Extension MEDIA

Gold Sponsor

MICROCHIP

Engineers' Guide Smartphone, Tablet & Wearables

FPGAs: Good Company in Consumer Mobile

Clear the Mobile Graphics Thicket

Image Stabilization for Tablet Camera Performance

www.eecatalog.com/smartphone

EECatalog

INTERNET OF THINGS
DEVELOPERS CONFERENCE

May 6-7, 2015
Hyatt Regency Santa Clara, CA

Resolving the Technical and Business Challenges of Getting Connected to the Internet of Things

Two full days of keynotes, training sessions and panel discussions. Topics covered include:

- Security
- Adopting legacy industrial connectivity solutions
- Communication gateways
- Differentiating Greenfield and Brownfield
- Embedded cloud computing
- Semiconductor technologies of the future for IoT
- Making money with the Internet of Things
- Networking, LTE equipment, and Femtocells
- Utilizing ultra-low power microcontrollers in IoT devices
- Software development environments for IoT and M2M applications

Plan now. Come to the **Internet of Things Developers Conference 2015!**

IoT-devcon.com

Welcome to the Engineers' Guide to Smartphone, Tablet & Wearables 2015

Contents

To subscribe to our series of Engineers' Guides for embedded developers and engineers, visit:
www.eecatalog.com/subscribe

EECatalog

Extension MEDIA

Engineers' Guide to Smartphone, Tablet & Wearables 2015

www.eecatalog.com/smartphone

Vice President & Publisher
Clair Bright
cbright@extensionmedia.com
(415) 255-0390 ext. 15

Editorial

Editor-in-Chief
Chris A. Ciufo
cciufo@extensionmedia.com

Managing Editor
Anne Fisher
afisher@extensionmedia.com

Creative/Production

Production Manager
Spryte Heithecker

Graphic Designers
Nicky Jacobson Senior Designer
Caldin Seides

Media Coordinator
Yishian Yao

Senior Web Developer
Mariam Moattari

Advertising/Reprint Sales

Vice President & Publisher
Embedded Electronics Media Group
Clair Bright
cbright@extensionmedia.com
(415) 255-0390 ext. 15

Sales Manager
Michael Cloward
mcloward@extensionmedia.com
(415) 255-0390 ext. 17

Marketing/Circulation
Jenna Johnson

To Subscribe
www.eecatalog.com/subscribe

Extension MEDIA

Extension Media, LLC Corporate Office

President and Publisher
Vince Ridley
vridley@extensionmedia.com

Vice President & Publisher
Embedded Electronics Media Group
Clair Bright
cbright@extensionmedia.com

Vice President, Marketing and Product Development
Karen Murray
kmurray@extensionmedia.com

Vice President, Business Development
Melissa Sterling
msterling@extensionmedia.com

Special Thanks to Our Sponsor:

MICROCHIP

FPGAs: Good Company in Consumer Mobile Devices Seeking Differentiation—Fast

Why a companion chip role for today's smaller, less power hungry FPGAs serves smartphones, smart watches, wearables and to-be-imagined mobile consumer devices well.

By Mauri Delostrinos, Lattice Semiconductor

Apple released its first generation iPhone on June 29, 2007; now we are using seventh-generation iPhone 5C and iPhone 5S products that have been available since last year. As I write this, news is breaking that Apple has sent out invites to an event that the company is hosting in Cupertino on September 9, with the tagline "Wish we could say more" to discuss the iPhone 6. The "i" brand—starting with iPods and then iPads before the emergence of the iPhone—has become so successful, first as an aspirational product, then as an ubiquitous item, that 100s of many different "i-__" products have been produced.

Similarly, a quick Internet search identifies that since 2009, Samsung has launched more than 100 smartphones, tablets, phablets, cameras and watches that bear the Galaxy brand alone. Products from other big name players such as LG, Sony and HTC also proliferate in the market, and then there are emerging makers in Asia looking to provide similar products, usually at a lower price premium.

What these well-known market research snippets ably demonstrate is the short lifecycle of any one product generation. Also in evidence: the reduced profitability window that any manufacturer of such consumer items has available to exploit and the essential speed of new product development. Also, with the huge choice on offer to consumers, manufacturers must constantly evolve new features and functions to differentiate themselves in a very competitive marketplace.

The Re-spin Conundrum

Early generation products in this sector used an architecture based on a processor and an application-specific signal processor (ASSP). This is a valid approach, especially in a cost-sensitive market. However, there are two major problems given the need for fast development and product differentiation. Developing any form of ASIC or ASSP-based solution requires a great deal of effort based on early marketing decisions, which can be costly and time-consuming to change in response to evolving competitor and consumer influences. Second, processors have limited ability to handle different I/O, memory types, display and sensor interfaces. Therefore, if the design calls for, say, a different type of sensor, either you'll need to change your processor or manage some form of bridging solution in an ASIC re-spin.

For these reasons—chiefly centering on speed of development and flexibility—manufacturers are now increasingly implementing mobile designs that use FPGAs in a companion chip role. Traditionally, FPGAs would have been considered too big, too expensive and too power-hungry for mobile

Figure 1. CSI2 Bridge Function

10th ANNUAL

MULTICORE DEVELOPERS CONFERENCE

May 6-7, 2015
Hyatt Regency Santa Clara, CA

Resolving the Technical Aspects and Business Challenges of Designing with Multicore Processors

Multicore Use in Application Specific Topics
- Android
- Automotive
- Cloud computing
- Networking and LTE equipment
- Mil/aero
- Real-time and safety considerations
- Security
- Smart phones, and more

Revolutionary Topics
- Manycore technologies
- GPGPUs

Business Topics
- Predictions
- Trends
- What to Avoid

Making you an Expert
- Applying virtualization technology
- Debug and Trace
- Legacy migration
- SoC design methodologies
- Software design strategies
- Performance analysis and optimization

Plan now to attend Multicore DevCon 2015!

www.multicore-devcon.com

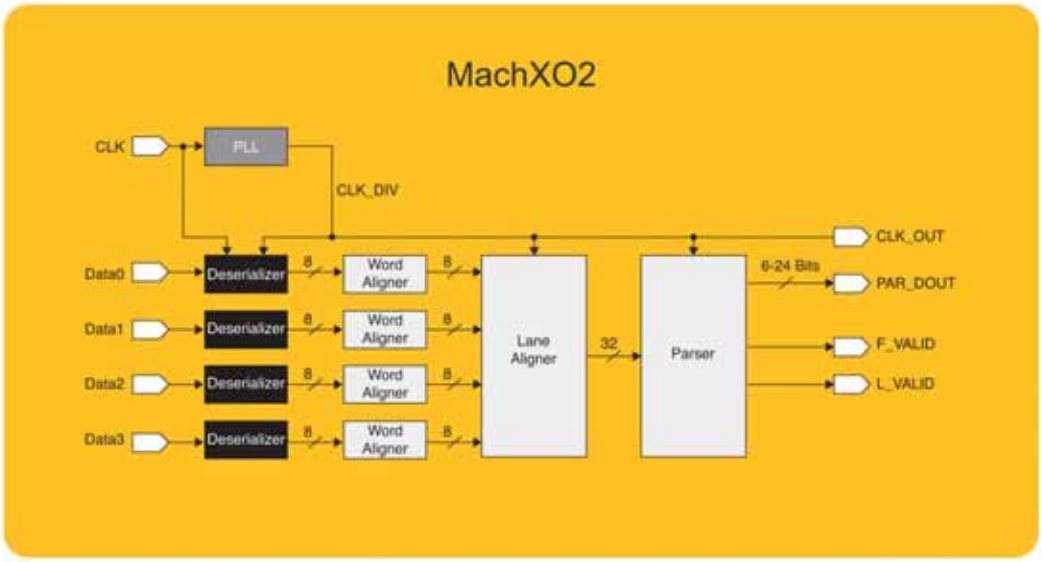

Figure 2. CSI2 Bridge System Block Diagram

consumer applications. However, with the advent of low gate count devices that are as small as 1.4 x 1.48mm, consume as little as 21 µW and cost only 50 cents the picture has changed. For example, a Chipworks teardown identifies a Lattice FPGA inside Samsung's Galaxy S5.

Easily Switching Display Types

Within this context of the arrival of smaller and less expensive FPGAs, Mobile Industry Processor Interface (MIPI) displays are a major application area for FPGAs. The majority of image sensors in the consumer market use the MIPI CSI2 interface. The MIPI has become the interface standard for the majority of components in consumer mobile devices. Camera Serial Interface 2 (CSI2) is the MIPI interface specification focused specifically on cameras. Often the ASSP used in smartphone and especially wearable electronics does not have a CSI2 interface. FPGAs can perform a bridging function to convert from CSI2 to parallel CMOS, enabling the manufacturer to easily switch display types and suppliers as is commercially advantageous.

Figure 1 shows this CSI2 to CMOS parallel function being performed by a member of Lattice's MachX02 FPGA family.

The CSI2 Bridge converts the CSI2 interface to a parallel sensor interface for an ISP. True LVDS input pads on the MachXO2 device handle the 200 mV common mode voltage of the MIPI DPHY high-speed interface. The CSI2 interface from the image sensor can be 1, 2 or 4 data lanes. To keep the FPGA density small, the CSI2 bridge is typically synthesized for a single CSI2 format. In most embedded applications the image sensor is typically configured for a single CSI2 output format at all times. However, multiple CSI2 formats can be supported for "on-the-fly" switching by adding multiple instantiations of the mipi_csi2_serial2parallel NGO in each desired format. Figure 2 depicts the system block diagram.

