

Testing Energy Efficient Designs

Primer

Table of Contents

Energy Savings: Good for the Designer, the Consumer and the Environment.....	3	System-Level Energy Management.....	16
An Example of Energy Efficient Design: The Mobile Handset.....	3	Power Dissipation in Digital Circuits.....	16
Power Conversion	4	Dynamic Power Management	17
Measuring Switching Loss	5	Dynamic Voltage and Frequency Scaling.....	18
Making Accurate Power Measurements: Probing Considerations.....	6	Energy Dissipation of Communications Buses.....	18
Eliminating Skew between Voltage and Current Probes	6	PCI Express Low Power Mode.....	19
Eliminating Probe Offset	6	Active State Power Management	19
Measuring Magnetic Power Loss	8	Link Speed and Link Width Changes	21
Magnetic Power Loss	8	Low Power Double Data Rate Memory	22
Magnetic Properties.....	9	LPDDR Power Saving Modes	23
Power Line Measurements	10	Power Saving Modes During Active Operation.....	23
Power Quality	10	Power Saving Modes During Inactive Operation.....	23
Current Harmonics.....	10	Validating LPDDR DRAM Devices	24
Power Consumption.....	11	Standby Power	26
Component-Level Optimization of Power Consumption ..	11	AC Standby Power and Current Measurements	26
Low Voltage Differential Signaling	11	Nanoampere DC Current Measurements.....	27
Single-ended vs. Differential Signaling	12	Conclusion.....	27
Probing Differential Signals.....	13	References.....	28
LED Backlights for Liquid Crystal Displays.....	13		
Driving LED Backlights.....	14		
Measuring Backlight Power.....	15		

Energy Savings: Good for the Designer, the Consumer and the Environment

Today, portable electronics can be found everywhere – homes, hospitals, schools, purses, and pockets. And, the list of portable applications seems endless, covering everything from mobile phones and laptops, digital cameras and MP3 players, to handheld gaming systems and more. With the explosion in portability, consumers have come to expect and demand long battery life. A decade ago, a typical mobile phone would operate for maybe 4 hours on one battery charge. Now, mobile phones will go days on a single charge.

Enabling the increase in battery life have been innovations in energy efficient design. When designs were simpler, the emphasis was on optimizing each hardware component, including the power supply, for highest performance with the lowest power dissipation.

As portable electronics reach higher levels of performance and greater levels of functionality in one product, the embedded systems that drive these products are significantly more sophisticated. These systems are far more complex than single components, consisting of several interacting heterogeneous components. Today's new energy management techniques look at the entire embedded system to drive down overall power consumption.

Line-powered electronics, including servers and personal computers, have also integrated these techniques, driven by business, consumer and government demands to reduce power consumption. As the number of people who plug in and power up electronics increases, and the number of electronics per person increases, the electricity costs faced by many consumers and businesses has dramatically increased. With global power consumption growing at an annual rate of approximately 2%, governments around the world are being forced to provide more capacity on their power grids to meet the growing demand. Improved energy efficiency is the quickest, cleanest and cheapest way to source additional energy.

Throughout the escalating demands and needs, engineers have responded with innovation – from new methods and materials to achieve the same functionality at lower power to creative new ways of managing power usage.

An Example of Energy Efficient Design: The Mobile Handset

A mobile handset offers a good example of the different techniques used today in energy efficient designs. Throughout the years, mobile handsets have undergone a transformation from pure voice-only devices to multimedia powerhouses featuring internet browsing, gaming, photo and music sharing, navigation and much more.

Every new feature added to a mobile handset requires a portion of the often limited energy budget, yet most handsets today rely on battery capacity similar to that of a voice-only design. Achieving such rich functionality has required design engineers to forge new ground in low power design.

The different circuits in the handset often require different supply voltages, necessitating that each circuit has a separate DC/DC converter. Traditional converters used linear regulators and lossy components, resulting in a power conversion efficiency of 50% or less. Today, switch-mode power supplies are the most common architecture, providing efficiencies of up to 90%.

Throughout the design, components are optimized to provide the lowest power usage. Displays that use specialized backlight LEDs provide power savings of up to 80%. The use of Low Voltage Differential Signaling (LVDS) for data transmission saves power while maintaining data integrity.

Techniques to manage energy at the system level have emerged as designs have become more complex. Processors and communication buses where frequency is scaled based on actual need combined with power management that scales down voltage with frequency offer significant reductions in energy consumption.

Finally, standby power has been minimized. It has been estimated that standby power – the power consumed by a product when it is switched off but continues to consume a reduced level of power – represents 10% or more of all residential electricity demand. For battery-powered products like mobile handsets, standby power is a result of bright displays, backlit buttons and instant-on power switches. Even with low-voltage, low-power logic and microprocessor “sleep mode” power management, standby power consumption can significantly limit battery life and the usability of the handset.

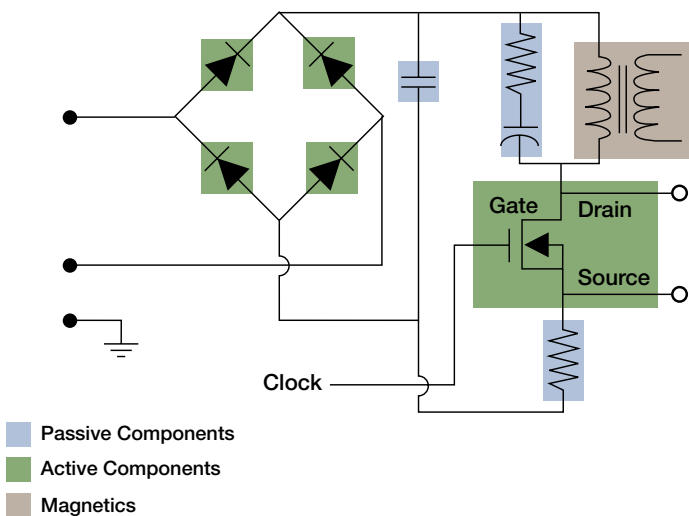


Figure 1. Switch-mode power supply simplified schematic.

Throughout the mobile handset, energy efficient design techniques have been implemented to provide a wide-range of features while maintaining long battery life. This same trend has been applied to other portable and line-power electronics. With these changes, test and measurement techniques have evolved. In this primer, several key energy efficient design techniques and their impact on the validation, debug and characterization test of new designs will be explored.

Power Conversion

Traditionally, the greatest energy loss in electronics has occurred during power conversion. Any line-powered electronic device has an AC/DC power supply which converts electrical power from alternating current (AC) utility power to direct current (DC) power. Furthermore, many electronics – line-powered and battery-operated – have DC/DC converters which convert from one DC voltage to another to accommodate the different supply voltages required by the different circuits within the design. Given that traditional linear power supplies have power conversion efficiencies of less than 50%, significant energy loss has historically occurred during these conversions.

Today, switch-mode power supplies (SMPS) are the prevailing architecture for power supplies, featuring conversion efficiencies of 80% to 90%. The SMPS minimizes the use of lossy components such as resistors and linear-mode transistors, and emphasizes components that are ideally lossless: switch-mode transistors, capacitors, and magnetics, as seen in Figure 1. Validating and characterizing a SMPS often requires measuring switching loss and magnetic power loss to determine the efficiency of the SMPS, and measuring power quality and line harmonics to understand the impact of the SMPS on the power line.

Measuring Switching Loss

The switch-mode transistors used in a SMPS often have fast switching times to minimize energy loss. Since these transistors dissipate very little power in either the On or Off states, the largest source of energy loss for a SMPS occurs during switching. During transitions, energy is lost due to the discharge of diode-stored charge and the discharge of energy stored in parasitic inductance and capacitance.

“Turn-off Loss” describes the loss when the device transitions from ON to OFF. “Turn-on Loss” describes the energy lost when the switching device transitions from OFF to ON. Energy loss during transitions can be expressed by the following equation:

$$E_{transition} = \int_{t_o}^{t_1} v_a(t) \cdot i_a(t) \cdot dt$$

where:

- $E_{transition}$ is the energy loss in the switch during the transition
- $v_a(t)$ is the instantaneous voltage across the switch
- $i_a(t)$ is the instantaneous current through the switch
- t_1 is when the transition is complete
- t_o is when the transition begins

The total energy loss for an entire switching cycle includes the switching losses (turn-on and turn-off losses) and conduction losses. The total loss is given by the formula:

$$E_{Loss} = E_{turn-on} + E_{on} + E_{turn-off}$$

where:

- E_{Loss} is the energy loss in the switch for a switching cycle
- $E_{turn-on}$ and $E_{turn-off}$ are the switching losses
- E_{on} is the conduction loss

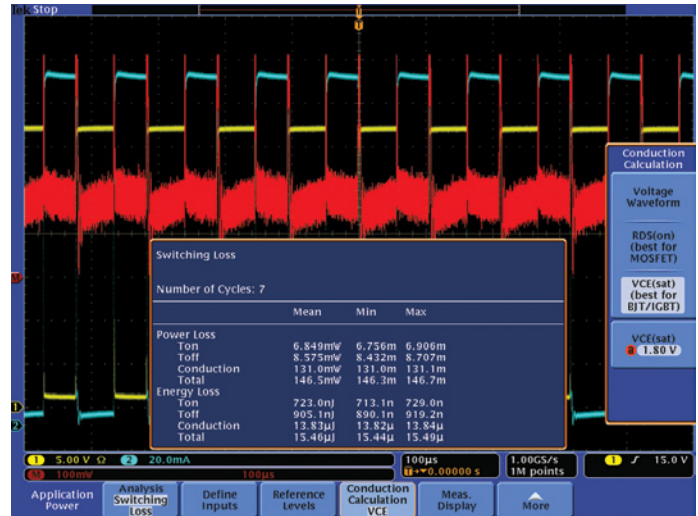


Figure 2. Switching loss measurements including statistics on the Tektronix MSO/DPO4000 Series oscilloscope and DPO4PWR power analysis software.

