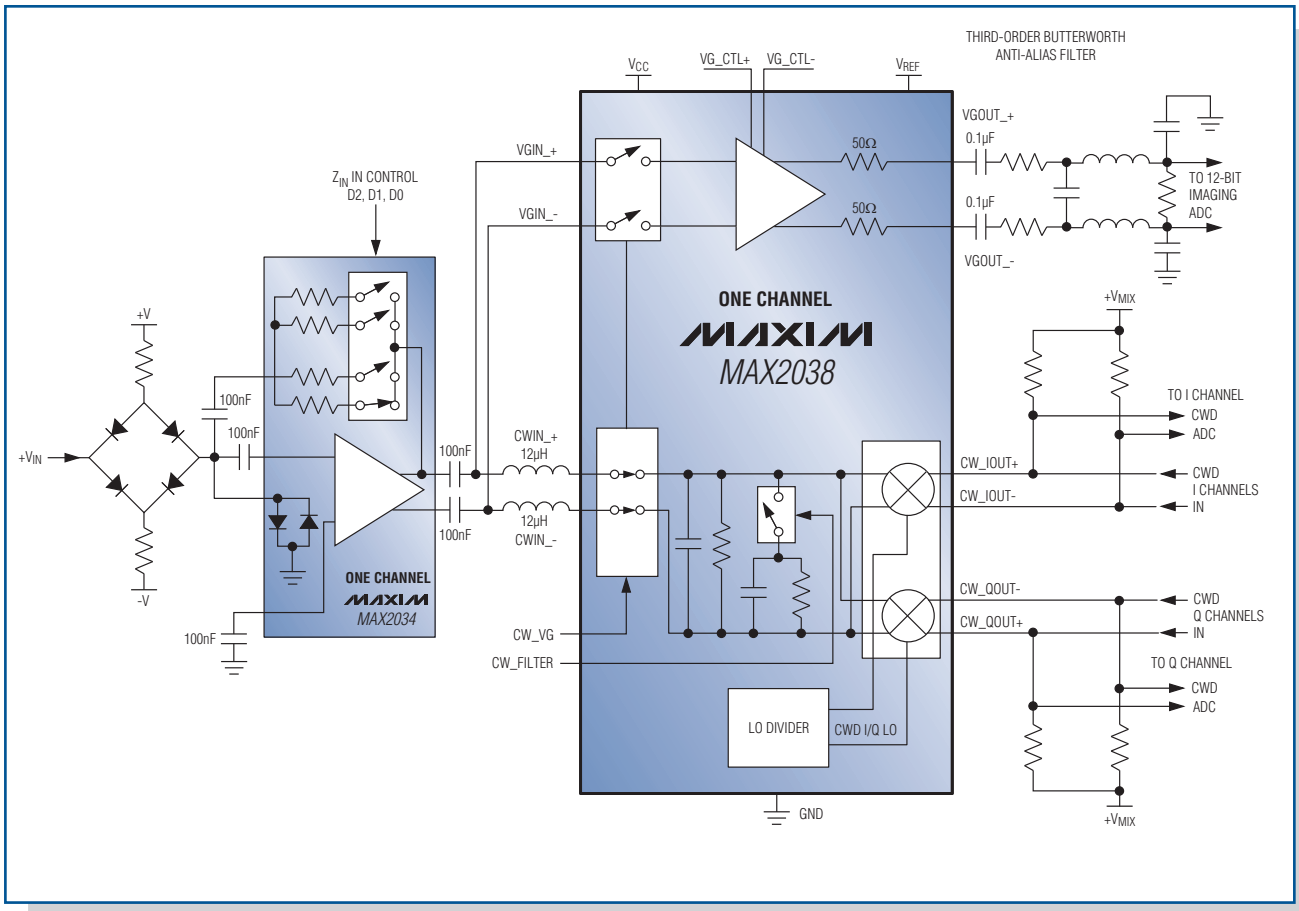

Figure 3. This simplified single-channel ultrasound receiver features the MAX2038 and MAX2034. The MAX2038 integrates eight VGA and CWD I/Q mixer/beamformer channels, and the MAX2034 integrates four LNA channels.