Support for Custom Functionality Efforts

It is tempting to be skeptical about how much can usefully be achieved using small—both in terms of actual size and gate count—FPGAs. In fact, many features such as IR Tx/Rx Control, Bar Code Emulation, Pedometer, Activity Monitoring, Sensor Pre-Processing and LED Control can be successfully implemented in devices such as Lattice's newly introduced iCE40 Ultra FPGA family. Integration of high current sink LED drivers, multipliers and accumulators optimizes custom function implementation, standard serial interfaces such as SPI & I2C and a whole host more of hardened IP. This ASSP-like integration reduces system power and speeds implementation so designers can spend more time on implementing their custom functionality.

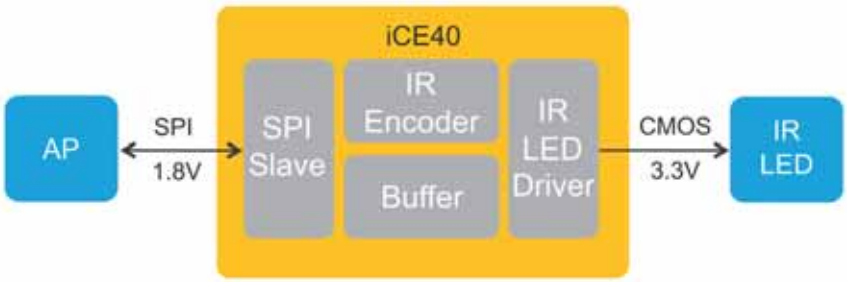

Figure 3. FPGA enables timing-critical IR-LED control.

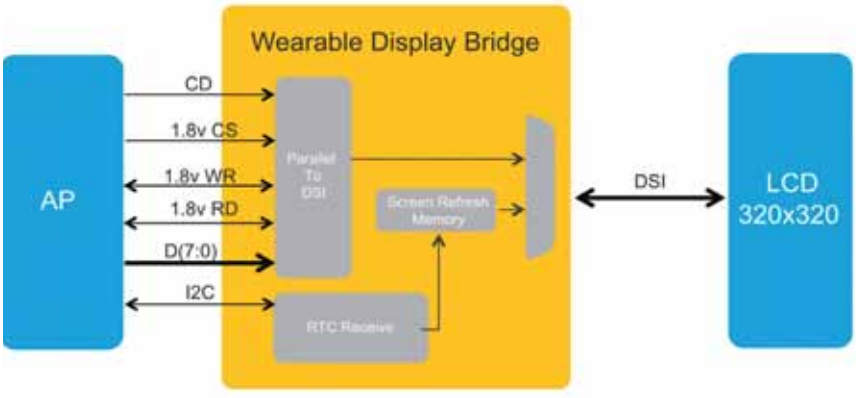

Figure 4. Wearable Display Bridge —Standard Parallel Bus to single to MIPI DSI

Figure 5: Smart Watch—Auto Time Calibration + IO Bridge

FPGAs can take on many other functions in the consumer space. For example, they can be used to capture LVDS video data at high speed and process it using the on-chip sysDSP block and embedded RAM. Another application adds a universal remote function within a smartphone using just a tiny iCE40 FPGA as shown in Figure 3.

Figure 4 details a standard parallel bus to single to MIPI DSI bridge that improves battery life in wearable displays by enabling the applications processor to remain in sleep mode for longer.

Figure 5 shows a design for a smart watch. In this application the FPGA is delivering an auto time calibration + IO bridge function. By so doing the FPGA overcomes the problem that some MCUs do not support 2.5V IOs, needed to interface with GPS modules enabling the watch to automatically reset time when the user travels to a new time zone.

FPGA flexibility enables the imagination to run riot. Smart glasses can include branding held in simple code, and intelligent lighting effects—whether user, context or sensor stimulated—can be simply realized using LEDs.

Designers working in the consumer market who have not considered using field programmable devices before may have concerns about the design methodology. Support is available in the form of tools for the following:

- Design Entry
- Synthesis
- Implementation
- Analysis
- On-chip Debug Hardware Analysis
- Simulation
- Programming
- Deployment

Support for third-party tools is also available.
Therefore, designers can develop, run and simulate their RTL code, then run the code and validate it in hardware. Software developers like this design methodology too as it enables C/C++ to be supported with soft-processor IP doing way with the need for developers to learn a new hardware language in order to implement a solution.

Lattice Diamond supports VHDL, Verilog, EDIF, schematics and multiple implementations. It also features an easy to use GUI, but as a script is sometimes the fastest way to do a task, full Tcl scripting support is also provided.

Conclusion

New generation FPGAs that have been architected for low power and small size, targeting mobile consumer applications, are currently being used in high-volume applications and are proving beneficial in adding flexibility to the development of products such as smartphones, tablets, eBooks and wearable electronics. Alongside devices that are fully uncommitted, FPGAs are also starting to be introduced that incorporate hardened features—memory, I/O, display and sensor interfaces, SERDES —which combine the flexibility of programmability with the efficiency of commonly demanded features. Either way, designers who have previously avoided taking the FPGA path may wish to reconsider.

Mauri Delostrinos is currently Consumer Account Manager at Lattice Semiconductor based in San Jose, California. He is responsible for developing strategic planning for selected Consumer Accounts and for supporting major consumer device manufacturers in the Bay Area and wider. A graduate of the Stanford University School of Business, Delostrinos has also held Applications and Field Applications Engineering roles with Lattice, during which time he participated in Product Definitions for next generation FPGA families and supported key accounts and leading consumer companies in the Silicon Valley.

Clear the Mobile Graphics Thicket

Embedded designers can follow a roadmap to alleviate graphics challenges when developing for mobile medical, smartphones/tablets, gaming, HDTV and more.

By Peter Harris, ARM

With today's mobile devices now offering as much computing power as some desktop computers, many consumers are using these devices as the primary means of consuming multimedia content. While this is great for consumers, it doesn't come without challenges for engineers designing the end devices.

Overcoming common design challenges faced during development is made easier by choosing the right GPU that offers the best power-to-energy-efficiency ratio and development tools to help spot and address potential problems during graphics optimization.

As graphics technologies continue to improve, new visual capabilities are being leveraged across all areas, from HD TVs to mobile gaming devices—even mobile medical devices. Advances in graphics technologies like removing idle time, pipeline throttling and increased shading capability are clearing the way for mobile graphics to continue to change lives.

In an effort to cut down the learning curve with graphics optimization on OpenGL ES, ARM has compiled a roadmap that developers can follow to navigate key graphics challenges including:

Pipelining: Collaborating the CPU and GPU

The first step in successfully starting your next graphics project is to understand the relationship between the application's function calls at the OpenGL ES API and the execution of the rendering operations those API calls require. The OpenGL ES API will act as a synchronous API from the application perspective. Since the API is synchronous, all API behavior after the draw call is specified to behave as if that rendering operation has already happened, but on nearly all hardware-accelerated OpenGL ES implementations this is an illusion maintained by the driver stack. Similar to the draw calls, the second illusion that is maintained by the driver is the end-of-frame buffer flip. Most developers first writing an OpenGL ES application will say that calling eglSwapBuffers swaps the front and back buffer for their application, which again maintains the illusion of driver synchronicity.

The reason for needing to create this illusion at all is in the interest of performance. If we forced the rendering operations to actually happen synchronously you would end up with the GPU and CPU idle at different points during the computing process, which negatively impacts performance.

To remove this idle time, designers can use the OpenGL ES driver to maintain the illusion of synchronous rendering behavior while actually processing rendering and frame swaps asynchronously. By running asynchronously designers can build a small backlog of work for the GPU, allowing a pipeline to be created where the GPU is processing older workloads from one end of the pipeline, while the CPU is busy pushing new work into the other, resulting in the best performance possible.

Removing this idle time is critical to a mobile device's ability to efficiently display the information needed. The resulting

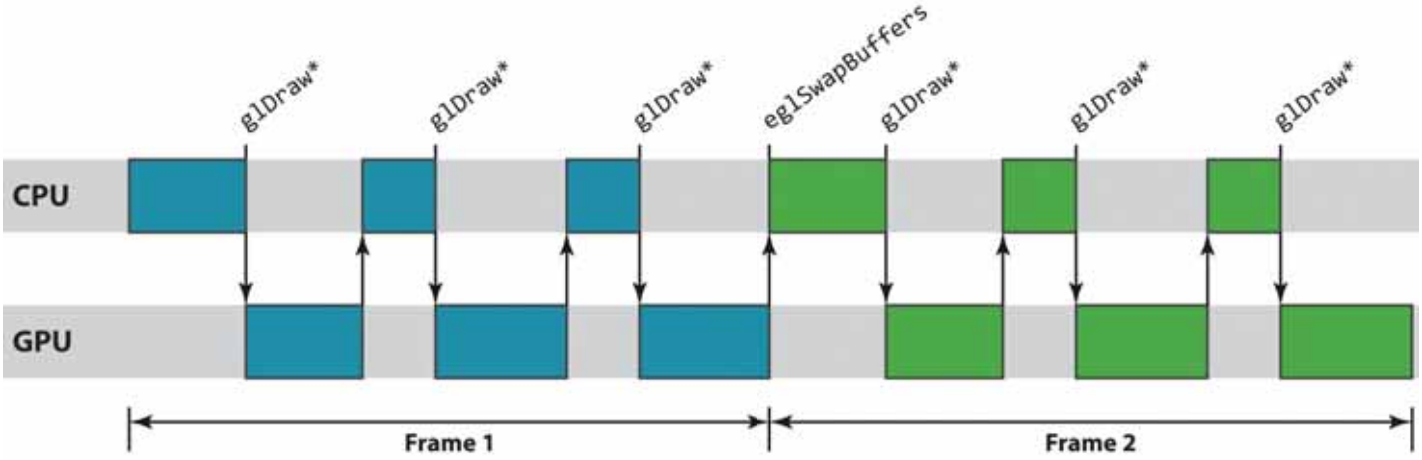

Figure 1. Creating a small backlog of work for the GPU lets the GPU and CPU work as a team.