A proper analysis of these losses is essential to characterize the supply and gauge its efficiency. Figure 2 shows switching loss measurements made with an oscilloscope. By using an oscilloscope with specialized power analysis software, switching losses and conduction losses can be measured across multiple switching cycles to determine device behavior over time. Measurement statistics easily show how the measurement results are changing.

Accurately measuring turn-on and turn-off losses can be a challenge since the losses occur over very short time periods, while the losses during the remainder of the switching cycle are minimal. This requires that the timing between the voltage and current waveforms be very precise and that measurement system offsets are minimized. Making accurate power measurements requires careful consideration of probing.




Figure 3. 9.4 ns skew between voltage and current signals.

Making Accurate Power Measurements: Probing Considerations

Eliminating Skew between Voltage and Current Probes

To make power measurements with a digital oscilloscope, it is necessary to measure voltage across and current through the drain-to-source of the MOSFET switching device or the collector-to-emitter of an IGBT. This task requires two separate probes: a high-voltage differential probe and a current probe. Each of these probes has its own characteristic propagation delay. The difference in these two delays, known as skew, causes inaccurate power measurements and distorted timing measurements.

It is important to understand the impact of the probes' propagation delays on maximum peak power and area measurements. After all, power is the product of voltage and current. If the two multiplied variables are not perfectly time aligned, then the result will be incorrect. The accuracy of measurements such as switching loss suffers when the probes are not properly de-skewed.

Figures 3 through 6 are actual oscilloscope screen views that demonstrate the effects of skew in probes. Figure 3 reveals




Figure 4. With skew, the peak amplitude of the power waveform is 5.141 W.

the skew between the voltage and current probes, while Figure 4 displays the results (5.141 W) of a measurement taken without first de-skewing the two probes.

Figure 5 shows the effect of de-skewing the probes. The two reference traces are overlapping, indicating that the delays have been equalized. The measurement results in Figure 6 illustrate the importance of proper de-skewing. As the example proves, skew introduced a measurement error of 5.3%. Accurate de-skew reduces error in peak-to-peak power loss measurements.

Some power measurement software will automatically de-skew the chosen probe combination. The software adjusts the delay between the voltage and current channels using live signals to remove the difference in propagation delay between the probes.

Also available is a static de-skew function that relies on the fact that certain voltage and current probes have constant and repeatable propagation delays. The static de-skew function automatically adjusts the delay between selected voltage and current channels based on an embedded table of propagation times for selected probes. This technique offers a quick and easy method to minimize skew.




Figure 5. Voltage and current signals aligned after de-skew process.




Figure 6. Peak amplitude has risen to 5.415 W (5.3% higher) after de-skew.

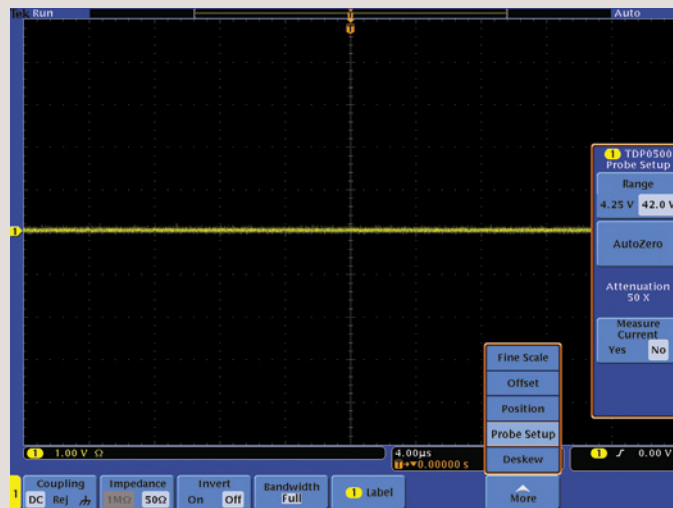


Figure 7. TekVPI probe menu displaying the AutoZero feature.

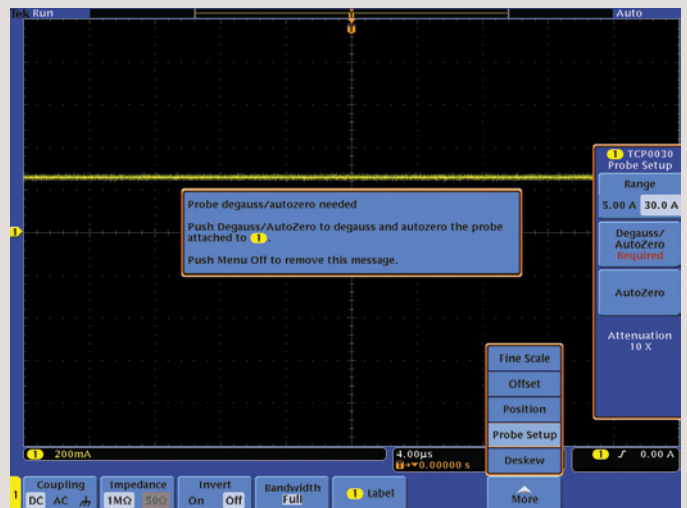


Figure 8. Degauss warning indicator alert from a TekVPI current probe.

Eliminating Probe Offset

Differential and current probes may have slight amplitude offset. This offset should be removed before taking measurements because it can affect accuracy. Some probes have a built-in, automated method for removing the offset.

A probe that is equipped with the Tektronix TekVPI™ Probe Interface works in conjunction with the oscilloscope to remove any DC offset errors in the signal path. Pushing the Menu button on a TekVPI probe brings up a Probe Controls box on the oscilloscope that displays the AutoZero feature, as shown in Figure 7. Selecting the AutoZero option will

automatically null out any DC offset error present in the measurement system.

A TekVPI current probe also has a Degauss/AutoZero button on the probe body. Depressing the AutoZero button will remove any DC offset error present in the measurement system. Degauss removes any residual DC flux in the core of the probe's transformer, which may be caused by a large amount of input current. Tektronix TekVPI current probes offer a Degauss warning indicator, as shown in Figure 8, which alerts the user to perform a degauss operation, preventing the probe from significantly drifting over time and impacting measurement accuracy.



Figure 9. Power loss of a single-winding inductor as measured by the DPO7000 Series Oscilloscope and DPOPWR power analysis software.

Measuring Magnetic Power Loss

Switch-mode power supplies usually employ inductors for filtering and transformers for changing voltage levels because ideally, these elements have low power loss.

Inductors exhibit increasing impedance with frequency, impeding higher frequencies more than lower frequencies. This makes them useful for filtering at the power supply input and the output.

Transformers couple AC voltage and current from a primary winding to a secondary winding, increasing or decreasing signal levels (either voltage or current but not both). Thus a transformer might accept 120 volts at its primary and step this down to 12 volts on the secondary with a proportional increase in current on the secondary. Note that this is not considered amplification because the signals net power does not increase. Because the transformer's primary and secondary are not electrically connected, they are also used to provide isolation between circuit elements.

Magnetic Power Loss

Magnetic power loss affects the efficiency, reliability, and thermal performance of the power supply. Two types of power losses are associated with magnetic elements: core loss in the ferrous core and copper loss in the copper windings.

$$\text{Magnetic Loss} = \text{Core Loss} + \text{Copper Loss}$$

Where:

- Core loss is composed of hysteresis loss and eddy current loss
- Copper loss is due to the resistance of the copper winding wire

The total power loss and the core loss can be quickly derived using information from the core vendor's data sheet and results from an oscilloscope running power measurement software. Copper loss can then be calculated from these two values. Knowing the different power loss components makes it possible to identify the cause for power loss at the magnetic component.

The method for calculating the magnetic component's total power loss depends in part on the type of component being measured. The device under test may be a single-winding inductor, a multiple-winding inductor, or a transformer. Figure 9 shows the measurement result for a single-winding inductor.

Channel 1 (yellow trace) is the voltage across the inductor and Channel 2 (blue trace) is the current, measured with a non-intrusive current probe, through the inductor. The power measurement software automatically computes and displays the power loss figure, here shown as 173.95 milliwatts.



Figure 10. B-H plot for single-winding inductor as measured by the DPO7000 Series Oscilloscope and DPOPWR power analysis software.

Magnetic Properties

For optimum performance, designers generally specify magnetic components using B-H (hysteresis) curves supplied by the manufacturers. These curves define the performance envelope of the magnetic's core material. To ensure key factors such as operating voltage and current are maintained within the linear region of the hysteresis curve during actual operation, it is essential to characterize the magnetic component while operating within the SMPS.