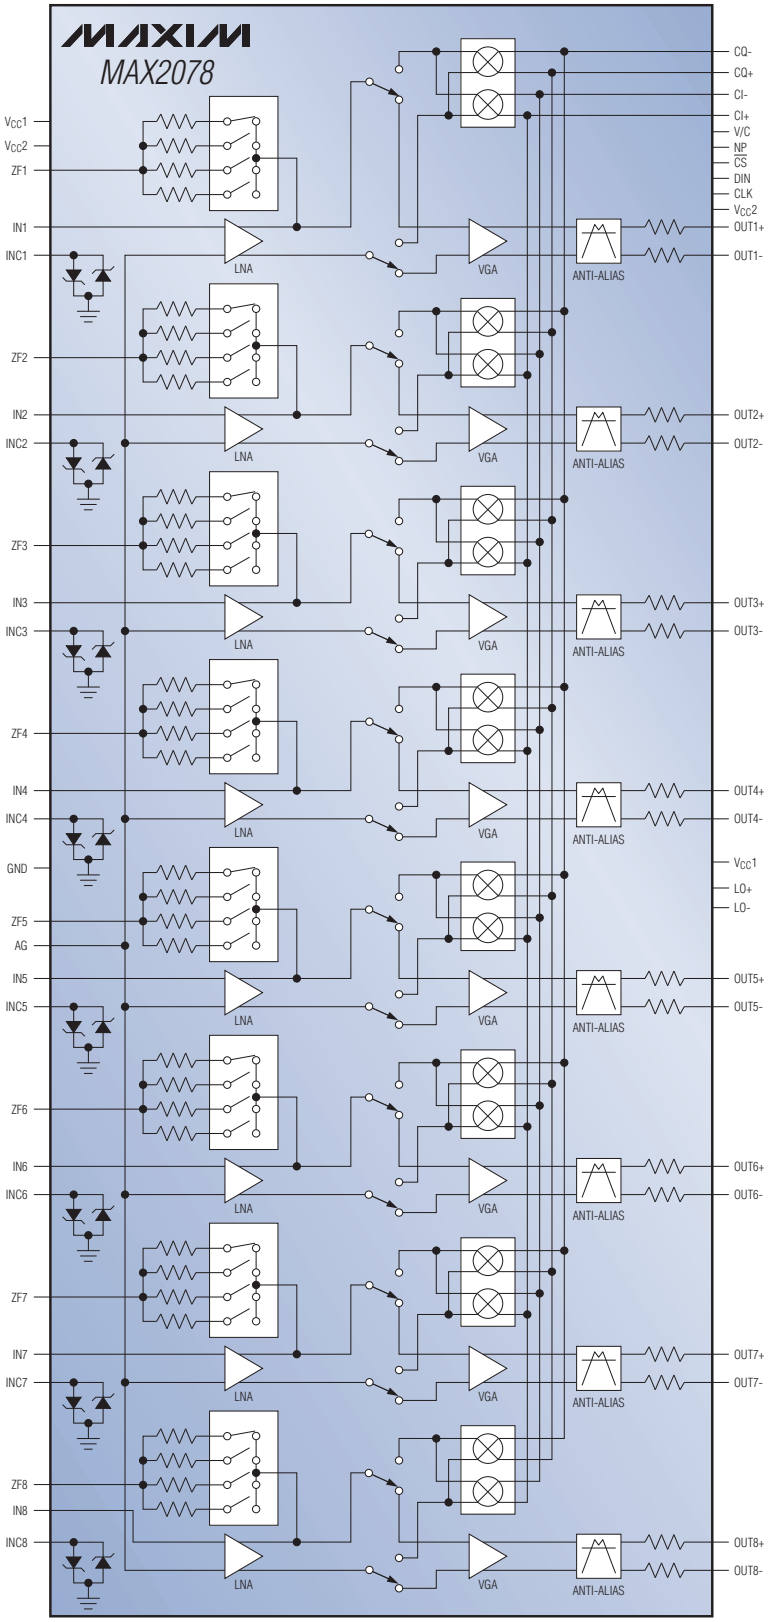

Figure 4. The MAX2078 ultra-low-power, octal ultrasound receiver with CWD beamformer integrates eight high-performance, low-power, ultrasound receive channels, which each includes an LNA, VGA, anti-aliasing filter, and fully programmable I/Q mixer/beamformer.