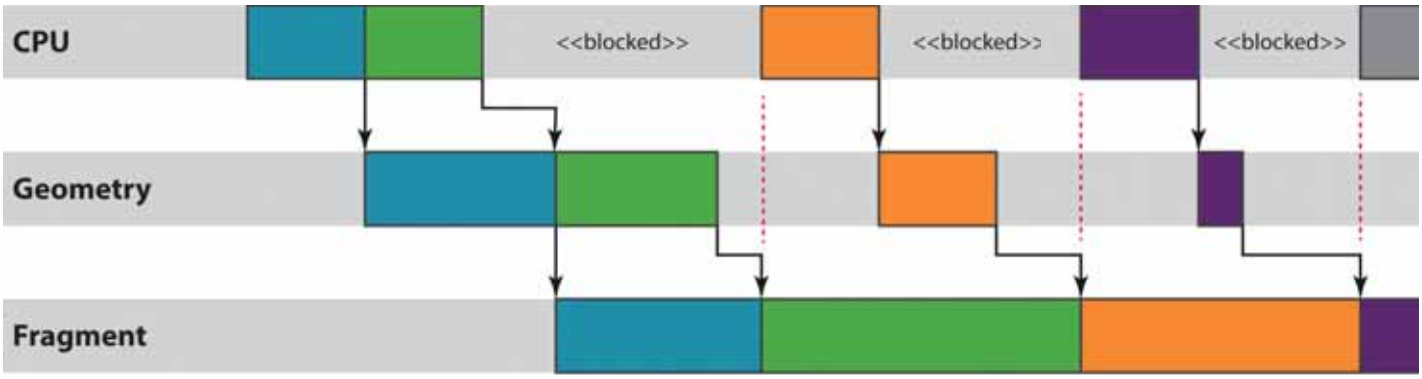

Figure 2. Implementing a throttling mechanism actually slows down the CPU thread periodically and stops it from queuing up work when the pipeline is already full.

smoother frame rate enables trouble-free analysis of images and as a side-effect of the clean pipelining, the optimal selection of both CPU and GPU operating frequencies will help extend battery life – allowing more detailed examinations and a larger number of patients being seen between charges.

Pipeline Throttling

Pipeline throttling is a strategy used to minimize latency between the CPU's work and frame rendering to avoid delays between user touch interaction with their device and the information displayed on the screen. Implementing a throttling mechanism actually slows down the CPU thread periodically and stops it from queuing up work when the pipeline is already full. This mechanism is normally provided by the host windowing system, rather than by the graphics driver itself. SurfaceFlinger — the Android window surface manager – can control the pipeline depth simply by refusing to return a buffer to an application's graphics stack if it already has more than "N" buffers queued for rendering. If this situation occurs you would expect to see the CPU going idle once per frame as soon as "N" is reached, blocking inside an EGL or OpenGL ES API function until the display consumes a pending buffer, freeing up one for new rendering operations.

This same scheme also limits the pipeline buffering if the graphics stack is running faster than the display refresh rate. In this scenario, content is "vsync limited" waiting for the vertical blank (vsync) signal which tells the display controller it can switch to the next front buffer. If the GPU is producing frames faster than the display can show them, then SurfaceFlinger will accumulate a number of buffers which have completed rendering but which still need to be shown on the screen.

The main objective of this strategy is to prevent the GPU from getting too far ahead of what is currently displayed on the screen. By only rendering work which is needed, less power is wasted, which once again extends battery life and allows diagnostic devices to be used in the field for longer.

The "Traditional" Approach

In a traditional mains-powered desktop GPU architecture — commonly called immediate mode architecture—the fragment shaders are executed on each primitive, in each draw call and in sequence. Each primitive is rendered to completion before starting the next one, with an algorithm which approximates to:

```
For each (primitive)
   For each (fragment)
      Render fragment
```

As any triangle in the stream may cover any part of the screen, the working set of data maintained by these renderers is large; typically at least a full-screen size color buffer, depth buffer and possibly a stencil buffer too. A typical working set for a modern device will be 32 bits-per-pixel (bpp) color and 32bpp packed depth/stencil. A 1080p smartphone display therefore has a working set of 16MB and a 4k2k TV has a working set of 64MB. Due to their size, these working buffers must be stored off-chip in a DRAM.

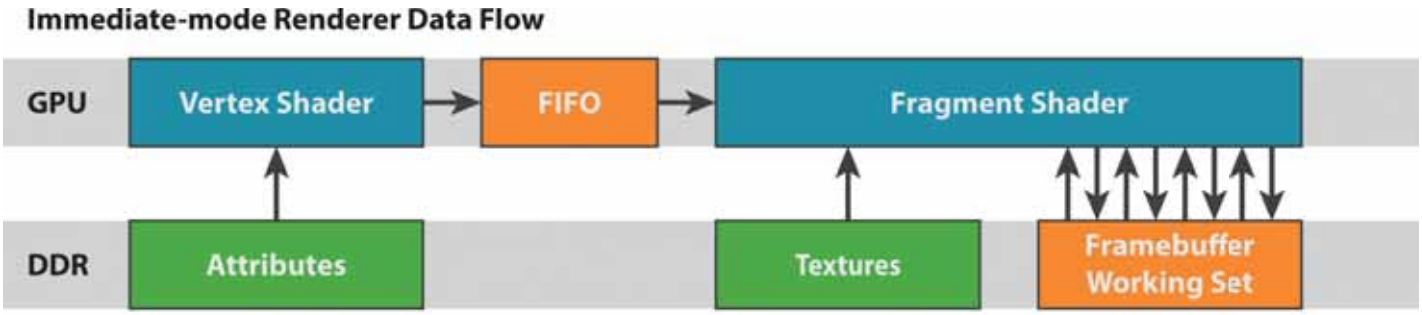

Figure 3. Immediate-mode Renderer Data Flow

Every blending, depth testing and stencil testing operation requires the current value of the data for the current fragment's pixel coordinate to be fetched from this working set. All fragments shaded will typically touch this working set, so at high resolutions the bandwidth load placed on this memory can be exceptionally high, with multiple read-modify-write operations per fragment, although caching can mitigate this slightly.

The ARM Mali GPU Approach

The Mali GPU family takes a very different approach, commonly called tile-based rendering, designed to minimize the amount of power-hungry external memory accesses, which are needed during rendering. The GPU uses a distinct two-pass rendering algorithm for each render target, first executing all of the geometry processing and then executing all of the fragment processing. During the geometry processing stage, the GPUs break up the screen into small 16x16 pixel tiles and construct a list of which rendering primitives are present in each tile. When the GPU fragment shading step runs, each shader core processes one 16x16 pixel tile at a time, rendering it to completion before starting the next one. For tile-based architectures the algorithm equates to:

```
For each (tile)
      For each (primitive in tile)
            For each (fragment in primitive in tile)
                  Render fragment
```

As a 16x16 tile is only a small fraction of the total screen area it is possible to keep the entire working set (color, depth, and stencil) for a whole tile in a fast RAM, which is tightly coupled with the GPU shader core. This tile-based approach has a number of advantages, specifically in terms of giving significant reductions in the bandwidth and power associated with framebuffer data, as well as being able to provide low-cost anti-aliasing in order to improve visual quality.

These benefits make Mali GPUs the ideal technology for mobile medical devices. Not only do they offer a range of performance and energy efficiency enhancements that extend battery life and enable higher screen resolutions, they also are ubiquitous and highly portable. Additionally, Mali GPUs are available all over the world in numerous form factors, and optimized for a range of different markets and requirements.

Why? For passenger entertainment. Think about minivans (shudder) and Suburbans loaded with kids.

Peter Harris is the Mali OpenGL ES Performance Architect at ARM, working on optimization of GPU hardware and software subsystems."

USB 3.0 SSIC: Low-Power Interconnect for Mobile Consumer Applications

Using an existing USB 3.0 software stack with the low-power capabilities of the MIPI M-PHY lets designers meet the increasing performance and battery life requirements of mobile or low-power electronics

Eric Huang and Hezi Saar, Product Marketing Managers, Synopsys, Inc.

While 6 billion people know about USB on the outside of smartphones, cameras, and laptops, only product designers are familiar with how USB is used inside of these products. In laptops, for example, the touchpad, webcam, and broadband modem often use standard USB parts—consuming standard USB power—internally. For example, a USB 3.0 camera will use a USB 3.0 PHY connection (with a cable) to a USB 3.0 PHY on the circuit board (internally). The camera maker can easily implement standard USB 3.0 drivers for both the embedded camera on the PCB system on chip (SoC) and the external camera. The disadvantage to the designer is that using two USB 3.0 PHYs drain the battery because they double the required power over the tiny distance inside the chassis.