Dedicated power measurement software can greatly simplify magnetic properties measurements with an oscilloscope. In



Figure 11. Inductance as measured by the DPO7000 Series Oscilloscope and DPOPWR power analysis software.

many instances, it is necessary only to measure the voltage and magnetizing current. The software performs the magnetic property measurement calculations for you. Figure 10 depicts the results of a magnetic property measurement on a single-winding inductor. The measurement can also be performed on a transformer with a primary and secondary current source.

In Figure 11, Channel 1 (yellow trace) is the voltage across the inductor and Channel 2 (blue trace) is the current through the inductor. After running the inductance measurement 100 times, the minimum, maximum, and mean measurement values are displayed.

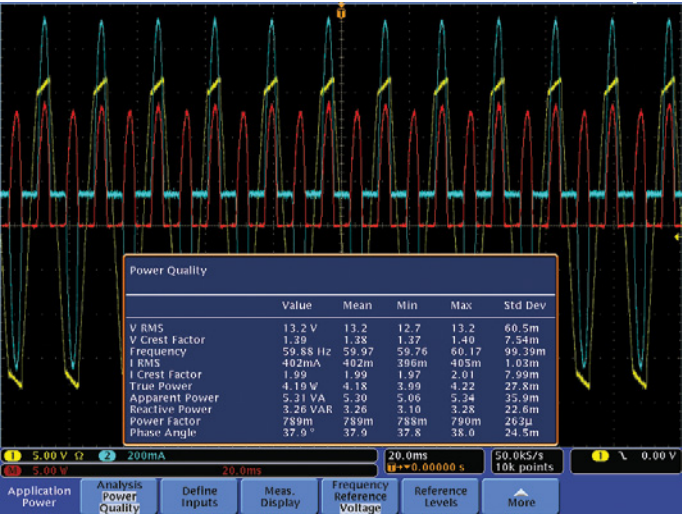


Figure 12. Power quality measurements taken with the MSO/DPO4000 Series oscilloscope and DPO4PWR power analysis software.

Power Line Measurements

For AC/DC power supplies, power line measurements such as power quality and line harmonics are important to characterize the interaction of the SMPS and its service environment. Real-world electrical power lines never supply ideal sine waves. There is always some distortion and impurity on the line. A switching power supply also presents a non-linear load to the source. Because of this, the voltage and current waveforms are not identical. Current is drawn for some portion of the input cycle, causing the generation of harmonics on the input current waveform. Determining the effects of these distortions is an important part of power engineering.

Power Quality

To determine the power consumption and distortion on the power line, power quality measurements are made at the input stage. Most power quality measurements can be automated by full-featured power measurement software running on an oscilloscope, performing lengthy procedures in seconds. By reducing the number of manual calculations, the oscilloscope acts as a very versatile and efficient power meter. Figure 12 shows an example of robust power measurement software.

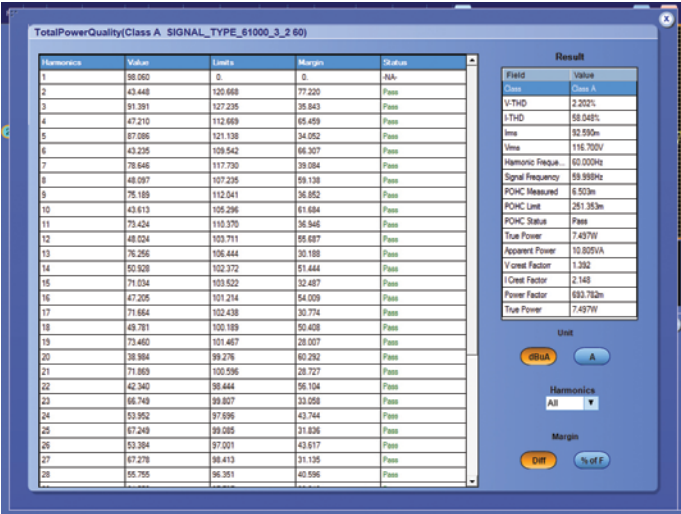


Figure 13. Current harmonics analysis using the DPO7000 Series oscilloscope with DPO4PWR power analysis software.

Current Harmonics

Switching power supplies tend to generate predominantly odd-order harmonics, which can find their way back into the power grid. The effect is cumulative, and as more and more switching supplies are connected to the grid (for example, as an office adds more desktop computers), the total percentage of harmonic distortion returned to the grid can rise. Since this distortion causes heat buildup in the cabling and transformers of the power grid, it is necessary to minimize harmonics. Regulatory standards such as EN/IEC61000-3-2 are in place to specify power quality for a particular non-linear load.

Performing harmonics analysis is as easy as taking an ordinary waveform measurement with the right oscilloscope and power measurement software. Figure 13 shows the result of a harmonic analysis on a power supply's load current. In this case, the software has automatically calculated the current harmonics and determined important values such as Total Harmonics Distortion (THD) relative to the fundamental and RMS values. These measurements are useful in analyzing compliance to standards such as EN/IEC61000-3-2 and MIL-STD-1399. Some software will automatically compare the measurement results to the standard for a quick check of device compliance.

Power Consumption

With power conversion efficiencies now reaching into the 90% range, additional increases are becoming difficult to realize. As conversion efficiencies plateau, attention has turned to minimizing power consumption within the embedded system.

One pathway to reduce power consumption is to improve the energy efficiency of each function within the design. One of the biggest consumers of energy in many electronics is the display. New display technologies that use LEDs for backlight have provided significant energy savings. Techniques to minimize power consumption during signal transmission is another example.

System-level energy management reduces power consumption further by exploiting idle times in system operation to reduce performance of digital systems, thereby saving energy. These energy management techniques are now readily available for microprocessors and other processing elements, as well as for communications buses. The latest specifications for standards like PCI Express®, Low Power DDR, MIPI® D-PHY and M-PHY, USB 3.0, Mobile PCI Express® Module (MXM) and others provide for low power modes.

Validating, debugging and characterizing today's energy efficient designs at times requires new and different test methodologies. The following sections will take a closer look at these changes.

Component-Level Optimization of Power Consumption

Low Voltage Differential Signaling

To lower power dissipation during signal transmission, Low Voltage Differential Signaling (LVDS) is a popular technique, used extensively in communication networks, computers, medical, automotive, and industrial applications.

Traditionally, transmission speeds were increased by lowering the load impedance and increasing the logic device's drive current, allowing the driver to overcome the loading effects of parasitic capacitance and inductance and switch states faster.

As design engineers have dealt with the challenge of reducing power consumption while increasing transmission speeds, LVDS has become common, especially for data rates from a few tens of MHz to about 5 Gb/s. LVDS provides several benefits:

- **Good noise immunity with low voltages.**

One of the simplest ways to improve the noise immunity of digital circuits is to use higher voltage signaling. In terms of signal to noise ratio, the large voltage swings of RS-232, CMOS, and TTL logic are one way to decrease the susceptibility to a given level of ambient noise. However, driven by low-power requirements, especially in battery-powered designs, many circuits need to operate on lower and lower power supply voltages, but without sacrificing product performance or reliability. Differential mode signaling, used in standards such as RS-422, RS-485, CAN, and FlexRay, can provide very good common-mode noise immunity with substantially lower voltage swings. LVDS uses a small-swing differential signal with low power consumption and excellent noise immunity.

- **Low Electro-Magnetic Interference (EMI) when operating at high speeds.**

The need to limit or reduce EMI also drives the trend to lower-voltage power supplies, since the electric field strength is proportional to the voltage swing. The bandwidth of the interference can also be reduced by controlling rise- and fall-times of the signals. Finally, EMI can be controlled with shielding. Differential signaling, such as with LVDS, typically uses much lower voltage swings than most single-ended logic families. Edge speeds are also usually limited. The signal pairs, carrying equal signal currents in opposite directions, are routed in parallel on the circuit board or over twisted pair cables, limiting the far-field interference.

- **High performance at low cost.**

Rising product complexity and density increase the need for simple, low-cost circuit implementations while providing high system performance. LVDS enables simple circuit implementation with a single-resistor load termination and DC coupling. And, because the signal current is always on, being steered between the differential lines, the power distribution circuitry can be simpler than that required for circuits with large voltage switching transients.




Figure 14. Differential signaling.

Single-ended vs. Differential Signaling

Single-ended data transmission uses one signal line for each information channel and a common ground return path. Single-ended receivers interpret the signal's logical state based upon the voltage of the input signal relative to ground. The main disadvantage of this technique is its relatively poor noise performance caused by noise coupling into the input signal and the noise on the common "ground" reference.

Differential signaling uses a pair of equal and opposite signals to transmit a digital signal, labeled V_+ and V_- in Figure 14.

The simplified LVDS circuit shown in Figure 15 shows how switching devices in the transmitter switch the 3.5 mA current source through the terminating resistor in one direction or the other. Since the receiving comparator has a high input impedance, the entire current flows through the termination resistor R_T (typically 100 Ω) resulting in a ± 350 mV voltage at the receiver inputs. This circuit delivers approximately 1.2 mW of power to the load.

The LVDS receiver measures the voltage difference ($V_+ - V_-$) between the signals and then compares the resulting voltage difference signal to a threshold of $0 \text{ V} \pm 100 \text{ mV}$ over a common mode voltage range from 0 V to 2.4 V, providing




Figure 15. Simplified LVDS differential transmitter/receiver schematic.

excellent noise margins and common-mode rejection. Because differential signals can be routed together, they are generally exposed to the same noise sources. By taking the difference between the two input voltages, the receiver cancels common-mode noise from the desired signal.

Differential circuits also radiate substantially less noise than single-ended circuits, since the constant current source is always on, eliminating the on/off switching spikes and the resulting EMI from high-current transistors. The signal amplitudes are also small, and the tightly-coupled conductors minimize the amount of radiated electromagnetic noise.

Finally, common-mode current circulating through the signal return path is minimized, so conducted noise is also lower.

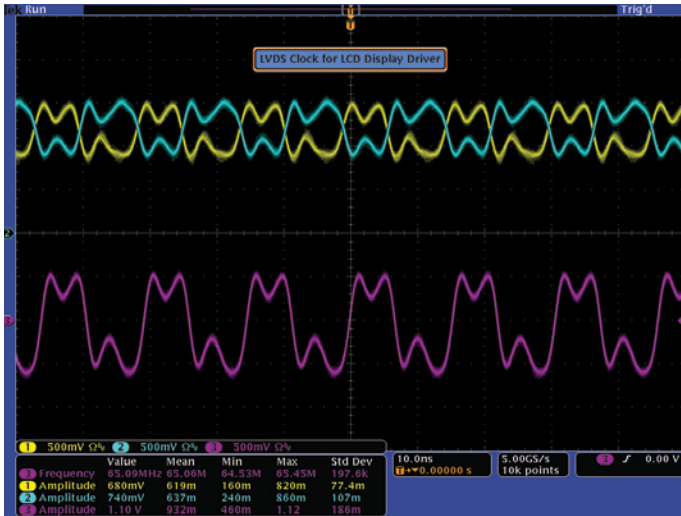


Figure 16. MSO4000 Series oscilloscope display of LVDS clock signals.

Probing Differential Signals

Because information is encoded in the mathematical difference between two physical-layer signals rather than simply the voltage of a signal relative to ground, a different measurement technique is required than for single-ended digital logic.

It is possible to measure a differential signal by making ground-referenced measurements to the two signals, and then mathematically calculating the difference. This technique requires careful matching of probe and oscilloscope channel gains, voltage offsets, propagation delays, and risetimes / bandwidths. Any differences can result in errors in amplitude and timing of the difference signal.

A more accurate and reliable method is to use an active differential probe. A differential probe places a difference amplifier out near the probe tips and the voltage difference is transmitted to the oscilloscope. Since there is only a single oscilloscope connection per channel, differential probes can double the number of signals measured by the oscilloscope.

Figure 16 shows the probing of V_+ and V_- signals with ground-referenced TAP1500 single-ended active probes on channels 1 (yellow) and 2 (cyan), and the differential signal, measured with a TDP1000 active differential probe, on channel 3 (magenta).

Low-power digital technologies such as LVDS meet many of the application requirements for high-speed data communication in modern product design, enabling needed reductions in product size, cost, and power consumption. These technologies, however, require careful consideration to determine the best measurement approach. Differential signaling is easily addressed with specially-designed differential probes, for example the Tektronix TDP Series, which enable the accurate measurement of LVDS signals without requiring manual adjustment for signal path differences.

LED Backlights for Liquid Crystal Displays

One of the most significant improvements in Liquid Crystal Display (LCD) technologies is the use of Light Emitting Diode (LED) backlights to reduce power consumption among other benefits. Backlights are used to illuminate the LCD panel, either from the side, top, or back of the panel.

Until recently, the majority of LCD displays used Cold Cathode Fluorescent Lamp (CCFL) backlights. For these designs, a typical display system includes a power inverter circuit to produce a high AC voltage, initially in the thousands of volts to strike an arc. The power supply then provides a voltage in the range of a few hundred volts to continue the current flow through the lamp's gas, which emits ultraviolet (UV) light. This UV energy is absorbed by the bulb's fluorescent coating, which re-radiates the energy as visible light.

In contrast, an LED backlight uses parallel rows of series-connected LEDs as the light source. For high-fidelity displays, the white backlights may consist of red, green, and blue (RGB) LEDs. For lower-performance applications, the white light can be generated with UV LEDs encapsulated with phosphorous, which re-radiates white light, or with blue LEDs and a color filter.




Figure 17. Backlight LED driver with current-limiting resistor.

LED backlights provide a number of significant advantages over CCFL backlights, including:

- **Thinner, lighter-weight LCD displays.**

LEDs are often smaller than their CCFL counterparts, enabling the use of thinner light guide plates or diffusers. This can result in a significant reduction in display thickness and mass.

- **Higher resistance to mechanical shock and vibration.**

CCFLs are glass tubes and vibration can significantly reduce the life expectancy of the devices.

- **Reduced power consumption.**

Although the luminous efficacy (measured in lumens/Watt) of LED and CCFL backlights are similar, LEDs emit light in only one direction, so the light can be coupled to the display more efficiently. This allows an LED backlight to provide the same brightness with less power. LED backlights also enable designers to vary the brightness of the backlight (or even the color of the backlight, when using RGB LEDs) in individual regions of the display to achieve significant additional power efficiency.




Figure 18. Backlight LED driver with constant-current sources.

- **Lower voltages and reduced electro-magnetic interference (EMI).**

A typical LED backlight might require a DC-DC power supply to provide 30-40 V DC, while a comparable CCFL backlight needs a power inverter to provide 300-400 V AC. Supplying this lower DC voltage over un-shielded cables within a design will likely generate less near-field EMI.

Driving LED Backlights

LED backlights can be electrically modeled as fixed voltage drops equal to the forward voltage of the LEDs (typically 2.5 – 3.5 V per LED, depending upon device and color), with the LED brightness controlled by the current flowing through the devices.

The simplest approach is to drive the backlight with a DC voltage through a current-limiting resistor, as seen in Figure 17. Because there can be significant variations in forward voltages between the LEDs, the circuit can be improved by replacing the simple resistor with a constant-current source for each series connection of LEDs, as seen in Figure 18.

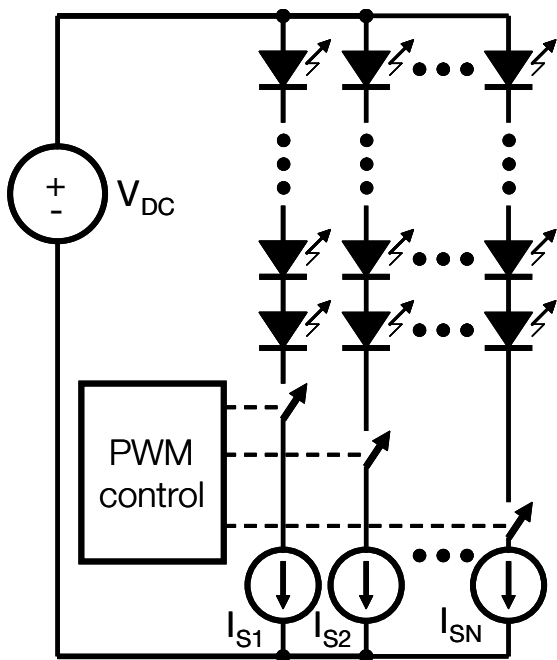


Figure 19. Backlight LED driver with modulated current sources.

If the application requires an extremely bright display, the lowest possible consumption, or the ability to control the brightness over a very wide range, another method is needed.

By exploiting two characteristics of the human visual system, an AC current can be substituted for the DC current. As long as the switching frequency is above about 100 Hz (above the “flicker fusion” rate), the backlight will not be perceived as flickering. Also, the human eye provides some image persistence, tending to “remember” the peaks of brightness for a short period of time, rather than responding to the average brightness over time. This allows the perceived brightness to be controlled by adjusting the duty cycle of the peak current.

A Pulse Width Modulated (PWM) or duty-cycle current control scheme, as seen in Figure 19, provides several advantages over the simple constant-current method. A higher perceived brightness can be achieved at a reduced average power level by operating the LEDs at relatively high peak currents. Because the brightness is controlled by adjusting the duty

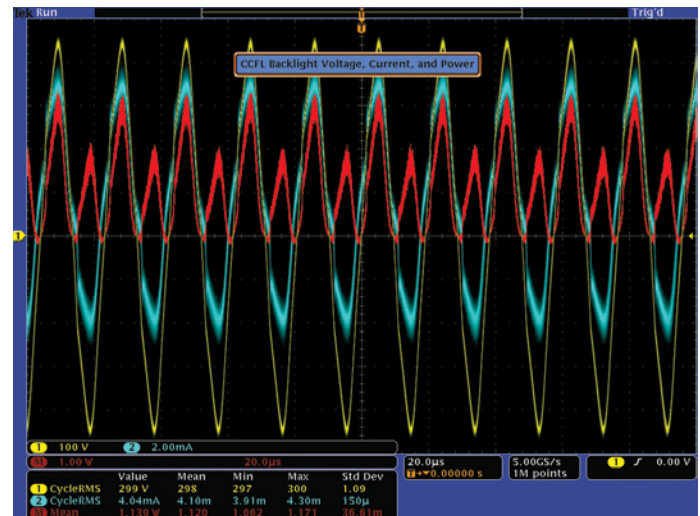


Figure 20. MSO4000 Series oscilloscope display of power measurements on CCFL Backlight.

cycle of the LED current, the average current and therefore the brightness can be controlled without changing the power supply voltage. This allows the designer to optimize the efficiency of the DC-DC converter for the circuit. Finally, because the peak current remains unchanged, the LEDs are operated at the same current levels (e.g. 15-20 mA for white LEDs) and exhibit consistent color.