Improving Speed, Power and Area with USB 3.0 SSIC and MIPI M-PHY

The USB Implementers Forum (USB-IF) introduced USB 3.0 in 2008 to increase speed and throughput, and soon after they introduced the SuperSpeed Inter-Chip (SSIC) standard for on-PCB communication to reduce power consumption. Less power is required when transmission distances shrink from meters to centimeters or millimeters, as they do inside mobile devices. SSIC uses the MIPI M-PHY to enable products implementing the standard to use as little as 20 percent of the power consumed by a USB 3.0 PHY.

To promote rapid adoption of SSIC, the USB-IF aligned SSIC with the MIPI Alliance's gigabit-speed, on-PCB, chip-to-chip PHY called the MIPI M-PHY. M-PHYs consume lower power and offer greater flexibility than USB 3.0 PHYs. M-PHYs can come in three speeds, called Gears. Gear1 operates at 1.25 or 1.45 Gbps, Gear2 at 2.5 to 2.9 Gbps, and Gear3 up to 5.8 Gbps. In addition, M-PHYs can have 1, 2, or 4 lanes. Each lane has x pins, so two lanes have 2x pins and four lanes have 4x pins. These lane configurations offer flexibility to run either in multiple parallel lanes at slower clock speeds to save power, or to run at faster speeds but consume fewer pins. Since many SoCs are pin and/or pad limited, designers often choose the faster Gear3 standard to save pins. A one-lane MIPI M-PHY has 16 pins. On the other hand, a standard USB 3.0 PHY has at least 15 pins including USB 2.0 D+ and D-; USB 3.0 Tx+, Tx- Rx+, Rx-, power, and ground pins.

The MIPI Alliance and USB-IF worked together to standardize the interface between USB 3.0 controllers and MIPI M-PHYs. According to the standard, the USB 3.0 controller uses a standard PIPE interface, which is the same interface for the USB 3.0 path to a USB 3.0 PHY. While the PIPE interface is preserved, the system still needs an interface to a standard M-PHY. The M-PHY v2.0 specification defines the SSIC interface to the M-PHY as the Reference M-PHY Module Interface (RMMI). The logic bridge between the USB 3.0 controller and the M-PHY is called the PHY Adapter. While it sounds simple, the PHY Adaptor is complex as it must synthesize and operate with the controller and the PHY. It must support USB 3.0 power savings modes (U1, U2, U3, and U4) while supporting 1, 2, or 4 lanes and/or Gear1, 2, or 3 speeds.

Using High-Speed Gear3 M-PHY with USB 3.0

The MIPI M-PHY v3.0 specification defines low-power implementation for chip-to-chip connectivity, including several high-speed gears that match the burst speed needed by a given application, as well as low-speed pulse width modulation gears that are used mostly for control. The specification also defines a variety of low-power modes that the link can utilize to enter and exit during long or brief idle times and rapidly get back to burst mode.

A MIPI M-PHY is about 50 percent smaller than a USB 3.0 PHY and consumes significantly less power, especially in Gear1, 1 lane operation. In this configuration, a MIPI M-PHY consumes

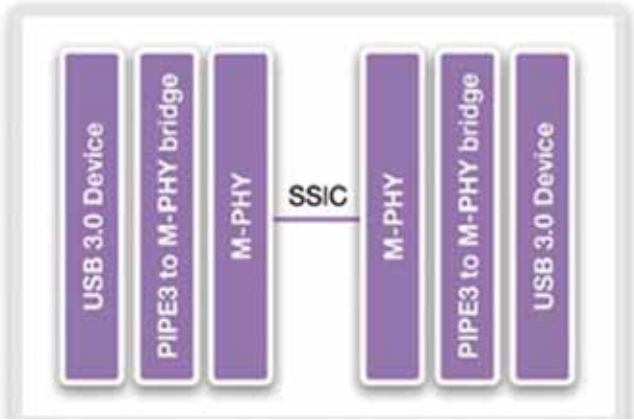

Figure 1: Standardized SSIC interface between USB 3.0 and MIPI M-PHY

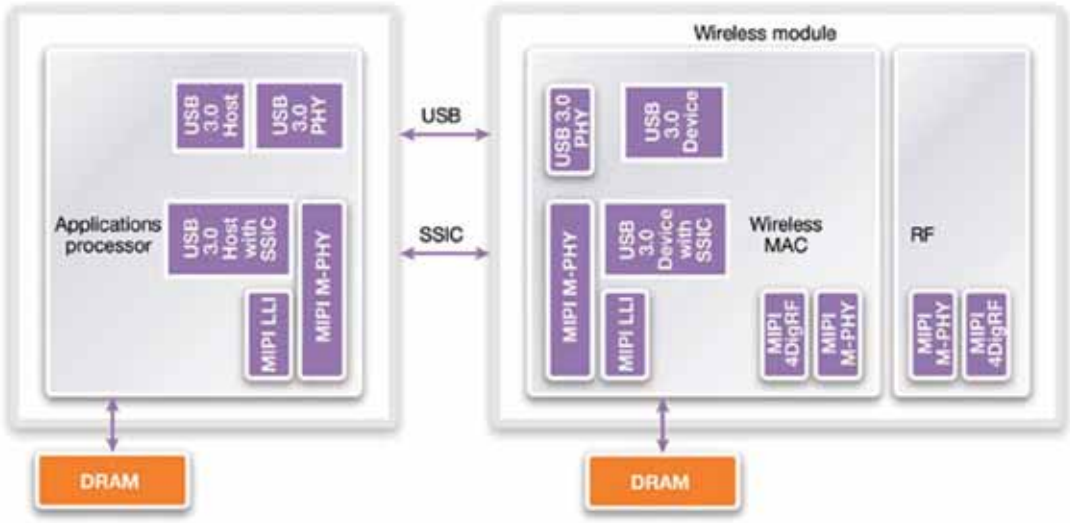

Figure 2: Example of using a single M-PHY with two controllers (USB 3.0 and LLI)

only 20 percent of the power of a USB 3.0 PHY. For two devices connected on PCB, this 80 percent power reduction during active operation at the system level is significant for portable devices. Part of this power savings is due to the smaller PHY, and part is because the Gear1, 1 lane M-PHY data rate is only 1.25 or 1.45 Gbps. As a USB 3.0 PHY always operates at 5 Gbps, the M-PHY allows for the lower data rate and power savings.

The high-speed MIPI M-PHY, working in conjunction with USB 3.0 SSIC, is tailored for mobile applications and is becoming a popular physical layer solution. With up to 5824 Mbps bandwidth, the High Speed Gear3 serialization speed meets devices' high bandwidth requirements. The M-PHY is designed to accommodate the intermittent nature of inter-chip communications and employs burst operation to toggle between data transmission and power saving states, effectively reducing power consumption.

Multiplexing with M-PHY

In most designs in the near future, a smartphone/tablet application processor SoC with a fully integrated USB 3.0 SSIC controller and M-PHY will connect on the PCB to a modem or WiFi SoC. The WiFi SoC also has an M-PHY and USB 3.0 controller. A second USB 3.0 controller and USB 3.0 PHY may be used for external USB connections. On an applications processor, this may be an external port for connecting to a USB 3.0 flash drive. In addition, the single M-PHY can be multiplexed with other MIPI functions.

In Figure 2, a single M-PHY is used with an LLI controller to allow the baseband of a wireless device to use the RAM for the applications processor. By multiplexing a single PHY with two digital controllers, designers save the area of an extra M-PHY, as long as only one digital controller is working at any one time. A MIPI M-PHY may also be MUXed with a PCIe controller to implement M-PCIe.

Using a MIPI M-PHY with SSIC controllers can result in up to an 80 percent power savings over USB 3.0 PHYs. The SSIC standard improves power efficiency while reducing area and maintaining throughput and preserving SSIC software compatibility. This makes SSIC attractive to designers of smartphones, tablets, and the wireless products they connect to, as SSIC offers bandwidth and power advantages in these highly competitive markets.

Summary

A fully integrated USB 3.0 SSIC controller and M-PHY enables low-power, efficient connectivity on a PCB between a smartphone/tablet application processor to a modem or WiFi SoC. Using an existing USB 3.0 software stack with the low-power capabilities of the MIPI M-PHY enables designers to meet the increasing performance and battery life requirements of mobile or low power electronics.

Eric Huang worked on USB at the beginning in 1995 with the world's first BIOS that supported USB keyboards and mice while at Award Software. After a departure into embedded systems software for real-time operating systems, Eric returned to USB cores and software at inSilicon, the leading supplier of USB IP in the world. inSilicon was acquired by Synopsys in 2002. Eric served as Chairman of the USB On-The-Go Working Group for the USB Implementers Forum from 2004-2006.

Eric Huang received an M.B.A. from Santa Clara University and an M.S. in Engineering from University of California Irvine, and a B.S. in Engineering from the University of Minnesota. He is a licensed Professional Engineer in Civil Engineering in the State of California

Hezi Saar serves as a staff product marketing manager at Synopsys and is responsible for its DesignWare MIPI controller and PHY IP product line. He brings more than 15 years of experience in the semiconductor and electronics industries in embedded systems. Prior to joining Synopsys, from 2004 to 2009, Saar served as senior product marketing manager leading Actel's Flash field-programmable-gate-array (FPGA) product lines. Previously, he served as a product marketing manager at ISD/Winbond and as a senior design engineer at RAD Data Communications. Saar holds a bachelor of science degree from Tel Aviv University in computer science and economics and an MBA from Columbia Southern University.