Measuring Backlight Power

Because of the dramatic differences in backlight driver voltage and current for CCFL circuits and LED circuits, different measurement equipment is needed.

Figure 20 shows the measurement results using an oscilloscope for a CCFL backlight circuit. The AC voltage (yellow waveform) is about 900 V_{pk-pk}, requiring a high-voltage differential probe. In this example, a Tektronix P5200 Series high-voltage differential probe is used. The current levels (cyan waveform) are in the milliAmp range, meaning a sensitive current probe is needed. Here, a Tektronix TCP Series current probe is used. The instantaneous power dissipation in the CCFL tube is shown with the red waveform, representing an average power of about 1.1 W.

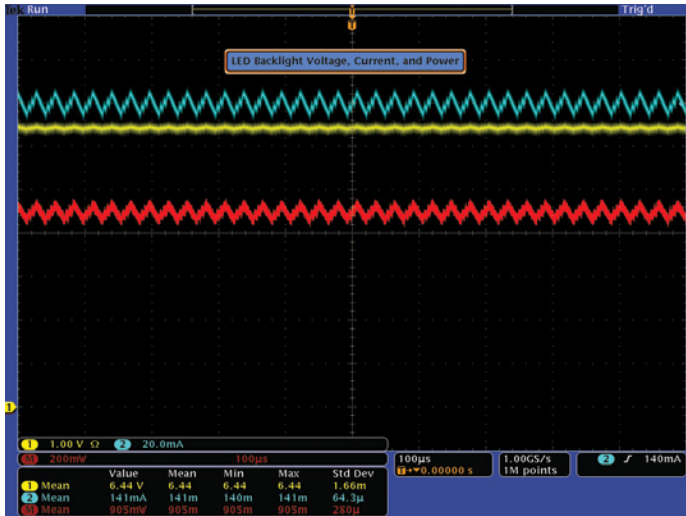


Figure 21. MSO4000 Series oscilloscope display of power measurements on LED Backlight.

In contrast, Figure 21 shows the measurement results when an oscilloscope is used to validate performance of a LED backlight circuit. As shown, the DC voltage (yellow waveform) is about 6.4 V. For this signal, a Tektronix TDP Series differential probe is used. The current levels (cyan waveform) are in the hundreds of milliAmps range. The instantaneous power dissipation in the LEDs is shown with the red waveform, representing an average power of about 900 mW. Compared to the CCFL backlight circuit, a power savings of about 200 mW is achieved.

As the trend for portable, often battery-operated products continues, LED-based LCDs offer significant advantages over their CCFL counterparts. LED backlights simplify product design and enable rugged, lighter displays that provide significant power savings. An oscilloscope provides the tools necessary to evaluate performance of LED backlights, including looking at measurement statistics to see how the device performance changes over time and the sensitivity to measure small changes in voltage and current.

System-Level Energy Management

Energy management techniques, such as Dynamic Power Management (DPM), Dynamic Voltage Scaling (DVS) and Dynamic Frequency Scaling (DFS), look at the system-level operation of a design for opportunities to reduce power consumption. By exploiting idle times in system operation or use case scenarios, power consumption can be reduced by shutting down idle components (DPM) or reducing performance of individual components (DVS and DFS) to provide just enough performance for a specific task. These energy management techniques can be applied to both individual processing elements (CPUs, FPGAs, ASICs), as well as the communication buses that transfer data between them.

Power Dissipation in Digital Circuits

The power dissipated by processing elements is caused by two main effects: static power dissipation that occurs whenever the processing element is turned on and dynamic power dissipation that occurs when computations are performed.

$$P_{PE} = P_{static} + P_{dynamic}$$

The static power dissipation occurs even when no computations are being performed. Two main factors, leakage power, P_{leak} , and bias power, P_{bias} , account for static power usage.

Dynamic power dissipation results from short-circuit power, P_{sc} , and switching power, P_{sw} . The short-circuit power consumption, P_{sc} , is proportional to the supply voltage. The switching power P_{sw} is dissipated when the parasitic capacitors of the transistor gates are charged and discharged during computation tasks.

$$P_{PE} = \underbrace{P_{leak} + P_{bias}}_{P_{static}} + \underbrace{P_{sc} + P_{sw}}_{P_{dynamic}}$$

The switching power, P_{sw} , is currently the dominant source of power usage, accounting for approximately 90% of the total power consumed by mainstream processing elements¹. The common equation for calculating the switching power of a processing element is ²:

$$P_{sw} = \alpha \cdot C_L \cdot V_{dd}^2 \cdot f$$

Where:

- α is a constant representing the switching activity required for the computational task
- C_L is the effective circuit load capacitance; assumed to be a constant determined by the complexity of the design and the circuit technology¹
- V_{dd} is the supply voltage
- f is the clock frequency

The above equation shows two methods for reducing switching power dissipation: reduce the processing element's frequency of operation or reduce the supply voltage. Since switching power is proportional to V_{dd}^2 , the largest energy savings can be achieved by reducing the circuit supply voltage.

In some cases, further energy savings are realized by also reducing the frequency of operation, in addition to the supply voltage. This technique must be carefully applied because computation time increases as the frequency of operation is reduced. Since energy dissipation is determined by power dissipation over time, there will be no energy savings if enough additional time is required to complete the task. However, careful application of both frequency and voltage scaling has been shown to offer higher energy savings than voltage scaling alone⁴.

Reducing the supply voltage will also reduce the leakage power consumption³, improving static power dissipation.

Dynamic Power Management

With DPM, processing elements and communication buses are put in to standby or sleeping modes whenever they are idle. Since it takes time and energy to reactivate these elements, DPM is carefully applied to ensure there are no violations in system operation or actually an increase in power consumption because of reactivation. Even when components are in standby mode, they dissipate energy as determined by the static power (P_{static}) of the component.




Figure 22. Example of a system schedule for an embedded system design. The slack time in the schedule for processing element 2 provides an opportunity for applying dynamic voltage scaling to reduce energy consumption without affecting system performance. Dynamic power management can be used to save energy during the idle times.

Dynamic Voltage and Frequency Scaling

With DVS and DFS, the switching power (P_{sw}) dissipation is reduced by scaling down the supply voltage and operational frequency of a processing element. Since reducing the operational frequency increases the computation time, DFS can only be applied when there is slack time in the system-level operation of the design.

An example of DVS and DFS is shown in Figure 22. In this case, the system schedule allows 30 ms for a task to be completed by processing element 2. However, processing element 2 is over performing and completing the task in 15 ms, leaving 15 ms of slack time. The supply voltage and operational frequency of this element is reduced until the task is completed in 30 ms, as shown in Figure 23, to match the system schedule. By doing so, the energy consumed by processing element 2 is reduced.

DPM can also be utilized in a design to maximize energy efficiency. Even if dynamic voltage and frequency scaling is applied to all components adapting their performance to the actual requirements of the system schedule and minimizing energy consumption, there will still be idle times. DPM can then be used to shut down components that are idle at a specific time for even further energy savings.




Figure 23. Using dynamic voltage scaling, the slack time in the schedule for processing element 2 has been eliminated by reducing the element's frequency of operation and supply voltage, reducing energy consumption.

Energy Dissipation of Communications Buses

In embedded systems with multiple processing elements, communication between elements is essential. With every data transfer over a communications bus, the line capacitance is charged and discharged, drawing current from the I/O pins of the elements. The power dissipated by these currents is given by¹:

$$P_{CL} = \beta \cdot C_{bus} \cdot f_{bus} \cdot V_{tr}^2$$

Where:

- $\beta \cdot C$ represents the switched load capacitance of the bus
- f_{bus} is the operational frequency of the bus
- V_{tr}^2 is the transmission voltage

In the case of communications buses, the transmission voltage can only be reduced by so much because of noise issues. Low voltage communications could be corrupted by noise, causing reliability problems.

Similar to DFS, the operational frequency or data transfer rate of the bus can be scaled down if the system schedule has slack time for bus communication. The bus can also be put in to a standby state during idle times, similar to DPM.

Link State	Description
L0-Active State	All Transactions are enabled and Link is operating in normal mode
L0s – A Low Resume Latency, Energy Saving Standby State	
L1 - Higher Latency, Low Power Standby State	All power supplies and reference clocks are active TLP & DLLP transmission is disabled
L2/L3 – Staging Point for Transition to L2 or L3	TLP & DLLP transmission is disabled
L2 – Aux Powered Link, Deep Energy Saving State	Main power supply and reference clocks are off
L3 – Link Off State	Link is in this state when no power is applied

Table 1. Available Power States for PCI Express.