LPDDR4: Meeting the Power Neutrality Challenge in Mobile Handsets

JEDEC has defined the fourth generation of low-power DDR (LPDDR) that can help developers achieve power neutrality in handset applications, as well as improving performance and cost.

By Daniel Skinner, Micron Technology, Inc.

As new capabilities and functionality are added to mobile handsets, they become more compute-intensive. For example, in higher-resolution cameras and screens, CPUs must perform more work and memory devices must move more data, faster. Consider ultra-high-definition (UHD) video capture, which is currently a differentiating feature in high-end handsets. To support UHD video, a 64-bit system requires 25.2 GB/s of peak bandwidth.

On the surface, UHD may seem to be of limited value—after all, handsets don't have 4K displays and, as yet, few users take advantage of features like streaming UHD video to a living room TV.

UHD becomes more compelling when users want to capture moments in real time, where the most immediately available camera is often on a handset. These moments are then shared across the Internet, sometimes viewed on devices like tablets, laptops and desktops that have UHD resolution-capable screens. No one wants their priceless moments to be low-quality or pixelated. Thus, even though a handset does not directly display UHD, many users consider it important that their devices are capable of capturing, and therefore processing, that number of pixels.

The Power Neutrality Challenge

Next-generation handsets must perform more work than earlier-generation devices without decreasing battery life. Users have come to expect extended battery life and demand that the next-generation handsets maintain the same battery life (at a minimum) as their predecessors—even when they have more/improved functionality.

This concept, called "power neutrality," is a major design consideration for mobile devices. For a next-generation device to perform more work using the same amount of power, the energy required per bit processed must be cut in half across the entire system. This can be extremely difficult given that handsets depend on multiple technologies that are not advancing quickly enough to keep pace with power neutrality.

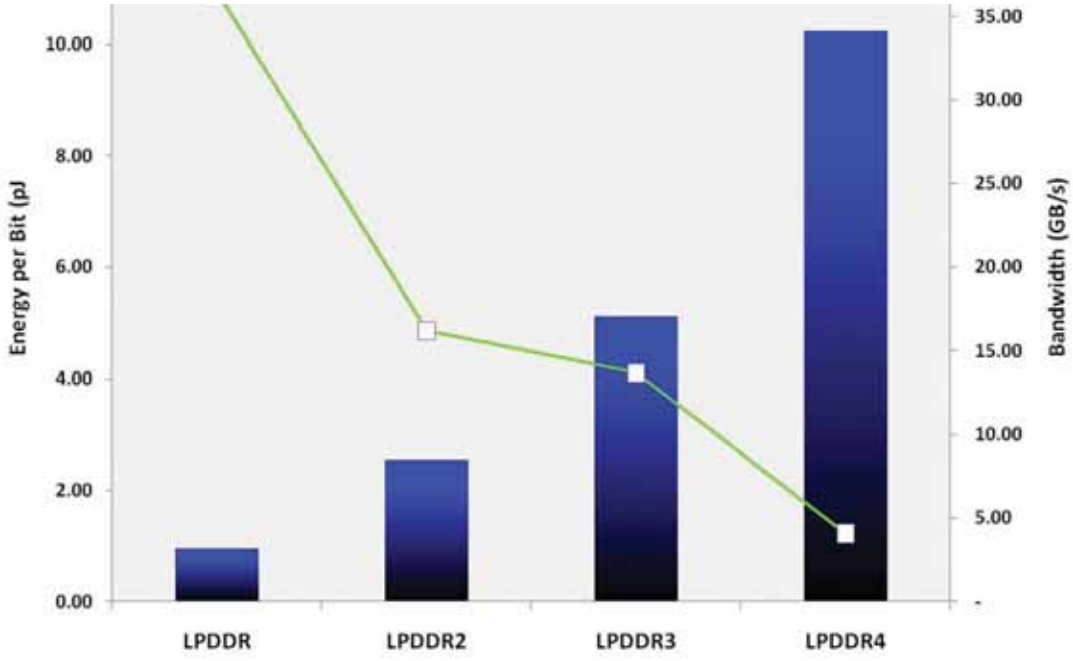

Figure 1: Mobile DRAM power requirements include active and stand-by power.

Specifically, the voltage rail has not dropped significantly—nor is battery capacity growing fast enough. The form factor for handsets is also fairly stable (with some exceptions), limiting overall battery size. Furthermore, the thermal design point (TDP) or envelope—which is the maximum amount of heat these systems can safely dissipate—is limited to 5 watts. Because the user is the final heat sink in many cases, these systems simply

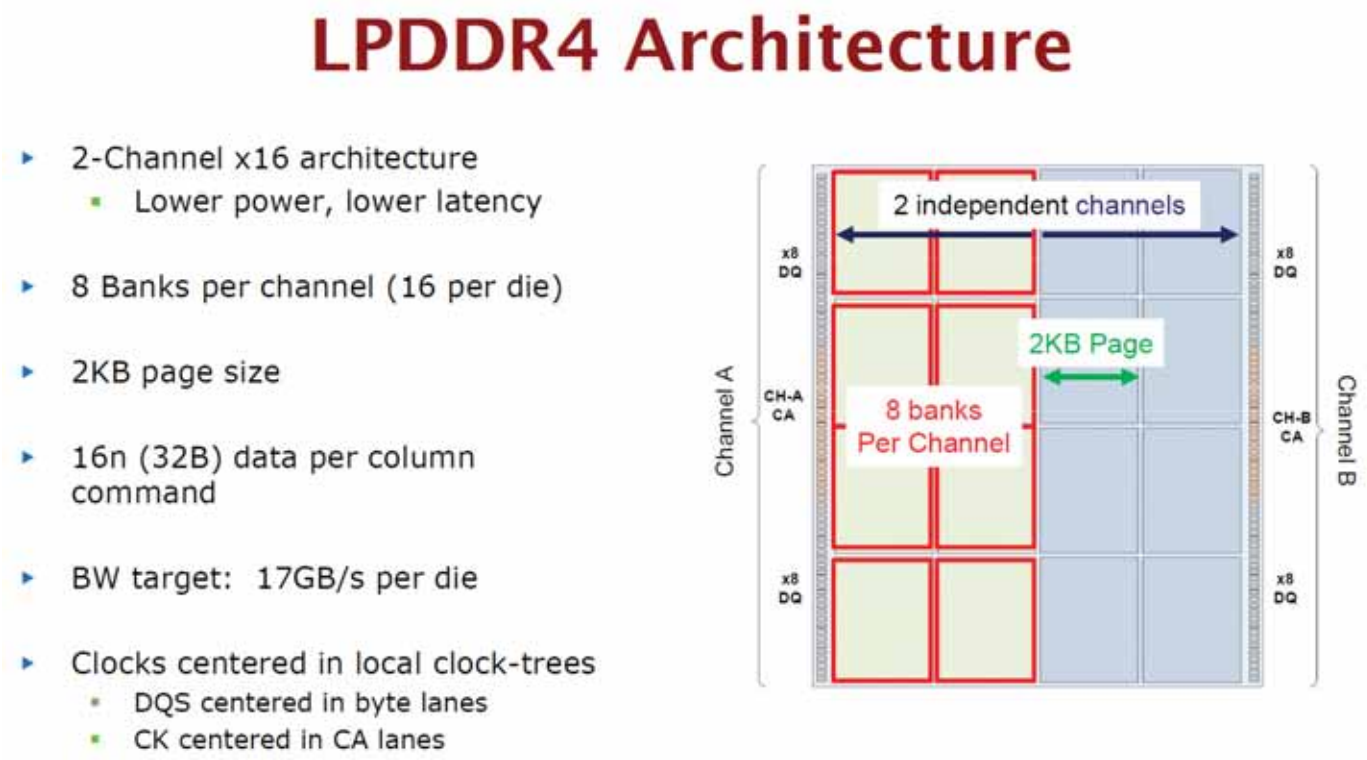

Figure 2: LPDDR4 is architected to meet the power, bandwidth, packaging, cost and compatibility requirements of the world's most advanced mobile systems.

cannot take on more heat. Because there is nothing on the horizon to suggest that these technologies are changing any time soon, even greater efficiency is required from the other parts of the system.

LPDDR4: A More Power-Efficient Memory Solution

In today's handsets, memory devices consume up to 30% of system power in standby modes. Thus, efficiencies in memory management play a substantial role in enabling manufacturers to achieve power neutrality.

To help developers improve a system's energy consumption per bit processed (see Figure 1), JEDEC has defined the fourth generation of low-power DDR (LPDDR). LPDDR4 provides more than just a speed upgrade from LPDDR3; it's an evolutionary step up thanks to enhanced functionality that can help developers achieve power neutrality in handset applications. LPDDR4 also doubles bandwidth performance, provides a low pin-count package, is backward compatible with previous generations of LPDDR and enables competitive pricing.

The LPDDR4 standard introduces several major architectural changes that are specifically designed to reduce the energy required per bit (see Figure 2):

2-Channel x 8-Bank Architecture: The internal LPDDR4 architecture contains two 16-bit channels (instead of one 32-bit channel), which reduces the effects of parasitic capacitance and results in lower active currents for READ and WRITE operations.

2K Page Size: The reduced DRAM page size (from 4K to 2K) decreases the amount of current required to activate a page when opened.

1.1V Supply Voltage: The reduced supply voltage (from 1.2V to 1.1V) provides a 20% decrease in switching power and approximately 10% savings in static power.

Advanced LVSTL Interface: The low-voltage swing (VOH) of LPDDR4's low-voltage swing-terminated logic (LVSTL) interface saves more than 50% power when switching I/O compared to LPDDR3.

The LVSTL Advantage

Low-voltage swing-terminated logic (LVSTL) is a significant enhancement to mobile memory technology. The LPDDR4 interface supports a programmable voltage level that divides the power supply rail for I/O by either 3.0 or 2.5. The 3.0 mode is intended for systems with better channel design and/or less loading. For more heavily loaded systems, or those with more channel losses, the 2.5 mode can be used to adjust power efficiency to increase signal integrity. This enables developers to balance cost, signal integrity and power. For example, a developer could select the 2.5 mode and use a lower-quality PCB material to reduce system cost.

LVSTL also provides a variety of termination settings, ranging from 40 ohm to 240 ohm. In general, the stronger the termination (40 ohm), the better the signal eye across the channel. However, stronger termination also consumes more

power. By adjusting termination settings, developers can tune systems for a variety of configurations, such as stronger termination for a lower-end handset. This flexibility also simplifies design because developers can initially set systems to the lowest power setting (240 ohm) and increase the termination level if greater signal integrity is required.

Simplified Design Process

It is important to note that LPDDR4 is not a drop-in replacement for LPDDR3. Systems need to leverage LPDDR4's power-saving features to fully optimize power efficiency. To simplify design for developers, memory suppliers like Micron have partnered with industry-leading controller manufacturers to provide optimized LPDDR4 memory, which enables developers to focus engineering resources on their own value-added innovations. As one of the creators of the LPDDR4 standard, Micron has the experience to help developers become familiar with the new memory technology as well as design custom solutions using it.

Memory Is No Longer a Bottleneck

Because LPDDR4 provides 2X the bandwidth (up to 34 GB/s) of LPDDR3 while using less energy per bit, memory is no longer a bottleneck for UHD and other compute-intensive applications in mobile devices. LPDDR4's power efficiency enables it to be an effective enabler of UHD technology, and thanks to its flexible implementation, LPDDR4 also gives developers more options for balancing signal integrity and power efficiency and, consequently, more control over cost versus power efficiency.

LPDDR4 provides more than just stellar bandwidth, improved power efficiency and power neutrality; developers can rely on packaging with lower pin counts for greater board density, low costs and backwards compatibility to previous generations of LPDDR—all without compromising performance.

The open design of today's mobile handsets mean they should be able to do everything we want them to. With LPDDR4, they can.

Dan Skinner joined Micron in 1989 as a product engineer in the Memory Application Group, and has since worked with DRAM, SRAM, flash memory and TCAM products. Throughout his career with Micron, Skinner has held management positions in engineering and marketing, most recently managing the CellularRAM™, Mobile SDRAM and RLDRAM™ product lines. He was appointed to his current position in 2006. Skinner holds a bachelor of science degree in electrical engineering from the University of Colorado and MBA from the Kellogg School of Management at Northwestern University.

Smartphones/Tablets Q&A with Microchip Technology

By Chris A. Ciufo, Editor-in-Chief, Embedded; Extension Media

Editor's note: Our thanks to Ken Nagai, Senior Product Marketing Manager for Microchip's USB and Networking Group (UNG), handling its USB Power Delivery solutions.

EECatalog: Where do Microchip products apply to mobile devices?

Ken Nagai: Microchip has shipped our USB solutions into mobile devices including both smartphones and tablets. Solutions include highly integrated and miniature USB transceivers and hub controllers.

Personal Electronics Device (PED) designs are used only in accessories for smartphones and tablets and interface using serial, USB and Bluetooth interfaces. Accessory designs range from audio docks, medical (glucose meter, blood pressure monitor, etc.), sports/fitness (chest straps, heart rate monitor, fuel bands) and general accessories (USB dongles for audio, MIDI, etc.) and wearables (smart watch). Microchip's Touch and Interface products provide a range of functionality including touch keys and grip detection. Additionally, Microchip's JukeBlox platform connects to mobile phones and tablets for audio streaming. Our products are not designed into phones or tablets.

EECatalog: The smartphone and mobile accessory market is exploding, based primarily on USB. What are some of the key trends here?

Ken Nagai: The ability to detect and switch between USB Host and USB Device mode is a key feature required by most accessory manufacturers to be able to support the broad spectrum of smartphone and tablet manufacturers. Some advanced features in the smartphone or tablet are only enabled when it is the USB Host. The ability of the accessory to identify and configure its USB port dynamically is absolutely required when interfacing to iOS, Android, Windows and other operating systems. However, the key trend is actually to move to wireless connectivity using Bluetooth Smart, Bluetooth Classic and Wi-Fi. Many manufacturers are only designing accessories with wireless interfaces for the previously mentioned applications.

EECatalog: USB 3.0 promises 5 Gbps data rates. Does this matter in these applications? Why or why not?

Ken Nagai: Yes, we are starting to see more mobile devices migrating to support USB3.0. USB is a ubiquitous port of connection for high-speed data transmission, and since mobile devices continue to increase in their storage capacity, the need to download/upload this data as quickly as possible is mandatory.

EECatalog: USB 3.0 added power profiles for charging myriad devices, and there are more charging profiles on the way. What is this all about, and what technologies are required here?

Ken Nagai: USB Power Delivery is a standard specification, which increases the power per port to 100W. Microchip has some new USB Power Delivery controller solutions to address this market need.

EECatalog: Speaking of power, how does this vary from OS to OS, and from processor to processor?

Ken Nagai: Since USB Power Delivery is a standard specification, it can attach to any processor, and can be used with any OS.

EECatalog: Mobile devices rely on batteries and low power is a given. What are the trends in power management?

Ken Nagai: Mobile devices always focus on the lowest power solutions possible. This is why Microchip has found success with our USB solutions. However, in addition, the delivery of higher power through a standard USB port is allowing faster battery charging times, as well as power sharing within the entire system configuration.

Chris A. Ciufo is editor-in-chief for embedded content at Extension Media, which includes the EECatalog print and digital publications and website, Embedded Intel® Solutions, and other related blogs and embedded channels. He has 29 years of embedded technology experience, and has degrees in electrical engineering, and in materials science, emphasizing solid state physics. He can be reached at cciufo@extensionmedia.com.

Advanced Image Stabilization Techniques for Tablet Camera Performance

By Mark Aaldering, ROHM Semiconductor

Intel processors play a leading role in the tablet and two-in-one device market, especially for those higher-performance devices targeted at business environments and high-end consumer applications. One of the more popular applications for these devices is still photography and video capture. Market research indicates that business users and consumers prefer to use their tablets to share high-quality photos or videos on Facebook, Instagram, Snapchat or other popular, visually oriented social media sites. In fact for many users, their tablet serves as a replacement for a digital still camera or inexpensive video camera.

Not surprisingly, Intel processors help make that possible. The latest generation of the Intel Atom processor, for example, not only improves overall performance and extends battery life, it also supports excellent graphics and video with integrated image signal processing for both still and video image capture. By coupling high-resolution screens with high pixel density, together with the graphics-processing capabilities embedded in Intel processors, many of today's tablets and two-in-one devices deliver extremely high-quality graphics and video.

Whether users are capturing still images or recording video, image stabilization plays a key role in producing a high-quality result by eliminating image distortion through pixel blurring and the creation of unwanted artifacts. Typically standalone cameras and mobile devices offering a photo or video function also add some form of image-stabilization capability to compensate for unintentional movements by the user. Intel-based tablets are no exception. The latest Atom processor adds multi-axis document image solution (DIS) and image alignment to help remove blur from moving objects.

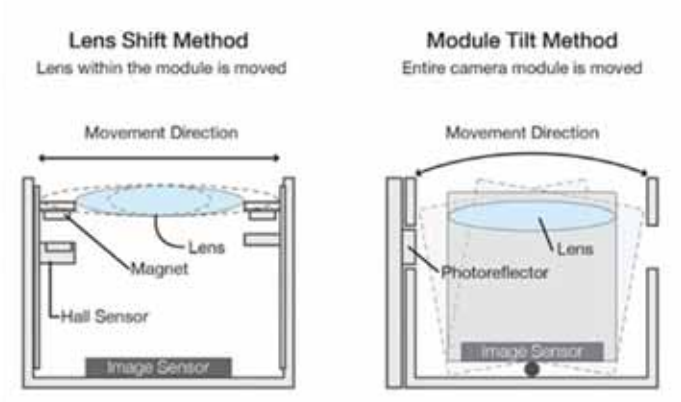

Figure 1: There are two primary methods of implementing optical image stabilization

However, as tablet and other mobile device developers move to ever-higher levels of resolution, demand is accelerating for more advanced image stabilization techniques. Two of the more common implementations—electronic image stabilization (EIS) and optical image stabilization (OIS)— are taking video and still image photography to a new level of performance.