Low Power DDR DRAM devices as well as several of today's communication buses offer low power modes, including PCI Express®, MIPI® D-PHY and M-PHY, USB 3.0, Mobile PCI Express® Module (MXM) and others. As will be seen in the following applications, debugging these buses can be challenging without the right measurement tools.

PCI Express Low Power Mode

PCI Express is a popular I/O interconnect standard in computers and other performance electronics today. Already in its second generation, it has replaced three other I/O buses - PCI, PCI-X and AGP. The PCI Express specification provides active state power management capabilities that conserve power by putting the bus in one of the power saving link states or dynamically configuring the link width or the link speed based on system requirements at any point in time.

Validation and verification of a PCI Express bus is more complicated because of these features. Problems can arise when the system enters or comes out of one of the power saving link states or when the link width or link speed changes dynamically in response to system requirements.

Active State Power Management

The PCI Express specification provides for managing the link power state based on system requirements. Table 1 lists the different power states in which the link can be present.

In order to maintain synchronization between the transmitter and the receiver, IDLE symbols need to be transmitted over the link when there is no data to be transmitted. The receiver decodes these IDLE symbols and discards them. To save power during these periods, the link can be put in one of the power saving states. The power savings, as well as the time to recover back to the L0 state, increases as the link transitions from the L0 to L3 state.

To understand the complexity added by active state power management to link validation and verification, consider a case where a PCI Express link is in the L0s state and transitions to the L0 state. Immediately after the transition, a TLP configuration write occurs that writes an incorrect value to a register causing the system to crash. To debug this problem, all the transactions that occur during the transition from the L0s to L0 state need to be acquired.

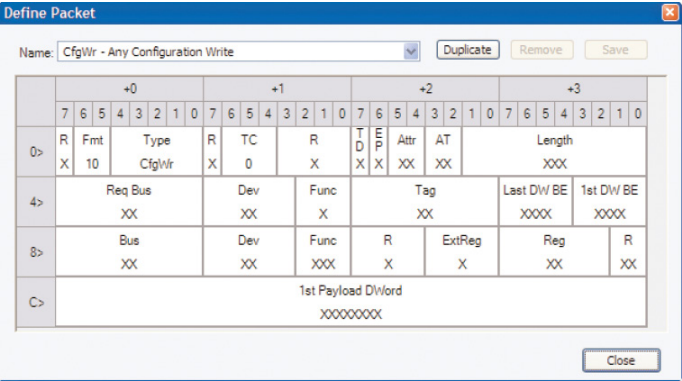


Figure 24. The TLA offers trigger capability with a layout similar to the definitions in the PCI Express standard.

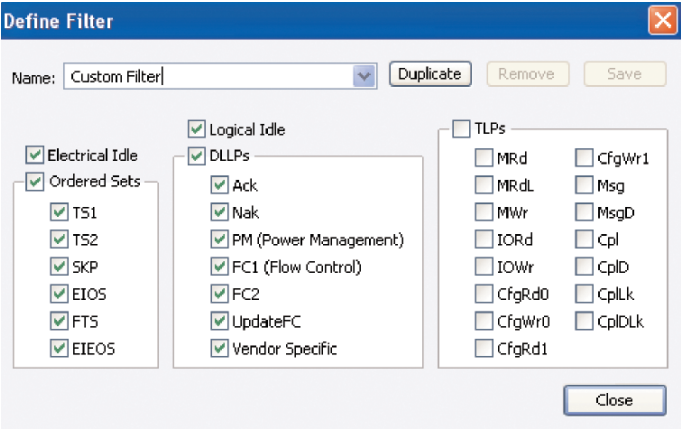


Figure 26. TLA offers filters to remove unwanted data in real time.

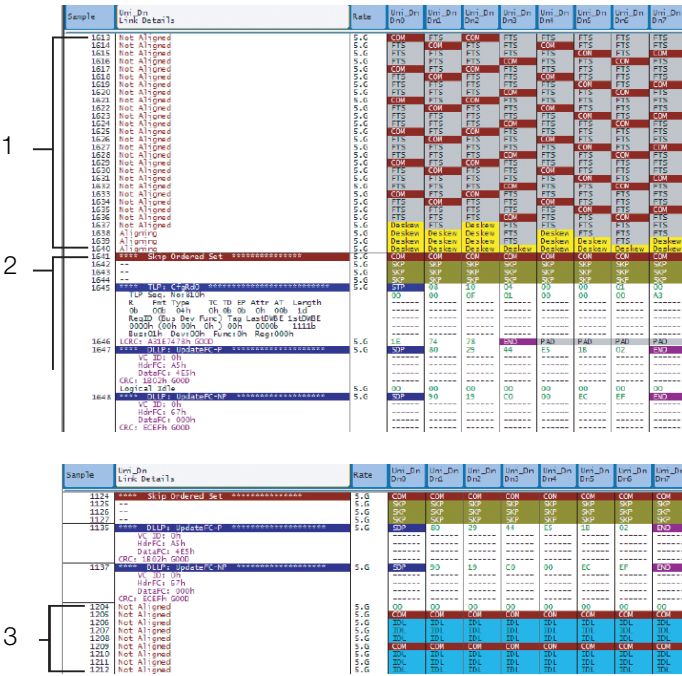


Figure 25. Tektronix TLA Series aligns data across all lanes of the bus.
1. TLA Displays the FTS Packets as the Link Exits Electrical Idle.
2. Data is aligned to the Skip Ordered Set.
3. Link Issues Electrical Idle Ordered Set and Enters L0s.

In this case, the Tektronix TLA Series of logic analyzers is used to troubleshoot the problem. The trigger capability on the TLA7Sxx Serial Module provides a layout similar to the definitions in the PCI Express specification. This enables the user to quickly and easily define conditions to trigger the logic analyzer on a specific event. For example, in the scenario just described, the logic analyzer can be setup to trigger on a TLP Configuration write based on the bus, device and the function number, as seen in Figure 24.

After the trigger is defined, the TLA7Sxx module is able to quickly bit lock and align the data across all the lanes of the bus after observing approximately 12 Fast Training Sequence (FTS) packets as the link exits the L0s and enters the L0 state, as seen in Figure 25. Since the TLA is able to quickly track the change in the link state, it is able to acquire all transactions that occur immediately after the bus enters the L0 state, providing insight in to the cause of the system crash.

The TLA also provides the ability to filter out unwanted data in real time, focusing data acquisition on problem areas (Figure 26). A common use of filtering is when the PCI Express bus is in the IDLE state. A filter can be defined to filter the IDLE symbols in real time, storing only the required data thus using the logic analyzer's memory more efficiently and capturing more relevant data that helps in debugging the problem.



1. End of Training Sequence to Initiate Link Width Change.
2. Links Enters Electrical Idle.
3. Training Sequences Completing the Link Width Change.

The PCI Express specification also provides for dynamically changing the link width or link speed based on the need to conserve power or provide performance in a given situation. These dynamic changes in the link condition can be challenging to debug.



Figure 28 shows the dialog that is used to setup the condition for triggering on a training sequence. In some cases, due to errors in the link, the required trigger conditions may not be found. For such instances, a custom sequence of up to 16 symbols can be built and set as a trigger condition on any particular lane. This triggering feature eliminates the need to manually look through the data to figure out the problems in the link.

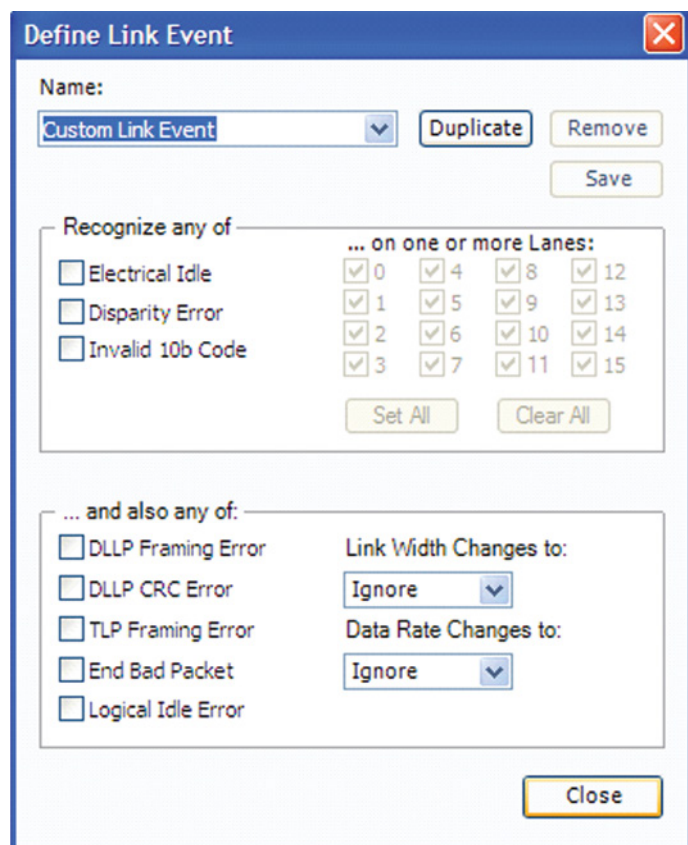


Figure 29. The link event trigger dialog to identify errors in the physical layer.