Basic Principles

Image stabilization techniques are designed to reduce blurring associated with relatively minor shaking of the camera within a few optical degrees while the image sensor is exposed to the capturing environment. These functions are not designed to prevent motion blur caused by movement of the target subject or extreme movements of the camera itself. This minor movement of the camera by the user is characterized by its pan and tilt components where the angular movements are known as yaw and pitch, respectively. Typically, these image stabilization functions cannot compensate for camera roll because rolling the lens doesn't actually change or compensate for the roll motion, and therefore does not have any effect on the image itself relative to the image sensor.

EIS is a digital image compensation technique which uses complex algorithms to compare frame contrast and pixel location for each changing frame. Pixels on the image border provide the buffer needed for motion compensation. An EIS algorithm calculates the subtle differences between each frame and the camera uses this information to interpolate new frames to reduce the sense of motion.

EIS offers distinct advantages and disadvantages. As an image-stabilization scheme, it offers developers a relatively compact and lower-cost option. However, image quality is limited due to image scaling and image signal post-processing artifacts and any incremental improvement in image quality requires additional power to capture additional images and perform image processing. In addition, EIS solutions do not perform well at full electronic zoom (long field-of-view) and under low-light conditions.

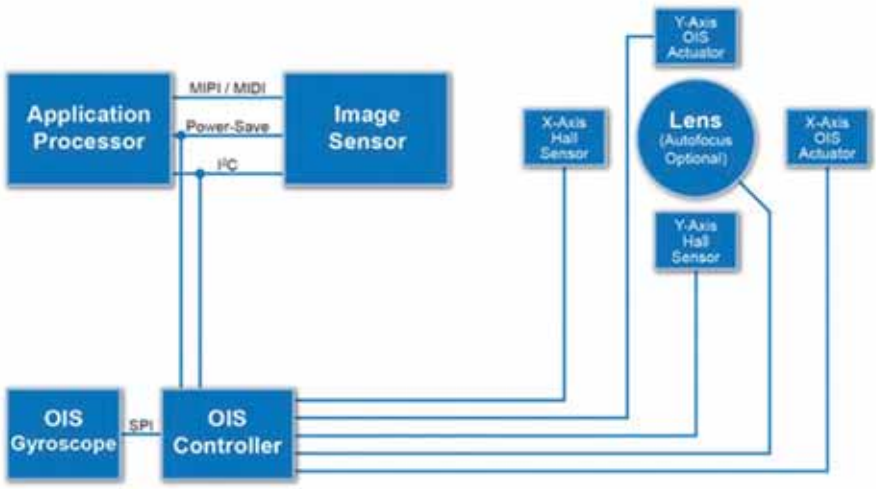

Figure 2. ROHM's OIS system uses a complete module of sensing, compensation and control components to accurately correct for unwanted camera movement.

In comparison, OIS is a mechanical technique used in imaging devices to stabilize the recording image by controlling the optical path to the image sensor. Two primary methods are used to implement OIS. One, called lens shift, involves moving the parts of the lens. The second, termed module tilt, moves the module itself (see Figure 1).

Camera movements by the user can cause misalignment of the optical path between the focusing lens and the center of the image sensor. In the OIS lens-shift method, only the lens within the camera module is controlled and used to realign the optical path to the center of the image sensor. The module tilt method, on the other hand, controls the movement of the entire module including the fixed lens and the image sensor. The module-tilt approach allows for greater range of movement compensation by the OIS system and achieves minimal image distortion because of the fixed focal length between the lens and the image sensor.

Compared to EIS solutions, OIS systems reduce blurring without significantly sacrificing image quality especially in low-light and long-range image capture. But unlike EIS which needs no additional hardware, OIS solutions require actuators and power driving sources that tend to require a larger footprint and higher cost.

Module Components

An OIS system relies on a complete module of sensing, compensation and control components to accurately correct for unwanted camera movement. This movement or vibration is characterized in the X/Y-plane, with yaw/pan and pitch/tilt movements detected by different types of isolated sensors. The lens shift method uses Hall sensors for lens movement detection while the module tilt method uses photodetectors to detect human movement. OIS controllers can use gyroscope data within a lens target-positioning circuit to predict where the lens needs to return in order to compensate for the user's natural movement. With lens shift, Hall sensors are used to detect real-time X/Y locations of the lens after taking into consideration actuator mechanical variances and the influence of gravity. The controller uses a separate internal servo system that combines the lens positioning data of the Hall sensors with the target lens position calculation from the gyroscope to calculate the exact driving power needed for the actuator to reposition the lens. The process is similar with module tilt but the module's location is measured and repositioned instead of just the lens. With both methods, the new lens position realigns the optical path to the center of the image sensor.

OIS control is designed to be very simple from the customer's standpoint, consisting simply of ON/OFF and enable/power-save modes. The only other commands are optional manual control of the lens in the X/Y plane or altering OIS performance based on ambient conditions such as day, night, sports, picture, video or viewfinder. This minimizes I2C traffic from the host processor to the OIS controller and simplifies software driver development for the end customer. All of the actual OIS control algorithms are performed autonomously on the controller itself, using the internal processor and RAM.

OIS Controller Considerations

Controller architectures for OIS applications vary significantly. Some combine a programmable core with custom programmable digital signal processing for gyroscope signal processing and servo control. Others integrate programmable gyroscope signal processing and servo control into the core itself. Typically all OIS memory and control calculations are performed on the OIS controller and do not require an external host processor's computational resources or external memory.

Developers looking for a controller for OIS applications should consider a number of issues. Does the controller offer full control

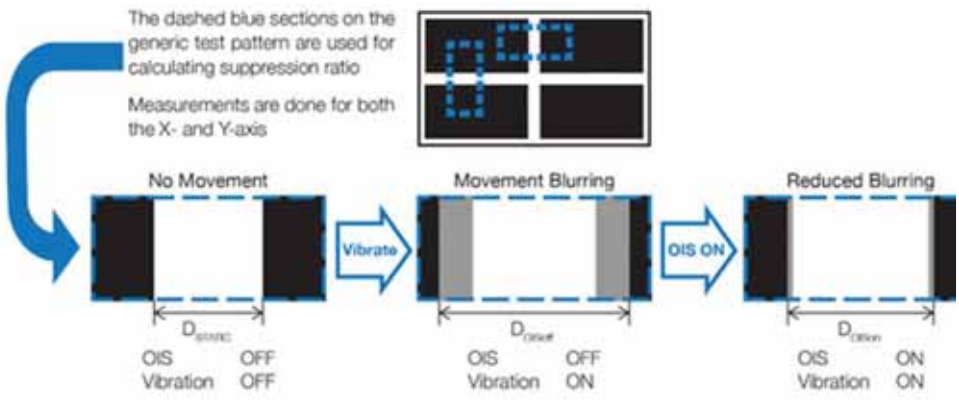

Figure 3. The DOISoff image exhibits much more blurring compared to the other images in generic test pattern.

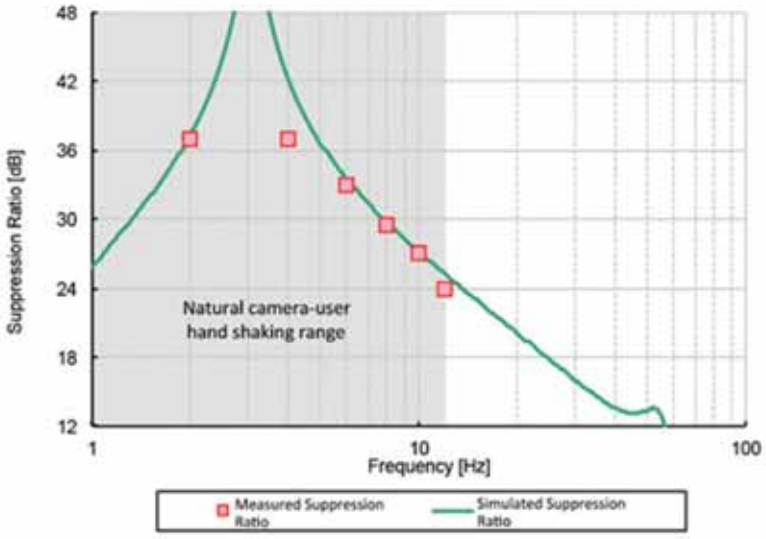

Figure 4. Graph compares real-world OIS performance vs. ROHM's simulated OIS performance.

of the X- and Y-axis voice coil motor (VCM) drivers, Hall amplifier and current drivers and photo-reflector drivers? Does it feature the wide variety of interfaces and peripherals needed for the application including I2C, ADCs, PLL oscillators, SPI master for digital gyroscopes and support for analog gyroscopes? Does the MCU support integrated drivers for autofocus, neutral density filters or shutter functions? Be aware that some controllers offer digital filter designs in their servo control and gyroscope signal processing circuits that can improve performance by dynamically compensating for gyroscope and actuator temperature drift while not removing intentional pan and tilt movement by the camera user. Others add custom control software for automatic lens control, automatic pan-tilt detection and access to different programmable capturing modes and calibration settings.

Measuring Image Stabilization

Image stabilization is measured by suppression ratio (SR) and is utilized to gauge OIS performance. SR is calculated using a spatial test chart with a target pattern. Images of the target pattern are captured with OIS ON/OFF and with/without vibration. The images with and without OIS are then compared to compute a ratio of the amount of blur in each image. This test is typically used to provide a final guarantee that all of the components in the OIS system are functioning properly.