To identify errors in the physical layer, the TLA also has a link event trigger. For this trigger, the event could be disparity or 8/10b error or error in framing the DLLPs or the TLPs.

The second generation of PCI Express provides power management features to conserve energy, including power saving states and the ability to dynamically change the link width or speed. These new features create challenges when debugging and validating a PCI Express link. Enhanced triggering and filtering capabilities, like those of the Tektronix TLA Series, simplify finding elusive problems during active state management on today's PCI Express links.

Low Power Double Data Rate Memory

Memory can consume a large portion of a design's energy budget. Low Power Double Data Rate (LPDDR) memory, also called Mobile DDR (MDDR), addresses this problem by providing more efficient device operation. This form of SDRAM (Synchronous Dynamic Random Access Memory) operates at 1.8 volts as opposed to the more traditional 2.5 volts, a key contributor to its energy savings.

LPDDR DRAM is most commonly used today in portable electronic devices, where a lower operating voltage translates directly into improved battery life. However in recent years, the advantages of lower-power DRAM devices have become attractive for applications beyond battery-powered mobile devices, as designers of line-powered electronics look at reducing energy consumption.

Reduction in operating voltage is a trend that extends beyond LPDDR to more mainstream memory technologies as well. DDR2, originally specified at 2.5 V, has seen later variants that lower the requirement to 1.8 V with further reductions in development. Similarly, DDR3 was introduced at a supply voltage of 1.5 V but will soon see that reduced to 1.35 V for some new components. LPDDR2, the newest entry in this power-reduction trend, requires only 1.2 V.

Reducing supply voltage is a clear way to reduce power consumption of a memory device. Additional energy savings are possible by reducing the performance of the device. LPDDR, as well as other DDR standards, specifies different power saving modes of operation that reduce performance depending on system requirements.

LPDDR Power Saving Modes

There are two primary categories of power saving features in a LPDDR DRAM. One group of features is used when the device is retaining data. The other group is when the device is not required to maintain data.

Power Saving Modes During Active Operation

The DRAM cell is a capacitor in which data is stored in the form of a charge. The capacitor will leak off charge meaning the contents need to be refreshed regularly during modes of operation in which data must be maintained.

The LPDDR DRAM specification calls for three refresh modes to minimize power dissipation while maintaining the required data states.

- Self refresh is the most basic refresh mode during which a low frequency internal clock is generated to maintain the contents of the DRAM.
- Temperature compensated self refresh automatically modifies the internal refresh clock frequency dependent on the temperature of the LPDDR DRAM device. During cooler operating temperatures, the refresh time can be made longer to save power.
- Partial array self refresh maintains data in only a portion of the DRAM.

Power Saving Modes During Inactive Operation

The power down mode can be used when the LPDDR DRAM device does not need to retain data and when access to the DRAM is not required for several seconds.

As noted earlier, the power consumed by a system is proportional to V_{dd}^2 and the frequency at which the clock is transitioning. Many of the power saving modes specified by the LPDDR standard leverage the frequency component of this power equation. The power saving refresh modes reduce the clock frequency to reduce power consumption. The power down mode is DPM, putting the DRAM in to standby mode during inactive periods. All of these power saving modes primarily impact static power consumption.

Dynamic power consumption can be reduced by optimizing data throughput, allowing the operating frequency of the device to be reduced while still meeting system performance requirements. The ability to do this is one of the key differentiators between different LPDDR DRAM devices.




Figure 30. Measurement Reference Levels for JEDEC Specification Conformance.

Validating LPDDR DRAM Devices

The jitter, timing, and electrical signal-quality tests required to validate memory devices have been specified in detail by JEDEC, the Joint Electron Device Engineering Committee. Parameters such as clock jitter, setup and hold timing, signal overshoot, undershoot, transition voltages etc are included in the comprehensive set of tests described in the JEDEC specifications for each memory technology. These specified tests are not only numerous but can also be complex to measure using general-purpose tools.

An example is measurement reference levels. JEDEC specifies certain voltage reference levels that must be used when making timing measurements. Figure 30 shows a graphic representation of the V_{ih} and V_{il} levels (both AC and DC) that are used for timing measurements on data signals. Note that levels for rising and falling edges are defined differently.

Because of the complexity inherent in the JEDEC-specified measurement methods - reference levels, pass/fail limits, etc. - it can be extremely valuable to have an application-specific measurement utility for DDR test. Using such a utility ensures that your measurements are configured properly and eliminates many hours of setup that would be required using general-purpose tools alone.



Figure 31. DDRA Setup Screen – Step 1 (DDR generation and speed selection).

Option DDRA for Tektronix real-time performance oscilloscopes is a software utility dedicated to automation and setup of measurements for testing DDR devices. The broad set of measurements available in DDRA all conform to the JEDEC specifications, but the user also has the option to customize many settings for measurement tasks on non-standard devices or system implementations.

The menu interface for DDRA has five steps which guide the user through a selection process. Step one of the interface is shown in Figure 31. Here the user selects the DDR generation to be tested (DDR, DDR2, LPDDR etc) and the speed grade of the memory. In addition to the default choices, the user can enter a custom speed setting, making the software easily adaptable to future technology advances, overclocking applications, and more. Once the generation and data rate have been selected, DDRA automatically configures the proper voltage references for measurements. Here again there is a “User Defined” setting, allowing the user to override the JEDEC defaults and enter custom values for V_{dd} and V_{ref} if desired.

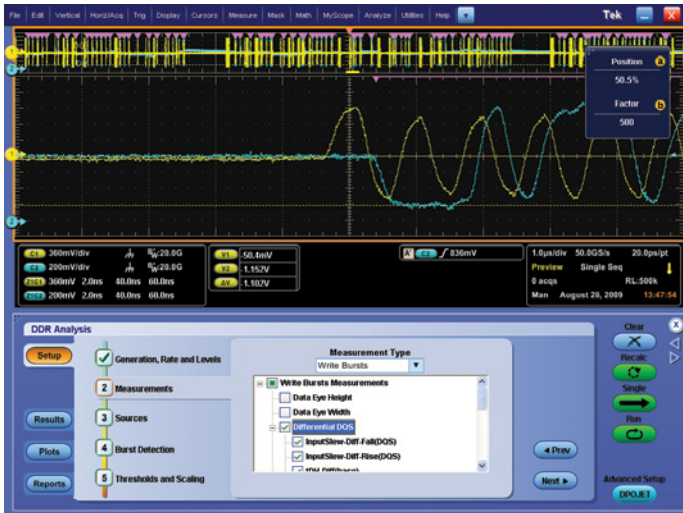


Figure 32. DDRA Setup Screen - Step 2 (measurement selection).

Step 2 allows the user to select which measurements to perform, as seen in Figure 32. The available measurements are grouped into drop-down menu selections according to which signals and probing connections are required. For example, measurements made on the Clock line are all grouped under a “Clock” drop-down menu. Read measurements, Write measurements and Address/Command measurements are similarly grouped into their own drop-down menus so that all measurements requiring a particular probing setup can be easily selected for a single test run.

The remaining steps 3, 4, and 5 in the DDRA menu interface guide the user on how to probe the needed signals and offer additional opportunities for customizing or adjusting parameters such as measurement reference levels.

Once the setup is complete, the oscilloscope will acquire the signals of interest, identify and mark data bursts if needed, and make the selected measurements. The DDRA “Results” panel (Figure 33) shows all measurement results with their statistical population, spec limits, pass/fail results and other data. If desired, a printed report can be generated, with an option to also save the waveform data that was used to make the measurements.

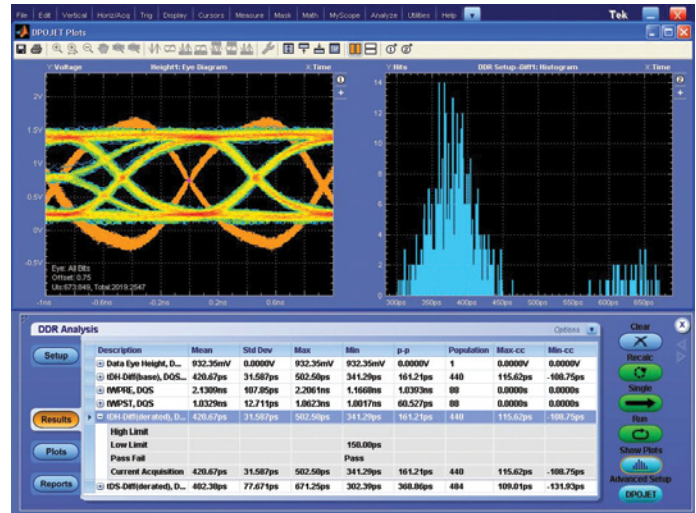


Figure 33. DDRA Results Screen showing two of the available plots.

Because all of the captured waveform data is available behind the measurement results, many options are available to the user beyond the results themselves. If a measurement fails the spec limits, it is possible to identify exactly where in the waveform record the failure occurred, then zoom in on the region of interest to investigate the exact signal details and characteristics at the time of failure. Several tools available in the software make it easy to analyze the captured data and to pinpoint regions of interest.