The figure below depicts examples of motion blur in the target pattern. The DSTATIC image represents an ideal result with no vibration or motion in the image. Ideally an OIS system attempts to match the quality of a still image with no motion blur and the DSTATIC image serves as a benchmark for calculating SR performance of the OIS system. In this example the DSTATIC image exhibits the shortest zoomed white area distance due to the absence of movement or blurring in the captured image. The DOISoff image represents the appearance of an image when it is vibrating or moving without using image stabilization. As a result, the DOISoff image exhibits much more blurring compared to the other images.

The observed amount of blur represents what needs to be corrected or suppressed to match the DOISoff image with the DSTATIC image. Therefore, the DOISon image represents the actual benefit of the OIS system. In this example, the DOISon image depicts an image that is vibrating or moving while image stabilization is enabled. The stabilization system suppresses blurring of the image and the distance of the zoomed white area is less than when compared to the DOISoff image. Once all three images have been captured, the blurring effect of each image is measured as a function of pixel count by calculating the number of pixels within the width of the zoomed white area and then using equation 1 (shown below diagram in Figure 3) to calculate final SR. This process is repeated for each image shaking frequency performance target and for each axis.

System Testing

Proper OIS operation requires simulating the entire system to ensure that all components interact correctly together. While most OIS controller suppliers can simulate the ideal performance of golden OIS components such as the actuator, ROHM has developed highly specialized simulation tools that allow not only for simulation of OIS components, but also provide real-world OIS component simulations as well. These real-world results help accelerate the development of custom firmware for customers integrating OIS into their design (see Figure 4).

OIS systems also require careful calibration to ensure proper operation. All of the components within the OIS system possess individual manufacturing variances and assembly misalignment variances. A properly functioning system, the OIS controller must know the subtle sensitivity variances introduced by the manufacturing and assembly processes. Once the calibration process is complete, the OIS controller can use the collected data to modify control of the system and its components.

Summary

As next-generation tablets and two-in-one devices migrate up the performance curve, users will increasingly demand higher performance image and video capture capabilities. High of users' list will be crisp, clear and blur-free images. By leveraging the latest advances in optical image stabilization, tablet and two-in-one device designers can meet those expectations.

Mark Aaldering is the senior director of technical product marketing at ROHM Semiconductor where his dedicated team drives new products into development and adoption in the computing, consumer, automotive and industrial markets.

IoT Standards Enable Interoperability

Experts weigh in on challenges in IoT device development, the value of standards, and the importance of standards groups such as the Industrial Internet Consortium.

By Kenton Williston, Intel® Internet of Things Alliance

Jens Wiegand, CTO of Kontrongeneric test pattern.

Ido Sarig, vice president and general manager of IoT Solutions Group at Wind River

Tony Magallanez, OEM Systems Engineer at McAfee

The Internet of Things (IoT) promises a future where everything is online. But today, a lack of standards makes it difficult to connect. To learn how developers can solve this problem, I spoke with three industry experts:

- Jens Wiegand, CTO of Kontron
- Ido Sarig, vice president and general manager of IoT Solutions Group at Wind River
- Tony Magallanez, OEM Systems Engineer at McAfee

Below are key excerpts from my interviews

Kenton Williston: What are the biggest challenges to deploying IoT solutions?

Jens Wiegand, Kontron: Currently the market is fragmented and characterized by incompatible systems and stovepiped solutions. IoT concepts like predictive maintenance, big data and analytics require a holistic approach, but there is a lack of cooperation between hardware and software suppliers, service providers and communication infrastructure vendors.

Ido Sarig, Wind River: Much of the industry's effort is focused on connecting legacy or "brownfield" devices that were not designed to be connected and even designed to make connectivity difficult in order to protect them from network-borne threats. Developers must figure out not only how to connect brownfield devices but how to safely connect them.

Another challenge is the lack of a single standard for connecting to networks. Many brownfield devices use proprietary protocols and will require gateways to connect with IP-based networks. And if theyare already IP-based, they may be using a wide variety of protocols. Developers will need to be able to build gateways that support virtually any communication protocol.

Tony Magallanez, McAfee: The major problem we see is the security of these devices. These devices tend to be not manned but often handle personally identifiable information. The question is how you protect the data both while it's on the system and while it's being transmitted between devices.

KW: How can developers address these issues? In particular, how do standards and multi-vendor solutions help?

Wiegand, Kontron: Developers should strive to build on solutions that adhere to industry standards on all levels, from communications protocols to cloud connectors. In particular, they should seek standards that are supported by multiple industry leaders in the form of application-ready concepts. Such standards can reduce complexity and risk, and provide a time-to-market advantage.

Sarig, Wind River: Delivering secure and reliable IoT solutions requires an end-to-end view that encompasses the endpoint device, the connectivity layer, the gateway and the application running in the cloud. For example, security challenges need to be factored in at every level. Virtually every known type of hardware and software security measure comes into play in IoT. Secure booting at the device level, access control and authentication, application

whitelisting and firewalls and intrusion prevention systems are just some of the tools at hand to respond to security threats.

The benefits of the IoT have been thus far constrained by the complexity of issues like this. As standards coalesce, market needs and business cases become more sharply defined, and operators and device manufacturers are freed to focus on the true value they can deliver: innovative new services and applications.

KW: What role do you see for standards bodies like the new Industrial Internet Consortium (IIC) (http://www.iiconsortium.org)?

Wiegand, Kontron: The IIC as well as Industry 4.0 (http://en.wikipedia.org/wiki/Industry_4.0) are good examples where a consortium of industry leaders drives towards a common goal: enabling business value for end customers by implementing standards and by developing the ecosystem to enable solutions.

Sarig, Wind River: Organizations such as the IIC bring together expertise and the tools to bring smart connectivity, high security and manageability to the market. These consortia will help expedite the realization of the IoT through specialized skills and expertise required to build intelligent devices which typically reside outside the core competency of operators and device manufacturers.

Magallanez, McAfee: Most people understand that they need security, but in many cases they lack expertise –and in far too many cases they end up doing nothing if they don't have to. That's where I see consortiums or standards-based organizations driving security.

Talking specifically about IIC, its recommendations are likely to overlap with the North American Electric Reliability Corporation Critical Infrastructure Protection (NERC CIP—http://www.nerc.com/pa/Stand/Pages/CIPStandards.aspx) recommendations. But one of the great things about these consortiums is that they go further than the regulations require. For example, they may recommend things like application whitelisting. In the regulatory bodies, whitelisting is still a bit of an outlying technology. So these standards bodies can help developers not only achieve compliance but also true security.

So how do you comply with these regulations? Typically, OEMs or ODMS have taken a buy or build mentality. However, it's very difficult to build your own security infrastructure. So it's wise to get help from vendors who are experts in security, and ask these security solution vendors to make sure the systems are secure.

KW: What role do you see for ecosystems like the Intel® IoT Solutions Alliance (http://iotsolutionsalliance.intel.com)?

Wiegand, Kontron: The Alliance is a great example of a large-scale ecosystem that enables rapid deployment with IoT solutions that are pre-integrated, verified and validated by system integrators and solution providers like Kontron. The benefits for our customers are flexibility, choice, velocity and the ability to focus on the development of innovative applications with less risk and pain.

Magallanez, McAfee: Alliance members will help provide the components to meet regulatory recommendations or requirements. They will give you the building blocks to get you to compliance and beyond. Without these groups, you will have different OEMs/ODMs and device owners all struggling to define what the security should look like. Not being security experts, they can miss out on some of the security opportunities.

KW: How you are using standards and ecosystem collaboration to create IoT solutions?

Sarig, Wind River: As part of Intel IoT Group, Wind River is collaborating with Intel on solutions like the Intel® Gateway Solutions for the Internet of Things (http://www.intel.com/content/www/us/en/internet-of-things/gateway-solutions.html), which serves as the software backbone for intelligent gateways. It is a complete software development environment that provides pre-integrated and fully tested ready-to-use components to secure, manage and connect intelligent gateways.

Magallanez, McAfee: What we are doing with the Intel Gateway Solutions for the IoT is providing a platform that will allow OEMs/ODMs to establish a base level of security and functionality in the device without having to do a lot of the development on their own. In addition, the hardware that's built in gives them a starting point that is easier than taking a huge SKU sheet and picking out the components individually.

Wiegand, Kontron: The Intel Gateway Solutions for the IoT is a good example of how Kontron builds on industry standards and capitalizes on the work of the Alliance by leveraging platform concepts to create IoT-enabled hardware and software stacks. [Editor's note: Kontron recently announced the Intel® Gateway Solutions for the IoT-based KBox A-201 mini.] But even more, we strive to develop standards further, grow the ecosystem and offer more value to our customers with IoT solutions that enhance the reach of our application-ready platforms. We see solution- and application-readiness as well as software expertise as key differentiators and innovation enablers.

This article first appeared in the Intel® Embedded Community (http://embedded.communities.intel.com/community/en), published by the Intel® Internet of Things Alliance.

Kenton Williston is editor-in-chief of the Embedded Innovator magazine, a publication of the Intel® Internet of Things Solutions Alliance (Intel® IoT Solutions Alliance). He is also a regular blogger for the Intel® Embedded Community, and contributes to several other publications as an Intel contractor.

WE DESIGNED THIS WEBSITE FOR YOU...

Complete Coverage of 30+ Key Embedded Technologies

Valuable Company and Product Information

Opinions and Insight

EXCLUSIVE Videos!

White Papers

Daily News

Featured Blogs

www.eecatalog.com