Today's DDR DRAM devices are evolving to use lower supply voltages to reduce power consumption. Validating these devices requires performing the numerous tests described in the JEDEC specifications, a time consuming and complicated task. With the DDRA software for Tektronix performance oscilloscopes, a broad set of automated measurements are available that conform to the JEDEC specifications, enabling simplified validation of memory devices. With the ability to customize measurement parameters such as voltage reference levels, even the next generation of low power devices can be tested.

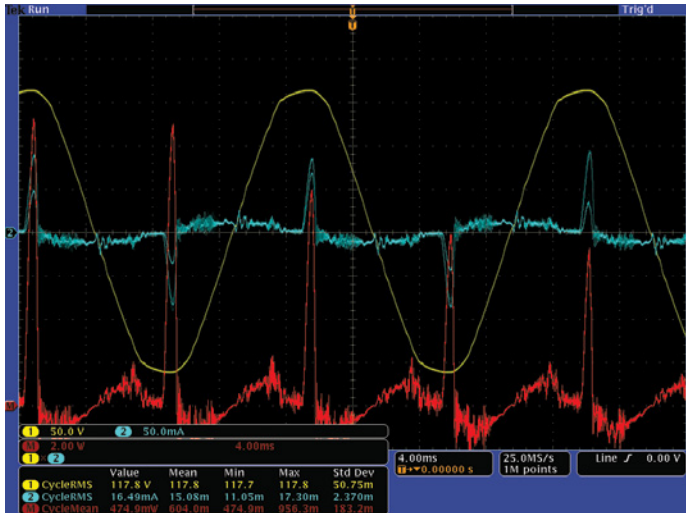


Figure 34. MSO4000 Series oscilloscope display of standby power measurements.

Standby Power

Standby power, also called “phantom load”, refers to the power consumed by a product when it is switched off but continues to consume a reduced level of power. Standby power is a side-effect of products providing “instant-on” capability, consumer electronics with always-on clock displays, as well as power conversion devices such as battery chargers. Because energy is power multiplied by time, even if the power level is low, significant energy can be consumed over long time periods. Standby power consumption affects products in two ways: dissipating AC line power and draining power stored in batteries.

To reduce standby power, engineers must carefully design these products with high-efficiency components and specialized power monitoring software to keep power consumption low while still providing features like “instant-on”. Measuring standby power during design validation and debug requires looking at low level signals over long time periods, a challenge without the right measurement tools.

AC Standby Power and Current Measurements

Measurements of device current when entering and exiting the standby state are usually best done with an oscilloscope and either a current probe to directly measure the AC line current or with a differential probe to measure the voltage drop across a shunt resistor in the AC line. The AC input voltage is most safely measured with a high-voltage differential probe or with a floating oscilloscope.




Figure 35. DMM Series multimeter display of standby current measurements shown in Excel.

An oscilloscope can mathematically calculate instantaneous power by multiplying the current and voltage waveforms. By using an oscilloscope with a high sample rate and long record length, the oscilloscope can use digital signal processing such as the HiRes boxcar averaging available on Tektronix oscilloscopes to provide high-resolution measurements over a long time period. In Figure 34, voltage (yellow waveform) and current (blue waveform) are measured on the input of a LCD monitor for a computer. The red waveform is the calculated instantaneous power. The oscilloscope calculates RMS voltage and current, as well as mean power, for each cycle of the waveform. Measurement statistics are also used to show worst-case variations of these measurements, plus the mean and standard deviation, over 32 consecutive sets of measurements.

Due to low signal levels in standby mode, a digital multimeter (DMM) may provide higher-resolution quiescent current measurements. The AC line current can be directly measured by inserting the DMM in series with the line, or the voltage drop across an external shunt resistor in the line can be measured, as shown in Figure 35. Because the current waveforms, especially in standby mode, can be very non-sinusoidal, it is very important to use a true-RMS DMM. Since the AC RMS voltage tends to be stable, an accurate AC current measurement can be used to reasonably predict the standby power.

DMMs are normally used to display a single numerical value, such as RMS current. However, some advanced bench DMMs such as the Tektronix DMM Series can log and export repetitive measurements, or display a graph of the measurement values over time right on the DMM display with the TrendPlot™ feature.




Figure 36. DMM Series multimeter nanoamp current measurement technique.

Nanoampere DC Current Measurements

For the most sensitive standby current measurements, such as standby battery current drain which is often in the low microamp range, a DMM such as the Tektronix DMM Series with a nanoamp DC current range setting will provide excellent measurement resolution for very low-level signals.

To eliminate the current measurement inaccuracy caused by adding a shunt resistor to the circuit, a current-to-voltage converter circuit is used for the 2000 μA and 200 μA ranges, as shown in Figure 36. The circuit provides a virtual (floating) ground at its input, and the output is a voltage proportional to current I_S ($V_{\text{OUT}} = -I_S \cdot R_F$), providing accurate measurements down to 1 nA resolution.

Whether verifying the specifications for a product, doing pre-compliance testing to industry standards, or estimating battery life, measuring standby power can be challenging. Accurately determining the power consumed by a design in standby mode means measuring a low level signal over a long time period. An oscilloscope with long record length and high sample rate, and sophisticated signal processing, enables today's designer to track many parameters of design performance over time. For focused analysis, a digital multimeter provides higher-resolution quiescent current measurements. Together, they provide a comprehensive set of measurement tools for standby power measurements.




Figure 37. A complete tool set for testing energy efficient designs may include an oscilloscope, logic analyzer, real-time spectrum analyzer, time-domain reflectometer, signal generator, digital multimeter, probes and analysis software.

Conclusion

Energy efficient design techniques have created new and complex test challenges, requiring design engineers to make numerous difficult measurements and to troubleshoot fast-changing signals, complicated protocols, and small changes in voltage and current.

To validate, debug and characterize these designs, a powerful and complete measurement tool set is needed; one that has the performance and time-saving features to address the latest energy saving design techniques. These tools include oscilloscopes, logic analyzers, probes, signal sources, and multimeters. Automated software packages that simplify complicated measurements are also an important part of the tool set.

As designs reach higher levels of energy efficiency, the right set of measurement tools can simplify and speed the often complicated tasks required for validating and debugging today's evolving designs.

References

1. Marcus T. Schmitz, Bashir M. Al-Hashimi and Petru Eles. System-level Design Techniques for Energy-Efficient Embedded Systems. Kluwer Academic Publishers, 2004.
2. Anantha P. Chandrakasan and Robert W. Brodersen. Low Power Digital CMOS Design. Kluwer Academic Publisher, 1995.
3. S. Martin, K. Flautner, T. Mudge, and D. Blaauw. Combined Dynamic Voltage Scaling and Adaptive Body Biasing for Lower Power Microprocessors under Dynamic Workloads. In Proceedings IEEE/ACM International Conference Computer-Aided Design (ICCAD-02), pages 721 – 725, 2002.
4. J. Taylor. Maximize Processor Energy Savings in Media-intensive Mobile Devices with Adaptive Voltage Scaling. Mobile Handset DesignLine, 2007.
5. L. Ternullo. Maximizing a Low-Power Memory Interface IP Subsystem. Chip Design Magazine, 2009.
6. M. Greenberg. LP-DDR: How Much power Will a Low-Power SDRAM Save You? Denali Application Note, 2009.

Contact Tektronix:

ASEAN / Australia (65) 6356 3900
Austria* 00800 2255 4835
Balkans, Israel, South Africa and other ISE Countries +41 52 675 3777
Belgium* 00800 2255 4835
Brazil +55 (11) 3759 7627
Canada 1 (800) 833-9200
Central East Europe and the Baltics +41 52 675 3777
Central Europe & Greece +41 52 675 3777
Denmark +45 80 88 1401
Finland +41 52 675 3777
France* 00800 2255 4835
Germany* 00800 2255 4835
Hong Kong 400-820-5835
Ireland* 00800 2255 4835
India +91-80-30792600
Italy* 00800 2255 4835
Japan 0120-441-046
Luxembourg +41 52 675 3777
Macau 400-820-5835
Mongolia 400-820-5835
Mexico, Central/South America & Caribbean 52 (55) 56 04 50 90
Middle East, Asia and North Africa +41 52 675 3777
The Netherlands* 00800 2255 4835
Norway 800 16098
People's Republic of China 400-820-5835
Poland +41 52 675 3777
Portugal 80 08 12370
Puerto Rico 1 (800) 833-9200
Republic of Korea +822-6917-5000
Russia +7 495 664 75 64
Singapore +65 6356-3900
South Africa +27 11 206 8360
Spain* 00800 2255 4835
Sweden* 00800 2255 4835
Switzerland* 00800 2255 4835
Taiwan 886-2-2656-6688
United Kingdom* 00800 2255 4835
USA 1 (800) 833-9200

* If the European phone number above is not accessible,
please call +41 52 675 3777

Contact List Updated June 2013

For Further Information

Tektronix maintains a comprehensive, constantly expanding collection of application notes, technical briefs and other resources to help engineers working on the cutting edge of technology. Please visit www.tektronix.com



Copyright © 2013, Tektronix. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.

12/13 EA/WWW

54W-23877-2

Tektronix®